# Faster Algorithms for Convex and Combinatorial Optimization 

by<br>Yin Tat Lee<br>Submitted to the Department of Mathematics in partial fulfillment of the requirements for the degree of<br>Doctor of Philosophy in Mathematics<br>at the<br>MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2016
(c) Massachusetts Institute of Technology 2016. All rights reserved.


#### Abstract

Author $\qquad$

May 13, 2016


Certified by $\qquad$
Jonathan A. Kelner
Associate Professor
Thesis Supervisor

Accepted by $\qquad$
Jonathan A. Kelner
Chairman, Applied Mathematics Committee

# Faster Algorithms for Convex and Combinatorial Optimization 

\author{


#### Abstract

Submitted to the Department of Mathematics on May 13, 2016, in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Mathematics


}


#### Abstract

In this thesis, we revisit three algorithmic techniques: sparsification, cutting and collapsing. We use


 them to obtain the following results on convex and combinatorial optimization:- Linear Programming: We obtain the first improvement to the running time for linear programming in 25 years. The convergence rate of this randomized algorithm nearly matches the universal barrier for interior point methods. As a corollary, we obtain the first $\tilde{O}(m \sqrt{n})$ time randomized algorithm for solving the maximum flow problem on directed graphs with $m$ edges and $n$ vertices. This improves upon the previous fastest running time of $\tilde{O}\left(m \min \left(n^{2 / 3}, m^{1 / 2}\right)\right)$, achieved over 15 years ago by Goldberg and Rao.
- Maximum Flow Problem: We obtain one of the first almost-linear time randomized algorithms for approximating the maximum flow in undirected graphs. As a corollary, we improve the running time of a wide range of algorithms that use the computation of maximum flows as a subroutine.
- Non-Smooth Convex Optimization: We obtain the first nearly-cubic time randomized algorithm for convex problems under the black box model. As a corollary, this implies a polynomially faster algorithm for three fundamental problems in computer science: submodular function minimization, matroid intersection, and semidefinite programming.
- Graph Sparsification: We obtain the first almost-linear time randomized algorithm for spectrally approximating any graph by one with just a linear number of edges. This sparse graph approximately preserves all cut values of the original graph and is useful for solving a wide range of combinatorial problems. This algorithm improves all previous linear-sized constructions, which required at least quadratic time.
- Numerical Linear Algebra: Multigrid is an efficient method for solving large-scale linear systems which arise from graphs in low dimensions. It has been used extensively for 30 years in scientific computing. Unlike the previous approaches that make assumptions on the graphs, we give the first generalization of the multigrid that provably solves Laplacian systems of any graphs in nearly-linear expected time.

Thesis Supervisor: Jonathan A. Kelner
Title: Associate Professor

## Acknowledgments

First, I would like to thank my advisor Jonathan Kelner. Jon is always full of ideas and provided me a lot of high-level inspiration. I am particularly thankful to him for encouraging me to learn interior point methods, which led to many of key results in this thesis. From the very beginning, Jon has given me lots of support, freedom and advice. Because of him, I have achieved what I never imagined. I will be forever grateful to Jon.

Next, I would like to thank my colleague, Aaron Sidford. Over the past 4 years, one of my most exciting activities has been discussing with Aaron in how to solve Maximum Flow faster. During thousands of hours of discussions, we obtained many of the results presented in this thesis. Although we still haven't gotten the ultimate result we want, I believe, someday, we will get it.

I would also like to thank Lap Chi Lau. Originally, I was planning to study scientific computing, but he made me change my mind. Since my undergraduate, I have been asking him for all kinds of career advice and, up to now, all of them are proven to be useful.

Next, I would like to thank Pravin Vaidya for his beautiful work on his interior point methods, cutting plane methods and Laplacian systems. Vaidya's results were definitely far ahead of his time and many results in this thesis are inspired by and built on his results.

Also, I would like to thank my academic friends and collaborators: Sébastien Bubeck, Michael Cohen, Laurent Demanet, Michel Goemans, Michael Kapralov, Tsz Chiu Kwok, Rasmus Kyng, Aleksander Mądry, Gary Miller, Cameron Musco, Christopher Musco, Lorenzo Orecchia, Shayan Oveis Gharan, Jakub Pachocki, Richard Peng, Satish Rao, Sushant Sachdeva, Ludwig Schmidt, Mohit Singh, Daniel Spielman, Nikhil Srivastava, Robert Strichartz, He Sun, Luca Trevisan, Santosh Vempala, Nisheeth Vishnoi, Matt Weinberg, Sam Chiu-wai Wong, and Zeyuan Allen Zhu.

Last but not the least, I would like to thank my parents Man Cheung and Sau Yuen, my siblings Ah Chi and Ah Tat, and my wife Yan Kit.

Yin Tat Lee
Cambridge, Massachusetts
May 2016
My work was partially supported by NSF awards 0843915 and 1111109. Some of my work was done while I was visiting the Simons Institute for the Theory of Computing, UC Berkeley.

## Contents

1 Introduction ..... 9
1.1 Techniques ..... 9
1.2 Results ..... 11
2 Preliminaries and Thesis Structure ..... 15
2.1 Problems ..... 15
2.2 Thesis Organization ..... 17
2.3 Notation ..... 24
I Sparsification ..... 29
$3 \ell_{2}$ Regression In Input Sparsity Time ..... 31
3.1 Introduction ..... 31
3.2 Previous Work ..... 34
3.3 Generalized Leverage Score ..... 35
3.4 Expectation Bound on Leverage Score Estimation ..... 35
3.5 Coherence-Reducing Weighting ..... 36
3.6 High Probability Bound on Leverage Score Estimation ..... 38
3.7 Main Result ..... 39
3.8 Appendix ..... 42
4 Linear-Sized Sparsifier In Almost-Linear Time ..... 45
4.1 Introduction ..... 45
4.2 Algorithm ..... 45
4.3 Analysis ..... 49
4.4 Appendix ..... 58
$5 \ell_{p}$ Regression and John Ellipsoid ..... 63
5.1 Introduction ..... 63
5.2 Lewis Weights ..... 63
5.3 Approximate Lewis Weights ..... 64
5.4 Approximate John Ellipsoids ..... 67
6 Faster Linear Programming ..... 69
6.1 Introduction ..... 69
6.2 Previous Work ..... 70
6.3 Simple Case ..... 72
6.4 Difficulties ..... 76
6.5 Weighted Path Finding ..... 77
6.6 Progressing Along Weighted Paths ..... 82
6.7 Weight Function ..... 91
6.8 The Algorithm ..... 100
6.9 Generalized Minimum Cost Flow ..... 104
6.10 Glossary ..... 106
6.11 Appendix: Technical Lemmas ..... 107
6.12 Appendix: Projection on Mixed Norm Ball ..... 109
6.13 Appendix: A $\tilde{O}(n)$-Self-Concordant Barrier Function ..... 110
7 Geometric Median In Nearly-Linear Time ..... 123
7.1 Introduction ..... 123
7.2 Preliminaries ..... 128
7.3 Properties of the Central Path ..... 129
7.4 Overview of the Algorithm ..... 130
7.5 Analysis of the Central Path ..... 133
7.6 Analysis of the Algorithm ..... 143
7.7 Pseudo Polynomial Time Algorithm ..... 148
7.8 Derivation of Penalty Function ..... 151
7.9 Technical Lemmas ..... 152
7.10 Weighted Geometric Median ..... 156
II Cutting ..... 158
8 Convex Minimization In Nearly-Cubic Time ..... 159
8.1 Introduction ..... 159
8.2 Preliminaries ..... 163
8.3 Our Cutting Plane Method ..... 164
8.4 Technical Tools ..... 186
8.5 Glossary ..... 192
9 Effective Use of Cutting Plane Method ..... 195
9.1 Introduction ..... 195
9.2 Preliminaries ..... 201
9.3 Convex Optimization ..... 202
9.4 Intersection of Convex Sets ..... 209
10 Submodular Minimization In Nearly-Cubic Time ..... 221
10.1 Introduction ..... 221
10.2 Preliminaries ..... 224
10.3 Improved Weakly Polynomial Algorithms for SFM ..... 227
10.4 Improved Strongly Polynomial Algorithms for SFM ..... 229
10.5 Discussion and Comparison with Previous Algorithms ..... 252
11 First Order Method vs Cutting Plane Method ..... 255
11.1 Introduction ..... 255
11.2 Intuition ..... 256
11.3 An Optimal Algorithm ..... 257
11.4 Politician ..... 259
11.5 Experiments ..... 264
11.6 Discussion ..... 268
11.7 Appendix: Convergence of SD+ ..... 270
III Collapsing ..... 276
12 Undirected MaxFlow In Almost-Linear Time ..... 277
12.1 Introduction ..... 277
12.2 Preliminaries ..... 282
12.3 Solving Max-Flow Using a Circulation Projection ..... 283
12.4 Oblivious Routing ..... 289
12.5 Flow Sparsifiers ..... 294
12.6 Removing Vertices in Oblivious Routing Construction ..... 304
12.7 Nonlinear Projection and Maximum Concurrent Flow ..... 309
12.8 Appendix ..... 315
13 Connection Laplacian In Nearly-Linear Time ..... 319
13.1 Introduction ..... 319
13.2 Preliminaries ..... 321
13.3 Overview of the Algorithms ..... 322
13.4 Summary ..... 328
13.5 Background ..... 329
13.6 Block Diagonally Dominant Matrices ..... 330
13.7 Schur Complement Chains ..... 332
13.8 Finding $\alpha$-bDD Subsets ..... 335
13.9 Jacobi Iteration on $\alpha$-bDD Matrices ..... 337
13.10 Existence of Linear-Sized Approximate Inverses ..... 339
13.11 Spectral Vertex Sparsification Algorithm ..... 341
13.12 Estimating Leverage Scores by Undersampling ..... 356
13.13 Recursive Construction of Schur Complement Chains ..... 358
13.14 Weighted Expander Constructions ..... 363
IV Geometry ..... 374
14 Sampling In Sub-Quadratic Steps ..... 375
14.1 Introduction ..... 375
14.2 Preliminaries ..... 379
14.3 Intuition ..... 386
14.4 Convergence of the Geodesic Walk ..... 389
14.5 Logarithmic Barrier ..... 407
14.6 Collocation Method for ODE ..... 431
14.7 Derivative Estimations ..... 437

## Chapter 1

## Introduction

Convex optimization has been studied extensively and is a prominent tool used in various areas such as combinatorial optimization, data analysis, operations research, and scientific computing. Each field has developed specialized tools including data structures, sampling methods, and dimension reduction. In this thesis, we combine and improve the optimization techniques from different fields to design provably faster optimization algorithms for several important problems, including linear programming, general non-smooth convex minimization, and maximum flow problems.

## ■ 1.1 Techniques

The techniques used in this thesis can be categorized into three types: sparsification, cutting and collapsing. Here, we explain these three techniques and highlight their applications to convex optimization and combinatorial optimization.

## ■ 1.1.1 Sparsification

Sparsification techniques originated from graph theory, and it refers to the task of approximating dense graphs by sparse graphs. There are different types of sparsifiers specifically designed for different graph problems, such as spanners for distance-related problems [49], cut sparsifiers for cut-related problems [33], spectral sparsifiers for spectral related problems [243, 242]. For instance, if we want to compute a minimum $s$ - $t$ cut on a given dense graph, instead of running our favorite $s$-t cut algorithm on the original graph, we can first compute a cut sparsifier of the original graph, then run the $s$-t cut algorithm on the sparsifier. These sparsifiers allow us to speed up graph algorithms for many problems on dense graphs by paying some error introduced by the sparsifiers.

Beyond graph related problems, this notion of spectral sparsifier has been generalized and applied to solve different matrix-related problems [76,170,56]. In the linear regression problem, we are given a matrix $\boldsymbol{A}$ and a vector $\boldsymbol{b}$, and we want to minimize $\|\boldsymbol{A} \boldsymbol{x}-\boldsymbol{b}\|_{2}$ over all $\boldsymbol{x} \in \mathbb{R}^{n}$. The sparsification concept allows us to compress a tall matrix $\boldsymbol{A}$ into an almost square matrix. This is useful in various contexts such as statistics problems with much more samples than variables. However, how can we find those sparsifiers in the first place without solving a dense problem? How sparse of a matrix can we effectively attain? How general is the idea of sparsification? In Part I, we answer these questions. In particular, we show how to solve overdetermined systems faster, how to find linear-sized sparsifiers of graphs in almost-linear time, and how to find geometric medians of point clouds in nearly-linear time. In one of the key results of this part, we extend the idea of sparsification to linear programming and give the first improvement to the running time for linear programming in 25 years.

## - 1.1.2 Cutting

Cutting refers to the task of iteratively refining a region that contains the solution. This technique is particularly useful for minimizing convex functions. In combinatorial optimization, Khachiyan's seminal result in 1980 [142] proved that the ellipsoid method, a cutting plane method ${ }^{1}$, solves linear programs in polynomial time. Since then, cutting plane methods have been crucial to solving discrete problems in polynomial time [110]. In continuous optimization, cutting plane methods have long played a critical role in convex optimization, where they are fundamental to the theory of non-smooth optimization [102].

Despite the key role that cutting plane methods have played historically in both combinatorial and convex optimization, over the past two decades, the theoretical running time of cutting plane methods for convex optimization has not been improved since the breakthrough result by Vaidya in 1989 [253, 253]. Moreover, for many of the key combinatorial applications of cutting plane methods, such as submodular minimization, matroid intersection, and submodular flow, the running time improvements over the past two decades have been primarily combinatorial; that is they have been achieved by discrete algorithms that do not use numerical machinery such as cutting plane methods.

In Part II, we make progress towards cutting plane methods on multiple directions. Firstly, we show how to improve on the running time of cutting plane methods. Secondly, we provide several frameworks for applying the cutting plane methods and illustrate the efficacy of these frameworks by obtaining faster running times for semidefinite programming, matroid intersection, and submodular flow. Thirdly, we show how to couple our approach with the problem-specific structure and obtain faster weakly and strongly polynomial running times for submodular function minimization, a problem of tremendous importance in combinatorial optimization. In both cases, our algorithms are faster than the previous best running times by a factor of roughly $\Omega\left(n^{2}\right)$. Finally, we give some variants of cutting plane methods that is suitable for large-scale problems and demonstrate empirically that our new methods compare favorably with the state-of-the-art algorithms.

## ■ 1.1.3 Collapsing

Collapsing refers to the task of removing variables from problems while (approximately) preserving solutions. There are different collapsings for different situations, such as mathematical induction for proofs, Gaussian elimination and Schur complement for matrices, and Fourier-Motzkin elimination for linear programs. Since collapsing often increases the density of the problems, it often does not lead to efficient algorithms.

In a seminal work of Spielman and Teng, they showed that this density issue can be avoided for graph Laplacians via sparsification. In Part III, we generalize this idea and show how to combine collapsing and sparsification for the maximum flow problem and various linear systems. As a result, we obtain the first almost-linear time algorithms for approximating the maximum flow in undirected graphs ${ }^{2}$ and the first nearly-linear-time algorithms for solving systems of equations in connection Laplacians, a generalization of Laplacian matrices that arises in many problems in image and signal processing.

In general, we assume $T(m, n)$ to be the time needed to solve a certain graph problem on a graph with $m$ edges and $n$ vertices. Suppose that we know $\mathcal{T}(m, n)=O(m)+O(\mathcal{T}(O(n), n))$ by sparsification and $\mathcal{T}(m, n)=O(m)+O(T(O(m), n / 2))$ by collapsing, then it is easy to show that $\mathcal{T}(m, n)=O(m)$

[^0]by alternatively sparsifying and collapsing. Although we do not know how to do sparsifying and collapsing in general, we believe that the concept of combining sparsifying and collapsing is a powerful technique that it will provide a foundation of future works.

## ■ 1.2 Results

By applying and improving the three techniques mentioned above, we have substantially advanced the state-of-the-art algorithms for solving many fundamental problems in computer science. Here, we describe three of the key results in this thesis.

## ■ 1.2.1 Convex Programming

Many optimization problems in theoretical computer science are shown to be polynomial-time-solvable by reducing to one of the following two general problems:

- The Feasibility Problem: Given a set $K \subset \mathbb{R}^{n}$ and a separation oracle that, given a point $\boldsymbol{x}$, either outputs that $\boldsymbol{x}$ is in $K$ or outputs a separating hyperplane. The problem is to find a point in $K$ or prove that $K$ is almost empty.
- The Intersection Problem: Given two sets $K_{1}, K_{2} \subset \mathbb{R}^{n}$, any vector $\boldsymbol{d}$ and an optimization oracle that, given a vector $\boldsymbol{c}$, outputs an $\boldsymbol{x}_{1}$ that maximizes $\boldsymbol{c}^{\top} \boldsymbol{x}_{1}$ over $K_{1}$ and $\boldsymbol{x}_{2}$ that maximizes $\boldsymbol{c}^{\top} \boldsymbol{x}_{2}$ over $K_{2}$. The problem is to find a point $\boldsymbol{x}$ nearly inside $K_{1} \cap K_{2}$ such that it approximately maximizes $\boldsymbol{d}^{\top} \boldsymbol{x}$.

For many problems, this reduction is either the best or the only known approach. Therefore, it is important to find an efficient algorithm for these two general problems. Since a breakthrough of Vaidya in 1989 [253], the feasibility problem can be solved in $\widetilde{O}\left(n T+n^{\omega+1}\right)$ time (Table 1.1), where $T$ is the cost of the oracle and $\omega \sim 2.373$. Furthermore, using the seminal works of Grötschel, Lovász, Schrijver and independently, Karp and Papadimitriou in 1981 and 1982 [110, 138], the intersection problem can be solved by solving $n$ feasibility problems. In this thesis, we show that both problems can be solved in $\widetilde{O}\left(n T+n^{3}\right)$ expected time ${ }^{3}$, the first improvement for both problems over 25 years.

As a corollary of these two theorems, we obtained polynomial improvements for many important problems, such as submodular function minimization (Table 1.2), matroid intersection (Table 1.3, 1.4), submodular flow (Table 1.5), and semidefinite programming (Table 1.6).

| Year | Algorithm | Complexity |
| :---: | :---: | :---: |
| 1979 | Ellipsoid Method [236, 274, 142] | $\widetilde{O}\left(n^{2} T+n^{4}\right)$ |
| 1988 | Inscribed Ellipsoid [144, 213] | $\widetilde{O}\left(n T+n^{4.5}\right)$ |
| 1989 | Volumetric Center [253] | $\widetilde{O}\left(n T+n^{\omega+1}\right)$ |
| 1995 | Analytic Center [20] | $\widetilde{O}\left(n T+n^{\omega+1}\right)$ |
| 2004 | Random Walk [36] | $\rightarrow \widetilde{O}\left(n T+n^{7}\right)$ |
| - | Chapter 8 | $\widetilde{O}\left(n T+n^{3}\right)$ |

Table 1.1: Algorithms for the Feasibility Problem. The arrow, $\rightarrow$, indicates that it solves a more general problem where only a membership oracle is given.

[^1]| Year | Author | Complexity |
| :---: | :---: | :---: |
| 1981,1988 | Grötschel, Lovász, Schrijver [110, 109] | $\widetilde{O}\left(n^{5} \cdot \mathrm{EO}+n^{7}\right)[189]$ |
| 1985 | Cunningham [61] | $O\left(M n^{6} \log (n M) \cdot \mathrm{EO}\right)$ |
| 2000 | Schrijver [231] | $O\left(n^{8} \cdot \mathrm{EO}+n^{9}\right)$ |
| 2000 | Iwata, Fleischer, Fujishige [121] | $O\left(n^{5} \cdot \mathrm{EO} \log M\right)$ |
| $O\left(n^{7} \log n \cdot \mathrm{EO}\right)$ |  |  |
| 2000 | Iwata, Fleischer [87] | $O\left(n^{7} \cdot \mathrm{EO}+n^{8}\right)$ |
| 2003 | Iwata [118] | $O\left(\left(n^{4} \cdot \mathrm{EO}+n^{5}\right) \log M\right)$ <br> $O\left(\left(n^{6} \cdot \mathrm{EO}+n^{7}\right) \log n\right)$ |
| 2003 | Vygen [264] | $O\left(n^{7} \cdot \mathrm{EO}+n^{8}\right)$ |
| 2007 | Orlin [216] | $O\left(n^{5} \cdot \mathrm{EO}+n^{6}\right)$ |
| 2009 | Iwata, Orlin [125] | $O\left(\left(n^{4} \cdot \mathrm{EO}+n^{5}\right) \log (n M)\right)$ |
| $O\left(\left(n^{5} \cdot \mathrm{EO}+n^{6}\right) \log n\right)$ |  |  |

Table 1.2: Algorithms for submodular function minimization. EO is the time for evaluation oracle of the submodular function and $M$ is the maximum absolute function value.

| Year | Author | Complexity |
| :---: | :---: | :---: |
| 1968 | Edmonds [81] | not stated |
| 1971 | Aigner, Dowling [6] | $O\left(n^{3} \mathcal{T}_{\text {ind }}\right)$ |
| 1975 | Lawler [157] | $O\left(n^{3} \mathcal{T}_{\text {ind }}\right)$ |
| 1986 | Cunningham [60] | $O\left(n^{2.5} \mathcal{T}_{\text {ind }}\right)$ |
| - | Chapter 9 | $\widetilde{O}\left(n^{2} \mathcal{T}_{\text {ind }}+n^{3}\right)$ |

Table 1.3: Algorithms for (unweighted) matroid intersection. $n$ is the size of the ground set and $\mathcal{T}_{\text {ind }}$ is the time to check if a set is independent.

| Year | Author | Complexity |
| :---: | :---: | :---: |
| 1978 | Fujishige [92] | not stated |
| 1981 | Grotschel, Lovasz, Schrijver[110] | weakly polynomial |
| 1987 | Frank, Tardos [91] | strongly polynomial |
| 1993 | McCormick, Ervolina [190] | $O\left(n^{7} h^{*} \log (n C U)\right)$ |
| 1994 | Wallacher, Zimmermann [266] | $O\left(n^{8} h \log (n C U)\right)$ |
| 1997 | Iwata [117] | $O\left(n^{7} h \log U\right)$ |
| 1998 | Iwata, McCormick, Shigeno [123] | $O\left(n^{4} h \min \left\{\log (n C), n^{2} \log n\right\}\right)$ |
| 1999 | Iwata, McCormick, Shigeno [124] | $O\left(n^{6} h \min \left\{\log (n U), n^{2} \log n\right\}\right)$ |
| 1999 | Fleischer, Iwata, McCormick[89] | $O\left(n^{4} h \min \left\{\log U, n^{2} \log n\right\}\right)$ |
| 1999 | Iwata, McCormick, Shigeno [122] | $O\left(n^{4} h \min \left\{\log C, n^{2} \log n\right\}\right)$ |
| 2000 | Fleischer, Iwata [88] | $O\left(m n^{5} \log (n U) \cdot \mathrm{EO}\right)$ |
| - | Chapter 9 | $O\left(n^{2} \log (n C U) \cdot \mathrm{EO}+n^{3} \log { }^{O(1)}(n C U)\right)$ |

Table 1.5: Algorithms for minimum cost submodular flow with $n$ vertices, maximum cost $C$ and maximum capacity $U$. The factor $h$ is the time for an exchange capacity oracle, $h^{*}$ is the time for a more complicated exchange capacity oracle, and EO is the time for evaluation oracle.

| Year | Author | Complexity |
| :---: | :---: | :---: |
| 1968 | Edmonds [81] | not stated |
| 1975 | Lawler [157] | $O\left(n^{3} \mathcal{T}_{\text {ind }}+n^{4}\right)$ |
| 1981 | Frank [90] | $O\left(n^{3}\left(\mathcal{T}_{\text {circuit }}+n\right)\right)$ |
| 1986 | Brezovec, Cornuéjols, Glover[39] | $O\left(n^{2}\left(\mathcal{T}_{\text {circuit }}+n\right)\right)$ |
| 1995 | Fujishige, Zhang [96] | $O\left(n^{2.5} \log (n M) \cdot \mathcal{T}_{\text {ind }}\right)$ |
| 1995 | Shigeno, Iwata [234] | $O\left(\left(n+\mathcal{T}_{\text {circuit }}\right) n^{1.5} \log (n M)\right)$ |
| - | Chapter 9 | $O\left(n^{2} \log (n M) \mathcal{T}_{\text {ind }}+n^{3} \log O(1)(n M)\right)$ |

Table 1.4: Algorithms for weighted matroid intersection. In addition to the notation in Table 1.3, $\mathcal{T}_{\text {circuit }}$ is the time needed to find a fundamental circuit and $M$ is the maximum weight.

| Year | Author | Complexity |
| :---: | :---: | :---: |
| 1992 | Nesterov, Nemirovsky[210] | $\tilde{O}\left(\sqrt{m}\left(n m^{\omega}+n^{\omega-1} m^{2}\right)\right)$ |
| 2000 | Anstreicher [13] | $\tilde{O}\left((m n)^{1 / 4}\left(n m^{\omega}+n^{\omega-1} m^{2}\right)\right)$ |
| 2003 | Krishnan, Mitchell [152] | $\tilde{O}\left(n\left(n^{\omega}+m^{\omega}+S\right)\right)$ (dual SDP) |
| - | Chapter 9 | $\tilde{O}\left(n\left(n^{2}+m^{\omega}+S\right)\right)$ |

Table 1.6: Algorithms for solving a $m \times m$ SDP with $n$ constraints and $S$ non-zero entries

## - 1.2.2 Linear Programming

 algorithm required either solving $\widetilde{O}(\sqrt{m})$ linear systems or inverting $\widetilde{O}\left((m n)^{1 / 4}\right)$ dense matrices [227, 253]. However, since a breakthrough of Nesterov and Nemirovski in 1994 [211], it has been known that $\widetilde{O}(\sqrt{n})$ iterations in fact suffice. Unfortunately, each iteration of their algorithm requires computing the center of gravity of a polytope, which is a problem that is harder than linear programming. In this thesis, we resolved this 25 -year-old computational question and developed an efficient algorithm that requires solving only $\widetilde{O}(\sqrt{n})$ linear systems (Table 1.7).

## ■ 1.2.3 Maximum Flow Problem

The maximum flow problem has been studied extensively for more than 60 years; it has a wide range of theoretical and practical applications and is widely used as a key subroutine in other algorithms. Until 5 years ago, work on the maximum flow and the more general problems of linear programming and first order convex optimization was typically performed independently. Nevertheless, all known

| Year | Author | Number of Iterations | Nature of iterations |
| :---: | :---: | :---: | :---: |
| 1984 | Karmarkar [136] | $\widetilde{O}(m)$ | Linear system solve |
| 1986 | Renegar $[227]$ | $\widetilde{O}(\sqrt{m})$ | Linear system solve |
| 1989 | Vaidya $[253]$ | $\widetilde{O}\left((m n)^{1 / 4}\right)$ | Expensive linear algebra |
| 1994 | Nesterov and Nemirovskii $[211]$ | $\widetilde{O}(\sqrt{n})$ | Volume computation |
| - | Chapter 6 | $\widetilde{O}(\sqrt{n})$ | $\widetilde{O}(1)$ Linear system solves |

Table 1.7: Algorithms for linear programs with $n$ variables and $m$ constraints

| Year | Author | Complexity | Directed | Weighted |
| :---: | :---: | :---: | :---: | :---: |
| 1998 | Goldberg \& Rao [104] | $\widetilde{O}\left(m \min \left(m^{1 / 2}, n^{2 / 3}\right)\right)$ | Yes | Yes |
| 1998 | Karger [134] | $\widetilde{O}\left(m \sqrt{n} \varepsilon^{-1}\right)$ | No | Yes |
| 2002 | Karger \& Levine [135] | $\widetilde{O}(n F)$ | No | Yes |
| 2011 | Christiano et al. [51] | $\widetilde{O}\left(m n^{1 / 3} \varepsilon^{-11 / 3}\right)$ | No | Yes |
| 2012 | Lee, Rao \& Srivastava [161] | $\widetilde{O}\left(m n^{1 / 3} \varepsilon^{-2 / 3}\right)$ | No | No |
| 2013 | Chapter 12 | $\widetilde{O}\left(m^{1+o(1)} \varepsilon^{-2}\right)$ | No | Yes |
|  | Sherman [233] | $\widetilde{O}\left(m^{10 / 7}\right)$ | Yes | No |
| 2013 | Mądry [183] | $\widetilde{O}(m \sqrt{n})$ | Yes | Yes |
| - | Chapter 6 |  |  |  |

Table 1.8: Algorithms for maximum flow problem with $n$ vertices, $m$ edges and maximum flow value $F$. Results before 1998 are omitted for brevity.
techniques, whether combinatorial or numerical, achieved essentially the same (or worse) running time, $\widetilde{O}\left(m^{1.5}\right)$ or $\widetilde{O}\left(m n^{2 / 3}\right)$ for solving the maximum flow on a weighted graph with $n$ vertices and $m$ edges. Consequently, the maximum flow problem poses a problem for testing the state-of-the-art methods in both convex and combinatorial optimization.

In this thesis, we obtain the first $\widetilde{O}(m \sqrt{n})$ time algorithm for solving the maximum flow problem on directed graphs with $m$ edges and $n$ vertices. This is an improvement over the previous fastest running time of $\widetilde{O}\left(m \min \left\{\sqrt{m}, n^{2 / 3}\right\}\right)$ achieved over 15 years ago by Goldberg and Rao [104] and the previous fastest running times for solving dense directed unit capacity graphs of $O\left(m \min \left\{\sqrt{m}, n^{2 / 3}\right\}\right)$ achieved by Even and Tarjan [85] over 35 years ago and of $\widetilde{O}\left(m^{10 / 7}\right)$ achieved recently by Mądry [183]. For undirected graphs, we give one of the first algorithm with running time $O\left(m^{1+o(1)} / \varepsilon^{2}\right)$ for approximating maximum flow.

## ■ 1.2.4 Philosophical Contributions

Besides the improvements on running time in many classical problems in computer science, the secondary goal of this thesis is to further promote the interaction between continuous optimization and discrete optimization.

For many decades, continuous optimization and discrete (combinatorial) optimization have been studied separately. In continuous optimization, researchers have developed optimal iterative methods under various black-box models, and in discrete optimization, researchers have developed various randomization techniques and data structures. Despite the different approaches, convex problems in both fields in fact share some common difficulties. In this thesis, we identify three important components from both fields and demonstrate the combination of both techniques is indeed useful for solving many fundamental problems in optimization.

Despite decades of research on convex optimization and combinatorial optimization, we still do not have nearly optimal time algorithm for the majority of non-trivial convex problems. We strongly believe that the combination of both fields will help advance future research in this area.

## Preliminaries and Thesis Structure

In this chapter, we first describe the problems we encounter throughout this thesis, then we discuss the structure of this thesis. At the end, we list some linear algebraic notations and some facts we commonly used throughout this thesis.

## ■ 2.1 Problems

## ■ 2.1.1 Black-Box Problems

We say that a set $K \subseteq \mathbb{R}^{n}$ is convex if for all $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{n}$ and all $t \in[0,1]$ it holds that $t \boldsymbol{x}+(1-t) \boldsymbol{y} \in K$. We say that a function $f: K \rightarrow \mathbb{R}$ is convex if $K$ is convex and for all $\boldsymbol{x}, \boldsymbol{y} \in K$ and all $t \in[0,1]$ it holds that $f(t \boldsymbol{x}+(1-t) \boldsymbol{y}) \leq t \cdot f(\boldsymbol{x})+(1-t) \cdot f(\boldsymbol{y})$. In this thesis, we are interested in the following minimization problem

$$
\min _{x \in K} f(\boldsymbol{x})
$$

where both $f$ and $K$ are convex and we call this a convex problem. There are two types of convex problems, black-box and structural problems, depending on if we have full information about the problem.

In black-box problems, the function $f$ and/or the set $K$ are unknown and one can only access them through queries to an oracle. Here, we give a list of black-box problems we are interested in. All of these problems below are very general and cover a wide range of problems in computer science.

- Convex Minimization: Given a convex function $f$ and a (sub)gradient oracle that computes $\nabla f(\boldsymbol{x})$ at any point $\boldsymbol{x}$. The problem is to find a point that minimizes $f$ over $\boldsymbol{x} \in \mathbb{R}^{n}$.
- The Feasibility Problem: Given a convex set $K \subset \mathbb{R}^{n}$ and a separation oracle that, given a point $\boldsymbol{x}$, either outputs that $\boldsymbol{x}$ is in $K$ or outputs a hyperplane separates $\boldsymbol{x}$ and $K$. The problem is to find a point in $K$ or prove that $K$ is almost empty.
- The Intersection Problem: Given two convex sets $K_{1}, K_{2} \subset \mathbb{R}^{n}$ and an optimization oracle that, given a vector $\boldsymbol{c}$, outputs a vector $\boldsymbol{x}_{1}$ that maximizes $\boldsymbol{c}^{\top} \boldsymbol{x}_{1}$ over $K_{1}$ and $\boldsymbol{x}_{2}$ that maximizes $\boldsymbol{c}^{\top} \boldsymbol{x}_{2}$ over $K_{2}$. For any vector $\boldsymbol{d}$, the problem is to find a point $\boldsymbol{x}$ nearly inside $K_{1} \cap K_{2}$ such that it approximately maximizes $\boldsymbol{d}^{\top} \boldsymbol{x}$.
- Submodular Minimization: Given an integer-valued function $f$ defined on all subsets of $n$ elements. $f$ is submodular if it satisfies the diminishing returns, i.e.

$$
f(X \cup\{\boldsymbol{x}\})-f(X) \geq f(Y \cup\{\boldsymbol{x}\})-f(Y)
$$

for any $X, Y$ such that $X \subset Y$ and any $\boldsymbol{x} \notin Y$. We are given an evaluation oracle that computes $f(S)$ for any subset $S$. The problem is to find a subset $S$ that minimizes $f$.

## ■ 2.1.2 Structural Problems

In structural problems, the function $f$ and the set $K$ are given directly by explicit formulas. In this and the next section, we list some structure problems we are interested in this thesis.

- $\ell_{p}$ regression: Given a matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ and a vector $\boldsymbol{b} \in \mathbb{R}^{m}$, the $\ell_{p}$ regression problem requires us to find the minimizer of

$$
\min _{\boldsymbol{x}}\|\boldsymbol{A} \boldsymbol{x}-\boldsymbol{b}\|_{p} .
$$

We are particularly interested in the case $p=2$ because many optimization algorithms need to solve a linear system per iteration. As we will see, techniques for solving $\ell_{2}$ regression problems can be sometimes used or generalized to general convex problems.

- Linear Programming: Given a matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ and vectors $\boldsymbol{b} \in \mathbb{R}^{m}, \boldsymbol{c} \in \mathbb{R}^{n}$, linear programming requires us to find the minimizer of

$$
\min _{A x \geq b} c^{\top} x
$$

This is one of the key problems in computer science. In both theory and practice, many problems can be reformulated as linear programs to take advantage of fast algorithms.

- Semidefinite Programming: Given a symmetric matrices $\boldsymbol{C} \in \mathbb{R}^{m \times m}, \boldsymbol{A}_{i} \in \mathbb{R}^{m \times m}$ for $i=1, \cdots, n$, the semidefinite program requires us to find the minimizer of

$$
\max _{\boldsymbol{X} \succeq \mathbf{0}} \boldsymbol{C} \bullet \boldsymbol{X} \text { s.t. } \boldsymbol{A}_{i} \bullet \boldsymbol{X}=b_{i} .
$$

This is a generalization of linear programming and has a lot of applications in control theory, polynomial minimization, and approximation algorithms.

- Geometric Median (Fermat-Weber problem): Given $n$ points $\left\{\boldsymbol{a}_{i}\right\}_{i=1}^{m}$ in $\mathbb{R}^{n}$, the geometric median problem requires us to find the minimizer of

$$
\min _{\boldsymbol{x}} \sum_{i=1}^{m}\left\|\boldsymbol{a}_{i}-\boldsymbol{x}\right\|_{2} .
$$

This is one of the oldest problems in computation geometry and has numerous applications such as facility location, clustering and statistics.

## ■ 2.1.3 Graph-Related Problems

Here, we define some graph-related problems that are used throughout in this thesis.
We let $G=(V, E)$ denote a graph with $n=|V|$ vertices, $m=|E|$ edges. We let $w_{e} \geq 0$ denote the weight of an edge. For all undirected graphs in this thesis, we assume an arbitrary orientation of the edges to clarify the meaning of vectors $f \in \mathbb{R}^{E}$.

Let $\boldsymbol{W} \in \mathbb{R}^{E \times E}$ denote the diagonal matrices associated with the weights and $\boldsymbol{B} \in \mathbb{R}^{E \times V}$ denote the graphs incidence matrix where for all $e=(a, b) \in E$ we have $\boldsymbol{B}^{\top} \mathbf{1}_{e}=\mathbf{1}_{a}-\mathbf{1}_{b}$. Let $\mathcal{L} \in \mathbb{R}^{V \times V}$ denote the graph Laplacian, i.e. $\mathcal{L} \stackrel{\text { def }}{=} \boldsymbol{B}^{\top} \boldsymbol{W} \boldsymbol{B}$. It is easy to see that

$$
\begin{equation*}
\boldsymbol{x}^{\top} \mathcal{L} \boldsymbol{x}=\sum_{u \sim v} \boldsymbol{w}_{u, v}\left(\boldsymbol{x}_{u}-\boldsymbol{x}_{v}\right)^{2} \tag{2.1}
\end{equation*}
$$

for any $\boldsymbol{x} \in \mathbb{R}^{n}$. Here are three important graph related problems:

Shortest Path Problem: The shortest path problem requires us to find a path $f$ between two vertices that minimizes the sum of the lengths $l$ on its constituent edges. Let $s$ be the starting point and $t$ be the ending point. This problem can be written as

$$
\min _{B^{\top} \boldsymbol{f}=\chi} \sum_{e} \boldsymbol{w}_{e}^{-1}\left|\boldsymbol{f}_{e}\right|
$$

where $\boldsymbol{\chi}=\mathbf{1}_{s}-\mathbf{1}_{t}$ and $\boldsymbol{w}_{e}=1 / \boldsymbol{l}_{e}$. We say $\boldsymbol{f} \in \mathbb{R}^{E}$ meets the demands $\boldsymbol{\chi}$ if $\boldsymbol{B}^{\top} \boldsymbol{f}=\boldsymbol{\chi}$.
Laplacian Systems: Given a graph with weights $\boldsymbol{w}$ and a vector $\chi \in \mathbb{R}^{E}$ with $\sum_{e} \chi_{e}=0$, our goal is to find $\boldsymbol{x}$ such that $\mathcal{L} \boldsymbol{x}=\boldsymbol{\chi}$. Setting $\boldsymbol{f}=\boldsymbol{W} \boldsymbol{B} \boldsymbol{x}$, this problem can be written as

$$
\min _{B^{\top} \boldsymbol{f}=\chi} \sum_{e} \boldsymbol{w}_{e}^{-1} \boldsymbol{f}_{e}^{2} .
$$

Undirected Maximum Flow Problem: Given a graph with capacities $\boldsymbol{w}$, the maximum flow problem requires us to find a flow $\boldsymbol{f}$ that routes as much flow as possible from a source vertex $s$ to a sink vertex $t$ while sending at most $\boldsymbol{w}_{e}$ units of flow over each edge $e$. Rescaling $\boldsymbol{f}$, this problem can be written as

$$
\min _{B^{\top} \boldsymbol{f}=\chi} \max _{e} \boldsymbol{w}_{e}^{-1}\left|\boldsymbol{f}_{e}\right|
$$

where $\boldsymbol{\chi}=\mathbf{1}_{s}-\mathbf{1}_{t}$.
Interestingly, these three problems respectively require us to find a flow that satisfies a certain demand and minimizes $\ell_{1}, \ell_{2}$, or $\ell_{\infty}$ norm.

## ■ 2.2 Thesis Organization

This thesis presents a thorough study of convex optimization and their applications to combinatorial optimization. In figure 2-1, we show the relationship between the chapters. Although the ideas used and the questions addressed in each chapter are all directly or indirectly related, we present the results in each chapter in a modular way so that they may be read in any order. We now discuss the results of each chapter in further details.

## - 2.2.1 Part I: Sparsification

In this part, we study several structural convex problems where the convex functions involve some tall matrices $\boldsymbol{A}$. Our goal is to find general techniques to reduce problems involve tall-and-skinny matrices $\boldsymbol{A}$ to problems involve only square matrices.

### 2.2.1.1 Chapter 3: $\ell_{2}$ Regression In Input Sparsity Time

In this chapter, we study the $\ell_{2}$ regression problem $\min _{\boldsymbol{x}}\|\boldsymbol{A} \boldsymbol{x}-\boldsymbol{b}\|_{2}$ for tall matrices $\boldsymbol{A} \in \mathbb{R}^{m \times n}$. In general, a small, manageable set of rows of $\boldsymbol{A}$ can be randomly selected to approximate a tall-and-skinny data matrix, improving processing time significantly. In particular, if we sample each row with probability proportional to its importance (called leverage score), $O(n \log n)$ rows suffices to approximate $\boldsymbol{A}^{\top} \boldsymbol{A}$ (Lemma 2.3.3). Unfortunately, the previous best approach to compute leverage scores involves solving some linear systems of the form $\boldsymbol{A}^{\top} \boldsymbol{A}$ (Lemma 2.3.4) and hence this is as difficult as the original problem. A simple alternative is to sample rows uniformly at random. While this often works, uniform sampling will eliminate critical row information for many natural instances.

We consider uniform sampling from a new persepective by examining what information it does preserve. Specifically, we show that uniform sampling yields a matrix that, in some sense, well ap-


Figure 2-1: This flow chart shows how the chapters are related between each other. A solid arrow indicates a direct usage of the result and a dotted arrow indicates a vague relation. The gray box indicates the key results in this thesis.
proximates a large fraction of the original matrix. While this weak form of approximation is not enough to solve linear regression directly, it is enough to compute a better approximation. This observation leads to simple iterative row sampling algorithms for matrix approximation that preserves row structure and sparsity at all intermediate steps. If it takes $\mathcal{T}(m, n)$ time to solve a $\ell_{2}$ regression problem with $m$ rows and $n$ columns, our result can be viewed as a reduction showing that

$$
\mathcal{T}(m, n)=\widetilde{O}(\operatorname{nnz}(\boldsymbol{A}))+\widetilde{O}(\mathcal{T}(O(n \log n), n))
$$

### 2.2.1.2 Chapter 4: Linear-Sized Sparsifier In Almost-Linear Time

Naturally, one may ask if $\Omega(n \log n)$ rows is necessary for spectral approximations. In a seminal work of Batson, Spielman, Srivastava, they showed that

Theorem 2.2.1 ([31]). For any matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}$, there is a diagonal matrix $\boldsymbol{D}$ with $O\left(n / \varepsilon^{2}\right)$ non-zeros such that

$$
(1-\varepsilon) \boldsymbol{A}^{\top} \boldsymbol{A} \preceq \boldsymbol{A}^{\top} \boldsymbol{D} \boldsymbol{A} \preceq(1+\varepsilon) \boldsymbol{A}^{\top} \boldsymbol{A} .
$$

However, all previous constructions of linear-sized spectral sparsification require solving $\widetilde{\Omega}(n)$ linear systems. In comparison, the sampling approaches use $O\left(n \log n / \varepsilon^{2}\right)$ rows, but the construction only requires solving $\widetilde{O}(1)$ linear systems. In this chapter, we give a much better trade-off for getting linear-sized spectral sparsifiers. For any integer $\rho$, we present an algorithm that uses $O\left(\rho n / \varepsilon^{2}\right)$ rows via solving $\tilde{O}\left(n^{1 / \rho}\right)$ linear systems. Hence, this gives a smooth trade off between sparsity and running time. A key component in our algorithm is a novel combination of two techniques used in literature for constructing spectral sparsification: random sampling by leverage scores and adaptive constructions based on barrier functions. Similar to the previous chapter, one can view this result as a reduction showing that

$$
\mathcal{T}(m, n)=O\left(\operatorname{nnz}(\boldsymbol{A}) n^{o(1)}\right)+\mathcal{T}(O(n), n)
$$

We believe the term $n^{o(1)}$ can be improved to $\log ^{O(1)} n$ and we left it as an open problem.

### 2.2.1.3 Chapter 4: $\ell_{p}$ Regression and John Ellipsoid

Similar to the $\ell_{2}$ regression, it is known that if we sample each row with certain probability (called $\ell_{p}$ Lewis weight), one can construct a matrix $\widetilde{\boldsymbol{A}}$ such that, for any $\boldsymbol{x},\|\widetilde{\boldsymbol{A}} \boldsymbol{x}\|_{p} \approx\|\boldsymbol{A} \boldsymbol{x}\|_{p}$. This is a useful primitive for solving $\ell_{p}$ regression problem. For $1 \leq p<4$, Cohen and Peng showed how to compute $\ell_{p}$ Lewis weight by solving $\widetilde{O}(1)$ linear systems [56] and they used it to construct an efficient approximation algorithm for $\ell_{p}$ regression. However, for $p \geq 4$, the current best algorithm for computing $\ell_{p}$ Lewis weight involves applying a general algorithm for convex problems.

In this chapter, we give a simple algorithm for $\ell_{p}$ Lewis weight for $p \geq 2$. Similar to [56], the running time of our algorithm is dominated by solving $\widetilde{O}(1)$ linear systems. Then, we discuss the relation between $\ell_{\infty}$ Lewis weight and John ellipsoid, the largest volume ellipsoid inside a given convex set. Finally, we show how to approximate John ellipsoid of symmetric polytopes by solving $\widetilde{O}(1)$ linear systems again.

### 2.2.1.4 Chapter 6: Faster Linear Programming

In this chapter, we show how to use leverage score and John ellipsoid to solve linear programs using only $\widetilde{O}(\sqrt{n} L)$ iterations where $\boldsymbol{A}$ is the constraint matrix of a linear program with $m$ constraints, $n$ variables, and bit complexity $L$. Each iteration of our method consists of solving $\widetilde{O}(1)$ linear systems and additional nearly-linear time computation. Our method improves upon the previous best iteration
bounds by a factor of $\widetilde{\Omega}\left((m / n)^{1 / 4}\right)$ for methods with polynomial time computable iterations [253] and by a factor of $\widetilde{\Omega}\left((m / n)^{1 / 2}\right)$ for methods which solve at most $\widetilde{O}(1)$ linear systems in each iteration achieved over 20 years ago [227]. This result shows that one can "sparsify" the linear program via John ellipsoid and make sure that the complexity of solving a linear program mainly depends on the number of variables $n$ instead of the number of constraints $m$.

Applying our techniques to the maximum flow yields an $\widetilde{O}\left(m \sqrt{n} \log ^{2} U\right)$ time algorithm for solving the maximum flow problem on directed graphs with $m$ edges, $n$ vertices, and capacity ratio $U$. This improves upon the previous fastest running time of $O\left(m \min \left\{m^{1 / 2}, n^{2 / 3}\right\} \log \left(n^{2} / m\right) \log (U)\right)$ achieved over 15 years ago by Goldberg and Rao [104] and improves upon the previous best running times for solving dense directed unit capacity graphs of $O\left(m \min \left\{m^{1 / 2}, n^{2 / 3}\right\}\right)$ achieved by Even and Tarjan [85] over 35 years ago and a running time of $\widetilde{O}\left(m^{10 / 7}\right)$ achieved recently by Mądry [183].

### 2.2.1.5 Chapter 7: Geometric Median In Nearly-Linear Time

From the results above, we see that a wide range of problems can be sparsified and the running time of solving a convex optimization problem mainly depends on the number of variables. In this chapter, we would like to argue it sometimes depends on the number of true "dimensions", which can be much smaller. For example, linear programs can be solved efficiently when the constraint matrix $\boldsymbol{A}$ is an identity matrix. Then, it is natural to ask if we can solve a linear program faster if the constraint matrix of the linear program can be divided into different blocks and all blocks are identity matrices. Unfortunately, we do not have such result yet.

In this chapter, we illustrate this belief via the geometric median problem: given $n$ points in $\mathbb{R}^{d}$ compute a point that minimizes the sum of Euclidean distances to the points. This is one of the oldest non-trivial problems in computational geometry yet despite an abundance of research the previous fastest algorithms for computing a $(1+\varepsilon)$-approximate geometric median were $O\left(d \cdot n^{4 / 3} \varepsilon^{-8 / 3}\right)$ by Chin et. al [50], $\tilde{O}\left(d \exp \left(\varepsilon^{-4} \log \varepsilon^{-1}\right)\right)$ by Badoiu et. al [25], $O\left(n d+\operatorname{poly}\left(d, \varepsilon^{-1}\right)\right.$ by Feldman and Langberg [86], and $O\left((n d)^{O(1)} \log \frac{1}{\varepsilon}\right)$ by Parrilo and Sturmfels [222] and Xue and Ye [271].

Different from all previous approaches, we rely on the fact that the Hessian of the objective function can be approximated by an identity matrix minus a rank 1 matrix, namely, the number of true "dimensions" of this problem is 2 . This allows us to develop a non-standard interior point method that computes a $(1+\varepsilon)$-approximate geometric median in time $O\left(n d \log ^{3} \frac{n}{\varepsilon}\right) .{ }^{1}$ Our result is one of very few cases we are aware of outperforming traditional interior point theory and the only we are aware of using interior point methods to obtain a nearly linear time algorithm for a canonical optimization problem that traditionally requires superlinear time.

## ■ 2.2.2 Part II: Cutting

In this part, we give a faster cutting plane method and show how to apply this cutting plane method efficiently in various problems.

### 2.2.2.1 Chapter 8: Convex Minimization In Nearly-Cubic Time

The central problem we consider in this chapter is the feasibility problem. We are promised that a set $K$ is contained a box of radius $R$ and a separation oracle, that given a point $\boldsymbol{x}$ in time SO, either outputs that $\boldsymbol{x}$ is in $K$ or outputs a separating hyperplane. We wish to either find a point in $K$ or prove that $K$ does not contain a ball of radius $\varepsilon$.

[^2]We show how to solve this problem in $O\left(n \mathrm{SO} \log (n R / \varepsilon)+n^{3} \log ^{O(1)}(n R / \varepsilon)\right)$ time. This is an improvement over the previous best running time of $O\left(n \mathrm{SO} \log (n R / \varepsilon)+n^{\omega+1} \log (n R / \varepsilon)\right)$ for the current best known bound of $\omega \sim 2.37$ [99]. Our key idea for achieving this running time improvement is a new straightforward technique for providing low variance unbiased estimates for changes in leverage scores that we hope will be of independent interest (See Section 8.4.1). We show how to use this technique along with ideas from $[20,257]$ and Chapter 6 to decrease the $\widetilde{O}\left(n^{\omega+1} \log (n R / \varepsilon)\right)$ overhead in the previous fastest algorithm [253].

### 2.2.2.2 Chapter 9: Effective Use of Cutting Plane Method

In this chapter, we provide two techniques for applying our cutting plane method (and cutting plane methods in general) to optimization problems and provide several applications of these techniques.

The first technique concerns reducing the number of dimensions through duality. For many problems, their dual is significantly simpler than itself (primal). We use semidefinite programming as a concrete example to show how to improve upon the running time for finding both primal and dual solution by using the cutting planes maintained by our cutting plane method.

The second technique concerns how to minimize a linear function over the intersection of convex sets using optimization oracle. We analyze a simple potential function which allows us to bypass the typical reduction between separation and optimization to achieve faster running times. This reduction provides an improvement over the reductions used previously in [110]. Moreover, we show how this technique allows us to achieve improved running times for matroid intersection and minimum cost submodular flow.

### 2.2.2.3 Chapter 10: Submodular Minimization In Nearly-Cubic Time

In this chapter, we consider the problem of submodular minimization - a fundamental problem in combinatorial optimization with many diverse applications in theoretical computer science, operations research, machine learning and economics. We show that by considering the interplay between the guarantees of our cutting plane algorithm and the primal-dual structure of submodular minimization, we can achieve improved running times in various settings.

First, we show that a direct application of our method yields an improved weakly polynomial time algorithm for submodular minimization. Then, we present a simple geometric argument that submodular function can be solved with $O\left(n^{3} \log n\right)$ oracle calls but with exponential running time. Finally, we show that by further studying the combinatorial structure of submodular minimization and a modification to our cutting plane algorithm we can obtain a fully improved strongly polynomial time algorithm for submodular minimization.

### 2.2.2.4 Chapter 11: First Order Method vs Cutting Plane Method

Cutting plane methods are often regarded as inefficient in practice because each iteration of these methods is costly and their convergence rate are sometimes independent of how simple the problem is. In this chapter, we obtain a cutting plane method enjoys the following properties

1. Each iteration takes only linear time for simple problems.
2. It converges faster for simpler problems.
3. It is a polynomial time algorithm in the worst case.

We first propose a new method, a combination of gradient descent and the ellipsoid method, that satisfies both properties 1 and 2. This method attains the optimal rate of convergence of Nesterov's accelerated gradient descent [203] and each iteration of this method takes exactly linear time. However, this method is not a polynomial time algorithm if the function is not smooth.

Then, we propose a new framework which leverages several concepts from convex optimization, from standard first-order methods to cutting plane methods. We show how to use our framework to derive a method that satisfies all three properties. We also demonstrate empirically that our new technique compares favorably with state of the art algorithms (such as accelerated gradient descent and BFGS). Therefore, this suggests that cutting plane methods can be as aggressive as first order methods if designed properly.

## - 2.2.3 Part III: Collapsing

In this part, we show how to collapse vertices in the maximum flow problem and the connection Laplacian problem.

### 2.2.3.1 Chapter 12: Undirected MaxFlow In Almost-Linear Time

Flow problems on trees are simple because we often know how to collapse the leaves on trees while maintaining the solution. Using this fact and various routing and optimization techniques, we introduce a new framework for approximately solving flow problems in capacitated, undirected graphs and apply it to provide asymptotically faster algorithms for the maximum $s$ - $t$ flow and maximum concurrent multicommodity flow problems. For graphs with $n$ vertices and $m$ edges, it allows us to find an $\varepsilon$-approximate maximum $s$ - $t$ flow in time $O\left(m^{1+o(1)} \varepsilon^{-2}\right)^{2}$, improving the previous best bound of $\widetilde{O}\left(m n^{1 / 3}\right.$ poly $\left.(1 / \varepsilon)\right)$. Applying the same framework in the multicommodity setting solves a maximum concurrent multicommodity flow problem with $k$ commodities in $O\left(m^{1+o(1)} \varepsilon^{-2} k^{2}\right)$ time, improving the existing bound of $\widetilde{O}\left(m^{4 / 3} \operatorname{poly}\left(k, \varepsilon^{-1}\right)\right)$.

Our algorithms utilize several new technical tools that we believe may be of independent interest:

- We show how to reduce approximate maximum flow and maximum concurrent flow to oblivious routing.
- We define and provide an efficient construction of a new type of flow sparsifier that allowing one to transfer flows from the sparse graph back to the original graph.
- We give the first almost-linear-time construction of an $O\left(m^{o(1)}\right)$-competitive oblivious routing scheme. No previous such algorithms ran in time better than $\widetilde{\Omega}(m n)$. By reducing the running time to almost-linear, our work provides a powerful new primitive for constructing very fast graph algorithms.


### 2.2.3.2 Chapter 13: Connection Laplacian In Nearly-Linear Time

The idea of collapsing can be extended to general convex problems as follows: Given a convex function $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$. For any $\boldsymbol{x} \in \mathbb{R}^{n}$, we write $\boldsymbol{x}=(\boldsymbol{y}, \boldsymbol{z})$ where $\boldsymbol{y}$ is the first $n / 2$ variables and $\boldsymbol{z}$ is the last $n / 2$ variables. We define the collapse of $f$ by

$$
g(\boldsymbol{y})=\min _{z \in \mathbb{R}^{n / 2}} f(\boldsymbol{x}, \boldsymbol{z}) .
$$

[^3]

Figure 2-2: Connections between interior point methods, cutting plane methods, first-order methods, learning algorithms for bandit problems, and sampling algorithms for convex sets.

In general, a convex function is difficult to minimize if some variables are intimately related to each other. If not, we can simply solve the problem coordinate-wise. However, if some variables are related and if we can solve the half of the these coordinates, the minimization problem often becomes trivial. Therefore, if we pick half of the variables in random as $\boldsymbol{y}, g$ should be easy to compute and this effectively reduce the dimension by half.

In this chapter, we show this is indeed true for connection Laplacian. As a result, we give the first nearly-linear time algorithms for solving systems of equations in connection Laplacians - a generalization of Laplacian matrices that arise in many problems in image and signal processing.

We also prove that every connection Laplacian has a linear sized approximate inverse. This is a LU factorization with a linear number of non-zero entries that is a spectral approximation of the original matrix. Using such a factorization, one can solve systems of equations in a connection Laplacian in linear time. Such a factorization was unknown even for ordinary graph Laplacians.

## ■ 2.2.4 Part IV: Geometry

In the end of this thesis, we would like to point out a large set of connections between interior point methods, cutting plane methods, first-order methods, learning algorithms for bandit problems, and sampling algorithms for convex sets as they relate to important problems across computer science, operation research, convex optimization and machine learning; see Figure 2-2 for a diagram of work that suggests this relation. One key technique spans all these algorithms is the geometry of convex set. Since these algorithms are widely used across various disciplines, we believe these connections can lead to breakthroughs in different areas. As an example, we only give an application to the sampling problem.

### 2.2.4.1 Chapter 13: Sampling In Sub-Quadratic Steps

The problems of sampling random points in convex set and computing volume of convex sets is a central problem in Markov chain Monte Carlo and leads to developments in geometric random walks. This problem has been studied for twenty years and the best algorithms take $\tilde{O}(n m)$ iterations for explicit polytopes. We introduce the geodesic walk for sampling Riemannian manifolds and apply it to the problem of generating uniform random points from polytopes. The walk is a discrete-time simulation of a stochastic differential equation (SDE) on the Riemannian manifold equipped with the metric induced by the Hessian of a convex function used for solving linear programming; each step is the solution of an ordinary differential equation (ODE). The resulting sampling algorithm for polytopes mixes in $\tilde{O}\left(m n^{\frac{3}{4}}\right)$ steps. This is the first walk that breaks the quadratic barrier for mixing in high dimension. We also show that each step of the geodesic walk (solving an ODE) can be implemented efficiently, thus obtaining the current fastest complexity for sampling polytopes.

Advances in algorithms often come from insights from different areas; I believe these new connections will make multifaceted impacts on computation.

## ■ 2.3 Notation

Here, we define some notations that are used throughout in this thesis.
Variables: We use bold lowercase, e.g $\boldsymbol{x}$, to denote a vector and bold uppercase, e.g. $\boldsymbol{A}$, to denote a matrix. For integers $z \in \mathbb{Z}$ we use $[z] \subseteq \mathbb{Z}$ to denote the set of integers from 1 to $z$. We let $\mathbf{1}_{i}$ denote the vector that has value 1 in coordinate $i$ and is 0 elsewhere. We let 1 denote the vector that has value 1 in all coordinates. We use $n n z(\boldsymbol{x})$ or $\operatorname{nnz}(\boldsymbol{A})$ to denote the number of nonzero entries in a vector or a matrix respectively. For a vector, $\boldsymbol{x}$, we let $\|\boldsymbol{x}\|_{M} \stackrel{\text { def }}{=} \sqrt{\boldsymbol{x}^{\top} \boldsymbol{M} \boldsymbol{x}}$. We use $\mathbb{R}_{>0}^{m}$ to denote the vectors in $\mathbb{R}^{m}$ where each coordinate is positive.

Vector Operations: We frequently apply scalar operations to vectors with the interpretation that these operations should be applied coordinate-wise. For example, for vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{n}$ we let $\boldsymbol{x} / \boldsymbol{y} \in \mathbb{R}^{n}$ with $[\boldsymbol{x} / \boldsymbol{y}]_{i} \xlongequal{\text { def }}\left(x_{i} / y_{i}\right)$ and $\log (\boldsymbol{x}) \in \mathbb{R}^{n}$ with $[\log (\boldsymbol{x})]_{i}=\log \left(x_{i}\right)$ for all $i \in[n]$.

Spectral Approximations: We call a symmetric matrix $\boldsymbol{A} \in \mathbb{R}^{n \times n}$ positive semidefinite (PSD) if $\boldsymbol{x}^{\top} \boldsymbol{A} \boldsymbol{x} \geq 0$ for all $\boldsymbol{x} \in \mathbb{R}^{n}$ and we call $\boldsymbol{A}$ positive definite ( PD ) if $\boldsymbol{x}^{\top} \boldsymbol{A} \boldsymbol{x}>0$ for all $\boldsymbol{x} \in \mathbb{R}^{n}$. For symmetric matrices $\boldsymbol{N}, \boldsymbol{M} \in \mathbb{R}^{n \times n}$, we write $\boldsymbol{N} \preceq \boldsymbol{M}$ to denote that $\boldsymbol{x}^{\top} \boldsymbol{N} \boldsymbol{x} \leq \boldsymbol{x}^{\top} \boldsymbol{M} \boldsymbol{x}$ for all $\boldsymbol{x} \in \mathbb{R}^{n}$, and we define $\boldsymbol{N} \succeq \boldsymbol{M}, \boldsymbol{N} \prec \boldsymbol{M}$ and $\boldsymbol{N} \succ \boldsymbol{M}$ analogously. For symmetric matrices, we call $\boldsymbol{B}$ is a $\varepsilon$-spectral approximation of $\boldsymbol{A}$ if $(1-\varepsilon) \boldsymbol{A} \preceq \boldsymbol{B} \preceq(1+\varepsilon) \boldsymbol{A}$. We denote this by $\boldsymbol{A} \approx_{\varepsilon} \boldsymbol{B}$. In chapter 13, we instead define it as $e^{-\varepsilon} \boldsymbol{A} \preceq \boldsymbol{B} \preceq e^{\varepsilon} \boldsymbol{A}$ for computational simplicity. For nonsymmetric matrices, we call $\boldsymbol{B}$ is a $\varepsilon$-spectral approximation of $\boldsymbol{A}$ if $\boldsymbol{B}^{\top} \boldsymbol{B}$ is a $\varepsilon$-spectral approximation of $\boldsymbol{A}^{\top} \boldsymbol{A}$.

Matrix Operations: For $\boldsymbol{A}, \boldsymbol{B} \in \mathbb{R}^{n \times m}$, we let $\boldsymbol{A} \circ \boldsymbol{B}$ denote the Schur product, i.e. $[\boldsymbol{A} \circ \boldsymbol{B}]_{i j} \stackrel{\text { def }}{=}$ $\boldsymbol{A}_{i j} \cdot \boldsymbol{B}_{i j}$ for all $i \in[n]$ and $j \in[m]$, and we let $\boldsymbol{A}^{(2)} \stackrel{\text { def }}{=} \boldsymbol{A} \circ \boldsymbol{A}$. For any norm $\|\cdot\|$ and matrix $\boldsymbol{M}$, the operator norm of $\boldsymbol{M}$ is defined by $\|\boldsymbol{M}\|=\sup _{\|\boldsymbol{x}\|=1}\|\boldsymbol{M} \boldsymbol{x}\|$. For any two matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ of equal dimensions, let the dot product of $\boldsymbol{A}$ and $\boldsymbol{B}$ defined by $\boldsymbol{A} \bullet \boldsymbol{B}=\operatorname{Tr}\left(\boldsymbol{A}^{\top} \boldsymbol{B}\right)$. For any matrix $\boldsymbol{M}$, the Frobenius norm of $\boldsymbol{M}$ is defined by $\|\boldsymbol{M}\|_{F}=\sqrt{\sum_{i, j} \boldsymbol{M}_{i j}^{2}}=\sqrt{\boldsymbol{M} \bullet \boldsymbol{M}}$.

Spectral Theory: For $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ with rank $r$, the singular value decomposition $\boldsymbol{A}$ is given by $\boldsymbol{U} \boldsymbol{\Sigma} \boldsymbol{V}^{\top}$ where $\boldsymbol{U} \in \mathbb{R}^{m \times r}$ and $\boldsymbol{V} \in \mathbb{R}^{n \times r}$ have orthonormal columns and $\boldsymbol{\Sigma} \in \mathbb{R}^{r \times r}$ is diagonal. We let $\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+}$denote the Moore-Penrose pseudoinverse of $\boldsymbol{A}^{\top} \boldsymbol{A}$, namely, $\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+}=\boldsymbol{V}\left(\boldsymbol{\Sigma}^{-1}\right)^{2} \boldsymbol{V}^{\top}$. For symmetric $\boldsymbol{A}$, the singular value decomposition $\boldsymbol{A}$ becomes $\boldsymbol{U} \boldsymbol{\Sigma} \boldsymbol{U}^{\top}$ where the diagonal entries of $\boldsymbol{\Sigma}$
are the eigenvalues of $\boldsymbol{A}$. We let $\lambda_{\max }(\boldsymbol{A})$ and $\lambda_{\min }(\boldsymbol{A})$ be the maximum and minimum eigenvalues of $\boldsymbol{A}$. The condition number of matrix $\boldsymbol{A}$ is defined by $\lambda_{\max }(\boldsymbol{A}) / \lambda_{\min }(\boldsymbol{A})$.

Diagonal Matrices: For $\boldsymbol{A} \in \mathbb{R}^{n \times n}$ we let $\operatorname{diag}(\boldsymbol{A}) \in \mathbb{R}^{n}$ denote the vector such that $\boldsymbol{\operatorname { d i a g }}(\boldsymbol{A})_{i}=$ $\boldsymbol{A}_{i i}$ for all $i \in[n]$. For $\boldsymbol{x} \in \mathbb{R}^{n}$ we let $\operatorname{Diag}(\boldsymbol{x}) \in \mathbb{R}^{n \times n}$ be the diagonal matrix such that $\boldsymbol{\operatorname { D i a g }}(\boldsymbol{\operatorname { D i a g }}(\boldsymbol{x}))=$ $\boldsymbol{x}$. For $\boldsymbol{A} \in \mathbb{R}^{n \times n}$ we let $\operatorname{Diag}(\boldsymbol{A})$ be the diagonal matrix such that $\operatorname{diag}(\operatorname{Diag}(\boldsymbol{A}))=\operatorname{diag}(\boldsymbol{A})$. For $\boldsymbol{x} \in \mathbb{R}^{n}$ when the meaning is clear from context we let $\boldsymbol{X} \in \mathbb{R}^{n \times n}$ denote $\boldsymbol{X} \xlongequal{\text { def }} \boldsymbol{\operatorname { D i a g }}(\boldsymbol{x})$.

Calculus: For $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ differentiable at $\boldsymbol{x} \in \mathbb{R}^{n}$, we let $\nabla f(\boldsymbol{x}) \in \mathbb{R}^{n}$ denote the gradient of $f$ at $\boldsymbol{x}$, i.e. $[\nabla f(\boldsymbol{x})]_{i}=\frac{\partial}{\partial x_{i}} f(\boldsymbol{x})$ for all $i \in[n]$. For $f \in \mathbb{R}^{n} \rightarrow \mathbb{R}$ twice differentiable at $\boldsymbol{x} \in \mathbb{R}^{n}$, we let $\nabla^{2} f(\boldsymbol{x})$ denote the Hessian of $f$ at $\boldsymbol{x}$, i.e. $[\nabla f(\boldsymbol{x})]_{i j}=\frac{\partial^{2}}{\partial x_{i} \partial x_{j}} f(\boldsymbol{x})$ for all $i, j \in[n]$. Sometimes, we will consider functions of two vectors, $g: \mathbb{R}^{n_{1} \times n_{2}} \rightarrow \mathbb{R}$, and wish to compute the gradient and Hessian of $g$ restricted to one of the two vectors. For $\boldsymbol{x} \in \mathbb{R}^{n_{1}}$ and $\boldsymbol{y} \in \mathbb{R}^{n_{2}}$ we let $\nabla_{x} \boldsymbol{g}(\boldsymbol{a}, \boldsymbol{b}) \in \mathbb{R}^{n_{1}}$ denote the gradient of $\boldsymbol{g}$ for fixed $\boldsymbol{y}$ at point $(\boldsymbol{a}, \boldsymbol{b}) \in \mathbb{R}^{n_{1} \times n_{2}}$. We define $\nabla_{\boldsymbol{y}}, \nabla_{\boldsymbol{x} \boldsymbol{x}}^{2}$, and $\nabla_{y y}^{2}$ similarly. For $h: \mathbb{R}^{n} \rightarrow \mathbb{R}^{m}$ differentiable at $\boldsymbol{x} \in \mathbb{R}^{n}$ we let $\boldsymbol{J}(\boldsymbol{h}(\boldsymbol{x})) \in \mathbb{R}^{m \times n}$ denote the Jacobian of $\boldsymbol{h}$ at $\boldsymbol{x}$ where for all $i \in[m]$ and $j \in[n]$ we let $[\boldsymbol{J}(\boldsymbol{h}(\boldsymbol{x}))]_{i j} \stackrel{\text { def }}{=} \frac{\partial}{\partial x_{j}} h(\boldsymbol{x})_{i}$. For functions of multiple vectors we use subscripts, e.g. $\boldsymbol{J}_{\boldsymbol{x}}$, to denote the Jacobian of the function restricted to the $\boldsymbol{x}$ variable.

Sets: We call a set $U$ is symmetric if $\boldsymbol{x} \in \mathbb{R}^{n} \Leftrightarrow-\boldsymbol{x} \in \mathbb{R}^{n}$. For any $\alpha>0$ and set $U \subseteq \mathbb{R}^{n}$ we let $\alpha U \stackrel{\text { def }}{=}\left\{\boldsymbol{x} \in \mathbb{R}^{n} \mid \alpha^{-1} \boldsymbol{x} \in U\right\}$. For any $p \in[1, \infty]$ and $r>0$ we refer to the set $B_{p}(r) \stackrel{\text { def }}{=}\left\{\boldsymbol{x} \in \mathbb{R}^{n} \mid\|\boldsymbol{x}\|_{p} \leq\right.$ $r\}$ as the $\ell_{p}$ ball of radius $r$. For brevity, we refer to $B_{2}(r)$ as a ball of radius $r$ and $B_{\infty}(r)$ as a box of radius $r$.

Matrix Multiplication Constant: We let $\omega<2.3728639$ [100] denote the matrix multiplication constant. It is known that the product a $m \times n$ matrix with a $n \times r$ matrix can be computed in $m n r \min (m, n, r)^{\omega-3+o(1)}$ time and the inverse of a $n \times n$ matrix can be computed in $n^{\omega+o(1)}$ time. For simplicity, we always write $\omega+o(1)$ as $\omega$.

Running Time: Unless mentioned specifically, we will disregard the issue of rounding error if the algorithm is obviously stable. Most of the algorithms we proposed are randomized algorithms that succeeds with high probability and the running time is often the expected running time. For any function $f$, we write $\widetilde{O}(f) \triangleq O\left(f \cdot \log ^{O(1)} f\right)$. Suppose $m$ is the length of the input and $n$ is the number of variables, we ssay that an algorithm runs in nearly-linear time if its running time is $\widetilde{O}(m)$, runs in almost linear time if its running time is $m^{1+o(1)}$, and runs in input sparsity time if its running time is $\widetilde{O}\left(m+n^{O(1)}\right)$.

## - 2.3.1 Sparsification and Leverage Score

In this section, we define spectral approximation and leverage score. They are two key concepts in this thesis, especially in Part I. First, we define leverage score which is a measure of importance. Given a matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}$. The leverage score of the $i^{\text {th }}$ row $\boldsymbol{a}_{i}$ of $\boldsymbol{A}$ is:

$$
\boldsymbol{\sigma}_{i}(\boldsymbol{A}) \stackrel{\text { def }}{=} \boldsymbol{a}_{i}^{\top}\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+} \boldsymbol{a}_{i}
$$

We also define the related cross leverage score as $\boldsymbol{\sigma}_{i j}(\boldsymbol{A}) \stackrel{\text { def }}{=} \boldsymbol{a}_{i}^{\top}\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+} \boldsymbol{a}_{j}$. Let $\boldsymbol{\sigma}(\boldsymbol{A})$ be a vector containing $\boldsymbol{A}$ 's $m$ leverage scores. $\boldsymbol{\sigma}(\boldsymbol{A})$ is the diagonal of $\boldsymbol{A}\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+} \boldsymbol{A}^{\top}$, which is a projection matrix. Thus, $\boldsymbol{\sigma}_{i}(\boldsymbol{A})=\mathbf{1}_{i}^{\top} \boldsymbol{A}\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+} \boldsymbol{A}^{\top} \mathbf{1}_{i} \leq 1$. Furthermore, since $\boldsymbol{A}\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+} \boldsymbol{A}^{\top}$ is a projection matrix, the sum of $\boldsymbol{A}$ 's leverage scores is equal to the matrix's rank:

$$
\begin{equation*}
\sum_{i=1}^{m} \boldsymbol{\sigma}_{i}(\boldsymbol{A})=\operatorname{Tr}\left(\boldsymbol{A}\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+} \boldsymbol{A}^{\top}\right)=\operatorname{rank}\left(\boldsymbol{A}\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+} \boldsymbol{A}^{\top}\right)=\operatorname{rank}(\boldsymbol{A}) \leq n \tag{2.2}
\end{equation*}
$$

A row's leverage score measures how important it is in composing the row space of $\boldsymbol{A}$. If a row has a component orthogonal to all other rows, its leverage score is 1 . Removing it would decrease the rank of $\boldsymbol{A}$, completely changing its row space. The coherence of $\boldsymbol{A}$ is $\|\boldsymbol{\sigma}(\boldsymbol{A})\|_{\infty}$. If $\boldsymbol{A}$ has low coherence, no particular row is especially important. If $\boldsymbol{A}$ has high coherence, it contains at least one row whose removal would significantly affect the composition of $\boldsymbol{A}$ 's row space. Two characterizations that helps with this intuition follows:

Lemma 2.3.1. For all $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ and $i \in[m]$ we have that

$$
\boldsymbol{\sigma}_{i}(\boldsymbol{A})=\min _{\boldsymbol{A}^{\top} \boldsymbol{x}=\boldsymbol{a}_{i}}\|\boldsymbol{x}\|_{2}^{2}
$$

where $\boldsymbol{a}_{i}$ is the $i^{t h}$ row of $\boldsymbol{A}$. Let $\boldsymbol{x}_{i}$ denote the optimal $\boldsymbol{x}$ for $\boldsymbol{a}_{i}$. The $j^{\text {th }}$ entry of $\boldsymbol{x}_{i}$ is given by $\boldsymbol{\sigma}_{i j}(\boldsymbol{A})$.
Lemma 2.3.2. For all $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ and $i \in[m]$, we have that $\boldsymbol{\sigma}_{i}(\boldsymbol{A})$ is the smallest $t$ such that

$$
\begin{equation*}
\boldsymbol{a}_{i} \boldsymbol{a}_{i}^{\top} \preceq t \cdot \boldsymbol{A}^{\top} \boldsymbol{A} . \tag{2.3}
\end{equation*}
$$

Sampling rows from $\boldsymbol{A}$ according to their exact leverage scores gives a spectral approximation for $\boldsymbol{A}$ with high probability. Sampling by leverage score overestimates also suffices. The following lemma is a corollary of matrix concentration result [251, Cor 5.2] and the inequality (2.3).

Lemma 2.3.3. Given an error parameter $0<\varepsilon<1$ and a vector $\boldsymbol{u}$ of leverage score overestimates, i.e., $\boldsymbol{\sigma}_{i}(\boldsymbol{A}) \leq \boldsymbol{u}_{i}$ for all $i$. For any large enough universal constant $c$, we define a probability distribution $\boldsymbol{p}_{i}=\min \left\{1, c \varepsilon^{-2} \boldsymbol{u}_{i} \log n\right\}$. Let $\boldsymbol{S}$ be the diagonal matrix with independently chosen entries. $\boldsymbol{S}_{i i}=\frac{1}{\sqrt{\boldsymbol{p}_{i}}}$ with probability $\boldsymbol{p}_{i}$ and 0 otherwise. Then, with probability at least $1-n^{-\Theta(c)}$, we have

$$
(1-\varepsilon) \boldsymbol{A}^{\top} \boldsymbol{A} \preceq \boldsymbol{A}^{\top} \boldsymbol{S}^{2} \boldsymbol{A} \preceq(1+\varepsilon) \boldsymbol{A}^{\top} \boldsymbol{A} .
$$

Although (2.2) shows that $\|\boldsymbol{\sigma}\|_{1}$ is small, computing $\boldsymbol{\sigma}$ exactly is too expensive for many purposes. In [242], they showed that we can compute leverage scores, $\boldsymbol{\sigma}$, approximately by solving only polylogarithmically many regression problems. This result uses the fact that

$$
\boldsymbol{\sigma}_{i}(\boldsymbol{A})=\left\|\boldsymbol{A}\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+} \boldsymbol{A}^{\top} \mathbf{1}_{i}\right\|_{2}
$$

and that by the Johnson-Lindenstrauss Lemma these lengths are persevered up to multiplicative error if we project these vectors onto certain random low dimensional subspaces.

Lemma 2.3.4 ([242]). Let $\mathcal{T}$ be the time to apply $\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+}$to an arbitrary vector. Then, for any $0<\theta<1$, it is possible to compute an estimate of $\sigma(\boldsymbol{A}), \boldsymbol{\sigma}^{(a p x)}$, in time $O\left(\theta^{-1}(\mathcal{T}+\mathrm{nnz}(\boldsymbol{A}))\right)$ such that w.h.p. in n, for all $i$,

$$
d^{-\max (\sqrt{\theta / \log n}, \theta)} \boldsymbol{\sigma}_{i}(\boldsymbol{A}) \leq \boldsymbol{\sigma}_{i}^{(a p x)} \leq d^{\max (\sqrt{\theta / \log n}, \theta)} \boldsymbol{\sigma}_{i}(\boldsymbol{A}) .
$$

Setting $\theta=\frac{\varepsilon^{2}}{\log n}$ gives $(1+\varepsilon)$ factor leverage score approximations in $\widetilde{O}\left((\mathcal{T}+\operatorname{nnz}(\boldsymbol{A})) / \varepsilon^{2}\right)$ time. Therefore, approximating leverage scores is as simple as solving $\widetilde{O}(1)$ linear systems.

## - 2.3.2 Separation Oracle

In this section, we define separation oracles. They are the key concepts in Part II. Our definitions are possibly non-standard and chosen to handle the different settings that occur in this part.

Definition 2.3.5 (Separation Oracle for a Set). Given a set $K \subset \mathbb{R}^{n}$ and $\delta \geq 0$, a $\delta$-separation oracle for $K$ is a function on $\mathbb{R}^{n}$ such that for any input $\boldsymbol{x} \in \mathbb{R}^{n}$, it either outputs "successful" or a half space of the form $H=\left\{\boldsymbol{z}: \boldsymbol{c}^{T} \boldsymbol{z} \leq \boldsymbol{c}^{T} \boldsymbol{x}+b\right\} \supseteq K$ with $b \leq \delta\|\boldsymbol{c}\|_{2}$ and $\boldsymbol{c} \neq \mathbf{0}$. We let $\mathrm{SO}_{\delta}(K)$ be the time complexity of this oracle.

The parameter $\delta$ indicates the accuracy of the oracle. For brevity we refer to a 0 -separation oracle for a set as just a separation oracle.

Note that in Definition 2.3.5 we do not assume that $K$ is convex. However, we remark that there is a separation oracle for a set if and only if it is convex and that there is a $\delta$ separation oracle if and only if the set is close to convex in some sense.

Definition 2.3.6 (Separation Oracle for a Function). For any function $f, \eta \geq 0$ and $\delta \geq 0$, a $(\eta, \delta)$ separation oracle on a set $\Gamma$ for $f$ is a function on $\mathbb{R}^{n}$ such that for any input $\boldsymbol{x} \in \Gamma$, it either asserts $f(\boldsymbol{x}) \leq \min _{\boldsymbol{y} \in \Gamma} f(\boldsymbol{y})+\eta$ or outputs a half space $H$ such that

$$
\begin{equation*}
\{\boldsymbol{z} \in \Gamma: f(\boldsymbol{z}) \leq f(\boldsymbol{x})\} \subset H \stackrel{\text { def }}{=}\left\{\boldsymbol{z}: \boldsymbol{c}^{T} \boldsymbol{z} \leq \boldsymbol{c}^{T} \boldsymbol{x}+b\right\} \tag{2.4}
\end{equation*}
$$

with $b \leq \delta\|\boldsymbol{c}\|$ and $\boldsymbol{c} \neq \mathbf{0}$. We let $\mathrm{SO}_{\eta, \delta}(f)$ be the time complexity of this oracle.

## - 2.3.3 Other Basic Facts

Here, we state some lemmas that are used throughout this thesis.
Lemma 2.3.7 (Sherman-Morrison Formula, [193, Thm 3]). Let $\boldsymbol{A} \in \mathbb{R}^{n \times n}$, and $\boldsymbol{u}, \boldsymbol{v} \perp \operatorname{ker}(\boldsymbol{A})$. Suppose that $1+\boldsymbol{v}^{\top} \boldsymbol{A}^{+} \boldsymbol{u} \neq 0$. Then, it holds that

$$
\left(\boldsymbol{A}+\boldsymbol{u} \boldsymbol{v}^{\top}\right)^{+}=\boldsymbol{A}^{+}-\frac{\boldsymbol{A}^{+} \boldsymbol{u} \boldsymbol{v}^{T} \boldsymbol{A}^{+}}{1+\boldsymbol{v}^{T} \boldsymbol{A}^{+} \boldsymbol{u}}
$$

Lemma 2.3.8. Let $\boldsymbol{B}_{t}$ be a differentiable family of invertible matrix. Then, we have

1. $\frac{d}{d t} \ln \operatorname{det}\left(\boldsymbol{B}_{t}\right)=\operatorname{Tr}\left(\boldsymbol{B}_{t}^{-1} \frac{d}{d t} \boldsymbol{B}_{t}\right)$.
2. $\frac{d}{d t} \boldsymbol{B}_{t}^{-1}=-\boldsymbol{B}_{t}^{-1}\left(\frac{d}{d t} \boldsymbol{B}_{t}\right) \boldsymbol{B}_{t}^{-1}$.

Lemma 2.3.9. Let $\boldsymbol{f}: \mathbb{R}^{n+m} \rightarrow \mathbb{R}^{m}$ be a continuously differentiable function. Suppose that $f(\boldsymbol{x}, \boldsymbol{y})=$ $\mathbf{0}$ for some $\boldsymbol{x} \in \mathbb{R}^{m}$ and $\boldsymbol{y} \in \mathbb{R}^{n}$ and that $\frac{\partial \boldsymbol{f}}{\partial \boldsymbol{y}}(\boldsymbol{x}, \boldsymbol{y})$ is invertible. Then, there is an open set $U$ containing $\boldsymbol{x}$, an open set $V$ containing $\boldsymbol{y}$, and a unique continuously differentiable function $\boldsymbol{g}: U \rightarrow V$ such that $f(\boldsymbol{x}, \boldsymbol{g}(\boldsymbol{x}))=\mathbf{0}$. Also, we have that

$$
\frac{\partial \boldsymbol{g}}{\partial \boldsymbol{x}}(\boldsymbol{x})=-\left(\frac{\partial \boldsymbol{f}}{\partial \boldsymbol{y}}(\boldsymbol{x}, \boldsymbol{g}(\boldsymbol{x}))\right)^{-1} \frac{\partial \boldsymbol{f}}{\partial \boldsymbol{x}}(\boldsymbol{x}, \boldsymbol{g}(\boldsymbol{x}))
$$

Lemma 2.3.10 (Matrix Chernoff Bound, [251, Cor 5.2]). Let $\left\{\boldsymbol{X}_{k}\right\}$ be a finite sequence of independent, random, and self-adjoint matrices with dimension $n$. Assume that each random matrix satisfies $\boldsymbol{X}_{k} \succeq$ $\mathbf{0}$, and $\lambda_{\max }\left(\boldsymbol{X}_{k}\right) \leq R$. Let $\mu_{\max } \geq \lambda_{\max }\left(\sum_{k} \mathbf{E} \boldsymbol{X}_{k}\right)$ and $\mu_{\min } \leq \lambda_{\min }\left(\sum_{k} \mathbf{E} \boldsymbol{X}_{k}\right)$. Then, it holds that

$$
\begin{aligned}
& \operatorname{Pr}\left[\lambda_{\max }\left(\sum_{k} \boldsymbol{X}_{k}\right) \geq(1+\delta) \mu_{\max }\right] \leq n \cdot\left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{\mu_{\max } / R} \text { for any } \delta \geq 0, \text { and } \\
& \quad \operatorname{Pr}\left[\lambda_{\min }\left(\sum_{k} \boldsymbol{X}_{k}\right) \leq(1-\delta) \mu_{\min }\right] \leq n \cdot\left(\frac{e^{-\delta}}{(1-\delta)^{1-\delta}}\right)^{\mu_{\min } / R} \quad \text { for any } \delta \in[0,1]
\end{aligned}
$$

Lemma 2.3.11 (Lieb Thirring Inequality, [171]). Let $\boldsymbol{A}$ and $\boldsymbol{B}$ be positive definite matrices and $q \geq 1$. Then it holds that

$$
\operatorname{Tr}(\boldsymbol{B} \boldsymbol{A} \boldsymbol{B})^{q} \leq \operatorname{Tr}\left(\boldsymbol{B}^{q} \boldsymbol{A}^{q} \boldsymbol{B}^{q}\right)
$$

Theorem 2.3.12 (Simple Constrained Minimization for Twice Differentiable Function [206]). Let $\boldsymbol{H}$ be a positive definite matrix, $K$ be a convex set and $L \geq \mu \geq 0$. Let $f(\boldsymbol{x}): K \rightarrow \mathbb{R}$ be a twice differentiable function such that $\mu \boldsymbol{H} \preceq \nabla^{2} f(\boldsymbol{x}) \preceq L \boldsymbol{H}$ for all $\boldsymbol{x} \in K$. If for some $\boldsymbol{x}^{(0)} \in K$ and all $k \geq 0$, we apply the update rule

$$
\boldsymbol{x}^{(k+1)}=\underset{\boldsymbol{x} \in K}{\arg \min }\left\langle\nabla f\left(\boldsymbol{x}^{(k)}\right), \boldsymbol{x}-\boldsymbol{x}^{(k)}\right\rangle+\frac{L}{2}\left\|\boldsymbol{x}-\boldsymbol{x}^{(k)}\right\|_{\boldsymbol{H}}^{2}
$$

then for all $k \geq 0$ we have

$$
\left\|\boldsymbol{x}^{(k)}-\boldsymbol{x}^{*}\right\|_{\boldsymbol{H}}^{2} \leq\left(1-\frac{\mu}{L}\right)^{k}\left\|\boldsymbol{x}^{(0)}-\boldsymbol{x}^{*}\right\|_{\boldsymbol{H}}^{2}
$$

Part I

## Sparsification



# $\ell_{2}$ Regression In Input Sparsity Time 

## - 3.1 Introduction

Many convex optimization problems can be written in the form of

$$
\min _{\boldsymbol{x} \in \mathbb{R}^{n}} \sum_{i=1}^{m} f_{i}\left(\boldsymbol{a}_{i}^{\top} \boldsymbol{x}\right)
$$

where $f_{i}$ are 1 dimension convex functions and $\boldsymbol{a}_{i} \in \mathbb{R}^{n}$. The goal of this part is to show that the running time for solving this type of problems, in general, should depends on $m$ linearly.

In this chapter, we first study the simplest case, $f_{i}$ being quadratic functions. In this case, the problem becomes the $\ell_{2}$ regression problem $\min _{\boldsymbol{x}}\|\boldsymbol{A} \boldsymbol{x}-\boldsymbol{b}\|$. Currently, there are two types of methods on this problem for the case $m \gg n$.

The first type of methods rely on spectral approximation. For a tall, narrow data matrix $\boldsymbol{A}$, these methods find a shorter approximate data matrix, $\widetilde{\boldsymbol{A}}$, such that, for all vectors $\boldsymbol{x},\|\widetilde{\boldsymbol{A}} \boldsymbol{x}\|_{2} \approx\|\boldsymbol{A} \boldsymbol{x}\|_{2}$. All of which rely on variations of Johnson-Lindenstrauss random projections for constructing $\widetilde{\boldsymbol{A}}$ [53, 185, 200, 170]. Such random projections do not maintain sparsity and row structure of $\boldsymbol{A}$. The second type of methods rely on uniform (or non-adaptive) sampling, such as stochastic descent. These methods works only for low coherence data and hence are less useful for theory development.

By re-examining uniform sampling, we give spectral approximation algorithms that avoid random projection entirely. Our methods are the first to match state-of-the-art runtimes while preserving row structure and sparsity in all matrix operations. Since it preserves row structure, this allows us to take advantage of the problem structure. Let $\mathcal{T}(m, n)$ be the time needed to solve a $\ell_{2}$ regression problem $\|\boldsymbol{A} \boldsymbol{x}-\boldsymbol{b}\|$ where $\boldsymbol{A}$ is a $m \times n$ matrix and each rows of $\boldsymbol{A}$ comes from a certain row dictionary. Roughly speaking, our method shows that

$$
\mathcal{T}(m, n)=\widetilde{O}(\mathrm{nnz}(\boldsymbol{A}))+\widetilde{O}(\mathcal{T}(\widetilde{O}(n), n))
$$

This powerful reduction allows us to develop the first nearly linear time algorithm for connection Laplacian (Chapter 13). In comparison, the random projection methods only shows that $\mathcal{T}(m, n)=$ $\widetilde{O}\left(\operatorname{nnz}(\boldsymbol{A})+n^{\omega}\right)$.

As we explained in Section 2.3.1, it is known that for a data matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}$, a spectral approximation can be obtained by sampling $O(n \log n)$ rows, each with probability proportional to its statistical leverage score. Recall that the leverage score of $\boldsymbol{A}$ 's $i^{\text {th }}$ row, $\boldsymbol{a}_{i}$, is $\boldsymbol{\sigma}_{i}=\boldsymbol{a}_{i}^{\top}\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+} \boldsymbol{a}_{i}$, where $\boldsymbol{M}^{+}$denotes the Moore-Penrose pseudoinverse of $\boldsymbol{M}$. A higher leverage score indicates that $\boldsymbol{a}_{i}$ is more important in composing the spectrum of $\boldsymbol{A}$.

Unfortunately, leverage scores are difficult to calculate - finding them involves computing the pseudoinverse $\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+}$, which is as slow as solving our regression problem in the first place. In practice, data is often assumed to have low coherence [195], in which case simply selecting rows uniformly at random works $[22,156]$. However, uniform sampling could be disastrous - if $\boldsymbol{A}$ contains
a row with some component orthogonal to all other rows, removing it will reduce the rank of $\boldsymbol{A}$ and thus we cannot possibly preserve all vector products ( $\|\tilde{\boldsymbol{A}} \boldsymbol{x}\|_{2}$ will start sending some vectors to 0 ). Any uniform sampling scheme is likely to drop any such single row. On the other hand, when leverage score sampling, such a row would have the highest possible leverage score of 1 .

Possible fixes include randomly "mixing" data points to avoid degeneracies [22]. However, this approach sacrifices sparsity and structure in our data matrix, increasing storage and runtime costs. Is there a more elegant fix? First note that sampling $\boldsymbol{A}$ by approximate leverage scores is fine, but we may need to select more than the optimal $O(n \log n)$ rows. With that in mind, consider the following straightforward algorithm for iterative sampling, inspired by [170]:

## 1. Reduce $\boldsymbol{A}$ significantly by sampling rows uniformly.

## 2. Use the smaller matrix to approximate $\left(A^{\top} A\right)^{+}$and leverage scores for $A$.

## 3. Resample rows of $A$ using these estimates to obtain a spectral approximation $\widetilde{A}$.

While intuitive, this scheme was not previously known to work. Our main result is proving that it does. This process will quickly converge on a small spectral approximation to $\boldsymbol{A}$, i.e. with $O(n \log n)$ rows.

A few results come close to an analysis of such a routine - in particular, two iterative sampling schemes are analyzed in [170]. However, the first ultimately requires Johnson-Lindenstrauss projections that mix rows, something we were hoping to avoid. The second almost maintains sparsity and row structure (except for possibly including rows of the identity in $\widetilde{\boldsymbol{A}}$ ), but its convergence rate depends on the condition number of $\boldsymbol{A}$.

More importantly, both of these results are similar in that they rely on the primitive that a (possibly poor) spectral approximation to $\boldsymbol{A}$ is sufficient for approximately computing leverage scores, which are in turn good enough for obtaining an even better spectral approximation. As mentioned, uniform sampling will not in general give a spectral approximation - it does not preserve information about all singular values. Our key contribution is a better understanding of what information uniform sampling does preserve. It turns out that, although weaker than a spectral approximation, the matrix obtained from uniform sampling can nonetheless give leverage score estimates that are good enough to obtain increasingly better approximations to $\boldsymbol{A}$.

## ■ 3.1.1 Our Approach

Suppose we compute a set of leverage score estimates, $\left\{\tilde{\boldsymbol{\sigma}}_{i}\right\}$, using $\left(\widetilde{\boldsymbol{A}}^{\top} \widetilde{\boldsymbol{A}}\right)^{+}$in place of $\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+}$for some already obtained matrix approximation $\widetilde{\boldsymbol{A}}$. As long as our leverage score approximations are upper bounds on the true scores ( $\tilde{\boldsymbol{\sigma}}_{i} \geq \boldsymbol{\sigma}_{i}$ ) we can use them for sampling and still obtain a spectral approximation to $\boldsymbol{A}$. The number of samples we take will be

$$
c \cdot \log n \cdot \sum_{i=1}^{m} \tilde{\boldsymbol{\sigma}}_{i}
$$

where $c$ is some fixed constant. When sampling by exact leverage scores, it can be shown that $\sum_{i} \boldsymbol{\sigma}_{i} \leq n$ so we take $c \cdot n \log n$ rows.

Thus, to prove that our proposed iterative algorithm works, we need to show that, if we uniformly sample a relatively small number of rows from $\boldsymbol{A}$ (Step 1) and estimate leverage scores using these rows (Step 2), then the sum of our estimates will be small. Then, when we sample by these estimated leverage scores in Step 3, we can sufficiently reduce the size of $\boldsymbol{A}$.

In prior work, the sum of overestimates was bounded by estimating each leverage score to within a multiplicative factor. This requires a spectral approximation, which is why previous iterative sampling schemes could only boost poor spectral approximations to better spectral approximations. Of course, a "for each" statement is not required, and we will not get one through uniform sampling. Thus, our core result avoids this technique. Specifically, we show,

Theorem 3.1.1. For any $k$, we can select $k$ rows uniformly at random from $\boldsymbol{A}$ to obtain $\widetilde{\boldsymbol{A}}$. Let $\left\{\tilde{\boldsymbol{\sigma}}_{i}\right\}$ be a set of leverage score estimates for $\boldsymbol{A}$ computed using $\widetilde{\boldsymbol{A}}^{1}$. Then, we have that $\tilde{\boldsymbol{\sigma}}_{i} \geq \boldsymbol{\sigma}_{i}$ for all $i$ and

$$
\mathbf{E}\left[\sum_{i=1}^{m} \tilde{\boldsymbol{\sigma}}_{i}\right] \leq \frac{m n}{k} .
$$

The validity of our proposed iterative sampling scheme immediately follows from Theorem 3.1.1. For example, if we uniformly sample $k=m / 2$ rows then $c \log n \sum \tilde{\boldsymbol{\sigma}}_{i} \leq O(n \log n)$, so we can cut our matrix down to $O(n \log n)$ rows. There is a convenient tradeoff - the more rows uniformly sampled in Step 1, the more we can cut $\boldsymbol{A}$ down by in Step 3. This tradeoff leads to natural recursive and iterative algorithms for row sampling.

We give a proof of Theorem 3.1.1 using a simple expectation argument that bounds $\mathbf{E} \sum \tilde{\boldsymbol{\sigma}}_{i}$. We also prove alternative versions of Theorem 3.1.1 with slightly different guarantees (Theorems 3.6.1) using a technique that we believe is of independent interest. It is well known that, if $\boldsymbol{A}$ has low coherence - that is, has a low maximum leverage score - then uniform sampling from the matrix is actually sufficient for obtaining a full spectral approximation. The uniform rate will upper bound the leverage score rate for every row. With this in mind, we show a powerful fact: while many matrices do not have low coherence, for any $\boldsymbol{A}$, we can decrease the weight on a small subset of rows to make the matrix have low coherence. Specifically,

Lemma 3.1.2 (Coherence Reducing Reweighting). For any $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ and any coherence upper bound $\alpha>0$, there exists a diagonal reweighting matrix $\boldsymbol{W} \in \mathbb{R}^{m \times m}$ with all entries in $[0,1]$ and just ( $n / \alpha$ ) entries not equal to 1 , such that:

$$
\forall i, \boldsymbol{\sigma}_{i}(\boldsymbol{W} \boldsymbol{A}) \leq \alpha
$$

Intuitively, this lemma shows that uniform sampling gives a matrix that spectrally approximates a large sub-matrix of the original data. It follows from our more general Theorem 3.5.1, which describes exactly how leverage scores of $\boldsymbol{A}$ can be manipulated through row reweighting.

We never actually find $\boldsymbol{W}$ explicitly - simply its existence implies our uniform sampling theorems. As explained, since $\boldsymbol{W} \boldsymbol{A}$ has low coherence, uniform sampling would give a spectral approximation to the reweighted matrix and thus a multiplicatively good approximation to each leverage score. Thus, the sum of estimated leverage scores for $\boldsymbol{W} \boldsymbol{A}$ will be low, i.e. $<O(n)$. It can be shown that, for any row that is not reweighted, the leverage score in $\boldsymbol{A}$ computed using a uniformly sampled $\tilde{\boldsymbol{A}}$, is never greater than the corresponding leverage score in $\boldsymbol{W} \boldsymbol{A}$ computed using a uniformly sampled $\widetilde{\boldsymbol{W A}}$. Thus, the sum of approximate leverage scores for rows in $\boldsymbol{A}$ that are not reweighted is small by comparison to their corresponding leverage scores in $\boldsymbol{W} \boldsymbol{A}$. How about the rows that are reweighted in $\boldsymbol{W} \boldsymbol{A}$ ? Lemma 3.1.2 claims there are not too many of these - we can trivially bound their leverage score estimates by 1 and even then the total sum of estimated leverage scores will be small.

This argument gives the result we need: even if a uniformly sampled $\widetilde{\boldsymbol{A}}$ cannot be used to obtain good per row leverage score upper bounds, it is sufficient for ensuring that the sum of all leverage score estimates is not too high.

[^4]
## ■ 3.2 Previous Work

## ■ 3.2.1 Randomized Numerical Linear Algebra

In the past decade, fast randomized algorithms for matrix problems have risen to prominence. Numerous results give improved running times for matrix multiplication, linear regression, and low rank approximation - helpful surveys of this work include [184] and [112]. In addition to asymptotic runtime gains, randomized alternatives to standard linear algebra tools tend to offer significant gains in terms of data access patterns and required working memory.

Algorithms for randomized linear algebra often work by generically reducing problem size - large matrices are compressed (using randomness) to smaller approximations which are processed deterministically via standard linear algebraic methods. Methods for matrix reduction divide roughly into two categories - random projection methods [54, 53, 185, 200, 229] and sampling methods [70, 71, 72, 74, 73, 170].

Random projection methods recombine rows or columns from a large matrix to form a much smaller problem that approximates the original. Descending from the Johnson-Lindenstrauss Lemma [130] and related results, these algorithms are impressive for their simplicity and speed - reducing a large matrix simply requires multiplication by an appropriately chosen random matrix.

Sampling methods, on the other hand, seek to approximate large matrices by judiciously selecting (and reweighting) few rows or columns. Sampling itself is even simpler and faster than random projection - the challenge becomes efficiently computing the correct measure of "importance" for rows or columns. More important rows or columns are selected with higher probability.

## ■ 3.2.2 Approximate Linear Regression

We focus on linear regression, i.e. solving overdetermined systems, which requires our matrix reduction step to produce a spectral approximation $\widetilde{\boldsymbol{A}}$ to the data matrix $\boldsymbol{A}$. One possibility is to obtain a $\varepsilon$ approximation with $O\left(n \log n / \varepsilon^{2}\right)$ rows and to solve regression on the smaller problem to give an approximate solution. To improve stability and achieve $\log (1 / \varepsilon)$ dependence, randomized schemes can be combined with known iterative regression algorithms. These methods only require a constant factor spectral approximation with $O(n \log n)$ rows [22, 55, 53, 192, 228].

When random projections are used, $\widetilde{\boldsymbol{A}}=\boldsymbol{\Pi} \boldsymbol{A}$ for some randomly generated matrix $\Pi$ which is known as a subspace embedding. Recent progress has significantly sped up the process of computing $\Pi \boldsymbol{A}$, leading to the first input-sparsity time algorithms for linear regression (or nearly input-sparsity time if iterative methods are employed) [53, 185, 200].

## - 3.2.3 Row Sampling

An alternative route to spectral approximation is importance sampling. Specifically, $O\left(n \log n / \varepsilon^{2}\right)$ rows can be sampled with probability proportional to their leverage scores [75, 76, 242]. In [242], Spielman and Srivastava specifically focus on spectral approximations for the edge-vertex incidence matrix of a graph. This is more commonly referred to as spectral graph sparsification, a primitive that has become important in research on graph algorithms. Each row in a graph's (potentially tall) edge-vertex incident matrix corresponds to an edge and the row's leverage score is exactly the edge's weighted effective resistance, which is used as the sampling probability in [242].

This application illustrates an important point: for spectral sparsification, it is critical that $\boldsymbol{A}$ is compressed via sampling instead of random projection. Sampling ensures that $\widetilde{\boldsymbol{A}}$ contains only reweighted rows from $\boldsymbol{A}$, so it remains an edge-vertex incidence matrix. In general, sampling is
interesting because it preserves row structure. Even if that structure is just a certain level of sparsity, it can reduce memory requirements and accelerate matrix operations.

While leverage scores for the edge-vertex incidence matrix of a graph can be computed quickly [147, 243], in general, computing leverage scores requires evaluating $\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+}$, which is as difficult as solving regression in the first place. Li, Miller, and Peng address this issue with methods for iteratively computing good row samples [170]. Their algorithms achieve input-sparsity time regression, but are fairly involved and rely on intermediate operations that ultimately require Johnson-Lindenstrauss projections, mixing rows and necessitating dense matrix operations. An alternative approach from [170] does preserve row structure (except for possible additions of rows from the identity to intermediate matrices) but converges in a number of steps that depends on $\boldsymbol{A}$ 's condition number.

## ■ 3.3 Generalized Leverage Score

We often approximate the leverage scores of $\boldsymbol{A}$ by computing them with respect to some other matrix $\boldsymbol{B} \in \mathbb{R}^{m^{\prime} \times n}$. We define the generalized leverage score $:$

$$
\sigma_{i}^{\boldsymbol{B}}(\boldsymbol{A}) \stackrel{\text { def }}{=} \begin{cases}\boldsymbol{A}_{i}^{\top}\left(\boldsymbol{B}^{\top} \boldsymbol{B}\right)^{+} \boldsymbol{A}_{i} & \text { if } \boldsymbol{A}_{i} \perp \operatorname{ker}(\boldsymbol{B})  \tag{3.1}\\ \infty & \text { otherwise }\end{cases}
$$

If $\boldsymbol{A}_{i}$ has an component in $\operatorname{ker}(\boldsymbol{B})$, we set its generalized leverage score to $\infty$, since it might be the only row in $\boldsymbol{A}$ pointing in this direction. Thus, when sampling rows, we cannot remove it. We could set the generalized leverage score to 1 , but using $\infty$ simplifies notation in some of our proofs. If $\boldsymbol{B}$ is a spectral approximation for $\boldsymbol{A}$, then every generalized leverage score is a good multiplicative approximation to its corresponding true leverage score:

Lemma 3.3.1 ([170, Lem 4.3]). If $\frac{1}{\lambda} \boldsymbol{A}^{\top} \boldsymbol{A} \preceq \boldsymbol{B}^{\top} \boldsymbol{B} \preceq \boldsymbol{A}^{\top} \boldsymbol{A}$, then $\sigma_{i}(\boldsymbol{A}) \leq \sigma_{i}^{\boldsymbol{B}}(\boldsymbol{A}) \leq \lambda \cdot \sigma_{i}(\boldsymbol{A})$.

## ■ 3.4 Expectation Bound on Leverage Score Estimation

In this section, we use a simple expectation argument to prove Theorem 3.1.1, which is restated in full below:

Theorem 3.1.1 (Full Statement). Given any $\boldsymbol{A} \in \mathbb{R}^{m \times n}$. Let $S$ denote a uniformly random sample of $k$ rows from $\boldsymbol{A}$ and let $\boldsymbol{S} \in \mathbb{R}^{m \times m}$ be its diagonal indicator matrix (i.e. $\boldsymbol{S}_{i i}=1$ for $i \in S, \boldsymbol{S}_{i i}=0$ otherwise). Define

$$
\tilde{\boldsymbol{\sigma}}_{i} \stackrel{\text { def }}{=} \begin{cases}\boldsymbol{\sigma}_{i}^{\boldsymbol{S A}}(\boldsymbol{A}) & \text { if } i \in S, \\ \frac{1}{1+\frac{1}{\boldsymbol{\sigma}_{i}^{S A}(\boldsymbol{A})}} & \text { otherwise. }\end{cases}
$$

Then, $\tilde{\boldsymbol{\sigma}}_{i} \geq \boldsymbol{\sigma}_{i}(\boldsymbol{A})$ for all $i$ and

$$
\mathbf{E}\left[\sum_{i=1}^{m} \tilde{\boldsymbol{\sigma}}_{i}\right] \leq \frac{m n}{k}
$$

Proof. First we show that our estimates are valid leverage score upper bounds, i.e. $\tilde{\boldsymbol{\sigma}}_{i} \geq \boldsymbol{\sigma}_{i}(\boldsymbol{A})$. Let $\boldsymbol{S}^{(i)}$ be the diagonal indicator matrix for $S \cup\{i\}$. We claim that, for all $i$,

$$
\begin{equation*}
\tilde{\boldsymbol{\sigma}}_{i}=\boldsymbol{\sigma}_{i}^{\boldsymbol{S}^{(i)} \boldsymbol{A}}(\boldsymbol{A}) \tag{3.2}
\end{equation*}
$$

This is proved case-by-case:

When $i \in S, \boldsymbol{S}=\boldsymbol{S}^{(i)}$ so it holds trivially.
When $i \notin S$ and $\boldsymbol{a}_{i} \not \perp \operatorname{ker}(\boldsymbol{S A})$, then by definition, $\boldsymbol{\sigma}_{i}^{\boldsymbol{S} \boldsymbol{A}}(\boldsymbol{A})=\infty$ and $\tilde{\boldsymbol{\sigma}}_{i}=\frac{1}{1+\frac{1}{\infty}}=1=\boldsymbol{\sigma}_{i}^{\boldsymbol{S}^{(i)} \boldsymbol{A}}(\boldsymbol{A})$.
When $i \notin S$ and $\boldsymbol{a}_{i} \perp \operatorname{ker}(\boldsymbol{S A})$ then by the Sherman-Morrison formula (Lemma 2.3.7),

$$
\begin{aligned}
& \boldsymbol{\sigma}_{i}^{\boldsymbol{S}^{(i)} \boldsymbol{A}}(\boldsymbol{A})=\boldsymbol{a}_{i}^{\top}\left(\boldsymbol{A}^{\top} \boldsymbol{S}^{2} \boldsymbol{A}+\boldsymbol{a}_{i} \boldsymbol{a}_{i}^{\top}\right)^{+} \boldsymbol{a}_{i} \\
& =\boldsymbol{a}_{i}^{\top}\left(\left(\boldsymbol{A}^{\top} \boldsymbol{S}^{2} \boldsymbol{A}\right)^{+}-\frac{\left(\boldsymbol{A}^{\top} \boldsymbol{S}^{2} \boldsymbol{A}\right)^{+} \boldsymbol{a}_{i} \boldsymbol{a}_{i}^{\top}\left(\boldsymbol{A}^{\top} \boldsymbol{S}^{2} \boldsymbol{A}\right)^{+}}{1+\boldsymbol{a}_{i}^{\top}\left(\boldsymbol{A}^{\top} \boldsymbol{S}^{2} \boldsymbol{A}\right)^{+} \boldsymbol{a}_{i}}\right) \boldsymbol{a}_{i} \\
& =\boldsymbol{\sigma}_{i}^{\boldsymbol{S A}}(\boldsymbol{A})-\frac{\boldsymbol{\sigma}_{i}^{\boldsymbol{S A}}(\boldsymbol{A})^{2}}{1+\boldsymbol{\sigma}_{i}^{\boldsymbol{S A}}(\boldsymbol{A})}=\frac{1}{1+\frac{1}{\boldsymbol{\sigma}_{i}^{S A}(\boldsymbol{A})}}=\tilde{\boldsymbol{\sigma}}_{i} .
\end{aligned}
$$

By (3.2) and the fact that $\boldsymbol{A}^{\top} \boldsymbol{S}^{(i)^{2}} \boldsymbol{A} \preceq \boldsymbol{A}^{\top} \boldsymbol{A}$ (see Lemma 3.3.1), we have $\tilde{\boldsymbol{\sigma}}_{i}=\boldsymbol{\sigma}_{i}^{\boldsymbol{S}^{(i)} \boldsymbol{A}}(\boldsymbol{A}) \geq \boldsymbol{\sigma}_{i}(\boldsymbol{A})$, so our estimates are upper bounds as desired. It remains to upper bound the expected sum of $\tilde{\boldsymbol{\sigma}}_{i}$. We can break down the sum as:

$$
\sum_{i=1}^{m} \tilde{\boldsymbol{\sigma}}_{i}=\sum_{i \in S} \tilde{\boldsymbol{\sigma}}_{i}+\sum_{i \notin S} \tilde{\boldsymbol{\sigma}}_{i} .
$$

The first term is simply the sum of $\boldsymbol{S} \boldsymbol{A}$ 's leverage scores, so it is equal to $\operatorname{rank}(\boldsymbol{S} \boldsymbol{A}) \leq n$ by (2.2). To bound the second term, consider a random process that first selects $\boldsymbol{S}$, then selects a random row $i$ $\notin S$ and returns $\tilde{\boldsymbol{\sigma}}_{i}$. There are always exactly $m-k$ rows $\notin S$, so the value returned by this random process is, in expectation, exactly equal to $\frac{1}{m-k} \cdot \mathbf{E} \sum_{i \notin S} \tilde{\sigma}_{i}$.

This random process is also equivalent to randomly selecting a set $S^{\prime}$ of $k+1$ rows, then randomly choosing a row $i \in S^{\prime}$ and returning its leverage score! In expectation it is therefore equal to the average leverage score in $\boldsymbol{S}^{\prime} \boldsymbol{A} . \boldsymbol{S}^{\prime} \boldsymbol{A}$ has $k+1$ rows and its leverage scores sum to its rank, so we can bound its average leverage score by $\frac{n}{k+1}$. Overall:

$$
\mathbf{E}\left[\sum_{i=1}^{m} \tilde{\boldsymbol{\sigma}}_{i}\right] \leq n+(m-k) \cdot \frac{n}{k+1} \leq \frac{n(m+1)}{k+1} \leq \frac{m n}{k} .
$$

## ■ 3.5 Coherence-Reducing Weighting

In this section, we prove Theorem 3.5.1, which shows that we can reweight a small number of rows in any matrix $\boldsymbol{A}$ to make it have low coherence. This structural result may be of independent interest. It is also fundamental in proving Theorem 3.6.1, a slightly stronger version of Theorem 3.1.1 that we will prove in Section 3.6.

Actually, for Theorem 3.5.1 we prove a more general statement, studying how to select a diagonal row reweighting matrix $\boldsymbol{W}$ to arbitrarily control the leverage scores of $\boldsymbol{W} \boldsymbol{A}$. One simple conjecture would be that, given a vector $\boldsymbol{u}$, there always exists a $\boldsymbol{W}$ such that $\boldsymbol{\sigma}_{i}(\boldsymbol{W} \boldsymbol{A})=\boldsymbol{u}_{i}$. This conjecture is unfortunately not true - if $\boldsymbol{A}$ is the identity matrix, then $\boldsymbol{\sigma}_{i}(\boldsymbol{W} \boldsymbol{A})=0$ if $\boldsymbol{W}_{i i}=0$ and $\boldsymbol{\sigma}_{i}(\boldsymbol{W} \boldsymbol{A})=1$ otherwise. Instead, we show the following:

Theorem 3.5.1. For any $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ and any vector $\boldsymbol{u} \in \mathbb{R}^{m}$ with $\boldsymbol{u}_{i}>0$ for all $i$, there exists a diagonal matrix $\boldsymbol{W} \in \mathbb{R}^{m \times m}$ with $\mathbf{0} \preceq \boldsymbol{W} \preceq \boldsymbol{I}$ such that:

$$
\begin{equation*}
\forall i, \boldsymbol{\sigma}_{i}(\boldsymbol{W} \boldsymbol{A}) \leq \boldsymbol{u}_{i} \tag{3.3}
\end{equation*}
$$

and

$$
\begin{equation*}
\sum_{i: \boldsymbol{W}_{i i} \neq 1} \boldsymbol{u}_{i} \leq n \tag{3.4}
\end{equation*}
$$

Note that (3.3) is easy to satisfy - it holds if we set $\boldsymbol{W}=\mathbf{0}$. Hence, the main result is the second claim. Not only does a $\boldsymbol{W}$ exist that gives the desired leverage score bounds, but it is only necessary to reweight rows in $\boldsymbol{A}$ with a low total weight in terms of $\boldsymbol{u}$.

For any incoherence parameter $\alpha$, if we set $\boldsymbol{u}_{i}=\alpha$ for all $i$, then this theorem shows the existence of a reweighting that reduces coherence to $\alpha$. Such a reweighting has $\sum_{i: W_{i i} \neq 1} \alpha \leq n$ and therefore $\left|\left\{i: \boldsymbol{W}_{i i} \neq 1\right\}\right| \leq \frac{n}{\alpha}$. So, we see that Lemma 3.1.2 follows as a special case of Theorem 3.5.1.

In order to prove Theorem 3.5.1, we first give two technical lemmas which are proved Section 3.8. Lemma 3.5.2 describes how the leverage scores of $\boldsymbol{A}$ evolve when a single row of $\boldsymbol{A}$ is reweighted. We show that, when we decrease the weight of a row, that row's leverage score decreases and the leverage score of all other rows increases.

Lemma 3.5.2 (Leverage Score Changes Under Rank 1 Updates). Given any $\boldsymbol{A} \in \mathbb{R}^{m \times n}, \gamma \in(0,1)$, and $i \in[m]$, let $\boldsymbol{W}$ be a diagonal matrix such that $\boldsymbol{W}_{i i}=\sqrt{1-\gamma}$ and $\boldsymbol{W}_{j j}=1$ for all $j \neq i$. Then,

$$
\boldsymbol{\sigma}_{i}(\boldsymbol{W} \boldsymbol{A})=\frac{(1-\gamma) \boldsymbol{\sigma}_{i}(\boldsymbol{A})}{1-\gamma \boldsymbol{\sigma}_{i}(\boldsymbol{A})} \leq \boldsymbol{\sigma}_{i}(\boldsymbol{A})
$$

and for all $j \neq i$,

$$
\boldsymbol{\sigma}_{j}(\boldsymbol{W} \boldsymbol{A})=\boldsymbol{\sigma}_{j}(\boldsymbol{A})+\frac{\gamma \boldsymbol{\sigma}_{i j}(\boldsymbol{A})^{2}}{1-\gamma \boldsymbol{\sigma}_{i}(\boldsymbol{A})} \geq \boldsymbol{\sigma}_{j}(\boldsymbol{A})
$$

Next we claim that, with respect to weightings of $\boldsymbol{A}$ 's rows, leverage scores are lower semicontinuous.

Lemma 3.5.3 (Leverage Scores are Lower Semi-continuous). $\boldsymbol{\sigma}(\boldsymbol{W} \boldsymbol{A})$ is lower semi-continuous in the diagonal matrix $\boldsymbol{W}$, i.e. for any sequence $\boldsymbol{W}^{(k)} \rightarrow \overline{\boldsymbol{W}}$ with $\boldsymbol{W}_{i i}^{(k)} \geq 0$ for all $k$ and $i$, we have

$$
\boldsymbol{\sigma}_{i}(\overline{\boldsymbol{W}} \boldsymbol{A}) \leq \liminf _{k \rightarrow \infty} \boldsymbol{\sigma}_{i}\left(\boldsymbol{W}^{(k)} \boldsymbol{A}\right)
$$

With Lemmas 3.5.2 and 3.5.3 in place, we are ready to prove the main reweighting theorem.
Proof of Theorem 3.5.1. We prove the existence of the required $\boldsymbol{W}$ by considering the limit of the Algorithm 1 for computing a reweighting matrix.

For all $k \geq 0$, let $\boldsymbol{W}^{(k)}$ be the value of $\boldsymbol{W}$ after the $k^{\text {th }}$ update to the weight. We show that $\overline{\boldsymbol{W}}=\lim _{k \rightarrow \infty} \boldsymbol{W}^{(k)}$ meets the conditions of Theorem 3.5.1. First note that Algorithm 1 is well defined and that all entries of $\boldsymbol{W}^{(k)}$ are non-negative for all $k \geq 0$. To see this, suppose we need to decrease $\boldsymbol{W}_{i i}^{(k)}$ so that $\boldsymbol{\sigma}_{i}\left(\boldsymbol{W}^{(k+1)} \boldsymbol{A}\right)=\boldsymbol{u}_{i}$. Note that the condition $\boldsymbol{\sigma}_{i}\left(\boldsymbol{W}^{(k)} \boldsymbol{A}\right)<1$ gives

$$
\lim _{\gamma \rightarrow 1} \frac{(1-\gamma) \boldsymbol{\sigma}_{i}\left(\boldsymbol{W}^{(k)} \boldsymbol{A}\right)}{1-\gamma \boldsymbol{\sigma}_{i}\left(\boldsymbol{W}^{(k)} \boldsymbol{A}\right)}=0
$$

Therefore, Lemma 3.5 .2 shows that we can make $\boldsymbol{\sigma}_{i}\left(\boldsymbol{W}^{(k+1)} \boldsymbol{A}\right)$ arbitrary small by setting $\gamma$ close enough to 1 . Since the leverage score for row $i$ is continuous, this implies that $\boldsymbol{W}^{(k+1)}$ exists as desired.

```
Algorithm 1: ComputeReweighting (a.k.a the whack-a-mole algorithm)
    Initialize \(\boldsymbol{W}=\boldsymbol{I}\).
    while true do
        for \(i=1\) to \(n\) do
            if \(\boldsymbol{\sigma}_{i}(\boldsymbol{W} \boldsymbol{A}) \geq \boldsymbol{u}_{i}\) then
            if \(\boldsymbol{\sigma}_{i}(\boldsymbol{W} \boldsymbol{A})<1\) then
                Decrease \(\boldsymbol{W}_{i i}\) so that \(\boldsymbol{\sigma}_{i}(\boldsymbol{W} \boldsymbol{A})=\boldsymbol{u}_{i}\).
            else
                Set \(\boldsymbol{W}_{i i}=0\).
            end
            end
        end
    end
    Output: \(W\).
```

Since, the entries of $\boldsymbol{W}^{(k)}$ are non-negative and decrease monotonically by construction, clearly $\overline{\boldsymbol{W}}$ exists. Furthermore, since setting $\boldsymbol{W}_{i i}=0$ makes $\boldsymbol{\sigma}_{i}(\boldsymbol{W} \boldsymbol{A})=0$, we see that, by construction,

$$
\liminf _{k \rightarrow \infty} \boldsymbol{\sigma}_{i}\left(\boldsymbol{W}^{(k)} \boldsymbol{A}\right) \leq \boldsymbol{u}_{i} \text { for all } i \in[n]
$$

Therefore, by Lemma 3.5.3 we have that $\boldsymbol{\sigma}_{i}(\overline{\boldsymbol{W}} \boldsymbol{A}) \leq u_{i}$.
It only remains to show that $\sum_{i: \overline{\boldsymbol{W}}_{i i} \neq 1} \boldsymbol{u}_{i} \leq n$. Let $k$ be the first iteration such that $\boldsymbol{W}_{i i}^{(k)} \neq 1$ for any $i$ such that $\overline{\boldsymbol{W}}_{i i} \neq 1$. Let $S \subseteq[m]$ be the set of rows such that $\boldsymbol{W}_{i i}^{(k)}=0$ and let $T=\left\{i: \overline{\boldsymbol{W}}_{i i} \neq\right.$ $1\}-S$. Since decreasing the weight of one row increases the leverage scores of all other rows, we have

$$
\begin{aligned}
\sum_{i \in T \cup S} \boldsymbol{u}_{i} & \leq \sum_{i \in T} \boldsymbol{\sigma}_{i}\left(\boldsymbol{W}^{(k)} \boldsymbol{A}\right)+\sum_{i \in S} 1 \\
& \leq \operatorname{rank}\left(\boldsymbol{W}^{(k)} \boldsymbol{A}\right)+|S|
\end{aligned}
$$

When we set $\boldsymbol{W}_{i i}=0$, it must be the case that $\boldsymbol{\sigma}_{i}(\boldsymbol{W} \boldsymbol{A})=1$. In this case, removing the $i^{\text {th }}$ row decreases the rank of $\boldsymbol{W} \boldsymbol{A}$ by 1 and hence $\operatorname{rank}\left(\boldsymbol{W}^{(k)} \boldsymbol{A}\right) \leq n-|S|$. Therefore,

$$
\sum_{i: \overline{\boldsymbol{W}_{i i}} \neq 1} \boldsymbol{u}_{i}=\sum_{i \in T \cup S} \boldsymbol{u}_{i} \leq n .
$$

## ■ 3.6 High Probability Bound on Leverage Score Estimation

Theorem 3.1.1 alone is enough to prove that a variety of iterative methods for spectral approximation work. In particular, we can use Lemma 2.3 .3 to show that sampling each row with probability $\Theta\left(\min \left\{1, \tilde{\boldsymbol{\sigma}}_{i} \log n\right\}\right)$ independently with suitable reweighting gives a spectral approximation of the original matrix. Since $\tilde{\boldsymbol{\sigma}}_{i}$ is a constant approximation of $\min \left\{1, \boldsymbol{\sigma}_{i}^{\boldsymbol{S} \boldsymbol{A}}(\boldsymbol{A})\right\}$, the total number of rows it takes is $\Theta\left(\sum_{i} \min \left\{1, \boldsymbol{\sigma}_{i}^{\boldsymbol{S} \boldsymbol{A}}(\boldsymbol{A}) \log n\right\}\right)$ which can be bounded by the $\sum_{i} \min \left\{1, \boldsymbol{\sigma}_{i}^{\boldsymbol{S} \boldsymbol{A}}(\boldsymbol{A})\right\}$ bound proved in Theorem 3.1.1.

In this section, we bound $\sum_{i} \min \left\{1, \boldsymbol{\sigma}_{i}^{\boldsymbol{S} \boldsymbol{A}}(\boldsymbol{A}) \log n\right\}$ directly and it allows us to bound the quantity with a high probability instead of merely in expectation. The proof relies on Theorem 3.5.1, which
intuitively shows that a large fraction of our matrix $\boldsymbol{A}$ has low coherence. Sampling rows uniformly will give a spectral approximation for this portion of our matrix. Then, since few rows are reweighted in $\boldsymbol{W} \boldsymbol{A}$, even loose upper bounds on the leverage scores for those rows will allow us to bound the total sum of estimated leverage scores when we sample uniformly.

Theorem 3.6.1. Given any $\boldsymbol{A} \in \mathbb{R}^{m \times n}$. Let $S$ denote a uniformly random sample of $k$ rows from $\boldsymbol{A}$ and let $\boldsymbol{S} \in \mathbb{R}^{m \times m}$ be its diagonal indicator matrix (i.e. $\boldsymbol{S}_{i i}=1$ for $i \in S$, $\boldsymbol{S}_{i i}=0$ otherwise). With high probability in $n$, we have that

$$
\sum_{i=1}^{n} \min \left\{1, \boldsymbol{\sigma}_{i}^{\boldsymbol{S A}}(\boldsymbol{A}) \log n\right\} \leq O\left(\frac{m n \log n}{k}\right)
$$

Resampling from $\boldsymbol{A}$ by these estimates will, with high probability, return a $\varepsilon$-spectral approximation to $\boldsymbol{A}$ with at most $O\left(\frac{m n \log n}{k \varepsilon^{2}}\right)$ rows.

Proof. Let $c$ be the constant $c$ in Lemma 2.3.3. By Theorem 3.5.1, there is a diagonal matrix $\boldsymbol{W} \preceq \boldsymbol{I}$ such that $\sigma_{i}(\boldsymbol{W} \boldsymbol{A}) \leq \frac{k}{c m \log n}$ for all $i$ and $\sum_{i: \boldsymbol{W}_{i i} \neq 1} \frac{k}{c m \log n} \leq n$. For this $\boldsymbol{W}$, we have

$$
\begin{align*}
\sum_{i=1}^{n} \min \left\{1, \boldsymbol{\sigma}_{i}^{\boldsymbol{S} \boldsymbol{A}}(\boldsymbol{A}) \log n\right\} & \leq \sum_{i: \boldsymbol{W}_{i i} \neq 1} 1+\log n \sum_{i: \boldsymbol{W}_{i i}=1} \boldsymbol{\sigma}_{i}^{\boldsymbol{S} \boldsymbol{A}}(\boldsymbol{A}) \\
& \leq \frac{c m n \log n}{k}+\log n \sum_{i: \boldsymbol{W}_{i i}=1} \boldsymbol{\sigma}_{i}^{\boldsymbol{S A}}(\boldsymbol{A}) \\
& =\frac{c m n \log n}{k}+\log n \sum_{i: \boldsymbol{W}_{i i}=1} \boldsymbol{\sigma}_{i}^{\boldsymbol{S} \boldsymbol{A}}(\boldsymbol{W} \boldsymbol{A}) . \tag{3.5}
\end{align*}
$$

Since $\boldsymbol{\sigma}_{i}(\boldsymbol{W} \boldsymbol{A}) \leq \frac{k}{c m \log n}$ and $\boldsymbol{S}$ is a uniformly random sample of $k$ rows, Lemma 2.3.3 shows that

$$
\begin{equation*}
\frac{1}{2} \boldsymbol{A}^{\top} \boldsymbol{W}^{2} \boldsymbol{A} \preceq \frac{m}{k} \boldsymbol{A}^{\top} \boldsymbol{W} \boldsymbol{S}^{2} \boldsymbol{W} \boldsymbol{A} . \tag{3.6}
\end{equation*}
$$

Hence (3.6) along with Lemma 3.3.1 shows that

$$
\boldsymbol{\sigma}_{i}^{S \boldsymbol{A}}(\boldsymbol{W} \boldsymbol{A}) \leq \frac{2 m}{k} \boldsymbol{\sigma}_{i}(\boldsymbol{W} \boldsymbol{A})
$$

Combining with (3.5), we have

$$
\begin{aligned}
\sum_{i=1}^{n} \min \left\{1, \boldsymbol{\sigma}_{i}^{\boldsymbol{S A}}(\boldsymbol{A}) \log n\right\} & \leq \frac{c m n \log n}{k}+\frac{2 m \log n}{k} \sum_{i: \boldsymbol{W}_{i i}=1} \boldsymbol{\sigma}_{i}(\boldsymbol{W} \boldsymbol{A}) \\
& \leq \frac{c m n \log n}{k}+\frac{2 m n \log n}{k}=O\left(\frac{m n \log n}{k}\right) .
\end{aligned}
$$

The last statement follows from Lemma 3.3.1 with sampling probabilities $\min \left\{1, \frac{c \sigma_{i}^{S \boldsymbol{A}}(\boldsymbol{A}) \log n}{\varepsilon^{2}}\right\}$.

## ■ 3.7 Main Result

As discussed in the introduction, Theorems 3.1.1 and 3.6.1 immediately yield new, extremely simple algorithms for spectral matrix approximation. For clarity, we initially present versions running in nearly input-sparsity time. However, we later explain how our first algorithm can be modified with
standard techniques to remove log factors, achieving input-sparsity time and thus matching state-of-the-art results [53, 185, 200]. Our algorithms rely solely on row sampling, which preserves matrix sparsity and structure, possibly improving space usage and runtime for intermediate system solves.

## - 3.7.1 Repeated Halving

The algorithm we present, RepeatedHalving, is a simple recursive procedure. We uniformly sample $\frac{m}{2}$ rows from $\boldsymbol{A}$ to obtain $\boldsymbol{A}^{\prime}$. By Theorems 3.1.1 and 3.6.1, estimating leverage scores of $\boldsymbol{A}$ with respect to this sample allows us to immediately find a spectral approximation to $\boldsymbol{A}$ with $O(n \log n)$ rows. Of course, $\boldsymbol{A}^{\prime}$ is still large, so computing these estimates would be slow. Thus, we recursively find a spectral approximation of $\boldsymbol{A}^{\prime}$ and use this to compute the estimated leverage scores.

## Algorithm 2: RepeatedHalving

Input: $m \times n$ matrix $\boldsymbol{A}$
Output: spectral approximation $\widetilde{\boldsymbol{A}}$ consisting of $O(n \log n)$ rescaled rows of $\boldsymbol{A}$
Uniformly sample $\frac{m}{2}$ rows of $\boldsymbol{A}$ to form $\boldsymbol{A}^{\prime}$.
If $\boldsymbol{A}^{\prime}$ has $>O(n \log n)$ rows, recursively compute a spectral approximation $\widetilde{\boldsymbol{A}}^{\prime}$ of $\boldsymbol{A}^{\prime}$.
Compute approximate generalized leverage scores of $\boldsymbol{A}$ w.r.t. $\widetilde{\boldsymbol{A}}^{\prime}$
Use these estimates to sample rows of $\boldsymbol{A}$ to form $\widetilde{\boldsymbol{A}}$.
Output $\widetilde{A}$.
In analyzing the runtimes of these algorithms, we assume $m=O(\operatorname{poly}(n))$, which is a reasonable assumption for any practical regression problem. ${ }^{2}$ Furthermore, we use the fact that a $n \times n$ system can be solved in time $n^{\omega}$, where $\omega$ is the matrix multiplication exponent. However, we emphasize that, depending on the structure and sparsity of $\boldsymbol{A}$, alternative system solving methods may yield faster results or runtimes with different trade offs. For example, if the rows of $\boldsymbol{A}$ are sparse, solving a system in $\widetilde{\boldsymbol{A}}$, where $\widetilde{\boldsymbol{A}}$ consists of $O(n \log n)$ rescaled rows from $\boldsymbol{A}$ may be accelerated by using iterative conjugate gradient, or other Krylov subspace methods (which can also avoid explicitly computing $\widetilde{\boldsymbol{A}}^{\top} \widetilde{\boldsymbol{A}}$ ). It is best to think of $n^{\omega}$ as the runtime of the fastest available system solver in your domain, and the quoted runtimes as general guidelines that will change somewhat depending on exactly how the above algorithms are implemented.

First, we give an important primitive showing that estimates of generalized leverage scores can be computed efficiently. Computing exact generalized leverage scores is slow and we only need constant factor approximations, which will only increase our sampling rates and hence number of rows sampled by a constant factor.

Lemma 3.7.1. Given B containing $O(n \log n)$ rescaled rows of $\boldsymbol{A}$, for any $1>\theta>0$, it is possible to compute an estimate of $\boldsymbol{\sigma}^{\boldsymbol{B}}(\boldsymbol{A}), \tilde{\boldsymbol{\sigma}}$, in $O\left(n^{\omega} \theta^{-1}+\operatorname{nnz}(\boldsymbol{A}) \theta^{-1}\right)$ time such that, w.h.p. in $n$, for all $i, \tilde{\boldsymbol{\sigma}}_{i} \geq \boldsymbol{\sigma}_{i}^{\boldsymbol{B}}(\boldsymbol{A})$ and $\tilde{\boldsymbol{\sigma}}_{i} \leq O\left(n^{\theta}\right) \boldsymbol{\sigma}_{i}^{\boldsymbol{B}}(\boldsymbol{A})$.

Setting $\theta=O\left(\frac{1}{\log n}\right)$ gives a constant factor approximation to generalized leverage score in $O\left(n^{\omega} \log n+\right.$ $n n z(\boldsymbol{A}) \log n)$ time.

Proof. The proof is basically same as Lemma 2.3.4 except that we use fast matrix multiplication to solve the corresponding linear systems and that we need to handle the case $\boldsymbol{a}_{i} \not \not \perp \operatorname{ker}(\boldsymbol{B})$. When $\boldsymbol{a}_{i} \not \perp \operatorname{ker}(\boldsymbol{B})$, its generalized leverage score should be $\infty-$ see (3.1). So, we need to check whether each $\boldsymbol{a}_{i}$ has a component in the null-space of $\boldsymbol{B}$. This can be done in a variety of ways. For example, we can choose a random gaussian vector $\boldsymbol{g}$ and compute $\boldsymbol{g}-\boldsymbol{B}\left(\boldsymbol{B}^{\top} \boldsymbol{B}\right)^{+} \boldsymbol{B}^{\top} \boldsymbol{g}$, which is the same as

[^5]$\boldsymbol{g}-\left(\boldsymbol{B}^{\top} \boldsymbol{B}\right)^{+} \boldsymbol{B}^{\top} \boldsymbol{B} \boldsymbol{g}$. This gives a random vector in the null space of $\boldsymbol{B}$, so computing its dot product with any row $\boldsymbol{a}_{i}$ will tell us (with probability 1) whether $\boldsymbol{a}_{i}$ is orthogonal to the null space or not.

With Lemma 3.7.1 in place, we can analyze the runtime of our algorithms. For simplicity, we give runtimes for computing a constant factor spectral approximation to $\boldsymbol{A}$, which can be used as a preconditioner in iterative regression algorithms $[22,53,228]$ or used to compute leverage scores of $\boldsymbol{A}$ up to a constant factor. We could sample $O\left(n \log n \varepsilon^{-2}\right)$ rows to directly obtain a $\varepsilon$-approximation. By Lemma 3.7.1 the runtime of this final refinement is just $O\left(\operatorname{nnz}(\boldsymbol{A}) \log n+n^{\omega} \log n\right)$.

Lemma 3.7.2. RepeatedHalving (Algorithm 2) runs in time $O\left(\mathrm{nnz}(\boldsymbol{A}) \log n+n^{\omega} \log (m / n) \log n\right)$, outputting a matrix with $\widetilde{\boldsymbol{A}}$ with $O(n \log n)$ rows such that $\widetilde{\boldsymbol{A}}^{\top} \widetilde{\boldsymbol{A}} \approx_{2} \boldsymbol{A}^{\top} \boldsymbol{A}$.

Proof. The proof is by induction - it suffices to show that the work at the top level is $O(\mathrm{nnz}(\boldsymbol{A}) \log n+$ $\left.n^{\omega} \log n\right)$. At each of the $O(\log (m / n))$ levels of recursion, we cut our matrix in half uniformly so $\mathrm{nnz}(\boldsymbol{A})$ will also be cut approximately in half with high probability.

By Theorem 3.6.1, sampling by $\boldsymbol{\sigma}_{i}^{\boldsymbol{A}^{\prime}}(\boldsymbol{A})$ allows us to obtain $\widetilde{\boldsymbol{A}}$ with $O(n \log n)$ rows. If we instead use $\widetilde{\boldsymbol{A}}^{\prime}$, our estimated leverage scores increase by at most a constant factor (since $\widetilde{\boldsymbol{A}}^{\prime}$ is a constant factor spectral approximation to $\boldsymbol{A}^{\prime}$ ). Furthermore, using Lemma 3.7.1 to approximate generalized leverage scores increases our estimates by another constant factor at most. Overall, $\widetilde{\boldsymbol{A}}$ will have $O(n \log n)$ rows as desired and our runtime at the top level is just the runtime of estimating leverage scores from Lemma 3.7.1 - $O\left(\operatorname{nnz}(\boldsymbol{A}) \log n+n^{\omega} \log n\right)$.

## ■ 3.7.2 Achieving Input Sparsity Time

We briefly note that, using techniques from [170], it is possible to remove the $\log n$ factor on the $\mathrm{nnz}(\boldsymbol{A})$ term to achieve true input-sparsity time with RepeatedHalving. Instead of using Lemma 3.7.1 to estimate generalized leverage scores from up to a constant factor using $\boldsymbol{A}^{\prime}$, we only estimate them up to a $O\left(n^{\theta}\right)$ factor for some constant $0<\theta<\omega-2$. Using these rough estimates, we obtain $\widetilde{\boldsymbol{A}}$ with $O\left(n^{1+\theta} \log n\right)$ rows. Then, for the rows in $\widetilde{\boldsymbol{A}}$, we can again compute generalized leverage scores with respect to $\boldsymbol{A}^{\prime}$, now up to constant factors, and reduce down to just $O(n \log n)$ rows. In total, each iteration will take time $O\left(\theta^{-1}\left(\mathrm{nnz}(\boldsymbol{A})+n^{\omega} \log n\right)\right)$, so obtaining a constant factor approximation to $\boldsymbol{A}$ takes time $O\left(\theta^{-1} \mathrm{nnz}(\boldsymbol{A})+\theta^{-1} n^{\omega} \log ^{2} n\right)$. Recall that we assume $m=\operatorname{poly}(n)$, so we have $\log (m / n)=O(\log n)$ iterations.

In order to obtain a $\varepsilon$-spectral approximation with only $O\left(n \log n \varepsilon^{-2}\right)$ rows, we first obtain a constant factor approximation, $\widetilde{\boldsymbol{A}}$, with $O(n \log n)$ rows. We then use leverage scores estimated with $\widetilde{\boldsymbol{A}}$ to compute a $\varepsilon / 2$-approximation to $\boldsymbol{A}$ with $O\left(n^{1+\theta} \log n \varepsilon^{-2}\right)$ rows. Finally, we again use leverage scores estimated with $\widetilde{\boldsymbol{A}}$ and Lemma 3.7.1 with $\theta=O\left(\varepsilon^{2} / \log n\right)$ to a compute a $\varepsilon / 2$-approximation to this smaller matrix with only $O\left(n \log n \varepsilon^{-2}\right)$ rows. This takes total time $O\left(\theta^{-1} \mathrm{nnz}(\boldsymbol{A})+\theta^{-1} n^{\omega} \log ^{2} n+\right.$ $n^{2+\theta} \log ^{2} n \varepsilon^{-2}$ ). The $n^{2+\theta} \log ^{3} n \varepsilon^{-2}$ comes from applying Lemma 3.7.1 to refine our second approximation, which has $O\left(n^{1+\theta} \log n \varepsilon^{-2}\right)$ rows and thus at most $O\left(n^{2+\theta} \log n \varepsilon^{-2}\right)$ nonzero entries. Overall, the technique yields:

Lemma 3.7.3. Given any constant $0<\theta \leq \omega-2$, and any error $0 \leq \varepsilon<1$, w.h.p. in $n$ we can compute a matrix $\widetilde{\boldsymbol{A}}$ with $O\left(n \log n \varepsilon^{-2}\right)$ rows such that $\widetilde{\boldsymbol{A}}^{\top} \widetilde{\boldsymbol{A}} \approx_{1+\varepsilon} \boldsymbol{A}^{\top} \boldsymbol{A}$ in $O\left(\theta^{-1} \mathrm{nnz}(\boldsymbol{A})+\right.$ $\left.\theta^{-1} n^{\omega} \log ^{2} n+n^{2+\theta} \varepsilon^{-2}\right)$ time .

Proof. As is standard, $\log n$ factors on the $n^{2+\theta}$ term are 'hidden' as we can just slightly increase the value of $\theta$ to subsume them.

## - 3.8 Appendix

## ■ 3.8.1 Rank 1 Updates

Here we prove Lemma 3.5.2, making critical use of the Sherman-Morrison formula for the MoorePenrose pseudoinverse (Lemma 2.3.7).
Lemma 3.5.2 (Leverage Score Changes Under Rank 1 Updates). Given any $\boldsymbol{A} \in \mathbb{R}^{m \times n}, \gamma \in(0,1)$, and $i \in[m]$, let $\boldsymbol{W}$ be a diagonal matrix such that $\boldsymbol{W}_{i i}=\sqrt{1-\gamma}$ and $\boldsymbol{W}_{j j}=1$ for all $j \neq i$. Then, we have

$$
\boldsymbol{\sigma}_{i}(\boldsymbol{W} \boldsymbol{A})=\frac{(1-\gamma) \boldsymbol{\sigma}_{i}(\boldsymbol{A})}{1-\gamma \boldsymbol{\sigma}_{i}(\boldsymbol{A})} \leq \boldsymbol{\sigma}_{i}(\boldsymbol{A})
$$

and for all $j \neq i$,

$$
\boldsymbol{\sigma}_{j}(\boldsymbol{W} \boldsymbol{A})=\boldsymbol{\sigma}_{j}(\boldsymbol{A})+\frac{\gamma \boldsymbol{\sigma}_{i j}(\boldsymbol{A})^{2}}{1-\gamma \boldsymbol{\sigma}_{i}(\boldsymbol{A})} \geq \boldsymbol{\sigma}_{j}(\boldsymbol{A})
$$

Proof.

$$
\begin{array}{rlr}
\boldsymbol{\sigma}_{i}(\boldsymbol{W} \boldsymbol{A}) & =\mathbf{1}_{i} \boldsymbol{W} \boldsymbol{A}\left(\boldsymbol{A}^{\top} \boldsymbol{W}^{2} \boldsymbol{A}\right)^{+} \boldsymbol{A}^{\top} \boldsymbol{W}^{\top} \mathbf{1}_{i}^{\top} & \text { (definition of leverage scores) } \\
& =(1-\gamma) \boldsymbol{a}_{i}^{\top}\left(\boldsymbol{A}^{\top} \boldsymbol{A}-\gamma \boldsymbol{a}_{i} \boldsymbol{a}_{i}^{\top}\right)^{+} \boldsymbol{a}_{i} & \text { (definition of } \boldsymbol{W} \text { ) } \\
& =(1-\gamma) \boldsymbol{a}_{i}^{\top}\left(\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+}+\gamma \frac{\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+} \boldsymbol{a}_{i} \boldsymbol{a}_{i}^{\top}\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+}}{1-\gamma \boldsymbol{a}_{i}^{\top}\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+} \boldsymbol{a}_{i}}\right) \boldsymbol{a}_{i} & \text { (Sherman-Morrison formula) } \\
& =(1-\gamma)\left(\boldsymbol{\sigma}_{i}(\boldsymbol{A})+\frac{\gamma \boldsymbol{\sigma}_{i}(\boldsymbol{A})^{2}}{1-\gamma \boldsymbol{\sigma}_{i}(\boldsymbol{A})}\right) \\
& =\frac{(1-\gamma) \boldsymbol{\sigma}_{i}(\boldsymbol{A})}{1-\gamma \boldsymbol{\sigma}_{i}(\boldsymbol{A})} \leq \boldsymbol{\sigma}_{i}(\boldsymbol{A}) . &
\end{array}
$$

Similarly,

$$
\begin{aligned}
\boldsymbol{\sigma}_{j}(\boldsymbol{W} \boldsymbol{A}) & =\boldsymbol{a}_{j}^{\top}\left(\boldsymbol{A}^{\top} \boldsymbol{A}-\gamma \boldsymbol{a}_{i} \boldsymbol{a}_{i}^{\top}\right)^{+} \boldsymbol{a}_{j} \\
& =\boldsymbol{a}_{j}^{\top}\left(\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+}+\gamma \frac{\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+} \boldsymbol{a}_{\boldsymbol{a}} \boldsymbol{a}_{i}^{\top}\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+}}{1-\gamma \boldsymbol{a}_{i}^{\top}\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{+} \boldsymbol{a}_{i}}\right) \boldsymbol{a}_{j} \\
& =\boldsymbol{\sigma}_{j}(\boldsymbol{A})+\frac{\gamma \boldsymbol{\sigma}_{i j}(\boldsymbol{A})^{2}}{1-\gamma \boldsymbol{\sigma}_{i}(\boldsymbol{A})} \geq \boldsymbol{\sigma}_{j}(\boldsymbol{A}) .
\end{aligned}
$$

## - 3.8.2 Lower Semi-continuity of Leverage Scores

Here we prove Lemma 3.5 .3 by providing a fairly general inequality, Lemma 3.8.1, for relating leverage scores under one set of weights to leverage scores under another.
Lemma 3.5.3 (Leverage Scores are Lower Semi-continuous). $\boldsymbol{\sigma}(\boldsymbol{W} \boldsymbol{A})$ is lower semi-continuous in the diagonal matrix $\boldsymbol{W}$, i.e. for any sequence $\boldsymbol{W}^{(k)} \rightarrow \overline{\boldsymbol{W}}$ with $\boldsymbol{W}_{i i}^{(k)} \geq 0$ for all $k$ and $i$, we have

$$
\begin{equation*}
\boldsymbol{\sigma}_{i}(\overline{\boldsymbol{W}} \boldsymbol{A}) \leq \liminf _{k \rightarrow \infty} \boldsymbol{\sigma}_{i}\left(\boldsymbol{W}^{(k)} \boldsymbol{A}\right) . \tag{3.7}
\end{equation*}
$$

Lemma 3.8.1 (Comparing Leverage Scores). Let $\boldsymbol{W}, \overline{\boldsymbol{W}} \in \mathbb{R}^{m \times m}$ be non-negative diagonal matrices and suppose that $\boldsymbol{W}_{i i}>0$ and $\overline{\boldsymbol{W}}_{i i}>0$ for some $i \in[m]$. Then

$$
\begin{equation*}
\boldsymbol{\sigma}_{i}(\overline{\boldsymbol{W}} \boldsymbol{A}) \leq \frac{\overline{\boldsymbol{W}}_{i i}^{2}}{\boldsymbol{W}_{i i}^{2}}\left(1+\sqrt{\lambda_{\max }\left(\boldsymbol{A}\left(\boldsymbol{A}^{\top} \overline{\boldsymbol{W}}^{2} \boldsymbol{A}\right)^{+} \boldsymbol{A}^{\top}\right)}\|\boldsymbol{W}-\overline{\boldsymbol{W}}\|_{\infty}\right)^{2} \boldsymbol{\sigma}_{i}(\boldsymbol{W} \boldsymbol{A}) \tag{3.8}
\end{equation*}
$$

Proof. Scaling the variables in Lemma 2.3.1 we have that there exists $\boldsymbol{x} \in \mathbb{R}^{m}$ such that

$$
\begin{equation*}
\boldsymbol{A}^{\top} \boldsymbol{W} \boldsymbol{x}=\boldsymbol{a}_{i} \quad \text { and } \quad\|\boldsymbol{x}\|_{2}^{2}=\frac{\boldsymbol{\sigma}_{i}(\boldsymbol{W} \boldsymbol{A})}{\boldsymbol{W}_{i i}^{2}} \tag{3.9}
\end{equation*}
$$

Note that $\boldsymbol{A}^{\top}(\boldsymbol{W}-\overline{\boldsymbol{W}}) \boldsymbol{x}$ is in the image of $\boldsymbol{A}^{\top} \overline{\boldsymbol{W}}$ as $\boldsymbol{A}^{\top} \boldsymbol{W} \boldsymbol{x}=\boldsymbol{a}_{i}$ and $\overline{\boldsymbol{W}}_{i i} \neq 0$. Consequently,

$$
\boldsymbol{A}^{\top} \overline{\boldsymbol{W}} \boldsymbol{y}=\boldsymbol{A}^{\top}(\boldsymbol{W}-\overline{\boldsymbol{W}}) \boldsymbol{x}^{(k)} \quad \text { for } \quad \boldsymbol{y} \stackrel{\text { def }}{=} \overline{\boldsymbol{W}} \boldsymbol{A}\left(\boldsymbol{A}^{\top} \overline{\boldsymbol{W}}^{2} \boldsymbol{A}\right)^{+} \boldsymbol{A}^{\top}(\boldsymbol{W}-\overline{\boldsymbol{W}}) \boldsymbol{x}
$$

Since $\boldsymbol{A}^{\top} \overline{\boldsymbol{W}}(\boldsymbol{x}+\boldsymbol{y})=\boldsymbol{a}_{i}$, Lemma 2.3.1 implies

$$
\begin{equation*}
\boldsymbol{\sigma}_{i}(\overline{\boldsymbol{W}} \boldsymbol{A}) \leq \overline{\boldsymbol{W}}_{i i}^{2}\|\boldsymbol{x}+\boldsymbol{y}\|_{2}^{2} \tag{3.10}
\end{equation*}
$$

We can bound the contribution of $\boldsymbol{y}$ by

$$
\begin{align*}
\|\boldsymbol{y}\|_{2}^{2} & \leq\left\|\overline{\boldsymbol{W}} \boldsymbol{A}\left(\boldsymbol{A}^{\top} \overline{\boldsymbol{W}}^{2} \boldsymbol{A}\right)^{+} \boldsymbol{A}^{\top}(\boldsymbol{W}-\overline{\boldsymbol{W}}) \boldsymbol{x}\right\|_{2}^{2} \\
& \leq \lambda_{\max }\left(\boldsymbol{A}\left(\boldsymbol{A}^{\top} \overline{\boldsymbol{W}}^{2} \boldsymbol{A}\right)^{+} \boldsymbol{A}^{\top}\right)\|\boldsymbol{W}-\overline{\boldsymbol{W}}\|_{\infty}^{2}\|\boldsymbol{x}\|_{2}^{2} \tag{3.11}
\end{align*}
$$

Applying triangle inequality to (3.9), (3.10), and (3.11) yields the result.
Proof of Lemma 3.5.3. For any $i \in[n]$ such that $\overline{\boldsymbol{W}}_{i i}=0$ (3.7) follows trivially from the fact that leverage scores are non-negative. For any $i \in[m]$ such that $\overline{\boldsymbol{W}}_{i i}>0$, since $\boldsymbol{W}^{(k)} \rightarrow \overline{\boldsymbol{W}}$ we know that, for all sufficiently large $k \geq N$ for some fixed value $N$, it is the case that $\boldsymbol{W}_{i i}^{(k)}>0$. Furthermore, this implies that as $k \rightarrow \infty$ we have $\overline{\boldsymbol{W}}_{i i}^{2} /\left(\boldsymbol{W}_{i i}^{(k)}\right)^{2} \rightarrow 1$ and $\left\|\boldsymbol{W}^{(k)}-\overline{\boldsymbol{W}}\right\|_{\infty} \rightarrow 0$. Applying Lemma 3.8.1 with $\boldsymbol{W}=\boldsymbol{W}^{(k)}$ and taking $\liminf _{k \rightarrow \infty}$ on both sides of (3.8) gives the result.

# Linear-Sized Sparsifier In Almost-Linear Time 

## ■ 4.1 Introduction

In the last chapter, we showed how to compute $O(n \log n)$ sized spectral sparsifier in input sparsity time. In this chapter, we show how to construct a spectral sparsifier with even smaller size. In particular, we present an almost-linear time algorithm for constructing linear-sized spectral sparsifiers for graphs, which improves all previous constructions that either require $\Omega\left(n^{2+\varepsilon}\right)$ time in order to produce linearsized sparsifiers [9, 31, 278], or $O\left(m \log ^{O(1)} n / \varepsilon^{2}\right)$ time but the number of edges in the sparsifiers is sub-optimal [242]. Our algorithm is conceptually simple, and is based on a novel combination of two techniques used in literature for constructing spectral sparsifiers: random sampling by leverage scores we used in the last chapter, and adaptive construction based on barrier functions [9, 31]. Our result is summarized as follows:
Theorem 4.1.1. Let $q \geq 1$ and $0<\varepsilon \leq 1$ be constants, and $\boldsymbol{I}=\sum_{i=1}^{m} \boldsymbol{v}_{i} \boldsymbol{v}_{i}^{T}$ be the sum of $m$ rank-1 PSD matrices. Then, there is an algorithm that outputs scalars $\left\{s_{i}\right\}_{i=1}^{m}$ with $\left|\left\{s_{i}: s_{i} \neq 0\right\}\right|=O\left(q n / \varepsilon^{2}\right)$ such that

$$
(1-\varepsilon) \cdot \boldsymbol{I} \preceq \sum_{i=1}^{m} s_{i} \boldsymbol{v}_{i} \boldsymbol{v}_{i}^{T} \preceq(1+\varepsilon) \cdot I .
$$

The algorithm runs in $\tilde{O}\left(Z+n^{\omega+1 / q} \varepsilon^{-4}\right)$ time, where $Z$ is the total number of non-zeros in $\boldsymbol{v}_{i}$.
Theorem 4.1.2. Let $q \geq 1$ and $0<\varepsilon \leq 1$ be constants, and $G=(V, E, \boldsymbol{w})$ be an undirected and weighted graph with $n$ vertices and $m$ edges. Then, there is an algorithm that outputs a $\varepsilon$-spectral sparsifier of $G$ with $O\left(q n / \varepsilon^{2}\right)$ edges. The algorithm runs in $\tilde{O}\left(m+n^{1+1 / q} \varepsilon^{-(6+1 / q)}\right)$ time.

## - 4.2 Algorithm

In this section we study the algorithm of sparsifying the sum of rank-1 PSD matrices. Remember that our goal is to, for any vectors $\boldsymbol{v}_{1}, \cdots \boldsymbol{v}_{m}$ with $\sum_{i=1}^{m} \boldsymbol{v}_{i} \boldsymbol{v}_{i}^{T}=I$, find scalars $\left\{s_{i}\right\}_{i=1}^{m}$ satisfying

$$
\left|\left\{s_{i}: s_{i} \neq 0\right\}\right|=O\left(\frac{q n}{\varepsilon^{2}}\right),
$$

such that

$$
(1-\varepsilon) \cdot \boldsymbol{I} \preceq \sum_{i=1}^{m} s_{i} \boldsymbol{v}_{i} \boldsymbol{v}_{i}^{T} \preceq(1+\varepsilon) \cdot \boldsymbol{I} .
$$

## - 4.2.1 Overview of our approach

At a high level, our algorithm can be viewed as an improved and randomized version of the algorithm presented in Batson et al. [31]. We refer their algorithm BSS for short, and first give a brief overview of the BSS algorithm.

The BSS algorithm proceeds by iterations, and the matrix $\boldsymbol{A}_{j}$ in iteration $j \geq 1$ is defined as the sum of $\boldsymbol{A}_{j-1}$ and a proper chosen rank-1 matrix. To control the spectral properties of matrix $\boldsymbol{A}_{j}$, the algorithm introduces two barrier values $u_{j}$ and $\ell_{j}$, where $u_{0}>0, \ell_{0}<0$ initially. It was proven that one can always find a vector in $\left\{\boldsymbol{v}_{i}\right\}_{i=1}^{m}$ and update $u_{j}, \ell_{j}$ in a proper manner in each iteration, such that the invariant

$$
\begin{equation*}
\ell_{j} \boldsymbol{I} \prec \boldsymbol{A}_{j} \prec u_{j} \boldsymbol{I} \tag{4.1}
\end{equation*}
$$

always holds [31]. To quantitatively measure "how close the eigenvalues of $\boldsymbol{A}$ are to the barriers $u$ and $\ell^{\prime \prime}$, Batson et al. [31] analyzed the potential function defined by

$$
\begin{equation*}
\Phi_{u, \ell}(\boldsymbol{A}) \triangleq \operatorname{Tr}(u \boldsymbol{I}-\boldsymbol{A})^{-1}+\operatorname{Tr}(\boldsymbol{A}-\ell \boldsymbol{I})^{-1} \tag{4.2}
\end{equation*}
$$

Notice that the value of $\Phi_{u, \ell}(\boldsymbol{A})$ is large if and only if some eigenvalue of $\boldsymbol{A}$ is close to $u$ or $\ell$. With the help of this potential function, it was shown that, when updating $\boldsymbol{A}_{j}$ and barrier values $u_{j}, \ell_{j}$ properly, it holds after $k=\Theta\left(n / \varepsilon^{2}\right)$ iterations that $\ell_{k} \geq c u_{k}$ for some constant $c$. This implies that the resulting matrix $\boldsymbol{A}_{k}$ is a linear-sized and $\boldsymbol{A}_{k} \approx_{O(\varepsilon)} \boldsymbol{I}$.

However, the BSS algorithm is deterministic, and its time complexity for finding a desired rank-1 matrix in each iteration is high. To improve the runtime and further discuss our approach, let us study the following randomized variant of the BSS algorithm: In each iteration, we choose a vector $\boldsymbol{v}_{i}$ with probability $p_{i}$, and add a rank- 1 matrix

$$
\boldsymbol{\Delta}_{\boldsymbol{A}} \triangleq \frac{\varepsilon}{\operatorname{Tr}(u \boldsymbol{I}-\boldsymbol{A})^{-1}+\operatorname{Tr}(\boldsymbol{A}-\ell \boldsymbol{I})^{-1}} \cdot \frac{1}{p_{i}} \cdot \boldsymbol{v}_{i} \boldsymbol{v}_{i}^{T}
$$

to the current matrix $A$. See Algorithm 3 for formal description.

```
Algorithm 3: Randomized BSS algorithm
    Set \(j=0, \ell_{0}=-8 n / \varepsilon, u_{0}=8 n / \varepsilon, \boldsymbol{A}_{0}=\mathbf{0}\).
    while \(u_{j}-\ell_{j}<8 n / \varepsilon\) do
        Let \(t=\operatorname{Tr}\left(u_{j} \boldsymbol{I}-\boldsymbol{A}_{j}\right)^{-1}+\operatorname{Tr}\left(\boldsymbol{A}_{j}-\ell_{j} \boldsymbol{I}\right)^{-1}\).
        Sample a vector \(\boldsymbol{v}_{i}\) with probability
                \(p_{i} \triangleq\left(\boldsymbol{v}_{i}^{T}\left(u_{j} \boldsymbol{I}-\boldsymbol{A}_{j}\right)^{-1} \boldsymbol{v}_{i}+\boldsymbol{v}_{i}^{T}\left(\boldsymbol{A}_{j}-\ell_{j} \boldsymbol{I}\right)^{-1} \boldsymbol{v}_{i}\right) / t\).
            \(\boldsymbol{A}_{j+1}=\boldsymbol{A}_{j}+\frac{\varepsilon}{t} \cdot \frac{1}{p_{i}} \cdot \boldsymbol{v}_{i} \boldsymbol{v}_{i}^{T} ;\)
        \(u_{j+1}=u_{j}+\frac{\varepsilon}{t \cdot(1-\varepsilon)}\) and \(\ell_{j+1}=\ell_{j}+\frac{\varepsilon}{t \cdot(1+\varepsilon)}\);
        \(j \leftarrow j+1 ;\)
    end
    Output: \(\boldsymbol{A}_{j}\).
```

Let us look at a fixed iteration $j$, and analyze how the added $\boldsymbol{\Delta}_{\boldsymbol{A}}$ impacts the potential function. We drop the subscript representing the iteration $j$ for simplicity. After adding $\boldsymbol{\Delta}_{\boldsymbol{A}}$, the first-order approximation of $\Phi_{u, \ell}(\boldsymbol{A})$ gives that

$$
\begin{equation*}
\Phi_{u, \ell}\left(\boldsymbol{A}+\boldsymbol{\Delta}_{\boldsymbol{A}}\right) \sim \Phi_{u, \ell}(\boldsymbol{A})+(u \boldsymbol{I}-\boldsymbol{A})^{-2} \bullet \boldsymbol{\Delta}_{\boldsymbol{A}}-(\boldsymbol{A}-\ell \boldsymbol{I})^{-2} \bullet \boldsymbol{\Delta}_{\boldsymbol{A}} . \tag{4.3}
\end{equation*}
$$

Since

$$
\mathbf{E}\left[\boldsymbol{\Delta}_{\boldsymbol{A}}\right]=\sum_{i=1}^{m} p_{i} \cdot\left(\frac{\varepsilon}{t} \cdot \frac{1}{p_{i}} \cdot \boldsymbol{v}_{i} \boldsymbol{v}_{i}^{T}\right)=\frac{\varepsilon}{t} \cdot \sum_{i=1}^{m} \boldsymbol{v}_{i} \boldsymbol{v}_{i}^{T}=\frac{\varepsilon}{t} \cdot I,
$$

we have that

$$
\begin{aligned}
\mathbf{E}\left[\Phi_{u, \ell}\left(\boldsymbol{A}+\boldsymbol{\Delta}_{\boldsymbol{A}}\right)\right] & \sim \Phi_{u, \ell}(\boldsymbol{A})+\frac{\varepsilon}{t} \cdot(u \boldsymbol{I}-\boldsymbol{A})^{-2} \bullet I-\frac{\varepsilon}{t} \cdot(\boldsymbol{A}-\ell \boldsymbol{I})^{-2} \bullet I \\
& =\Phi_{u, \ell}(\boldsymbol{A})+\frac{\varepsilon}{t} \cdot \operatorname{Tr}(u \boldsymbol{I}-\boldsymbol{A})^{-2}-\frac{\varepsilon}{t} \cdot \operatorname{Tr}(\boldsymbol{A}-\ell \boldsymbol{I})^{-2} \\
& =\Phi_{u, \ell}(\boldsymbol{A})-\frac{\varepsilon}{t} \cdot \frac{\mathrm{~d}}{\mathrm{~d} u} \Phi_{u, \ell}(\boldsymbol{A})-\frac{\varepsilon}{t} \cdot \frac{\mathrm{~d}}{\mathrm{~d} \ell} \Phi_{u, \ell}(\boldsymbol{A})
\end{aligned}
$$

Notice that if we increase $u$ by $\varepsilon / t$ and $\ell$ by $\varepsilon / t, \Phi_{u, \ell}$ approximately increases by

$$
\frac{\varepsilon}{t} \cdot \frac{\mathrm{~d}}{\mathrm{~d} u} \Phi_{u, \ell}(\boldsymbol{A})+\frac{\varepsilon}{t} \cdot \frac{\mathrm{~d}}{\mathrm{~d} \ell} \Phi_{u, \ell}(\boldsymbol{A})
$$

Hence, comparing $\Phi_{u+\varepsilon / t, \ell+\varepsilon / t}\left(\boldsymbol{A}+\boldsymbol{\Delta}_{\boldsymbol{A}}\right)$ with $\Phi_{u, \ell}(\boldsymbol{A})$, the increase of the potential function due to the change of barrier values is approximately compensated by the decrease of the potential function by the effect of $\boldsymbol{\Delta}_{\boldsymbol{A}}$. For a more rigorous analysis, we need to analyze the higher-order terms and increase $u$ slightly more than $\ell$. Batson et al. [31] gives the following estimate:
Lemma 4.2.1 ([31], proof of Lemma 3.3 and 3.4). Let $\boldsymbol{A} \in \mathbb{R}^{n \times n}$, and u, $\ell$ be parameters satisfying $\ell \boldsymbol{I} \prec \boldsymbol{A} \prec u \boldsymbol{I}$. Suppose that $\boldsymbol{w} \in \mathbb{R}^{n}$ satisfies $\boldsymbol{w} \boldsymbol{w}^{T} \preceq \delta(u \boldsymbol{I}-\boldsymbol{A})$ and $\boldsymbol{w} \boldsymbol{w}^{T} \preceq \delta(\boldsymbol{A}-\ell \boldsymbol{I})$ for some $0<\delta<1$. Then, it holds that

$$
\Phi_{u, \ell}\left(\boldsymbol{A}+\boldsymbol{w} \boldsymbol{w}^{T}\right) \leq \Phi_{u, \ell}(\boldsymbol{A})+\frac{\boldsymbol{w}^{T}(u \boldsymbol{I}-\boldsymbol{A})^{-2} \boldsymbol{w}}{1-\delta}-\frac{\boldsymbol{w}^{T}(\boldsymbol{A}-\ell \boldsymbol{I})^{-2} \boldsymbol{w}}{1+\delta}
$$

The estimate above shows that the first-order approximation (4.3) is good as long as $\boldsymbol{w} \boldsymbol{w}^{T} \preceq$ $\delta(u \boldsymbol{I}-\boldsymbol{A})$ and $\boldsymbol{w} \boldsymbol{w}^{T} \preceq \delta(\boldsymbol{A}-\ell \boldsymbol{I})$ for small $\delta$. By setting $\delta=\varepsilon$, it is easy to see that the added matrix $\boldsymbol{\Delta}_{\boldsymbol{A}}$ satisfies the preconditions in Lemma 4.2.1, since

$$
\frac{\varepsilon}{t} \cdot \frac{1}{p_{i}} \cdot \boldsymbol{v}_{i} \boldsymbol{v}_{i}^{T}=\frac{\varepsilon \cdot \boldsymbol{v}_{i} \boldsymbol{v}_{i}^{T}}{v_{i}^{T}(u \boldsymbol{I}-\boldsymbol{A})^{-1} v_{i}+\boldsymbol{v}_{i}^{T}(\boldsymbol{A}-\ell \boldsymbol{I})^{-1} \boldsymbol{v}_{i}} \preceq \frac{\varepsilon \cdot v_{i} v_{i}^{T}}{v_{i}^{T}(u \boldsymbol{I}-\boldsymbol{A})^{-1} v_{i}} \preceq \varepsilon(u \boldsymbol{I}-\boldsymbol{A}) .
$$

Here we used the fact that $v v^{T} \preceq\left(v^{T} B^{-1} v\right) B$ for any vector $v$ and PSD matrix $B$. Similarly, we have that

$$
\frac{\varepsilon}{t} \cdot \frac{1}{p_{i}} \cdot v_{i} v_{i}^{T} \preceq \varepsilon(A-\ell I) .
$$

Hence, if the initial value of the potential function is small, $\Phi_{u, \ell}(A)$ is small throughout the executions of the whole algorithm. Up to a constant factor, this gives the same result as [31], and Algorithm 3 constructs an $\Theta\left(n / \varepsilon^{2}\right)$-sized $O(\varepsilon)$-spectral sparsifier.

However, Algorithm 3 runs for $\Theta\left(n / \varepsilon^{2}\right)$ iterations, and in each iteration the algorithm re-computes the probability distribution for sampling vectors. To improve the runtime of Algorithm 3, we need to overcome two bottlenecks: (1) we need a fast algorithm to approximate the sampling probability $\left\{p_{i}\right\}_{i=1}^{m}$ of vectors; (2) we need to reduce the number of iterations required by the algorithm, i.e. re-using the sampling probability $\left\{p_{i}\right\}_{i=1}^{m}$ for few iterations.

For fast approximation of the sampling probabilities, we adopt the idea proposed in [9]: instead of defining the potential function by (4.2), we define the potential function by

$$
\begin{equation*}
\Phi_{u, \ell}(\boldsymbol{A}) \triangleq \operatorname{Tr}(u \boldsymbol{I}-\boldsymbol{A})^{-q}+\operatorname{Tr}(\boldsymbol{A}-\ell \boldsymbol{I})^{-q} \tag{4.4}
\end{equation*}
$$

Since $q$ is a large constant, the value of the potential function becomes larger when some eigenvalue of $A$ is closer to $u$ or $\ell$. Hence, a bounded value of $\Phi_{u, \ell}(A)$ insures that the eigenvalues of $A$ never get too close to $u$ or $\ell$, which allows us to compute the sampling probabilities $\left\{p_{i}\right\}_{i=1}^{m}$ efficiently simply by Taylor expansion. Moreover, with our new potential function (4.4) one can prove a similar result
as Lemma 4.2.1. This gives an alternative analysis of the algorithm presented in [9], which is the first almost-quadratic time algorithm for constructing linear-sized spectral sparsifiers.

To overcome the second bottleneck, we re-compute $\left\{p_{i}\right\}_{i=1}^{m}$ after every $\Theta\left(n^{1-1 / q}\right)$ iterations: we show that as long as the sampling probability satisfies

$$
p_{i} \geq C \cdot \frac{\boldsymbol{v}_{i}^{T}(u \boldsymbol{I}-\boldsymbol{A})^{-1} \boldsymbol{v}_{i}+\boldsymbol{v}_{i}^{T}(\boldsymbol{A}-\ell \boldsymbol{I})^{-1} \boldsymbol{v}_{i}}{\sum_{i=1}^{m}\left(\boldsymbol{v}_{i}^{T}(u \boldsymbol{I}-\boldsymbol{A})^{-1} \boldsymbol{v}_{i}+\boldsymbol{v}_{i}^{T}(\boldsymbol{A}-\ell \boldsymbol{I})^{-1} \boldsymbol{v}_{i}\right)}
$$

for some constant $C>0$, we can still sample $v_{i}$ with probability $p_{i}$ and get the same guarantee on the potential function. The reason is as follows: assume that $\boldsymbol{\Delta}_{\boldsymbol{A}}=\sum_{i=1}^{T} \boldsymbol{\Delta}_{\boldsymbol{A}, i}$ is the sum of the sampled matrices within $T=O\left(n^{1-1 / q}\right)$ iterations. If a randomly chosen matrix $\boldsymbol{\Delta}_{\boldsymbol{A}, i}$ satisfies $\boldsymbol{\Delta}_{\boldsymbol{A}, i} \preceq$ $\frac{1}{C q}(u \boldsymbol{I}-\boldsymbol{A})$, then by the matrix Chernoff bound $\boldsymbol{\Delta}_{\boldsymbol{A}} \preceq \frac{1}{2}(u \boldsymbol{I}-\boldsymbol{A})$ holds with high probability. By scaling every sampled rank-1 matrix $q$ times smaller, the sampling probability only changes by a constant factor within $T$ iterations. Since we choose $\Theta\left(n / \varepsilon^{2}\right)$ vectors in total, our algorithm only recomputes the sampling probabilities $\Theta\left(n^{1 / q} / \varepsilon^{2}\right)$ times.

## - 4.2.2 Algorithm description

Our actual algorithm follows the same framework as Algorithm 3, and proceeds by iterations. Initially, the algorithm sets

$$
u_{0} \triangleq(2 n)^{1 / q}, \quad \ell_{0} \triangleq-(2 n)^{1 / q}, \quad A_{0} \triangleq \mathbf{0} .
$$

After iteration $j$ the algorithm updates $u_{j}, \ell_{j}$ by $\Delta_{u, j}, \Delta_{\ell, j}$ respectively, i.e.,

$$
u_{j+1} \triangleq u_{j}+\Delta_{u, j}, \quad \ell_{j+1} \triangleq \ell_{j}+\Delta_{\ell, j},
$$

and updates $A_{j}$ with respect to the chosen matrix in iteration $j$. The choice of $\Delta_{u, j}$ and $\Delta_{\ell, j}$ insures that the invariant

$$
\ell_{j} I \prec \boldsymbol{A}_{j} \prec u_{j} I
$$

always holds for any iteration $j$. In iteration $j$, the algorithm computes the relative leverage score of vectors $\left\{v_{i}\right\}_{i=1}^{m}$ defined by

$$
R_{i}\left(\boldsymbol{A}_{j}, u_{j}, \ell_{j}\right) \triangleq \boldsymbol{v}_{i}^{T}\left(u_{j} \boldsymbol{I}-\boldsymbol{A}_{j}\right)^{-1} \boldsymbol{v}_{i}+\boldsymbol{v}_{i}^{T}\left(\boldsymbol{A}_{j}-\ell_{j} \boldsymbol{I}\right)^{-1} \boldsymbol{v}_{i},
$$

and samples $N_{j}$ vectors independently with replacement, where vector $v_{i}$ is chosen with probability proportional to $R_{i}\left(A_{j}, u_{j}, \ell_{j}\right)$, and

$$
N_{j} \triangleq \frac{1}{n^{2 / q}}\left(\sum_{i=1}^{m} R_{i}\left(\boldsymbol{A}_{j}, u_{j}, \ell_{j}\right)\right) \min \left\{\lambda_{\min }\left(u_{j} \boldsymbol{I}-\boldsymbol{A}_{j}\right), \lambda_{\min }\left(\boldsymbol{A}_{j}-\ell_{j} \boldsymbol{I}\right)\right\} .
$$

The algorithm sets $\boldsymbol{A}_{j+1}$ to be the sum of $\boldsymbol{A}_{j}$ and sampled $\boldsymbol{v}_{i} \boldsymbol{v}_{i}^{T}$ with proper reweighting. For technical reasons, we define $\Delta_{u, j}$ and $\Delta_{\ell, j}$ by

$$
\Delta_{u, j} \triangleq(1+2 \varepsilon) \cdot \frac{\varepsilon \cdot N_{j}}{q \cdot \sum_{i=1}^{m} R_{i}\left(\boldsymbol{A}_{j}, u_{j}, \ell_{j}\right)}, \quad \Delta_{\ell, j} \triangleq(1-2 \varepsilon) \cdot \frac{\varepsilon \cdot N_{j}}{q \cdot \sum_{i=1}^{m} R_{i}\left(\boldsymbol{A}_{j}, u_{j}, \ell_{j}\right)} .
$$

See Algorithm 4 for formal description.
We remark that, although exact values of $N_{j}$ and relative leverage scores are difficult to compute in almost-linear time, we can use approximated values of $\left\{R_{i}\right\}_{i=1}^{m}$ and $N_{j}$ instead. It is easy to see that in each iteration an over estimate of the relative leverage score for every vector $\boldsymbol{v}_{i}$, and an under estimate of $N_{j}$ with constant-factor approximation suffice for our purpose.

```
Algorithm 4: Algorithm for constructing spectral sparsifiers
    Assume: \(0<\varepsilon \leq 1 / 120, q \geq 10\).
    Set \(j=0, \ell_{0}=-(2 n)^{1 / q}, u_{0}=(2 n)^{1 / q}, \boldsymbol{A}_{0}=\mathbf{0}\).
    while \(u_{j}-\ell_{j}<4 \cdot(2 n)^{1 / q}\) do
        Let \(\boldsymbol{W}_{j}=\mathbf{0}\).
        Compute \(R_{i}\left(\boldsymbol{A}_{j}, u_{j}, \ell_{j}\right)\) for all vectors \(v_{i}\).
        Sample \(N_{j}\) vectors independently with replacement, where every \(v_{i}\) is chosen with
        probability proportional to \(R_{i}\left(\boldsymbol{A}_{j}, u_{j}, \ell_{j}\right)\).
        For every sampled \(v\), add \(\varepsilon / q \cdot\left(R_{i}\left(\boldsymbol{A}_{j}, u_{j}, \ell_{j}\right)\right)^{-1} \cdot v v^{T}\) to \(\boldsymbol{W}_{j}\).
        \(\boldsymbol{A}_{j+1}=\boldsymbol{A}_{j}+\boldsymbol{W}_{j}\).
        \(u_{j+1}=u_{j}+\Delta_{u, j}, \ell_{j+1}=\ell_{j}+\Delta_{\ell, j}\).
        \(j=j+1\).
    end
    Output: \(\boldsymbol{A}_{j}\).
```


### 4.3 Analysis

In this section we prove that the output matrix returned by Algorithm 4 is a $\varepsilon$-spectral sparsifier, and analyze the algorithm's runtime. To simplify our analysis, we always assume in the rest of the chapter that $\varepsilon$ and $q$ are constants such that $0<\varepsilon \leq 1 / 120$, and $q \geq 10$.

This section is organized as follows: we will first show how the potential function (4.4) evolves after each iteration in Section 4.3.1. Combing this with the ending condition of the algorithm, we will prove in Section 4.3.2 that the algorithm outputs a linear-sized spectral sparsifier. Finally, we will prove Theorem 4.1.2 and Theorem 4.1.1 in Section 4.3.3. In appendix, we show almost-linear time algorithms for approximately computing all required quantities of Algorithm 4 within each iteration when constructing graph sparsifiers.

## ■ 4.3.1 Analysis of a single iteration

We analyze the sampling scheme within a single iteration, and drop the subscript representing the iteration $j$ for simplicity. Recall that in each iteration the algorithm samples $N$ vectors independently from $\left\{\boldsymbol{v}_{i}\right\}_{i=1}^{m}$ satisfying $\sum_{i=1}^{m} \boldsymbol{v}_{i} \boldsymbol{v}_{i}^{T}=\boldsymbol{I}$, where every vector $\boldsymbol{v}_{i}$ is sampled with probability

$$
\frac{R_{i}(\boldsymbol{A}, u, \ell)}{\sum_{j=1}^{m} R_{j}(\boldsymbol{A}, u, \ell)} .
$$

We use $\boldsymbol{v}_{1}, \cdots, \boldsymbol{v}_{N}$ to denote these $N$ sampled vectors, and define the reweighted vectors by

$$
\boldsymbol{w}_{i} \triangleq \sqrt{\frac{\varepsilon}{q \cdot R_{i}(\boldsymbol{A}, u, \ell)}} \cdot \boldsymbol{v}_{i},
$$

for any $1 \leq i \leq N$. Let

$$
\boldsymbol{W} \triangleq \sum_{i=1}^{N} \boldsymbol{w}_{i} \boldsymbol{w}_{i}^{T}
$$

and we use $\boldsymbol{W} \sim \mathcal{D}(\boldsymbol{A}, u, \ell)$ to represent that $\boldsymbol{W}$ is sampled in this way with parameters $\boldsymbol{A}, u$ and $\ell$. We will show that with high probability matrix $\boldsymbol{W}$ satisfies $\mathbf{0} \preceq \boldsymbol{W} \preceq \frac{1}{2}(u \boldsymbol{I}-\boldsymbol{A})$. This is basically follows from Matrix Chernoff Bound (Lemma 2.3.10).

Lemma 4.3.1. Assume that the number of sampled vectors satisfies

$$
N<\frac{2}{n^{2 / q}}\left(\sum_{i=1}^{m} R_{i}(\boldsymbol{A}, u, \ell)\right) \cdot \lambda_{\min }(u \boldsymbol{I}-\boldsymbol{A}) .
$$

Then, it holds that $\mathbf{E}[\boldsymbol{W}]=\frac{\varepsilon}{q} \cdot \frac{N}{\sum_{i=1}^{m} R_{i}(\boldsymbol{A}, u, \ell)} \cdot \boldsymbol{I}$ and

$$
\operatorname{Pr}\left[\mathbf{0} \preceq \boldsymbol{W} \preceq \frac{1}{2} \cdot(u \boldsymbol{I}-\boldsymbol{A})\right] \geq 1-\frac{\varepsilon^{2}}{100 q n} .
$$

Proof. By the description of the sampling procedure, the first statement follows from the calculation

$$
\mathbf{E}\left[\boldsymbol{w}_{i} \boldsymbol{w}_{i}^{T}\right]=\sum_{j=1}^{m} \frac{R_{j}(\boldsymbol{A}, u, \ell)}{\sum_{t=1}^{m} R_{t}(\boldsymbol{A}, u, \ell)} \cdot \frac{\varepsilon}{q} \cdot \frac{\boldsymbol{v}_{j} \boldsymbol{v}_{j}^{T}}{R_{j}(\boldsymbol{A}, u, \ell)}=\frac{\varepsilon}{q} \cdot \frac{1}{\sum_{t=1}^{m} R_{t}(\boldsymbol{A}, u, \ell)} \cdot \boldsymbol{I},
$$

and

$$
\mathbf{E}[\boldsymbol{W}]=\mathbf{E}\left[\sum_{i=1}^{N} \boldsymbol{w}_{i} \boldsymbol{w}_{i}^{T}\right]=\frac{\varepsilon}{q} \cdot \frac{N}{\sum_{i=1}^{m} R_{i}(\boldsymbol{A}, u, \ell)} \cdot \boldsymbol{I} .
$$

Now for the second statement. Let $\boldsymbol{z}_{i}=(u \boldsymbol{I}-\boldsymbol{A})^{-1 / 2} \boldsymbol{w}_{i}$. It holds that

$$
\begin{aligned}
\operatorname{Tr}\left(\boldsymbol{z}_{i} \boldsymbol{z}_{i}^{T}\right) & =\frac{\varepsilon}{q} \cdot \frac{\operatorname{Tr}\left((u \boldsymbol{I}-\boldsymbol{A})^{-1 / 2} \boldsymbol{v}_{i} \boldsymbol{v}_{i}^{T}(u \boldsymbol{I}-\boldsymbol{A})^{-1 / 2}\right)}{R_{i}(A, u, \ell)} \\
& \leq \frac{\varepsilon}{q} \cdot \frac{\boldsymbol{v}_{i}^{T}(u \boldsymbol{I}-\boldsymbol{A})^{-1} \boldsymbol{v}_{i}}{\boldsymbol{v}_{i}^{T}(u \boldsymbol{I}-\boldsymbol{A})^{-1} \boldsymbol{v}_{i}+\boldsymbol{v}_{i}^{T}(\boldsymbol{A}-\ell \boldsymbol{I})^{-1} \boldsymbol{v}_{i}} \leq \frac{\varepsilon}{q},
\end{aligned}
$$

and $\lambda_{\max }\left(\boldsymbol{z}_{i} \boldsymbol{z}_{i}^{T}\right) \leq \frac{\varepsilon}{q}$. Moreover, it holds that

$$
\begin{equation*}
\mathbf{E}\left[\sum_{i=1}^{N} \boldsymbol{z}_{i} \boldsymbol{z}_{i}^{T}\right] \preceq \frac{\varepsilon}{q} \cdot \frac{N}{\sum_{t=1}^{m} R_{t}(\boldsymbol{A}, u, \ell)} \cdot \lambda_{\max }\left(\frac{1}{u \boldsymbol{I}-\boldsymbol{A}}\right) \cdot \boldsymbol{I} . \tag{4.5}
\end{equation*}
$$

This implies that

$$
\lambda_{\max }\left(\mathbf{E}\left[\sum_{i=1}^{N} \boldsymbol{z}_{i} \boldsymbol{z}_{i}^{T}\right]\right) \leq \frac{\varepsilon}{q} \cdot \frac{N}{\sum_{t=1}^{m} R_{t}(\boldsymbol{A}, u, \ell)} \cdot \lambda_{\max }\left(\frac{1}{u \boldsymbol{I}-\boldsymbol{A}}\right) .
$$

Setting $\mu=\frac{\varepsilon}{q} \cdot \frac{N}{\sum_{i=1}^{m} R_{i}(\boldsymbol{A}, u, \ell)} \cdot \lambda_{\max }\left(\frac{1}{u \boldsymbol{I}-\boldsymbol{A}}\right)$, the Matrix Chernoff Bound (Lemma 2.3.10) shows that

$$
\operatorname{Pr}\left[\lambda_{\max }\left(\sum_{i=1}^{N} \boldsymbol{z}_{i} \boldsymbol{z}_{i}^{T}\right) \geq(1+\delta) \mu\right] \leq n \cdot\left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{\mu \cdot q / \varepsilon}
$$

Set the value of $1+\delta$ to be

$$
\begin{aligned}
1+\delta & =\frac{1}{2 \mu}=\frac{q}{2 \varepsilon N} \cdot\left(\sum_{j=1}^{m} R_{j}(\boldsymbol{A}, u, \ell)\right) \cdot \frac{1}{\lambda_{\max }\left(\frac{1}{u \boldsymbol{I}-\boldsymbol{A}}\right)} \\
& =\frac{q}{2 \varepsilon N} \cdot\left(\sum_{j=1}^{m} R_{j}(\boldsymbol{A}, u, \ell)\right) \cdot \lambda_{\min }(u \boldsymbol{I}-\boldsymbol{A}) \geq \frac{q}{4 \varepsilon} \cdot n^{2 / q},
\end{aligned}
$$

where the last inequality follows from the condition on $N$. Hence, with probability at least

$$
1-n \cdot\left(\frac{e^{\delta}}{(1+\delta)^{1+\delta}}\right)^{\mu \cdot q / \varepsilon} \geq 1-n \cdot\left(\frac{e}{1+\delta}\right)^{(1+\delta) \cdot \mu \cdot q / \varepsilon} \geq 1-n\left(\frac{e}{1+\delta}\right)^{\frac{q}{2 \varepsilon}} \geq 1-\frac{\varepsilon^{2}}{100 q n}
$$

we have that

$$
\lambda_{\max }\left(\sum_{i=1}^{N} z_{i} z_{i}^{T}\right) \leq(1+\delta) \cdot \mu=\frac{1}{2}
$$

which implies that $\mathbf{0} \preceq \sum_{i=1}^{N} \boldsymbol{z}_{i} \boldsymbol{z}_{i}^{T} \preceq \frac{1}{2} \cdot I$ and $\mathbf{0} \preceq W \preceq \frac{1}{2} \cdot(u I-A)$.
Now we analyze the change of the potential function after each iteration, and show that the expected value of the potential function decreases over time. By Lemma 4.3.1, with probability at least $1-\frac{\varepsilon^{2}}{100 q n}$, it holds that

$$
\mathbf{0} \preceq \boldsymbol{W} \preceq \frac{1}{2}(u \boldsymbol{I}-\boldsymbol{A}) .
$$

We define

$$
\widetilde{\mathbf{E}}[f(\boldsymbol{W})] \triangleq \sum_{\boldsymbol{W} \sim \mathcal{D}(\boldsymbol{A}, u, \ell)} \operatorname{Pr}\left[\boldsymbol{W} \text { is chosen and } \boldsymbol{W} \preceq \frac{1}{2}(u \boldsymbol{I}-\boldsymbol{A})\right] \cdot f(\boldsymbol{W}) .
$$

Lemma 4.3.2 below shows how the potential function changes after each iteration, and plays a key role in our analysis. This lemma was first proved in [31] for the case of $q=1$, and similar lemma was proved in [9].

Lemma 4.3.2 ([9]). Let $q \geq 10$ and $\varepsilon \leq 1 / 10$. Suppose that $\boldsymbol{w}^{T}(u \boldsymbol{I}-\boldsymbol{A})^{-1} \boldsymbol{w} \leq \varepsilon / q$ and $\boldsymbol{w}^{T}(\boldsymbol{A}-$ $\ell \boldsymbol{I})^{-1} \boldsymbol{w} \leq \varepsilon / q$. Then, it holds that

$$
\operatorname{Tr}\left(\boldsymbol{A}+\boldsymbol{w} \boldsymbol{w}^{T}-\ell \boldsymbol{I}\right)^{-q} \leq \operatorname{Tr}(\boldsymbol{A}-\ell \boldsymbol{I})^{-q}-q(1-\varepsilon) \boldsymbol{w}^{T}(\boldsymbol{A}-\ell \boldsymbol{I})^{-(q+1)} \boldsymbol{w}
$$

and

$$
\operatorname{Tr}\left(u \boldsymbol{I}-\boldsymbol{A}-\boldsymbol{w} \boldsymbol{w}^{T}\right)^{-q} \leq \operatorname{Tr}(u \boldsymbol{I}-\boldsymbol{A})^{-q}+q(1+\varepsilon) \boldsymbol{w}^{T}(u \boldsymbol{I}-\boldsymbol{A})^{-(q+1)} \boldsymbol{w}
$$

Proof. Let $\boldsymbol{Y}=\boldsymbol{A}-\ell \boldsymbol{I}$. By the Sherman-Morrison Formula (Lemma 2.3.7), it holds that

$$
\begin{equation*}
\operatorname{Tr}\left(\boldsymbol{Y}+\boldsymbol{w} \boldsymbol{w}^{T}\right)^{-q}=\operatorname{Tr}\left(\boldsymbol{Y}^{-1}-\frac{\boldsymbol{Y}^{-1} w w^{T} \boldsymbol{Y}^{-1}}{1+w^{T} \boldsymbol{Y}^{-1} w}\right)^{q} \tag{4.6}
\end{equation*}
$$

By the assumption of $w^{T} \boldsymbol{Y}^{-1} w \leq \varepsilon / q$, we have that

$$
\begin{align*}
\operatorname{Tr}\left(\boldsymbol{Y}+w w^{T}\right)^{-q} & \leq \operatorname{Tr}\left(\boldsymbol{Y}^{-1}-\frac{\boldsymbol{Y}^{-1} w w^{T} \boldsymbol{Y}^{-1}}{1+\varepsilon / q}\right)^{q}  \tag{4.7}\\
& =\operatorname{Tr}\left(\boldsymbol{Y}^{-1 / 2}\left(\boldsymbol{I}-\frac{\boldsymbol{Y}^{-1 / 2} w w^{T} \boldsymbol{Y}^{-1 / 2}}{1+\varepsilon / q}\right) \boldsymbol{Y}^{-1 / 2}\right)^{q} \\
& \leq \operatorname{Tr}\left(\boldsymbol{Y}^{-q / 2}\left(\boldsymbol{I}-\frac{\boldsymbol{Y}^{-1 / 2} w w^{T} \boldsymbol{Y}^{-1 / 2}}{1+\varepsilon / q}\right)^{q} \boldsymbol{Y}^{-q / 2}\right)  \tag{4.8}\\
& =\operatorname{Tr}\left(\boldsymbol{Y}^{-q}\left(\boldsymbol{I}-\frac{\boldsymbol{Y}^{-1 / 2} w w^{T} \boldsymbol{Y}^{-1 / 2}}{1+\varepsilon / q}\right)^{q}\right) \tag{4.9}
\end{align*}
$$

where (4.7) uses the fact that $\boldsymbol{A} \preceq \boldsymbol{B}$ implies that $\operatorname{Tr}\left(\boldsymbol{A}^{q}\right) \leq \operatorname{Tr}\left(\boldsymbol{B}^{q}\right)$, (4.8) follows from the LiebThirring inequality (Lemma 2.3.11), and (4.8) uses the fact that the trace is invariant under cyclic
permutations.
Let $\boldsymbol{D}=\frac{\boldsymbol{Y}^{-1 / 2} w w^{T} \boldsymbol{Y}^{-1 / 2}}{1+\varepsilon / q}$. Note that $0 \preceq D \preceq \frac{\varepsilon}{q} \cdot \boldsymbol{I}$, and

$$
\begin{aligned}
(I-D)^{q} & \preceq I-q D+\frac{q(q-1)}{2} D^{2} \\
& \preceq I-\left(q-\frac{\varepsilon(q-1)}{2}\right) D .
\end{aligned}
$$

Therefore, we have that

$$
\begin{aligned}
\left(I-\frac{\boldsymbol{Y}^{-1 / 2} w w^{T} \boldsymbol{Y}^{-1 / 2}}{1+\varepsilon / q}\right)^{q} & \preceq I-\left(q-\frac{\varepsilon(q-1)}{2}\right) \frac{\boldsymbol{Y}^{-1 / 2} w w^{T} \boldsymbol{Y}^{-1 / 2}}{1+\varepsilon / q} \\
& \preceq I-\left(q-\frac{\varepsilon(q-1)}{2}\right)\left(1-\frac{\varepsilon}{q}\right) \boldsymbol{Y}^{-1 / 2} w w^{T} \boldsymbol{Y}^{-1 / 2} \\
& \preceq I-q\left(1-\frac{\varepsilon(q+1)}{2 q}\right) \boldsymbol{Y}^{-1 / 2} w w^{T} \boldsymbol{Y}^{-1 / 2} \\
& \preceq I-q(1-\varepsilon) \boldsymbol{Y}^{-1 / 2} w w^{T} \boldsymbol{Y}^{-1 / 2} .
\end{aligned}
$$

This implies that

$$
\begin{aligned}
\operatorname{Tr}\left(Y+w w^{T}\right)^{-q} & \leq \operatorname{Tr}\left(Y^{-q}\left(I-q(1-\varepsilon) Y^{-1 / 2} w w^{T} Y^{-1 / 2}\right)\right) \\
& \leq \operatorname{Tr}\left(Y^{-q}\right)-q(1-\varepsilon) w^{T} Y^{-(q+1)} w,
\end{aligned}
$$

which proves the first statement.

Now for the second inequality. Let $Z=u I-A$. By the Sherman-Morrison Formula (Lemma 2.3.7), it holds that

$$
\operatorname{Tr}\left(\boldsymbol{Z}-w w^{T}\right)^{-q}=\operatorname{Tr}\left(\boldsymbol{Z}^{-1}+\frac{\boldsymbol{Z}^{-1} w w^{T} \boldsymbol{Z}^{-1}}{1-w^{T} \boldsymbol{Z}^{-1} w}\right)^{q}
$$

By the assumption of $w^{T} \boldsymbol{Z}^{-1} w \leq \varepsilon / q$, it holds that

$$
\begin{align*}
\operatorname{Tr}\left(\boldsymbol{Z}-w w^{T}\right)^{-q} & \leq \operatorname{Tr}\left(\boldsymbol{Z}^{-1}+\frac{\boldsymbol{Z}^{-1} w w^{T} \boldsymbol{Z}^{-1}}{1-\varepsilon / q}\right)^{q}  \tag{4.10}\\
& =\operatorname{Tr}\left(\boldsymbol{Z}^{-1 / 2}\left(\boldsymbol{I}+\frac{\boldsymbol{Z}^{-1 / 2} w w^{T} \boldsymbol{Z}^{-1 / 2}}{1-\varepsilon / q}\right) \boldsymbol{Z}^{-1 / 2}\right)^{q} \\
& \leq \operatorname{Tr}\left(\boldsymbol{Z}^{-q / 2}\left(\boldsymbol{I}+\frac{\boldsymbol{Z}^{-1 / 2} w w^{T} \boldsymbol{Z}^{-1 / 2}}{1-\varepsilon / q}\right)^{q} \boldsymbol{Z}^{-q / 2}\right)  \tag{4.11}\\
& =\operatorname{Tr}\left(\boldsymbol{Z}^{-q}\left(\boldsymbol{I}+\frac{\boldsymbol{Z}^{-1 / 2} w w^{T} \boldsymbol{Z}^{-1 / 2}}{1-\varepsilon / q}\right)^{q}\right) \tag{4.12}
\end{align*}
$$

where (4.10) uses the fact that $\boldsymbol{A} \preceq \boldsymbol{B}$ implies $\operatorname{Tr}\left(\boldsymbol{A}^{q}\right) \leq \operatorname{Tr}\left(\boldsymbol{B}^{q}\right)$, (4.11) follows from the LiebThirring inequality (Lemma 2.3.11), and (4.12) uses the fact that the trace is invariant under cyclic permutations.

Let $\boldsymbol{E}=\boldsymbol{Z}^{-1 / 2} w w^{T} \boldsymbol{Z}^{-1 / 2}$. Combing $E \preceq \frac{\varepsilon}{q} \cdot \boldsymbol{I}$ with the assumption that $q \geq 10$ and $\varepsilon \leq 1 / 10$, we
have that

$$
\begin{aligned}
\left(I+\frac{E}{1-\varepsilon / q}\right)^{q} & \preceq I+\frac{q E}{1-\varepsilon / q}+\frac{q(q-1)}{2}\left(1+\frac{\varepsilon / q}{1-\varepsilon / q}\right)^{q-2}\left(\frac{E}{1-\varepsilon / q}\right)^{2} \\
& \preceq I+q\left(1+1.1 \frac{\varepsilon}{q}\right) E+1.4 \frac{q(q-1)}{2} E^{2} \\
& \preceq I+q(1+0.3 \varepsilon) E+0.7 \varepsilon q E \\
& \preceq I+q(1+\varepsilon) E .
\end{aligned}
$$

Therefore, we have that

$$
\operatorname{Tr}\left(Z-w w^{T}\right)^{-q} \leq \operatorname{Tr}\left(Z^{-q}\right)+q(1+\varepsilon) w^{T} Z^{-(q+1)} w,
$$

which proves the second statement.

The following lemma shows that, with our choice of updated matrices and barrier values, the value of the potential function will never increase.

Lemma 4.3.3. It holds for any iteration $j$ that $\Phi_{u_{j+1}, \ell_{j+1}}\left(A_{j+1}\right) \leq \Phi_{u_{j}, \ell_{j}}\left(A_{j}\right)$.
Proof. Let $w_{1} w_{1}^{T}, \cdots, w_{N_{j}} w_{N_{j}}^{T}$ be the matrices picked in iteration $j$, and define for any $0 \leq i \leq N_{j}$ that $B_{i}=A_{j}+\sum_{t=1}^{i} w_{t} w_{t}^{T}$. We study the change of the potential function after adding a rank-1 matrix within each iteration. For this reason, we use

$$
\bar{\Delta}_{u}=\frac{\Delta_{u, j}}{N_{j}}=(1+2 \varepsilon) \cdot \frac{\varepsilon}{q \cdot \sum_{t=1}^{m} R_{t}\left(A_{j}, u_{j}, \ell_{j}\right)},
$$

and

$$
\bar{\Delta}_{\ell}=\frac{\Delta_{\ell, j}}{N_{j}}=(1-2 \varepsilon) \cdot \frac{\varepsilon}{q \cdot \sum_{t=1}^{m} R_{t}\left(A_{j}, u_{j}, \ell_{j}\right)}
$$

to express the average change of the barrier values $\Delta_{u, j}$ and $\Delta_{\ell, j}$. We further define for $0 \leq j \leq N_{j}$ that $\hat{u}_{i}=u_{j}+i \cdot \bar{\Delta}_{u}, \quad \hat{\ell}_{i}=\ell_{j}+i \cdot \bar{\Delta}_{\ell}$.

Assuming $\mathbf{0} \preceq W_{j} \preceq \frac{1}{2}\left(u_{j} \boldsymbol{I}-\boldsymbol{A}_{\boldsymbol{j}}\right)$, we claim that

$$
\begin{equation*}
w_{i} w_{i}^{T} \preceq \frac{2 \varepsilon}{q} \cdot\left(\hat{u}_{i} I-B_{i-1}\right) \text { and } w_{i} w_{i}^{T} \preceq \frac{2 \varepsilon}{q} \cdot\left(B_{i-1}-\hat{\ell}_{i} I\right), \tag{4.13}
\end{equation*}
$$

for any $1 \leq i \leq N_{j}$. Based on this, we apply Lemma 4.3.2 and get that

$$
\begin{gather*}
\Phi_{\hat{u}_{i}, \hat{\ell}_{i}}\left(B_{i-1}+w_{i} w_{i}^{T}\right) \leq \\
\Phi_{\hat{u}_{i}, \hat{\ell}_{i}}\left(B_{i-1}\right)+q(1+2 \varepsilon) \operatorname{Tr}\left(\left(\hat{u}_{i} I-B_{i-1}\right)^{-(q+1)} \mathbf{E}\left[w_{i} w_{i}^{T}\right]\right) \\
\quad-q(1-2 \varepsilon) \operatorname{Tr}\left(\left(B_{i-1}-\hat{\ell}_{i} I\right)^{-(q+1)} \mathbf{E}\left[w_{i} w_{i}^{T}\right]\right) \\
=  \tag{4.14}\\
\Phi_{\hat{u}_{i}, \hat{\ell}_{i}}\left(B_{i-1}\right)+q \cdot \bar{\Delta}_{u} \cdot \operatorname{Tr}\left(\left(\hat{u}_{i} I-B_{i-1}\right)^{-(q+1)}\right) \\
\quad-q \cdot \bar{\Delta}_{\ell} \cdot \operatorname{Tr}\left(\left(B_{i-1}-\hat{\ell}_{i} I\right)^{-(q+1)}\right) .
\end{gather*}
$$

We define a function $f_{i}: \mathbb{R} \rightarrow \mathbb{R}$ by

$$
f_{i}(x)=\operatorname{Tr}\left(\left(\hat{u}_{i-1}+x \cdot \bar{\Delta}_{u}\right) I-B_{i-1}\right)^{-q}+\operatorname{Tr}\left(B_{i-1}-\left(\hat{\ell}_{i-1}+x \cdot \bar{\Delta}_{\ell}\right) I\right)^{-q}
$$

Notice that

$$
\begin{aligned}
\frac{\mathrm{d} f_{i}(x)}{\mathrm{d} x}=-q \cdot & \bar{\Delta}_{u} \cdot \operatorname{Tr}\left(\left(\hat{u}_{i-1}+x \cdot \bar{\Delta}_{u}\right) I-B_{i-1}\right)^{-(q+1)} \\
& +q \cdot \bar{\Delta}_{\ell} \cdot \operatorname{Tr}\left(B_{i-1}-\left(\hat{\ell}_{i-1}+x \cdot \bar{\Delta}_{\ell}\right) I\right)^{-(q+1)} .
\end{aligned}
$$

Since $f$ is convex, we have that

$$
\begin{equation*}
\left.\frac{\mathrm{d} f_{i}(x)}{\mathrm{d} x}\right|_{x=1} \geq f_{i}(1)-f_{i}(0)=\Phi_{\hat{u}_{i}, \hat{\ell}_{i}}\left(B_{i-1}\right)-\Phi_{\hat{u}_{i-1}, \hat{\ell}_{i-1}}\left(B_{i-1}\right) . \tag{4.15}
\end{equation*}
$$

Putting (4.14) and (4.15) together, we have that

$$
\Phi_{\hat{u}_{i}, \hat{\ell}_{i}}\left(B_{i}\right) \leq \Phi_{\hat{u}_{i}, \hat{\ell}_{i}}\left(B_{i-1}\right)-\left.\frac{\mathrm{d} f_{i}(x)}{\mathrm{d} x}\right|_{x=1} \leq \Phi_{\hat{u}_{i-1}, \hat{\ell}_{i-1}}\left(B_{i-1}\right) .
$$

Repeating this argument we have that

$$
\Phi_{u_{j+1}, \ell_{j+1}}\left(A_{j+1}\right)=\Phi_{\hat{u}_{N_{j}}, \hat{\ell}_{N_{j}}}\left(B_{N_{j}}\right) \leq \Phi_{\hat{u}_{0}, \hat{\ell}_{0}}\left(B_{0}\right)=\Phi_{u_{j}, \ell_{j}}\left(A_{j}\right),
$$

which proves the statement.
So, it suffices to prove (4.13). Since $v v^{T} \preceq\left(v^{T} B^{-1} v\right) B$ for any vector $v$ and PSD matrix $B$, we have that

$$
\frac{v_{i} v_{i}^{T}}{R_{i}\left(A_{j}, u_{j}, \ell_{j}\right)} \preceq \frac{v_{i} v_{i}^{T}}{v_{i}^{T}\left(u_{j} I-A_{j}\right)^{-1} v_{i}} \preceq u_{j} I-A_{j} .
$$

By the assumption of $W_{j} \preceq \frac{1}{2}\left(u_{j} I-A_{j}\right)$, it holds that

$$
w_{i} w_{i}^{T}=\frac{\varepsilon}{q R_{i}\left(A_{j}, u_{j}, \ell_{j}\right)} v_{i} v_{i}^{T} \preceq \frac{\varepsilon}{q}\left(u_{j} I-A_{j}\right) \preceq \frac{2 \varepsilon}{q}\left(\hat{u}_{i} I-B_{i-1}\right) .
$$

This proves the first statement of the claim.
For the second statement, notice that

$$
\ell_{j+1}-\ell_{j}=\Delta_{\ell, j} \leq \frac{\varepsilon N_{j}}{q \sum_{t=1}^{m} R_{t}\left(A_{j}, u_{j}, \ell_{j}\right)} \leq \frac{1}{2} \cdot \lambda_{\min }\left(A_{j}-\ell_{j} I\right)
$$

and hence

$$
w_{i} w_{i}^{T} \preceq \frac{\varepsilon}{q}\left(A_{j}-\ell_{j} I\right) \preceq \frac{2 \varepsilon}{q}\left(A_{j}-\hat{\ell}_{i} I\right) \preceq \frac{2 \varepsilon}{q}\left(B_{i-1}-\hat{\ell}_{i} I\right) .
$$

### 4.3.2 Analysis of the approximation guarantee

In this subsection we will prove that the algorithm produces a linear-sized $O(\varepsilon)$-spectral sparsifier. We assume that the algorithm finishes after $k$ iterations, and will prove that the condition number of $A_{k}$ is small, which follows from our setting of parameters.
Lemma 4.3.4. The output matrix $A_{k}$ has condition number at most $1+O(\varepsilon)$.
Proof. Since the condition number of $A_{k}$ is at most $\frac{u_{k}}{\ell_{k}}=\left(1-\frac{u_{k}-\ell_{k}}{u_{k}}\right)^{-1}$, it suffices to prove that $\left(u_{k}-\ell_{k}\right) / u_{k}=O(\varepsilon)$.

Since the increase rate of $\Delta_{u, j}-\Delta_{\ell, j}$ with respect to $\Delta_{u, j}$ for any iteration $j$ is

$$
\frac{\Delta_{u, j}-\Delta_{\ell, j}}{\Delta_{u, j}}=\frac{(1+2 \varepsilon)-(1-2 \varepsilon)}{1+2 \varepsilon}=\frac{4 \varepsilon}{1+2 \varepsilon} \leq 4 \varepsilon,
$$

we have that

$$
\begin{aligned}
\frac{u_{k}-\ell_{k}}{u_{k}} & =\frac{2 \cdot(2 n)^{1 / q}+\sum_{j=0}^{k-1}\left(\Delta_{u, j}-\Delta_{\ell, j}\right)}{(2 n)^{1 / q}+\sum_{j=0}^{k-1} \Delta_{u, j}} \\
& \leq \frac{2 \cdot(2 n)^{1 / q}+\sum_{j=0}^{k-1}\left(\Delta_{u, j}-\Delta_{\ell, j}\right)}{(2 n)^{1 / q}+(4 \varepsilon)^{-1} \sum_{j=0}^{k-1}\left(\Delta_{u, j}-\Delta_{\ell, j}\right)}
\end{aligned}
$$

By the ending condition of the algorithm, it holds that $u_{k}-\ell_{k} \geq 4 \cdot(2 n)^{1 / q}$, i.e.

$$
\sum_{j=0}^{k-1}\left(\Delta_{u, j}-\Delta_{\ell, j}\right) \geq 2 \cdot(2 n)^{1 / q}
$$

Hence, it holds that

$$
\frac{u_{k}-\ell_{k}}{u_{k}} \leq \frac{2 \cdot(2 n)^{1 / q}+2 \cdot(2 n)^{1 / q}}{(2 n)^{1 / q}+(4 \varepsilon)^{-1} 2 \cdot(2 n)^{1 / q}} \leq 8 \varepsilon,
$$

which finishes the proof.
Now we prove that the algorithm finishes in $O\left(\frac{q n^{3 / q}}{\varepsilon^{2}}\right)$ iterations, and picks $O\left(q n / \varepsilon^{2}\right)$ vectors in total.
Lemma 4.3.5. The following statements hold:

1. With probability at least $4 / 5$, the algorithm finishes in $\frac{10 q n^{3 / q}}{\varepsilon^{2}}$ iterations.
2. With probability at least $4 / 5$, the algorithm chooses at most $10 q n / \varepsilon^{2}$ vectors.

Proof. Notice that after iteration $j$ the barrier gap $u_{j}-\ell_{j}$ is increased by

$$
\begin{aligned}
\Delta_{u, j}-\Delta_{\ell, j} & =\frac{4 \varepsilon^{2}}{q} \frac{N_{j}}{\sum_{i=1}^{m} R_{i}\left(A_{j}, u_{j}, \ell_{j}\right)} \\
& =\frac{4 \varepsilon^{2}}{q} \frac{1}{n^{2 / q}} \cdot \min \left\{\lambda_{\min }\left(u_{j} I-A_{j}\right), \lambda_{\min }\left(A_{j}-\ell_{j} I\right)\right\} \\
& \geq \frac{4 \varepsilon^{2}}{q} \frac{1}{n^{2 / q}} \cdot\left(\Phi_{u_{j}, \ell_{j}}\left(A_{j}\right)\right)^{-1 / q} .
\end{aligned}
$$

Since the algorithm finishes within $k$ iterations if $\sum_{j=0}^{k-1}\left(\Delta_{u, j}-\Delta_{\ell, j}\right) \geq 2 \cdot(2 n)^{1 / q}$, it holds that

$$
\begin{aligned}
\operatorname{Pr}[\text { algorithm finishes within } k \text { iterations }] & \geq \operatorname{Pr}\left[\sum_{j=0}^{k-1}\left(\Delta_{u, j}-\Delta_{\ell, j}\right) \geq 2 \cdot(2 n)^{1 / q}\right] \\
& \geq \operatorname{Pr}\left[\sum_{j=0}^{k-1} \frac{4 \varepsilon^{2}}{q n^{2 / q}} \cdot\left(\Phi_{u_{j}, \ell_{j}}\left(A_{j}\right)\right)^{-1 / q} \geq 2 \cdot(2 n)^{1 / q}\right] \\
& =\operatorname{Pr}\left[\sum_{j=0}^{k-1}\left(\Phi_{u_{j}, \ell_{j}}\left(A_{j}\right)\right)^{-1 / q} \geq \frac{q}{2 \varepsilon^{2}} \cdot\left(2 n^{3}\right)^{1 / q}\right] \\
& \geq \operatorname{Pr}\left[\sum_{j=0}^{k-1}\left(\Phi_{u_{j}, \ell_{j}}\left(A_{j}\right)\right)^{1 / q} \leq 2 \cdot \frac{k^{2} \varepsilon^{2}}{q} \cdot\left(\frac{1}{2 n^{3}}\right)^{1 / q}\right],
\end{aligned}
$$

where the last inequality follows from the fact that

$$
\left(\sum_{j=0}^{k-1}\left(\Phi_{u_{j}, \ell_{j}}\left(A_{j}\right)\right)^{-1 / q}\right) \cdot\left(\sum_{j=0}^{k-1}\left(\Phi_{u_{j}, \ell_{j}}\left(A_{j}\right)\right)^{1 / q}\right) \geq k^{2} .
$$

By Lemma 4.3.1, every picked matrix $W_{j}$ in iteration $j$ satisfies $0 \preceq W_{j} \preceq \frac{1}{2} \cdot\left(u_{j} I-A\right)$ with probability at least $1-\frac{\varepsilon^{2}}{100 q n}$, and with probability $9 / 10$ all matrices picked in $k=10 q n / \varepsilon^{2}$ iterations satisfy the condition above. Also, by Lemma 4.3 .3 we have that

$$
\begin{equation*}
\sum_{j=0}^{k-1}\left(\Phi_{u_{j}, \ell_{j}}\left(A_{j}\right)\right)^{1 / q}=\sum_{j=0}^{k-1}\left(\Phi_{u_{j}, \ell_{j}}\left(A_{j}\right)\right)^{1 / q} \leq \sum_{j=0}^{k-1}\left(\Phi_{u_{j}, \ell_{j}}\left(A_{j}\right)\right)^{1 / q} \leq k, \tag{4.16}
\end{equation*}
$$

since the initial value of the potential function is at most 1 . Therefore, it holds that

$$
\begin{aligned}
& \operatorname{Pr} \text { [algorithm finishes in more than } k \text { iterations] } \\
& \leq \operatorname{Pr}\left[\sum_{j=0}^{k-1}\left(\Phi_{u_{j}, \ell_{j}}\left(A_{j}\right)\right)^{1 / q} \geq 2 \cdot \frac{k^{2} \varepsilon^{2}}{q} \cdot\left(\frac{1}{2 n^{3}}\right)^{1 / q}\right] \\
& \leq \operatorname{Pr}\left[\sum_{j=0}^{k-1}\left(\Phi_{u_{j}, \ell_{j}}\left(A_{j}\right)\right)^{1 / q} \geq 2 \cdot \frac{k^{2} \varepsilon^{2}}{q} \cdot\left(\frac{1}{2 n^{3}}\right)^{1 / q} \text { and } \forall j: W_{j} \preceq \frac{1}{2}\left(u_{j} I-A_{j}\right)\right] \\
& \quad+\operatorname{Pr}\left[\exists j: W_{j} \npreceq \frac{1}{2}\left(u_{j} I-A_{j}\right)\right] \leq \frac{q}{2 \cdot k \varepsilon^{2}} \cdot\left(2 n^{3}\right)^{1 / q}+1 / 10 \leq 1 / 5,
\end{aligned}
$$

where the second last inequity follows from Markov's inequality and (4.16), and the last inequality follows by our choice of $k$. This proves the first statement.

Now for the second statement. Notice that for every vector chosen in iteration $j$, the barrier gap $\Delta_{u, j}-\Delta_{\ell, j}$ is increased on average by

$$
\frac{\Delta_{u, j}-\Delta_{\ell, j}}{N_{j}}=\frac{4 \varepsilon^{2}}{q \sum_{i=1}^{m} R_{i}\left(A_{j}, u_{j}, \ell_{j}\right)} .
$$

To bound $R_{i}\left(A_{j}, u_{j}, \ell_{j}\right)$, let the eigenvalues of matrix $A_{j}$ be $\lambda_{1}, \cdots, \lambda_{n}$. Then, it holds that

$$
\begin{aligned}
\sum_{i=1}^{m} R_{i}\left(\boldsymbol{A}_{\boldsymbol{j}}, u_{j}, \ell_{j}\right) & =\sum_{i=1}^{m} \vec{v}_{i}^{T}\left(u_{j} \boldsymbol{I}-A_{j}\right)^{-1} \boldsymbol{v}_{i}+\sum_{i=1}^{m} \vec{v}_{i}^{T}\left(\boldsymbol{A}_{\boldsymbol{j}}-\ell_{j} \boldsymbol{I}\right)^{-1} \vec{v}_{i} \\
& =\sum_{i=1}^{n} \frac{1}{u_{j}-\lambda_{i}}+\sum_{i=1}^{n} \frac{1}{\lambda_{i}-\ell_{j}} \\
& \leq\left(\sum_{i=1}^{n}\left(u_{j}-\lambda_{i}\right)^{-q}+\sum_{i=1}^{n}\left(\lambda_{i}-\ell_{j}\right)^{-q}\right)^{1 / q}(2 n)^{1-1 / q} \\
& =\left(\Phi_{u_{j}, \ell_{j}}\left(\boldsymbol{A}_{\boldsymbol{j}}\right)\right)^{1 / q} \cdot(2 n)^{1-1 / q} .
\end{aligned}
$$

Therefore, we have that

$$
\begin{equation*}
\frac{\Delta_{u, j}-\Delta_{\ell, j}}{N_{j}} \geq \frac{4 \varepsilon^{2}}{q} \cdot \frac{1}{(2 n)^{1-1 / q} \cdot\left(\Phi_{u_{j}, \ell_{j}}\left(A_{j}\right)\right)^{1 / q}} . \tag{4.17}
\end{equation*}
$$

Let $v_{1}, \cdots, v_{z}$ be the vectors sampled by the algorithm, and $v_{j}$ is picked in iteration $\tau_{j}$, where
$1 \leq j \leq z$. We first assume that the algorithm could check the ending condition after adding every single vector. In such case, it holds that
$\operatorname{Pr}$ [algorithm finishes after choosing $z$ vectors]

$$
\begin{aligned}
& \geq \operatorname{Pr}\left[\sum_{j=1}^{z} \frac{4 \varepsilon^{2}}{q} \cdot \frac{1}{(2 n)^{1-1 / q} \cdot\left(\Phi_{u_{\tau_{j}}, \ell_{\tau_{j}}}\left(A_{\tau_{j}}\right)\right)^{1 / q}} \geq 2 \cdot(2 n)^{1 / q}\right] \\
& =\operatorname{Pr}\left[\sum_{j=1}^{z}\left(\Phi_{u_{\tau_{j}}, \ell_{\tau_{j}}}\left(A_{\tau_{j}}\right)\right)^{-1 / q} \geq q n / \varepsilon^{2}\right] .
\end{aligned}
$$

Following the same proof as the first part and noticing that in the final iteration the algorithm chooses at most $O(n)$ extra vectors, we obtain the second statement.

## - 4.3.3 Proof of the main results

Now we analyze the runtime of the algorithm, and prove the main results. We first analyze the algorithm for sparsifying the sum of rank-1 PSD matrices, and prove Theorem 4.1.1.

Proof of Theorem 4.1.1. By Theorem 3.7.3, we can first find a $\varepsilon$-approximation with $O\left(n \log n / \varepsilon^{2}\right)$ vectors in time $\tilde{O}\left(Z+n^{\omega} \varepsilon^{-2}\right)$. Note that if $q<\log n$, we are done.

Otherwise, by Lemma 4.3.5, with probability at least $4 / 5$ the algorithm chooses at most $10 \mathrm{qn} / \varepsilon^{2}$ vectors, and by Lemma 4.3.4 the condition number of $A_{k}$ is at most $1+O(\varepsilon)$, implying that the matrix $A_{k}$ is a $(1+O(\varepsilon))$-approximation of $I$. These two results together prove that $A_{k}$ is a linear-sized spectral sparsifier.

For the runtime, Lemma 4.3 .5 proves that the algorithm finishes in $\frac{10 q n^{3 / q}}{\varepsilon^{2}}$ iterations, and it is easy to see that all the required quantities in each iteration can be approximately computed in $\tilde{O}\left(m \cdot n^{\omega-1}\right)$ time using fast matrix multiplication. Therefore, the total runtime of the algorithm is $\tilde{O}\left(Z+n^{\omega} \varepsilon^{-2}+\frac{q \cdot m}{\varepsilon^{2}} \cdot n^{\omega-1+3 / q}\right)$. Using $q<\log n$ and $m=O\left(n \log n / \varepsilon^{2}\right)$, we have

$$
\tilde{O}\left(Z+n^{\omega+3 / q} \varepsilon^{-4}\right) .
$$

Next we apply our algorithm in the graph setting, and prove Theorem 4.1.2.
Proof of Theorem 4.1.2. For any edge $e=\{u, v\}$ in graph $G=(V, E)$, we define $b_{e} \in \mathbb{R}^{n}$, where $b_{e}(w)=1$ if $w=u, b_{e}(w)=-1$ if $w=v$, and $b_{e}(w)=0$ otherwise. Then the Laplacian matrix of $G$ can be written as $L=\sum_{e \in E[G]} b_{e} b_{e}^{T}$. By setting $v_{e}=L^{-1 / 2} b_{e}$ for $e \in E[G]$, it is easy to see that constructing a spectral sparsifier of $G$ is equivalent to sparsifying the matrix $\sum_{e \in E[G]} v_{e} v_{e}^{T}$. We apply the same analysis as in the proof of Theorem 4.1.1, and hence the output is a linear-sized spectral sparsifier of graph $G$. So it suffices to analyze the runtime of the algorithm.

By Lemma 4.3.1 and the Union Bound, with probability at least $9 / 10$ all the matrices picked in $k=\frac{10 q n^{3 / q}}{\varepsilon^{2}}$ iterations satisfy $W_{j} \preceq \frac{1}{2}\left(u_{j} I-A_{j}\right)$. Conditioning on the event, with constant probability $\mathbf{E}\left[\Phi_{u_{j}, \ell_{j}}\left(A_{j}\right)\right] \leq 2$ for all iterations $j$, and by Markov's inequality with high probability it holds that $\Phi_{u_{j}, \ell_{j}}\left(A_{j}\right)=O\left(q n / \varepsilon^{2}\right)$ for all iterations $j$.

On the other hand, notice that it holds for any $1 \leq j \leq n$ that

$$
\left(u-\lambda_{j}\right)^{-q} \leq \sum_{i=1}^{n}\left(u-\lambda_{i}\right)^{-q}<\Phi_{u, \ell}(\boldsymbol{A}),
$$

where $\lambda_{1}, \ldots, \lambda_{n}$ are the eigenvalues of $A$. This implies that $\lambda_{j}<u-\left(\Phi_{u, \ell}(\boldsymbol{A})\right)^{-1 / q}$. Similarly, it holds that $\lambda_{j}>\ell+\left(\Phi_{u, \ell}(\boldsymbol{A})\right)^{-1 / q}$ for any $1 \leq j \leq n$. Therefore, we have that

$$
\left(\ell_{j}+O\left(\left(\frac{\varepsilon^{2}}{q n}\right)^{1 / q}\right)\right) \boldsymbol{I} \prec \boldsymbol{A}_{\boldsymbol{j}} \prec\left(u_{j}-O\left(\left(\frac{\varepsilon^{2}}{q n}\right)^{1 / q}\right)\right) \boldsymbol{I} .
$$

Since both of $u_{j}$ and $\ell_{j}$ are of the order $O\left(n^{1 / q}\right)$, we set $\eta=O\left((\varepsilon / n)^{2 / q}\right)$ and obtain that

$$
\left(\ell_{j}+\left|\ell_{j}\right| \eta\right) I \prec A_{j} \prec(1-\eta) u_{j} I .
$$

Hence, we apply Lemma 4.4.3 and Lemma 4.4.4 to compute all required quantities in each iteration up to constant approximation in time

$$
\tilde{O}\left(\frac{m}{\varepsilon^{2} \cdot \eta}\right)=\tilde{O}\left(\frac{m \cdot n^{2 / q}}{\varepsilon^{2+2 / q}}\right) .
$$

Since by Lemma 4.3.5 the algorithm finishes in $\frac{10 q n^{3 / q}}{\varepsilon^{2}}$ iterations with probability at least $4 / 5$, the total runtime of the algorithm is

$$
\tilde{O}\left(\frac{q \cdot m \cdot n^{5 / q}}{\varepsilon^{4+4 / q}}\right) .
$$

Again, we can assume $m=\tilde{O}\left(n \log n / \varepsilon^{2}\right)$ and $q<\log n$ by existing algorithm by using existing algorithm for graph sparsification. Hence, we have the promised runtime.

## ■ 4.4 Appendix

In this section we study fast approximation of the required quantities for constructing graph sparsifiers, and show that the number of sampled vectors $N_{j}$ and the relative leverage scores $\left\{R_{i}\left(A_{j}, u_{j}, \ell_{j}\right)\right\}_{i=1}^{m}$ used in each iteration can be approximately computed in almost-linear time. For simplicity we drop the subscript $j$ expressing the iterations in this subsection. We will assume that the following assumption holds on $A$, and show that the matrix under consideration always satisfies this condition.
Assumption 4.4.1. Let $L$ and $\tilde{L}$ be the Laplacian matrices of graph $G$ and its subgraph after reweighting. Let $A=L^{-1 / 2} \tilde{L} L^{-1 / 2}$, and assume that

$$
(\ell+|\ell| \eta) \cdot I \prec A \prec(1-\eta) u \cdot I
$$

holds for some $0<\eta<1$.
Lemma 4.4.2. Under Assumption 4.4.1, the following statements hold:

1. We can construct a matrix $S_{u}$ such that

$$
S_{u} \approx_{\varepsilon / 10}(u I-A)^{-1 / 2},
$$

and $S_{u}=p(A)$ for a polynomial $p$ of degree $O\left(\frac{\log (1 / \varepsilon \eta)}{\eta}\right)$.
2. We can construct a matrix $S_{\ell}$ such that

$$
S_{\ell} \approx_{\varepsilon / 10}(A-\ell I)^{-1 / 2}
$$

Moreover, $S_{\ell}$ is of the form $\left(A^{\prime}\right)^{-1 / 2} q\left(\left(A^{\prime}\right)^{-1}\right)$, where $q$ is a polynomial of degree $O\left(\frac{\log (1 / \varepsilon \eta)}{\eta}\right)$ and $A^{\prime}=L^{-1 / 2} L^{\prime} L^{-1 / 2}$ for some Laplacian matrix $L^{\prime}$.

Proof. By Taylor expansion, it holds that

$$
(1-x)^{-1 / 2}=1+\sum_{k=1}^{\infty} \prod_{j=0}^{k-1}\left(j+\frac{1}{2}\right) \frac{x^{k}}{k!}
$$

We define for any $T \in \mathbb{N}$ that

$$
p_{T}(x)=1+\sum_{k=1}^{T} \prod_{j=0}^{k-1}\left(j+\frac{1}{2}\right) \frac{x^{k}}{k!}
$$

Then, it holds for any $0<x<1-\eta$ that

$$
\begin{aligned}
p_{T}(x) \leq(1-x)^{-1 / 2} & =p_{T}(x)+\sum_{k=T+1}^{\infty} \prod_{j=0}^{k-1}\left(j+\frac{1}{2}\right) \frac{x^{k}}{k!} \\
& \leq p_{T}(x)+\sum_{k=T+1}^{\infty} x^{k} \\
& \leq p_{T}(x)+\frac{(1-\eta)^{T+1}}{\eta}
\end{aligned}
$$

Hence, it holds that

$$
(u I-A)^{-1 / 2}=u^{-1 / 2}\left(I-u^{-1} A\right)^{-1 / 2} \succeq u^{-1 / 2} p_{T}\left(u^{-1} A\right),
$$

and

$$
(u I-A)^{-1 / 2} \preceq u^{-1 / 2}\left(p_{T}\left(u^{-1} A\right)+\frac{(1-\eta)^{T+1}}{\eta} \cdot I\right),
$$

since $u^{-1} A \preceq(1-\eta) I$. Notice that $u^{-1 / 2} I \preceq(u I-A)^{-1 / 2}$, and therefore

$$
(u I-A)^{-1 / 2} \preceq u^{-1 / 2} p_{T}\left(u^{-1} A\right)+\frac{(1-\eta)^{T+1}}{\eta} \cdot(u I-A)^{-1 / 2} .
$$

Setting $T=\frac{c \log (1 /(\varepsilon \eta))}{\eta}$ for some constant $c$ and defining $S_{u}=u^{-1 / 2} p_{T}\left(u^{-1} A\right)$ gives us that

$$
S_{u} \approx_{\varepsilon / 10}(u I-A)^{-1 / 2}
$$

Now for the second statement. Our construction of $S_{\ell}$ is based on the case distinction ( $\ell>0$, and $\ell \leq 0$ ).

Case $1(\ell>0)$ : Notice that

$$
(A-\ell I)^{-1 / 2}=A^{-1 / 2}\left(I-\ell A^{-1}\right)^{-1 / 2}
$$

and

$$
p_{T}\left(\ell A^{-1}\right) \preceq\left(I-\ell A^{-1}\right)^{-1 / 2} \preceq p_{T}\left(\ell A^{-1}\right)+\frac{(1-\eta / 2)^{T+1}}{\eta / 2} \cdot I .
$$

Using the same analysis as before, we have that

$$
A^{-1 / 2}\left(I-\ell A^{-1}\right)^{-1 / 2} \approx_{\varepsilon / 10} A^{-1 / 2} p_{T}\left(\ell A^{-1}\right)
$$

By defining $S_{\ell}=A^{-1 / 2} p_{T}\left(\ell A^{-1}\right)$, i.e., $A^{\prime}=A$ and $q\left(\left(A^{\prime}\right)^{-1}\right)=p_{T}\left(\ell A^{-1}\right)$, we have that

$$
S_{\ell} \approx_{\varepsilon / 10}(A-\ell I)^{-1 / 2}
$$

Case $2(\ell \leq 0)$ : We look at the matrix

$$
A-\ell I=L^{-1 / 2} \tilde{L} L^{-1 / 2}-\ell I=L^{-1 / 2}(\tilde{L}-\ell L) L^{-1 / 2}
$$

Notice that $\tilde{L}-\ell L$ is a Laplacian matrix, and hence this reduces to the case of $\ell=0$, for which we simply set $S_{\ell}=(A-\ell I)^{-1 / 2}$. Therefore, we can write $S_{\ell}$ as a desired form, where $A^{\prime}=A-\ell I$ and polynomial $q=1$.

The following two lemmas present nearly-linear time algorithms for computing the required quantities within each iteration. We remark that an almost-linear time algorithm for computing similar quantities was shown in [9].

Lemma 4.4.3. Let $A=\sum_{i=1}^{m} v_{i} v_{i}^{T}$, and suppose that $A$ satisfies Assumption 4.4.1. Then, we can compute $\left\{r_{i}\right\}_{i=1}^{m}$ and $\left\{t_{i}\right\}_{i=1}^{m}$ in $\tilde{O}\left(\frac{m}{\varepsilon^{2} \eta}\right)$ time such that

$$
(1-\varepsilon) r_{i} \leq v_{i}^{T}(u I-A)^{-1} v_{i} \leq(1+\varepsilon) r_{i},
$$

and

$$
(1-\varepsilon) t_{i} \leq v_{i}^{T}(A-\ell I)^{-1} v_{i} \leq(1+\varepsilon) t_{i}
$$

Proof. Define $u_{i}=L^{1 / 2} v_{i}$ for any $1 \leq i \leq m$. By Lemma 4.4.2, we have that

$$
\begin{aligned}
v_{i}^{T}(u I-A)^{-1} v_{i} & \approx_{3 \varepsilon / 10}\left\|p(A) v_{i}\right\|^{2} \\
& =\left\|p\left(L^{-1 / 2} \tilde{L} L^{-1 / 2}\right) L^{-1 / 2} u_{i}\right\|^{2} \\
& =\left\|L^{1 / 2} p\left(L^{-1} \tilde{L}\right) L^{-1} u_{i}\right\|^{2} .
\end{aligned}
$$

Let $L=B^{T} B$ for some $B \in \mathbb{R}^{m \times n}$. Then, it holds that

$$
v_{i}^{T}(u I-A)^{-1} v_{i} \approx_{3 \varepsilon / 10}\left\|B p\left(L^{-1} \tilde{L}\right) L^{-1} u_{i}\right\|^{2}
$$

We invoke the Johnson-Lindenstrauss Lemma and find a random matrix $Q \in \mathbb{R}^{O\left(\log n / \varepsilon^{2}\right) \times m}$ : With high probability, it holds that

$$
v_{i}^{T}(u I-A)^{-1} v_{i} \approx_{4 \varepsilon / 10}\left\|Q B p\left(L^{-1} \tilde{L}\right) L^{-1} u_{i}\right\|^{2}
$$

We apply a nearly-linear time Laplacian solver to compute $\left\|Q B p\left(L^{-1} \tilde{L}\right) L^{-1} u_{i}\right\|^{2}$ for all $\left\{u_{i}\right\}_{i=1}^{m}$ up to $(1 \pm \varepsilon / 10)$-multiplicative error in time $\tilde{O}\left(\frac{m}{\varepsilon^{2} \eta}\right)$. This gives the desired $\left\{r_{i}\right\}_{i=1}^{m}$.

The computation for $\left\{t_{i}\right\}_{i=1}^{m}$ is similar. By Lemma 4.4.2, it holds for any $1 \leq i \leq m$ that

$$
\begin{aligned}
v_{i}^{T}(A-\ell I)^{-1} v_{i} & \approx_{3 \varepsilon / 10}\left\|\left(A^{\prime}\right)^{-1 / 2} q\left(\left(A^{\prime}\right)^{-1}\right) v_{i}\right\|^{2} \\
& =\left\|\left(A^{\prime}\right)^{-1 / 2} q\left(L^{1 / 2}\left(L^{\prime}\right)^{-1} L^{1 / 2}\right) L^{-1 / 2} u_{i}\right\|^{2} \\
& =\left\|\left(A^{\prime}\right)^{-1 / 2} L^{-1 / 2} q\left(L\left(L^{\prime}\right)^{-1}\right) u_{i}\right\|^{2} .
\end{aligned}
$$

Let $L^{\prime}=\left(B^{\prime}\right)^{T}\left(B^{\prime}\right)$ for some $B^{\prime} \in \mathbb{R}^{m \times n}$. Then, it holds that

$$
\begin{aligned}
v_{i}^{T}(A-\ell I)^{-1} v_{i} & \approx_{3 \varepsilon / 10}\left\|\left(L^{\prime}\right)^{-1 / 2} q\left(L\left(L^{\prime}\right)^{-1}\right) u_{i}\right\|^{2} \\
& =\left\|\left(L^{\prime}\right)^{1 / 2}\left(L^{\prime}\right)^{-1} q\left(L\left(L^{\prime}\right)^{-1}\right) u_{i}\right\|^{2} \\
& =\left\|B^{\prime}\left(L^{\prime}\right)^{-1} q\left(L\left(L^{\prime}\right)^{-1}\right) u_{i}\right\|^{2}
\end{aligned}
$$

We invoke the Johnson-Lindenstrauss Lemma and a nearly-linear time Laplacian solver as before to obtain required $\left\{t_{i}\right\}_{i=1}^{m}$. The total runtime is $\tilde{O}\left(\frac{m}{\eta \varepsilon^{2}}\right)$.

Lemma 4.4.4. Under Assumption 4.4.1, we can compute values $\alpha, \beta$ in $\tilde{O}\left(\frac{m}{\eta \varepsilon^{3}}\right)$ time such that

$$
(1-\varepsilon) \alpha \leq \lambda_{\min }(u I-A) \leq(1+\varepsilon) \alpha
$$

and

$$
(1-\varepsilon) \beta \leq \lambda_{\min }(A-\ell I) \leq(1+\varepsilon) \beta
$$

Proof. By Lemma 4.4.2, we have that $S_{u} \approx_{\varepsilon / 10}(u I-A)^{-1 / 2}$. Hence, $\lambda_{\max }\left(S_{u}\right)^{-2} \approx_{3 \varepsilon / 10} \lambda_{\min }(u I-A)$, and it suffices to estimate $\lambda_{\max }\left(S_{u}\right)$. Since

$$
\lambda_{\max }\left(S_{u}\right) \leq\left(\operatorname{Tr}\left(S_{u}^{2 k}\right)\right)^{1 / 2 k} \leq n^{1 / 2 k} \lambda_{\max }\left(S_{u}\right)
$$

by picking $k=\log n / \varepsilon$ we have that $\left(\operatorname{Tr}\left(S_{u}^{2 k}\right)\right)^{1 / 2 k} \approx_{\varepsilon / 2} \lambda_{\max }\left(S_{u}\right)$. Notice that

$$
\operatorname{Tr}\left(S_{u}^{2 k}\right)=\operatorname{Tr}\left(p^{2 k}\left(L^{-1 / 2} \tilde{L} L^{-1 / 2}\right)\right)=\operatorname{Tr}\left(p^{2 k}\left(L^{-1} \tilde{L}\right)\right)
$$

Set $\tilde{L}=\tilde{B}^{T} \tilde{B}$ for some matrix $\tilde{B} \in \mathbb{R}^{m \times n}$, and we have that

$$
\operatorname{Tr}\left(S_{u}^{2 k}\right)=\operatorname{Tr}\left(p^{2 k}\left(\tilde{B} L^{-1} \tilde{B}^{T}\right)\right)
$$

Since we can apply $p^{k}\left(\tilde{B} L^{-1} \tilde{B}^{T}\right)$ to vectors in $\tilde{O}\left(\frac{m}{\eta \varepsilon}\right)$ time, we invoke the Johnson-Lindenstrauss Lemma and approximate $\operatorname{Tr}\left(S_{u}^{2 k}\right)$ in $\tilde{O}\left(\frac{m}{\eta \varepsilon^{3}}\right)$ time.

We approximate $\lambda_{\min }(A-\ell I)$ in a similar way. Notice that

$$
\begin{aligned}
\operatorname{Tr}\left(S_{\ell}^{4 k}\right) & =\operatorname{Tr}\left(\left(A^{\prime}\right)^{-1 / 2} q\left(\left(A^{\prime}\right)^{-1}\right)\right)^{4 k} \\
& =\operatorname{Tr}\left(q\left(\left(A^{\prime}\right)^{-1}\right)\left(A^{\prime}\right)^{-1} q\left(\left(A^{\prime}\right)^{-1}\right)\right)^{2 k}
\end{aligned}
$$

Let $z$ be a polynomial defined by $z(x)=x q^{2}(x)$ and $L^{\prime}=\left(B^{\prime}\right)^{T}\left(B^{\prime}\right)$. Then, we have that

$$
\operatorname{Tr}\left(S_{\ell}^{4 k}\right)=\operatorname{Tr}\left(z^{2 k}\left(\left(A^{\prime}\right)^{-1}\right)\right)=\operatorname{Tr}\left(z^{2 k}\left(L^{1 / 2}\left(L^{\prime}\right)^{-1} L^{1 / 2}\right)\right)
$$

Applying the same analysis as before, we can estimate the trace in $\tilde{O}\left(\frac{m}{\eta \varepsilon^{3}}\right)$ time.

## $\ell_{p}$ Regression and John Ellipsoid

## ■ 5.1 Introduction

In this chapter, we study sparsification for $\ell_{p}$ regression problems.
For $\ell_{2}$ regression, Lemma 2.3 .3 shows that given a tall matrix, we can sample $O(n \log n)$ rows of $\boldsymbol{A}$ to form a matrix $\widetilde{\boldsymbol{A}}$ such that, for all vectors $\boldsymbol{x},\|\widetilde{\boldsymbol{A}} \boldsymbol{x}\|_{2} \approx\|\boldsymbol{A} \boldsymbol{x}\|_{2}$. Similarly, for $\ell_{p}$ regression, it is known that an approximation can be obtained by sampling rows with a certain probability (called $\ell_{p}$ Lewis weights). The main goal of this chapter is to show how to approximate $\ell_{p}$ Lewis weights by solving $\widetilde{O}(1)$ linear systems (Theorem 5.3.4). Since we know how to solve linear systems in input sparsity time, this allows us to compute $\ell_{p}$ Lewis Weights in nearly input sparsity time.

This chapter is organized as follows: In Section 5.2, we define $\ell_{p}$ Lewis weights; in Section 5.3, we show how to approximate $\ell_{p}$ Lewis weights efficiently; in Section 5.4, we discuss some geometric implications of $\ell_{p}$ Lewis weights.

## ■ 5.2 Lewis Weights

Unlike leverage scores, $\ell_{p}$ Lewis weights do not have closed formulas and are defined by the following non-linear equation.

Definition 5.2.1. For a matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ and $1 \leq p<\infty$, the $\ell_{p}$ Lewis weight is the unique vector $\boldsymbol{w} \in \mathbb{R}^{m}$ such that

$$
\boldsymbol{w}_{i}=\boldsymbol{\sigma}_{i}\left(\boldsymbol{W}^{1 / 2-1 / p} \boldsymbol{A}\right) \quad \text { for all } i \in[m]
$$

where $\boldsymbol{\sigma}_{i}(\boldsymbol{A})=\left(\boldsymbol{A}\left(\boldsymbol{A}^{\top} \boldsymbol{A}\right)^{-1} \boldsymbol{A}\right)_{i i}$.
The following result show how to use $\ell_{p}$ Lewis weights to construct sparsifiers for $\ell_{p}$ regression. They follow from a collection of results for Banach spaces (See [56] for the summary of each result).

Theorem 5.2.2 ([169, 38, 245, 246, 56]). Let $\boldsymbol{w}$ be the $\ell_{p}$ Lewis weight of some given matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ and $1 \leq p<\infty$. Given an error parameter $0<\varepsilon<1$ and sampling valuesp that

$$
\boldsymbol{p}_{i} \geq f(p, n, \varepsilon, \delta) \boldsymbol{w}_{i}
$$

| $p$ | Success probability: $1-\delta$ | Number of rows: $n \cdot f(p, n, \varepsilon, \delta)$ |
| :---: | :---: | :---: |
| $p=1$ | $1-1 / n^{C}$ | $O\left(n \log (n / \varepsilon) / \varepsilon^{2}\right)$ |
| $p=1$ | $1-1 / C$ | $O\left(n \log n / \varepsilon^{2}\right)$ |
| $1<p<2$ | $1-1 / C$ | $O\left(n \log (n / \varepsilon) \log (\log n / \varepsilon)^{2} / \varepsilon^{2}\right)$ |
| $p>2$ | $1-1 / n^{C}$ | $O\left(n^{p / 2} \log n \log (1 / \varepsilon) / \varepsilon^{5}\right)$ |

Table 5.1: $n \cdot f(p, n, \varepsilon, \delta)$ : Number of rows needed for $\ell_{p}$ sampling. The constants in big O notation depend on $C$.
where $f(p, n, \varepsilon, \delta)$ is shown in the Table 5.1. Let $\boldsymbol{S}$ be the diagonal matrix with independently chosen entries. $\boldsymbol{S}_{i i}=\frac{1}{\boldsymbol{p}_{i}^{1 / p}}$ with probability $\boldsymbol{p}_{i}$ and 0 otherwise. Then, with probability at least $1-\delta$, we have that

$$
(1-\varepsilon)\|\boldsymbol{A} \boldsymbol{x}\|_{p} \leq\|\boldsymbol{S} \boldsymbol{A} \boldsymbol{x}\|_{p} \leq(1+\varepsilon)\|\boldsymbol{A} \boldsymbol{x}\|_{p}
$$

for all $\boldsymbol{x} \in \mathbb{R}^{n}$.
This theorem shows if we have a good upper estimate of $\ell_{p}$ Lewis weights, then we can simple $\widetilde{O}\left(\max \left(n, n^{p / 2}\right)\right)$ rows to approximate the original matrix. Since $\|\boldsymbol{A} \boldsymbol{x}-\boldsymbol{b}\|_{p}=\left\|[\boldsymbol{A}, \boldsymbol{b}]\binom{\boldsymbol{x}}{-1}\right\|_{p}$, Lewis weights allow us to reduce a $\ell_{p}$ regression problem involving a tall matrix to a matrix with only $\widetilde{O}\left(\max \left(n, n^{p / 2}\right)\right)$ rows.

## ■ 5.3 Approximate Lewis Weights

In this section, we show how to compute $\ell_{p}$ Lewis weights by solving $\widetilde{O}(1)$ linear systems for $1 \leq p<\infty$. For $p<4$, this can be done by the following algorithm [56].

```
Algorithm 5: Approximate \(\ell_{p}\) Lewis weights for \(p<4\)
    Let \(\boldsymbol{v}_{i}^{(1)}=n / m\) for all \(i \in[m]\) and \(T=\Theta\left(\frac{\log (m / \varepsilon)}{4-p}\right)\).
    for \(k=1, \cdots, T\) do
        constant \(c_{1}\).
        \(\boldsymbol{v}_{i}^{(k+1)}=\left(\left(\boldsymbol{v}_{i}^{(k)}\right)^{2 / p-1} \tilde{\boldsymbol{\sigma}}_{i}\right)^{p / 2}\).
    end
    Output: \(\boldsymbol{v}^{(T+1)}\).
```

    Input: a matrix \(\boldsymbol{A} \in \mathbb{R}^{m \times n}\), accuracy \(0<\varepsilon<1,1 \leq p<4\).
        Let \(\tilde{\boldsymbol{\sigma}}\) be a \(c_{1}(4-p) \varepsilon\)-approximation of \(\boldsymbol{\sigma}_{i}\left(\left(\boldsymbol{V}^{(k)}\right)^{1 / 2-1 / p} \boldsymbol{A}\right)\) for some small enough
    Theorem 5.3.1 ([56]). For any matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}$, accuracy $0<\varepsilon<1$ and $1 \leq p<4$, the algorithm 5 outputs a vector $\boldsymbol{v}$ in time $\widetilde{O}\left(\mathcal{T} \varepsilon^{-2}\right)$ with high probability such that

$$
\begin{equation*}
(1-\varepsilon) \boldsymbol{w}_{i} \leq \boldsymbol{v}_{i} \leq(1+\varepsilon) \boldsymbol{w}_{i} \tag{5.1}
\end{equation*}
$$

where $\boldsymbol{w}$ is the $\ell_{p}$ Lewis weight for $\boldsymbol{A}$ and $\mathcal{T}$ be the time needed to apply $\left(\boldsymbol{A}^{\top} \boldsymbol{D} \boldsymbol{A}\right)^{-1}$ to a vector for any diagonal matrix $\boldsymbol{D}$.

This algorithm relies on the stability of $\ell_{p}$ Lewis weights for $1 \leq p<4$, namely, if we rescale each row of $\boldsymbol{A}$ by a constant, the Lewis weight of $\boldsymbol{A}$ changes only by a constant. However, this is not true for large $p$. For $p \sim \infty$, the Lewis weights can change arbitrarily even with a tiny rescaling and hence it is almost impossible to approximate Lewis weight in the sense of (5.1). To avoid this inability issue, we define the $\varepsilon$-approximate $\ell^{p}$ Lewis weight as follows:

Definition 5.3.2. For a matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ and $1 \leq p<\infty$, we call a vector $\boldsymbol{w} \in \mathbb{R}^{m}$ is an $\varepsilon$ approximate $\ell^{p}$ Lewis weight for $\boldsymbol{A}$ if

$$
(1-\varepsilon) \boldsymbol{\sigma}_{i}\left(\boldsymbol{W}^{1 / 2-1 / p} \boldsymbol{A}\right) \leq \boldsymbol{w}_{i} \leq(1+\varepsilon) \boldsymbol{\sigma}_{i}\left(\boldsymbol{W}^{1 / 2-1 / p} \boldsymbol{A}\right) .
$$

For $1 \leq p<4$, an $\varepsilon$-approximate $\ell^{p}$ Lewis weight is indeed a $O(\varepsilon)$ approximate of the true $\ell^{p}$ Lewis weight in the sense of (5.1). Although this is not true for larger $p$, the following lemma shows that $\varepsilon$-approximate $\ell^{p}$ Lewis weights suffice for the sampling purpose.

Lemma 5.3.3. Given an error parameter $0<\varepsilon<\frac{1}{2}$. Let $\boldsymbol{w}$ be an $\varepsilon$-approximate $\ell_{p}$ Lewis weight of some given matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ and $1 \leq p<\infty$. Let $\boldsymbol{p}$ be the sampling values such that

$$
\boldsymbol{p}_{i} \geq f(p, n, \varepsilon, \delta) \boldsymbol{w}_{i}
$$

where $f(p, n, \varepsilon, \delta)$ is shown in the Table 5.1. Let $\boldsymbol{S}$ be the diagonal matrix with independently chosen entries. $\boldsymbol{S}_{i i}=\frac{1}{\boldsymbol{p}_{i}^{1 / p}}$ with probability $\boldsymbol{p}_{i}$ and 0 otherwise. Then, with probability at least $1-\delta$, we have that

$$
(1-2 \varepsilon)\|\boldsymbol{A} \boldsymbol{x}\|_{p} \leq\|\boldsymbol{S} \boldsymbol{A} \boldsymbol{x}\|_{p} \leq(1+2 \varepsilon)\|\boldsymbol{A} \boldsymbol{x}\|_{p}
$$

for all $\boldsymbol{x} \in \mathbb{R}^{n}$.
Proof. Let $\boldsymbol{d}_{i}=\boldsymbol{\sigma}_{i}\left(\boldsymbol{W}^{1 / 2-1 / p} \boldsymbol{A}\right) / \boldsymbol{w}_{i}, \widetilde{\boldsymbol{A}}=\boldsymbol{D}^{-1 / 2+1 / p} \boldsymbol{A}$ and $\widetilde{\boldsymbol{W}}=\boldsymbol{D} \boldsymbol{W}$. From the definition of $\boldsymbol{\sigma}_{i}$ (Definition 5.2.1), we see that

$$
\begin{aligned}
\boldsymbol{\sigma}_{i}\left(\widetilde{\boldsymbol{W}}^{1 / 2-1 / p} \widetilde{\boldsymbol{A}}\right) & =\boldsymbol{\sigma}_{i}\left(\boldsymbol{W}^{1 / 2-1 / p} \boldsymbol{A}\right) \\
& =\boldsymbol{d}_{i} \boldsymbol{w}_{i}=\widetilde{\boldsymbol{w}}_{i} .
\end{aligned}
$$

Hence, $\widetilde{\boldsymbol{w}}$ is the $\ell_{p}$ Lewis weight of $\widetilde{\boldsymbol{A}}$ and theorem 5.2.2 shows that

$$
(1-\varepsilon)\|\tilde{\boldsymbol{A}} \boldsymbol{x}\|_{p} \leq\|\boldsymbol{S} \widetilde{\boldsymbol{A}} \boldsymbol{x}\|_{p} \leq(1+\varepsilon)\|\widetilde{\boldsymbol{A}} \boldsymbol{x}\|_{p}
$$

From the definition of $\varepsilon$-approximate $\ell_{p}$ Lewis weight, we know that $1-\varepsilon \leq \boldsymbol{d}_{i} \leq 1+\varepsilon$ and hence

$$
(1-2 \varepsilon)\|\boldsymbol{A} \boldsymbol{x}\|_{p} \leq\|\boldsymbol{S} \boldsymbol{A} \boldsymbol{x}\|_{p} \leq(1+2 \varepsilon)\|\boldsymbol{A} \boldsymbol{x}\|_{p} .
$$

Note that if we have a constant approximate of leverage scores, we can construct an arbitrarily good approximation by sampling more rows. However, this is not true for $\ell_{p}$ Lewis weights; we need a good approximation of Lewis weight to get a good approximate of the original matrix.

Now, we have everything to describe our result, a simple multiplicative weight update algorithm.

```
Algorithm 6: Approximate \(\ell_{p}\) Lewis weights for \(p \geq 2\)
    Input: a matrix \(\boldsymbol{A} \in \mathbb{R}^{m \times n}\), accuracy \(0<\varepsilon<1, p \geq 2\).
    Let \(\boldsymbol{v}_{i}^{(1)}=n / m\) for all \(i \in[m]\) and \(T=\Theta\left(\frac{\log (m / n)}{\varepsilon}\right)\).
    for \(k=1, \cdots, T\) do
        Let \(\tilde{\boldsymbol{\sigma}}\) be a \(\frac{\varepsilon}{4}\)-approximation of \(\boldsymbol{\sigma}_{i}\left(\left(\boldsymbol{V}^{(k)}\right)^{1 / 2-1 / p} \boldsymbol{A}\right)\).
        \(\boldsymbol{v}_{i}^{(k+1)}=\tilde{\boldsymbol{\sigma}}_{i}\).
    end
    Output: \(\frac{1}{T} \sum_{k=1}^{T} \boldsymbol{v}^{(k)}\).
```

Theorem 5.3.4. For any matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}$, accuracy $0<\varepsilon<1$ and $p \geq 2$, the algorithm 6 outputs an $\varepsilon$-approximate $\ell_{p}$ Lewis weight in time $\widetilde{O}\left(\mathcal{T} \varepsilon^{-3}\right)$ where $\mathcal{T}$ be the time needed to apply $\left(\boldsymbol{A}^{\top} \boldsymbol{D} \boldsymbol{A}\right)^{-1}$ to a vector for any diagonal matrix $\boldsymbol{D}$.

Since we know how to solve linear systems in input sparsity time (Chapter 3), this theorem gives an input sparsity time algorithm for approximating Lewis weights. This algorithm is based on the following convexity result.

Lemma 5.3.5. For any matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}, i \in[m]$ and $p \geq 2$, the function

$$
\phi_{i}(\boldsymbol{w})=\log \left(\boldsymbol{a}_{i}^{\top}\left(\boldsymbol{A}^{\top} \boldsymbol{W}^{1-2 / p} \boldsymbol{A}\right)^{-1} \boldsymbol{a}_{i} \boldsymbol{w}_{i}^{-2 / p}\right)
$$

is convex.

Proof. For any vector $\boldsymbol{h} \in \mathbb{R}^{m}$, we define $\phi_{i}(t)=\phi_{i}(\boldsymbol{w}+t \boldsymbol{h})$. Let $\boldsymbol{M}=\left(\boldsymbol{A}^{\top} \boldsymbol{W}^{1-2 / p} \boldsymbol{A}\right)^{-1}$, we have that

$$
\begin{aligned}
\left.\frac{d}{d t}\right|_{t=0} \phi_{i}(t)= & -\frac{2}{p} \frac{\boldsymbol{h}_{i}}{\boldsymbol{w}_{i}}-\frac{(1-2 / p) \boldsymbol{a}_{i}^{\top} \boldsymbol{M} \boldsymbol{A}^{\top} \boldsymbol{H} \boldsymbol{W}^{-2 / p} \boldsymbol{A} \boldsymbol{M} \boldsymbol{a}_{i}}{\boldsymbol{a}_{i}^{\top} \boldsymbol{M} \boldsymbol{a}_{i}} \\
\left.\frac{d^{2}}{d t^{2}}\right|_{t=0} \phi_{i}(t)= & \frac{2}{p} \frac{\boldsymbol{h}_{i}^{2}}{\boldsymbol{w}_{i}^{2}}-\frac{(1-2 / p)^{2}\left(\boldsymbol{a}_{i}^{\top} \boldsymbol{M} \boldsymbol{A}^{\top} \boldsymbol{H} \boldsymbol{W}^{-2 / p} \boldsymbol{A} \boldsymbol{M} \boldsymbol{a}_{i}\right)^{2}}{\left(\boldsymbol{a}_{i}^{\top} \boldsymbol{M} \boldsymbol{a}_{i}\right)^{2}} \\
& -\frac{(-2 / p)(1-2 / p) \boldsymbol{a}_{i}^{\top} \boldsymbol{M} \boldsymbol{A}^{\top} \boldsymbol{H}^{2} \boldsymbol{W}^{-1-2 / p} \boldsymbol{A} \boldsymbol{M} \boldsymbol{a}_{i}}{\boldsymbol{a}_{i}^{\top} \boldsymbol{M} \boldsymbol{a}_{i}} \\
& +\frac{2(1-2 / p)^{2} \boldsymbol{a}_{i}^{\top} \boldsymbol{M} \boldsymbol{A}^{\top} \boldsymbol{H} \boldsymbol{W}^{-2 / p} \boldsymbol{A} \boldsymbol{M} \boldsymbol{A}^{\top} \boldsymbol{H} \boldsymbol{W}^{-2 / p} \boldsymbol{A} \boldsymbol{M} \boldsymbol{a}_{i}}{\boldsymbol{a}_{i}^{\top} \boldsymbol{M} \boldsymbol{a}_{i}}
\end{aligned}
$$

By Cauchy-Schwarz inequality, we have that

$$
\left(\boldsymbol{a}_{i}^{\top} \boldsymbol{M} \boldsymbol{A}^{\top} \boldsymbol{H} \boldsymbol{W}^{-2 / p} \boldsymbol{A} \boldsymbol{M} \boldsymbol{a}_{i}\right)^{2} \leq\left(\boldsymbol{a}_{i}^{\top} \boldsymbol{M} \boldsymbol{A}^{\top} \boldsymbol{H} \boldsymbol{W}^{-2 / p} \boldsymbol{A} \boldsymbol{M} \boldsymbol{A}^{\top} \boldsymbol{H} \boldsymbol{W}^{-2 / p} \boldsymbol{A} \boldsymbol{M} \boldsymbol{a}_{i}\right)\left(\boldsymbol{a}_{i}^{\top} \boldsymbol{M} \boldsymbol{a}_{i}\right)
$$

and hence

$$
\begin{aligned}
\left.\frac{d^{2}}{d t^{2}}\right|_{t=0} \phi_{i}(t) \geq & \frac{2}{p}\left(\frac{\boldsymbol{h}_{i}^{2}}{\boldsymbol{w}_{i}^{2}}+\frac{(1-2 / p) \boldsymbol{a}_{i}^{\top} \boldsymbol{M} \boldsymbol{A}^{\top} \boldsymbol{H}^{2} \boldsymbol{W}^{-1-2 / p} \boldsymbol{A} \boldsymbol{M} \boldsymbol{a}_{i}}{\boldsymbol{a}_{i}^{\top} \boldsymbol{M} \boldsymbol{a}_{i}}\right) \\
& +\frac{(1-2 / p)^{2} \boldsymbol{a}_{i}^{\top} \boldsymbol{M} \boldsymbol{A}^{\top} \boldsymbol{H} \boldsymbol{W}^{-2 / p} \boldsymbol{A} \boldsymbol{M} \boldsymbol{A}^{\top} \boldsymbol{H} \boldsymbol{W}^{-2 / p} \boldsymbol{A} \boldsymbol{M} \boldsymbol{a}_{i}}{\boldsymbol{a}_{i}^{\top} \boldsymbol{M} \boldsymbol{a}_{i}} \\
\geq & 0
\end{aligned}
$$

Since this holds for any any vector $\boldsymbol{h} \in \mathbb{R}^{m}$, this shows that $\phi_{i}$ is convex.

Now, we prove Theorem 5.3.4.

Proof of Theorem 5.3.4. Let $\overline{\boldsymbol{v}}=\frac{1}{T} \sum_{k=1}^{T} \boldsymbol{v}_{i}^{(k)}$. Let $\boldsymbol{d}_{i}^{(k)}=\tilde{\boldsymbol{\sigma}}_{i} / \boldsymbol{\sigma}_{i}\left(\left(\boldsymbol{V}^{(k)}\right)^{1 / 2-1 / p} \boldsymbol{A}\right)$. Then, we have that

$$
\begin{aligned}
\boldsymbol{v}_{i}^{(k+1)} & =\boldsymbol{d}_{i}^{(k)} \boldsymbol{\sigma}_{i}\left(\left(\boldsymbol{V}^{(k)}\right)^{1 / 2-1 / p} \boldsymbol{A}\right) \\
& =\boldsymbol{d}_{i}^{(k)}\left(\boldsymbol{v}_{i}^{(k)}\right)^{1-2 / p} \boldsymbol{a}_{i}^{\top}\left(\boldsymbol{A}^{\top}\left(\boldsymbol{V}^{(k)}\right)^{1-2 / p} \boldsymbol{A}\right)^{-1} \boldsymbol{a}_{i}
\end{aligned}
$$

Hence, we have

$$
\log \boldsymbol{v}_{i}^{(k+1)}=\log \boldsymbol{v}_{i}^{(k)}+\log \left(\boldsymbol{a}_{i}^{\top}\left(\boldsymbol{A}^{\top}\left(\boldsymbol{V}^{(k)}\right)^{1-2 / p} \boldsymbol{A}\right)^{-1} a_{i}\left(\boldsymbol{v}_{i}^{(k)}\right)^{-2 / p}\right)+\log \left(\boldsymbol{d}_{i}^{(k)}\right)
$$

Using $\boldsymbol{v}_{i}^{(1)}=\frac{n}{m}$ and $\boldsymbol{v}_{i}^{(T+1)} \leq 1$, we have that

$$
\begin{aligned}
0 & \geq \log \boldsymbol{v}_{i}^{(T+1)} \\
& =\log \boldsymbol{v}_{i}^{(1)}+\sum_{k=1}^{T} \log \left(\boldsymbol{a}_{i}^{\top}\left(\boldsymbol{A}^{\top}\left(\boldsymbol{V}^{(k)}\right)^{1-2 / p} \boldsymbol{A}\right)^{-1} \boldsymbol{a}_{i}\left(\boldsymbol{v}_{i}^{(k)}\right)^{-2 / p}\right)+\sum_{k=1}^{T} \log \left(\boldsymbol{d}_{i}^{(k)}\right) .
\end{aligned}
$$

By Lemma 5.3.5, we have that

$$
\log \left(\boldsymbol{a}_{i}^{\top}\left(\boldsymbol{A}^{\top} \overline{\boldsymbol{V}}^{1-2 / p} \boldsymbol{A}\right)^{-1} \boldsymbol{a}_{i} \overline{\boldsymbol{v}}_{i}^{-2 / p}\right) \leq \frac{1}{T} \sum_{k=1}^{T} \log \left(\boldsymbol{a}_{i}^{\top}\left(\boldsymbol{A}^{\top}\left(\boldsymbol{V}^{(k)}\right)^{1-2 / p} \boldsymbol{A}\right)^{-1} \boldsymbol{a}_{i}\left(\boldsymbol{v}_{i}^{(k)}\right)^{-2 / p}\right) .
$$

Hence, we have that

$$
\begin{aligned}
\log \left(\frac{\boldsymbol{\sigma}_{i}\left(\overline{\boldsymbol{V}}^{1 / 2-1 / p} \boldsymbol{A}\right)}{\overline{\boldsymbol{v}}_{i}}\right) & =\log \left(\boldsymbol{a}_{i}^{\top}\left(\boldsymbol{A}^{\top} \overline{\boldsymbol{V}}^{1-2 / p} \boldsymbol{A}\right)^{-1} \boldsymbol{a}_{i} \overline{\boldsymbol{v}}_{i}^{-2 / p}\right) \\
& \leq \frac{\log \left(\frac{m}{n}\right)}{T}-\frac{1}{T} \sum_{i=1}^{T} \log \left(\boldsymbol{d}_{i}^{(k)}\right) \\
& \leq \frac{\log \left(\frac{m}{n}\right)}{T}+\frac{1}{2} \varepsilon
\end{aligned}
$$

By setting $T$ large enough, we have that

$$
(1-\varepsilon) \boldsymbol{\sigma}_{i}\left(\overline{\boldsymbol{V}}^{1 / 2-1 / p} \boldsymbol{A}\right) \leq \overline{\boldsymbol{v}}_{i} \leq(1+\varepsilon) \boldsymbol{\sigma}_{i}\left(\overline{\boldsymbol{V}}^{1 / 2-1 / p} \boldsymbol{A}\right)
$$

## ■ 5.4 Approximate John Ellipsoids

In this section, we discuss some geometric implications of $\ell_{p}$ Lewis weight. [270] shows that the $\ell_{p}$ Lewis weight is given by the following optimization problem.

Lemma 5.4.1. Given a matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}$. Let $\boldsymbol{M}$ be the maximizer of the following optimization problem

$$
\begin{equation*}
\max _{\boldsymbol{M}} \log \operatorname{det} \boldsymbol{M} \text { subject to } \sum_{i=1}^{m}\left(\boldsymbol{a}_{i}^{\top} \boldsymbol{M} \boldsymbol{a}_{i}\right)^{p / 2} \leq n \text { and } \boldsymbol{M} \succeq \mathbf{0} \text {. } \tag{5.2}
\end{equation*}
$$

Then, the $\ell_{p}$ Lewis weight of $\boldsymbol{A}$ is given by $\boldsymbol{w}_{i}=\left(\boldsymbol{a}_{i}^{\top} \boldsymbol{M} \boldsymbol{a}_{i}\right)^{p / 2}$. Furthermore, we have that $\boldsymbol{M}^{-1}=$ $\sum_{i} \boldsymbol{w}_{i}^{1-2 / p} \boldsymbol{a}_{i} \boldsymbol{a}_{i}^{\top}$.

To understand this optimization problem, we define the $\ell_{p}$ Lewis Ellipsoids and John ellipsoids as follows:

Definition 5.4.2. For any matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}$, we define the $\ell_{p}$ Lewis Ellipsoid of $\boldsymbol{A}$ by $E=\{\boldsymbol{x} \in$ $\mathbb{R}^{n}$ such that $\left.\boldsymbol{x}^{\top} \boldsymbol{A}^{\top} \boldsymbol{W}^{1-2 / p} \boldsymbol{A} \boldsymbol{x} \leq 1\right\}$ where $\boldsymbol{w}$ is the $\ell_{p}$ Lewis weight of $\boldsymbol{A}$. For any convex set $K$, we define the John ellipsoid of $K$ is the maximum volume ellipsoid contained inside $K$

We note that $\sum_{i=1}^{m}\left(\boldsymbol{a}_{i}^{\top} \boldsymbol{M} \boldsymbol{a}_{i}\right)^{p / 2}$ is proportional to the average value of $\|\boldsymbol{A} \boldsymbol{x}\|_{p}^{p}$ inside the ellipsoid $\left\{x^{\top} \boldsymbol{M}^{-1} \boldsymbol{x} \leq 1\right\}$. Therefore, the $\ell_{p}$ Lewis Ellipsoids is the maximum volume ellipsoid $E$ such that the average value of $\|\boldsymbol{A} \boldsymbol{x}\|_{p}^{p}$ inside that ellipsoid is less than certain value. In particular, the $\ell_{\infty}$ Lewis

Ellipsoid of $\boldsymbol{A}$ exactly corresponds to the John ellipsoid of the symmetric polytope $\left\{\|\boldsymbol{A}\|_{\infty} \leq 1\right\}$. Therefore, our algorithm for computing Lewis weights gives an efficient algorithm for approximating John ellipsoids for symmetric polytopes.

Lemma 5.4.3. Given a symmetric polytope $P=\left\{\boldsymbol{x}\right.$ such that $\left.\|\boldsymbol{A} \boldsymbol{x}\|_{\infty} \leq 1\right\}$. Then, we can find $\boldsymbol{v}$ such that

$$
\boldsymbol{\sigma}_{i}\left(\boldsymbol{V}^{1 / 2} \boldsymbol{A}\right) \leq \boldsymbol{v}_{i} \leq(1+\varepsilon) \boldsymbol{\sigma}_{i}\left(\boldsymbol{V}^{1 / 2} \boldsymbol{A}\right)
$$

in time $\widetilde{O}\left(\mathcal{T} \varepsilon^{-3}\right)$ where $\mathcal{T}$ be the time needed to apply $\left(\boldsymbol{A}^{\top} \boldsymbol{D} \boldsymbol{A}\right)^{-1}$ to a vector for any diagonal matrix D. Furthermore, the ellipsoid $E=\left\{\boldsymbol{x} \in \mathbb{R}^{n}\right.$ such that $\left.\boldsymbol{x}^{\top} \boldsymbol{A}^{\top} \boldsymbol{V} \boldsymbol{A} \boldsymbol{x} \leq 1\right\}$ satisfies the rounding condition

$$
\begin{equation*}
E \subset P \subset \sqrt{(1+\varepsilon) n} E . \tag{5.3}
\end{equation*}
$$

Proof. By Lemma 5.3.4, we can find $\boldsymbol{v}$ such that $\left(1-\frac{\varepsilon}{3}\right) \boldsymbol{\sigma}_{i}\left(\boldsymbol{V}^{1 / 2} \boldsymbol{A}\right) \leq \boldsymbol{v}_{i} \leq\left(1+\frac{\varepsilon}{3}\right) \boldsymbol{\sigma}_{i}\left(\boldsymbol{V}^{1 / 2} \boldsymbol{A}\right)$ in time $\widetilde{O}\left(\mathcal{T} \varepsilon^{-3}\right)$. By scaling $\boldsymbol{v}$ up $1 /(1-\varepsilon / 3)$, we get the first guarantee.

By $\boldsymbol{\sigma}_{i}\left(\boldsymbol{V}^{1 / 2} \boldsymbol{A}\right) \leq \boldsymbol{v}_{i}$, we know that for all $i$

$$
\boldsymbol{a}_{i}^{\top}\left(\boldsymbol{A}^{\top} \boldsymbol{V} \boldsymbol{A}\right)^{-1} \boldsymbol{a}_{i} \leq 1
$$

Since $\left(\boldsymbol{a}_{i}^{\top} \boldsymbol{x}\right)^{2} \leq\left(\boldsymbol{a}_{i}^{\top}\left(\boldsymbol{A}^{\top} \boldsymbol{V} \boldsymbol{A}\right)^{-1} \boldsymbol{a}_{i}\right)\left(\boldsymbol{x}^{\top} \boldsymbol{A}^{\top} \boldsymbol{V} \boldsymbol{A} \boldsymbol{x}\right)$, this shows that $E \subset P$.
On the other hand, for any $\boldsymbol{x} \in P$, we have that

$$
\begin{aligned}
\boldsymbol{x}^{\top} \boldsymbol{A}^{\top} \boldsymbol{V} \boldsymbol{A} \boldsymbol{x} & \leq \sum_{i=1}^{m} \boldsymbol{v}_{i}\|\boldsymbol{A} \boldsymbol{x}\|_{\infty}^{2} \\
& \leq \sum_{i=1}^{m}(1+\varepsilon) \boldsymbol{\sigma}_{i}\left(\boldsymbol{V}^{1 / 2} \boldsymbol{A}\right) \\
& =(1+\varepsilon) n .
\end{aligned}
$$

Therefore, we have that $P \subset \sqrt{(1+\varepsilon) n} E$.
Remark 5.4.4. We call the ellipsoid found in this theorem is an approximate John ellipsoid because it almost satisfies the optimality condition of John ellipsoid, i.e. $\boldsymbol{v}_{i}=\boldsymbol{\sigma}_{i}\left(\boldsymbol{V}^{1 / 2} \boldsymbol{A}\right)$. Although this does not show that the ellipsoid is actually close to the true John ellipsoid, this has similar rounding guarantee (5.3) as John ellipsoid and hence satisfies the requirement for most of the applications such as volume computation [262] and solving linear systems (next chapter).

# Faster Linear Programming 

## ■ 6.1 Introduction

In this chapter, we present a new algorithm for the following linear programs

$$
\begin{equation*}
\underset{\substack{\boldsymbol{x} \in \mathbb{R}^{m} \\ \forall i \in[m]: \boldsymbol{A}^{T} \boldsymbol{x}=\boldsymbol{b} \\: l_{i} \leq x_{i} \leq u_{i}}}{ } \boldsymbol{c}^{T} \boldsymbol{x} . \tag{6.1}
\end{equation*}
$$

Our algorithm is based on interior point methods, a framework for solving linear programming. Roughly speaking, interior point methods approximate the domain by a sequence of ellipsoids and the number of iterations depends on how good are these approximations. In Chapter 5, we showed a family of ellipsoids that is easy to compute and gives a good approximations for arbitrary symmetric polytopes. Using these ellipsoids, we develop an interior point method that takes $\widetilde{O}(\sqrt{\operatorname{rank}(\boldsymbol{A})} L)^{1}$ iterations (Theorem 6.8.1). Furthermore, we show how to achieve this convergence rate while only solving $\widetilde{O}(1)$ linear systems and performing additional $\widetilde{O}(\operatorname{nnz}(\boldsymbol{A}))$ work in each iteration. This is the first polynomial improvement on the running time of solving linear programming over 25 years.

Applying our techniques to the linear programming formulation of the maximum flow problem we achieve a running time of $\widetilde{O}(|E| \sqrt{|V|} \log (U))$ for solving the problem on a graph with $|E|$ edges, $|V|$ vertices, and capacity ratio $U$ (Theorem 6.9.3). This improves upon the previous fastest running time of $\widetilde{O}\left(|E| \min \left\{|E|^{1 / 2},|V|^{2 / 3}\right\} \log (U)\right)$ achieved over 15 years ago by Goldberg and Rao $[104]^{2}$ and the previous fastest running times for solving dense directed unit capacity graphs of $O\left(|E| \min \left\{|E|^{1 / 2},|V|^{2 / 3}\right\}\right)$ achieved by Even and Tarjan [85] over 35 years ago and of $\widetilde{O}\left(|E|^{10 / 7}\right)$ achieved recently by Mądry [183].

## ■ 6.1.1 Organization

The rest of this chapter is structured as follows. In Section 6.3, we provide some background information on linear programming and interior point methods and show how to obtain the "first" $\widetilde{O}(\sqrt{\operatorname{rank}(\boldsymbol{A})} L)$ iterations algorithm for the standard dual linear program (6.2). In Section 6.4, we explain why this algorithm is not sufficient for the maximum flow problem. In Section 6.5 we introduce our framework to use our alternative barrier function more efficiently and in Section 6.6 we present the key lemmas used to analyze progress along paths and in Section 6.7 we introduce the weight function we use to find paths. In Section 6.8 provide our linear programming algorithm for the linear program (6.1). In Section 6.9 we use these results to achieve our desired running times for the maximum flow

[^6]problem and its generalizations.

## ■ 6.2 Previous Work

Linear programming and maximum flow are extremely well studied problems with long histories. Here we focus on recent history that particularly influenced our work.

## ■ 6.2.1 Convergence Rate of Interior Point Methods

Given a matrix, $\boldsymbol{A} \in \mathbb{R}^{m \times n}$, and vectors, $\boldsymbol{b} \in \mathbb{R}^{m}$ and $\boldsymbol{c} \in \mathbb{R}^{n}$, solving the linear program ${ }^{3}$

$$
\begin{equation*}
\min _{x \in \mathbb{R}^{n}: A x \geq \boldsymbol{b}} \boldsymbol{c}^{T} \boldsymbol{x} \tag{6.2}
\end{equation*}
$$

is a core algorithmic task for both the theory and practice of computer science.
In 1984, Karmarkar [136] provided the first proof of an interior point method running in polynomial time. This method required $O(m L)$ iterations where the running time of each iteration was dominated by the time needed to solve a linear system of the form $\left(\boldsymbol{A}^{T} \boldsymbol{D} \boldsymbol{A}\right) \boldsymbol{x}=\boldsymbol{y}$ for some positive diagonal matrix $\boldsymbol{D} \in \mathbb{R}^{m \times m}$ and some $\boldsymbol{y} \in \mathbb{R}^{n}$. Using low rank matrix updates and preconditioning Karmarkar achieved a running time of $O\left(m^{3.5} L\right)$.

In 1988, Renegar improved the convergence rate of interior point methods to $O(\sqrt{m} L)$ iterations [227]. He presented a path following method where the iterations had comparable complexity to Karmarkar's method. Using a combination of techniques involving low rank updates, preconditioning and fast matrix multiplication, the amortized complexity of each iteration was improved [254, 107, 211] yielding the current state of the art running time of $\tilde{O}\left(m^{1.5} n L\right)$ [257].

In 1989, Vaidya [258] proposed two barrier functions which were shown to yield $O\left((\operatorname{mrank}(\boldsymbol{A}))^{1 / 4} L\right)$ and $O(\operatorname{rank}(\boldsymbol{A}) L)$ iteration linear programming algorithms [255, 258, 253]. Unfortunately each iteration of these methods required explicit computation of the projection matrix $\boldsymbol{D}^{1 / 2} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{D} \boldsymbol{A}\right)^{-1} \boldsymbol{A}^{T} \boldsymbol{D}^{1 / 2}$ for a positive diagonal matrix $\boldsymbol{D} \in \mathbb{R}^{m \times m}$. This was improved by Anstreicher [15] who showed it sufficed to compute the diagonal of this projection matrix. Unfortunately both methods do not yield faster running times unless $m \gg n$.

In 1994 Nesterov and Nemirovski [211] showed that path-following methods can in principle be applied to any solve convex optimization problem, given a suitable barrier function. The number of iterations of their method depends on a square root of a quantity known as the self-concordance of the barrier. They showed that for any convex set in $\mathbb{R}^{n}$, there exists a barrier function, called the universal barrier function, with self-concordance $O(n)$ and therefore in principle any convex optimization problem with $n$ variables can be solved in $O(\sqrt{n} L)$ iterations. However, this result is generally considered to be only of theoretical interest as the universal barrier function is defined as the volume of certain polytopes, a problem which in full-genearality is NP-hard and can only be computed approximately in $O\left(n^{c}\right)$ for some large constant $c$ [180].

These results suggested that you can solve linear programs closer to the $\tilde{O}(\sqrt{\operatorname{rank}(\boldsymbol{A})} L)$ bound achieved by the universal barrier only if you pay more in each iteration. In this chapter we show that this is not the case and up to polylogarithmic factors we achieve the convergence rate of the universal barrier function while only having iterations of cost comparable to that of Renegar's algorithm.

[^7]| Year | Author | Number of Iterations | Nature of iterations |
| :---: | :--- | :---: | :---: |
| 1984 | Karmarkar [136] | $O(m L)$ | Linear system solve |
| 1986 | Renegar [227] | $O(\sqrt{m} L)$ | Linear system solve |
| 1989 | Vaidya [253] | $O\left((m \operatorname{rank}(\boldsymbol{A}))^{1 / 4} L\right)$ | Expensive linear algebra |
| 1994 | Nesterov and Nemirovskii [211] | $O(\sqrt{\operatorname{rank}(\boldsymbol{A})} L)$ | Volume computation |
| - | This chapter | $\tilde{O}(\sqrt{\operatorname{rank}(\boldsymbol{A})} L)$ | $\widetilde{O}(1)$ Linear system solves |

Figure 6-1: Interior Point Iteration Improvements

## ■ 6.2.2 Recent Breakthroughs on the Maximum Flow Problem

A beautiful line of work on solving the maximum flow on undirected graphs began with a result of Benzcur and Karger in which they showed how to reduce approximately computing minimum s-t cuts on arbitrary undirected graphs to the same problem on sparse graphs, i.e. those with $|E|=\widetilde{O}(|V|)$ [33]. Pushing these ideas further Karger showed how to perform a similar reduction for the maximum flow problem [134] and Karger and Levine showed how to compute the exact maximum flow in an unweighted undirected graph with maximum flow value $F$ in $\widetilde{O}(|E|+|V| F)$ time [135].

In a breakthrough result of Spielman and Teng [243] they showed that a particular class of linear systems, Laplacians, can be solved in nearly linear time. Leveraging these fast Laplacian system solvers Christiano, Kelner, Mądry, and Spielman [51] showed how to compute $(1-\varepsilon)$ approximations to the maximum flow problem on undirected graphs in $\widetilde{O}\left(|E| \cdot|V|^{1 / 3} \varepsilon^{-11 / 3}\right)$. Later Lee, Rao and Srivastava [161] improved the running time to $\widetilde{O}\left(|E| \cdot|V|^{1 / 3} \varepsilon^{-2 / 3}\right)$ for unweighted undirected graphs. This line of work culminated in recent results of Sherman [233] and Kelner, Lee, Orecchia, and Sidford (Chapter 12) who showed how to solve the problem in $O\left(|E|^{1+o(1)} \varepsilon^{-2}\right)$, using congestion-approximators, oblivious routings, and techniques developed by Mądry [182].

In 2008, Daitch and Spielman [63] showed that, by careful application of interior point techniques, fast Laplacian system solvers [243], and a novel method for solving M-matrices, they could match (up to polylogarithmic factors) the running time of Goldberg Rao and achieve a running time of $\widetilde{O}\left(|E|^{3 / 2} \log ^{2}(U)\right)$ not just for maximum flow but also for the minimum cost flow and lossy generalized minimum cost flow problems (see Fig 6-2).

Very recently Mądry [183] achieved an astounding running time of $\widetilde{O}\left(|E|^{10 / 7}\right)$ for solving the maximum flow problem on un-capacitated directed graphs by a novel application and modification of interior point methods. This shattered numerous barriers providing the first general improvement over the running time of $O\left(|E| \min \left\{|E|^{1 / 2},|V|^{2 / 3}\right\}\right)$ for solving unit capacity graphs proven over 35 years ago by Even and Tarjan [85] in 1975.

While our algorithm for solving the maximum flow problem is new, we make extensive use of these breakthroughs. We use sampling techniques first discovered in the context of graph sparsification (previous two chapters) to re-weight the graph so that we make progress at a rate commensurate with the number of vertices and not the number of edges. We use fast Laplacian system solvers as in $[51,161]$ to make the cost of interior point iterations cheap when applied to the linear program formulations analyzed by Daitch and Spielman [63]. Furthermore, as in Mądry [183] we use weights to change the central path (albeit for a slightly different purpose). We believe this further emphasizes the power of these tools as general purpose techniques for algorithm design.

| Year | Author | Running Time |
| :---: | :--- | :---: |
| 1972 | Edmonds, Karp [84] | $\tilde{O}\left(\|E\|^{2} \log (U)\right)$ |
| 1984 | Tardos [247] | $O\left(\|E\|^{4}\right)$ |
| 1984 | Orlin [217] | $\tilde{O}\left(\|E\|^{3}\right)$ |
| 1986 | Galil, Tardos [98] | $\tilde{O}\left(\|E\|\|V\|^{2}\right)$ |
| 1987 | Goldberg, Tarjan [106] | $\tilde{O}(\|E\|\|V\| \log (U))$ |
| 1988 | Orlin [215] | $\tilde{O}\left(\|E\|^{2}\right)$ |
| 2008 | Daitch, Spielman [63] | $\tilde{O}\left(\|E\|^{3 / 2} \log ^{2}(U)\right)$ |
| - | This chapter | $\tilde{O}\left(\|E\| \sqrt{\|V\|} \log ^{2}(U)\right)$ |

Figure 6-2: Minimum Cost Flow Running Times Improvements

## ■ 6.3 Simple Case

In this section, we introduce the notation of $\nu$-self-concordant barrier functions and explain how such barrier functions can be used to develop an $O(\sqrt{\nu} L)$ iterations linear programming algorithm. Then, we construct an $\tilde{O}(n)$-self-concordant barrier function which can be computed in polynomial time. This barrier function will only be used to give the intuition on how we are getting our main algorithm. The reader well versed in path following methods can likely skip to subsection 6.3 .5 and to the more curious reader we encourage them to consider some of the many wonderful expositions on this subject [209, 273, 107] for further reading.

## ■ 6.3.1 The Setup

Given a matrix, $\boldsymbol{A} \in \mathbb{R}^{m \times n}$, and vectors, $\boldsymbol{b} \in \mathbb{R}^{m}$ and $\boldsymbol{c} \in \mathbb{R}^{n}$, the goal of this section is to solve the following linear program

$$
\begin{equation*}
\min _{\boldsymbol{x} \in \mathbb{R}^{n}: \boldsymbol{A} \boldsymbol{x} \geq \boldsymbol{b}} \boldsymbol{c}^{T} \boldsymbol{x} \tag{6.3}
\end{equation*}
$$

This is called the dual of the standard form of a linear program and all linear programs can be expressed in this form. We call a vector $\boldsymbol{x} \in \mathbb{R}^{m}$ feasible if $\boldsymbol{A} \boldsymbol{x} \geq \boldsymbol{b}$, we call $\boldsymbol{c}^{T} \boldsymbol{x}$ the cost of such a vector. therefore our goal is to either compute a minimum cost feasible vector or determine that none exists.

We assume that $\boldsymbol{A}$ is full $\operatorname{rank}$, i.e. $\operatorname{rank}(\boldsymbol{A})=n$, and that $m \geq n$. Nevertheless, we still write many of our results using $\operatorname{rank}(\boldsymbol{A})$ rather than $n$ for two reasons. First, this notation makes clear that $\operatorname{rank}(\boldsymbol{A})$ is referring to the smaller of the two quantities $m$ and $n$. Second, if $\operatorname{rank}(\boldsymbol{A})<n$, then we can reduce the number of variables to $\operatorname{rank}(\boldsymbol{A})$ by a change of basis. Hence, we only need to solve linear programs in the full rank version.

## ■ 6.3.2 Central Path

Interior point methods solve (6.3) by maintaining a point $\boldsymbol{x}$ that is in the interior of the feasible region, i.e. $\boldsymbol{x} \in \Omega$ where

$$
\Omega \stackrel{\text { def }}{=}\left\{\boldsymbol{x} \in \mathbb{R}^{n}: \boldsymbol{A} \boldsymbol{x}>\boldsymbol{b}\right\} .
$$

These methods attempt to iteratively decrease the cost of $\boldsymbol{x}$ while maintaining strict feasibility. This is often done by considering some measurement of the distance to feasibility such as $\boldsymbol{s} \stackrel{\text { def }}{=} \boldsymbol{A x}-\boldsymbol{b}$, called the slacks, and creating some penalty for these distances approaching 0 . By carefully balancing penalties for small $\boldsymbol{s}(\boldsymbol{x})$ and penalties for large $\boldsymbol{c}^{T} \boldsymbol{x}$, these methods eventually compute a point close
enough to the optimum solution that it can be computed exactly.
Path following methods fix ratios between the the penalty for large $\boldsymbol{c}^{T} \boldsymbol{x}$ and the penalty for small $s$ and alternate between steps of optimizing with respect to this ratio and changing the ratio. These methods typically encode the penalties through a barrier function $\phi: \Omega \rightarrow \mathbb{R}$ such that $\phi(\boldsymbol{x}) \rightarrow \infty$ as $\boldsymbol{s}(\boldsymbol{x})_{i} \rightarrow 0$ for any $i \in[m]$ and they encode the ratio through some parameter $t>0$. Formally, they attempt to solve optimization problems of the following form for increasing values of $t$

$$
\begin{equation*}
\min _{\boldsymbol{x} \in \mathbb{R}^{m}} f_{t}(\boldsymbol{x}) \quad \text { where } \quad f_{t}(\boldsymbol{x}) \stackrel{\text { def }}{=} t \cdot \boldsymbol{c}^{T} \boldsymbol{x}+\phi(\boldsymbol{x}) \tag{6.4}
\end{equation*}
$$

Since $\phi(\boldsymbol{x}) \rightarrow \infty$ as $\boldsymbol{s}(\boldsymbol{x})_{i} \rightarrow 0$, the minimizer of $f_{t}(\boldsymbol{x})$, denoted $\boldsymbol{x}^{*}(t)$, is in $\Omega$ for all $t$. As $t$ increases the effect of the cost vector on $\boldsymbol{x}^{*}(t)$ increases and the distance from the boundary of the feasible region as measured by $\boldsymbol{s}(\boldsymbol{x})$ decreases. One can think of the points $\left\{\boldsymbol{x}^{*}(t) \mid t>0\right\}$ as a path in $\mathbb{R}^{n}$, called the central path, where $\boldsymbol{x}^{*}(t)$ approaches a solution to (6.3) as $t \rightarrow \infty$. A standard choice of barrier is the standard $\log$ barrier, $\phi(\boldsymbol{x}) \stackrel{\text { def }}{=}-\sum_{i=1}^{m} \log \left(\boldsymbol{s}(\boldsymbol{x})_{i}\right)$ and for this choice of barrier we refer to $\left\{\boldsymbol{x}^{*}(t) \mid t>0\right\}$ as the standard central path.

Path following methods typically follow the following framework:
(1) Compute Initial Point: Compute an approximation $\boldsymbol{x}^{*}(t)$ for some $t$.
(2) Follow the central path: Repeatedly increase $t$ and compute an approximation to $\boldsymbol{x}^{*}(t)$.
(3) Round to optimal solution: Use the approximation to $\boldsymbol{x}^{*}(t)$ to compute the solution to (6.3).

Usually, steps (1) and (3) are not the bottleneck of interior point methods (including our algorithm) and can be carried out by standard interior point techniques. In this chapter, we ignore the details for step (1) and (3) and refer the interested reader to the paper version of this chapter [165, 166]. In the following subsection, we provide a simple technique for performing (2) and serves as the foundation for the algorithms considered in the remainder of this section.

## - 6.3.3 Following the Path

Typically, path following methods compute an approximation to $\boldsymbol{x}^{*}(t)$ using Newton's method. While for an arbitrary current point $\boldsymbol{x} \in \Omega$ and $t>0$, the function $f_{t}(\boldsymbol{x})$ can be ill-behaved, in a region near $\boldsymbol{x}^{*}(t)$, the Hessian $\nabla^{2} f_{t}(\boldsymbol{x})$ changes fairly slowly. More precisely, if one considers the second order approximation of $f_{t}(\boldsymbol{z})$ around some point $\boldsymbol{x} \in \Omega$ "close enough" to $\boldsymbol{x}^{*}(t)$

$$
f_{t}(\boldsymbol{z}) \approx f_{t}(\boldsymbol{x})+\left\langle\nabla f_{t}(\boldsymbol{x}), \boldsymbol{z}-\boldsymbol{x}\right\rangle+\frac{1}{2}(\boldsymbol{z}-\boldsymbol{x})^{T}\left(\nabla^{2} f_{t}(\boldsymbol{x})\right)(\boldsymbol{z}-\boldsymbol{x})
$$

and applies one step of Newton's method, i.e. minimizes this quadratic approximation to compute

$$
\boldsymbol{x}^{(n e w)}:=\boldsymbol{x}-\left(\nabla^{2} f_{t}(\boldsymbol{x})\right)^{-1} \nabla f_{t}(\boldsymbol{x})
$$

for $\boldsymbol{s} \stackrel{\text { def }}{=} \boldsymbol{s}(\boldsymbol{x})$, then this procedure rapidly converges to $\boldsymbol{x}^{*}(t)$.
To quantify this, we measure centrality, i.e. how close the current point $\boldsymbol{x} \in \Omega$ is to $\boldsymbol{x}^{*}(t)$, by the size of this Newton step in the Hessian induced norm. For $\boldsymbol{x} \in \Omega$ and Newton step $\boldsymbol{h}_{t}(\boldsymbol{x}) \stackrel{\text { def }}{=}$ $\left(\nabla^{2} f_{t}(\boldsymbol{x})\right)^{-1} \nabla f_{t}(\boldsymbol{x})$, we denote centrality by $\delta_{t}(\boldsymbol{x}) \stackrel{\text { def }}{=}\left\|\boldsymbol{h}_{t}(\boldsymbol{x})\right\|_{\nabla^{2} f_{t}(\boldsymbol{x})}$.

For the standard $\log$ barrier function, $\phi(\boldsymbol{x}) \stackrel{\text { def }}{=}-\sum_{i \in[m]} \log \left(\boldsymbol{s}(\boldsymbol{x})_{i}\right)$, a standard analysis of Newton's method shows that if $\delta_{t}(\boldsymbol{x}) \leq \frac{1}{2}$, then for $\boldsymbol{x}^{(n e w)}:=\boldsymbol{x}-\boldsymbol{h}(\boldsymbol{x})$, we have $\delta_{t}\left(\boldsymbol{x}^{(n e w)}\right) \leq 2 \delta_{t}(\boldsymbol{x})^{2}$. Furthermore, if $\delta_{t}(\boldsymbol{x}) \leq \frac{1}{4}$, it is not hard to show that for $t^{\prime}=t\left(1+m^{-1 / 2} / 4\right)$ we have $\delta_{t^{\prime}}(\boldsymbol{x}) \leq 1 / 2$. Combining these facts yields that in $O(\sqrt{m})$ iterations we can double $t$ while maintaining a nearly centered $\boldsymbol{x}$,
i.e. $\delta_{t}(\boldsymbol{x})$ at most a constant. With some additional work, it can be shown that by maintaining a nearly centered $\boldsymbol{x}$ and increasing $t$ at most $O(\sqrt{m} L)$ times one can compute a solution to (6.3) exactly. Therefore, this barrier function gives an $O(\sqrt{m} L)$ iterations algorithm for solving (6.3).

### 6.3.4 Self-Concordant Barrier Functions

In a seminal result of Nesterov and Nemirovski [211], they studied the sufficient conditions (called self-concordance) on $\phi$ that allows us to follow the path efficiently. They showed that given the ability to construct a $\nu$-self-concordant barrier for a convex set, one can minimize linear functions over that convex set with a convergence rate of $O(\sqrt{\nu})$.

Definition 6.3.1. Given a open convex set $K \subset \mathbb{R}^{n}$. We call $\phi$ is a $\nu$-self-concordant barrier function for $K$ if the following conditions are satisfies

- $\phi$ is convex and thrice continuously differentiable,
- $\phi\left(\boldsymbol{x}_{i}\right) \rightarrow \infty$ for any sequence $\boldsymbol{x}_{i} \in K$ converging to boundary of $K$.
- $\left|D^{3} \phi(\boldsymbol{x})[\boldsymbol{h}, \boldsymbol{h}, \boldsymbol{h}]\right| \leq 2\left|D^{2} \phi(\boldsymbol{x})[\boldsymbol{h}, \boldsymbol{h}]\right|^{3 / 2}$ for any $\boldsymbol{x} \in K, \boldsymbol{h} \in \mathbb{R}^{n}$,
- $|D F(\boldsymbol{x})[\boldsymbol{h}]| \leq \sqrt{\nu}\left|D^{2} \phi(\boldsymbol{x})[\boldsymbol{h}, \boldsymbol{h}]\right|^{1 / 2}$ for any $\boldsymbol{x} \in K, \boldsymbol{h} \in \mathbb{R}^{n}$.

For example, the standard $\log$ barrier function is a $m$-self concordant barrier function for polytopes. Nesterov and Nemirovski further showed that, for any convex set in $\mathbb{R}^{n}$, there exists a barrier function, called the universal barrier function, with self-concordance $O(n)$ [211]. However, the universal barrier function is defined via the volume of certain polytopes, which is very difficult to compute.

### 6.3.5 A $\tilde{O}(n)$-Self-Concordant Barrier Function

In this section, we showed that for any polytope, there is an easy-to-compute barrier function with self-concordance $\tilde{O}(n)$. This barrier function is based on the following intuition on interior point methods. In high level, interior point methods replace the key difficulty of linear programming, the hard constraints, by barrier functions and solve the linear programs by a sequence of linear systems induced by barrier functions. Since these linear system corresponds to minimizing linear functions over ellipsoids and these ellipsoids comes from the second order approximation of barrier functions, interior point methods, to a certain degree, use a sequence of ellipsoids to approximate polytopes. Self-concordance can be thought as a geometric condition that relates how good those ellipsoids approximate the domain. In particular, the following lemma shows that the second order approximation of the barrier function at the minimum point indeed has a geometric implication, it approximates the domain.

Theorem 6.3.2 ([206, Thm 4.2.6]). Given a $\nu$-self-concordant barrier function $\phi$ for some convex set $\Omega \subset \mathbb{R}^{n}$. Let $x_{\phi}$ be the minimizer of $\phi$ and let the Dikin ellipse $E=\left\{x \in \mathbb{R}^{n}:\left(x-x_{\phi}\right)^{T} \nabla^{2} \phi\left(x_{\phi}\right)(x-\right.$ $\left.\left.x_{\phi}\right) \leq 1\right\}$. Then, we have that

$$
\begin{equation*}
E \subset \Omega \subset(\nu+2 \sqrt{\nu}) E \tag{6.5}
\end{equation*}
$$

This shows that in order to construct a $\tilde{O}(n)$-self-concordant barrier function, we must able to find an ellipse that gives a $\tilde{O}(n)$ approximation of the convex set $\Omega$. There are quite a few ellipsoids satisfying the condition (6.5). In contrast to other known ellipsoids that gives an approximation guarantee such as second moment, John ellipsoid is defined by a convex optimization problem and therefore can be computed in weakly polynomial time.

Lemma 6.3.3 ([128]). Given a convex $\Omega \subset \mathbb{R}^{n}$. We define the John ellipsoid $J(\Omega)$ as the largest volume ellipsoid contained inside $\Omega$. Then, we have that

$$
J(\Omega) \subset K \subset n J(\Omega)
$$

There are different ways to write down the John ellipsoid as a convex problem. Our barrier function is motivated by the following formulation, called D-optimal design.

Lemma 6.3.4 ([143]). For any polytope $\Omega=\{\boldsymbol{A} \boldsymbol{x}>\boldsymbol{b}\}$, the John ellipsoid $J(K)$ is given by $\{\boldsymbol{y} \in$ $\left.\mathbb{R}^{n}:(\boldsymbol{y}-\boldsymbol{x})^{T} \boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}(\boldsymbol{y}-\boldsymbol{x}) \leq 1\right\}$ where $(\boldsymbol{x}, \boldsymbol{w})$ is the saddle point of the following convex concave problem

$$
\begin{equation*}
\min _{\boldsymbol{x}} \max _{\sum w_{i}=n, w_{i} \geq 0} \ln \operatorname{det}\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-1} \boldsymbol{W} \boldsymbol{S}^{-1} \boldsymbol{A}\right) \tag{6.6}
\end{equation*}
$$

where $\boldsymbol{S}$ and $\boldsymbol{W}$ are diagonal $m \times m$ matrices with $\boldsymbol{S}_{i i}=a_{i}^{T} \boldsymbol{x}-b_{i}$ and $\boldsymbol{W}_{i i}=w_{i}$.
Motivated by Theorem 6.3.2, we want to find a barrier function whose minimizer is exactly the center of John ellipsoid. From (6.6), we see that the function

$$
\phi_{\infty}(\boldsymbol{x}) \stackrel{\text { def }}{=} \max _{\sum w_{i}=n, w_{i} \geq 0} \ln \operatorname{det}\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-1} \boldsymbol{W} \boldsymbol{S}^{-1} \boldsymbol{A}\right)
$$

satisfies this requirement exactly.
However, $\phi_{\infty}(\boldsymbol{x})$ is not even continuously differentiable. To see this, we define $J(\Omega, x)$ be the maximum volume ellipsoid inside $K$ and centered at $x$. We note that $\phi_{\infty}(\boldsymbol{x})=c \log (\operatorname{vol}(J(\Omega, x)))$ for some universal constant $c$. Consider the convex set $\Omega=[0,1]$. For any $x<1 / 2, J(\Omega, x)$ is the "ellipsoid" $[0,2 x]$ and in some sense this ellipsoid does not see the constraint $x \leq 1$. Therefore, the function $\phi_{\infty}(x)$ is not continuously differentiable at $1 / 2$.

To make $\phi_{\infty}(\boldsymbol{x})$ smooth, we employ the standard approach, adding a strongly concave term into the objective function $\ln \operatorname{det}\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-1} \boldsymbol{W} \boldsymbol{S}^{-1} \boldsymbol{A}\right)$. In general, if a smooth function $f(x, y)$ is strongly concave in $y$, then $\max _{y} f(x, y)$ is smooth. Since the domain of $w$ is a simplex, a natural choice for the strongly concave term is the entropy function. Therefore, we consider the following function

$$
\max _{\sum w_{i}=n, w_{i} \geq 0} \ln \operatorname{det}\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-1} \boldsymbol{W} \boldsymbol{S}^{-1} \boldsymbol{A}\right)+\rho(\boldsymbol{w})
$$

with $\rho(\boldsymbol{w})=-\sum_{i=1}^{m} w_{i} \ln w_{i}$. By some direct calculations, one can show this is indeed a $\tilde{O}(n)$ selfconcordant barrier function. In fact, there are many other choices of $\rho$ that give similar result. For example, we considered $\rho(\boldsymbol{w})=-\sum_{i=1}^{m} w_{i}^{1+1 / \log (m / n)}$ in our original paper. In the appendix, we proved the following:

Theorem 6.3.5. For any polytope $\Omega=\{\boldsymbol{A} x>b\}$, we define the barrier $p(x)$ for $\Omega$ as $\max _{w_{i} \geq 0} f(x, w)$ where

$$
\begin{equation*}
f(x, w)=\ln \operatorname{det}\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-1} \boldsymbol{W} \boldsymbol{S}^{-1} \boldsymbol{A}\right)-\frac{n}{m} \sum_{i=1}^{m} w_{i} \ln w_{i}-\frac{n}{m} \sum_{i=1}^{m} \ln s_{i} \tag{6.7}
\end{equation*}
$$

and $\boldsymbol{S}$ and $\boldsymbol{W}$ are diagonal $m \times m$ matrices with $\boldsymbol{S}_{i i}=a_{i}^{T} \boldsymbol{x}-b_{i}$ and $\boldsymbol{W}_{i i}=w_{i}$. Then, $p$ is an $O\left(n \ln ^{6}\left(\frac{m e}{n}\right)\right)$-self concordant barrier function (after rescaling). More precisely, we have

1. For any $x \in \Omega^{\circ},(\nabla p(x))^{T}\left(\nabla^{2} p(x)\right)^{-1}(\nabla p(x)) \leq O(n)$.
2. For any $x \in \Omega^{\circ}$, $h \in \mathbb{R}^{n}$, we have $\left|\nabla^{3} p(x)[h, h, h]\right| \leq O\left(\ln ^{3}\left(\frac{m e}{n}\right)\right)\left|\nabla^{2} p(x)[h, h]\right|^{3 / 2}$.

We remark that both the scaling $n / m$ and the extra term $\sum_{i=1}^{m} \ln s_{i}$ is not essential. In fact, one consider the $\ell_{p}$ Lewis ellipsoids instead and use the following barrier function

$$
\phi_{p}(x) \stackrel{\text { def }}{=}{\underset{M i n}{M}}^{\boldsymbol{M}} \log \operatorname{det} \boldsymbol{M} \text { subject to } \sum_{i=1}^{m}\left(\frac{\boldsymbol{a}_{i}^{\top} \boldsymbol{M} \boldsymbol{a}_{i}}{s_{i}^{2}(x)}\right)^{p / 2} \leq n \text { and } \boldsymbol{M} \succeq \mathbf{0} .
$$

One can show that $\phi_{\log (m)}$ is an $\widetilde{O}(n)$-self concordant barrier function. However, this barrier function is more difficult to analyze and hence we left it as an 15 pages calculus exercise.

## ■ 6.4 Difficulties

In this section, we discuss two issues arises from using the barrier function defined in the previous section and outline our approaches used in the rest of this chapter for fixing them.

## ■ 6.4.1 Avoiding Matrix Multiplication

The minimax formula involves the maximizer $\boldsymbol{w}$ in (6.7), which we call weight function $g(\boldsymbol{x})$. Computing $g(\boldsymbol{x})$ requires computing the diagonal of some projection matrices as in Vaidya and Anstreicher's work [257, 15] and is therefore not efficient enough for our purposes. Although we can compute $g(\boldsymbol{x})$ approximately up to certain multiplicative coordinate-wise error using dimension reduction techniques (Section 6.7.1), this error is still too much for path following to handle the directly as multiplicatively changing weights can hurt our measures of progress too much.

To avoid the minimax formulation, we note that the barrier function $p(\boldsymbol{x})$ has similar behaviors as $-\sum \boldsymbol{w}_{i} \ln \boldsymbol{s}_{i}(\boldsymbol{x})$ around $\boldsymbol{x}$ if $\boldsymbol{w}$ is close to the maximizer in (6.7) for that given $\boldsymbol{x}$. Therefore, rather than using the minimax formula, we maintain separate weights $\boldsymbol{w}$ and current point $\boldsymbol{x}$ and use the barrier

$$
\phi(\boldsymbol{x}, \boldsymbol{w})=-\sum_{i \in[m]} \boldsymbol{w}_{i} \ln \boldsymbol{s}_{i}(\boldsymbol{x}) .
$$

We then maintain two invariants, (1) $\boldsymbol{x}$ is centered, i.e. $\boldsymbol{x}$ close to the minimum point of $t \cdot \boldsymbol{c}^{T} \boldsymbol{x}+\phi(\boldsymbol{x}, \boldsymbol{w})$ and (2) $\boldsymbol{w}$ close to $\boldsymbol{g}(\boldsymbol{x})$ multiplicatively.

We separate the problem of maintaining these invariants into two steps. First, we show that a Newton step for $\boldsymbol{x}$ improves centrality without moving $\boldsymbol{w}$ too far away from $\boldsymbol{g}(\boldsymbol{x})$. Second, we show that given a multiplicative approximation to $\boldsymbol{g}(\boldsymbol{x})$ and bounds for how much $\boldsymbol{g}(\boldsymbol{x})$ may have changed, we can maintain the invariant that $\boldsymbol{g}(\boldsymbol{x})$ is close to $\boldsymbol{w}$ multiplicatively without moving $\boldsymbol{w}$ too much. We formulate this as a general two player game and prove that there is an efficient strategy to maintain our desired invariant. Combining these and standard techniques in path-following methods, we obtain an $\tilde{O}(\sqrt{\operatorname{rank}(\boldsymbol{A})} L)$ iterations path-following algorithm for solving (6.2) where each iterations consists of $\tilde{O}(1)$ linear system solves.

## ■ 6.4.2 Generalizing The Result

Unfortunately, this result is insufficient to produce faster algorithms for the maximum flow problem and its generalizations. Given an arbitrary minimum cost maximum flow instance there is a natural linear program that one can use to express the problem:

$$
\begin{equation*}
\underset{\substack{\boldsymbol{x} \in \mathbb{R}^{m} \\ \forall i \in[m]: \boldsymbol{A}^{T} \boldsymbol{x}=\boldsymbol{b} \\ \min _{i} \leq x_{i} \leq u_{i}}}{\boldsymbol{c}^{T} \boldsymbol{x}} \tag{6.8}
\end{equation*}
$$

where the variables $x_{i}$ denote the flow on an edge, the $l_{i}$ and $u_{i}$ denote lower and upper bounds on how much flow we can put on the edge, and $\boldsymbol{A}$ is the incidence matrix associated with the graph [63]. In this formulation, $\operatorname{rank}(\boldsymbol{A})$ is less than the number of vertices in the graph and using fast Laplacian system solvers $[243,147,146,141,162,170,223]$ we can solve linear systems involving $\boldsymbol{A}$ in time nearly linear in the number of edges in the graph. Thus, if we could solve (6.8) in time comparable to that we achieve for solving (6.2) this would immediately yield a $\widetilde{O}\left(m \sqrt{n} \log ^{O(1)}(U)\right)$ algorithm for the maximum flow problem. Unfortunately, it is not clear how to apply our previous results in this more general setting and naive attempts to write (6.8) in the form of (6.2) without increasing $\operatorname{rank}(\boldsymbol{A})$ fail.

Even more troubling, achieving a faster than $\widetilde{O}(\sqrt{m} L)$ iterations interior point method for solving general linear programs in this form would break a long-standing barrier for the convergence rate of interior point methods. To the best of the authors knowledge, there is no general purpose interior point method that achieves a convergence rate faster than the self concordance of the best barrier of the feasible region. Furthermore, using lower bounds results of Nesterov and Nemirovski [211, Proposition 2.3.6], it is not hard to see that any general barrier for (6.8) must have self-concordance $\Omega(m)$.

After all, linear programs of the form (6.2) are easier because each step of Newton steps lies in some ellipsoids that is strictly contained inside the domain $\boldsymbol{A} \boldsymbol{x} \geq \boldsymbol{b}$. However, such restriction implies that it takes at least $O(\sqrt{m})$ iterations to go from one point in the box $\left\{l_{i} \leq x_{i} \leq u_{i}\right\}$ to another point and that is too slow. Therefore, we need to explicitly bound the size of the Newton step in both the $\ell_{\infty}$ norm and a weighted $\ell_{2}$ norm. Hence, we measure the progress by the size of the Newton step in a mixed norm of the form $\|\cdot\|=\|\cdot\|_{\infty}+C_{\text {norm }}\|\cdot\|_{W}$ to keep track of these two quantities simultaneously.

Measuring of Newton step size both with respect to the mixed norm helps to explain how our method outperforms the self-concordance of the best barrier for the space. Self-concordance is based on $\ell_{2}$ analysis and the lower bounds for self-concordance are precisely the failure of the sphere to approximate a box. While ideally we would just perform optimization over the $\ell_{\infty}$ box directly, $\ell_{\infty}$ is ripe with degeneracies that makes this analysis difficult. Nevertheless, unconstrained minimization over a box is quite simple and by working with this mixed norm and choosing weights to improve the conditioning, we are taking advantage of the simplicity of minimizing $\ell_{\infty}$ over most of the domain and only paying for the $n$-self-concordance of a barrier for the smaller subspace induce by the requirement that $\boldsymbol{A}^{T} \boldsymbol{x}=\boldsymbol{b}$.

## ■ 6.5 Weighted Path Finding

In this section, we formally define the formulation of linear program we solve (Subsection 6.5.1), introduce the weighted central path we follow (Subsection 6.5.3), define key properties of the path (Subsection 6.5.4) and the weights (Subsection 6.5.5) that we will use to produce an efficient path finding scheme.

## ■ 6.5.1 The Problem

As we explained in subsection 6.4.2, the standard dual linear program is not suitable for the maximum flow application. From now on, our goal is to efficiently solve the following linear program

$$
\begin{equation*}
\underset{\substack{\boldsymbol{x} \in \mathbb{R}^{m} \\ \forall i \in[m]: \boldsymbol{A}^{T} \boldsymbol{x}=\boldsymbol{b} \\: l_{i} \leq x_{i} \leq u_{i}}}{ } \boldsymbol{c}^{T} \boldsymbol{x} \tag{6.9}
\end{equation*}
$$

where $\boldsymbol{A} \in \mathbb{R}^{m \times n}, \boldsymbol{b} \in \mathbb{R}^{n}, \boldsymbol{c} \in \mathbb{R}^{m}, l_{i} \in \mathbb{R} \cup\{-\infty\}$, and $u_{i} \in \mathbb{R} \cup\{+\infty\} .{ }^{4}$ We assume that for all $i \in[m]$ the domain of variable $x_{i}$, $\operatorname{dom}\left(x_{i}\right) \stackrel{\text { def }}{=}\left\{x: l_{i} \leq x \leq u_{i}\right\}$, is non-degenerate. In particular we assume that $\operatorname{dom}\left(x_{i}\right)$ is not the empty set, a singleton, or the entire real line, i.e. $l_{i}<u_{i}$ and either $l_{i} \neq-\infty$ or $u_{i} \neq+\infty$. Furthermore we make the standard assumptions that $\boldsymbol{A}$ has full column rank, and therefore $m \geq n$, and we assume that the interior of the polytope, $\Omega \stackrel{\text { def }}{=}\left\{\boldsymbol{x} \in \mathbb{R}^{m}: \boldsymbol{A}^{T} \boldsymbol{x}=\boldsymbol{b}, l_{i}<x_{i}<u_{i}\right\}$, is non-empty.

The linear program (6.9) is a generalization of standard form, the case where for all $i \in[m]$ we have $l_{i}=0$ and $u_{i}=+\infty$. While it is well known that all linear programs can be written in standard form, the transformations to rewrite (6.9) in standard form may increase the rank of $\boldsymbol{A}$ and therefore we solve (6.9) directly.

## ■ 6.5.2 Coordinate Barrier Functions

Rather than working directly with the different domain of the $x_{i}$ we take a slightly more general approach and for the remainder of the chapter assume that for all $i \in[m]$ we have a barrier function, $\phi_{i}: \operatorname{dom}\left(x_{i}\right) \rightarrow \mathbb{R}$, such that

$$
\lim _{x \rightarrow l_{i}} \phi_{i}(x)=\lim _{x \rightarrow u_{i}} \phi_{i}(x)=+\infty
$$

More precisely, we assume that each $\phi_{i}$ is a 1-self-concordant barrier function. Recall from definition 6.3.1, $\phi$ is a 1 -self-concordant barrier function if

$$
\begin{equation*}
\left|\phi^{\prime \prime \prime}(x)\right| \leq 2\left(\phi^{\prime \prime}(x)\right)^{3 / 2} \text { for all } x \in \operatorname{dom}(\phi) \tag{6.10}
\end{equation*}
$$

and

$$
\begin{equation*}
\left|\phi^{\prime}(x)\right| \leq \sqrt{\phi^{\prime \prime}(x)} \text { for all } x \in \operatorname{dom}(\phi) \tag{6.11}
\end{equation*}
$$

The first condition (6.10) bounds how quickly the second order approximation to the function can change and the second condition (6.11) bounds how much force the barrier can exert.

The existence of a self-concordant barrier for the domain is a standard assumption for interior point methods [211]. However, for completeness, here we show how for each possible setting of the $l_{i}$ and $u_{i}$ there is an explicit 1 -self-concordant barrier function we can use:

- Case (1): $l_{i}$ finite and $u_{i}=\infty$ : Here we use a $\log$ barrier defined as $\phi_{i}(x) \stackrel{\text { def }}{=}-\log \left(x-l_{i}\right)$. For this barrier we have

$$
\phi_{i}^{\prime}(x)=-\frac{1}{x-l_{i}} \quad, \quad \phi_{i}^{\prime \prime}(x)=\frac{1}{\left(x-l_{i}\right)^{2}} \quad, \quad \text { and } \quad \phi_{i}^{\prime \prime \prime}(x)=-\frac{2}{\left(x-l_{i}\right)^{3}}
$$

and therefore clearly $\left|\phi_{i}^{\prime \prime \prime}(x)\right|=2\left(\phi_{i}^{\prime \prime}(x)\right)^{3 / 2},\left|\phi_{i}^{\prime}(x)\right|=\sqrt{\phi_{i}^{\prime \prime}(x)}$, and $\lim _{x \rightarrow l_{i}} \phi_{i}(x)=\infty$.

- Case (2): $l_{i}=-\infty$ and $u_{i}$ finite: Here we use a $\log$ barrier defined as $\phi_{i}(x) \stackrel{\text { def }}{=}-\log \left(u_{i}-x\right)$. For this barrier we have

$$
\phi_{i}^{\prime}(x)=\frac{1}{u_{i}-x} \quad, \quad \phi_{i}^{\prime \prime}(x)=\frac{1}{\left(u_{i}-x\right)^{2}} \quad, \quad \text { and } \quad \phi_{i}^{\prime \prime \prime}(x)=-\frac{2}{\left(u_{i}-x\right)^{3}}
$$

and therefore clearly $\left|\phi_{i}^{\prime \prime \prime}(x)\right|=2\left(\phi_{i}^{\prime \prime}(x)\right)^{3 / 2},\left|\phi_{i}^{\prime}(x)\right|=\sqrt{\phi_{i}^{\prime \prime}(x)}$, and $\lim _{x \rightarrow u_{i}} \phi_{i}(x)=\infty$.

[^8]- Case (3): $l_{i}$ finite and $u_{i}$ finite: Here we use a trigonometric barrier ${ }^{5}$ defined as $\phi_{i}(x) \stackrel{\text { def }}{=}$ $-\log \cos \left(a_{i} x+b_{i}\right)$ for $a_{i}=\frac{\pi}{u_{i}-l_{i}}$ and $b_{i}=-\frac{\pi}{2} \frac{u_{i}+l_{i}}{u_{i}-l_{i}}$. Note for this choice as $x \rightarrow u_{i}$ we have $a_{i} x+b_{i} \rightarrow \frac{\pi}{2}$ and as $x \rightarrow l_{i}$ we have $a_{i} x+b_{i} \rightarrow \frac{-\pi}{2}$ and in both cases $\phi_{i}(x) \rightarrow \infty$. Furthermor $e$,

$$
\phi_{i}^{\prime}(x)=a_{i} \tan \left(a_{i} x+b_{i}\right) \quad, \quad \phi_{i}^{\prime \prime}(x)=\frac{a_{i}^{2}}{\cos ^{2}\left(a_{i} x+b_{i}\right)} \quad \text {, and } \quad \phi_{i}^{\prime \prime \prime}=\frac{2 a_{i}^{3} \sin \left(a_{i} x+b_{i}\right)}{\cos ^{3}\left(a_{i} x+b_{i}\right)} .
$$

Therefore, we have $\left|\phi_{i}^{\prime \prime \prime}(x)\right| \leq 2\left(\phi^{\prime \prime}(x)\right)^{3 / 2}$ and $\left|\phi_{i}^{\prime}(x)\right| \leq \sqrt{\phi_{i}^{\prime \prime}(x)}$.
For the remainder of this chapter we will simply assume that we have a 1 self-concordant barrier $\phi_{i}$ for each of the $\operatorname{dom}\left(\phi_{i}\right)$ and will not use any more structure about the barriers.

While there is much theory regarding properties of self-concordant barrier functions we will primarily use two common properties about self-concordant barriers functions. The first property, Lemma 6.5.1, shows that the Hessian of the barrier cannot change to quickly, and the second property, Lemma 6.5.2 we use to reason about how the force exerted by the barrier changes over the domain.

Lemma 6.5.1 ([206, Theorem 4.1.6]). Suppose $\phi$ is a 1-self-concordant barrier function. For all $s \in \operatorname{dom}(\phi)$ if $r \stackrel{\text { def }}{=} \sqrt{\phi^{\prime \prime}(s)}|s-t|<1$ then $t \in \operatorname{dom}(\phi)$ and

$$
(1-r) \sqrt{\phi^{\prime \prime}(s)} \leq \sqrt{\phi^{\prime \prime}(t)} \leq \frac{\sqrt{\phi^{\prime \prime}(s)}}{1-r} .
$$

Lemma 6.5.2 ([206, Theorem 4.2.4]). Suppose $\phi$ is a 1-self-concordant barrier function. For all $x, y \in \operatorname{dom}(\phi)$, we have

$$
\phi^{\prime}(x) \cdot(y-x) \leq 1
$$

## ■ 6.5.3 The Weighted Central Path

Our linear programming algorithm maintains a feasible point $\boldsymbol{x} \in \Omega$, weights $\boldsymbol{w} \in \mathbb{R}_{>0}^{m}$, and minimizes the following penalized objective function

$$
\begin{equation*}
\min _{\boldsymbol{A}^{T} \boldsymbol{x}=\boldsymbol{b}} f_{t}(\boldsymbol{x}, \boldsymbol{w}) \quad \text { where } \quad f_{t}(\boldsymbol{x}, \boldsymbol{w}) \stackrel{\text { def }}{=} t \cdot \boldsymbol{c}^{T} \boldsymbol{x}+\sum_{i \in[m]} w_{i} \phi_{i}\left(\boldsymbol{x}_{i}\right) \tag{6.12}
\end{equation*}
$$

for increasing $t$ and small $\boldsymbol{w}$. For every fixed set of weights, $\boldsymbol{w} \in \mathbb{R}_{>0}^{m}$ the set of points $\boldsymbol{x}_{\boldsymbol{w}}(t)=$ $\arg \min _{x \in \Omega} f_{t}(\boldsymbol{x}, \boldsymbol{w})$ for $t \in[0, \infty)$ forms a path through the interior of the polytope that we call the weighted central path. We call $\boldsymbol{x}_{\boldsymbol{w}}(0)$ a weighted center of the polytope and note that $\lim _{t \rightarrow \infty} \boldsymbol{x}_{\boldsymbol{w}}(0)$ is a solution to the linear program.

Our weighted path finding algorithm follows a simple iterative scheme. We assume we have a feasible point $\{\boldsymbol{x}, \boldsymbol{w}\} \in\left\{\Omega \times \mathbb{R}_{>0}^{m}\right\}$ and a weight function $\boldsymbol{g}(\boldsymbol{x}): \Omega \rightarrow \mathbb{R}_{>0}^{m}$, such that for any point $\boldsymbol{x} \in \mathbb{R}_{>0}^{m}$ the function $\boldsymbol{g}(\boldsymbol{x})$ returns a good set of weights that suggest a possibly better weighted path. Although the weight function we use is motivated by the barrier function introduced in Theorem 6.3.5, we do not assume anything on $\boldsymbol{g}$ in this section. Our algorithm then repeats the following.

1. If $\boldsymbol{x}$ close to $\arg \min _{\boldsymbol{y} \in \Omega} f_{t}(\boldsymbol{y}, \boldsymbol{w})$, then increase $t$.
2. Otherwise, use projected Newton step to update $\boldsymbol{x}$ and move $\boldsymbol{w}$ closer to $\boldsymbol{g}(\boldsymbol{x})$.
3. Repeat.
[^9]In the remainder of this section we present how we measure both the quality of a current feasible point $\{\boldsymbol{x}, \boldsymbol{w}\} \in\left\{\Omega \times \mathbb{R}_{>0}^{m}\right\}$ and the quality of the weight function. In Section 6.5.4 we derive and present both how we measure how close $\{\boldsymbol{x}, \boldsymbol{w}\}$ is to the weighted central path and the step we take to improve this centrality. Then in Section 6.5 .5 we present how we measure the quality of a weight function, i.e. how good the weighted paths it finds are.

## ■ 6.5.4 Measuring Centrality.

Here we explain how we measure the distance from $\boldsymbol{x}$ to the minimum of $f_{t}(\boldsymbol{x}, \boldsymbol{w})$ for fixed $\boldsymbol{w}$. This distance is a measure of how close $\boldsymbol{x}$ is to the weighted central path and we refer to it as the centrality of $\boldsymbol{x}$, denote $\delta_{t}(\boldsymbol{x}, \boldsymbol{w})$.

To motivate our centrality measure, we first compute a projected Newton step for $\boldsymbol{x}$. For all $\boldsymbol{x} \in \Omega$, we define $\boldsymbol{\phi}(\boldsymbol{x}) \in \mathbb{R}^{m}$ by $\phi(\boldsymbol{x})_{i}=\phi_{i}\left(\boldsymbol{x}_{i}\right)$ for $i \in[m]$. We define $\boldsymbol{\phi}^{\prime}(\boldsymbol{x}), \phi^{\prime \prime}(\boldsymbol{x})$, and $\phi^{\prime \prime \prime}(\boldsymbol{x})$ similarly and let $\boldsymbol{\Phi}, \boldsymbol{\Phi}^{\prime}, \boldsymbol{\Phi}^{\prime \prime}, \boldsymbol{\Phi}^{\prime \prime \prime}$ denote the diagonal matrices corresponding to these matrices. Using this, we have ${ }^{6}$

$$
\nabla_{x} f_{t}(\boldsymbol{x}, \boldsymbol{w})=t \cdot \boldsymbol{c}+\boldsymbol{w} \boldsymbol{\phi}^{\prime}(\boldsymbol{x}) \quad \text { and } \quad \nabla_{x x} f_{t}(\boldsymbol{x}, \boldsymbol{w})=\boldsymbol{W} \boldsymbol{\Phi}^{\prime \prime}(\boldsymbol{x}) .
$$

Therefore, a Newton step for $\boldsymbol{x}$ is given by

$$
\begin{align*}
\boldsymbol{h}_{t}(\boldsymbol{x}, \boldsymbol{w}) & =-\left(\boldsymbol{W} \boldsymbol{\Phi}^{\prime \prime}(\boldsymbol{x})\right)^{-1 / 2} \boldsymbol{P}_{\boldsymbol{A}^{T}\left(\boldsymbol{W} \boldsymbol{\Phi}^{\prime \prime}(\boldsymbol{x})\right)^{-1 / 2}}\left(\boldsymbol{W} \boldsymbol{\Phi}^{\prime \prime}(\boldsymbol{x})\right)^{-1 / 2} \nabla_{x} f_{t}(\boldsymbol{x}, \boldsymbol{w}) \\
& =-\boldsymbol{\Phi}^{\prime \prime}(\boldsymbol{x})^{-1 / 2} \boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}} \boldsymbol{W}^{-1} \boldsymbol{\Phi}^{\prime \prime}(\boldsymbol{x})^{-1 / 2} \nabla_{x} f_{t}(\boldsymbol{x}, \boldsymbol{w}) \tag{6.13}
\end{align*}
$$

where $\boldsymbol{P}_{\boldsymbol{A}^{T}\left(\boldsymbol{W} \boldsymbol{\Phi}^{\prime \prime}(\boldsymbol{x})\right)^{-1 / 2}}$ is the orthogonal projection onto the kernel of $\boldsymbol{A}^{T}\left(\boldsymbol{W} \boldsymbol{\Phi}^{\prime \prime}(\boldsymbol{x})\right)^{-1 / 2}$ and $\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}}$ is the orthogonal projection onto the kernel of $\boldsymbol{A}^{T}\left(\boldsymbol{\Phi}^{\prime \prime}(\boldsymbol{x})\right)^{-1 / 2}$ with respect to the norm $\|\cdot\|_{\boldsymbol{W}}$, i.e.

$$
\begin{equation*}
\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}} \stackrel{\text { def }}{=} \boldsymbol{I}-\boldsymbol{W}^{-1} \boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W}^{-1} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}_{x}^{T} \quad \text { for } \quad \boldsymbol{A}_{x} \stackrel{\text { def }}{=} \boldsymbol{\Phi}^{\prime \prime}(\boldsymbol{x})^{-1 / 2} \boldsymbol{A} . \tag{6.14}
\end{equation*}
$$

As with standard convergence analysis of Newton's method, we wish to keep the Newton step size in the Hessian norm, i.e. $\left\|\boldsymbol{h}_{t}(\boldsymbol{x}, \boldsymbol{w})\right\|_{\boldsymbol{W}^{\prime \prime}(\boldsymbol{x})}=\left\|\sqrt{\boldsymbol{\phi}^{\prime \prime}(\boldsymbol{x})} \boldsymbol{h}_{t}(\boldsymbol{x}, \boldsymbol{w})\right\|_{\boldsymbol{W}}$, small and the multiplicative change in the Hessian, $\left\|\sqrt{\phi^{\prime \prime}(\boldsymbol{x})} \boldsymbol{h}_{t}(\boldsymbol{x}, \boldsymbol{w})\right\|_{\infty}$, small (See Lemma 6.5.2). While in the unweighted case we can bound the multiplicative change by the change in the hessian norm (since $\|\cdot\|_{\infty} \leq\|\cdot\|_{2}$ ), here we would like to use small weights and this comparison would be insufficient.

To track both these quantities simultaneously, we define the mixed norm for all $\boldsymbol{y} \in \mathbb{R}^{m}$ by

$$
\begin{equation*}
\|\boldsymbol{y}\|_{\boldsymbol{w}+\infty} \stackrel{\text { def }}{=}\|\boldsymbol{y}\|_{\infty}+C_{\text {norm }}\|\boldsymbol{y}\|_{\boldsymbol{W}} \tag{6.15}
\end{equation*}
$$

for some $C_{\text {norm }}>0$ that we define later. Note that $\|\cdot\|_{\boldsymbol{w}+\infty}$ is indeed a norm for $\boldsymbol{w} \in \mathbb{R}_{>0}^{m}$ as in this case both $\|\cdot\|_{\infty}$ and $\|\cdot\|_{\boldsymbol{W}}$ are norms. However, rather than measuring centrality by the quantity $\left\|\sqrt{\phi^{\prime \prime}(\boldsymbol{x})} \boldsymbol{h}_{t}(\boldsymbol{x}, \boldsymbol{w})\right\|_{\boldsymbol{w}+\infty}=\| \boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}}\left(\frac{\nabla_{x} f_{t}(\boldsymbol{x}, \boldsymbol{w})}{\left.\boldsymbol{w} \sqrt{\overline{\phi^{\prime \prime}}}\right) \|_{\boldsymbol{w}+\infty} \text {, we instead find it more convenient to use the }}\right.$ following idealized form

$$
\delta_{t}(\boldsymbol{x}, \boldsymbol{w}) \stackrel{\text { def }}{=} \min _{\eta \in \mathbb{R}^{n}}\left\|\frac{\nabla_{x} f_{t}(\boldsymbol{x}, \boldsymbol{w})-\boldsymbol{A} \boldsymbol{\eta}}{\boldsymbol{w} \sqrt{\phi^{\prime \prime}(\boldsymbol{x})}}\right\|_{\boldsymbol{w}+\infty}
$$

We justify this definition by showing these two quantities differ by at most a multiplicative factor of $\left\|\boldsymbol{P}_{x, w}\right\|_{w+\infty}$ as follows

$$
\begin{equation*}
\delta_{t}(\boldsymbol{x}, \boldsymbol{w}) \leq\left\|\sqrt{\boldsymbol{\phi}^{\prime \prime}(\boldsymbol{x})} \boldsymbol{h}_{t}(\boldsymbol{x}, \boldsymbol{w})\right\|_{\boldsymbol{w}+\infty} \leq\left\|\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}}\right\|_{\boldsymbol{w}+\infty} \cdot \delta_{t}(\boldsymbol{x}, \boldsymbol{w}) . \tag{6.16}
\end{equation*}
$$

[^10]This a direct consequence of the more general Lemma 6.11 .1 that we prove in the appendix.
We summarize this section with the following definition.
Definition 6.5.3 (Centrality Measure). For $\{\boldsymbol{x}, \boldsymbol{w}\} \in\left\{\Omega \times \mathbb{R}_{>0}^{m}\right\}$ and $t \geq 0$, we let $\boldsymbol{h}_{t}(\boldsymbol{x}, \boldsymbol{w})$ denote the projected newton step for $\boldsymbol{x}$ on the penalized objective $f_{t}$ given by

$$
\boldsymbol{h}_{t}(\boldsymbol{x}, \boldsymbol{w}) \stackrel{\text { def }}{=}-\frac{1}{\sqrt{\vec{\phi}^{\prime \prime}(\boldsymbol{x})}} \boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}}\left(\frac{\nabla_{x} f_{t}(\boldsymbol{x}, \boldsymbol{w})}{\boldsymbol{w} \sqrt{\vec{\phi}^{\prime \prime}(\boldsymbol{x})}}\right)
$$

where $\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}}$ is the orthogonal projection onto the kernel of $\boldsymbol{A}^{T}\left(\boldsymbol{\Phi}^{\prime \prime}\right)^{-1 / 2}$ with respect to the norm $\|\cdot\|_{\boldsymbol{W}}$ (see (6.15)). We measure the centrality of $\{\boldsymbol{x}, \boldsymbol{w}\}$ by

$$
\begin{equation*}
\delta_{t}(\boldsymbol{x}, \boldsymbol{w}) \stackrel{\text { def }}{=} \min _{\boldsymbol{\eta} \in \mathbb{R}^{n}}\left\|\frac{\nabla_{x} f_{t}(\boldsymbol{x}, \boldsymbol{w})-\boldsymbol{A} \boldsymbol{\eta}}{\boldsymbol{w} \sqrt{\boldsymbol{\phi}^{\prime \prime}(\boldsymbol{x})}}\right\|_{\boldsymbol{w}+\infty} \tag{6.17}
\end{equation*}
$$

where for all $\boldsymbol{y} \in \mathbb{R}^{m}$ we let $\|\boldsymbol{y}\|_{\boldsymbol{w}+\infty} \stackrel{\text { def }}{=}\|\boldsymbol{y}\|_{\infty}+C_{\text {norm }}\|\boldsymbol{y}\|_{\boldsymbol{W}}$ for some $C_{\text {norm }}>0$ we define later.

### 6.5.5 The Weight Function

With the Newton step and centrality conditions defined, the specification of our algorithm becomes more clear. Our algorithm is as follows

1. If $\delta_{t}(\boldsymbol{x}, \boldsymbol{w})$ is small, then increase $t$.
2. Set $\boldsymbol{x}^{(\text {new })} \leftarrow \boldsymbol{x}+\boldsymbol{h}_{t}(\boldsymbol{x}, \boldsymbol{w})$ and move $\boldsymbol{w}^{(\text {new })}$ towards $\boldsymbol{g}\left(\boldsymbol{x}^{(\text {new })}\right)$.
3. Repeat.

To prove this algorithm converges, we need to show what happens to $\delta_{t}(\boldsymbol{x}, \boldsymbol{w})$ when we change $t$, $\boldsymbol{x}$ and $\boldsymbol{w}$. At the heart of this chapter is understanding what conditions we need to impose on the weight function $\boldsymbol{g}(\boldsymbol{x}): \Omega \rightarrow \mathbb{R}_{>0}^{m}$ so that we can bound this change in $\delta_{t}(\boldsymbol{x}, \boldsymbol{w})$ and hence achieve fast converge rates. In Lemma 6.6 .1 we show that the effect of changing $t$ on $\delta_{t}$ is bounded by $C_{\text {norm }}$ and $\|\boldsymbol{g}(\boldsymbol{x})\|_{1}$, in Lemma 6.6 .2 we show that the effect that a Newton Step on $\boldsymbol{x}$ has on $\delta_{t}$ is bounded by $\left\|\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{g}(\boldsymbol{x})}\right\|_{\boldsymbol{g}(\boldsymbol{x})+\infty}$, and in Lemma 6.6.3 and 6.6.4 we show the change of $\boldsymbol{w}$ as $\boldsymbol{g}(\boldsymbol{x})$ changes is bounded by $\left\|\boldsymbol{G}(\boldsymbol{x})^{-1} \boldsymbol{G}^{\prime}(\boldsymbol{x})\left(\boldsymbol{\Phi}^{\prime \prime}(\boldsymbol{x})\right)^{-1 / 2}\right\|_{\boldsymbol{g}(\boldsymbol{x})+\infty}$.

Hence for the remainder of the chapter we assume we have a weight function $\boldsymbol{g}(\boldsymbol{x}): \Omega \rightarrow \mathbb{R}_{>0}^{m}$ and make the following assumptions regarding our weight function. In Section 6.7 we prove that such weight function exists.

Definition 6.5.4 (Weight Function). A weight function is a differentiable function from $\boldsymbol{g}: \Omega \rightarrow \mathbb{R}_{>0}^{m}$ such that for constants $c_{1}(\boldsymbol{g}), c_{\gamma}(\boldsymbol{g})$, and $c_{\delta}(\boldsymbol{g})$, we have the following for all $\boldsymbol{x} \in \Omega$ :

- Size : The size $c_{1}(\boldsymbol{g})=\|\boldsymbol{g}(\boldsymbol{x})\|_{1}$.
- Slack Sensitivity: The slack sensitivity $c_{\gamma}(\boldsymbol{g})$ satisfies $1 \leq c_{\gamma}(\boldsymbol{g}) \leq \frac{5}{4}$ and $\left\|\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}}\right\|_{\boldsymbol{w}+\infty} \leq c_{\gamma}(\boldsymbol{g})$ for any $\boldsymbol{w}$ such that $\frac{4}{5} \boldsymbol{g}(\boldsymbol{x}) \leq \boldsymbol{w} \leq \frac{5}{4} \boldsymbol{g}(\boldsymbol{x})$.
- Step Consistency : The step consistency $c_{\delta}(\boldsymbol{g})$ satisfies $c_{\delta}(\boldsymbol{g}) \cdot c_{\gamma}(\boldsymbol{g})<1$ and

$$
\left\|\boldsymbol{G}(\boldsymbol{x})^{-1} \boldsymbol{G}^{\prime}(\boldsymbol{x})\left(\boldsymbol{\Phi}^{\prime \prime}(\boldsymbol{x})\right)^{-1 / 2}\right\|_{\boldsymbol{g}(\boldsymbol{x})+\infty} \leq c_{\delta} \leq 1
$$

- Uniformity : The weight function satisfies $\|\boldsymbol{g}(\boldsymbol{x})\|_{\infty} \leq 2$.


## ■ 6.5.6 Notation

Here we introduce various notation that we will use throughout the chapter (excluding Section 6.13). This should be used primarily for reference as we reintroduce notation as needed later in the chapter. For a matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ and vector $\boldsymbol{x} \in \mathbb{R}_{>0}^{m}$ we define the following matrices and vectors

- Projection matrix $\boldsymbol{P}_{\boldsymbol{A}}(\boldsymbol{x}) \in \mathbb{R}^{m \times m}: \boldsymbol{P}_{\boldsymbol{A}}(\boldsymbol{x}) \stackrel{\text { def }}{=} \boldsymbol{X}^{1 / 2} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{X} \boldsymbol{A}\right)^{-1} \boldsymbol{A}^{T} \boldsymbol{X}^{1 / 2}$.
- Leverage scores $\sigma_{\boldsymbol{A}}(\boldsymbol{x}) \in \mathbb{R}^{m}: \boldsymbol{\sigma}_{\boldsymbol{A}}(\boldsymbol{x}) \stackrel{\text { def }}{=} \boldsymbol{\operatorname { d i a g }}\left(\boldsymbol{P}_{\boldsymbol{A}}(\boldsymbol{x})\right)$.
- Leverage matrix $\boldsymbol{\Sigma}_{\boldsymbol{A}}(\boldsymbol{x}) \in \mathbb{R}^{m \times m}: \boldsymbol{\Sigma}_{\boldsymbol{A}}(\boldsymbol{x}) \stackrel{\text { def }}{=} \boldsymbol{\operatorname { D i a g }}\left(\boldsymbol{P}_{\boldsymbol{A}}(\boldsymbol{x})\right)$.
- Projection Laplacian $\boldsymbol{\Lambda}_{\boldsymbol{A}}(\boldsymbol{x}) \in \mathbb{R}^{m \times m}: \boldsymbol{\Lambda}_{\boldsymbol{A}}(\boldsymbol{x}) \stackrel{\text { def }}{=} \boldsymbol{\Sigma}_{\boldsymbol{A}}(\boldsymbol{x})-\boldsymbol{P}_{\boldsymbol{A}}(\boldsymbol{x})^{(2)}$.


## ■ 6.6 Progressing Along Weighted Paths

In this section, we provide the main lemmas we need for an $\tilde{O}(\sqrt{\operatorname{rank}(\boldsymbol{A})} \log (U / \varepsilon))$ iterations weighted path following algorithm for (6.9) assuming a weight function satisfying Definition 6.5.5. In Section 6.6.1, 6.6.2, and 6.6 .3 we show how centrality, $\delta_{t}(\boldsymbol{x}, \boldsymbol{w})$, is affected by changing $t, \boldsymbol{x} \in \Omega$, and $\boldsymbol{w} \in \mathbb{R}_{>0}^{m}$ respectively. In Section 6.6 .5 we then show how to use these Lemmas to improve centrality using approximate computations of the weight function, $\boldsymbol{g}: \Omega \rightarrow \mathbb{R}_{>0}^{m}$.

## ■ 6.6.1 Changing $t$

Here we bound how much centrality increases as we increase $t$. We show that this rate of increase is governed by $C_{\text {norm }}$ and $\|\boldsymbol{w}\|_{1}$.

Lemma 6.6.1. For all $\{\boldsymbol{x}, \boldsymbol{w}\} \in\left\{\Omega \times \mathbb{R}_{>0}^{m}\right\}, t>0$ and $\alpha \geq 0$, we have

$$
\delta_{(1+\alpha) t}(\boldsymbol{x}, \boldsymbol{w}) \leq(1+\alpha) \delta_{t}(\boldsymbol{x}, \boldsymbol{w})+\alpha\left(1+C_{n o r m} \sqrt{\|\boldsymbol{w}\|_{1}}\right) .
$$

Proof. Let $\boldsymbol{\eta}_{t} \in \mathbb{R}^{n}$ be such that

$$
\delta_{t}(\boldsymbol{x}, \boldsymbol{w})=\left\|\frac{\nabla_{x} f_{t}(\boldsymbol{x}, \boldsymbol{w})+\boldsymbol{A} \boldsymbol{\eta}_{t}}{\boldsymbol{w} \sqrt{\phi^{\prime \prime}(\boldsymbol{x})}}\right\|_{\boldsymbol{w}+\infty}=\left\|\frac{t \cdot \boldsymbol{c}+\boldsymbol{w} \boldsymbol{\phi}^{\prime}(\boldsymbol{x})+\boldsymbol{A} \boldsymbol{\eta}_{t}}{\boldsymbol{w} \sqrt{\phi^{\prime \prime}(\boldsymbol{x})}}\right\|_{\boldsymbol{w}+\infty}
$$

Applying this to the definition of $\delta_{(1+\alpha) t}$ and using that $\|\cdot\|_{w+\infty}$ is a norm then yields

$$
\begin{aligned}
\delta_{(1+\alpha) t}(\boldsymbol{x}, \boldsymbol{w}) & =\min _{\eta \in \mathbb{R}^{n}}\left\|\frac{(1+\alpha) t \cdot \boldsymbol{c}+\boldsymbol{w} \boldsymbol{\phi}^{\prime}(\boldsymbol{x})+\boldsymbol{A} \boldsymbol{\eta}}{\boldsymbol{w} \sqrt{\phi^{\prime \prime}(\boldsymbol{x})}}\right\|_{\boldsymbol{w}+\infty} \\
& \leq\left\|\frac{(1+\alpha) t \cdot \boldsymbol{c}+\boldsymbol{w} \phi^{\prime}(\boldsymbol{x})+\boldsymbol{A}(1+\alpha) \boldsymbol{\eta}_{t}}{\boldsymbol{w} \sqrt{\phi^{\prime \prime}(\boldsymbol{x})}}\right\|_{\boldsymbol{w}+\infty} \\
& \leq(1+\alpha)\left\|\frac{t \cdot \boldsymbol{c}+\boldsymbol{w} \boldsymbol{\phi}^{\prime}(\boldsymbol{x})+\boldsymbol{A} \boldsymbol{\eta}_{t}}{\boldsymbol{w} \sqrt{\phi^{\prime \prime}(\boldsymbol{x})}}\right\|_{\boldsymbol{w}+\infty}+\alpha\left\|\frac{\boldsymbol{w} \boldsymbol{\phi}^{\prime}(\boldsymbol{x})}{\boldsymbol{w} \sqrt{\boldsymbol{\phi}^{\prime \prime}(\boldsymbol{x})}}\right\|_{\boldsymbol{w}+\infty} \\
& =(1+\alpha) \delta_{t}(\boldsymbol{x}, \boldsymbol{w})+\alpha\left(\left\|\frac{\boldsymbol{\phi}^{\prime}(\boldsymbol{x})}{\sqrt{\boldsymbol{\phi}^{\prime \prime}(\boldsymbol{x})}}\right\|_{\infty}+C_{\mathrm{norm}}\left\|\frac{\phi^{\prime}(\boldsymbol{x})}{\sqrt{\boldsymbol{\phi}^{\prime \prime}(\boldsymbol{x})}}\right\|_{\boldsymbol{W}}\right)
\end{aligned}
$$

Using that $\left|\phi_{i}^{\prime}(\boldsymbol{x})\right| \leq \sqrt{\phi_{i}^{\prime \prime}(\boldsymbol{x})}$ for all $i \in[m]$ and $\boldsymbol{x} \in \mathbb{R}^{m}$ by Equation (6.11) yields the result.

## ■ 6.6.2 Changing $x$

Here we analyze the effect of a Newton step on $\boldsymbol{x}$ on centrality. We show for sufficiently central $\{\boldsymbol{x}, \boldsymbol{w}\} \in\left\{\Omega \times \mathbb{R}_{>0}^{m}\right\}$ and $\boldsymbol{w}$ sufficiently close to $\boldsymbol{g}(\boldsymbol{x})$ Newton steps converge quadratically.
Lemma 6.6.2. Let $\left\{\boldsymbol{x}_{0}, \boldsymbol{w}\right\} \in\left\{\Omega \times \mathbb{R}_{>0}^{m}\right\}$ such that $\delta_{t}\left(\boldsymbol{x}_{0}, \boldsymbol{w}\right) \leq \frac{1}{10}$ and $\frac{4}{5} \boldsymbol{g}(\boldsymbol{x}) \leq \boldsymbol{w} \leq \frac{5}{4} \boldsymbol{g}(\boldsymbol{x})$ and consider a Newton step $\boldsymbol{x}_{1}=\boldsymbol{x}_{0}+\boldsymbol{h}_{t}(\boldsymbol{x}, \boldsymbol{w})$. Then, $\delta_{t}\left(\boldsymbol{x}_{1}, \boldsymbol{w}\right) \leq 4\left(\delta_{t}\left(\boldsymbol{x}_{0}, \boldsymbol{w}\right)\right)^{2}$.

Proof. Let $\boldsymbol{\phi}_{0} \stackrel{\text { def }}{=} \boldsymbol{\phi}\left(\boldsymbol{x}_{0}\right)$ and let $\boldsymbol{\phi}_{1} \stackrel{\text { def }}{=} \boldsymbol{\phi}\left(\boldsymbol{x}_{1}\right)$. By the definition of $\boldsymbol{h}_{t}\left(\boldsymbol{x}_{0}, \boldsymbol{w}\right)$ and the formula of $\boldsymbol{P}_{\boldsymbol{x}_{0}, \boldsymbol{w}}$ we know that there is some $\boldsymbol{\eta}_{0} \in \mathbb{R}^{n}$ such that

$$
-\sqrt{\overrightarrow{\phi_{0}^{\prime \prime}}} \boldsymbol{h}_{t}\left(\boldsymbol{x}_{0}, \boldsymbol{w}\right)=\frac{t \cdot \boldsymbol{c}+\boldsymbol{w} \phi_{0}^{\prime}-\boldsymbol{A} \boldsymbol{\eta}_{0}}{\boldsymbol{w} \sqrt{\boldsymbol{\phi}_{0}^{\prime \prime}}} .
$$

Therefore, $\boldsymbol{A} \boldsymbol{\eta}_{0}=\boldsymbol{c}+\boldsymbol{w} \phi_{0}^{\prime}+\boldsymbol{w} \phi_{0}^{\prime \prime} h_{t}\left(\boldsymbol{x}_{0}, \boldsymbol{w}\right)$. Recalling the definition of $\delta_{t}$ this implies that

$$
\begin{aligned}
\delta_{t}\left(\boldsymbol{x}_{1}, \boldsymbol{w}\right) & =\min _{\eta \in \mathbb{R}^{n}}\left\|\frac{t \cdot \boldsymbol{c}+\boldsymbol{w} \vec{\phi}_{1}^{\prime}-\boldsymbol{A} \boldsymbol{\eta}}{\boldsymbol{w} \sqrt{\vec{\phi}_{1}^{\prime \prime}}}\right\|_{\boldsymbol{w}+\infty} \leq\left\|\frac{t \cdot \boldsymbol{c}+\boldsymbol{w} \vec{\phi}_{1}^{\prime}-\boldsymbol{A} \boldsymbol{\eta}_{0}}{\boldsymbol{w} \sqrt{\vec{\phi}_{1}^{\prime \prime}}}\right\|_{\boldsymbol{w}+\infty} \\
& \leq\left\|\frac{\boldsymbol{w}\left(\overrightarrow{\phi_{1}^{\prime}}-\boldsymbol{\phi}_{0}^{\prime}\right)-\boldsymbol{w} \boldsymbol{\phi}_{0}^{\prime \prime} \boldsymbol{h}_{t}\left(\boldsymbol{x}_{0}, \boldsymbol{w}\right)}{\boldsymbol{w} \sqrt{\vec{\phi}_{1}^{\prime \prime}}}\right\|_{\boldsymbol{w}+\infty}=\left\|\frac{\left(\overrightarrow{\phi_{1}^{\prime}}-\boldsymbol{\phi}_{0}^{\prime}\right)-\boldsymbol{\phi}_{0}^{\prime \prime} \boldsymbol{h}_{t}\left(\boldsymbol{x}_{0}, \boldsymbol{w}\right)}{\sqrt{\vec{\phi}_{1}^{\prime \prime}}}\right\|_{\boldsymbol{w}+\boldsymbol{\infty}}
\end{aligned}
$$

By the mean value theorem, we have $\vec{\phi}_{1}^{\prime}-\boldsymbol{\phi}_{0}^{\prime}=\overrightarrow{\phi^{\prime \prime}}(\vec{\theta}) \boldsymbol{h}_{t}\left(\boldsymbol{x}_{0}, \boldsymbol{w}\right)$ for some $\vec{\theta}$ between $\boldsymbol{x}_{0}$ and $\boldsymbol{x}_{1}$ coordinate-wise. Hence,

$$
\begin{aligned}
\delta_{t}\left(\boldsymbol{x}_{1}, \boldsymbol{w}\right) & \leq\left\|\frac{\vec{\phi}^{\prime \prime}(\vec{\theta}) \boldsymbol{h}_{t}\left(\boldsymbol{x}_{0}, \boldsymbol{w}\right)-\boldsymbol{\phi}_{0}^{\prime \prime} \boldsymbol{h}_{t}\left(\boldsymbol{x}_{0}, \boldsymbol{w}\right)}{\sqrt{\vec{\phi}_{1}^{\prime \prime}}}\right\|_{\boldsymbol{w}+\infty}=\left\|\frac{\left(\overrightarrow{\phi^{\prime \prime}}(\vec{\theta})-\overrightarrow{\phi_{0}^{\prime \prime}}\right)}{\sqrt{\vec{\phi}_{1}^{\prime \prime}} \sqrt{\phi_{0}^{\prime \prime}}}\left(\sqrt{\phi_{0}^{\prime \prime}} \boldsymbol{h}_{t}\left(\boldsymbol{x}_{0}, \boldsymbol{w}\right)\right)\right\|_{\boldsymbol{w}+\infty} \\
& \leq\left\|\frac{\vec{\phi}^{\prime \prime}(\vec{\theta})-\overrightarrow{\phi_{0}^{\prime \prime}}}{\sqrt{\overrightarrow{\phi_{1}^{\prime \prime}}} \sqrt{\phi_{0}^{\prime \prime}}}\right\|_{\infty} \cdot\left\|\sqrt{\phi_{0}^{\prime \prime}} \boldsymbol{h}_{t}\left(\boldsymbol{x}_{0}, \boldsymbol{w}\right)\right\|_{\boldsymbol{w}+\infty} .
\end{aligned}
$$

To bound the first term, we use Lemma 6.5.1 as follows

$$
\begin{aligned}
\left\|\frac{\left(\overrightarrow{\phi^{\prime \prime}}(\vec{\theta})-\overrightarrow{\phi_{0}^{\prime \prime}}\right)}{\sqrt{\overrightarrow{\phi_{1}^{\prime \prime}}} \sqrt{\vec{\phi}_{0}^{\prime \prime}}}\right\|_{\infty} & \leq\left\|\frac{\phi^{\prime \prime}(\vec{\theta})}{\phi_{0}^{\prime \prime}}-1\right\|_{\infty} \cdot\left\|\frac{\sqrt{\vec{\phi}_{0}^{\prime \prime}}}{\sqrt{\overrightarrow{\phi_{1}^{\prime \prime}}}}\right\|_{\infty} \\
& \leq \|\left(1-\left\|\sqrt{\vec{\phi}_{0}^{\prime \prime}} \boldsymbol{h}_{t}\left(\boldsymbol{x}_{0}, \boldsymbol{w}\right)\right\|_{\infty}\right)^{-2}-1 \mid \cdot\left(1-\left\|\sqrt{\vec{\phi}_{0}^{\prime \prime}} \boldsymbol{h}_{t}\left(\boldsymbol{x}_{0}, \boldsymbol{w}\right)\right\|_{\infty}\right)^{-1} .
\end{aligned}
$$

Using (6.16), i.e. Lemma 6.11.1, the bound $c_{\gamma} \leq 2$, and the assumption on $\delta_{t}\left(\boldsymbol{x}_{0}, \boldsymbol{w}\right)$, we have

$$
\left\|\sqrt{\vec{\phi}_{0}^{\prime \prime}} \boldsymbol{h}_{t}\left(\boldsymbol{x}_{0}, \boldsymbol{w}\right)\right\|_{\infty} \leq\left\|\sqrt{\vec{\phi}_{0}^{\prime \prime}} \boldsymbol{h}_{t}\left(\boldsymbol{x}_{0}, \boldsymbol{w}\right)\right\|_{\boldsymbol{w}+\infty} \leq c_{\gamma} \cdot \delta_{t}\left(\boldsymbol{x}_{0}, \boldsymbol{w}\right) \leq \frac{1}{5} .
$$

Using $\left((1-t)^{-2}-1\right) \cdot(1-t)^{-1} \leq 4 t$ for $t \leq 1 / 5$, we have

$$
\left\|\frac{\left(\vec{\phi}^{\prime \prime}(\vec{\theta})-\vec{\phi}_{0}^{\prime \prime}\right)}{\sqrt{\vec{\phi}_{1}^{\prime \prime}} \sqrt{\vec{\phi}_{0}^{\prime \prime}}}\right\|_{\infty} \leq 4\left\|\sqrt{\vec{\phi}_{0}^{\prime \prime}} h_{t}\left(\boldsymbol{x}_{0}, \boldsymbol{w}\right)\right\|_{\infty}
$$

Combining the above formulas yields that $\delta_{t}\left(\boldsymbol{x}_{1}, \boldsymbol{w}\right) \leq 4\left(\delta_{t}\left(\boldsymbol{x}_{0}, \boldsymbol{w}\right)\right)^{2}$ as desired.

## ■ 6.6.3 Changing $w$

In the previous subsection we used the assumption that the weights, $\boldsymbol{w}$, were multiplicatively close to the output of the weight function, $\boldsymbol{g}(\boldsymbol{x})$, for the current point $\boldsymbol{x} \in \Omega$. In order to maintain this invariant when we change $\boldsymbol{x}$ we will need to change $\boldsymbol{w}$ to move it closer to $\boldsymbol{g}(\boldsymbol{x})$. Here we bound how much $\boldsymbol{g}(\boldsymbol{x})$ can move as we move $\boldsymbol{x}$ (Lemma 6.6.3) and we bound how much changing $\boldsymbol{w}$ can hurt centrality (Lemma 6.6.4). Together these lemmas will allow us to show that we can keep $\boldsymbol{w}$ close to $\boldsymbol{g}(\boldsymbol{x})$ while still improving centrality (Section 6.6.5).

Lemma 6.6.3. For all $t \in[0,1]$, let $\boldsymbol{x}_{t} \stackrel{\text { def }}{=} \boldsymbol{x}_{0}+t \boldsymbol{\Delta}_{x}$ for $\boldsymbol{\Delta}_{x} \in \mathbb{R}^{m}, \boldsymbol{x}_{t} \in \Omega, \boldsymbol{g}_{t}=\boldsymbol{g}\left(\boldsymbol{x}_{t}\right)$ and $\varepsilon=\left\|\sqrt{\overrightarrow{\phi_{0}^{\prime \prime}}} \Delta_{x}\right\|_{g_{0}+\infty} \leq 0.1$. Then

$$
\left\|\log \left(\boldsymbol{g}_{1}\right)-\log \left(\boldsymbol{g}_{0}\right)\right\|_{\boldsymbol{g}_{0}+\infty} \leq c_{\delta} \varepsilon(1+4 \varepsilon) \leq 0.2
$$

and for all $s, t \in[0,1]$ and for all $\boldsymbol{y} \in \mathbb{R}^{m}$ we have

$$
\begin{equation*}
\|\boldsymbol{y}\|_{g_{s}+\infty} \leq(1+2 \varepsilon)\|\boldsymbol{y}\|_{g_{t}+\infty} . \tag{6.18}
\end{equation*}
$$

Proof. Let $\boldsymbol{q}:[0,1] \rightarrow \mathbb{R}^{m}$ be given by $\boldsymbol{q}(t) \stackrel{\text { def }}{=} \log \left(\boldsymbol{g}_{t}\right)$ for all $t \in[0,1]$. Then, we have

$$
\boldsymbol{q}^{\prime}(t)=\boldsymbol{G}_{t}^{-1} \boldsymbol{G}_{t}^{\prime} \boldsymbol{\Delta}_{x} .
$$

Let $Q(t) \stackrel{\text { def }}{=}\|\vec{q}(t)-\vec{q}(0)\|_{g_{0}+\infty}$. Using Jensen's inequality we have that for all $u \in[0,1]$,

$$
Q(u) \leq \bar{Q}(u) \stackrel{\text { def }}{=} \int_{0}^{u}\left\|\boldsymbol{G}_{t}^{-1} \boldsymbol{G}_{t}^{\prime}\left(\vec{\phi}_{t}^{\prime \prime}\right)^{-1 / 2}\right\|_{\boldsymbol{g}_{0}+\infty}\left\|\sqrt{\vec{\phi}_{t}^{\prime \prime}} \Delta_{x}\right\|_{\boldsymbol{g}_{0}+\infty} d t .
$$

Using Lemma 6.5 .1 and $\varepsilon \leq \frac{1}{10}$, we have for all $t \in[0,1]$,

$$
\begin{aligned}
\left\|\sqrt{\vec{\phi}_{t}^{\prime \prime}} \Delta_{x}\right\|_{g_{0}+\infty} & \leq\left\|\sqrt{\vec{\phi}_{t}^{\prime \prime}} \sqrt{\overrightarrow{\phi_{0}^{\prime \prime}}}\right\|_{\infty}\left\|\sqrt{\vec{\phi}_{0}^{\prime \prime}} \Delta_{x}\right\|_{g_{0}+\infty} \\
& \leq\left(1-\left\|\sqrt{\vec{\phi}_{0}^{\prime \prime}} \Delta_{x}\right\|_{\infty}\right)^{-1}\left\|\sqrt{\vec{\phi}_{0}^{\prime \prime}} \Delta_{x}\right\|_{g_{0}+\infty} \leq \frac{\varepsilon}{1-\varepsilon} .
\end{aligned}
$$

Thus, we have

$$
\begin{equation*}
\bar{Q}(u) \leq \frac{\varepsilon}{1-\varepsilon} \int_{0}^{u}\left\|\boldsymbol{G}_{t}^{-1} \boldsymbol{G}_{t}^{\prime}\left(\vec{\phi}_{t}^{\prime \prime}\right)^{-1 / 2}\right\|_{\boldsymbol{g}_{0}+\infty} d t \tag{6.19}
\end{equation*}
$$

Note that $\bar{Q}$ is monotonically increasing. Let $\theta=\sup _{u \in[0,1]}\left\{\bar{Q}(u) \leq c_{\delta} \varepsilon(1+4 \varepsilon)\right\}$. Since $\bar{Q}(\theta) \leq \frac{1}{2}$, we know that for all $s, t \in[0, \theta]$, we have

$$
\left\|\frac{\boldsymbol{g}\left(\boldsymbol{x}_{s}\right)-\boldsymbol{g}\left(\boldsymbol{x}_{t}\right)}{\boldsymbol{g}\left(\boldsymbol{x}_{t}\right)}\right\|_{\infty} \leq\|\vec{q}(s)-\vec{q}(t)\|_{\infty}+\|\vec{q}(s)-\vec{q}(t)\|_{\infty}^{2}
$$

and therefore

$$
\left\|\boldsymbol{g}_{s} / \boldsymbol{g}_{t}\right\|_{\infty} \leq\left(1+\|\vec{q}(s)-\vec{q}(t)\|_{\infty}+\|\vec{q}(s)-\vec{q}(t)\|_{\infty}^{2}\right)^{2} \leq\left(1+c_{\delta} \varepsilon(1+4 \varepsilon)\right)^{2}
$$

Consequently,

$$
\|\boldsymbol{y}\|_{g_{s}+\infty} \leq\left(1+c_{\delta} \varepsilon(1+4 \varepsilon)\right)\|\boldsymbol{y}\|_{g_{t}+\infty} \leq(1+2 \varepsilon)\|\boldsymbol{y}\|_{g_{t}+\infty} .
$$

Using (6.19), we have for all $u \in[0, \theta]$,

$$
\begin{aligned}
Q(u) \leq \bar{Q}(u) & \leq \frac{\varepsilon}{1-\varepsilon} \int_{0}^{u}\left\|\boldsymbol{G}_{t}^{-1} \boldsymbol{G}_{t}^{\prime}\left(\vec{\phi}_{t}^{\prime \prime}\right)^{-1 / 2}\right\|_{\boldsymbol{g}_{0}+\infty} d t \\
& \leq \frac{\varepsilon}{1-\varepsilon} \int_{0}^{u}(1+2 \varepsilon)\left\|\boldsymbol{G}_{t}^{-1} \boldsymbol{G}_{t}^{\prime}\left(\vec{\phi}_{t}^{\prime \prime}\right)^{-1 / 2}\right\|_{\boldsymbol{g}_{t}+\infty} d t \\
& \leq \frac{\varepsilon}{1-\varepsilon}(1+2 \varepsilon) c_{\delta} \theta \leq c_{\delta} \varepsilon(1+4 \varepsilon)
\end{aligned}
$$

Consequently, we have that $\theta=1$ and we have the desired result with $Q(1) \leq c_{\delta} \varepsilon(1+4 \varepsilon)<\frac{1}{5}$.
Lemma 6.6.4. Let $\boldsymbol{v}, \boldsymbol{w} \in \mathbb{R}_{>0}^{m}$ such that $\varepsilon=\|\log (\boldsymbol{w})-\log (\boldsymbol{v})\|_{\boldsymbol{w}+\infty} \leq 0.1$. Then for $\boldsymbol{x} \in \Omega$ we have

$$
\delta_{t}(\boldsymbol{x}, \boldsymbol{v}) \leq(1+4 \varepsilon)\left(\delta_{t}(\boldsymbol{x}, \boldsymbol{w})+\varepsilon\right) .
$$

Proof. Let $\boldsymbol{\eta}_{w}$ be such that

$$
\begin{equation*}
\delta_{t}(\boldsymbol{x}, \boldsymbol{w})=\left\|\frac{\boldsymbol{c}+\boldsymbol{w} \vec{\phi}^{\prime}(\boldsymbol{x})-\boldsymbol{A} \boldsymbol{\eta}_{w}}{\boldsymbol{w} \sqrt{\vec{\phi}^{\prime \prime}(\boldsymbol{x})}}\right\|_{\boldsymbol{w}+\infty} \tag{6.20}
\end{equation*}
$$

Furthermore, the assumption shows that $(1+\varepsilon)^{-2} \boldsymbol{w}_{i} \leq \boldsymbol{v}_{i} \leq(1+\varepsilon)^{2} \boldsymbol{w}_{i}$ for all $i$. Using these, we bound the energy with the new weights as follows

$$
\begin{aligned}
\delta_{t}(\boldsymbol{x}, \boldsymbol{v}) & =\min _{\eta}\left\|\frac{\boldsymbol{c}+\boldsymbol{v} \vec{\phi}^{\prime}(\boldsymbol{x})-\boldsymbol{A} \boldsymbol{\eta}}{\boldsymbol{v} \sqrt{\vec{\phi}^{\prime \prime}(\boldsymbol{x})}}\right\|_{\boldsymbol{v}+\infty} \leq\left\|\frac{\boldsymbol{c}+\boldsymbol{v} \vec{\phi}^{\prime}(\boldsymbol{x})-\boldsymbol{A} \boldsymbol{\eta}_{w}}{\boldsymbol{v} \sqrt{\vec{\phi}^{\prime \prime}(\boldsymbol{x})}}\right\|_{\boldsymbol{v}+\infty} \\
& \leq(1+\varepsilon)\left\|\frac{\boldsymbol{c}+\boldsymbol{v} \vec{\phi}^{\prime}(\boldsymbol{x})-\boldsymbol{A} \boldsymbol{\eta}_{w}}{\boldsymbol{v} \sqrt{\vec{\phi}^{\prime \prime}(\boldsymbol{x})}}\right\|_{\boldsymbol{w}+\infty} \\
& \leq(1+\varepsilon) \cdot\left(\left\|\frac{\boldsymbol{c}+\boldsymbol{w} \vec{\phi}^{\prime}(\boldsymbol{x})-\boldsymbol{A} \boldsymbol{\eta}_{w}}{\boldsymbol{v} \sqrt{\vec{\phi}^{\prime \prime}(\boldsymbol{x})}}\right\|_{\boldsymbol{w}+\infty}+\left\|\frac{(\boldsymbol{v}-\boldsymbol{w}) \overrightarrow{\phi^{\prime}}(\boldsymbol{x})}{\boldsymbol{v} \sqrt{\vec{\phi}^{\prime \prime}(\boldsymbol{x})}}\right\|_{\boldsymbol{w}+\infty}\right) \\
& \leq(1+\varepsilon)^{3} \delta_{t}(\boldsymbol{x}, \boldsymbol{w})+(1+\varepsilon) \cdot\left\|\frac{\phi^{\prime}(\boldsymbol{x})}{\sqrt{\boldsymbol{\phi}^{\prime \prime}(\boldsymbol{x})}}\right\|_{\infty} \cdot\left\|\frac{(\boldsymbol{v}-\boldsymbol{w})}{\boldsymbol{v}}\right\|_{\boldsymbol{w}+\infty}
\end{aligned}
$$

Using that $\left|\phi_{i}^{\prime}(\boldsymbol{x})\right| \leq \sqrt{\phi_{i}^{\prime \prime}(\boldsymbol{x})}$ for all $i \in[m]$ by Equation (6.11), we have that

$$
\begin{aligned}
\delta_{t}(\boldsymbol{x}, \boldsymbol{v}) & \leq(1+\varepsilon)^{3} \delta_{t}(\boldsymbol{x}, \boldsymbol{w})+(1+\varepsilon)^{2} \varepsilon \\
& \leq(1+4 \varepsilon)\left(\delta_{t}(\boldsymbol{x}, \boldsymbol{w})+\varepsilon\right)
\end{aligned}
$$

## ■ 6.6.4 The Chasing 0 Game

In the previous subsection, we saw how much the weight function, $\boldsymbol{g}(\boldsymbol{x})$, can change after a Newton step on $\boldsymbol{x}$ and we bounded how much we can move the weights without affecting centrality too much.

Here we study how to correctly move the weights even when we cannot compute the weight function exactly. To effectively use multiplicative approximations to the weight function, we need to smooth out changes to the weights by using some slowly changing approximation to the weight function. In this subsection, we present the smoothing problem in a general form that we call the chasing 0 game and we provide an effective strategy for playing this game. In next subsection, we show how to use this strategy to produce a efficient weighted path following scheme.

The chasing 0 game is as follows. There is a player, an adversary, and a point $\boldsymbol{x} \in \mathbb{R}^{m}$. The goal of the player is to keep the point close to $\mathbf{0}$ in $\ell_{\infty}$ norm and the goal of the adversary tries to move $\boldsymbol{x}$ away from $\mathbf{0} \in \mathbb{R}^{m}$. The game proceeds for an infinite number of iterations where in each iteration the adversary moves the current point $\boldsymbol{x}^{(k)} \in \mathbb{R}^{m}$ to some new point $\boldsymbol{y}^{(k)} \in \mathbb{R}^{m}$ and the player needs to respond. The player does not know $\boldsymbol{x}^{(k)}, \boldsymbol{y}^{(k)}$, or the move of the adversary. All the player knows is that the adversary moved the point within some convex set $U^{(k)}$ and the player knows some $\boldsymbol{z}^{(k)} \in \mathbb{R}^{n}$ that is close to $\boldsymbol{y}^{(k)}$ in $\ell_{\infty}$ norm. With this information the player is allowed to move the point a little more than the adversary. Formally, the player is allowed to set the next point to $\boldsymbol{x}^{(k+1)} \in \mathbb{R}^{m}$ such that $\boldsymbol{\Delta}^{(k)} \stackrel{\text { def }}{=} \boldsymbol{x}^{(k+1)}-\boldsymbol{y}^{(k)} \in(1+\varepsilon) U$ for some fixed $\varepsilon>0$.

The question we would like to address is, how close the player can keep $\boldsymbol{x}^{(k+1)}$ to $\mathbf{0}$ in $\ell_{\infty}$ norm? In particular, we would like an efficient strategy for computing $\boldsymbol{\Delta}^{(k)}$ such that $\left\|\boldsymbol{x}^{(k)}\right\|_{\infty}$ is bounded for all $k \geq 0$.

```
Algorithm 7: Chasing 0 Game
    Given \(R>0, \varepsilon>0, \boldsymbol{x}^{(0)} \in \mathbb{R}^{m}\).
    for \(k=1,2, \cdots\) do
        The adversary announces symmetric convex set \(U^{(k)} \subseteq \mathbb{R}^{n}\) and \(\boldsymbol{u}^{(k)} \in U^{(k)}\).
        The adversary sets \(\boldsymbol{y}^{(k)}:=\boldsymbol{x}^{(k)}+\boldsymbol{u}^{(k)}\).
        The adversary announces \(\boldsymbol{z}^{(k)}\) such that \(\left\|\boldsymbol{z}^{(k)}-\boldsymbol{y}^{(k)}\right\|_{\infty} \leq R\).
        The player chooses \(\boldsymbol{\Delta}^{(k)} \in(1+\varepsilon) U^{(k)}\).
        The player sets \(\boldsymbol{x}^{(k+1)}=\boldsymbol{y}^{(k)}+\boldsymbol{\Delta}^{(k)}\).
    end
```

We show that assuming that the $U^{(k)}$ are sufficiently bounded then there is strategy that the player can follow to ensure that that $\left\|\boldsymbol{x}^{(k)}\right\|_{\infty}$ is never too large. Our strategy simply consists of taking "gradient steps" using the following potential function.

Definition 6.6.5. For any $\mu \geq 0$ let $p_{\mu}: \mathbb{R} \rightarrow \mathbb{R}$ and $\Phi_{\mu}: \mathbb{R}^{m} \rightarrow \mathbb{R}$ be given by

$$
\forall x \in \mathbb{R} \quad: \quad p_{\mu}(x) \stackrel{\text { def }}{=} e^{\mu x}+e^{-\mu x} \quad \text { and } \quad \Phi_{\mu}(x) \stackrel{\text { def }}{=} \sum_{i \in[m]} p_{\mu}\left(x_{i}\right) .
$$

In other words, for all $k$ we simply set $\boldsymbol{\Delta}^{(k)}$ to be the vector in $(1+\varepsilon) U^{(k)}$ that best minimizes the potential function of the observed position, i.e. $\Phi_{\mu}\left(\boldsymbol{z}^{(k)}\right)$ for an appropriate choice of $\mu$. In the following theorem we show that this suffices to keep $\Phi_{\mu}\left(\boldsymbol{x}^{(k)}\right)$ small and that small $\Phi_{\mu}\left(\boldsymbol{x}^{(k)}\right)$ implies small $\left\|\boldsymbol{x}^{(k)}\right\|_{\infty}$ and hence has the desired properties.

Theorem 6.6.6. Suppose that each $U^{(k)}$ is a symmetric convex set that contains an $\ell_{\infty}$ ball of radius $r_{k}$ and is contained in a $\ell_{\infty}$ ball of radius $R_{k} \leq R$. Let $0<\varepsilon<\frac{1}{5}$ and consider the strategy

$$
\boldsymbol{\Delta}^{(k)}=(1+\varepsilon) \underset{\Delta \in U^{(k)}}{\arg \min }\left\langle\nabla \Phi_{\mu}\left(\boldsymbol{z}^{(k)}\right), \boldsymbol{\Delta}\right\rangle \quad \text { where } \quad \mu=\frac{\varepsilon}{12 R} .
$$

Let $\tau \stackrel{\text { def }}{=} \max _{k} \frac{R_{k}}{r_{k}}$ and suppose $\Phi_{\mu}\left(\boldsymbol{x}^{(0)}\right) \leq \frac{12 m \tau}{\varepsilon}$ (or more specifically $\left\|\boldsymbol{x}^{(0)}\right\|_{\infty} \leq \frac{12 R}{\varepsilon} \log \left(\frac{6 \tau}{\varepsilon}\right)$ ) then

$$
\forall k \geq 0 \quad: \quad \Phi_{\mu}\left(\boldsymbol{x}^{(k+1)}\right) \leq\left(1-\frac{\varepsilon^{2} r_{k}}{24 R}\right) \Phi_{\mu}\left(\boldsymbol{x}^{(k)}\right)+\varepsilon m \frac{R_{k}}{2 R} \leq \frac{12 m \tau}{\varepsilon}
$$

In particular, we have $\left\|\boldsymbol{x}^{(k)}\right\|_{\infty} \leq \frac{12 R}{\varepsilon} \log \left(\frac{12 m \tau}{\varepsilon}\right)$.
To prove Theorem 6.6 .6 we first provide the following lemma regarding properties of the potential function $\Phi_{\mu}$.

Lemma 6.6.7. For all $\boldsymbol{x} \in \mathbb{R}^{m}$, we have

$$
\begin{equation*}
e^{\mu\|\boldsymbol{x}\|_{\infty}} \leq \Phi_{\mu}(\boldsymbol{x}) \leq 2 m e^{\mu\|\boldsymbol{x}\|_{\infty}} \quad \text { and } \quad \mu \Phi_{\mu}(\boldsymbol{x})-2 \mu m \leq\left\|\nabla \Phi_{\mu}(\boldsymbol{x})\right\|_{1} \tag{6.21}
\end{equation*}
$$

Furthermore, for any symmetric convex set $U \subseteq \mathbb{R}^{m}$ and any $\boldsymbol{x} \in \mathbb{R}^{m}$, let $\boldsymbol{x}^{b} \stackrel{\text { def }}{=} \arg \max _{\boldsymbol{y} \in U}\langle\boldsymbol{x}, \boldsymbol{y}\rangle$ and $\|\boldsymbol{x}\|_{U} \stackrel{\text { def }}{=} \max _{\boldsymbol{y} \in U}\langle\boldsymbol{x}, \boldsymbol{y}\rangle$. Then for all $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{m}$ with $\|\boldsymbol{x}-\boldsymbol{y}\|_{\infty} \leq \delta \leq \frac{1}{5 \mu}$ we have

$$
\begin{equation*}
e^{-\mu \delta}\left\|\nabla \Phi_{\mu}(\boldsymbol{y})\right\|_{U}-\mu\left\|\nabla \Phi_{\mu}(\boldsymbol{y})^{b}\right\|_{1} \leq\left\langle\nabla \Phi_{\mu}(\boldsymbol{x}), \nabla \Phi_{\mu}(\boldsymbol{y})^{b}\right\rangle \leq e^{\mu \delta}\left\|\nabla \Phi_{\mu}(\boldsymbol{y})\right\|_{U}+\mu e^{\mu \delta}\left\|\nabla \Phi_{\mu}(\boldsymbol{y})^{b}\right\|_{1} \tag{6.22}
\end{equation*}
$$

If additionally $U$ is contained in a $\ell_{\infty}$ ball of radius $R$ then

$$
\begin{equation*}
e^{-\mu \delta}\left\|\nabla \Phi_{\mu}(\boldsymbol{y})\right\|_{U}-\mu m R \leq\left\|\nabla \Phi_{\mu}(\boldsymbol{x})\right\|_{U} \leq e^{\mu \delta}\left\|\nabla \Phi_{\mu}(\boldsymbol{y})\right\|_{U}+\mu e^{\mu \delta} m R \tag{6.23}
\end{equation*}
$$

Proof. First we note that for all $x \in \mathbb{R}$ we have

$$
e^{\mu|x|} \leq p_{\mu}(x) \leq 2 e^{\mu|x|} \quad \text { and } \quad p_{\mu}^{\prime}(x)=\mu \operatorname{sign}(x)\left(e^{\mu|x|}-e^{-\mu|x|}\right)
$$

and therefore we have (6.21).
Next let $x, y \in \mathbb{R}$ such that $|x-y| \leq \delta$. Note that $\left|p_{\mu}^{\prime}(x)\right|=p_{\mu}^{\prime}(|x|)=\mu\left(e^{\mu|x|}-e^{-\mu|x|}\right)$ and since $|x-y| \leq \delta$ we have that $|x|=|y|+z$ for some $z \in(-\delta, \delta)$. Using that $p^{\prime}(|x|)$ is monotonic in $|x|$ we then have

$$
\begin{align*}
\left|p_{\mu}^{\prime}(x)\right| & =p_{\mu}^{\prime}(|x|)=p_{\mu}^{\prime}(|y|+z) \leq p_{\mu}^{\prime}(|y|+\delta) \\
& =\mu\left(e^{\mu|y|+\mu \delta}-e^{-\mu|y|-\mu \delta}\right)=e^{\mu \delta} p^{\prime}(|y|)+\mu\left(e^{\mu \delta-\mu|y|}-e^{-\mu|y|-\mu \delta}\right) \\
& \leq e^{\mu \delta}\left|p^{\prime}(y)\right|+\mu e^{\mu \delta} \tag{6.24}
\end{align*}
$$

By symmetry (i.e. replacing $x$ and $y$ ) this implies that

$$
\begin{equation*}
\left|p_{\mu}^{\prime}(x)\right| \geq e^{-\mu \delta}\left|p^{\prime}(y)\right|-\mu \tag{6.25}
\end{equation*}
$$

Since $U$ is symmetric this implies that for all $i \in[m]$ we have $\operatorname{sign}\left(\nabla \Phi_{\mu}(\boldsymbol{y})^{b}\right)_{i}=\nabla \Phi_{\mu}(\boldsymbol{y})_{i}=\operatorname{sign}\left(y_{i}\right)$ . Therefore, if for all $i \in[n]$ we have $\operatorname{sign}\left(x_{i}\right)=\operatorname{sign}\left(y_{i}\right)$, by (6.24), we see that

$$
\begin{aligned}
\left\langle\nabla \Phi_{\mu}(\boldsymbol{x}), \nabla \Phi_{\mu}(\boldsymbol{y})^{b}\right\rangle & =\sum_{i} p_{\mu}^{\prime}\left(x_{i}\right) \nabla \Phi_{\mu}(\boldsymbol{y})_{i}^{b} \\
& \leq \sum_{i}\left(e^{\mu \delta} p_{\mu}^{\prime}\left(y_{i}\right)+\mu e^{\mu \delta}\right) \nabla \Phi_{\mu}(\boldsymbol{y})_{i}^{b} \\
& \leq e^{\mu \delta}\left\langle\nabla \Phi_{\mu}(\boldsymbol{y}), \nabla \Phi_{\mu}(\boldsymbol{y})^{b}\right\rangle+\mu e^{\mu \delta}\left\|\nabla \Phi_{\mu}(\boldsymbol{y})^{b}\right\|_{1} \\
& =e^{\mu \delta}\left\|\nabla \Phi_{\mu}(\boldsymbol{y})\right\|_{U}+\mu e^{\mu \delta}\left\|\nabla \Phi_{\mu}(\boldsymbol{y})^{b}\right\|_{1}
\end{aligned}
$$

Similarly, using (6.25), we have $e^{-\mu \delta}\left\|\nabla \Phi_{\mu}(\boldsymbol{y})\right\|_{U}-\mu\left\|\nabla \Phi_{\mu}(\boldsymbol{y})^{b}\right\|_{1} \leq\left\langle\nabla \Phi_{\mu}(\boldsymbol{x}), \nabla \Phi_{\mu}(\boldsymbol{y})^{b}\right\rangle$ and hence
(6.22) holds. On the other hand if $\operatorname{sign}\left(x_{i}\right) \neq \operatorname{sign}\left(y_{i}\right)$ then we know that $\left|x_{i}\right| \leq \delta$ and consequently $\left|p_{\mu}^{\prime}\left(x_{i}\right)\right| \leq \mu\left(e^{\mu \delta}-e^{-\mu \delta}\right) \leq \frac{\mu}{2}$ since $\delta \leq \frac{1}{5 \mu}$. Thus, we have

$$
e^{-\mu \delta}\left|p_{\mu}^{\prime}\left(y_{i}\right)\right|-\mu \leq-\frac{\mu}{2} \leq \operatorname{sign}\left(y_{i}\right) p_{\mu}^{\prime}\left(x_{i}\right) \leq 0 \leq e^{\mu \delta}\left|p_{\mu}^{\prime}\left(y_{i}\right)\right|+\mu e^{\mu \delta}
$$

Taking inner product on both sides with $\nabla \Phi_{\mu}(\boldsymbol{y})_{i}^{b}$ and using definition of $\|\cdot\|_{U}$ and ${ }^{b}$, we get (6.22). Thus, (6.22) holds in general.

Finally we note that since $U$ is contained in a $\ell_{\infty}$ ball of radius $R$, we have $\left\|\boldsymbol{y}^{b}\right\|_{1} \leq m R$ for all $\boldsymbol{y}$. Using this fact, (6.22), and the definition of $\|\cdot\|_{U}$, we obtain

$$
e^{-\mu \delta}\left\|\nabla \Phi_{\mu}(\boldsymbol{y})\right\|_{U}-\mu m R \leq\left\langle\nabla \Phi_{\mu}(\boldsymbol{x}), \nabla \Phi_{\mu}(\boldsymbol{y})^{b}\right\rangle \leq\left\|\nabla \Phi_{\mu}(\boldsymbol{x})\right\|_{U}
$$

where the last line comes from the fact $\nabla \Phi_{\mu}(\boldsymbol{y})^{b} \in U$ and the definition of $\|\cdot\|_{U}$. By symmetry (6.23) follows.

Using Lemma 6.6 .7 we prove Theorem 6.6.6.

Theorem 6.6.6. Let $\|\boldsymbol{x}\|_{U^{(k)}}=\max _{\boldsymbol{y} \in U^{(k)}}\langle\boldsymbol{x}, \boldsymbol{y}\rangle$ and $\boldsymbol{x}^{\text {b }}{ }^{(k)}=\arg \max _{\boldsymbol{y} \in U^{(k)}}\langle\boldsymbol{x}, \boldsymbol{y}\rangle$. Since $U^{(k)}$ is symmetric, we know that $\boldsymbol{\Delta}^{(k)}=-(1+\varepsilon)\left(\nabla \Phi_{\mu}\left(\boldsymbol{z}^{(k)}\right)\right)^{b(k)}$ and therefore by applying the mean value theorem twice we have that

$$
\begin{aligned}
\Phi_{\mu}\left(\boldsymbol{x}^{(k+1)}\right) & =\Phi_{\mu}\left(\boldsymbol{y}^{(k)}\right)+\left\langle\nabla \Phi_{\mu}(\boldsymbol{z}), \boldsymbol{x}^{(k+1)}-\boldsymbol{y}^{(k)}\right\rangle \\
& =\Phi_{\mu}\left(\boldsymbol{x}^{(k)}\right)+\left\langle\nabla \Phi_{\mu}(\boldsymbol{y}), \boldsymbol{y}^{(k)}-\boldsymbol{x}^{(k)}\right\rangle+\left\langle\nabla \Phi_{\mu}(\boldsymbol{z}), \boldsymbol{x}^{(k+1)}-\boldsymbol{y}^{(k)}\right\rangle
\end{aligned}
$$

for some $\boldsymbol{y}$ between $\boldsymbol{y}^{(k)}$ and $\boldsymbol{x}^{(k)}$ and some $\boldsymbol{z}$ between $\boldsymbol{x}^{(k+1)}$ and $\boldsymbol{y}^{(k)}$. Now, using that $\boldsymbol{y}^{(k)}-\boldsymbol{x}^{(k)} \in$ $U^{(k)}$ and that $\boldsymbol{x}^{(k+1)}-\boldsymbol{y}^{(k)}=\boldsymbol{\Delta}^{(k)}$ we have

$$
\begin{equation*}
\Phi_{\mu}\left(\boldsymbol{x}^{(k+1)}\right) \leq \Phi_{\mu}\left(\boldsymbol{x}^{(k)}\right)+\left\|\nabla \Phi_{\mu}(\boldsymbol{y})\right\|_{U^{(k)}}-(1+\varepsilon)\left\langle\nabla \Phi_{\mu}(\boldsymbol{z}),\left(\nabla \Phi_{\mu}\left(\boldsymbol{z}^{(k)}\right)\right)^{b(k)}\right\rangle \tag{6.26}
\end{equation*}
$$

Since $U^{k}$ is contained within the $\ell_{\infty}$ ball of radius $R_{k}$ Lemma 6.6 .7 shows that

$$
\begin{equation*}
\left\|\nabla \Phi_{\mu}(\boldsymbol{y})\right\|_{U^{(k)}} \leq e^{\mu R_{k}}\left\|\nabla \Phi_{\mu}\left(\boldsymbol{x}^{(k)}\right)\right\|_{U^{(k)}}+m \mu R_{k} e^{\mu R_{k}} \tag{6.27}
\end{equation*}
$$

Furthermore, since $\varepsilon<\frac{1}{5}$ and $R_{k} \leq R$, by triangle inequality we have $\left\|\boldsymbol{z}-\boldsymbol{z}^{(k)}\right\|_{\infty} \leq(1+\varepsilon) R_{k}+R \leq 3 R$ and $\left\|\boldsymbol{z}^{(k)}-\boldsymbol{x}^{(k)}\right\|_{\infty} \leq 2 R$. Therefore, applying Lemma 6.6.7 twice yields that

$$
\begin{align*}
\left\langle\nabla \Phi_{\mu}(\boldsymbol{z}), \nabla \Phi_{\mu}\left(\boldsymbol{z}^{(k)}\right)^{b}(k)\right. & \geq e^{-3 \mu R}\left\|\nabla \Phi_{\mu}\left(\boldsymbol{z}^{(k)}\right)\right\|_{U^{(k)}}-\mu m R_{k} \\
& \geq e^{-5 \mu R}\left\|\nabla \Phi_{\mu}\left(\boldsymbol{x}^{(k)}\right)\right\|_{U^{(k)}}-2 \mu m R_{k} . \tag{6.28}
\end{align*}
$$

Combining (6.26), (6.27), and (6.28) then yields that

$$
\Phi_{\mu}\left(\boldsymbol{x}^{(k+1)}\right) \leq \Phi_{\mu}\left(\boldsymbol{x}^{(k)}\right)-\left((1+\varepsilon) e^{-5 \mu R}-e^{\mu R}\right)\left\|\nabla \Phi_{\mu}\left(\boldsymbol{x}^{(k)}\right)\right\|_{U^{(k)}}+m \mu R_{k} e^{\mu R}+2(1+\varepsilon) m \mu R_{k} .
$$

Since we chose $\mu=\frac{\varepsilon}{12 R}$, we have $1+\varepsilon \leq \frac{\varepsilon}{2}+(1+6 \mu R) \leq \frac{\varepsilon}{2} e^{5 \mu R}+e^{6 \mu R}$. Hence, we have $(1+\varepsilon) e^{-5 \mu R}-$ $e^{\mu R} \leq \frac{\varepsilon}{2}$. Also, since $0<\varepsilon<\frac{1}{5}$ we have

$$
m \mu R_{k} e^{\mu R}+2(1+\varepsilon) m \mu R_{k} \leq\left(e^{\mu R}+2(1+\varepsilon)\right) m \mu R_{k} \leq \varepsilon m \frac{7 R_{k}}{24 R}
$$

Thus, we have

$$
\Phi_{\mu}\left(\boldsymbol{x}^{(k+1)}\right) \leq \Phi_{\mu}\left(\boldsymbol{x}^{(k)}\right)-\frac{\varepsilon}{2}\left\|\nabla \Phi_{\mu}\left(\boldsymbol{x}^{(k)}\right)\right\|_{U^{(k)}}+\varepsilon m \frac{7 R_{k}}{24 R} .
$$

Using Lemma 6.6.7 and the fact that $U_{k}$ contains a $\ell_{\infty}$ ball of radius $r_{k}$, we have

$$
\left\|\nabla \Phi_{\mu}\left(\boldsymbol{x}^{(k)}\right)\right\|_{U^{(k)}} \geq r_{k}\left\|\nabla \Phi_{\mu}\left(\boldsymbol{x}^{(k)}\right)\right\|_{1} \geq \frac{\varepsilon r_{k}}{12 R}\left(\Phi_{\mu}\left(\boldsymbol{x}^{(k)}\right)-2 m\right) .
$$

Therefore, we have that

$$
\begin{aligned}
\Phi_{\mu}\left(\boldsymbol{x}^{(k+1)}\right) & \leq\left(1-\frac{\varepsilon^{2} r_{k}}{24 R}\right) \Phi_{\mu}\left(\boldsymbol{x}^{(k)}\right)+\frac{\varepsilon r_{k}}{12 R} m+\varepsilon m \frac{7 R_{k}}{24 R} \\
& \leq\left(1-\frac{\varepsilon^{2} r_{k}}{24 R}\right) \Phi_{\mu}\left(\boldsymbol{x}^{(k)}\right)+\varepsilon m \frac{R_{k}}{2 R} .
\end{aligned}
$$

Hence, if $\Phi_{\mu}\left(\boldsymbol{x}^{(k)}\right) \leq \frac{12 m \tau}{\varepsilon}$, we have $\Phi_{\mu}\left(\boldsymbol{x}^{(k+1)}\right) \leq \frac{12 m \tau}{\varepsilon}$. Since $\Phi_{\mu}\left(\boldsymbol{x}^{(0)}\right) \leq \frac{12 m \tau}{\varepsilon}$ by assumption we have by induction that $\Phi_{\mu}\left(\boldsymbol{x}^{(k)}\right) \leq \frac{12 m \tau}{\varepsilon}$ for all $k$. The necessary bound on $\left\|\boldsymbol{x}^{(k)}\right\|_{\infty}$ then follows immediately from Lemma 6.6.7.

## - 6.6.5 Centering

Here, we show how to use the results of the previous subsection to perform weighted path following given access only to a multiplicative approximation of the weight function. To apply the results of the previous subsection, we can think updating weight is playing this game, we want to make sure the error between $\boldsymbol{w}$ and $\boldsymbol{g}(\boldsymbol{x})$ is close to 0 while the adversary control the next point $\boldsymbol{g}(\boldsymbol{x})$ and the noise in the approximate $\boldsymbol{g}(\boldsymbol{x})$. Theorem 6.6.6 shows that we can control the error to be small in $\ell_{\infty}$ if we can approximate $\boldsymbol{g}(\boldsymbol{x})$ with small $\ell_{\infty}$ error.

Formally, we will measure its distance from the optimal weights in log scale by

$$
\begin{equation*}
\boldsymbol{\Psi}(\boldsymbol{x}, \boldsymbol{w}) \stackrel{\text { def }}{=} \log (\boldsymbol{g}(\boldsymbol{x}))-\log (\boldsymbol{w}) . \tag{6.29}
\end{equation*}
$$

Our goal will be to keep $\|\Psi(\boldsymbol{x}, \boldsymbol{w})\|_{w+\infty} \leq K$ for some error $K$ that is just small enough to not impair our ability to decrease $\delta_{t}$ linearly and not to impair our ability to approximate $\boldsymbol{g}$. We will attempt to do this without moving $\boldsymbol{w}$ too much in $\|\cdot\|_{w+\infty}$.

```
Algorithm 8: \(\left(\boldsymbol{x}^{\text {(new) }}, \boldsymbol{w}^{(\text {new })}\right)=\) centeringInexact \((\boldsymbol{x}, \boldsymbol{w}, K)\)
    Let \(c_{k}=\frac{1}{1-c_{\delta} c_{\gamma}}, R=\frac{K}{48 c_{k} \log (400 m)}, \delta_{t}=\delta_{t}(\boldsymbol{x}, \boldsymbol{w})\) and \(\varepsilon=\frac{1}{2 c_{k}}\).
    Set \(\boldsymbol{x}^{(\mathrm{new})}=\boldsymbol{x}-\frac{1}{\sqrt{\widehat{\phi^{\prime \prime}}(\boldsymbol{x})}} \boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}}\left(\frac{\left.\boldsymbol{t} \boldsymbol{c}-\boldsymbol{w}{\overrightarrow{\phi^{\prime}}}^{\boldsymbol{w}} \boldsymbol{x}\right)}{\boldsymbol{w} \sqrt{\boldsymbol{\phi}^{\prime \prime}(\boldsymbol{x})}}\right)\).
    Let \(U=\left\{\boldsymbol{x} \in \mathbb{R}^{m} \left\lvert\,\|\boldsymbol{x}\|_{\boldsymbol{w}+\infty} \leq\left(1-\frac{7}{8 c_{k}}\right) \delta_{t}\right.\right\}\).
    Find \(\boldsymbol{z}\) such that \(\| \boldsymbol{z}-\log \left(\boldsymbol{g}\left(\boldsymbol{x}^{\text {new })}\right) \|_{\infty} \leq R\right.\).
    Set \(\boldsymbol{w}^{(\text {new })}=\exp \left(\log (\boldsymbol{w})+(1+\varepsilon) \arg \min _{\boldsymbol{u} \in U}\left\langle\nabla \Phi_{\frac{\varepsilon}{12 R}}(\boldsymbol{z}-\log (\boldsymbol{w})), \boldsymbol{u}\right\rangle\right)\).
    Output ( \(\boldsymbol{x}^{(\text {new })}, \boldsymbol{w}^{(\text {new })}\) ).
```

The minimization problem in last step in centeringInexact is simply a projection onto the convex set $U$ and it can be done in $\tilde{O}(1)$ depth and $\tilde{O}(m)$ work. See section 6.12 for details.

Theorem 6.6.8. Assume that $24 m^{1 / 4} \geq c_{k} \stackrel{\text { def }}{=} \frac{1}{1-c_{\delta} c_{\gamma}} \geq 5,1 \leq C_{n o r m} \leq 2 c_{k}$ and $K \leq \frac{1}{20 c_{k}}$. Let
$\boldsymbol{\Psi}(\boldsymbol{x}, \boldsymbol{w}) \stackrel{\text { def }}{=} \log (\boldsymbol{g}(\boldsymbol{x}))-\log (\boldsymbol{w})$. Suppose that

$$
\delta \stackrel{\text { def }}{=} \delta_{t}(\boldsymbol{x}, \boldsymbol{w}) \leq \frac{K}{48 c_{k} \log (400 m)} \quad \text { and } \quad \Phi_{\mu}(\boldsymbol{\Psi}(\boldsymbol{x}, \boldsymbol{w})) \leq(400 m)^{2}
$$

where $\mu=\frac{\varepsilon}{12 R}=2 \log (400 m) / K$. Let $\left(\boldsymbol{x}^{(n e w)}, \boldsymbol{w}^{(\text {new })}\right)=$ centeringInexact $(\boldsymbol{x}, \boldsymbol{w}, K)$, then

$$
\delta_{t}\left(\boldsymbol{x}^{(n e w)}, \boldsymbol{w}^{(n e w)}\right) \leq\left(1-\frac{1}{4 c_{k}}\right) \delta \quad \text { and } \quad \Phi_{\mu}\left(\boldsymbol{\Psi}\left(\boldsymbol{x}^{(n e w)}, \boldsymbol{w}^{(n e w)}\right)\right) \leq(400 m)^{2} .
$$

Also, we have $\left\|\log \left(\boldsymbol{g}\left(\boldsymbol{x}^{(\text {new })}\right)\right)-\log (\boldsymbol{w})\right\|_{\infty} \leq K$.

Proof. By Lemma 6.6.3, inequality (6.16), $c_{\delta} c_{\gamma} \leq 1$ and $c_{\gamma} \leq \frac{5}{4}$ (see Def 6.5.4), we have

$$
\begin{aligned}
\left\|\log \left(\boldsymbol{g}\left(\boldsymbol{x}^{(\mathrm{new})}\right)\right)-\log (\boldsymbol{g}(\boldsymbol{x}))\right\|_{\boldsymbol{g}(\boldsymbol{x})+\infty} & \leq c_{\delta} c_{\gamma} \delta\left(1+4 c_{\gamma} \delta\right) \\
& \leq c_{\delta} c_{\gamma} \delta+5 \delta^{2} \\
& \leq\left(1-\frac{15}{16 c_{k}}\right) \delta
\end{aligned}
$$

Using $K \leq \frac{1}{20 c_{k}}$, we have

$$
\left\|\frac{\boldsymbol{w}-\boldsymbol{g}(\boldsymbol{x})}{\boldsymbol{g}(\boldsymbol{x})}\right\|_{\infty} \leq\|\log (\boldsymbol{w})-\log (\boldsymbol{g}(\boldsymbol{x}))\|_{\infty}+\|\log (\boldsymbol{w})-\log (\boldsymbol{g}(\boldsymbol{x}))\|_{\infty}^{2} \leq \frac{21}{20} K \leq \frac{1}{16 c_{k}} .
$$

Hence, we have

$$
\begin{aligned}
\left\|\log \left(\boldsymbol{g}\left(\boldsymbol{x}^{(\mathrm{new})}\right)\right)-\log (\boldsymbol{g}(\boldsymbol{x}))\right\|_{\boldsymbol{w}+\infty} & \leq\left(1+\frac{1}{16 c_{k}}\right)\left(1-\frac{15}{16 c_{k}}\right) \delta \\
& \leq\left(1-\frac{7}{8 c_{k}}\right) \delta
\end{aligned}
$$

Therefore, we know that for the Newton step, we have $\boldsymbol{\Psi}\left(\boldsymbol{x}^{(\mathrm{new})}, \boldsymbol{w}\right)-\boldsymbol{\Psi}(\boldsymbol{x}, \boldsymbol{w}) \in U$ where $U$ is the symmetric convex set given by

$$
U \stackrel{\text { def }}{=}\left\{\boldsymbol{x} \in \mathbb{R}^{n} \mid\|\boldsymbol{x}\|_{\boldsymbol{w}+\infty} \leq C\right\}
$$

where $C=\left(1-\frac{7}{8 c_{k}}\right) \delta$. Note that from our assumption on $\delta$, we have

$$
C \leq \delta \leq \frac{K}{48 c_{k} \log (400 m)}=R
$$

It ensures that $U$ are contained in some $\ell_{\infty}$ ball of radius $R$. Therefore, we can play the chasing 0 game on $\boldsymbol{\Psi}(\boldsymbol{x}, \boldsymbol{w})$ attempting to maintain the invariant that $\|\boldsymbol{\Psi}(\boldsymbol{x}, \boldsymbol{w})\|_{\infty} \leq K$ without taking steps that are more than $1+\varepsilon$ times the size of $U$ where we pick $\varepsilon=\frac{1}{2 c_{k}}$ so to not interfere with our ability to decrease $\delta_{t}$ linearly.

However, to do this with the chasing 0 game, we need to ensure that $R$ satisfying the following

$$
\frac{12 R}{\varepsilon} \log \left(\frac{12 m \tau}{\varepsilon}\right) \leq K
$$

where here $\tau$ is as defined in Theorem 6.6.6.
To bound $\tau$, we need to lower bound the radius of $\ell_{\infty}$ ball it contains. Since by assumption
$\|\boldsymbol{g}(\boldsymbol{x})\|_{\infty} \leq 2$ and $\|\Psi(\boldsymbol{x}, \boldsymbol{w})\|_{\infty} \leq \frac{1}{8}$, we have that $\|\boldsymbol{w}\|_{\infty} \leq 3$. Hence, we have

$$
\forall u \in \mathbb{R}^{m} \quad: \quad\|\boldsymbol{u}\|_{\infty}^{2} \geq \frac{1}{3 m}\|\boldsymbol{u}\|_{\boldsymbol{w}}^{2}
$$

Consequently, if $\|\boldsymbol{u}\|_{\infty} \leq \frac{\delta}{5 C_{\text {norm }} \sqrt{m}}$, then $\boldsymbol{u} \in U$. So, we have that $U$ contains a box of radius $\frac{\delta}{5 C_{\text {norm }} \sqrt{m}}$ and since $U$ is contained in a box of radius $\delta$, we have that

$$
\tau \leq 5 C_{\text {norm }} \sqrt{m} \leq 10 c_{k} \sqrt{m}
$$

Using $c_{k} \leq 24 m^{1 / 4}$, we have

$$
\begin{aligned}
\frac{12 R}{\varepsilon} \log \left(\frac{12 m \tau}{\varepsilon}\right) & \leq 24 c_{k} R \log \left(240 m^{3 / 2} c_{k}^{2}\right) \\
& \leq 48 c_{k} R \log (400 m)=K
\end{aligned}
$$

and

$$
\frac{12 m \tau}{\varepsilon} \leq 240 m^{3 / 2} c_{k}^{2} \leq(400 m)^{2}
$$

This proves that we meet the conditions of Theorem 6.6.6. Consequently, $\left\|\boldsymbol{\Psi}\left(\boldsymbol{x}^{(\text {new })}, \boldsymbol{w}^{(\text {new })}\right)\right\|_{\infty} \leq K$ and $\Phi_{\alpha}\left(\boldsymbol{\Psi}\left(\boldsymbol{x}^{(\text {new })}, \boldsymbol{w}^{(\text {new })}\right)\right) \leq(400 m)^{2}$.

Since $K \leq \frac{1}{4}$, Lemma 6.6 .2 shows that

$$
\delta_{t}\left(\boldsymbol{x}^{(\text {new })}, \boldsymbol{w}\right) \leq 4\left(\delta_{t}(\boldsymbol{x}, \boldsymbol{w})\right)^{2}
$$

Our choice of $\boldsymbol{w}^{(\text {new })}$ shows that

$$
\begin{aligned}
\left\|\log (\boldsymbol{w})-\log \left(\boldsymbol{w}^{(\mathrm{new})}\right)\right\|_{\boldsymbol{w}+\infty} & \leq\left(1+\frac{1}{2 c_{k}}\right)\left(1-\frac{7}{8 c_{k}}\right) \delta \\
& \leq\left(1-\frac{3}{8 c_{k}}\right) \delta
\end{aligned}
$$

Using $\delta \leq \frac{1}{80 c_{k}}$, the Lemma 6.6 .4 shows that

$$
\begin{aligned}
\delta_{t}\left(\boldsymbol{x}^{(\text {new })}, \boldsymbol{w}^{(\mathrm{new})}\right) & \leq\left(1+4\left(1-\frac{3}{8 c_{k}}\right) \delta\right)\left(\delta_{t}\left(\boldsymbol{x}^{(\mathrm{new})}, \boldsymbol{w}\right)+\left(1-\frac{3}{8 c_{k}}\right) \delta\right) \\
& \leq(1+4 \delta)\left(4 \delta^{2}+\left(1-\frac{3}{8 c_{k}}\right) \delta\right) \\
& \leq\left(1-\frac{3}{8 c_{k}}\right) \delta+4 \delta^{2}+16 \delta^{3}+4 \delta^{2} \\
& \leq\left(1-\frac{1}{4 c_{k}}\right) \delta
\end{aligned}
$$

## ■ 6.7 Weight Function

In this section we present the weight function that we use to achieve our $\widetilde{O}(\sqrt{\operatorname{rank}(\boldsymbol{A})} \log (U / \varepsilon))$ iteration linear program solver. This weight function is motivated from the barrier introduced in

Subsection 6.3.5. We define the weight function $\boldsymbol{g}: \Omega \rightarrow \mathbb{R}_{>0}^{m}$ for all $\boldsymbol{x} \in \mathbb{R}_{>0}^{m}$ as follows

$$
\begin{equation*}
\boldsymbol{g}(\boldsymbol{x}) \stackrel{\text { def }}{=} \underset{\boldsymbol{w} \in \mathbb{R}_{>0}^{m}}{\arg \min } \hat{f}(\boldsymbol{x}, \boldsymbol{w}) \quad \text { where } \quad \hat{f}(\boldsymbol{x}, \boldsymbol{w}) \stackrel{\text { def }}{=} \mathbf{1}^{T} \boldsymbol{w}+\frac{1}{\alpha} \log \operatorname{det}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W}^{-\alpha} \boldsymbol{A}_{x}\right)-\beta \sum_{i \in[m]} \log w_{i} . \tag{6.30}
\end{equation*}
$$

where here and in the remainder of the subsection we let $\boldsymbol{A}_{x} \stackrel{\text { def }}{=}\left(\boldsymbol{\Phi}^{\prime \prime}(\boldsymbol{x})\right)^{-1 / 2} \boldsymbol{A}$ and the parameters $\alpha, \beta$ are chosen later such that the following hold

$$
\begin{equation*}
\alpha \in[1,2), \beta \in(0,1), \text { and } \beta^{1-\alpha} \leq 2 . \tag{6.31}
\end{equation*}
$$

Here we choose $\beta$ small and $\alpha$ just slightly larger than 1 .
We start by computing the gradient and Hessian of $\hat{f}(\boldsymbol{x}, \boldsymbol{w})$ with respect to $\boldsymbol{w}$.
Lemma 6.7.1. For all $\boldsymbol{x} \in \Omega$ and $\boldsymbol{w} \in \mathbb{R}_{>0}^{m}$, we have

$$
\nabla_{w} \hat{f}(\boldsymbol{x}, \boldsymbol{w})=\left(\boldsymbol{I}-\boldsymbol{\Sigma} \boldsymbol{W}^{-1}-\beta \boldsymbol{W}^{-1}\right) \mathbf{1} \quad \text { and } \quad \nabla_{w w}^{2} \hat{f}(\boldsymbol{x}, \boldsymbol{w})=\boldsymbol{W}^{-1}(\boldsymbol{\Sigma}+\beta \boldsymbol{I}+\alpha \boldsymbol{\Lambda}) \boldsymbol{W}^{-1}
$$

where $\boldsymbol{\Sigma} \stackrel{\text { def }}{=} \boldsymbol{\Sigma}_{\boldsymbol{A}_{x}}\left(\boldsymbol{w}^{-\alpha}\right)$ and $\boldsymbol{\Lambda} \stackrel{\text { def }}{=} \boldsymbol{\Lambda}_{\boldsymbol{A}_{x}}\left(\boldsymbol{w}^{-\alpha}\right)$.
Proof. Using Lemma 6.11 .2 and the chain rule we compute the gradient of $\nabla_{w} \hat{f}(\boldsymbol{x}, \boldsymbol{w})$ as follows

$$
\begin{aligned}
\nabla_{w} \hat{f}(\boldsymbol{x}, \boldsymbol{w}) & =\mathbf{1}+\frac{1}{\alpha} \boldsymbol{\Sigma} \boldsymbol{W}^{\alpha}\left(-\alpha \boldsymbol{W}^{-\alpha-1}\right)-\beta \boldsymbol{W}^{-1} \mathbf{1} \\
& =\left(\boldsymbol{I}-\boldsymbol{\Sigma} \boldsymbol{W}^{-1}-\beta \boldsymbol{W}^{-1}\right) \mathbf{1}
\end{aligned}
$$

Next, using Lemma 6.11.2 and chain rule, we compute the following for all $i, j \in[m]$

$$
\begin{aligned}
\frac{\partial\left(\nabla_{w} \hat{f}(\boldsymbol{x}, \boldsymbol{w})\right)_{i}}{\partial \boldsymbol{w}_{j}} & =-\frac{\boldsymbol{w}_{i} \boldsymbol{\Lambda}_{i j} \boldsymbol{w}_{j}^{\alpha}\left(-\alpha \boldsymbol{w}_{j}^{-\alpha-1}\right)-\boldsymbol{\Sigma}_{i j} \mathbf{1}_{i=j}}{\boldsymbol{w}_{i}^{2}}+\beta \mathbf{1}_{i=j}\left\{\boldsymbol{w}_{i}^{-2}\right\} \\
& =\frac{\boldsymbol{\Sigma}_{i j}}{\boldsymbol{w}_{i} \boldsymbol{w}_{j}}+\alpha \frac{\boldsymbol{\Lambda}_{i j}}{\boldsymbol{w}_{i} \boldsymbol{w}_{j}}+\frac{\beta \mathbf{1}_{i=j}}{\boldsymbol{w}_{i}^{2}} . \quad \text { (Using that } \boldsymbol{\Sigma} \text { is diagonal) }
\end{aligned}
$$

Consequently, $\nabla_{w w}^{2} \hat{f}(\boldsymbol{x}, \boldsymbol{w})=\boldsymbol{W}^{-1}(\boldsymbol{\Sigma}+\beta \boldsymbol{I}+\alpha \boldsymbol{\Lambda}) \boldsymbol{W}^{-1}$ as desired.
Lemma 6.7.2. For all $\boldsymbol{x} \in \Omega$, the weight function $\boldsymbol{g}(\boldsymbol{x})$ is a well defined with

$$
\beta \leq g_{i}(\boldsymbol{s}) \leq 1+\beta \quad \text { and } \quad\|\boldsymbol{g}(\boldsymbol{x})\|_{1}=\operatorname{rank}(\boldsymbol{A})+\beta \cdot m .
$$

Furthermore, for all $\boldsymbol{x} \in \Omega$, the weight function obeys the following equations

$$
\boldsymbol{G}(\boldsymbol{x})=(\boldsymbol{\Sigma}+\beta \boldsymbol{I}) \mathbf{1} \quad, \text { and } \quad \boldsymbol{G}^{\prime}(\boldsymbol{x})=-\boldsymbol{G}(\boldsymbol{x})(\boldsymbol{G}(\boldsymbol{x})+\alpha \boldsymbol{\Lambda})^{-1} \boldsymbol{\Lambda}\left(\boldsymbol{\Phi}^{\prime \prime}(\boldsymbol{x})\right)^{-1} \boldsymbol{\Phi}^{\prime \prime \prime}(\boldsymbol{x})
$$

where $\boldsymbol{\Sigma} \stackrel{\text { def }}{=} \boldsymbol{\Sigma}_{\boldsymbol{A}_{x}}\left(\boldsymbol{g}^{-\alpha}(\boldsymbol{x})\right), \boldsymbol{\Lambda} \stackrel{\text { def }}{=} \boldsymbol{\Lambda}_{\boldsymbol{A}_{x}}\left(\boldsymbol{g}^{-\alpha}(\boldsymbol{x})\right)$, and $\boldsymbol{G}^{\prime}(\boldsymbol{x})$ is the Jacobian matrix of $\boldsymbol{g}$ at $\boldsymbol{x}$.
Proof. By Lemma 6.11.2 we have that $\boldsymbol{\Sigma} \succeq \boldsymbol{\Lambda} \succeq \mathbf{0}$. Therefore, by Lemma 6.7.1, we have that $\nabla_{w w}^{2} \hat{f}(\boldsymbol{x}, \boldsymbol{w}) \succeq \beta \boldsymbol{W}^{-2}$ and $\hat{f}(\boldsymbol{x}, \boldsymbol{w})$ is convex. Using the formula for the gradient in Lemma 6.7.1, we see that that for all $i \in[m]$ it is the case that

$$
\left[\nabla_{w} \hat{f}(\boldsymbol{x}, \boldsymbol{w})\right]_{i}=\frac{1}{w_{i}}\left(w_{i}-\boldsymbol{\Sigma}_{i i}-\beta\right)
$$

Using that $0 \leq \sigma_{i} \leq 1$ for all $i$ by Lemma 6.11 .2 and $\beta \in(0,1)$ by (6.31), we see that if $\boldsymbol{w}_{i} \in(0, \beta)$ then $\left[\nabla_{w} \hat{f}(\boldsymbol{x}, \boldsymbol{w})\right]_{i}$ is strictly negative and if $\boldsymbol{w}_{i} \in(1+\beta, \infty)$ then $\left[\nabla_{w} \hat{f}(\boldsymbol{x}, \boldsymbol{w})\right]_{i}$ is strictly positive.

Therefore, for any $\boldsymbol{x} \in \Omega$, the $\boldsymbol{w}$ that minimizes this convex function $\hat{f}(\boldsymbol{x}, \boldsymbol{w})$ lies between the box between $\beta$ to $1+\beta$. Since $\hat{f}$ is strongly convex in this region, the minimizer is unique.

The formula for $\boldsymbol{G}(\boldsymbol{x})$ follows by setting $\nabla_{w} \hat{f}(\boldsymbol{x}, \boldsymbol{w})=\mathbf{0}$ and the size of $g(\boldsymbol{x})$ follows from the fact that $\|\boldsymbol{\sigma}\|_{1}=\operatorname{tr}\left(\boldsymbol{P}_{\boldsymbol{A}_{x}}\left(\boldsymbol{g}^{-\alpha}(\boldsymbol{x})\right)\right)$. Since $\boldsymbol{P}_{\boldsymbol{A}_{x}}\left(\boldsymbol{g}^{-\alpha}(\boldsymbol{x})\right)$ is a projection onto the image of $\boldsymbol{G}(\boldsymbol{x})^{-\alpha / 2} \boldsymbol{A}_{x}$ and since $\boldsymbol{g}(\boldsymbol{x})>\mathbf{0}$ and $\vec{\phi}^{\prime \prime}(\boldsymbol{x})>\mathbf{0}$, we have that the dimension of the image of $\boldsymbol{G}(\boldsymbol{x})^{-\alpha / 2} \boldsymbol{A}_{x}$ is the rank of $\boldsymbol{A}$. Hence, we have that $\|\boldsymbol{g}(\boldsymbol{x})\|_{1}=\operatorname{rank}(\boldsymbol{A})+\beta \cdot m$.

By Lemma 6.11.2 and chain rule, we get the following for all $i, j \in[m]$

$$
\frac{\partial\left(\nabla_{\boldsymbol{w}} \hat{f}(\boldsymbol{x}, \boldsymbol{w})\right)_{i}}{\partial \boldsymbol{x}_{j}}=-\boldsymbol{w}_{i}^{-1} \boldsymbol{\Lambda}_{i j} \vec{\phi}_{j}^{\prime \prime}(\boldsymbol{x})\left(-\left(\vec{\phi}_{j}^{\prime \prime}(\boldsymbol{x})\right)^{-2} \vec{\phi}_{j}^{\prime \prime \prime}(\boldsymbol{x})\right)=\boldsymbol{w}_{i}^{-1} \boldsymbol{\Lambda}_{i j}\left(\vec{\phi}_{j}^{\prime \prime}(\boldsymbol{x})\right)^{-1} \vec{\phi}_{j}^{\prime \prime \prime}(\boldsymbol{x}) .
$$

Consequently, $\boldsymbol{J}_{\boldsymbol{x}}\left(\nabla_{\boldsymbol{w}} \hat{f}(\boldsymbol{x}, \boldsymbol{w})\right)=\boldsymbol{W}^{-1} \boldsymbol{\Lambda}\left(\boldsymbol{\Phi}^{\prime \prime}(\boldsymbol{x})\right)^{-1} \boldsymbol{\Phi}^{\prime \prime \prime}(\boldsymbol{x})$ where $\boldsymbol{J}_{\boldsymbol{x}}$ denotes the Jacobian matrix of the function $\nabla_{\boldsymbol{w}} \hat{f}(\boldsymbol{x}, \boldsymbol{w})$ with respect to $\boldsymbol{x}$. Since we have already know that $\boldsymbol{J}_{\boldsymbol{w}}\left(\nabla_{\boldsymbol{w}} \hat{f}(\boldsymbol{x}, \boldsymbol{w})\right)=$ $\nabla_{\boldsymbol{w} \boldsymbol{w}}^{2} f_{t}(\boldsymbol{x}, \boldsymbol{w})=\boldsymbol{W}^{-1}(\boldsymbol{\Sigma}+\beta \boldsymbol{I}+\alpha \boldsymbol{\Lambda}) \boldsymbol{W}^{-1}$ is positive definite (and hence invertible), by applying the implicit function theorem (Lemma 2.3.9) to the specification of $\boldsymbol{g}(\boldsymbol{x})$ as the solution to $\nabla_{\boldsymbol{w}} \hat{f}(\boldsymbol{x}, \boldsymbol{w})=\mathbf{0}$, we have

$$
\boldsymbol{G}^{\prime}(\boldsymbol{x})=-\left(\boldsymbol{J}_{\boldsymbol{w}}\left(\nabla_{w} \hat{f}(\boldsymbol{x}, \boldsymbol{w})\right)\right)^{-1}\left(\boldsymbol{J}_{\boldsymbol{x}}\left(\nabla_{\boldsymbol{w}} \hat{f}(\boldsymbol{x}, \boldsymbol{w})\right)\right)=-\boldsymbol{G}(\boldsymbol{x})(\boldsymbol{G}(\boldsymbol{x})+\alpha \boldsymbol{\Lambda})^{-1} \boldsymbol{\Lambda}\left(\boldsymbol{\Phi}^{\prime \prime}(\boldsymbol{x})\right)^{-1} \boldsymbol{\Phi}^{\prime \prime \prime}(\boldsymbol{x})
$$

Now we show the step consistency of $\boldsymbol{g}$.
Lemma 6.7.3 (Step Consistency). For all $\boldsymbol{x} \in \Omega$ and $\boldsymbol{y} \in \mathbb{R}^{m}$, and

$$
\boldsymbol{B} \stackrel{\text { def }}{=} \boldsymbol{G}(\boldsymbol{x})^{-1} \boldsymbol{G}^{\prime}(\boldsymbol{x})\left(\boldsymbol{\phi}(\boldsymbol{x})^{\prime \prime}\right)^{-1 / 2},
$$

we have

$$
\|\boldsymbol{B} \boldsymbol{y}\|_{\boldsymbol{G}(\boldsymbol{x})} \leq \frac{2}{1+\alpha}\|\boldsymbol{y}\|_{\boldsymbol{G}(\boldsymbol{x})} \quad \text { and } \quad\|\boldsymbol{B} \boldsymbol{y}\|_{\infty} \leq \frac{2}{1+\alpha}\left(\|\boldsymbol{y}\|_{\infty}+\frac{1+2 \alpha}{1+\alpha}\|\boldsymbol{y}\|_{\boldsymbol{G}(\boldsymbol{x})}\right) .
$$

Therefore

$$
\|\boldsymbol{B}\|_{g+\infty} \leq \frac{2}{1+\alpha}\left(1+\frac{2}{C_{\text {norm }}}\right) .
$$

Proof. Fix an arbitrary $\boldsymbol{x} \in \Omega$ and let $\boldsymbol{g} \stackrel{\text { def }}{=} \boldsymbol{g}(\boldsymbol{x}), \boldsymbol{\sigma} \stackrel{\text { def }}{=} \boldsymbol{\sigma}_{\boldsymbol{A}_{\boldsymbol{x}}}\left(\boldsymbol{g}^{-\alpha}(\boldsymbol{x})\right), \boldsymbol{\Sigma} \stackrel{\text { def }}{=} \boldsymbol{\Sigma}_{\boldsymbol{A}_{x}}\left(\boldsymbol{g}^{-\alpha}(\boldsymbol{x})\right), \boldsymbol{P} \stackrel{\text { def }}{=}$ $\boldsymbol{P}_{\boldsymbol{A}_{\boldsymbol{x}}}\left(\boldsymbol{g}^{-\alpha}(\boldsymbol{x})\right), \boldsymbol{\Lambda} \stackrel{\text { def }}{=} \boldsymbol{\Lambda}_{\boldsymbol{A}_{\boldsymbol{x}}}\left(\boldsymbol{g}^{-\alpha}(\boldsymbol{x})\right)$. Also, fix an arbitrary $\boldsymbol{y} \in \mathbb{R}^{m}$ and let $\boldsymbol{z} \stackrel{\text { def }}{=} \boldsymbol{B} \boldsymbol{y}$.

By Lemma 6.7.2, $\boldsymbol{G}^{\prime}=-\boldsymbol{G}(\boldsymbol{G}+\alpha \boldsymbol{\Lambda})^{-1} \boldsymbol{\Lambda}\left(\boldsymbol{\Phi}^{\prime \prime}\right)^{-1} \boldsymbol{\Phi}^{\prime \prime \prime}$ and therefore

$$
\begin{aligned}
\boldsymbol{B} & =-\boldsymbol{G}^{-1}\left(\boldsymbol{G}(\boldsymbol{G}+\alpha \boldsymbol{\Lambda})^{-1} \boldsymbol{\Lambda}\left(\boldsymbol{\Phi}^{\prime \prime}\right)^{-1} \boldsymbol{\Phi}^{\prime \prime \prime}\right)\left(\boldsymbol{\Phi}^{\prime \prime}\right)^{-1 / 2} \\
& =(\boldsymbol{G}+\alpha \boldsymbol{\Lambda})^{-1}(2 \boldsymbol{\Lambda}) \boldsymbol{D i a g}\left(\frac{-\boldsymbol{\phi}^{\prime \prime \prime}}{2\left(\phi^{\prime \prime}\right)^{3 / 2}}\right)
\end{aligned}
$$

Let $\boldsymbol{C} \stackrel{\text { def }}{=}(\boldsymbol{G}+\alpha \boldsymbol{\Lambda})^{-1}(2 \boldsymbol{\Lambda})$ and let $\boldsymbol{y}^{\prime} \stackrel{\text { def }}{=} \boldsymbol{\operatorname { D i a g }}\left(\frac{-\phi^{\prime \prime \prime}}{2\left(\phi^{\prime \prime}\right)^{3 / 2}}\right) \boldsymbol{y}$. By the self concordance of $\boldsymbol{\phi}(6.11)$, we know that $\left\|\boldsymbol{y}^{\prime}\right\| \leq\|\boldsymbol{y}\|$ for both $\|\cdot\|_{\boldsymbol{G}}$ and $\|\cdot\|_{\infty}$. Since $\boldsymbol{z}=\boldsymbol{B} \boldsymbol{y}=\boldsymbol{C} \boldsymbol{y}^{\prime}$, it suffices to bound $\left\|\boldsymbol{C} \boldsymbol{y}^{\prime}\right\|$ in terms of $\left\|\boldsymbol{y}^{\prime}\right\|$ for the necessary norms.

Letting $\overline{\boldsymbol{\Lambda}} \xlongequal{\text { def }} \boldsymbol{G}^{-1 / 2} \boldsymbol{\Lambda} \boldsymbol{G}^{-1 / 2}$, we simplify the equation further and note that

$$
\|\boldsymbol{C}\|_{\boldsymbol{G}}=\left\|\boldsymbol{G}^{1 / 2}(\boldsymbol{G}+\alpha \boldsymbol{\Lambda})^{-1}(2 \boldsymbol{\Lambda}) \boldsymbol{G}^{-1 / 2}\right\|_{2}=\left\|(\boldsymbol{I}+\alpha \overline{\boldsymbol{\Lambda}})^{-1}(2 \overline{\boldsymbol{\Lambda}})\right\|_{2} .
$$

Now, for any eigenvector, $\boldsymbol{v}$, of $\overline{\boldsymbol{\Lambda}}$ with eigenvalue $\lambda$, we see that $\boldsymbol{v}$ is an eigenvector of $(\boldsymbol{I}+\alpha \overline{\boldsymbol{\Lambda}})^{-1}(2 \overline{\boldsymbol{\Lambda}})$ with eigenvalue $2 \lambda /(1+\alpha \lambda)$. Furthermore, since $\mathbf{0} \preceq \overline{\boldsymbol{\Lambda}} \preceq \boldsymbol{I}$, we have that $\|\boldsymbol{C}\|_{\boldsymbol{G}} \leq 2 /(1+\alpha)$ and hence $\|\boldsymbol{z}\|_{\boldsymbol{G}} \leq 2(1+\alpha)^{-1}\left\|\boldsymbol{y}^{\prime}\right\|_{\boldsymbol{G}} \leq 2(1+\alpha)^{-1}\|\boldsymbol{y}\|_{\boldsymbol{G}}$ as desired.

To bound $\|\boldsymbol{z}\|_{\infty}$, we use that $(\boldsymbol{G}+\alpha \boldsymbol{\Lambda}) \boldsymbol{z}=2 \boldsymbol{\Lambda} \boldsymbol{y}^{\prime}, \boldsymbol{\Lambda}=\boldsymbol{\Sigma}-\boldsymbol{P}^{(2)}$, and $\boldsymbol{G}=\boldsymbol{\Sigma}+\beta \boldsymbol{I}$ to derive

$$
(1+\alpha) \boldsymbol{\Sigma} \boldsymbol{z}+\beta \boldsymbol{z}-\alpha \boldsymbol{P}^{(2)} \boldsymbol{z}=2 \boldsymbol{\Sigma} \boldsymbol{y}^{\prime}-2 \boldsymbol{P}^{(2)} \boldsymbol{y}^{\prime}
$$

Looking at the $i^{\text {th }}$ coordinate of both sides and using that $\boldsymbol{\sigma}_{i} \geq 0$, we have

$$
\begin{aligned}
&\left((1+\alpha) \boldsymbol{\sigma}_{i}+\beta\right)\left|\boldsymbol{z}_{i}\right| \\
& \leq \\
& \leq\left|\left[\boldsymbol{P}^{(2)} \boldsymbol{z}\right]_{i}\right|+2 \boldsymbol{\sigma}_{i}\left\|\boldsymbol{y}^{\prime}\right\|_{\infty}+2\left|\left[\boldsymbol{P}^{(2)} \boldsymbol{y}^{\prime}\right]_{i}\right| \\
& \leq \alpha \boldsymbol{\sigma}_{i}\|\boldsymbol{z}\|_{\boldsymbol{\Sigma}}+2 \boldsymbol{\sigma}_{i}\left\|\boldsymbol{y}^{\prime}\right\|_{\infty}+2 \boldsymbol{\sigma}_{i}\left\|\boldsymbol{y}^{\prime}\right\|_{\boldsymbol{\Sigma}} \\
& \leq 2 \boldsymbol{\sigma}_{i}\left\|\boldsymbol{y}^{\prime}\right\|_{\infty}+\boldsymbol{\sigma}_{i}\left(\frac{2 \alpha}{1+\alpha}+2\right)\left\|\boldsymbol{y}^{\prime}\right\|_{\boldsymbol{G}} \quad \quad \text { (Lemma 6.11.2) }
\end{aligned}
$$

Hence, we have

$$
\begin{aligned}
\left|\boldsymbol{z}_{i}\right| & \leq \frac{2}{1+\alpha}\left\|\boldsymbol{y}^{\prime}\right\|_{\infty}+\frac{1}{1+\alpha}\left(\frac{2 \alpha}{1+\alpha}+2\right)\left\|\boldsymbol{y}^{\prime}\right\|_{\boldsymbol{G}} \\
& \leq \frac{2}{1+\alpha}\left[\left\|\boldsymbol{y}^{\prime}\right\|_{\infty}+2\left\|\boldsymbol{y}^{\prime}\right\|_{\boldsymbol{G}}\right] .
\end{aligned}
$$

Therefore, $\|\boldsymbol{B} \boldsymbol{y}\|_{\infty}=\|\boldsymbol{z}\|_{\infty} \leq 2(1+\alpha)^{-1}\left(\left\|\boldsymbol{y}^{\prime}\right\|_{\infty}+2\left\|\boldsymbol{y}^{\prime}\right\|_{\boldsymbol{G}}\right)$. Finally, we note that

$$
\begin{align*}
\|\boldsymbol{B} \boldsymbol{y}\|_{\boldsymbol{g}+\infty} & =\|\boldsymbol{B} \boldsymbol{y}\|_{\infty}+C_{\text {norm }}\|\boldsymbol{B} \boldsymbol{y}\|_{\boldsymbol{G}}  \tag{Definition}\\
& \leq \frac{2}{1+\alpha}\|\boldsymbol{y}\|_{\infty}+\frac{2}{1+\alpha} \cdot 2\|\boldsymbol{y}\|_{\boldsymbol{G}}+\frac{2}{1+\alpha} C_{\text {norm }}\|\boldsymbol{y}\|_{\boldsymbol{G}} \\
& \leq \frac{2}{1+\alpha}\left(1+\frac{2}{C_{\text {norm }}}\right)\|\boldsymbol{y}\|_{\boldsymbol{g}+\infty} .
\end{align*}
$$

Theorem 6.7.4. Choosing parameters

$$
\alpha=1+\frac{1}{\log _{2}\left(\frac{2 m}{\operatorname{rank}(\boldsymbol{A})}\right)} \quad, \quad \beta=\frac{\operatorname{rank}(\boldsymbol{A})}{2 m} \quad \text {, and } \quad C_{\text {norm }}=18 \log _{2}\left(\frac{2 m}{\operatorname{rank}(\boldsymbol{A})}\right)
$$

yields

$$
c_{1}(\boldsymbol{g})=2 \operatorname{rank}(\boldsymbol{A}) \quad, \quad c_{\gamma}(\boldsymbol{g})=1+\frac{1}{9 \log _{2}\left(\frac{2 m}{\operatorname{rank}(\boldsymbol{A})}\right)} \quad, \quad \text { and } \quad c_{\delta}(\boldsymbol{g})=1-\frac{2}{9 \log _{2}\left(\frac{2 m}{\operatorname{rank}(\boldsymbol{A})}\right)} .
$$

In particular, we have

$$
c_{\gamma}(\boldsymbol{g}) c_{\delta}(\boldsymbol{g}) \leq 1-\frac{1}{9 \log _{2}\left(\frac{2 m}{\operatorname{rank}(\boldsymbol{A})}\right)} .
$$

Proof. The bounds on $c_{1}(\boldsymbol{g})$ and $c_{\delta}(\boldsymbol{g})$ follow immediately from Lemma 6.7.2 and Lemma 6.7.3. Now, we estimate the $c_{\gamma}(\boldsymbol{g})$ and let $\frac{4}{5} \boldsymbol{g} \leq \boldsymbol{w} \leq \frac{5}{4} \boldsymbol{g}$. Fix an arbitrary $\boldsymbol{x} \in \Omega$ and let $\boldsymbol{g} \xlongequal{\text { def }} \boldsymbol{g}(\boldsymbol{x})$. Recall that by Lemma 6.7.2, we have $\boldsymbol{g} \geq \beta$. Furthermore, since $\boldsymbol{g}^{-1}=\boldsymbol{g}^{\alpha-1} \boldsymbol{g}^{-\alpha}$ and $\beta^{\alpha-1} \geq \frac{1}{2}$, the following holds

$$
\begin{equation*}
\frac{4}{10} \boldsymbol{g}_{i}^{-\alpha} \leq \frac{4}{5} \beta^{\alpha-1} \boldsymbol{g}_{i}^{-\alpha} \leq \frac{4}{10} \boldsymbol{g}_{i}^{-1} \leq \boldsymbol{w}_{i}^{-1} \tag{6.32}
\end{equation*}
$$

for all $i$. Applying this and using the definition of $\boldsymbol{P}_{\boldsymbol{A}_{x}}$ yields

$$
\begin{equation*}
\boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W}^{-1} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}_{x}^{T} \preceq \frac{10}{4} \boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{G}^{-\alpha} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}_{x}^{T}=\frac{10}{4} \boldsymbol{G}^{\alpha / 2} \boldsymbol{P}_{\boldsymbol{A}_{x}}\left(\boldsymbol{g}^{-\alpha}\right) \boldsymbol{G}^{\alpha / 2} . \tag{6.33}
\end{equation*}
$$

Hence, we have

$$
\begin{aligned}
\frac{\boldsymbol{\sigma}_{i}\left(\frac{1}{\boldsymbol{w} \vec{\phi}^{\prime \prime}}\right)}{\boldsymbol{w}_{i}} & =\frac{\mathbf{1}_{i}^{T} \boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W}^{-1} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}_{x}^{T} \mathbf{1}_{i}}{\boldsymbol{w}_{i}^{2}} \\
& \leq \frac{10}{4} \frac{\mathbf{1}_{i}^{T} \boldsymbol{G}^{\alpha / 2} \boldsymbol{P}_{\boldsymbol{A}_{x}}\left(\boldsymbol{g}^{-\alpha}\right) \boldsymbol{G}^{\alpha / 2} \mathbf{1}_{i}}{\boldsymbol{w}_{i}^{2}} \\
& \leq \frac{10}{4}\left(\frac{5}{4}\right)^{2} \frac{\boldsymbol{\sigma}_{i}\left(\frac{1}{\boldsymbol{g}^{\alpha} \vec{\phi}^{\prime \prime}}\right)}{\boldsymbol{g}_{i}^{-2 \alpha}}<4 .
\end{aligned}
$$

Since $\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}}$ is an orthogonal projection in $\|\cdot\|_{\boldsymbol{w}}$, we have $\left\|\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}}\right\|_{\boldsymbol{w} \rightarrow \boldsymbol{w}}=1$. Let $\overline{\boldsymbol{P}}_{\boldsymbol{x}, \boldsymbol{w}} \stackrel{\text { def }}{=} \boldsymbol{I}-\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}}$, we have

$$
\begin{aligned}
\left\|\overline{\boldsymbol{P}}_{\boldsymbol{x}, \boldsymbol{w}}\right\|_{\boldsymbol{w} \rightarrow \infty} & =\max _{i \in[m]} \max _{\|\boldsymbol{y}\| \boldsymbol{w} \leq 1} \mathbf{1}_{i}^{T} \overline{\boldsymbol{P}}_{\boldsymbol{x}, \boldsymbol{w}} \boldsymbol{y} \\
& \leq \max _{i \in[m]}\left\|(\boldsymbol{w})^{-1 / 2} \overline{\boldsymbol{P}}_{\boldsymbol{x}, \boldsymbol{w}}^{T} \mathbf{1}_{i}\right\|^{2} \\
& =\max _{i \in[m]} \sqrt{\mathbf{1}_{i} \boldsymbol{W}^{-1} \boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W}^{-1} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}_{x}^{T} \boldsymbol{W}^{-1} \mathbf{1}_{i} .} \\
& =\max _{i \in[m]} \sqrt{\frac{\sigma_{i}\left(\frac{1}{\boldsymbol{w} \vec{\phi}^{\prime \prime}}\right)}{w_{i}}} \leq 2 .
\end{aligned}
$$

For any $\boldsymbol{y}$, we have

$$
\begin{aligned}
\left\|\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}} \boldsymbol{y}\right\|_{w+\infty} & \leq\left\|\boldsymbol{P}_{x, \boldsymbol{w}} \boldsymbol{y}\right\|_{\infty}+C_{\text {norm }}\left\|\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}} \boldsymbol{y}\right\|_{\boldsymbol{w}} \\
& \leq\|\boldsymbol{y}\|_{\infty}+\left\|\overline{\boldsymbol{P}}_{x, w} \boldsymbol{y}\right\|_{\infty}+C_{\text {norm }}\|\boldsymbol{y}\|_{\boldsymbol{w}} \\
& \leq\|\boldsymbol{y}\|_{\infty}+\left(2+C_{\text {norm }}\right)\|\boldsymbol{y}\|_{\boldsymbol{w}} \\
& \leq \frac{C_{\text {norm }}+2}{C_{\text {norm }}}\|\boldsymbol{y}\|_{w+\infty}
\end{aligned}
$$

Hence, we have $c_{\gamma} \leq \frac{C_{\text {norm }}+2}{C_{\text {norm }}}$. Thus, we have picked $C_{\text {norm }}=\frac{18}{\alpha-1}$ and have $c_{\gamma} \leq 1+\frac{\alpha-1}{9}$.

## ■ 6.7.1 Weight computation

Here, we describe how to efficiently compute approximations to the weight function $\boldsymbol{g}: \Omega \rightarrow \mathbb{R}_{>0}^{m}$ as given by (6.30). The two main technical tools we use towards this end are the gradient descent method, Theorem 2.3.12, a standard result in convex optimization, and fast numerical methods for estimating leverage scores using the Johnson-Lindenstrauss Lemma, Theorem 2.3.4, a powerful tool in randomized numerical linear algebra.

Since the weight function, $\boldsymbol{g}$, is defined as the minimizer of a convex optimization problem, we could use the gradient descent method directly to minimize $\hat{f}$ and hence compute $\boldsymbol{g}$. Indeed, in Lemma 6.7.6 we show how applying the gradient descent method in a carefully scaled space allows us to compute $\boldsymbol{g}(\boldsymbol{x})$ to high accuracy in $\widetilde{O}(1)$ iterations. Unfortunately, this result makes two assumptions to compute $\boldsymbol{g}(\boldsymbol{x}):(1)$ we are given a weight $\boldsymbol{w} \in \Omega$ that is not too far from $\boldsymbol{g}(\boldsymbol{x})$ and (2) we compute the gradient
of $\hat{f}$ exactly.
Assumption (1) is not an issue as we always ensure that $\boldsymbol{g}$ does not change too much between calls to compute $\boldsymbol{g}$ and therefore can always use our previous weights as the approximation to $\boldsymbol{g}(\boldsymbol{x})$. However, naively computing the gradient of $\hat{f}$ is computationally expensive and hence assumption (2) is problematic. To deal with this issue we use the fact that by careful application of Johnson-Lindenstrauss one can compute a multiplicative approximation to the gradient efficiently and in Theorem 6.7.7 we show that this suffices to compute an approximation to $\boldsymbol{g}$ that suffices to use in our weighted path following scheme.

Recall from Theorem 2.3.12 that if we take repeated projected gradient steps then we can achieve linear convergence up to bounds on how much the hessian of the function changes over the domain of interest. To apply this Theorem 2.3 .12 to compute $\boldsymbol{g}(\boldsymbol{s})$ we first need to show that there is a region around the optimal point $\boldsymbol{g}(\boldsymbol{s})$ such that the Hessian of $\hat{f}$ does not change too much.
Lemma 6.7.5 (Hessian Approximation). For $\left\|\boldsymbol{W}_{(0)}^{-1}\left(\boldsymbol{g}(\boldsymbol{x})-\boldsymbol{w}^{(0)}\right)\right\|_{\infty} \leq \frac{1}{24}$ we have

$$
0.84 \boldsymbol{W}^{-1} \preceq \nabla_{\boldsymbol{w} \boldsymbol{w}}^{2} f_{t}(\boldsymbol{x}, \boldsymbol{w}) \preceq 3.6 \boldsymbol{W}^{-1} .
$$

Proof. From Lemma 6.7.1, we know that

$$
\nabla_{\boldsymbol{w} \boldsymbol{w}}^{2} f_{t}(\boldsymbol{x}, \boldsymbol{w})=\boldsymbol{W}^{-1}(\boldsymbol{\Sigma}+\beta \boldsymbol{I}+\alpha \boldsymbol{\Lambda}) \boldsymbol{W}^{-1}
$$

where $\boldsymbol{\Sigma}=\boldsymbol{\Sigma}_{\boldsymbol{A}_{x}}\left(\boldsymbol{w}^{-\alpha}\right)$ and $\boldsymbol{\Lambda}=\boldsymbol{\Lambda}_{\boldsymbol{A}_{x}}\left(\boldsymbol{w}^{-\alpha}\right)$. Using $\mathbf{0} \preceq \boldsymbol{\Lambda} \preceq \boldsymbol{\Sigma}$, we have

$$
\boldsymbol{W}^{-1}(\boldsymbol{\Sigma}+\beta \boldsymbol{I}) \boldsymbol{W}^{-1} \preceq \nabla_{\boldsymbol{w} \boldsymbol{w}}^{2} f_{t}(\boldsymbol{x}, \boldsymbol{w}) \preceq 3 \boldsymbol{W}^{-1}(\boldsymbol{\Sigma}+\beta \boldsymbol{I}) \boldsymbol{W}^{-1}
$$

Using that $\left\|\boldsymbol{W}_{(0)}^{-1}\left(\boldsymbol{g}(\boldsymbol{x})-\boldsymbol{w}^{(0)}\right)\right\|_{\infty} \leq \frac{1}{24}$, we have

$$
\boldsymbol{\Sigma}+\beta \boldsymbol{I} \preceq\left(1-\frac{1}{24}\right)^{-4} \boldsymbol{\Sigma}_{\boldsymbol{A}_{x}}\left(\boldsymbol{g}^{-\alpha}\right)+\beta \boldsymbol{I} \preceq\left(1-\frac{1}{24}\right)^{-4} \boldsymbol{G} \preceq 1.19 \boldsymbol{W}
$$

and

$$
\boldsymbol{\Sigma}+\beta \boldsymbol{I} \succeq\left(1-\frac{1}{24}\right)^{4} \boldsymbol{\Sigma}_{\boldsymbol{A}_{x}}\left(\boldsymbol{g}^{-\alpha}\right)+\beta \boldsymbol{I} \succeq\left(1-\frac{1}{24}\right)^{4} \boldsymbol{G} \succeq 0.84 \boldsymbol{W}
$$

Combining Theorem 2.3.12 and Lemma 6.7.5, we get the following algorithm to compute the weight function using the exact computation of the gradient of $\hat{f}$. Note that this algorithm applies Theorem 2.3.12 multiple times as in each iteration we are taking a gradient step with respect to a different norm.
Lemma 6.7.6 (Exact Weight Computation). Given $\boldsymbol{w}^{(0)} \in \mathbb{R}_{>0}^{m}$ such that $\left\|\boldsymbol{W}_{(0)}^{-1}\left(\boldsymbol{g}(\boldsymbol{x})-\boldsymbol{w}^{(0)}\right)\right\|_{\infty} \leq \frac{1}{48}$. Let

$$
Q=\left\{\boldsymbol{w} \in \mathbb{R}^{m} \left\lvert\,\left\|\boldsymbol{W}_{(0)}^{-1}\left(\boldsymbol{g}(\boldsymbol{x})-\boldsymbol{w}^{(0)}\right)\right\|_{\infty} \leq \frac{1}{48}\right.\right\} .
$$

For all $k \geq 0$ let

$$
\boldsymbol{w}^{(k+1)}=\underset{\boldsymbol{w} \in Q}{\arg \min }\left\|\boldsymbol{w}-\frac{1}{2}\left(\boldsymbol{w}^{(k)}+\boldsymbol{\sigma}_{\boldsymbol{A}_{\boldsymbol{x}}}\left(\left(\boldsymbol{w}^{(k)}\right)^{-\alpha}\right)+\beta\right)\right\|_{\boldsymbol{W}_{(k)}^{-1}}^{2}
$$

This implies that for all $k$,

$$
\left\|\boldsymbol{g}(\boldsymbol{x})^{-1}\left(\boldsymbol{g}(\boldsymbol{x})-\boldsymbol{w}^{(k)}\right)\right\|_{\infty}^{2} \leq 4 m^{2}\left(1-\frac{1}{10}\right)^{k}
$$

Proof. Note that iterations of Theorem 2.3.12 can be rewritten as

$$
\begin{aligned}
\boldsymbol{w}^{(k+1)} & =\underset{\boldsymbol{w} \in Q}{\arg \min }\left\langle\left(\boldsymbol{I}-\boldsymbol{\Sigma}_{\boldsymbol{A}_{\boldsymbol{x}}}\left(\left(\boldsymbol{w}^{(k)}\right)^{-\alpha}\right) \boldsymbol{W}_{(k)}^{-1}-\beta \boldsymbol{W}_{(k)}^{-1}\right) \mathbf{1}, \boldsymbol{w}\right\rangle+\frac{4}{2}\left\|\boldsymbol{w}-\boldsymbol{w}^{(k)}\right\|_{\boldsymbol{W}_{(k)}^{-1}}^{2} \\
& =\underset{\boldsymbol{w} \in Q}{\arg \min }\left\|\boldsymbol{w}-\frac{1}{4}\left(\boldsymbol{w}^{(k)}+\boldsymbol{\sigma}_{\boldsymbol{A}_{x}}\left(\left(\boldsymbol{w}^{(k)}\right)^{-\alpha}\right)+\beta\right)\right\|_{\boldsymbol{W}_{(k)}^{-1}}^{2}
\end{aligned}
$$

which is the same as in the statement of this lemma. To apply Theorem 2.3.12 we note that for any $\boldsymbol{w} \in Q$ the definition of $Q$ and the fact that $\alpha \in(0,1)$ implies that $\left(1-\frac{1}{48}\right) \boldsymbol{W}_{(0)} \preceq \boldsymbol{W} \preceq\left(1+\frac{1}{48}\right) \boldsymbol{W}_{(0)}$. Therefore Lemma 6.7 .5 shows that for all $\boldsymbol{w}^{(k)} \in Q$,

$$
\begin{equation*}
\frac{4}{5} \boldsymbol{W}_{(k)}^{-1} \preceq 0.84 \boldsymbol{W}_{(0)}^{-1} \preceq \nabla_{\boldsymbol{w} \boldsymbol{w}}^{2} f_{t}(\boldsymbol{x}, \boldsymbol{w}) \preceq 3.6 \boldsymbol{W}_{(0)}^{-1} \preceq 4 \boldsymbol{W}_{(k)}^{-1} . \tag{6.34}
\end{equation*}
$$

Hence, Theorem 2.3.12 and inequality (6.34) shows that

$$
\left\|\boldsymbol{w}^{(k+1)}-\boldsymbol{g}\right\|_{\boldsymbol{W}_{(k)}^{-1}}^{2} \leq\left(1-\frac{1}{5}\right)\left\|\boldsymbol{w}^{(k)}-\boldsymbol{g}\right\|_{\boldsymbol{W}_{(k)}^{-1}}^{2}
$$

Since $\left\|\boldsymbol{W}_{(0)}^{-1}\left(\boldsymbol{g}(\boldsymbol{s})-\boldsymbol{w}^{(0)}\right)\right\|_{\infty} \leq \frac{1}{48}$ and $\boldsymbol{w}^{(k)} \in Q$ we know that $\boldsymbol{G} \succeq\left(1-\frac{1}{48}\right)^{2} \boldsymbol{W}_{(k)}$. Hence, we have

$$
\begin{aligned}
\left\|\boldsymbol{w}^{(k)}-\boldsymbol{g}\right\|_{\boldsymbol{G}^{-1}}^{2} & \leq\left(1-\frac{1}{48}\right)^{-2}\left(1-\frac{1}{5}\right)\left\|\boldsymbol{w}^{(k-1)}-\boldsymbol{g}\right\|_{\boldsymbol{G}^{-1}}^{2} \\
& \leq\left(1-\frac{1}{10}\right)\left\|\boldsymbol{w}^{(k-1)}-\boldsymbol{g}\right\|_{\boldsymbol{G}^{-1}}^{2} \\
& \leq\left(1-\frac{1}{10}\right)^{k}\left\|\boldsymbol{w}^{(0)}-\boldsymbol{g}\right\|_{\boldsymbol{G}^{-1}}^{2}
\end{aligned}
$$

The result follows from the facts that

$$
\left\|\boldsymbol{w}^{(0)}-\boldsymbol{g}\right\|_{\boldsymbol{G}^{-1}}^{2} \leq m\|\boldsymbol{g}\|_{\infty}\left\|\boldsymbol{g}^{-1}\left(\boldsymbol{g}-\boldsymbol{w}^{(0)}\right)\right\|_{\infty}^{2} \leq \frac{m(1+\beta)}{\left(1-\frac{1}{48}\right)^{2}}\left\|\boldsymbol{w}_{(0)}^{-1}\left(\boldsymbol{g}-\boldsymbol{w}^{(0)}\right)\right\|_{\infty}^{2}
$$

and $\left\|\boldsymbol{g}^{-1}\left(\boldsymbol{w}^{(k)}-\boldsymbol{g}\right)\right\|_{\infty}^{2} \leq \beta^{-1}\left\|\boldsymbol{w}^{(k)}-\boldsymbol{g}\right\|_{\boldsymbol{G}^{-1}}^{2}$ where $\beta=\frac{\operatorname{rank}(\boldsymbol{A})}{2 m}$.
Recall that Theorem 2.3.4 showed that we can compute leverage scores, $\boldsymbol{\sigma}_{\boldsymbol{A}_{x}}$, approximately by solving only $\widetilde{O}(1)$ many regression problems. Now, we show that we can modify Lemma 6.7.6 and use Lemma 2.3.4. Our weight computation and the analysis is as follows.

## Algorithm 9: $\boldsymbol{w}=$ computeWeight $\left(\boldsymbol{x}, \boldsymbol{w}^{(0)}, K\right)$

Let $\alpha=1+\frac{1}{\log _{2}\left(\frac{2 m}{\operatorname{rank}(\boldsymbol{A})}\right)}, \beta=\frac{\operatorname{rank}(\boldsymbol{A})}{2 m}, \varepsilon=\frac{K}{40 c_{r} \log \left(\frac{2 m}{K}\right)}$.
Let $Q=\left\{\boldsymbol{w} \in \mathbb{R}^{m} \left\lvert\,\left\|\boldsymbol{W}_{(0)}^{-1}\left(\boldsymbol{w}-\boldsymbol{w}^{(0)}\right)\right\|_{\infty} \leq \frac{1}{48}\right.\right\}$.
for $j=1$ to $k$ where $k=\left\lceil 20 \log \left(\frac{4 m}{K}\right)\right\rceil$ do
Let $\boldsymbol{\sigma}^{(j)}$ be a $(1+\varepsilon)$ approximation of $\boldsymbol{\sigma}_{\boldsymbol{A}}\left(\left(\boldsymbol{w}^{(j)}\right)^{-\alpha}\right)$. (Theorem 2.3.4)
Set $\boldsymbol{w}^{(j)}=\arg \min _{\boldsymbol{w} \in Q}\left\|\boldsymbol{w}-\frac{1}{4}\left(\boldsymbol{w}^{(j-1)}+\boldsymbol{\sigma}^{(j)}+\beta \mathbf{1}\right)\right\|_{\boldsymbol{w}_{(j-1)}^{-1}}^{2}$
end
Output $\boldsymbol{w}^{(j)}$.
Note that the convex set $Q$ is aligned with standard basis and hence the step on $\boldsymbol{w}^{(j)}$ can be
computed by explicit formula (6.36).
Theorem 6.7.7 (Approximate Weight Computation). Let $\boldsymbol{x} \in \Omega,\left\|\boldsymbol{W}_{(0)}^{-1}\left(\boldsymbol{g}(\boldsymbol{x})-\boldsymbol{w}^{(0)}\right)\right\|_{\infty} \leq \frac{1}{48}$ and $K \in(0,1)$. The algorithm computeWeight $\left(\boldsymbol{x}, \boldsymbol{w}^{(0)}, K\right)$ returns $\boldsymbol{w}$ such that

$$
\left\|\boldsymbol{g}(\boldsymbol{x})^{-1}(\boldsymbol{g}(\boldsymbol{x})-\boldsymbol{w})\right\|_{\infty} \leq K
$$

with probability $1-\frac{\tilde{O}(1)}{m}$.
The running time is dominated by the time needed to solve $\tilde{O}\left(\log ^{3}(1 / K) / K^{2}\right)$ linear systems.
Proof. Consider an execution of computeWeight $\left(\boldsymbol{x}, \boldsymbol{w}^{(0)}, K\right)$ where $\boldsymbol{\sigma}^{(j)}=\boldsymbol{\sigma}_{\boldsymbol{A}_{s}}\left(\left(\boldsymbol{w}^{(j)}\right)^{-\alpha}\right)$ exactly, i.e. $\quad \boldsymbol{\sigma}^{(j)}=\boldsymbol{\sigma}_{\boldsymbol{A}_{s}}\left((\boldsymbol{w})^{-\alpha}\right)$, and let $\boldsymbol{v}^{(j)}$ denote the $\boldsymbol{w}$ computed during this idealized execution of computeWeight.

Now suppose that for all $i \in[m]$ we have

$$
\begin{equation*}
(1-\varepsilon)^{M} \boldsymbol{v}_{i}^{(j)} \leq \boldsymbol{w}_{i}^{(j)} \leq(1+\varepsilon)^{M} \boldsymbol{v}_{i}^{(j)} \tag{6.35}
\end{equation*}
$$

for some $M \geq 0$ and $j \in[k-1]$. Since the objective function and the constraint $Q$ are axis-aligned we can compute $\boldsymbol{w}^{(j)}$ coordinate-wise and we see that

$$
\begin{equation*}
\boldsymbol{w}^{(j+1)}=\operatorname{median}\left(\left(1-\frac{1}{48}\right) \boldsymbol{w}^{(0)}, \frac{1}{4}\left(\boldsymbol{\sigma}_{\boldsymbol{A}_{s}}\left(\left(\boldsymbol{w}^{(j)}\right)^{\alpha}\right)+\beta\right),\left(1+\frac{1}{48}\right) \boldsymbol{w}^{(0)}\right) \tag{6.36}
\end{equation*}
$$

where $[\operatorname{median}(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})]_{i}$ is equal to the median of $x_{i}, y_{i}$ and $z_{i}$ for all $i \in[m]$. By (6.35), (6.36), and the fact that $(1-\varepsilon) \sigma_{\boldsymbol{A}_{x}}\left(\left(\boldsymbol{w}^{(j+1)}\right)^{\alpha}\right)_{i} \leq \boldsymbol{\sigma}_{i}^{(j+1)} \leq(1+\varepsilon) \sigma_{\boldsymbol{A}_{x}}\left(\left(\boldsymbol{w}^{(j+1)}\right)^{\alpha}\right)_{i}$ for all $i \in[m]$, we have that

$$
(1-\varepsilon)^{M+1} \boldsymbol{v}_{i}^{(j+1)} \leq \boldsymbol{w}_{i}^{(j+1)} \leq(1+\varepsilon)^{M+1} \boldsymbol{v}_{i}^{(j+1)}
$$

Since $\boldsymbol{v}^{(0)}=\boldsymbol{w}^{(0)}$ and since $j \in[k-1]$ was arbitrary we can apply induction and we have that for all $j \in[k]$

$$
(1-\varepsilon)^{j} \boldsymbol{v}_{i}^{(j)} \leq \boldsymbol{w}_{i}^{(j)} \leq(1+\varepsilon)^{j} \boldsymbol{v}_{i}^{(j)} .
$$

Note that $k \varepsilon \leq \frac{1}{8}$ and therefore by Taylor series expansion we have $\left\|\boldsymbol{V}_{(k)}^{-1}\left(\boldsymbol{w}^{(k)}-\boldsymbol{v}^{(k)}\right)\right\|_{\infty} \leq \frac{9}{8} \varepsilon k$. Furthermore since $\boldsymbol{v}^{(k)} \in Q$ we know that $\boldsymbol{G}(s) \succeq\left(1-\frac{1}{24}\right) \boldsymbol{V}_{(k)}$. Putting these together, applying Lemma 6.7.6, and recalling that $k=\left\lceil 20 \log \left(\frac{4 m}{K}\right)\right\rceil$ we have

$$
\begin{aligned}
\left\|\boldsymbol{g}(\boldsymbol{x})^{-1}\left(\boldsymbol{g}(\boldsymbol{x})-\boldsymbol{w}^{(k)}\right)\right\|_{\infty} & \leq\left\|\boldsymbol{g}(\boldsymbol{x})^{-1}\left(\boldsymbol{g}(\boldsymbol{x})-\boldsymbol{v}^{(k)}\right)\right\|_{\infty}+\left\|\boldsymbol{g}(\boldsymbol{x})^{-1}\left(\boldsymbol{v}^{(k)}-\boldsymbol{w}^{(k)}\right)\right\|_{\infty} \\
& \leq 2 m\left(1-\frac{1}{10}\right)^{\frac{k}{2}}+\left(1-\frac{1}{24}\right)^{-1}\left\|\boldsymbol{V}_{(k)}^{-1}\left(\boldsymbol{v}^{(k)}-\boldsymbol{w}^{(k)}\right)\right\|_{\infty} \\
& \leq 2 m \cdot \exp \left(-\frac{k}{20}\right)+1.5 k \varepsilon \\
& \leq \frac{K}{2}+1.5 \varepsilon\left\lceil 20 \log \left(\frac{4 m}{K}\right)\right\rceil \leq K
\end{aligned}
$$

Finally, we show how to compute an initial weight without having an approximate weight to help the computation. The algorithm computeInitialWeight $(\boldsymbol{x}, K)$ computes an initial weight in $\tilde{O}(\sqrt{\operatorname{rank} \boldsymbol{A}})$ iterations of computeWeight by computing $\boldsymbol{g}$ for a large enough value of $\beta$ that the
calculation is trivial and then decreasing $\beta$ gradually.

```
Algorithm 10: \(\boldsymbol{w}=\) computeInitialWeight \((s, K)\)
    Let \(\alpha=1+\frac{1}{\log _{2}\left(\frac{2 m}{\operatorname{rank}(\boldsymbol{A})}\right)}, \beta=48\) and \(\boldsymbol{w}=\beta \mathbf{1}\).
    while \(\beta>\frac{\operatorname{rank}(\boldsymbol{A})}{2 m}\) do
            Set \(\boldsymbol{w}=\) computeWeight \(\left(\boldsymbol{x}, \boldsymbol{w}, \frac{1}{100}\right)\).
            Set \(\beta=\max \left\{\left(1-\frac{1}{1000 \sqrt{\operatorname{rank}(\boldsymbol{A})}}\right) \beta, \frac{\operatorname{rank}(\boldsymbol{A})}{2 m}\right\}\).
    end
    Output computeWeight \((\boldsymbol{x}, \boldsymbol{w}, K)\).
```

Theorem 6.7.8 (Computating Initial Weights). For $\boldsymbol{x} \in \Omega$ and $K>0$, with constant probability the algorithm computeInitialWeight $(\boldsymbol{x}, K)$ returns $\boldsymbol{w} \in \mathbb{R}_{>0}^{m}$ such that

$$
\left\|\boldsymbol{g}(\boldsymbol{x})^{-1}(\boldsymbol{g}(\boldsymbol{x})-\boldsymbol{w})\right\|_{\infty} \leq K
$$

The total running time of computeInitialWeight $(\boldsymbol{x}, K)$ is dominated by the time needed to solve $\tilde{O}\left(\sqrt{\operatorname{rank}(\boldsymbol{A})} \log ^{3}(1 / K) / K^{2}\right)$ linear systems.

Proof. Fix $\boldsymbol{x} \in \Omega$. Let $\boldsymbol{g}: \Omega \rightarrow \mathbb{R}^{m}$ be defined by ${ }^{7}$

$$
\boldsymbol{g}(\beta) \stackrel{\text { def }}{=} \underset{\boldsymbol{w} \in \mathbb{R}_{>0}^{m}}{\arg \min } \mathbf{1}^{T} \boldsymbol{w}+\frac{1}{\alpha} \log \operatorname{det}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W}^{-\alpha} \boldsymbol{A}_{x}\right)-\beta \sum_{i \in[m]} \log w_{i}
$$

We first prove by induction that at the beginning of each while iteration, we have that

$$
\begin{equation*}
\left\|\boldsymbol{W}^{-1}(\boldsymbol{g}(\beta)-\boldsymbol{w})\right\|_{\infty} \leq \frac{1}{48} \tag{6.37}
\end{equation*}
$$

First, we prove that (6.37) holds at the first iteration. Since $\boldsymbol{g}(\beta)=\boldsymbol{\sigma}(\beta)+\beta$ where $\boldsymbol{\sigma}(\beta) \stackrel{\text { def }}{=}$ $\boldsymbol{\sigma}_{\boldsymbol{A}_{x}}\left(\boldsymbol{g}^{-\alpha}(\beta)\right)$, we have that for all $i \in[m]$

$$
\beta \leq g(\beta)_{i} \leq 1+\beta
$$

Therefore, the initial weight, $\boldsymbol{w}=\beta \mathbf{1} \in \mathbb{R}_{>0}^{m}$ satisfies the invariant (6.37).
Now, we assume (6.37) holds at the beginning of a certain iteration. After the step of computeWeight, by Theorem 6.7.7 we have

$$
\begin{equation*}
\left\|\boldsymbol{G}(\beta)^{-1}(\boldsymbol{g}(\beta)-\boldsymbol{w})\right\|_{\infty} \leq \frac{1}{100} \tag{6.38}
\end{equation*}
$$

Therefore, it suffices to prove that $\boldsymbol{g}(\beta)$ is close to $\boldsymbol{g}(\beta-\theta)$ for small $\theta$.
To bound how much $\boldsymbol{g}(\beta)$ changes for small changes in $\beta$ we proceed similarly to Lemma 6.7.2. First by the implicit function theorem (Lemma 2.3.9) and direct calculation we know that

$$
\begin{equation*}
\frac{d \boldsymbol{g}}{d \beta}=-\left(\boldsymbol{J}_{\boldsymbol{w}}\left(\nabla_{w} \hat{f}(\boldsymbol{x}, \boldsymbol{w})\right)\right)^{-1}\left(\boldsymbol{J}_{\beta}\left(\nabla_{\boldsymbol{w}} \hat{f}(\boldsymbol{x}, \boldsymbol{w})\right)\right)=\boldsymbol{G}(\beta)\left(\boldsymbol{G}(\beta)+\alpha \boldsymbol{\Lambda}_{g}\right)^{-1} \mathbf{1} \tag{6.39}
\end{equation*}
$$

where $\boldsymbol{\Lambda}_{g} \stackrel{\text { def }}{=} \boldsymbol{\Lambda}_{\boldsymbol{A}_{x}}\left(\boldsymbol{G}(\beta)^{-\alpha} \mathbf{1}\right)$. Next to estimate how fast $\boldsymbol{g}$ can change as a function of $\beta$ we estimate (6.39) in a similar manner to Lemma 6.7.3. Note that

$$
\boldsymbol{G}(\beta)+\alpha \boldsymbol{\Lambda}_{g} \succeq \boldsymbol{\Sigma}(\beta)
$$

[^11]where $\boldsymbol{\Sigma}(\beta) \stackrel{\text { def }}{=} \boldsymbol{\Sigma}_{\boldsymbol{A}_{s}}\left(\boldsymbol{g}^{\alpha}(\beta)\right)$. Consequently,
\[

$$
\begin{align*}
\left\|\boldsymbol{G}(\beta)^{-1} \frac{d \boldsymbol{g}}{d \beta}\right\|_{\boldsymbol{\Sigma}(\beta)}^{2} & \leq\left\|\left(\boldsymbol{G}(\beta)+\alpha \boldsymbol{\Lambda}_{g}\right)^{-1} \mathbf{1}\right\|_{\boldsymbol{\Sigma}(\beta)}^{2} \\
& \leq\|\mathbf{1}\|_{\boldsymbol{\Sigma}(\beta)}^{2}=\operatorname{rank}(\boldsymbol{A}) \tag{6.40}
\end{align*}
$$
\]

Using this estimate of how much $\boldsymbol{g}$ changes in the $\boldsymbol{\Sigma}(\beta)$ norm, we now estimate how much $\boldsymbol{g}$ changes in the $\ell_{\infty}$ norm. Let $\boldsymbol{z} \stackrel{\text { def }}{=}\left(\boldsymbol{G}(\beta)+\alpha \boldsymbol{\Lambda}_{g}\right)^{-1} \mathbf{1}$. Then, we have

$$
\begin{aligned}
\left((1+\alpha) \boldsymbol{\sigma}_{i}(\beta)+\beta\right)\left|\boldsymbol{z}_{i}\right| & \leq\left|\alpha \mathbf{1}_{i}^{T} \boldsymbol{P}^{(2)} \boldsymbol{z}\right|+1 \\
& \leq \alpha \boldsymbol{\sigma}_{i}(\beta)\|\boldsymbol{z}\|_{\boldsymbol{\Sigma}(\beta)}+1
\end{aligned}
$$

Using (6.40) and $\alpha<1$, we have

$$
\left\|\frac{d \ln \boldsymbol{g}}{d \beta}\right\|_{\infty}=\|\boldsymbol{z}\|_{\infty} \leq \max \left(\frac{\alpha\|\boldsymbol{z}\|_{\boldsymbol{\Sigma}(\beta)}}{1+\alpha}, \frac{1}{\beta}\right) \leq \max \left(\sqrt{\operatorname{rank}(\boldsymbol{A})}, \frac{1}{\beta}\right)
$$

Using (6.38), we have that

$$
\left\|\boldsymbol{G}(\beta-\theta)^{-1}(\boldsymbol{g}(\beta-\theta)-\boldsymbol{w})\right\|_{\infty} \leq \frac{1}{48}
$$

for $\theta \leq \frac{\beta}{1000 \sqrt{\operatorname{rank}(\boldsymbol{A})}}$.
Hence, this proves the while loop preserves the invariant (6.37). Therefore, the assumptions needed for Theorem 6.7.7 throughout and computeWeight works as desired. Since each iteration $\beta$ decreased by $\tilde{O}(1 / \sqrt{\operatorname{rank}(\boldsymbol{A})})$ portion and the initial $\beta$ is $\tilde{O}(1)$ we see that the algorithm requires only $\tilde{O}(\sqrt{\operatorname{rank}(\boldsymbol{A})})$ iterations. Using Theorem 6.7 .7 to bound the total number of linear systems solved then yields the result.

## ■ 6.8 The Algorithm

Here we show how to use the results of previous sections to solve (6.9) using exact linear system solver. In the next section we will discuss how to relax this assumption. The central goal of this section is to develop an algorithm, LPSolve, for which we can prove the following theorem

Theorem 6.8.1. Suppose we have an interior point $\boldsymbol{x}_{0} \in \Omega$ for the linear program (6.9).Then, the algorithm LPSolve outputs $\boldsymbol{x}$ such that $\boldsymbol{c}^{T} \boldsymbol{x} \leq O P T+\varepsilon$ in $\tilde{O}\left(\sqrt{\operatorname{rank}(\boldsymbol{A})}\left(\mathcal{T}_{w}+\operatorname{nnz}(\boldsymbol{A})\right) \log (U / \varepsilon)\right)$ work and $\tilde{O}\left(\sqrt{\operatorname{rank}(\boldsymbol{A})} \mathcal{T}_{d} \log (U / \varepsilon)\right)$ depth where $U=\max \left(\left\|\frac{\boldsymbol{u}-\boldsymbol{l}}{\boldsymbol{u}-\boldsymbol{x}_{0}}\right\|_{\infty},\left\|\frac{\boldsymbol{u}-\boldsymbol{l}}{\boldsymbol{x}_{0}-\boldsymbol{l}}\right\|_{\infty},\|\boldsymbol{u}-\boldsymbol{l}\|_{\infty},\|\boldsymbol{c}\|_{\infty}\right)$ and $\mathcal{T}_{w}$ and $\mathcal{T}_{d}$ is the work and depth needed to compute $\left(\boldsymbol{A}^{T} \boldsymbol{D} \boldsymbol{A}\right)^{-1} \boldsymbol{q}$ for input positive definite diagonal matrix $\boldsymbol{D}$ and vector $\boldsymbol{q}$.

We break this proof into several parts. First we provide Lemma 6.8.2, and adaptation of a proof from [206, Thm 4.2.7] that allows us to reason about the effects of making progress along the weighted central path. Then we provide Lemma 6.8.3 that we use to bound the distance to the weighted central path in terms of centrality. After that in Lemma 6.8.3, we analyze a subroutine, pathFollowing, for following the weighted central path. Using these lemmas we conclude by describing our LPSolve algorithm and proving Theorem 6.8.1.

Lemma 6.8.2 ([206, Theorem 4.2.7]). Let $x^{*} \in \mathbb{R}^{m}$ denote an optimal solution to (6.9) and $\boldsymbol{x}_{t}=$ $\arg \min f_{t}(\boldsymbol{x}, \boldsymbol{w})$ for some $t>0$ and $\boldsymbol{w} \in \mathbb{R}_{>0}^{m}$. Then the following holds

$$
\boldsymbol{c}^{T} \boldsymbol{x}_{t}(\boldsymbol{w})-\boldsymbol{c}^{T} \boldsymbol{x}^{*} \leq \frac{\|\boldsymbol{w}\|_{1}}{t} .
$$

Proof. By the optimality conditions of (6.9) we know that $\nabla_{x} f_{t}\left(\boldsymbol{x}_{t}(\boldsymbol{w})\right)=t \cdot \boldsymbol{c}+\boldsymbol{w} \vec{\phi}^{\prime}\left(\boldsymbol{x}_{t}(\boldsymbol{w})\right)$ is orthogonal to the kernel of $\boldsymbol{A}^{T}$. Furthermore since $\boldsymbol{x}_{t}(\boldsymbol{w})-\boldsymbol{x}^{*} \in \operatorname{ker}\left(\boldsymbol{A}^{T}\right)$ we have

$$
\left(t \cdot \boldsymbol{c}+\boldsymbol{w} \vec{\phi}^{\prime}\left(\boldsymbol{x}_{t}(\boldsymbol{w})\right)\right)^{T}\left(\boldsymbol{x}_{t}(\boldsymbol{w})-\boldsymbol{x}^{*}\right)=0
$$

Using that $\phi_{i}^{\prime}\left(x_{t}(\boldsymbol{w})_{i}\right) \cdot\left(x_{i}^{*}-x_{t}(\boldsymbol{w})_{i}\right) \leq 1$ by Lemma 6.5.2 then yields

$$
\boldsymbol{c}^{T}\left(\boldsymbol{x}_{t}(\boldsymbol{w})-\boldsymbol{x}^{*}\right)=\frac{1}{t} \sum_{i \in[m]} w_{i} \cdot \phi_{i}^{\prime}\left(x_{t}(\boldsymbol{w})_{i}\right) \cdot\left(x_{i}^{*}-x_{t}(\boldsymbol{w})_{i}\right) \leq \frac{\|\boldsymbol{w}\|_{1}}{t} .
$$

Lemma 6.8.3. For $\delta_{t}\left(\boldsymbol{x}^{(1)}, \boldsymbol{g}\left(\boldsymbol{x}^{(1)}\right)\right) \leq \frac{1}{960 c_{k}^{2} \log (400 m)}$ and $\boldsymbol{x}_{t} \stackrel{\text { def }}{=} \arg \min f_{t}(\boldsymbol{x}, \boldsymbol{w})$ we have

$$
\left\|\sqrt{\vec{\phi}^{\prime \prime}\left(\boldsymbol{x}_{t}\right)}\left(\boldsymbol{x}^{(1)}-\boldsymbol{x}_{t}\right)\right\|_{\infty} \leq 16 c_{\gamma} c_{k} \delta_{t}\left(\boldsymbol{x}^{(1)}, \boldsymbol{g}\left(\boldsymbol{x}^{(1)}\right)\right)
$$

Proof. We use Theorem 6.6 .8 with exact weight computation and start with $\boldsymbol{x}^{(1)}$ and $\boldsymbol{w}^{(1)}=\boldsymbol{g}\left(\boldsymbol{x}^{(1)}\right)$. In each iteration, $\delta_{t}$ is decreased by a factor of $\left(1-\frac{1}{4 c_{k}}\right) \cdot(6.16)$ shows that

$$
\left\|\sqrt{\vec{\phi}^{\prime \prime}\left(\boldsymbol{x}^{(k)}\right)}\left(\boldsymbol{x}^{(k+1)}-\boldsymbol{x}^{(k)}\right)\right\|_{\infty} \leq c_{\gamma} \delta_{t}\left(\boldsymbol{x}^{(k)}, \boldsymbol{w}^{(k)}\right)
$$

The Lemma 6.5 .1 shows that

$$
\begin{aligned}
\left\|\log \left(\phi^{\prime \prime}\left(\boldsymbol{x}^{(k)}\right)\right)-\log \left(\phi^{\prime \prime}\left(\boldsymbol{x}^{(k+1)}\right)\right)\right\|_{\infty} & \leq\left(1-2 c_{\gamma} \delta_{t}\left(\boldsymbol{x}^{(k)}, \boldsymbol{w}^{(k)}\right)\right)^{-1} \\
& \leq e^{4 c_{\gamma} \delta_{t}\left(\boldsymbol{x}^{(k)}, \boldsymbol{w}^{(k)}\right)}
\end{aligned}
$$

Therefore, for any $k$, we have

$$
\left\|\log \left(\phi^{\prime \prime}\left(\boldsymbol{x}^{(1)}\right)\right)-\log \left(\phi^{\prime \prime}\left(\boldsymbol{x}^{(k)}\right)\right)\right\|_{\infty} \leq e^{4 c_{\gamma} \sum \delta_{t}\left(\boldsymbol{x}^{(k)}, \boldsymbol{w}^{(k)}\right)} \leq e^{32 c_{k} c_{\gamma} \delta_{t}\left(\boldsymbol{x}^{(1)}, \boldsymbol{g}\left(\boldsymbol{x}^{(1)}\right)\right)} \leq 2 .
$$

Hence, for any $k$, we have

$$
\begin{aligned}
\left\|\sqrt{\vec{\phi}^{\prime \prime}\left(\boldsymbol{x}_{t}\right)}\left(\boldsymbol{x}^{(1)}-\boldsymbol{x}^{(k)}\right)\right\|_{\infty} & \leq \sum 2 c_{\gamma} \delta_{t}\left(\boldsymbol{x}^{(k)}, \boldsymbol{w}^{(k)}\right) \\
& \leq 16 c_{\gamma} c_{k} \delta_{t}\left(\boldsymbol{x}^{(1)}, \boldsymbol{g}\left(\boldsymbol{x}^{(1)}\right)\right) .
\end{aligned}
$$

It is clear now $\boldsymbol{x}^{(k)}$ forms a Cauchy sequence and converges to $\boldsymbol{x}_{t}$ because $\delta_{t}$ continuous and $\boldsymbol{x}_{t}$ is the unique point such that $\delta_{t}=0$.

Next, we put together the results of Section 6.6 and analyze the following algorithm for following
the weighted central path.

```
Algorithm 11: \(\left(\boldsymbol{x}^{(\text {final })}, \boldsymbol{w}^{(\text {new })}\right)=\operatorname{pathFollowing}\left(\boldsymbol{x}, \boldsymbol{w}, t_{\text {start }}, t_{\text {end }}, \varepsilon\right)\)
    Let \(c_{k}=9 \log _{2}\left(\frac{2 m}{\operatorname{rank}(\boldsymbol{A})}\right), t=t_{\text {start }}, K=\frac{1}{20 c_{k}}\).
    while \(\left(t<t_{\text {end }}\right.\) if \(\left.t_{\text {start }}<t_{\text {end }}\right)\) or \(\left(t>t_{\text {end }}\right.\) if \(\left.t_{\text {start }}>t_{\text {end }}\right)\) do
        Set \(\left(\boldsymbol{x}^{(\text {new })}, \boldsymbol{w}^{(\text {new })}\right)=\) centeringInexact \((\boldsymbol{x}, \boldsymbol{w}, K)\) where it uses the function
            computeWeight to find the approximation of \(\boldsymbol{g}(\boldsymbol{x})\).
        Set \(t \leftarrow t\left(1 \pm \frac{1}{10^{5} c_{k}^{4} \log (400 \mathrm{~m}) \sqrt{\operatorname{rank}(\boldsymbol{A})}}\right)\) where the sign of \(\pm\) is the sign of \(t_{\text {end }}-t_{\text {start }}\).
        Set \(\boldsymbol{x} \leftarrow \boldsymbol{x}^{(\text {final })}\) and \(\boldsymbol{w} \leftarrow \boldsymbol{w}^{(\text {new })}\).
    end
    for \(i=1, \cdots, 4 c_{k} \log (1 / \varepsilon)\) do
        Set \((\boldsymbol{x}, \boldsymbol{w})=\) centeringInexact \((\boldsymbol{x}, \boldsymbol{w}, K)\) where it uses the function computeWeight to find
        the approximation of \(\boldsymbol{g}(\boldsymbol{x})\).
    end
    Output \((\boldsymbol{x}, \boldsymbol{w})\).
```

Theorem 6.8.4. Suppose that

$$
\delta_{t_{s t a r t}}(\boldsymbol{x}, \boldsymbol{w}) \leq \frac{1}{960 c_{k}^{2} \log (400 m)} \quad \text { and } \quad \Phi_{\mu}(\boldsymbol{\Psi}(\boldsymbol{x}, \boldsymbol{w})) \leq(400 m)^{2}
$$

where $\mu=2 \log (400 m) / K$. Let $\left(\boldsymbol{x}^{(\text {final })}, \boldsymbol{w}^{(\text {new })}\right)=\operatorname{pathFollowing~}\left(\boldsymbol{x}, \boldsymbol{w}, t_{\text {start }}, t_{\text {end }}\right)$, then

$$
\delta_{t_{\text {end }}}\left(\boldsymbol{x}^{(\text {final })}, \boldsymbol{w}^{(\text {new })}\right) \leq \varepsilon \quad \text { and } \quad \Phi_{\mu}\left(\boldsymbol{\Psi}\left(\boldsymbol{x}^{(\text {final })}, \boldsymbol{w}^{(\text {new })}\right)\right) \leq(400 m)^{2} .
$$

Furthermore, pathFollowing $\left(\boldsymbol{x}, \boldsymbol{w}, t_{\text {start }}, t_{\text {end }}\right)$ takes time

$$
\tilde{O}\left(\sqrt{\operatorname{rank}(\boldsymbol{A})}\left(\left|\log \left(\frac{t_{\text {end }}}{t_{\text {start }}}\right)\right|+\log (1 / \varepsilon)\right)(\mathcal{T}+m)\right)
$$

where $\mathcal{T}$ is the time needed to solve on linear system.
Proof. We use induction to prove the invariant that $\delta_{t}(\boldsymbol{x}, \boldsymbol{w}) \leq \frac{1}{960 c_{k}^{2} \log (400 \mathrm{~m})}$ and $\Phi_{\alpha}(\boldsymbol{\Psi}(\boldsymbol{x}, \boldsymbol{w})) \leq$ $(400 \mathrm{~m})^{2}$ on the beginning of while iteration. The initial condition holds by assumption.

Now, assume the invariant holds at a certain iteration. Theorem 6.6 .8 shows that

$$
\begin{equation*}
\left\|\log \left(\boldsymbol{g}\left(\boldsymbol{x}^{(\mathrm{new})}\right)\right)-\log (\boldsymbol{w})\right\|_{\infty} \leq K \leq \frac{1}{20 c_{k}} \tag{6.41}
\end{equation*}
$$

Thus, the weight satisfies the condition of Theorem 6.7.7 and the algorithm centeringInexact can use the function computeWeight to find the approximation of $\boldsymbol{g}\left(\boldsymbol{x}^{(\text {new })}\right)$. Consequently,

$$
\delta_{t}\left(\boldsymbol{x}^{(\mathrm{final})}, \boldsymbol{w}^{(\mathrm{new})}\right) \leq\left(1-\frac{1}{4 c_{k}}\right) \delta_{t} \quad \text { and } \quad \Phi_{\alpha}\left(\boldsymbol{\Psi}\left(\boldsymbol{x}^{(\mathrm{final})}, \boldsymbol{w}^{(\mathrm{new})}\right)\right) \leq(400 m)^{2}
$$

Using Lemma 6.6.1, (6.41) and Theorem 6.7.4, we have

$$
\delta_{t}\left(\boldsymbol{x}^{(\mathrm{final})}, \boldsymbol{w}^{(\mathrm{new})}\right) \leq \frac{1}{960 c_{k}^{2} \log (400 m)}
$$

Hence, we proved that the invariant holds after one iteration. The $\delta_{t}<\varepsilon$ bounds follows from the last loop.

Algorithm 12: $\boldsymbol{x}^{\text {(final) }}=$ LPSolve $(\boldsymbol{x}, \varepsilon)$
Let $\beta=\frac{\operatorname{rank}(\boldsymbol{A})}{2 m}, \boldsymbol{w}=\operatorname{computeInitialWeight}\left(\boldsymbol{x}, \frac{1}{10^{5} \log ^{5}(400 m)}\right), d=-\boldsymbol{w}_{i} \phi_{i}^{\prime}(\boldsymbol{x})$.
Let $t_{1}=\left(10^{10} U^{2} m^{3}\right)^{-1}, t_{2}=3 \mathrm{~m} / \varepsilon, \varepsilon_{1}=\frac{1}{2000 c_{k}^{2} \log (400 m)}, \varepsilon_{2}=\frac{\varepsilon}{100^{3} m^{3} U^{2}}$.
Set $\left(\boldsymbol{x}^{(\text {new })}, \boldsymbol{w}^{(\text {new })}\right)=\operatorname{pathFollowing}\left(\boldsymbol{x}, \boldsymbol{w}, 1, t_{1}, \varepsilon_{1}\right)$ with cost vector $\boldsymbol{d}$.
$\operatorname{Set}\left(\boldsymbol{x}^{(\text {final })}, \boldsymbol{w}^{(\mathrm{final})}\right)=\operatorname{pathFollowing}\left(\boldsymbol{x}^{(\mathrm{new})}, \boldsymbol{w}^{(\text {new })}, t_{1}, t_{2}, \varepsilon_{2}\right)$ with cost vector $\boldsymbol{c}$.
Output $\boldsymbol{x}^{(\text {final })}$.

Proof of Theorem 6.8.1. By Theorem 6.7.8, we know the initial weight satisfies

$$
\left\|\boldsymbol{G}(\boldsymbol{x})^{-1}(\boldsymbol{g}(\boldsymbol{x})-\boldsymbol{w})\right\|_{\infty} \leq \frac{1}{10^{5} \log ^{5}(400 m)}
$$

By the definition of $\boldsymbol{d}$, we have $\boldsymbol{x}$ is the minimum of $\min \boldsymbol{d}^{T} \boldsymbol{x}-\sum \boldsymbol{w}_{i} \phi_{i}(\boldsymbol{x})$ given $\boldsymbol{A}^{T} \boldsymbol{x}=\boldsymbol{b}$. Therefore, $(\boldsymbol{x}, \boldsymbol{w})$ satisfies the assumption of Theorem 6.8 .4 because $\delta_{t}=0$ and $\Phi_{\alpha}$ is small enough. Hence, we have

$$
\delta_{t_{1}}\left(\boldsymbol{x}^{(\mathrm{new})}, \boldsymbol{w}^{(\mathrm{new})}\right) \leq \frac{1}{2000 c_{k}^{2} \log (400 m)} \quad \text { and } \quad \Phi_{\alpha}\left(\boldsymbol{\Psi}\left(\boldsymbol{x}^{(\mathrm{new})}, \boldsymbol{w}^{(\mathrm{new})}\right)\right) \leq(400 m)^{2}
$$

Lemma 6.5 .2 shows that $\left\|\phi_{i}^{\prime}(\boldsymbol{x})\right\|_{\infty} \leq U$ and hence $\|\boldsymbol{c}-\boldsymbol{d}\|_{\infty} \leq 2 U$. Also, Lemma 6.5 .1 shows that $\min _{\boldsymbol{y}} \sqrt{\overrightarrow{\phi^{\prime \prime}}(\boldsymbol{y})} \geq \frac{1}{U}$. Therefore, we have

$$
\begin{aligned}
\delta_{t_{1}}^{\boldsymbol{c}}\left(\boldsymbol{x}^{(\mathrm{new})}, \boldsymbol{w}^{(\mathrm{new})}\right) & =\min _{\boldsymbol{\eta} \in \mathbb{R}^{n}}\left\|\frac{t_{1} \boldsymbol{c}+\boldsymbol{w} \vec{\phi}^{\prime}\left(\boldsymbol{x}^{(\mathrm{new})}\right)-\boldsymbol{A} \boldsymbol{\eta}}{\boldsymbol{w}^{(\mathrm{new})} \sqrt{\vec{\phi}^{\prime \prime}\left(\boldsymbol{x}^{(\mathrm{new})}\right)}}\right\|_{\boldsymbol{w}^{(\mathrm{new})}+\infty} \\
& \leq \min _{\boldsymbol{\eta} \in \mathbb{R}^{n}} \| \frac{t_{1} \boldsymbol{d}+\boldsymbol{w}{\overrightarrow{\phi^{\prime}}\left(\boldsymbol{x}^{(\mathrm{new})}\right)-\boldsymbol{A} \boldsymbol{\eta}}_{\boldsymbol{w}^{(\mathrm{new})} \sqrt{\vec{\phi}^{\prime \prime}\left(\boldsymbol{x}^{(\mathrm{new})}\right)}}\left\|_{\boldsymbol{w}^{(\mathrm{new})}+\infty}+t_{1}\right\| \frac{\boldsymbol{c}-\boldsymbol{d}}{\boldsymbol{w}^{(\mathrm{new})} \sqrt{\vec{\phi}^{\prime \prime}\left(\boldsymbol{x}^{(\mathrm{new})}\right)}} \|_{\boldsymbol{w}^{(\mathrm{new})}+\infty}}{} \\
& \leq \delta_{t_{1}}^{\boldsymbol{d}}\left(\boldsymbol{x}^{(\mathrm{new})}, \boldsymbol{w}^{(\mathrm{new})}\right)+4 U^{2} t_{1}\|\mathbf{1}\|_{\boldsymbol{w}+\infty} \\
& =\delta_{t_{1}}^{\boldsymbol{d}}\left(\boldsymbol{x}^{(\mathrm{new})}, \boldsymbol{w}^{(\mathrm{new})}\right)+100 m U^{2} t_{1}
\end{aligned}
$$

Since we have chosen $t_{1}$ small enough, we have $\delta_{t_{1}}^{c}\left(\boldsymbol{x}^{(n e w)}, \boldsymbol{w}^{(n e w)}\right)$ is small enough to satisfy the assumption of Theorem 6.8.4. So, we only need to prove how large $t_{2}$ should be and how small $\varepsilon_{2}$ should be in order to get $\boldsymbol{x}$ such that $\boldsymbol{c}^{T} \boldsymbol{x} \leq$ OPT $+\varepsilon$. By Lemma 6.8 .2 and $\left\|\boldsymbol{w}^{(\text {final })}\right\| \leq 3 m$, we have

$$
\boldsymbol{c}^{T} \boldsymbol{x}_{t_{2}} \leq \mathrm{OPT}+\frac{3 m}{t_{2}}
$$

Also, Lemma 6.8 .3 shows that we have

$$
\left\|\sqrt{\vec{\phi}^{\prime \prime}\left(\boldsymbol{x}_{t_{2}}\right)}\left(\boldsymbol{x}^{(\mathrm{final})}-\boldsymbol{x}_{t_{2}}\right)\right\|_{\infty} \leq 32 \varepsilon_{2} c_{k}
$$

Using $\min _{\boldsymbol{y}} \sqrt{\vec{\phi}^{\prime \prime}(\boldsymbol{y})} \geq \frac{1}{U}$, we have $\left\|\boldsymbol{x}^{(\text {final })}-\boldsymbol{x}_{t_{2}}\right\|_{\infty} \leq 32 \varepsilon_{2} c_{k} U$ and hence our choice of $t_{2}$ and $\varepsilon_{2}$ gives the result

$$
\boldsymbol{c}^{T} \boldsymbol{x}^{(\text {final })} \leq \mathrm{OPT}+\frac{3 m}{t_{2}}+32 \varepsilon_{2} c_{k} U^{2} \leq \mathrm{OPT}+\varepsilon
$$

## ■ 6.9 Generalized Minimum Cost Flow

In this section we show how to use the interior point method in Section 6.8 to solve the maximum flow problem in time $\widetilde{O}\left(m \sqrt{n} \log ^{O(1)}(U)\right)$, to solve the minimum cost flow problem in time, $\widetilde{O}\left(m \sqrt{n} \log ^{O(1)}(U)\right)$, and to compute $\varepsilon$-approximate solutions to the lossy generalized minimum cost flow problem in time $\widetilde{O}\left(m \sqrt{n} \log ^{O(1)}(U / \varepsilon)\right)$. Our algorithm for the generalized minimum cost flow problem is essentially the same as our algorithm for the simpler specific case of minimum cost flow and maximum flow and therefore, we present the algorithm for the generalized minimum cost flow problem directly.

The generalized minimum cost flow problem [63] is as follows. Let $G=(V, E)$ be a connected directed graph where each edge $e$ has capacity $c_{e}>0$ and multiplier $1 \geq \gamma_{e}>0$. For each edge $e$, there can be only at most $c_{e}$ units of flow on that edge and the flow on that edge must be non-negative. Also, for each unit of flow entering edge $e$, there are only $\gamma_{e}$ units of flow going out. The generalized maximum flow problem is to compute how much flow can be sent into $t$ given a unlimited source $s$. The generalized minimum cost flow is to ask what is the minimum cost of sending the maximum flow given the cost of each edge is $q_{e}$. The maximum flow and the minimum cost flow are the case with $\gamma_{e}=1$ for all edges $e$.

Since the generalized minimum cost flow includes all of these cases, we focus on this general formulation. The problem can be written as the following linear program

$$
\min _{\mathbf{0} \leq \boldsymbol{x} \leq \boldsymbol{c}} \boldsymbol{q}^{T} \boldsymbol{x} \text { such that } \boldsymbol{A} \boldsymbol{x}=F \mathbf{1}_{t}
$$

where $F$ is the generalized maximum flow value, $\mathbf{1}_{t}$ is a indicator vector of size $(n-1)$ that is non-zero at vertices $t$ and $\boldsymbol{A}$ is a $|V \backslash\{s\}| \times|E|$ matrix such that for each edge $e$, we have

$$
\begin{aligned}
\boldsymbol{A}\left(e_{\text {head }}, e\right) & =\gamma(e) \\
\boldsymbol{A}\left(e_{\text {tail }}, e\right) & =-1
\end{aligned}
$$

In order words, the constraint $\boldsymbol{A} x=F \mathbf{1}_{t}$ requires the flow to satisfies the flow conversation at all vertices except $s$ and $t$ and requires it flows $F$ unit of flow to $t$. We assume $c_{e}$ are integer and $\gamma_{e}$ is a rational number. Let $U$ be the maximum of $c_{e}, q_{e}$, the numerator of $\gamma_{e}$ and the denominator of $\gamma_{e}$. For the generalized flow problems, getting an efficient exact algorithm is difficult and we aim for approximation algorithms only.

Definition 6.9.1. We call a flow an $\varepsilon$-approximate generalized maximum flow if it is a flow satisfies the flow conservation and the flow value is larger than maximum flow value minus $\varepsilon$. We call a flow is an $\varepsilon$-approximate generalized minimum cost maximum flow if it is an $\varepsilon$-approximate maximum flow and has cost not greater than the minimum cost maximum flow value.

Note that $\operatorname{rank}(\boldsymbol{A})=n-1$ because the graph is connected and hence our algorithm takes only $\tilde{O}(\sqrt{n} L)$ iterations. Therefore, the problems remaining are to compute $L$ and bound how much time is required to solve the linear systems involved. However, $L$ is large in the most general setting and hence we cannot use the standard theory to say how to get the initial point, how to round to the vertex. Furthermore, the condition number of $\boldsymbol{A}^{T} \boldsymbol{A}$ can be very bad.

In [63], they used dual path following to solve the generalized minimum cost flow problem with the caveats that the dual polytope is not bounded, the problem of getting the initial flow, the problem of rounding it to the a feasible flow. We use there analysis to formulate the problem in a manner more amenable to our algorithms. Since we are doing the primal path following, we will state a reformulation of the LP slightly different.

Theorem 6.9.2 ([63]). Given a directed graph $G$. We can find a new directed graph $\tilde{G}$ with $O(m)$ edges and $O(n)$ vertices in $\tilde{O}(m)$ time such that the modified linear program

$$
\min _{0 \leq x_{i} \leq c_{i}, 0 \leq y_{i} \leq 4 m U^{2}, 0 \leq z_{i} \leq 4 m U^{2}} \boldsymbol{q}^{T} \boldsymbol{x}+\frac{256 m^{5} U^{5}}{\varepsilon^{2}}\left(\mathbf{1}^{T} \boldsymbol{y}+\mathbf{1}^{T} \boldsymbol{z}\right) \text { such that } \boldsymbol{A} \boldsymbol{x}+\boldsymbol{y}-\boldsymbol{z}=F \mathbf{1}_{t}
$$

satisfies the following conditions:

1. $\boldsymbol{x}=\frac{c}{2} \mathbf{1}, \boldsymbol{y}=2 m U^{2} \mathbf{1}-\left(\boldsymbol{A} \frac{c}{2} \mathbf{1}\right)^{-}+F \overrightarrow{1_{t}}, z=2 m U^{2} \mathbf{1}+\left(\boldsymbol{A} \frac{c}{2} \mathbf{1}\right)^{+}$is an interior point of the linear program.
2. Given any $(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})$ such that $\|\boldsymbol{A} \boldsymbol{x}+\boldsymbol{y}-\boldsymbol{z}\|_{2} \leq \frac{\varepsilon^{2}}{128 m^{2} n^{2} U^{3}}$ and with cost value within $\frac{\varepsilon^{2}}{128 m^{2} n^{2} U^{3}}$ of the optimum. Then, one can compute an $\varepsilon$-approximate minimum cost maximum flow for graph $G$ in time $\tilde{O}(m)$.
3. The linear system of the linear program is well-conditioned, i.e., the condition number of

$$
\left[\begin{array}{lll}
\boldsymbol{A} & \boldsymbol{I} & -\boldsymbol{I}
\end{array}\right]\left[\begin{array}{c}
\boldsymbol{A}^{T} \\
\boldsymbol{I} \\
-\boldsymbol{I}
\end{array}\right]
$$

is $O(m U)$.
4. The linear system of the linear program can be solve in nearly linear time, i.e. for any diagonal matrix $\boldsymbol{S}$ with condition number $\kappa$ and vector $b$, it takes $\tilde{O}\left(m \log \left(\frac{\kappa U}{\delta}\right)\right)$ time to find $x$ such that

$$
\left\|x-\boldsymbol{L}^{-1} b\right\|_{L} \leq \delta\|x\|_{L}
$$

where $\boldsymbol{L}=\left[\begin{array}{lll}\boldsymbol{A} & \boldsymbol{I} & -\boldsymbol{I}\end{array}\right] \boldsymbol{S}\left[\begin{array}{c}\boldsymbol{A}^{T} \\ \boldsymbol{I} \\ -\boldsymbol{I}\end{array}\right]$.
The main difference between what stated in [63] and here is that

1. Our linear program solver can support constraint $l_{i} \leq x_{i} \leq u_{i}$ and hence we do not need to split the flow variable to positive part and negative part.
2. Our linear program solver is primal path following and hence we add the constraint $y_{i} \leq 4 m U^{2}$ and $z_{i} \leq 4 m U^{2}$. Since the maximum flow value is at most $m U^{2}$, it does not affect the optimal solution of the linear program.
3. We remove the variable $\mathbf{x}_{3}$ in [63] because the purpose of that is to make the dual polytope is bounded and we do not need it here.

Using the reduction mentioned above, one can obtain the promised generalized minimum cost flow algorithm.

Theorem 6.9.3. There is a randomized algorithm to compute an $\varepsilon$-approximate generalized minimum cost maximum flow in $\tilde{O}\left(\sqrt{n} \log ^{O(1)}(U / \varepsilon)\right)$ depth $\widetilde{O}\left(m \sqrt{n} \log ^{O(1)}(U / \varepsilon)\right)$ total work (see Definition 6.9.1). Furthermore, there is an algorithm to compute an exact standard minimum cost maximum flow in $\tilde{O}\left(\sqrt{n} \log ^{O(1)}(U)\right)$ depth and $\widetilde{O}\left(m \sqrt{n} \log ^{O(1)}(U)\right)$ total work.

Proof. Using the reduction above and Theorem 6.8.1, we get an algorithm of generalized minimum cost flow by solving $\widetilde{O}(\sqrt{n})$ linear systems to $\tilde{\mathcal{O}}(1)$ bit accuracy and the condition number of those systems
are poly $(m U / \varepsilon)$. In [63], they showed that the linear system involved can be reduced to $\tilde{O}(\log (U / \varepsilon))$ many Laplacian systems and hence we can use a recent nearly linear work polylogarithmic depth Laplacian system solver of Spielman and Peng [223]. In total, it takes $\widetilde{O}\left(m \log ^{O(1)}\left(\frac{U}{\varepsilon}\right)\right)$ time to solve each systems.

For the standard minimum cost maximum flow problem, it is known that the solution set is a convex polytope with integer coordinates and we can use Isolation lemma to make sure there is unique minimum. Hence, we only need to take $\varepsilon=\operatorname{poly}(1 / m U)$ and round the solution to the closest integer. See Section 3.5 in [63] for details.

### 6.10 Glossary

Here we summarize various linear programming specific notation that we use throughout the chapter. For many quantities we included the typical order of magnitude as they appear during our algorithms.

- Linear program related: constraint matrix $\boldsymbol{A} \in \mathbb{R}^{m \times n}$, cost vector $\boldsymbol{c} \in \mathbb{R}^{m}$, constraint vector $\boldsymbol{b} \in \mathbb{R}^{n}$, solution $\boldsymbol{x} \in \mathbb{R}^{m}$, weights of constraints $\boldsymbol{w} \in \mathbb{R}^{m}$ where $m$ is the number of variables and $n$ is the number of constraints.
- Matrix version of variables: $\boldsymbol{S}$ is the diagonal matrix corresponds to $s, W$ corresponds to $\boldsymbol{w}, \boldsymbol{\Phi}$ corresponds to $\phi$.
- Penalized objective function (6.12): $f_{t}(\boldsymbol{x}, \boldsymbol{w})=t \cdot \boldsymbol{c}^{T} \boldsymbol{x}+\sum_{i \in[m]} \boldsymbol{w}_{i} \phi_{i}\left(\boldsymbol{x}_{i}\right)$.
- Barrier functions (Sec 6.5.1): For $[l, \infty)$, we use $\phi(x)=-\log (x-l)$. For $(-\infty, u]$, we use $\phi(x)=-\log (u-x)$. For $[l, u]$, we use $\phi(x)=-\log (a x+b)$ where $a=\frac{\pi}{u-l}$ and $b=-\frac{\pi}{2} \frac{u+l}{u-l}$.
- The projection matrix $\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}}$ (6.14): $\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}}=\boldsymbol{I}-\boldsymbol{W}^{-1} \boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W}^{-1} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}_{x}^{T}$ where $\boldsymbol{A}_{\boldsymbol{x}} \stackrel{\text { def }}{=}$ $\boldsymbol{\Phi}^{\prime \prime}(\boldsymbol{x})^{-1 / 2} \boldsymbol{A}$.
- Newton step (6.13): $\boldsymbol{h}_{t}(\boldsymbol{x}, \boldsymbol{w})=-\boldsymbol{\Phi}^{\prime \prime}(\boldsymbol{x})^{-1 / 2} \boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}} \boldsymbol{W}^{-1} \boldsymbol{\Phi}^{\prime \prime}(\boldsymbol{x})^{-1 / 2} \nabla_{x} f_{t}(\boldsymbol{x}, \boldsymbol{w})$.
- The mixed norm (6.15): $\|\boldsymbol{y}\|_{\boldsymbol{w}+\infty}=\|\boldsymbol{y}\|_{\infty}+C_{\text {norm }}\|\boldsymbol{y}\|_{\boldsymbol{W}}$ where $C_{\text {norm }} \approx \operatorname{polylog}(m)$.
- Centrality (6.17): $\delta_{t}(\boldsymbol{x}, \boldsymbol{w})=\min _{\boldsymbol{\eta} \in \mathbb{R}^{n}}\left\|\frac{\nabla_{x} f_{t}(\boldsymbol{x}, \boldsymbol{w})-\boldsymbol{A} \eta}{\boldsymbol{w} \sqrt{\phi^{\prime \prime}(\boldsymbol{x})}}\right\|_{\boldsymbol{w}+\infty} \approx \frac{1}{\operatorname{polylog}(m)}$.
- Properties of weight function (Def 6.5.4): size $c_{1}(\boldsymbol{g})=\|\boldsymbol{g}(\boldsymbol{x})\|_{1} \approx \operatorname{rank}(\boldsymbol{A})$, slack sensitivity $c_{\gamma}(\boldsymbol{g})=\left\|\boldsymbol{P}_{x, \boldsymbol{w}}\right\|_{\boldsymbol{w}+\infty} \approx 1+\frac{1}{\operatorname{polylog}(m)}$, step consistency $c_{\delta}(\boldsymbol{g}) \approx 1-\frac{1}{\operatorname{polylog}(m)}$.
- Difference between $\boldsymbol{g}$ and $\boldsymbol{w}(6.29): \mathbf{\Psi}(\boldsymbol{x}, \boldsymbol{w})=\log (\boldsymbol{g}(\boldsymbol{x}))-\log (\boldsymbol{w})$.
- Potential function for tracing 0 (Thm 6.6.6): $\Phi_{\mu}(\boldsymbol{x})=e^{\mu x}+e^{-\mu x} \approx \operatorname{poly}(m)$.
- The weight function proposed (6.30):

$$
\boldsymbol{g}(\boldsymbol{x})=\underset{\boldsymbol{w} \in \mathbb{R}_{>0}^{m}}{\arg \min } \hat{f}(\boldsymbol{x}, \boldsymbol{w}) \quad \text { where } \quad \hat{f}(\boldsymbol{x}, \boldsymbol{w})=\mathbf{1}^{T} \boldsymbol{w}+\frac{1}{\alpha} \log \operatorname{det}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W}^{-\alpha} \boldsymbol{A}_{x}\right)-\beta \sum_{i} \log w_{i}
$$

where $\boldsymbol{A}_{x}=\left(\boldsymbol{\Phi}^{\prime \prime}(\boldsymbol{x})\right)^{-1 / 2} \boldsymbol{A}, \alpha \approx 1+1 / \log _{2}\left(\frac{m}{\operatorname{rank}(\boldsymbol{A})}\right), \beta \approx \operatorname{rank}(\boldsymbol{A}) / m$.

## ■ 6.11 Appendix: Technical Lemmas

Lemma 6.11.1. For any norm $\|\cdot\|$ and $\|\boldsymbol{y}\|_{Q} \xlongequal{\text { def }} \min _{\boldsymbol{\eta} \in \mathbb{R}^{n}}\left\|\boldsymbol{y}-\frac{\boldsymbol{A} \eta}{\left.\boldsymbol{w} \sqrt{\bar{\phi}^{\prime \prime}(\boldsymbol{x}}\right)}\right\|$, we have

$$
\|\boldsymbol{y}\|_{Q} \leq\left\|\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}} \boldsymbol{y}\right\| \leq\left\|\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}}\right\| \cdot\|\boldsymbol{y}\|_{Q} .
$$

Proof. By definition $\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}} \boldsymbol{y}=\boldsymbol{y}-\frac{A \eta_{y}}{\boldsymbol{w} \sqrt{\bar{\phi}^{\prime \prime}(\boldsymbol{x})}}$ for some $\boldsymbol{\eta}_{y} \in \mathbb{R}^{n}$. Consequently,

$$
\|\boldsymbol{y}\|_{Q}=\min _{\eta \in \mathbb{R}^{n}}\left\|\boldsymbol{y}-\frac{\boldsymbol{A} \eta}{\boldsymbol{w} \sqrt{\vec{\phi}^{\prime \prime}(\boldsymbol{x})}}\right\| \leq\left\|\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}} \boldsymbol{y}\right\| .
$$

On the other hand, let $\boldsymbol{\eta}_{q}$ such that $\|\boldsymbol{y}\|_{Q}=\left\|\boldsymbol{y}-\frac{\boldsymbol{A} \boldsymbol{\eta}_{q}}{\boldsymbol{w} \sqrt{\bar{\phi}^{\prime \prime}(\boldsymbol{x})}}\right\|$. Since $\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}} \boldsymbol{W}^{-1}\left(\boldsymbol{\Phi}^{\prime \prime}\right)^{-1 / 2} \boldsymbol{A}=\mathbf{0}$, we have

$$
\left\|\boldsymbol{P}_{x, w} \boldsymbol{y}\right\|=\left\|\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}}\left(\boldsymbol{y}-\frac{\boldsymbol{A} \boldsymbol{\eta}_{q}}{\boldsymbol{w} \sqrt{\overrightarrow{\phi^{\prime \prime}}}}\right)\right\| \leq\left\|\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}}\right\| \cdot\left\|\boldsymbol{y}-\frac{\boldsymbol{A} \boldsymbol{\eta}_{q}}{\boldsymbol{w} \sqrt{\overrightarrow{\phi^{\prime \prime}}}}\right\|=\left\|\boldsymbol{P}_{\boldsymbol{x}, \boldsymbol{w}}\right\| \cdot\|\boldsymbol{y}\|_{Q} .
$$

Lemma 6.11.2. For any projection matrix $\boldsymbol{P} \in \mathbb{R}^{m \times m}, \boldsymbol{\Sigma}=\boldsymbol{\operatorname { D i a g }}(\boldsymbol{P}), i, j \in[m], \boldsymbol{x} \in \mathbb{R}^{m}$, and $\boldsymbol{w} \in \mathbb{R}_{>0}^{m}$ we have

1. $\boldsymbol{\Sigma}_{i i}=\sum_{j \in[m]} \boldsymbol{P}_{i j}^{(2)}$,
2. $\mathbf{0} \preceq \boldsymbol{P}^{(2)} \preceq \boldsymbol{\Sigma} \preceq \boldsymbol{I}$,
3. $\boldsymbol{P}_{i j}^{(2)} \leq \boldsymbol{\Sigma}_{i i} \boldsymbol{\Sigma}_{j j}$,
4. $\left|\mathbf{1}_{i}^{T} \boldsymbol{P}^{(2)} \boldsymbol{x}\right| \leq \boldsymbol{\Sigma}_{i i}\|\boldsymbol{x}\|_{\boldsymbol{\Sigma}}$.
5. $\nabla_{\boldsymbol{w}} \log \operatorname{det}\left(\boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}\right)=\boldsymbol{\Sigma}_{\boldsymbol{A}}(\boldsymbol{w}) \boldsymbol{w}^{-1}$.
6. $\boldsymbol{J}_{\boldsymbol{w}}\left(\boldsymbol{\sigma}_{\boldsymbol{A}}(\boldsymbol{w})\right)=\boldsymbol{\Lambda}_{\boldsymbol{A}}(\boldsymbol{w}) \boldsymbol{W}^{-1}$.

Proof. To prove (1), we simply note that by definition of a projection matrix $\boldsymbol{P}=\boldsymbol{P} \boldsymbol{P}$ and therefore

$$
\boldsymbol{\Sigma}_{i i}=\boldsymbol{P}_{i i}=\mathbf{1}_{i}^{T} \boldsymbol{P} \mathbf{1}_{i}=\mathbf{1}_{i}^{T} \boldsymbol{P} \boldsymbol{P} \mathbf{1}_{i}=\sum_{j \in[n]} \boldsymbol{P}_{i j}^{2}=\sum_{j \in[n]} \boldsymbol{P}_{i j}^{(2)}
$$

To prove (2), we observe that since $\boldsymbol{P}$ is a projection matrix, all its eigenvectors are either 0 or 1. Therefore, $\boldsymbol{\Sigma} \preceq \boldsymbol{I}$ and by (1) $\boldsymbol{\Sigma}-\boldsymbol{P}^{(2)}$ is diagonally dominant. Consequently, $\boldsymbol{\Sigma}-\boldsymbol{P}^{(2)} \succeq 0$. Rearranging terms and using the well known fact that the shur product of two positive semi-definite matrices is positive semi-definite yields (2).

To prove (3), we use $\boldsymbol{P}=\boldsymbol{P} \boldsymbol{P}$, Cauchy-Schwarz, and (1) to derive

$$
\boldsymbol{P}_{i j}=\sum_{k \in[n]} \boldsymbol{P}_{i k} \boldsymbol{P}_{k j} \leq \sqrt{\left(\sum_{k \in[n]} \boldsymbol{P}_{i k}^{2}\right)\left(\sum_{k \in[n]} \boldsymbol{P}_{k j}^{2}\right)}=\sqrt{\boldsymbol{\Sigma}_{i i} \boldsymbol{\Sigma}_{j j}} .
$$

Squaring then yields (3).
To prove (4), we note that by the definition of $\boldsymbol{P}^{(2)}$ and Cauchy-Schwarz, we have

$$
\begin{equation*}
\left|\mathbf{1}_{i}^{T} \boldsymbol{P}^{(2)} \boldsymbol{x}\right|=\left|\sum_{j \in[n]} \boldsymbol{P}_{i j}^{(2)} \vec{x}_{j}\right| \leq \sqrt{\left(\sum_{j \in[n]} \boldsymbol{\Sigma}_{j j} \vec{x}_{j}^{2}\right) \cdot \sum_{j \in[n]} \frac{\boldsymbol{P}_{i j}^{(4)}}{\boldsymbol{\Sigma}_{j j}}} \tag{6.42}
\end{equation*}
$$

Now, by (1) and (3), we know that

$$
\begin{equation*}
\sum_{j \in[n]} \frac{\boldsymbol{P}_{i j}^{4}}{\boldsymbol{\Sigma}_{j j}} \leq \sum_{j \in[n]} \frac{\boldsymbol{P}_{i j}^{2} \boldsymbol{\Sigma}_{i i} \boldsymbol{\Sigma}_{j j}}{\boldsymbol{\Sigma}_{j j}}=\boldsymbol{\Sigma}_{i i} \sum_{j \in[n]} \boldsymbol{P}_{i j}^{2}=\boldsymbol{\Sigma}_{i i}^{2} \tag{6.43}
\end{equation*}
$$

Since $\|\boldsymbol{x}\|_{\boldsymbol{\Sigma}} \stackrel{\text { def }}{=} \sqrt{\sum_{j \in[n]} \boldsymbol{\Sigma}_{j j} \boldsymbol{x}_{j}^{2}}$, combining (6.42) and (6.43) yields $\left|\mathbf{1}_{i}^{T} \boldsymbol{P}^{(2)} \boldsymbol{x}\right| \leq \boldsymbol{\Sigma}_{i i}\|\boldsymbol{x}\|_{\boldsymbol{\Sigma}}$ as desired.
To prove (5), we let $f(\boldsymbol{w})=\log \operatorname{det}\left(\boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}\right)$. Lemma 2.3.8 shows that

$$
\begin{aligned}
\frac{\partial}{\partial w_{i}} f(\boldsymbol{w}) & =\operatorname{Tr}\left(\left(\boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}\right)^{-1} \boldsymbol{a}_{i} \boldsymbol{a}_{i}^{T}\right) \\
& =\boldsymbol{\Sigma}_{\boldsymbol{A}}(\boldsymbol{w})_{i i} w_{i}^{-1}
\end{aligned}
$$

To prove (6), we let $\boldsymbol{R}_{\boldsymbol{A}}(\boldsymbol{w})_{i j} \stackrel{\text { def }}{=} \boldsymbol{a}_{i}^{T}\left(\boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}\right)^{-1} \boldsymbol{a}_{j}$. Lemma 2.3.8 shows that

$$
\frac{\partial}{\partial w_{i}}\left(\boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}\right)^{-1}=-\left(\boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}\right)^{-1} \boldsymbol{a}_{i} \boldsymbol{a}_{i}^{T}\left(\boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}\right)^{-1}
$$

and hence $\frac{\partial}{\partial \boldsymbol{w}_{k}}\left[\boldsymbol{R}_{\boldsymbol{A}}(\boldsymbol{w})\right]_{i j}=-\boldsymbol{R}_{\boldsymbol{A}}(\boldsymbol{w})_{i k} \boldsymbol{R}_{\boldsymbol{A}}(\boldsymbol{w})_{k j}$. Since by definition $\boldsymbol{\sigma}_{\boldsymbol{A}}(\boldsymbol{w})_{i}=\boldsymbol{w}_{i} \boldsymbol{R}_{\boldsymbol{A}}(\boldsymbol{w})_{i i}$ by the previous lemma, we have that

$$
\frac{\partial}{\partial w_{j}} \boldsymbol{\sigma}_{\boldsymbol{A}}(\boldsymbol{w})_{i}=\mathbf{1}_{i=j} \boldsymbol{R}_{\boldsymbol{A}}(\boldsymbol{w})_{i i}-w_{i} \boldsymbol{R}_{\boldsymbol{A}}(\boldsymbol{w})_{i j}^{(2)}
$$

Writing this in matrix form and recalling the definition of the Jacobian then yields

$$
\boldsymbol{J}_{\boldsymbol{w}}\left(\boldsymbol{\sigma}_{\boldsymbol{A}}(\boldsymbol{w})\right)=\boldsymbol{\operatorname { D i a g }}\left(\boldsymbol{R}_{\boldsymbol{A}}(\boldsymbol{w})\right)-\boldsymbol{W} \boldsymbol{R}_{\boldsymbol{A}}(\boldsymbol{w})^{(2)} .
$$

Right multiplying by $\boldsymbol{I}=\boldsymbol{W} \boldsymbol{W}^{-1}$ and recalling the definition of $\boldsymbol{\Lambda}_{\boldsymbol{A}}$ then yields the result.

Lemma 6.11.3. For any $x, \varepsilon$ and $\lambda>0$, we have

$$
\frac{\pi}{2}|x|-\pi \varepsilon-\frac{1}{\lambda} \leq x \tan ^{-1}(\lambda(x+\varepsilon)) \leq \frac{\pi}{2}|x| .
$$

Proof. We first consider the case $\varepsilon=0$. Note that

$$
x \tan ^{-1}(\lambda x) \leq \frac{\pi}{2}|x|
$$

Also, we note that

$$
x \tan ^{-1}(\lambda x) \geq|x|\left(\frac{\pi}{2}-\frac{1}{\lambda|x|}\right)
$$

because

$$
\left|\tan \left(\frac{\pi}{2}-\frac{1}{\lambda|x|}\right)\right|=\left|\frac{\cos \left(\frac{1}{\lambda|x|}\right)}{\sin \left(\frac{1}{\lambda \mid x}\right)}\right| \leq \lambda|x| .
$$

Hence, we have

$$
\frac{\pi}{2}|x|-\frac{1}{\lambda} \leq x \tan ^{-1}(\lambda x) \leq \frac{\pi}{2}|x|
$$

For $\varepsilon \neq 0$, we have

$$
\frac{\pi}{2}|x+\varepsilon|-\frac{1}{\lambda} \leq(x+\varepsilon) \tan ^{-1}(\lambda(x+\varepsilon)) \leq \frac{\pi}{2}|x+\varepsilon|
$$

Thus, we have

$$
\frac{\pi}{2}|x|-\pi \varepsilon-\frac{1}{\lambda} \leq x \tan ^{-1}(\lambda(x+\varepsilon)) \leq \frac{\pi}{2}|x|
$$

## ■ 6.12 Appendix: Projection on Mixed Norm Ball

In this section, we give an algorithm to solve the problem

$$
\begin{equation*}
\max _{\|\boldsymbol{x}\|_{2}+\left\|\boldsymbol{l}^{-1} \boldsymbol{x}\right\|_{\infty} \leq 1}\langle\boldsymbol{a}, \boldsymbol{x}\rangle \tag{6.44}
\end{equation*}
$$

for some given vector $\boldsymbol{l}$ and $\boldsymbol{a}$. This is used in Section 6.7.1 to compute the weights. We note that

$$
\begin{align*}
\max _{\|\boldsymbol{x}\|_{2}+\left\|\boldsymbol{l}^{-1} \boldsymbol{x}\right\|_{\infty} \leq 1}\langle\boldsymbol{a}, \boldsymbol{x}\rangle & =\max _{0 \leq t \leq 1} \|_{\| \|_{2} \leq 1-t \text { and }-t l_{i} \leq x_{i} \leq t l_{i}}\langle\boldsymbol{a}, \boldsymbol{x}\rangle \\
& =\max _{0 \leq t \leq 1}(1-t) \max _{\|\boldsymbol{x}\|_{2} \leq 1 \text { and }-\frac{t}{1-t} l_{i} \leq x_{i} \leq \frac{t}{1-t} l_{i}}\langle\boldsymbol{a}, \boldsymbol{x}\rangle \\
& =\max _{0 \leq t \leq 1}(1-t) f\left(\frac{t}{1-t}\right) \tag{6.45}
\end{align*}
$$

where

$$
f\left(\frac{t}{1-t}\right)=\max _{\|\boldsymbol{x}\|_{2} \leq 1,-t l_{i} \leq x_{i} \leq t l_{i}}\langle\boldsymbol{a}, \boldsymbol{x}\rangle .
$$

In [165], we showed that the maximizer $\boldsymbol{x}$ in the problem $f$ must be of the form

$$
\boldsymbol{x}^{i_{t}}= \begin{cases}\frac{t}{1-t} \operatorname{sign}\left(a_{j}\right) l_{j} & \text { if } j \in\left\{1,2, \cdots, i_{t}\right\}  \tag{6.46}\\ \sqrt{\frac{1-\left(\frac{t}{1-t}\right)^{2} \sum_{k=0}^{i_{t}} l_{k}^{2}}{1-\sum_{k=0}^{i_{t}} a_{k}^{2}}} \boldsymbol{a}_{j} & \text { otherwise }\end{cases}
$$

where $i_{t}$ be the first coordinate such that

$$
\frac{1-\left(\frac{t}{1-t}\right)^{2} \sum_{k=0}^{i} l_{k}^{2}}{1-\sum_{k=0}^{i} a_{k}^{2}} \leq \frac{\left(\frac{t}{1-t}\right)^{2} l_{i}^{2}}{a_{i}^{2}}
$$

Note that $i_{t} \geq i_{s}$ if $t \leq s$. Therefore, we have that the set of $t$ such that $i_{t}=j$ is simply an interval given by ${ }^{8}$

$$
\begin{equation*}
\frac{\left|a_{j}\right|}{\sqrt{l_{j}^{2}\left(1-\sum_{k=0}^{j} a_{k}^{2}\right)+a_{j}^{2} \sum_{k=0}^{j} l_{k}^{2}}} \leq \frac{t}{1-t}<\frac{\left|a_{j-1}\right|}{\sqrt{l_{j-1}^{2}\left(1-\sum_{k=0}^{j-1} a_{k}^{2}\right)+a_{j-1}^{2} \sum_{k=0}^{j-1} l_{k}^{2}}} \tag{6.47}
\end{equation*}
$$

[^12]Therefore, we know that

$$
\begin{aligned}
f\left(\frac{t}{1-t}\right) & =\left\langle\boldsymbol{a}, \boldsymbol{x}^{\left(i_{t}\right)}\right\rangle \\
& =\frac{t}{1-t} \sum_{j=1}^{i_{t}}\left|a_{j}\right|\left|l_{j}\right|+\sqrt{1-\left(\frac{t}{1-t}\right)^{2} \sum_{k=0}^{i_{t}} l_{k}^{2}} \sqrt{1-\sum_{k=0}^{i_{t}} a_{k}^{2}} .
\end{aligned}
$$

Putting this into 6.45 , we have that

$$
\max _{\|\boldsymbol{x}\|_{2}+\left\|l^{-1} \boldsymbol{x}\right\|_{\infty} \leq 1}\langle\boldsymbol{a}, \boldsymbol{x}\rangle=\max _{0 \leq t \leq 1} t \sum_{j=1}^{i_{t}}\left|a_{j}\right|\left|l_{j}\right|+\sqrt{(1-t)^{2}-t^{2} \sum_{k=0}^{i_{t}} l_{k}^{2}} \sqrt{1-\sum_{k=0}^{i_{t}} a_{k}^{2}} .
$$

Note that the function $t \sum_{j=1}^{i}\left|a_{j}\right|\left|l_{j}\right|+\sqrt{(1-t)^{2}-t^{2} \sum_{k=0}^{i} l_{k}^{2}} \sqrt{1-\sum_{k=0}^{i} a_{k}^{2}}$ is concave and the solution has a close form. Therefore, one can compute the maximum value for each interval of $t$ (6.47) and find which is the best. Hence, we get the following algorithm.

| $\boldsymbol{x}=$ projectOntoMixedNormBallParallel $(\boldsymbol{a}, \boldsymbol{l})$ |
| :--- |
| 1. Set $\boldsymbol{a}=\boldsymbol{a} /\\|\boldsymbol{a}\\|_{2}$. |
| 2. Sort the coordinate such that $\left\|a_{i}\right\| / l_{i}$ is in descending order. |
| 3. Precompute $\sum_{k=0}^{i} l_{k}^{2}, \sum_{k=0}^{i} a_{k}^{2}$ and $\sum_{j=1}^{i}\left\|a_{j}\right\|\left\|l_{j}\right\|$ for all $i$. |
| 4. Let $g_{i}(t)=t \sum_{j=1}^{i}\left\|a_{j}\right\|\left\|l_{j}\right\|+\sqrt{(1-t)^{2}-t^{2} \sum_{k=0}^{i} l_{k}^{2}} \sqrt{1-\sum_{k=0}^{i} a_{k}^{2}}$. |
| 5. For each $j \in\{1, \cdots, n\}$, Find $t_{j}=\arg \max _{i_{t}=j} g_{j}(t)$ using $(6.47)$ |
| 6. Find $i=\arg \max _{i} g_{i}\left(t_{i}\right)$. |
| 7. Output $\left(1-t_{i}\right) \boldsymbol{x}^{(i)}$ defined by 6.46. |

The discussion above leads to the following theorem.
Theorem 6.12.1. The algorithm projectOntoMixedNormBallParallel outputs a solution to

$$
\max _{\|x\|_{2}+\left\|l^{-1} x\right\|_{\infty} \leq 1}\langle\boldsymbol{a}, \boldsymbol{x}\rangle
$$

in total work $\tilde{O}(m)$ and depth $\tilde{O}(1)$.

## ■ 6.13 Appendix: A $\tilde{O}(n)$-Self-Concordant Barrier Function

In this section, we proved the barrier function introduced in Subsection 6.3.5 is an $\tilde{O}(n)$-self-concordant barrier. We separate the proof of Theorem 6.3.5 into two lemma, Lemma 6.13.3 and Lemma 6.13.8. The proof is pretty straightforward, but it requires some familiarity with matrix calculus.

## ■ 6.13.1 Notation

Here, we define some notations that is used throughout in this section. They different from previous sections because this section our domain is $\Omega \stackrel{\text { def }}{=}\{\boldsymbol{A} \boldsymbol{x}>\boldsymbol{b}\}$ instead of $\boldsymbol{A}^{\top} \boldsymbol{x}=\boldsymbol{b}$.

1. Let $f(x, w)=\ln \operatorname{det}\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-1} \boldsymbol{W} \boldsymbol{S}^{-1} \boldsymbol{A}\right)-\frac{n}{m} \sum_{i=1}^{m} w_{i} \ln w_{i}-\frac{n}{m} \sum_{i=1}^{m} \ln s_{i}$.
2. Let $\boldsymbol{S}=\boldsymbol{\operatorname { D i a g }} \boldsymbol{A x}-\boldsymbol{b}$ and $\boldsymbol{A}_{x}=\boldsymbol{S}^{-1} \boldsymbol{A}$.
3. Let $\boldsymbol{P}(x, w)=\sqrt{\boldsymbol{W}} \boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}_{x}^{T} \sqrt{\boldsymbol{W}}, \sigma(x, w)=\boldsymbol{\operatorname { d i a g }}(\boldsymbol{P}(x, w))$.
4. Let $\boldsymbol{P}^{(2)}(x, w)$ be a $m \times m$ matrix such that $\left(\boldsymbol{P}^{(2)}(x, w)\right)_{i j}=(\boldsymbol{P}(x, w))_{i j}^{2}$.
5. Let $\boldsymbol{\Lambda}(x, w)=\boldsymbol{\Sigma}(x, w)-\boldsymbol{P}^{(2)}(x, w)$.
6. Let $w_{x}$ be the unique maximizer of $f(x, w)$, i.e. $w_{x} \stackrel{\text { def }}{=} \arg \max _{w_{i} \geq 0} f(x, w)$.
7. Let $\boldsymbol{P}_{x}=\boldsymbol{P}(x, w), \sigma_{x}=\sigma\left(x, w_{x}\right), \boldsymbol{\Sigma}_{x}=\boldsymbol{\Sigma}\left(x, w_{x}\right), \boldsymbol{P}_{x}^{(2)}=\boldsymbol{P}^{(2)}\left(x, w_{x}\right)$ and $\boldsymbol{\Lambda}_{x}=\boldsymbol{\Lambda}\left(x, w_{x}\right)$.

Similar to lemma 6.11.2, we have $\mathbf{0} \preceq \boldsymbol{P}^{(2)}(x, w) \preceq \boldsymbol{\Sigma}(x, w) \preceq \boldsymbol{I}, \mathbf{0} \preceq \boldsymbol{\Lambda}(x, w)$ and $\sum_{i} \sigma_{i}(x, w)=n$.

## ■ 6.13.2 Basic properties of $f$ and $p$

First, we compute the gradient and Hessian of $f$ and $p$.
Lemma 6.13.1. For any interior point $x \in \Omega$ and any positive weight $w>0$, we have

$$
\begin{align*}
\nabla_{x} f(x, w) & =-2 \boldsymbol{A}_{x}^{T} \sigma(x, w)-\frac{n}{m} \boldsymbol{A}_{x}^{T} 1, \\
\nabla_{w} f(x, w) & =\boldsymbol{W}^{-1} \sigma(x, w)-\frac{n}{m} \ln (e w)  \tag{6.48}\\
\nabla_{x x}^{2} f(x, w) & =\boldsymbol{A}_{x}^{T}\left(6 \boldsymbol{\Sigma}(x, w)-4 \boldsymbol{P}^{(2)}(x, w)+\frac{n}{m} \boldsymbol{I}\right) \boldsymbol{A}_{x}, \\
\nabla_{w w}^{2} f(x, w) & =-\boldsymbol{W}^{-1} \boldsymbol{P}^{(2)}(x, w) \boldsymbol{W}^{-1}-\frac{n}{m} \boldsymbol{W}^{-1} \\
\nabla_{x w}^{2} f(x, w) & =-2 \boldsymbol{A}_{x}^{T}\left(\boldsymbol{\Sigma}(x, w)-\boldsymbol{P}^{(2)}(x, w)\right) \boldsymbol{W}^{-1}
\end{align*}
$$

For $x \in \Omega$, $w_{x}$ uniquely exist in the region $\left\{\frac{1}{e} \leq w_{i} \leq \frac{m}{n}\right\}$ and $p$ satisfies

$$
\begin{align*}
\nabla p(x) & =-2 \boldsymbol{A}_{x}^{T} \sigma_{x}-\frac{n}{m} \boldsymbol{A}_{x}^{T} 1 \\
\nabla^{2} p(x) & =\boldsymbol{A}_{x}^{T}\left(6 \boldsymbol{\Sigma}_{x}-4 \boldsymbol{P}_{x}^{(2)}+\frac{n}{m} I\right) \boldsymbol{A}_{x}+4 \boldsymbol{A}_{x}^{T} \boldsymbol{\Lambda}_{x}\left(\boldsymbol{P}_{x}^{(2)}+\frac{n}{m} \boldsymbol{W}_{x}\right)^{-1} \boldsymbol{\Lambda}_{x} \boldsymbol{A}_{x} . \tag{6.49}
\end{align*}
$$

Also, we have

$$
\begin{equation*}
\boldsymbol{J}_{x} w=-2 \boldsymbol{W}\left(\boldsymbol{P}^{(2)}(x, w)+\frac{n}{m} \boldsymbol{W}\right)^{-1}\left(\boldsymbol{\Sigma}(x, w)-\boldsymbol{P}^{(2)}(x, w)\right) \boldsymbol{A}_{x} . \tag{6.50}
\end{equation*}
$$

Proof. For $\nabla_{x} f$, we have

$$
\begin{aligned}
\frac{\partial f}{\partial s_{i}} & =\operatorname{Tr}\left(\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-1} \boldsymbol{W} \boldsymbol{S}^{-1} \boldsymbol{A}\right)^{-1}\left(\frac{\partial}{\partial s_{i}}\right)\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-1} \boldsymbol{W} \boldsymbol{S}^{-1} \boldsymbol{A}\right)\right)-\frac{n}{m} \frac{1}{s_{i}} \\
& =-2 \operatorname{Tr}\left(\boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-1} \boldsymbol{W} \boldsymbol{S}^{-1} \boldsymbol{A}\right)^{-1} \boldsymbol{A}^{T} \frac{w_{i}}{s_{i}^{3}} 1_{i} 1_{i}^{T}\right)-\frac{n}{m} \frac{1}{s_{i}} \\
& =-2 \frac{\sigma_{i}(x, w)}{s_{i}}-\frac{n}{m} \frac{1}{s_{i}}
\end{aligned}
$$

By Chain rule, we have

$$
\nabla_{x} f(x, w)=-2 \boldsymbol{A}_{x}^{T} \sigma(x, w)-\frac{n}{m} \boldsymbol{A}_{x}^{T} 1 .
$$

For $\nabla_{w} f$, we have

$$
\begin{aligned}
\frac{\partial f}{\partial w_{i}} & =\operatorname{Tr}\left(\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-1} \boldsymbol{W} \boldsymbol{S}^{-1} \boldsymbol{A}\right)^{-1}\left(\frac{\partial}{\partial w_{i}}\right)\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-1} \boldsymbol{W} \boldsymbol{S}^{-1} \boldsymbol{A}\right)\right)-\frac{n}{m} \ln w_{i}-\frac{n}{m} \\
& =\frac{\sigma_{i}(x, w)}{w_{i}}-\frac{n}{m} \ln \left(e w_{i}\right)
\end{aligned}
$$

Hence, we have

$$
\nabla_{w} f(x, w)=\boldsymbol{W}^{-1} \sigma-\frac{n}{m} \ln (e w) .
$$

For $\nabla_{x x} f$, taking partial derivative on $\frac{\partial f}{\partial s_{i}}$, we have

$$
\begin{aligned}
\frac{\partial^{2} f}{\partial s_{i} \partial s_{j}}= & \frac{\partial}{\partial s_{j}}\left(-2 \operatorname{Tr}\left(\boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-1} \boldsymbol{W} \boldsymbol{S}^{-1} \boldsymbol{A}\right)^{-1} \boldsymbol{A}^{T} \frac{w_{i}}{s_{i}^{3}} 1_{i} 1_{i}^{T}\right)-\frac{n}{m} \frac{1}{s_{i}}\right) \\
= & -4 \operatorname{Tr}\left(\boldsymbol{A}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}^{T} \frac{w_{j}}{s_{j}^{3}} 1_{j} 1_{j}^{T} \boldsymbol{A}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W}_{x} \boldsymbol{A}\right)^{-1} \boldsymbol{A}^{T} \frac{w_{i}}{s_{i}^{3}} 1_{i} 1_{i}^{T}\right) \\
& 6 \operatorname{Tr}\left(\boldsymbol{A}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}^{T} \frac{w_{i}}{s_{i}^{4}} 1_{i} 1_{i}^{T} 1_{i=j}\right)+\frac{n}{m} \frac{1}{s_{i}^{2}} 1_{i=j} \\
= & 6 \frac{\sigma_{i}(x, w)}{s_{i}^{2}} 1_{i=j}-4 \frac{(\boldsymbol{P}(x, w))_{i j}^{2}}{s_{i} s_{j}}+\frac{n}{m} \frac{1}{s_{i}^{2}} 1_{i=j} .
\end{aligned}
$$

By Chain rule, we have

$$
\nabla_{x x} f(x, w)=\boldsymbol{A}_{x}^{T}\left(6 \boldsymbol{\Sigma}(x, w)-4 \boldsymbol{P}^{(2)}(x, w)+\frac{n}{m} \boldsymbol{I}\right) \boldsymbol{A}_{x} .
$$

For $\nabla_{w w} f$, taking partial derivative on $\frac{\partial f}{\partial w_{i}}$, we have

$$
\begin{aligned}
\frac{\partial^{2} f}{\partial w_{i} \partial w_{j}} & =\frac{\partial}{\partial w_{j}}\left(\operatorname{Tr}\left(\boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}_{x}^{T} e_{i} e_{i}^{T}\right)-\frac{n}{m} \ln w_{i}-\frac{n}{m}\right) \\
& =-\operatorname{Tr}\left(\boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}_{x}^{T} e_{j} e_{j}^{T} \boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}_{x}^{T} e_{i} e_{i}^{T}\right)-\frac{n}{m} \frac{1}{w_{j}} 1_{i=j} \\
& =-\frac{(\boldsymbol{P}(x, w))_{i j}^{2}}{w_{i} w_{j}}-\frac{n}{m} \frac{1}{w_{j}} 1_{i=j} .
\end{aligned}
$$

Hence, we have

$$
\nabla_{w w} f(x, w)=-\boldsymbol{W}^{-1} \boldsymbol{P}^{(2)}(x, w) \boldsymbol{W}^{-1}-\frac{n}{m} \boldsymbol{W}^{-1}
$$

For $\nabla_{x w} f$, taking partial derivative on $\frac{\partial f}{\partial w_{i}}$, we have

$$
\begin{aligned}
\frac{\partial^{2} f}{\partial s_{j} \partial w_{i}}= & \frac{\partial}{\partial s_{j}}\left(\operatorname{Tr}\left(\boldsymbol{A}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}^{T} \frac{e_{i} e_{i}^{T}}{s_{i}^{2}}\right)-\frac{n}{m} \ln w_{i}-\frac{n}{m}\right) \\
= & 2 \operatorname{Tr}\left(\boldsymbol{A}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}^{T} \frac{w_{j} e_{j} e_{j}^{T}}{s_{j}^{3}} \boldsymbol{A}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}^{T} \frac{e_{i} e_{i}^{T}}{s_{i}^{2}}\right) \\
& -2 \operatorname{Tr}\left(\boldsymbol{A}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}^{T} \frac{e_{i} e_{i}^{T}}{s_{i}^{3}} 1_{i=j}\right) \\
= & 2 \frac{(\boldsymbol{P}(x, w))_{i j}}{s_{j} w_{i}}-2 \frac{\sigma_{i}(x, w)}{s_{j} w_{i}} 1_{i=j} .
\end{aligned}
$$

By Chain rule, we have

$$
\nabla_{x w} f(x, w)=-2 \boldsymbol{A}_{x}^{T}\left(\boldsymbol{\Sigma}(x, w)-\boldsymbol{P}^{(2)}(x, w)\right) \boldsymbol{W}^{-1}
$$

Now, we prove that $w_{x}$ uniquely exists in the region $\left\{\frac{1}{e} \leq w_{i} \leq \frac{m}{n}\right\}$. For any $w$ such that $w_{i}>\frac{m}{n}$ for some $i$, we have

$$
\begin{aligned}
\frac{\partial f}{\partial w_{i}} & =\frac{\sigma_{i}(x, w)}{w_{i}}-\frac{n}{m} \ln \left(e w_{i}\right) \\
& \leq \frac{1}{w_{i}}-\frac{n}{m} \ln \left(e w_{i}\right) \\
& \leq \frac{n}{m}-\frac{n}{m} \ln \left(e \frac{m}{n}\right)<0 .
\end{aligned}
$$

Hence, we have $\sup _{w_{i} \geq 0} f(x, w)=\sup _{\frac{m}{n} \geq w_{i} \geq 0} f(x, w)$. Since the region $0 \leq w_{i} \leq \frac{m}{n}$ is a compact set and $f$ is a continuous function, there is $w_{x}$ such that attains $\sup _{w_{i} \geq 0} f(x, w)$. For the uniqueness, we note that $\nabla_{w w} f(x, w) \prec \mathbf{0}$ for all $w$ and hence $f$ is strictly concave and the maximizer is unique. Now, we note that for any $w$ such that $w_{i}<\frac{1}{e}$ for some $i$, we have

$$
\frac{\partial f}{\partial w_{i}}=\frac{\sigma_{i}(x, w)}{w_{i}}-\frac{n}{m} \ln \left(e w_{i}\right)>0 .
$$

Therefore, the minimizer is in the region $\left\{\frac{1}{e} \leq w_{i} \leq \frac{m}{n}\right\}$.
To compute $p$, we note that $f(x, w)$ is concave in $w$ and the maximizer is in the interior of the set $\left\{w_{i} \geq 0\right\}$ and therefore the optimality conditions imply that

$$
\begin{equation*}
\nabla_{w} f\left(x, w_{x}\right)=0 \tag{6.51}
\end{equation*}
$$

Since $\nabla_{w w}^{2} f\left(x, w_{x}\right)$ is invertible, implicit function theorem (Lemma 2.3.9) shows that $w_{x}$ is differentiable and satisfies

$$
\boldsymbol{J}_{x} w=-\left(\nabla_{w w}^{2} f\left(x, w_{x}\right)\right)^{-1}\left(\nabla_{w x}^{2} f\left(x, w_{x}\right)\right) .
$$

This gives (6.50).
Using that $p(x)=f\left(x, w_{x}\right)$ and applying the chain rule, we have

$$
\begin{aligned}
\nabla p(x) & =\nabla_{x} f\left(x, w_{x}\right)+\nabla_{w} f\left(x, w_{x}\right) \boldsymbol{J}_{x} w_{x} \\
& =\nabla_{x} f\left(x, w_{x}\right)
\end{aligned}
$$

because the optimality condition. Taking derivative again, we have

$$
\begin{aligned}
\nabla^{2} p(x) & =\nabla_{x x}^{2} f\left(x, w_{x}\right)+\nabla_{x w}^{2} f\left(x, w_{x}\right) \boldsymbol{J}_{x} w_{x} \\
& =\nabla_{x x}^{2} f\left(x, w_{x}\right)-\nabla_{x w}^{2} f\left(x, w_{x}\right)\left(\nabla_{w w}^{2} f\left(x, w_{x}\right)\right)^{-1}\left(\nabla_{w x}^{2} f\left(x, w_{x}\right)\right)
\end{aligned}
$$

Substituting in the computed values for $\nabla_{x x}^{2} f\left(x, w_{x}\right), \nabla_{x w}^{2} f\left(x, w_{x}\right)$, and $\nabla_{w x}^{2} f\left(x, w_{x}\right)$ then yields the result.

Now, we discuss properties of the optimal weight $w_{x}$.
Lemma 6.13.2. For $x \in \Omega$, we have $\sigma_{x}=\frac{n}{m} w_{x} \ln \left(e w_{x}\right)$. Also, we have $\max _{i} \frac{\sigma_{x, i}}{w_{x, i}} \leq \frac{n}{m} \ln \left(\frac{m e}{n}\right)$.
Proof. The optimality conditions (6.51) implies

$$
\sigma_{x}=\frac{n}{m} w_{x} \ln \left(e w_{x}\right) .
$$

Lemma 6.13 .1 shows that $\frac{1}{e} \leq w_{x} \leq \frac{m}{n}$ and hence

$$
\frac{\sigma_{x, i}}{w_{x, i}} \leq \frac{n}{m} \ln \left(\frac{m e}{n}\right)
$$

Using these, we can bound $(\nabla p)^{T}\left(\nabla^{2} p\right)^{-1}(\nabla p)$.
Lemma 6.13.3. For $x \in \Omega$, we have

$$
(\nabla p(x))^{T}\left(\nabla^{2} p(x)\right)^{-1}(\nabla p(x)) \leq 6 n
$$

Proof. Using $\nabla p(x)=-2 \boldsymbol{A}_{x}^{T} \sigma_{x}-\frac{n}{m} \boldsymbol{A}_{x}^{T} 1$, we have

$$
\begin{aligned}
& (\nabla p(x))^{T}\left(\nabla^{2} p(x)\right)^{-1}(\nabla p(x)) \\
\leq & 8 \sigma_{x}^{T} \boldsymbol{A}_{x}\left(\nabla^{2} p(x)\right)^{-1} \boldsymbol{A}_{x}^{T} \sigma_{x}+2\left(\frac{n}{m}\right)^{2} 1^{T} \boldsymbol{A}_{x}\left(\nabla^{2} p(x)\right)^{-1} \boldsymbol{A}_{x}^{T} 1
\end{aligned}
$$

For the first term, we use $\boldsymbol{\Lambda}_{x} \succeq 0$ and $\boldsymbol{P}_{x}^{(2)} \succeq 0$ and get

$$
\begin{aligned}
& \sigma_{x}^{T} \boldsymbol{A}_{x}\left(\nabla^{2} p(x)\right)^{-1} \boldsymbol{A}_{x}^{T} \sigma_{x} \\
\leq & \sigma_{x}^{T} \boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T}\left(6 \boldsymbol{\Sigma}_{x}-4 \boldsymbol{P}_{x}^{(2)}+\frac{n}{m} \boldsymbol{I}_{m}\right) \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}_{x}^{T} \sigma_{x} \\
\leq & \frac{1}{2} \sigma_{x}^{T} \boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{\Sigma}_{x} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}_{x}^{T} \sigma_{x} \\
\leq & \frac{1}{2} \sum_{i=1}^{m} \sigma_{x, i} \\
= & \frac{n}{2}
\end{aligned}
$$

For the second term, we have

$$
\begin{aligned}
& 1_{m}^{T} \boldsymbol{A}_{x}\left(\nabla^{2} p(x)\right)^{-1} \boldsymbol{A}_{x}^{T} 1_{m} \\
\leq & 1_{m}^{T} \boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T}\left(6 \boldsymbol{\Sigma}_{x}-4 \boldsymbol{P}_{x}^{(2)}+\frac{n}{m} \boldsymbol{I}_{m}\right) \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}_{x}^{T} 1_{m} \\
\leq & \frac{m}{n} 1_{m}^{T} \boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}_{x}^{T} 1_{m} \\
\leq & \frac{m^{2}}{n}
\end{aligned}
$$

Combining both terms, we have the result.

### 6.13.3 $w_{x}, \Sigma_{x}$ and $\boldsymbol{P}_{x}^{(2)}$ are stable

In this subsection, we write $x_{t}=x+t h, \boldsymbol{A}_{t}=\boldsymbol{A}_{x_{t}}, s_{t}=\boldsymbol{A}_{t} x_{t}-b, \boldsymbol{S}_{t}=\boldsymbol{D i a g}\left(s_{t}\right), w_{t}=w_{x_{t}}, \boldsymbol{W}_{t}=$ $\boldsymbol{D i a g}\left(w_{t}\right), \boldsymbol{P}_{t}=\boldsymbol{P}\left(x_{t}, w_{t}\right), \sigma_{t}=\sigma\left(x_{t}, w_{t}\right), \boldsymbol{\Sigma}_{t}=\boldsymbol{\Sigma}\left(x_{t}, w_{t}\right), \boldsymbol{P}_{t}^{(2)}=\boldsymbol{P}^{(2)}\left(x_{t}, w_{t}\right)$ and $\boldsymbol{\Lambda}_{t}=\boldsymbol{\Lambda}\left(x_{t}, w_{t}\right)$. First of all, we show that $\boldsymbol{A}_{x} h$ is small.

Lemma 6.13.4. For any $x \in \Omega$ and $h \in \mathbb{R}^{n}$, we have

$$
\left\|\boldsymbol{A}_{x} h\right\|_{\infty} \leq \frac{1}{\sqrt{2}} \ln \left(\frac{m e}{n}\right)\|h\|_{\nabla^{2} p(x)}
$$

Therefore, we have

$$
\left\|\boldsymbol{S}_{t}^{-1} \frac{d}{d t} s_{t}\right\|_{\infty} \leq \frac{1}{\sqrt{2}} \ln \left(\frac{m e}{n}\right)\|h\|_{\nabla^{2} p\left(x_{t}\right)}
$$

Proof. Note that

$$
\begin{aligned}
\left\|\boldsymbol{A}_{x} h\right\|_{\infty} & =\max _{i}\left\langle e_{i}, \boldsymbol{A}_{x} h\right\rangle \\
& \leq \max _{i}\left\langle\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W} \boldsymbol{A}_{x}\right)^{-1 / 2} \boldsymbol{A}_{x}^{T} 1_{i},\left(\boldsymbol{A}_{x}^{T} \boldsymbol{W} \boldsymbol{A}_{x}\right)^{1 / 2} h\right\rangle \\
& \leq\left(\max _{i} \frac{\sigma_{x, i}}{w_{x, i}}\right)\|h\|_{\boldsymbol{A}_{x}^{T} \boldsymbol{W} \boldsymbol{A}_{x}}
\end{aligned}
$$

Lemma 6.13 .2 shows that $\max _{i} \frac{\sigma_{x, i}}{w_{x, i}} \leq \frac{n}{m} \ln \left(\frac{m e}{n}\right)$ and hence

$$
\left\|\boldsymbol{A}_{x} h\right\|_{\infty} \leq \frac{n}{m} \ln \left(\frac{m e}{n}\right)\|h\|_{\boldsymbol{A}_{x}^{T} \boldsymbol{W} \boldsymbol{A}_{x}}
$$

On the other hand, we have

$$
\begin{aligned}
\nabla^{2} p & =\boldsymbol{A}_{x}^{T}\left(6 \boldsymbol{\Sigma}_{x}-4 \boldsymbol{P}_{x}^{(2)}+\frac{n}{m} \boldsymbol{I}\right) \boldsymbol{A}_{x}+4 \boldsymbol{A}_{x}^{T} \boldsymbol{\Lambda}_{x}\left(\frac{n}{m} \boldsymbol{W}_{x}+\boldsymbol{P}_{x}^{(2)}\right)^{-1} \boldsymbol{\Lambda}_{x} \boldsymbol{A}_{x} \\
& \succeq \boldsymbol{A}_{x}^{T}\left(2 \boldsymbol{\Sigma}_{x}+\frac{n}{m} \boldsymbol{I}\right) \boldsymbol{A}_{x}
\end{aligned}
$$

Lemma 6.13 .2 shows that $\sigma_{x}=\frac{n}{m} w \ln (e w)$ and hence

$$
2 \sigma+\frac{n}{m}=\frac{n}{m}(2 w \ln (e w)+1)
$$

Using $\frac{2 w \ln (e w)+1}{w} \geq 4-\ln (4) \geq 2$ for all $w>0$, we have

$$
\nabla^{2} p(x) \succeq 2 \boldsymbol{A}_{x}^{T} \boldsymbol{W}_{x} \boldsymbol{A}_{x}
$$

Hence, we have

$$
\left\|\boldsymbol{A}_{x} h\right\|_{\infty} \leq \frac{1}{\sqrt{2}} \ln \left(\frac{m e}{n}\right)\|h\|_{\nabla^{2} p(x)}
$$

The last claim follows from the equality

$$
\boldsymbol{S}_{t}^{-1} \frac{d}{d t} s_{t}=\boldsymbol{A}_{t} h
$$

In the next lemma, we show that $w_{t}$ is stable both $\boldsymbol{P}^{(2)}$ norm, $\boldsymbol{\Sigma}$ norm and $\ell^{\infty}$ norm.
Lemma 6.13.5. We have

$$
\begin{aligned}
\left\|\boldsymbol{W}_{t}^{-1} \frac{d}{d t} w_{t}\right\|_{\boldsymbol{P}_{t}^{(2)}} & \leq \frac{1}{2}\|h\|_{\nabla^{2} p\left(x_{t}\right)} \\
\left\|\boldsymbol{W}_{t}^{-1} \frac{d}{d t} w_{t}\right\|_{\boldsymbol{\Sigma}_{t}} & \leq \sqrt{\ln \left(\frac{m}{n} e\right)}\|h\|_{\nabla^{2} p\left(x_{t}\right)} \\
\left\|\boldsymbol{W}_{t}^{-1} \frac{d}{d t} w_{t}\right\|_{\infty} & \leq \ln \left(\frac{m e}{n}\right)\left(1+\sqrt{2} \ln \left(\frac{m e^{2}}{n}\right)\right)\|h\|_{\nabla^{2} p\left(x_{t}\right)}
\end{aligned}
$$

Proof. Using (6.50), we have

$$
\begin{equation*}
\frac{d}{d t} w_{t}=-2 \boldsymbol{W}_{t}\left(\frac{n}{m} \boldsymbol{W}_{t}+\boldsymbol{P}_{t}^{(2)}\right)^{-1} \boldsymbol{\Lambda}_{t} \boldsymbol{A}_{t} h \tag{6.52}
\end{equation*}
$$

Let $z=\left(\frac{n}{m} \boldsymbol{W}_{t}+\boldsymbol{P}_{t}^{(2)}\right)^{-1} \boldsymbol{\Lambda}_{t} \boldsymbol{A}_{t} h$. First, we note that

$$
\|z\|_{\boldsymbol{P}_{t}^{(2)}}^{2}=h^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\frac{n}{m} \boldsymbol{W}_{t}+\boldsymbol{P}_{t}^{(2)}\right)^{-1} \boldsymbol{P}_{t}^{(2)}\left(\frac{n}{m} \boldsymbol{W}_{t}+\boldsymbol{P}_{t}^{(2)}\right)^{-1} \boldsymbol{\Lambda}_{t} \boldsymbol{A}_{t} h
$$

Since $\boldsymbol{P}_{t}^{(2)} \preceq \frac{n}{m} \boldsymbol{W}_{t}+\boldsymbol{P}_{t}^{(2)}$ and $\nabla^{2} p \succeq 4 \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\frac{n}{m} \boldsymbol{W}_{t}+\boldsymbol{P}_{t}^{(2)}\right)^{-1} \boldsymbol{\Lambda}_{t} \boldsymbol{A}_{t}$ (6.49), we have

$$
\|z\|_{\boldsymbol{P}_{t}^{(2)}}^{2} \leq h^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\frac{n}{m} \boldsymbol{W}_{t}+\boldsymbol{P}_{t}^{(2)}\right)^{-1} \boldsymbol{\Lambda}_{t} \boldsymbol{A}_{t} h \leq \frac{1}{4}\|h\|_{\nabla^{2} p\left(x_{t}\right)}^{2}
$$

For the $\boldsymbol{\Sigma}_{t}$ bound, note that

$$
\|z\|_{\boldsymbol{\Sigma}_{t}}^{2}=h^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\frac{n}{m} \boldsymbol{W}_{t}+\boldsymbol{P}_{t}^{(2)}\right)^{-1} \boldsymbol{\Sigma}_{t}\left(\frac{n}{m} \boldsymbol{W}_{t}+\boldsymbol{P}_{t}^{(2)}\right)^{-1} \boldsymbol{\Lambda}_{t} \boldsymbol{A}_{t} h
$$

Using Lemma 6.13.2, we have $\max _{i} \frac{\sigma_{t, i}}{w_{t, i}} \leq \frac{n}{m} \ln \left(\frac{m e}{n}\right)$ and hence

$$
\begin{aligned}
\|z\|_{\boldsymbol{\Sigma}_{t}}^{2} & \leq \ln \left(\frac{m e}{n}\right) h^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\frac{n}{m} \boldsymbol{W}_{t}+\boldsymbol{P}_{t}^{(2)}\right)^{-1}\left(\frac{n}{m} \boldsymbol{W}_{t}\right)\left(\frac{n}{m} \boldsymbol{W}_{t}+\boldsymbol{P}_{t}^{(2)}\right)^{-1} \boldsymbol{\Lambda}_{t} \boldsymbol{A}_{t} h \\
& \leq \ln \left(\frac{m e}{n}\right) h^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\frac{n}{m} \boldsymbol{W}_{t}+\boldsymbol{P}_{t}^{(2)}\right)^{-1} \boldsymbol{\Lambda}_{t} \boldsymbol{A}_{t} h \leq \frac{\ln \left(\frac{m e}{n}\right)}{4}\|h\|_{\nabla^{2} p\left(x_{t}\right)}^{2}
\end{aligned}
$$

For the $\ell^{\infty}$ bound, note that

$$
\left(\frac{n}{m} \boldsymbol{W}_{t}+\boldsymbol{P}_{t}^{(2)}\right) z=\boldsymbol{\Sigma}_{t} \boldsymbol{A}_{t} h-\boldsymbol{P}_{t}^{(2)} \boldsymbol{A}_{t} h
$$

Hence, we have

$$
\begin{equation*}
\frac{n}{m} w_{t, i}\left|z_{i}\right| \leq\left|\boldsymbol{P}_{t}^{(2)} z\right|_{i}+\sigma_{t, i}\left|\boldsymbol{A}_{t} h\right|_{i}+\left|\boldsymbol{P}_{t}^{(2)} \boldsymbol{A}_{t} h\right|_{i} \tag{6.53}
\end{equation*}
$$

For the term $\left|\boldsymbol{P}_{t}^{(2)} z\right|_{i}$, note that

$$
\left|\boldsymbol{P}_{t}^{(2)} z\right|_{i}=\sqrt{1_{i}^{T} \boldsymbol{P}_{t}^{(2)} 1_{i}} \sqrt{\|z\|_{\boldsymbol{P}_{t}^{(2)}}^{2}}=\sigma_{t, i}\|z\|_{\boldsymbol{P}_{t}^{(2)}} \leq \frac{\sigma_{t, i}}{2}\|h\|_{\nabla^{2} p\left(x_{t}\right)}
$$

Therefore, $\left|\boldsymbol{P}_{t}^{(2)} z\right|_{i} \leq \frac{\sigma_{t, i}}{2}\|h\|_{\nabla^{2} p\left(x_{t}\right)}$.
For the term $\left|\boldsymbol{P}_{t}^{(2)} \boldsymbol{A}_{t} h\right|_{i}$, using $\nabla^{2} p\left(x_{t}\right) \succeq 2 \boldsymbol{A}_{t}^{t} \boldsymbol{\Sigma}_{t} \boldsymbol{A}_{t}$ and $\boldsymbol{\Sigma}_{t} \succeq \boldsymbol{P}_{t}^{(2)}$, we have

$$
\begin{aligned}
\left|\boldsymbol{P}_{t}^{(2)} \boldsymbol{A}_{t} h\right|_{i} & =\sqrt{1_{i}^{T} \boldsymbol{P}_{t}^{(2)} 1_{i}} \sqrt{\left\|\boldsymbol{A}_{t} h\right\|_{\boldsymbol{P}_{t}^{(2)}}^{2}} \\
& \leq \sigma_{t, i}\left\|\boldsymbol{A}_{t} h\right\|_{\boldsymbol{\Sigma}_{t}} \\
& \leq \frac{\sigma_{t, i}}{\sqrt{2}}\|h\|_{\nabla^{2} p\left(x_{t}\right)}
\end{aligned}
$$

For the term $\left|\boldsymbol{A}_{t} h\right|_{i}$, Lemma 6.13 .4 shows that $\left|\boldsymbol{A}_{t} h\right|_{i} \leq \frac{1}{\sqrt{2}} \ln \left(\frac{m e}{n}\right)\|h\|_{\nabla^{2} p\left(x_{t}\right)}$.
Combining these terms into (6.53), we have

$$
\frac{n}{m} w_{t, i}\left|z_{i}\right| \leq \sigma_{t, i}\left(\frac{1}{2}+\frac{1}{\sqrt{2}}+\frac{1}{\sqrt{2}} \ln \left(\frac{m e}{n}\right)\right)\|h\|_{\nabla^{2} p\left(x_{t}\right)}
$$

Using Lemma 6.13.2, we have $\max _{i} \frac{\sigma_{t, i}}{w_{t, i}} \leq \frac{n}{m} \ln \left(\frac{m e}{n}\right)$ and hence

$$
\left|z_{i}\right| \leq \ln \left(\frac{m e}{n}\right)\left(\frac{1}{2}+\frac{1}{\sqrt{2}}+\frac{1}{\sqrt{2}} \ln \left(\frac{m e}{n}\right)\right)\|h\|_{\nabla^{2} p\left(x_{t}\right)} .
$$

Using the formula (6.52) for $w_{t}$, we have the result.

In the next lemma, we show that $\boldsymbol{\Sigma}_{t}$ is stable.

Lemma 6.13.6. We have

$$
\left|\frac{d}{d t} \boldsymbol{\Sigma}_{t}\right| \preceq 2 \ln \left(\frac{m e}{n}\right)\left(1+\sqrt{2} \ln \left(\frac{m e^{2}}{n}\right)\right)\|h\|_{\nabla^{2} p\left(x_{t}\right)} \boldsymbol{\Sigma}_{t} .
$$

Proof. Fix any $y \in \mathbb{R}^{n}$. Let $\alpha(t)=y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Sigma}_{t} \boldsymbol{A}_{t} y$. The optimality condition (6.48) gives

$$
\sigma_{t}=\frac{n}{m} w_{t} \ln \left(e w_{t}\right) .
$$

Taking derivative on both side, we have

$$
\begin{aligned}
\frac{d}{d t} \sigma_{t} & =\frac{n}{m}\left(\frac{d}{d t} w_{t}\right) \ln \left(e w_{t}\right)+\frac{n}{m}\left(\frac{d}{d t} w_{t}\right) \\
& =\frac{n}{m}\left(\frac{d}{d t} w_{t}\right) \ln \left(e^{2} w_{t}\right)
\end{aligned}
$$

Using Lemma 6.13.5, we have

$$
\left\|\boldsymbol{W}_{t}^{-1} \frac{d}{d t} w_{t}\right\|_{\infty} \leq \ln \left(\frac{m e}{n}\right)\left(1+\sqrt{2} \ln \left(\frac{m e^{2}}{n}\right)\right)\|h\|_{\nabla^{2} p\left(x_{t}\right)} .
$$

Therefore, for all $i$, we have

$$
\begin{aligned}
\left|\frac{d}{d t} \sigma_{t, i}\right| & \leq \ln \left(\frac{m e}{n}\right)\left(1+\sqrt{2} \ln \left(\frac{m e^{2}}{n}\right)\right) \frac{n}{m} w_{t, i} \ln \left(e^{2} w_{t}\right) \\
& \leq 2 \ln \left(\frac{m e}{n}\right)\left(1+\sqrt{2} \ln \left(\frac{m e^{2}}{n}\right)\right) \sigma_{t, i}\|h\|_{\nabla^{2} p\left(x_{t}\right)} .
\end{aligned}
$$

This gives the result.

Now, we want to prove $\boldsymbol{P}_{t}^{(2)}$ is stable.
Lemma 6.13.7. We have

$$
\left|\frac{d}{d t} y^{T} \boldsymbol{P}_{t}^{(2)} y\right| \preceq\left(\ln \left(\frac{m e}{n}\right)\left(1+\sqrt{2} \ln \left(\frac{m e^{3}}{n}\right)\right)+2 \sqrt{\ln \left(\frac{m}{n} e\right)}+4\right)\|y\|_{\boldsymbol{\Sigma}_{t}}^{2}\|h\|_{\nabla^{2} p\left(x_{t}\right)} .
$$

Proof. Fix any $y \in \mathbb{R}^{n}$. Let $\alpha(t)=y^{T} \boldsymbol{P}_{t}^{(2)} y$. Then, we have

$$
\frac{d \alpha}{d t}=\frac{d}{d t} \sum_{i, j}\left(\boldsymbol{P}_{t}\right)_{i, j}^{2} y_{i} y_{j}=2 \sum_{i, j}\left(\boldsymbol{P}_{t}\right)_{i, j} \frac{d}{d t}\left(\boldsymbol{P}_{t}\right)_{i, j} y_{i} y_{j} .
$$

Let $z_{t}=\frac{d}{d t} \ln \left(w_{t} / s_{t}^{2}\right)$. Note that

$$
\begin{aligned}
\frac{d}{d t}\left(\boldsymbol{P}_{t}\right)_{i, j} & =\frac{d}{d t} 1_{i}^{T} \sqrt{\boldsymbol{W}_{t}} \boldsymbol{A}_{t}\left(\boldsymbol{A}_{t}^{T} \boldsymbol{W}_{t} \boldsymbol{A}_{t}\right)^{-1} \boldsymbol{A}_{t}^{T} \sqrt{\boldsymbol{W}_{t}} 1_{j} \\
& =\frac{1}{2} \frac{d z_{t, i}}{d t}\left(\boldsymbol{P}_{t}\right)_{i, j}+\frac{1}{2} \frac{d z_{t, j}}{d t}\left(\boldsymbol{P}_{t}\right)_{i, j}-1_{i}^{T} \cdot \boldsymbol{P}_{t} \cdot \boldsymbol{D i a g} \frac{d z_{t, k}}{d t} \cdot \boldsymbol{P}_{t} \cdot 1_{j}
\end{aligned}
$$

Hence, we have

$$
\begin{align*}
\frac{d \alpha}{d t}= & \sum_{i, j}\left(\boldsymbol{P}_{t}\right)_{i, j} \frac{d z_{t, i}}{d t}\left(\boldsymbol{P}_{t}\right)_{i, j} y_{i} y_{j}+\sum_{i, j}\left(\boldsymbol{P}_{t}\right)_{i, j} \frac{d z_{t, j}}{d t}\left(\boldsymbol{P}_{t}\right)_{i, j} y_{i} y_{j} \\
& -2 \sum_{i, j}\left(\boldsymbol{P}_{t}\right)_{i, j} 1_{i}^{T} \cdot \boldsymbol{P}_{t} \cdot \boldsymbol{D i a g}\left(\frac{d z_{t, k}}{d t}\right) \cdot \boldsymbol{P}_{t} \cdot 1_{j} y_{i} y_{j} \\
= & 2 y^{T} \boldsymbol{P}_{t}^{(2)} \boldsymbol{D i a g}\left(\frac{d z_{t, j}}{d t}\right) y-2 \sum_{i, j, k}\left(\boldsymbol{P}_{t}\right)_{i, j}\left(\boldsymbol{P}_{t}\right)_{i, k}\left(\boldsymbol{P}_{t}\right)_{j, k} y_{i} y_{j} \frac{d z_{t, k}}{d t} . \tag{6.54}
\end{align*}
$$

For the first term $y^{T} \boldsymbol{P}_{t}^{(2)} \boldsymbol{\operatorname { d i a g }}\left(\frac{d z_{t, j}}{d t}\right) y$, using $\boldsymbol{P}_{t}^{(2)} \preceq \boldsymbol{\Sigma}_{t}$, we have

$$
\left|y^{T} \boldsymbol{P}_{t}^{(2)} \boldsymbol{D i a g}\left(\frac{d z_{t, j}}{d t}\right) y\right| \leq\|y\|_{\boldsymbol{P}_{t}^{(2)}}\left\|\boldsymbol{D i a g}\left(\frac{d z_{t, j}}{d t}\right) y\right\|_{\boldsymbol{P}_{t}^{(2)}} \leq\|y\|_{\boldsymbol{\Sigma}_{t}}^{2}\left\|\frac{d z_{t}}{d t}\right\|_{\infty} .
$$

Using Lemma 6.13.4 and 6.13.5, we have

$$
\begin{aligned}
\left\|\frac{d z_{t}}{d t}\right\|_{\infty} & \leq\left\|\frac{d}{d t} \ln \left(w_{t}\right)\right\|_{\infty}+2\left\|\frac{d}{d t} \ln \left(s_{t}\right)\right\|_{\infty} \\
& \leq \ln \left(\frac{m e}{n}\right)\left(1+\sqrt{2} \ln \left(\frac{m e^{3}}{n}\right)\right)\|h\|_{\nabla^{2} p\left(x_{t}\right)}
\end{aligned}
$$

Hence, we have

$$
\left|y^{T} \boldsymbol{P}_{t}^{(2)} \boldsymbol{D i a g}\left(\frac{d z_{t, j}}{d t}\right) y\right| \leq\left(\ln \left(\frac{m e}{n}\right)\left(1+\sqrt{2} \ln \left(\frac{m e^{3}}{n}\right)\right)\right)\|y\|_{\boldsymbol{\Sigma}_{t}}^{2}\|h\|_{\nabla^{2} p\left(x_{t}\right)} .
$$

For the second term $\sum_{i, j, k}\left(\boldsymbol{P}_{t}\right)_{i, j}\left(\boldsymbol{P}_{t}\right)_{i, k}\left(\boldsymbol{P}_{t}\right)_{j, k} y_{i} y_{j} \frac{d}{d t} z_{t, k}$, we let $a_{i, j}=\left(\boldsymbol{P}_{t}\right)_{i, j} \sqrt{\left|y_{i} y_{j}\right|}, b_{i, k}=$ $\left(\boldsymbol{P}_{t}\right)_{i, k} \sqrt{\left|y_{i} \frac{d}{d t} z_{t, k}\right|}$ and $c_{j, k}=\left(\boldsymbol{P}_{t}\right)_{j, k} \sqrt{\left|y_{j} \frac{d}{d t} z_{t, k}\right|}$ and get

$$
\left|\sum_{i, j, k}\left(\boldsymbol{P}_{t}\right)_{i, j}\left(\boldsymbol{P}_{t}\right)_{i, k}\left(\boldsymbol{P}_{t}\right)_{j, k} y_{i} y_{j} \frac{d}{d t} z_{t, k}\right| \leq \sum_{i, j, k}\left|a_{i, j} b_{i, k} c_{j, k}\right| .
$$

Using Cauchy Schwarz Inequality twice and $\boldsymbol{P}_{t}^{(2)} \preceq \boldsymbol{\Sigma}_{t}$, we have

$$
\begin{aligned}
\left|\sum_{i, j, k}\left(\boldsymbol{P}_{t}\right)_{i, j}\left(\boldsymbol{P}_{t}\right)_{i, k}\left(\boldsymbol{P}_{t}\right)_{j, k} y_{i} y_{j} \frac{d}{d t} z_{t, k}\right| & \leq \sqrt{\sum_{i, j} a_{i, j}^{2}} \sqrt{\sum_{i, k} b_{i, k}^{2}} \sqrt{\sum_{j, k} c_{j, k}^{2}} \\
& =\||y|\|_{\boldsymbol{P}_{t}^{(2)}}\langle | y\left|,\left|\frac{d}{d t} z_{t}\right|\right\rangle_{\boldsymbol{P}_{t}^{(2)}}^{2} \\
& \leq\|y\|_{\boldsymbol{\Sigma}_{t}}^{2}\left\|\frac{d}{d t} z_{t}\right\|_{\boldsymbol{\Sigma}_{t}} .
\end{aligned}
$$

Using Lemma 6.13.5, we have

$$
\left|\sum_{i, j, k}\left(\boldsymbol{P}_{t}\right)_{i, j}\left(\boldsymbol{P}_{t}\right)_{i, k}\left(\boldsymbol{P}_{t}\right)_{j, k} y_{i} y_{j} \frac{d}{d t} z_{t, k}\right| \leq\left(\sqrt{\ln \left(\frac{m}{n} e\right)}+2\right)\|y\|_{\boldsymbol{\Sigma}_{t}}^{2}\|h\|_{\nabla^{2} p\left(x_{t}\right)}
$$

Combining two cases into (6.54), we have

$$
\left|\frac{d \alpha}{d t}\right| \leq\left(\ln \left(\frac{m e}{n}\right)\left(1+\sqrt{2} \ln \left(\frac{m e^{3}}{n}\right)\right)+2 \sqrt{\ln \left(\frac{m}{n} e\right)}+4\right)\|y\|_{\Sigma_{t}}^{2}\|h\|_{\nabla^{2} p\left(x_{t}\right)} .
$$

## $\square$ 6.13.4 $\nabla^{2} p$ is stable

Now, we can combine all lemmas before to prove the wanted statement.
Lemma 6.13.8. We have

$$
\left|\frac{d}{d t} \nabla^{2} p\right| \preceq O\left(\ln ^{3}\left(\frac{m e}{n}\right)\right)\|h\|_{\nabla^{2} p\left(x_{t}\right)} \nabla^{2} p .
$$

Proof. Note that

$$
\nabla^{2} p\left(x_{t}\right)=\boldsymbol{A}_{t}^{T}\left(6 \boldsymbol{\Sigma}_{t}-4 \boldsymbol{P}_{t}^{(2)}+\frac{n}{m} I\right) \boldsymbol{A}_{t}+4 \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1} \boldsymbol{\Lambda}_{t} \boldsymbol{A}_{t}
$$

Fix any $y \in \mathbb{R}^{n}$. Let $\alpha(t)=y^{T} \nabla^{2} p\left(x_{t}\right) y$. Then, we have

$$
\begin{aligned}
\left|\frac{d \alpha}{d t}\right| \leq & 6\left|\frac{d}{d t} y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Sigma}_{t} \boldsymbol{A}_{t} y\right|+4\left|\frac{d}{d t} y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{P}_{t}^{(2)} \boldsymbol{A}_{t} y\right| \\
& +4\left|\frac{d}{d t} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1} \boldsymbol{\Lambda}_{t} \boldsymbol{A}_{t}\right| \\
\leq & 12\left|y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Sigma}_{t}\left(\frac{d}{d t} \boldsymbol{A}_{t}\right) y\right|+6\left|y^{T} \boldsymbol{A}_{t}^{T}\left(\frac{d}{d t} \boldsymbol{\Sigma}_{t}\right) \boldsymbol{A}_{t} y\right| \\
& +8\left|y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{P}_{t}^{(2)}\left(\frac{d}{d t} \boldsymbol{A}_{t}\right) y\right|+4\left|y^{T} \boldsymbol{A}_{t}^{T}\left(\frac{d}{d t} \boldsymbol{P}_{t}^{(2)}\right) \boldsymbol{A}_{t} y\right| \\
& +8\left|y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1} \boldsymbol{\Lambda}_{t}\left(\frac{d}{d t} \boldsymbol{A}_{t}\right) y\right| \\
& +8\left|y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1}\left(\frac{d}{d t} \boldsymbol{P}_{t}^{(2)}\right) \boldsymbol{A}_{t} y\right| \\
& +8\left|y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1}\left(\frac{d}{d t} \boldsymbol{\Sigma}_{t}\right) \boldsymbol{A}_{t} y\right| \\
& +4\left|y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1}\left(\frac{d}{d t} \boldsymbol{P}_{t}^{(2)}\right)\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1} \boldsymbol{\Lambda}_{t} \boldsymbol{A}_{t} y\right| \\
& +4\left|y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1}\left(\frac{d}{d t} \frac{n}{m} \boldsymbol{W}_{t}\right)\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1} \boldsymbol{\Lambda}_{t} \boldsymbol{A}_{t} y\right| .
\end{aligned}
$$

Now, we need to bound the terms above one by one. Using Lemma 6.13.4, we have

$$
\begin{aligned}
\left|y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Sigma}_{t}\left(\frac{d}{d t} \boldsymbol{A}_{t}\right) y\right| & \leq O\left(\ln \left(\frac{m e}{n}\right)\right)\|h\|_{\nabla^{2} p\left(x_{t}\right)}\left|y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Sigma}_{t} \boldsymbol{A}_{t} y\right| \\
& =O\left(\ln \left(\frac{m e}{n}\right)\right)\|h\|_{\nabla^{2} p\left(x_{t}\right)}\|y\|_{\nabla^{2} p\left(x_{t}\right)}^{2} .
\end{aligned}
$$

Using Lemma 6.13.6, we have

$$
\begin{aligned}
\left|y^{T} \boldsymbol{A}_{t}^{T}\left(\frac{d}{d t} \boldsymbol{\Sigma}_{t}\right) \boldsymbol{A}_{t} y\right| & \leq O\left(\ln ^{2}\left(\frac{m e}{n}\right)\right)\|h\|_{\nabla^{2} p\left(x_{t}\right)}\left|y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Sigma}_{t} \boldsymbol{A}_{t} y\right| \\
& \leq O\left(\ln ^{2}\left(\frac{m e}{n}\right)\right)\|h\|_{\nabla^{2} p\left(x_{t}\right)}\|y\|_{\nabla^{2} p\left(x_{t}\right)}^{2} .
\end{aligned}
$$

Using Lemma 6.13.4, we have

$$
\begin{aligned}
\left|y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{P}_{t}^{(2)}\left(\frac{d}{d t} \boldsymbol{A}_{t}\right) y\right| & \leq \sqrt{y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{P}_{t}^{(2)} \boldsymbol{A}_{t} y} \sqrt{y^{T}\left(\frac{d}{d t} \boldsymbol{A}_{t}\right) \boldsymbol{P}_{t}^{(2)}\left(\frac{d}{d t} \boldsymbol{A}_{t}\right) y} \\
& \leq O(1) \sqrt{\|y\|_{\nabla^{2} p\left(x_{t}\right)}^{2}} \sqrt{\left(\ln \left(\frac{m e}{n}\right)\|h\|_{\nabla^{2} p\left(x_{t}\right)}\right)^{2}\|y\|_{\nabla^{2} p\left(x_{t}\right)}^{2}} \\
& =O\left(\ln \left(\frac{m e}{n}\right)\right)\|h\|_{\nabla^{2} p\left(x_{t}\right)}\|y\|_{\nabla^{2} p\left(x_{t}\right)}^{2} .
\end{aligned}
$$

Using Lemma 6.13.7, we have

$$
\begin{aligned}
\left|y^{T} \boldsymbol{A}_{t}^{T}\left(\frac{d}{d t} \boldsymbol{P}_{t}^{(2)}\right) \boldsymbol{A}_{t} y\right| & \leq O\left(\ln ^{2}\left(\frac{m e}{n}\right)\right)\left\|\boldsymbol{A}_{t} y\right\|_{\boldsymbol{\Sigma}_{t}}^{2}\|h\|_{\nabla^{2} p\left(x_{t}\right)} \\
& =O\left(\ln ^{2}\left(\frac{m e}{n}\right)\right)\|h\|_{\nabla^{2} p\left(x_{t}\right)}\|y\|_{\nabla^{2} p\left(x_{t}\right)}^{2}
\end{aligned}
$$

Using Lemma 6.13.4, we have

$$
\begin{aligned}
& \left|y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1} \boldsymbol{\Lambda}_{t}\left(\frac{d}{d t} \boldsymbol{A}_{t}\right) y\right| \\
\leq & \sqrt{y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1} \boldsymbol{\Lambda}_{t}\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1} \boldsymbol{A}_{t}^{T} y} \\
& \times \sqrt{y^{T}\left(\frac{d}{d t} \boldsymbol{A}_{t}\right)^{T} \boldsymbol{\Lambda}_{t}\left(\frac{d}{d t} \boldsymbol{A}_{t}\right) y} \\
= & O(1) \sqrt{\ln \left(\frac{m e}{n}\right)\|y\|_{\nabla^{2} p\left(x_{t}\right)}^{2}} \sqrt{y^{T}\left(\frac{d}{d t} \boldsymbol{A}_{t}\right)^{T} \Sigma_{t}\left(\frac{d}{d t} \boldsymbol{A}_{t}\right) y} \\
= & O\left(\ln ^{3 / 2}\left(\frac{m e}{n}\right)\right)\|h\|_{\nabla^{2} p\left(x_{t}\right)}\|y\|_{\nabla^{2} p\left(x_{t}\right)}^{2}
\end{aligned}
$$

Using Lemma 6.13.6, we have

$$
\begin{aligned}
& \left|y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1}\left(\frac{d}{d t} \boldsymbol{\Sigma}_{t}\right) \boldsymbol{A}_{t} y\right| \\
\leq & \sqrt{y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1}\left(\frac{d}{d t} \boldsymbol{\Sigma}_{t}\right)\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1} \boldsymbol{A}_{t}^{T} y} \\
& \times \sqrt{y^{T} \boldsymbol{A}_{t}^{T}\left(\frac{d}{d t} \boldsymbol{\Sigma}_{t}\right) \boldsymbol{A}_{t} y} \\
\leq & O\left(\ln ^{3 / 2}\left(\frac{m e}{n}\right)\right)\|h\|_{\nabla^{2} p\left(x_{t}\right)}\|y\|_{\nabla^{2} p\left(x_{t}\right)}^{2}
\end{aligned}
$$

Let $z=\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1} \boldsymbol{\Lambda}_{t} \boldsymbol{A}_{t} y$, using Lemma 6.13.7, we have

$$
\begin{aligned}
& \left|y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1}\left(\frac{d}{d t} \boldsymbol{P}_{t}^{(2)}\right)\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1} \boldsymbol{\Lambda}_{t} \boldsymbol{A}_{t} y\right| \\
\leq & O\left(\ln ^{2}\left(\frac{m e}{n}\right)\right)\|z\|_{\boldsymbol{\Sigma}_{t}}^{2}\|h\|_{\nabla^{2} p\left(x_{t}\right)} .
\end{aligned}
$$

Note that

$$
\begin{aligned}
\|z\|_{\boldsymbol{\Sigma}_{t}}^{2} & =y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1} \boldsymbol{\Sigma}_{t}\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1} \boldsymbol{\Lambda}_{t} \boldsymbol{A}_{t} y \\
& \leq \ln \left(\frac{m e}{n}\right) y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1} \boldsymbol{\Lambda}_{t} \boldsymbol{A}_{t} y \leq O\left(\ln \left(\frac{m e}{n}\right)\right)\|y\|_{\nabla^{2} p\left(x_{t}\right)}^{2} .
\end{aligned}
$$

Hence, we have

$$
\begin{aligned}
& \left|y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1}\left(\frac{d}{d t} \boldsymbol{P}_{t}^{(2)}\right)\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1} \boldsymbol{\Lambda}_{t} \boldsymbol{A}_{t} y\right| \\
\leq & O\left(\ln ^{3}\left(\frac{m e}{n}\right)\right)\|y\|_{\nabla^{2} p\left(x_{t}\right)}^{2}\|h\|_{\nabla^{2} p\left(x_{t}\right)} .
\end{aligned}
$$

Using Lemma 6.13.5, we have

$$
\begin{aligned}
& \left|y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1}\left(\frac{d}{d t} \frac{n}{m} \boldsymbol{W}_{t}\right)\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1} \boldsymbol{\Lambda}_{t} \boldsymbol{A}_{t} y\right| \\
\leq & O\left(\ln ^{2}\left(\frac{m e}{n}\right)\right)\|y\|_{\nabla^{2} p\left(x_{t}\right)}^{2}\|h\|_{\nabla^{2} p\left(x_{t}\right)}
\end{aligned}
$$

Using Lemma 6.13.7, we have

$$
\begin{aligned}
& \left|y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1}\left(\frac{d}{d t} \boldsymbol{P}_{t}^{(2)}\right) \boldsymbol{A}_{t} y\right| \\
\leq & \sqrt{y^{T} \boldsymbol{A}_{t}^{T} \boldsymbol{\Lambda}_{t}\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1}\left(\frac{d}{d t} \boldsymbol{P}_{t}^{(2)}\right)\left(\boldsymbol{P}_{t}^{(2)}+\frac{n}{m} \boldsymbol{W}_{t}\right)^{-1} \boldsymbol{\Lambda}_{t} \boldsymbol{A}_{t} y} \sqrt{y^{T} \boldsymbol{A}_{t}^{T}\left(\frac{d}{d t} \boldsymbol{P}_{t}^{(2)}\right) \boldsymbol{A}_{t} y} \\
\leq & \sqrt{O\left(\ln ^{3}\left(\frac{m e}{n}\right)\right)\|y\|_{\nabla^{2} p\left(x_{t}\right)}^{2}\|h\|_{\nabla^{2} p\left(x_{t}\right)}} \sqrt{O\left(\ln ^{2}\left(\frac{m e}{n}\right)\|y\|_{\nabla^{2} p\left(x_{t}\right)}^{2}\|h\|_{\nabla^{2} p\left(x_{t}\right)}\right.} \\
= & O\left(\ln ^{5 / 2}\left(\frac{m e}{n}\right)\right)\|y\|_{\nabla^{2} p\left(x_{t}\right)}^{2}\|h\|_{\nabla^{2} p\left(x_{t}\right)} .
\end{aligned}
$$

Combining all terms together, we have $\left|\frac{d \alpha}{d t}\right|=O\left(\ln ^{3}\left(\frac{m e}{n}\right)\right)\|y\|_{\nabla^{2} p\left(x_{t}\right)}^{2}\|h\|_{\nabla^{2} p\left(x_{t}\right)}$.

# Geometric Median In Nearly-Linear Time 

## ■ 7.1 Introduction

In this chapter, we study Fermat-Weber problem, one of the oldest easily-stated nontrivial problems in computational geometry. Given a set of $n$ points in dimensions, $a^{(1)}, \ldots, a^{(n)} \in \mathbb{R}^{d}$, find a point $x_{*} \in \mathbb{R}^{d}$ that minimizes the sum of Euclidean distances to them:

$$
x_{*} \in \underset{x \in \mathbb{R}^{d}}{\arg \min } f(x) \quad \text { where } \quad f(x) \xlongequal{\text { def }} \sum_{i \in[n]}\left\|x-a^{(i)}\right\|_{2}
$$

This problem, also known as the geometric median problem, is well studied and has numerous applications. It is often considered over low dimensional spaces in the context of the facility location problem [267] and over higher dimensional spaces it has applications to clustering in machine learning and data analysis. For example, computing the geometric median is a subroutine in popular expectation maximization heuristics for $k$-medians clustering.

The problem is also important to robust estimation, where we like to find a point representative of given set of points that is resistant to outliers. The geometric median is a rotation and translation invariant estimator that achieves the optimal breakdown point of 0.5 , i.e. it is a good estimator even when up to half of the input data is arbitrarily corrupted [172]. Moreover, if a large constant fraction of the points lie in a ball of diameter $\varepsilon$ then the geometric median lies in that ball with diameter $O(\varepsilon)$ (see Lemma 7.7.1). Consequently, the geometric median can be used to turn expected results into high probability results: e.g. if the $a^{(i)}$ are drawn independently such that $\mathbf{E}\left\|x-x_{*}\right\|_{2} \leq \varepsilon$ for some $\varepsilon>0$ and $x \in \mathbb{R}^{d}$ then this fact, Markov bound, and Chernoff Bound, imply $\left\|x_{*}-a^{(i)}\right\|_{2}=O(\varepsilon)$ with high probability in $n$.

Despite the ancient nature of the Fermat-Weber problem and its many uses there are relatively few theoretical guarantees for solving it (see Table 7.1). To compute a $(1+\varepsilon)$-approximate solution, i.e. $x \in \mathbb{R}^{d}$ with $f(x) \leq(1+\varepsilon) f\left(x_{*}\right)$, the previous fastest running times were either $O\left(d \cdot n^{4 / 3} \varepsilon^{-8 / 3}\right)$ by [50], $\tilde{O}\left(d \exp \varepsilon^{-4} \log \varepsilon^{-1}\right)$ by [25], $\tilde{O}\left(n d+\operatorname{poly}\left(d, \varepsilon^{-1}\right)\right)$ by [86], or $O\left((n d)^{O(1)} \log \frac{1}{\varepsilon}\right)$ time by [222, 271]. In this chapter, we improve upon these running times by providing an $O\left(n d \log ^{3} \frac{n}{\varepsilon}\right)$ time algorithm ${ }^{1}$ as well as an $O\left(d / \varepsilon^{2}\right)$ time algorithm, provided we have an oracle for sampling a random $a^{(i)}$. Picking the faster algorithm for the particular value of $\varepsilon$ improves the running time to $O\left(n d \log ^{3} \frac{1}{\varepsilon}\right)$. We also extend these results to compute a $(1+\varepsilon)$-approximate solution to the more general Weber's problem, $\min _{x \in \mathbb{R}^{d}} \sum_{i \in[n]} w_{i}\left\|x-a^{(i)}\right\|_{2}$ for non-negative $w_{i}$, in time $O\left(n d \log ^{3} \frac{1}{\varepsilon}\right)$ (see Section 7.10).

Our $O\left(n d \log ^{3} \frac{n}{\varepsilon}\right)$ time algorithm is a careful modification of standard interior point methods for solving the geometric median problem. We provide a long step interior point method tailored to the geometric median problem for which we can implement every iteration in nearly linear time. While our analysis starts with a simple $O\left((n d)^{O(1)} \log \frac{1}{\varepsilon}\right)$ time interior point method and shows how to improve it,

[^13]our final algorithm is quite non-standard from the perspective of interior point literature. Our result is one of very few cases we are aware of outperforming traditional interior point theory (Chapter 6 and [183]) and the only we are aware of using interior point methods to obtain a nearly linear time algorithm for a canonical optimization problem that traditionally requires superlinear time. We hope our work leads to further improvements in this line of research.

Our $O\left(d \varepsilon^{-2}\right)$ algorithm is a relatively straightforward application of sampling techniques and stochastic subgradient descent. Some additional insight is required simply to provide a rigorous analysis of the robustness of the geometric median and use this to streamline our application of stochastic subgradient descent. We include it for completeness however, we defer its proof to Section 7.7. The bulk of the work in this chapter is focused on developing our $O\left(n d \log ^{3} \frac{n}{\varepsilon}\right)$ time algorithm which we believe uses a set of techniques of independent interest.

## ■ 7.1.1 Previous Work

The geometric median problem was first formulated for the case of three points in the early 1600 s by Pierre de Fermat [148, 69]. A simple elegant ruler and compass construction was given in the same century by Evangelista Torricelli. Such a construction does not generalize when a larger number of points is considered: Bajaj has shown the even for five points, the geometric median is not expressible by radicals over the rationals [27]. Hence, the $(1+\varepsilon)$-approximate problem has been studied for larger values of $n$.

Many authors have proposed algorithms with runtime polynomial in $n, d$ and $1 / \varepsilon$. The most cited and used algorithm is Weiszfeld's 1937 algorithm [268]. Unfortunately Weiszfeld's algorithm may not converge and if it does it may do so very slowly. There have been many proposed modifications to Weiszfeld's algorithm [58, 224, 220, 28, 260, 155] that generally give non-asymptotic runtime guarantees. In light of more modern multiplicative weights methods his algorithm can be viewed as a re-weighted least squares iteration. Chin et al. [50] considered the more general $L_{2}$ embedding problem: placing the vertices of a graph into $\mathbb{R}^{d}$, where some of the vertices have fixed positions while the remaining vertices are allowed to float, with the objective of minimizing the sum of the Euclidean edge lengths. Using the multiplicative weights method, they obtained a run time of $O\left(d \cdot n^{4 / 3} \varepsilon^{-8 / 3}\right)$ for a broad class of problems, including the geometric median problem. ${ }^{2}$

Many authors consider problems that generalize the Fermat-Weber problem, and obtain algorithms for finding the geometric median as a specialization. For example, Badoiu et al. give an approximate $k$-median algorithm by sub-sampling with the runtime for $k=1$ of $\widetilde{O}\left(d \cdot \exp \left(O\left(\varepsilon^{-4}\right)\right)\right)$ [25]. Parrilo and Sturmfels demonstrated that the problem can be reduced to semidefinite programming, thus obtaining a runtime of $\widetilde{O}\left(\operatorname{poly}(n, d) \log \varepsilon^{-1}\right)$ [222]. Furthermore, Bose et al. gave a linear time algorithm for fixed $d$ and $\varepsilon^{-1}$, based on low-dimensional data structures [37] and it has been show how to obtain running times of $\widetilde{O}\left(n d+\operatorname{poly}\left(d, \varepsilon^{-1}\right)\right)$ for this problem and a more general class of problems.[113, 86].

An approach very related to ours was studied by Xue and Ye [271]. They give an interior point method with barrier analysis that runs in time $\tilde{O}\left(\left(d^{3}+d^{2} n\right) \sqrt{n} \log \varepsilon^{-1}\right)$.

## ■ 7.1.2 Overview of $O\left(n d \log ^{3} \frac{n}{\varepsilon}\right)$ Time Algorithm

## Interior Point Primer

Our algorithm is broadly inspired by interior point methods, a broad class of methods for efficiently solving convex optimization problems [273, 211]. Given an instance of the geometric median problem

[^14]| Year | Authors | Runtime | Comments |
| :---: | :---: | :---: | :---: |
| 1659 | Torricelli $[263]$ | - | Assuming $n=3$ |
| 1937 | Weiszfeld [268] | - | Does not always converge |
| 1990 | Chandrasekaran and Tamir[47] | $\widetilde{O}\left(n \cdot \operatorname{poly}(d) \log \varepsilon^{-1}\right)$ | Ellipsoid method |
| 1997 | Xue and Ye [271] | $\widetilde{O}\left(\left(d^{3}+d^{2} n\right) \sqrt{n} \log \varepsilon^{-1}\right)$ | Interior point with barrier method |
| 2000 | Indyk [116] | $\widetilde{O}\left(d n \cdot \varepsilon^{-2}\right)$ | Optimizes only over $x$ in the input |
| 2001 | Parrilo and Sturmfels $[222]$ | $\widetilde{O}\left(\operatorname{poly}(n, d) \log \varepsilon^{-1}\right)$ | Reduction to SDP |
| 2002 | Badoiu et al. [25] | $\widetilde{O}\left(d \cdot \exp \left(O\left(\varepsilon^{-4}\right)\right)\right)$ | Sampling |
| 2003 | Bose et al. $[37]$ | $\widetilde{O}(n)$ | Assuming $d, \varepsilon^{-1}=O(1)$ |
| 2005 | Har-Peled and Kushal $[113]$ | $\widetilde{O}\left(n+\operatorname{poly}\left(\varepsilon^{-1}\right)\right)$ | Assuming $d=O(1)$ |
| 2011 | Feldman and Langberg $[86]$ | $\widetilde{O}\left(n d+\operatorname{poly}\left(d, \varepsilon^{-1}\right)\right)$ | Coreset |
| 2013 | Chin et al. [50] | $\widetilde{O}\left(d n^{4 / 3} \cdot \varepsilon^{-8 / 3}\right)$ | Multiplicative weights |
| - | This Chapter | $O\left(n d \log ^{3}(n / \varepsilon)\right)$ | Interior point with custom analysis |
| - | This Chapter | $O\left(d \varepsilon^{-2}\right)$ | Stochastic gradient descent |

Table 7.1: Selected Previous Results.
we first put the problem in a more natural form for applying interior point methods. Rather than writing the problem as minimizing a convex function over $\mathbb{R}^{d}$

$$
\begin{equation*}
\min _{x \in \mathbb{R}^{d}} f(x) \quad \text { where } \quad f(x) \xlongequal{\text { def }} \sum_{i \in[n]}\left\|x-a^{(i)}\right\|_{2} \tag{7.1}
\end{equation*}
$$

we instead write the problem as minimizing a linear function over a larger convex space:

$$
\begin{equation*}
\min _{\{\alpha, x\} \in S} 1^{\top} \alpha \quad \text { where } S=\left\{\alpha \in \mathbb{R}^{n}, x \in \mathbb{R}^{d} \quad \mid \quad\left\|x^{(i)}-a^{(i)}\right\|_{2} \leq \alpha_{i} \quad \text { for all } \quad i \in[n]\right\} . \tag{7.2}
\end{equation*}
$$

Clearly, these problems are the same as at optimality $\alpha_{i}=\left\|x^{(i)}-a^{(i)}\right\|_{2}$.
To solve problems of the form (7.2) interior point methods relax the constraint $\{\alpha, x\} \in S$ through the introduction of a barrier function. In particular they assume that there is a real valued function $p$ such that as $\{\alpha, x\}$ moves towards the boundary of $S$ the value of $p$ goes to infinity. A popular class of interior point methods known as path following methods [227, 107], they consider relaxations of (7.2) of the form $\min _{\{\alpha, x\} \in \mathbb{R}^{n} \times \mathbb{R}^{d}} t \cdot 1^{\top} \alpha+p(\alpha, x)$. The minimizers of this function form a path, known as the central path, parameterized by $t$. The methods then use variants of Newton's method to follow the path until $t$ is large enough that a high quality approximate solution is obtained. The number of iterations of these methods are then typically governed by a property of $p$ known as its self concordance $\nu$. Given a $\nu$-self concordant barrier, typically interior point methods require $O\left(\sqrt{\nu} \log \frac{1}{\varepsilon}\right)$ iterations to compute a $(1+\varepsilon)$-approximate solution.

For our particular convex set, the construction of our barrier function is particularly simple, we consider each constraint $\left\|x-a^{(i)}\right\|_{2} \leq \alpha_{i}$ individually. In particular, it is known that the function $p^{(i)}(\alpha, x)=-\ln \alpha_{i}-\ln \left(\alpha_{i}^{2}-\left\|x-a^{(i)}\right\|_{2}\right)$ is a $O(1)$ self-concordant barrier function for the set $S^{(i)}=$ $\left\{x \in \mathbb{R}^{d}, \alpha \in \mathbb{R}^{n} \mid\left\|x-a^{(i)}\right\|_{2} \leq \alpha_{i}\right\}[206]$. Since $\cap_{i \in[n]} S^{(i)}=S$ we can use the barrier $\sum_{i \in[n]} p^{(i)}(\alpha, x)$ for $p(\alpha, x)$ and standard self-concordance theory shows that this is an $O(n)$ self concordant barrier for $S$. Consequently, this easily yields an interior point method for solving the geometric median problem in $O\left((n d)^{O(1)} \log \frac{1}{\varepsilon}\right)$ time.

## Difficulties

Unfortunately obtaining a nearly linear time algorithm for geometric median using interior point methods as presented poses numerous difficulties. Particularly troubling is the number of iterations required by standard interior point algorithms. The approach outlined in the previous section produced an $O(n)$-self concordant barrier and even if we use more advanced self concordance machinery, i.e. the universal barrier [211], the best known self concordance of barrier for the convex set $\sum_{i \in[n]}\left\|x-a^{(i)}\right\|_{2} \leq$ $c$ is $O(d)$. An interesting open question still left open by our work is to determine what is the minimal self concordance of a barrier for this set.

Consequently, even if we could implement every iteration of an interior point scheme in nearly linear time it is unclear whether one should hope for a nearly linear time interior point algorithm for the geometric median. While there are a instances of outperforming standard self-concordance analysis (Chapter 6), these instances are few, complex, and to varying degrees specialized to the problems they solve. Moreover, we are unaware of any interior point scheme providing a provable nearly linear time for a general nontrivial convex optimization problem.

## Beyond Standard Interior Point

Despite these difficulties we do obtain a nearly linear time interior point based algorithm that only requires $O\left(\log \frac{n}{\varepsilon}\right)$ iterations, i.e. increases to the path parameter. After choosing the natural penalty functions $p^{(i)}$ described above, we optimize in closed form over the $\alpha_{i}$ to obtain the following penalized objective function: ${ }^{3}$

$$
\min _{x} f_{t}(x) \quad \text { where } \quad f_{t}(x)=\sum_{i \in[n]} \sqrt{1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}-\ln \left[1+\sqrt{1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}\right]
$$

We then approximately minimize $f_{t}(x)$ for increasing $t$. We let $x_{t} \stackrel{\text { def }}{=} \arg \min _{x \in \mathbb{R}^{d}} f_{t}(x)$ for $x \geq 0$, and thinking of $\left\{x_{t}: t \geq 0\right\}$ as a continuous curve known as the central path, we show how to approximately follow this path. As $\lim _{t \rightarrow \infty} x_{t}=x_{*}$ this approach yields a $(1+\varepsilon)$-approximation.

So far our analysis is standard and interior point theory yields an $\Omega(\sqrt{n})$ iteration interior point scheme. To overcome this we take a more detailed look at $x_{t}$. We note that for any $t$ if there is any rapid change in $x_{t}$ it must occur in the direction of the smallest eigenvector of $\nabla^{2} f_{t}(x)$, denoted $v_{t}$, what we henceforth may refer to as the bad direction at $x_{t}$. Furthermore, we show that this bad direction is stable in the sense that for all directions $d \perp v_{t}$ it is the case that $d^{\top}\left(x_{t}-x_{t^{\prime}}\right)$ is small for $t^{\prime} \leq c t$ for a small constant $c$.

In fact, we show that this movement over such a long step, i.e. constant increase in $t$, in the directions orthogonal to the bad direction is small enough that for any movement around a ball of this size the Hessian of $f_{t}$ only changes by a small multiplicative constant. In short, starting at $x_{t}$ there exists a point $y$ obtained just by moving from $x_{t}$ in the bad direction, such that $y$ is close enough to $x_{t^{\prime}}$ that standard first order method will converge quickly to $x_{t^{\prime}}$ ! Thus, we might hope to find such a $y$, quickly converge to $x_{t^{\prime}}$ and repeat. If we increase $t$ by a multiplicative constant in every such iterations, standard interior point theory suggests that $O\left(\log \frac{n}{\varepsilon}\right)$ iterations suffices.

## Building an Algorithm

To turn the structural result in the previous section into a fast algorithm there are several further issues we need to address. We need to

[^15]- (1) Show how to find the point along the bad direction that is close to $x_{t^{\prime}}$
- (2) Show how to solve linear systems in the Hessian to actually converge quickly to $x_{t^{\prime}}$
- (3) Show how to find the bad direction
- (4) Bound the accuracy required by these computations

Deferring (1) for the moment, our solution to the rest are relatively straightforward. Careful inspection of the Hessian of $f_{t}$ reveals that it is well approximated by a multiple of the identity matrix minus a rank 1 matrix. Consequently using explicit formulas for the inverse of of matrix under rank 1 updates, i.e. the Sherman-Morrison formula, we can solve such systems in nearly linear time thereby addressing (2). For (3), we show that the well known power method carefully applied to the Hessian yields the bad direction if it exists. Finally, for (4) we show that a constant approximate geometric median is near enough to the central path for $t=\Theta\left(\frac{1}{f\left(x_{*}\right)}\right)$ and that it suffices to compute a central path point at $t=O\left(\frac{n}{f\left(x_{*}\right) \varepsilon}\right)$ to compute a $1+\varepsilon$-geometric median. Moreover, for these values of $t$, the precision needed in other operations is clear.

The more difficult operation is (1). Given $x_{t}$ and the bad direction exactly, it is still not clear how to find the point along the bad direction line from $x_{t}$ that is close to $x_{t^{\prime}}$. Just performing binary search on the objective function a priori might not yield such a point due to discrepancies between a ball in Euclidean norm and a ball in hessian norm and the size of the distance from the optimal point in euclidean norm. To overcome this issue we still line search on the bad direction, however rather than simply using $f\left(x_{t}+\alpha \cdot v_{t}\right)$ as the objective function to line search on, we use the function $g(\alpha)=\min _{\left\|x-x_{t}-\alpha \cdot v_{t}\right\|_{2} \leq c} f(x)$ for some constant $c$, that is given an $\alpha$ we move $\alpha$ in the bad direction and take the best objective function value in a ball around that point. For appropriate choice of $c$ the minimizers of $\alpha$ will include the optimal point we are looking for. Moreover, we can show that $g$ is convex and that it suffices to perform the minimization approximately.

Putting these pieces together yields our result. We perform $O\left(\log \frac{n}{\varepsilon}\right)$ iterations of interior point (i.e. increasing $t$ ), where in each iteration we spend $O\left(n d \log \frac{n}{\varepsilon}\right)$ time to compute a high quality approximation to the bad direction, and then we perform $O\left(\log \frac{n}{\varepsilon}\right)$ approximate evaluations on $g(\alpha)$ to binary search on the bad direction line, and then to approximately evaluate $g$ we perform gradient descent in approximate Hessian norm to high precision which again takes $O\left(n d \log \frac{n}{\varepsilon}\right)$ time. Altogether this yields a $O\left(n d \log ^{3} \frac{n}{\varepsilon}\right)$ time algorithm to compute a $1+\varepsilon$ geometric median. Here we made minimal effort to improve the log factors and plan to investigate this further in future work.

## - 7.1.3 Overview of $O\left(d \varepsilon^{-2}\right)$ Time Algorithm

In addition to providing a nearly linear time algorithm we provide a stand alone result on quickly computing a crude $(1+\varepsilon)$-approximate geometric median in Section 7.7. In particular, given an oracle for sampling a random $a^{(i)}$ we provide an $O\left(d \varepsilon^{-2}\right)$, i.e. sublinear, time algorithm that computes such an approximate median. Our algorithm for this result is fairly straightforward. First, we show that random sampling can be used to obtain some constant approximate information about the optimal point in constant time. In particular we show how this can be used to deduce an Euclidean ball which contains the optimal point. Second, we perform stochastic subgradient descent within this ball to achieve our desired result.

## ■ 7.1.4 Organization

The rest of the chapter is structured as follows. After covering preliminaries in Section 2.3, in Section 7.3 we provide various results about the central path that we use to derive our nearly linear time
algorithm. In Section 7.4 we then provide our nearly linear time algorithm. All the proofs and supporting lemmas for these sections are deferred to Section 7.5 and Section 7.6. In Section 7.7 we provide our $O\left(d / \varepsilon^{2}\right)$ algorithm, in Section 7.8 we provide the derivation of our penalized objective function, in Section 7.9 we provide general technical machinery we use throughout, and in Section 7.10 we show how to extend our results to Weber's problem, i.e. weighted geometric median.

### 7.2 Preliminaries

## ■ 7.2.1 General Notation

For a symmetric positive semidefinite matrix (PSD), $\boldsymbol{A}$, we let $\lambda_{1}(\boldsymbol{A}) \geq \ldots \geq \lambda_{n}(\boldsymbol{A}) \geq 0$ denote the eigenvalues of $\boldsymbol{A}$ and let $v_{1}(\boldsymbol{A}), \ldots, v_{n}(\boldsymbol{A})$ denote corresponding eigenvectors.

## ■ 7.2.2 Problem Notation

The central problem of this chapter is as follows: we are given points $a^{(1)}, \ldots, a^{(n)} \in \mathbb{R}^{d}$ and we wish to compute a geometric median, i.e. $x_{*} \in \arg \min _{x \in \mathbb{R}^{d}} f(x)$ where $f(x)=\sum_{i \in[n]}\left\|a^{(i)}-x\right\|_{2}$. We call a point $x \in \mathbb{R}^{d}$ an $(1+\varepsilon)$-approximate geometric median if $f(x) \leq(1+\varepsilon) f\left(x_{*}\right)$.

## ■ 7.2.3 Penalized Objective Notation

To solve this problem we relax the objective function $f$ and instead consider the following family of penalized objective functions parameterized by $t>0$

$$
\min _{x \in \mathbb{R}^{d}} f_{t}(x) \quad \text { where } \quad f_{t}(x)=\sum_{i \in[n]} \sqrt{1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}-\ln \left[1+\sqrt{1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}\right]
$$

This penalized objective function is derived from a natural interior point formulation of the geometric median problem (See Section 7.8). For all path parameters $t>0$, we let $x_{t} \stackrel{\text { def }}{=} \arg \min _{x} f_{t}(x)$. Our primary goal is to obtain good approximations to the central path $\left\{x_{t}: t>0\right\}$ for increasing values of $t$.

We let $g_{t}^{(i)}(x) \stackrel{\text { def }}{=} \sqrt{1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}$ and $f_{t}^{(i)}(x) \stackrel{\text { def }}{=} g_{t}^{(i)}(x)-\ln \left(1+g_{t}^{(i)}(x)\right)$ so $f_{t}(x)=\sum_{i \in[n]} f_{t}^{(i)}(x)$. We refer to the quantity $w_{t}(x) \stackrel{\text { def }}{=} \sum_{i \in[n]} \frac{1}{1+g_{t}^{(i)}(x)}$ as weight as it is a natural measure of total contribution of the $a^{(i)}$ to $\nabla^{2} f_{t}(x)$. We let

$$
\bar{g}_{t}(x) \stackrel{\text { def }}{=} w_{t}(x)\left[\sum_{i \in[n]} \frac{1}{\left(1+g_{t}^{(i)}\left(x_{t}\right)\right) g_{t}^{(i)}\left(x_{t}\right)}\right]^{-1}
$$

denote a natural term that helps upper bound the rate of change of the central path. Furthermore, we let $u^{(i)}(x)$ denote the unit vector corresonding to $x-a^{(i)}$, i.e. $u^{(i)}(x) \xlongequal{\text { def }} x-a^{(i)} /\left\|x-a^{(i)}\right\|_{2}$ when $\left\|x-a^{(i)}\right\|_{2} \neq 0$ and $u^{(i)}(x) \stackrel{\text { def }}{=} 0$ otherwise. Finally we let $\mu_{t}(x) \stackrel{\text { def }}{=} \lambda_{d}\left(\nabla^{2} f_{t}(x)\right)$, denote the minimum eigenvalue of $\nabla^{2} f_{t}(x)$, and let $v_{t}(x)$ denote a corresponding eigenvector. To simplify notation we often drop the $(x)$ in these definitions when $x=x_{t}$ and $t$ is clear from context.

## ■ 7.3 Properties of the Central Path

Here provide various facts regarding the penalized objective function and the central path. While we use the lemmas in this section throughout the chapter, the main contribution of this section is Lemma 7.3.5 in Section 7.3.3. There we prove that with the exception of a single direction, the change in the central path is small over a constant multiplicative change in the path parameter. In addition, we show that our penalized objective function is stable under changes in a $O\left(\frac{1}{t}\right)$ Euclidean ball (Section 7.3.1), we bound the change in the Hessian over the central path (Section 7.3.2). and we relate $f\left(x_{t}\right)$ to $f\left(x_{*}\right)$ (Section 7.3.4).

## ■ 7.3.1 How Much Does the Hessian Change in General?

Here, we show that the Hessian of the penalized objective function is stable under changes in a $O\left(\frac{1}{t}\right)$ sized Euclidean ball. This shows that if we have a point which is close to a central path point in Euclidean norm, then we can use Newton method to find it.
Lemma 7.3.1. Suppose that $\|x-y\|_{2} \leq \frac{\varepsilon}{t}$ with $\varepsilon \leq \frac{1}{20}$. Then, we have

$$
\left(1-6 \varepsilon^{2 / 3}\right) \nabla^{2} f_{t}(x) \preceq \nabla^{2} f_{t}(y) \preceq\left(1+6 \varepsilon^{2 / 3}\right) \nabla^{2} f_{t}(x) .
$$

## ■ 7.3.2 How Much Does the Hessian Change Along the Path?

Here we bound how much the Hessian of the penalized objective function can change along the central path. First we provide the following lemma bound several aspects of the penalized objective function and proving that the weight, $w_{t}$, only changes by a small amount multiplicatively given small multiplicative changes in the path parameter, $t$.
Lemma 7.3.2. For all $t \geq 0$ and $i \in[n]$ the following hold

$$
\left\|\frac{d}{d t} x_{t}\right\|_{2} \leq \frac{1}{t^{2}} \bar{g}_{t}\left(x_{t}\right) \quad, \quad\left|\frac{d}{d t} g_{t}^{(i)}\left(x_{t}\right)\right| \leq \frac{1}{t}\left(g_{t}^{(i)}\left(x_{t}\right)+\bar{g}_{t}\right) \quad, \text { and } \quad\left|\frac{d}{d t} w_{t}\right| \leq \frac{2}{t} w_{t}
$$

Consequently, for all $t^{\prime} \geq t$ we have that $\left(\frac{t}{t^{\prime}}\right)^{2} w_{t} \leq w_{t^{\prime}} \leq\left(\frac{t^{\prime}}{t}\right)^{2} w_{t}$.
Next we use this lemma to bound the change in the Hessian with respect to $t$.
Lemma 7.3.3. For all $t \geq 0$ we have

$$
\begin{equation*}
-12 \cdot t \cdot w_{t} \boldsymbol{I} \preceq \frac{d}{d t}\left[\nabla^{2} f_{t}\left(x_{t}\right)\right] \preceq 12 \cdot t \cdot w_{t} \boldsymbol{I} \tag{7.3}
\end{equation*}
$$

and therefore for all $\beta \in\left[0, \frac{1}{8}\right]$

$$
\begin{equation*}
\nabla^{2} f\left(x_{t}\right)-15 \beta t^{2} w_{t} \boldsymbol{I} \preceq \nabla^{2} f\left(x_{t(1+\beta)}\right) \preceq \nabla^{2} f\left(x_{t}\right)+15 \beta t^{2} w_{t} \boldsymbol{I} . \tag{7.4}
\end{equation*}
$$

## ■ 7.3.3 Where is the Next Optimal Point?

Here we prove our main result of this section. We prove that over a long step the central path moves very little in directions orthogonal to the smallest eigenvector of the Hessian. We begin by noting the Hessian is approximately a scaled identity minus a rank 1 matrix.
Lemma 7.3.4. For all $t$ we have $\frac{1}{2}\left[t^{2} \cdot w_{t} \boldsymbol{I}-\left(t^{2} \cdot w_{t}-\mu_{t}\right) v_{t} v_{t}^{\top}\right] \preceq \nabla^{2} f_{t}\left(x_{t}\right) \preceq t^{2} \cdot w_{t} \boldsymbol{I}-\left(t^{2} \cdot w_{t}-\mu_{t}\right) v_{t} v_{t}^{\top}$
Using this and the lemmas of the previous section we bound the amount $x_{t}$ can move in every direction far from $v_{t}$.

Lemma 7.3.5 (The Central Path is Almost Straight). For all $t \geq 0, \beta \in\left[0, \frac{1}{600}\right]$, and any unit vector $y$ with $\left|\left\langle y, v_{t}\right\rangle\right| \leq \frac{1}{t^{2} \cdot \kappa}$ where $\kappa=\max _{\delta \in[t,(1+\beta) t]} \frac{w_{\delta}}{\mu_{\delta}}$, we have $y^{\top}\left(x_{(1+\beta) t}-x_{t}\right) \leq \frac{6 \beta}{t}$.

## - 7.3.4 Where is the End?

In this section, we bound the quality of the central path with respect to the geometric median objective. In particular, we show that if we can solve the problem for some $t=\frac{2 n}{\varepsilon f\left(x_{*}\right)}$ then we obtain an $(1+\varepsilon)-$ approximate solution. As our algorithm ultimately starts from an initial $t=1 / O\left(f\left(x_{*}\right)\right)$ and increases $t$ by a multiplicative constant in every iteration, this yields an $O\left(\log \frac{n}{\varepsilon}\right)$ iteration algorithm.
Lemma 7.3.6. $f\left(x_{t}\right)-f\left(x_{*}\right) \leq \frac{2 n}{t}$ for all $t>0$.

## ■ 7.4 Overview of the Algorithm

```
Algorithm 13: AccurateMedian( \(\varepsilon\) )
    Input: desired accuracy \(\varepsilon \in(0,1)\)
    // Compute a 2-approximate geometric median and use it to center
    Compute \(x^{(0)}:=\frac{1}{n} \sum_{i \in[n]} a^{(i)}\) and \(\widetilde{f}_{*}:=f\left(x^{(0)}\right) / /\) Note \(\tilde{f}_{*} \leq 2 f\left(x_{*}\right)\) by Lemma 7.5.6
    Let \(t_{i}=\frac{1}{400 \tilde{f}_{*}}\left(1+\frac{1}{600}\right)^{i-1}, \tilde{\varepsilon}_{*}=\frac{1}{3} \varepsilon\), and \(\tilde{t}_{*}=\frac{2 n}{\tilde{\varepsilon} \cdot \tilde{f}_{*}}\)
    Let \(\varepsilon_{v}=\frac{1}{8}\left(\frac{\tilde{\varepsilon}_{*}}{7 n}\right)^{2}\) and let \(\varepsilon_{c}=\left(\frac{\varepsilon_{v}}{36}\right)^{\frac{3}{2}}\)
    \(x^{(1)}=\operatorname{LineSearch}\left(x^{(0)}, t_{1}, t_{1}, 0, \varepsilon_{c}\right)\)
    // Iteratively improve quality of approximation
    Let \(k=\max _{i \in \mathbb{Z}} t_{i} \leq \tilde{t}_{*}\)
    for \(i \in[1, k]\) do
        // Compute \(\varepsilon_{v}\)-approximate minimum eigenvalue and eigenvector of \(\nabla^{2} f_{t_{i}}\left(x^{(i)}\right)\)
        \(\left(\lambda^{(i)}, u^{(i)}\right)=\operatorname{ApproxMinEig}\left(x^{(i)}, t_{i}, \varepsilon_{v}\right)\).
        // Line search to find \(x^{(i+1)}\) such that \(\left\|x^{(i+1)}-x_{t_{i+1}}\right\|_{2} \leq \frac{\varepsilon_{c}}{t_{i+1}}\)
        \(x^{(i+1)}=\operatorname{LineSearch}\left(x^{(i)}, t_{i}, t_{i+1}, u^{(i)}, \varepsilon_{c}\right) \mathrm{w}\).
    end
    Output: \(\varepsilon\)-approximate geometric median \(x^{(k)}\)
```

Here we show how to use the structural results from the previous section to obtain a nearly linear time algorithm for computing the geometric median. Our algorithm follows a simple structure (See Algorithm 13). First we use simply average the $a^{(i)}$ to compute a ( $1+\frac{1}{n}$ )-approximate median, denoted $x^{(0)}$. Then for a number of iterations we repeatedly move closer to $x_{t}$ for some path parameter $t$, compute the minimum eigenvector of the Hessian, and line search in that direction to find an approximation to a point further along the central path. Ultimately, this yields a point $x^{(k)}$ that is precise enough approximation to a point along the central path with large enough $t$ that we can simply out $x^{(k)}$ as our $(1+\varepsilon)$-approximate geometric median.

We split the remainder of the algorithm specification and its analysis into several parts. First in Section 7.4.1 we show how to compute an approximate minimum eigenvector and eigenvalue of the Hessian of the penalized objective function. Then in Section 7.4 .2 we show how to use this eigenvector to line search for the next central path point. Finally, in Section 7.4 .3 we put these results together to obtain our nearly linear time algorithm. Throughout this section we will want an upper bound to
$f\left(x_{*}\right)$ and a slight lower bound on $\varepsilon$, the geometric median accuracy we are aiming for. We use an easily computed $\tilde{f}_{*} \leq 2 f\left(x_{*}\right)$ for the former and $\tilde{\varepsilon}_{*}=\frac{1}{3} \varepsilon$ throughout the section.

## - 7.4.1 Eigenvector Computation and Hessian Approximation

Here we show how to compute the minimum eigenvector of $\nabla^{2} f_{t}(x)$ and thereby obtain a concise approximation to $\nabla^{2} f_{t}(x)$. Our main algorithmic tool is the well known power method and the fact that it converges quickly on a matrix with a large eigenvalue gap. To improve our logarithmic terms we need a slightly non-standard analysis of the method and therefore we provide and analyze this method for completeness in Section 7.6.1. Using this tool we estimate the top eigenvector as follows.

```
Algorithm 14: ApproxMinEig \((x, t, \varepsilon)\)
    Input: Point \(x \in \mathbb{R}^{d}\), path parameter \(t\), and target accuracy \(\varepsilon\).
    Let \(\boldsymbol{A}=\sum_{i \in[n]} \frac{t^{4}\left(x-a^{(i)}\right)\left(x-a^{(i)}\right)^{\top}}{\left(1+g_{t}^{(i)}(x)\right)^{2} g_{t}^{(i)}(x)}\)
    Let \(u:=\operatorname{PowerMethod}\left(\boldsymbol{A}, \Theta\left(\log \left(\frac{n}{\varepsilon}\right)\right)\right)\)
    Let \(\lambda=u^{\top} \nabla^{2} f_{t}(x) u\)
    Output: \((\lambda, u)\)
```

Lemma 7.4.1 (Computing Hessian Approximation). Let $x \in \mathbb{R}^{d}, t>0$, and $\varepsilon \in\left(0, \frac{1}{4}\right)$. The algorithm ApproxMinEig $(x, t, \varepsilon)$ outputs $(\lambda, u)$ in $O\left(n d \log \frac{n}{\varepsilon}\right)$ time such that if $\mu_{t}(x) \leq \frac{1}{4} t^{2} w_{t}(x)$ then $\left\langle v_{t}(x), u\right\rangle^{2} \geq 1-\varepsilon$ with high probability in $n$. Furthermore, if $\varepsilon \leq \frac{\mu_{t}(x)}{8 t^{2} \cdot w_{t}(x)}$ then $\frac{1}{4} \boldsymbol{Q} \preceq \nabla^{2} f_{t}(x) \preceq 4 \boldsymbol{Q}$ with high probability in $n$ where $\boldsymbol{Q} \stackrel{\text { def }}{=} t^{2} \cdot w_{t}(x)-\left(t^{2} \cdot w_{t}(x)-\lambda\right) u u^{\top}$.

Furthermore, we show that the $v^{(i)}$ computed by this algorithm is sufficiently close to the bad direction. Combining 7.4.1 with the structural results from the previous section and Lemma 7.9.4, a minor technical lemma regarding the transitivity of large inner products, we provide the following lemma.
Lemma 7.4.2. Let $u=\operatorname{ApproxMinEig}\left(x, t, \varepsilon_{v}\right)$ for $\varepsilon_{v}<\frac{1}{8}$ and $x$ such that $\left\|x-x_{t}\right\|_{2} \leq \frac{\varepsilon_{c}}{t}$ for $\varepsilon_{c} \leq\left(\frac{\varepsilon_{v}}{36}\right)^{\frac{3}{2}}$. $\mu_{t} \leq \frac{1}{4} t^{2} \cdot w_{t}$. For all unit vectors $y \perp u$, we have $\left\langle y, v_{t}\right\rangle^{2} \leq 8 \varepsilon_{v}$.

Note that this lemma assumes $\mu_{t}$ is small. When $\mu_{t}$ is large, we instead show that the next central path point is close to the current point and hence we do not need to compute the bad direction to center quickly.
Lemma 7.4.3. Suppose $\mu_{t} \geq \frac{1}{4} t^{2} \cdot w_{t}$ and let $t^{\prime} \in\left[t,\left(1+\frac{1}{600}\right) t\right]$ then $\left\|x_{t^{\prime}}-x_{t}\right\|_{2} \leq \frac{1}{100 t}$.

## ■ 7.4.2 Line Searching

Here we show how to line search along the bad direction to find the next point on the central path. Unfortunately, it is not clear if you can binary search on the objective function directly. If we search over $\alpha$ to minimize $f_{t_{i+1}}\left(y^{(i)}+\alpha v^{(i)}\right)$ directly it is unclear if we actually obtain a point close to $x_{t+1}$. It might be the case that even after minimizing $\alpha$ we would be unable to move towards $x_{t+1}$ efficiently.

To overcome this difficulty, we use the fact that over the region $\|x-y\|_{2}=O\left(\frac{1}{t}\right)$ the Hessian changes by at most a constant and therefore we can minimize $f_{t}(x)$ over this region extremely quickly. Therefore, we instead line search on the following function

$$
\begin{equation*}
g_{t, y, v}(\alpha) \stackrel{\text { def }}{=} \min _{\|x-(y+\alpha v)\|_{2} \leq \frac{1}{49 t}} f_{t}(x) \tag{7.5}
\end{equation*}
$$

and use that we can evaluate $g_{t, y, v}(\alpha)$ approximately by using an appropriate centering procedure. We can show (See Lemma 7.9.6) that $g_{t, y, v}(\alpha)$ is convex and therefore we can minimize it efficiently just by doing an appropriate binary search. By finding the approximately minimizing $\alpha$ and outputting the corresponding approximately minimizing $x$, we can obtain $x^{(i+1)}$ that is close enough to $x_{t_{i+1}}$. For notational convenience, we simply write $g(\alpha)$ if $t, y, v$ is clear from the context.

First, we show how we can locally center and provide error analysis for that algorithm.

```
Algorithm 15: LocalCenter \((y, t, \varepsilon)\)
    Input: Point \(y \in \mathbb{R}^{d}\), path parameter \(t>0\), target accuracy \(\varepsilon>0\).
    Let \((\lambda, v):=\operatorname{ApproxMinEig}(x, t, \varepsilon)\).
    Let \(\boldsymbol{Q}=t^{2} \cdot w_{t}(y) \boldsymbol{I}-\left(t^{2} \cdot w_{t}(y)-\lambda\right) v v^{\top}\)
    Let \(x^{(0)}=y\)
    for \(i=1, \ldots, k=64 \log \frac{1}{\varepsilon}\) do
        Let \(x^{(i)}=\min _{\|x-y\|_{2} \leq \frac{1}{49 t}} f\left(x^{(i-1)}\right)+\left\langle\nabla f_{t}\left(x^{(i-1)}\right), x-x^{(i-1)}\right\rangle+4\left\|x-x^{(i-1)}\right\|_{\boldsymbol{Q}}^{2}\).
    end
    Output: \(x^{(k)}\)
```

Lemma 7.4.4. Given some $y \in \mathbb{R}^{d}, t>0$ and $\varepsilon \in\left(0, \min \left\{\frac{1}{4}, \frac{\mu_{t}(x)}{8 t^{2} \cdot w_{t}(x)}\right\}\right)$. In $O\left(n d \log \left(\frac{n}{\varepsilon}\right)\right)$ time LocalCenter $(y, t, \varepsilon)$ computes $x^{(k)}$ such that with high probability in $n / \varepsilon$.

$$
f_{t}\left(x^{(k)}\right)-\min _{\|x-y\|_{2} \leq \frac{1}{49 t}} f_{t}(x) \leq \varepsilon\left(f_{t}(y)-\min _{\|x-y\|_{2} \leq \frac{1}{49 t}} f_{t}(x)\right) .
$$

Using this local centering algorithm as well as a general result for minimizing one dimensional convex functions using a noisy oracle (See Section 7.9.3) we obtain our line search algorithm.

```
Algorithm 16: LineSearch \(\left(y, t, t^{\prime}, u, \varepsilon\right)\)
    Input: Point \(y \in \mathbb{R}^{d}\), current path parameter \(t\), next path parameter \(t^{\prime}\), bad direction \(u\), target
    accuracy \(\varepsilon\)
    Let \(\varepsilon_{O}=\frac{\varepsilon^{2} \cdot \tilde{\varepsilon}_{*}^{2}}{300 n}, \ell=-6 \widetilde{f}_{*}, u=6 \widetilde{f}_{*}\).
    Define the oracle \(q: \mathbb{R} \rightarrow \mathbb{R}\) by \(q(\alpha)=f_{t^{\prime}}\left(\right.\) LocalCenter \(\left.\left(y+\alpha u, t^{\prime}, \varepsilon_{O}\right)\right)\)
    Let \(\alpha^{\prime}=\) OneDimMinimizer \(\left(\ell, u, \varepsilon_{O}, q, t n\right)\)
    Output: \(x^{\prime}=\) LocalCenter \(\left(y+\alpha u, t^{\prime}, \varepsilon_{O}\right)\)
```

Lemma 7.4.5. Let $\frac{1}{400 f\left(x_{*}\right)} \leq t \leq t^{\prime} \leq\left(1+\frac{1}{600}\right) t \leq \frac{2 n}{\tilde{\varepsilon} \cdot \tilde{f}_{*}}$ and let $u=\operatorname{ApproxMinEig}\left(y, t, \varepsilon_{v}\right)$ for $\varepsilon_{v} \leq$ $\frac{1}{8}\left(\frac{\tilde{\varepsilon}_{*}}{7 n}\right)^{2}$ and $y \in \mathbb{R}^{d}$ such that $\left\|y-x_{t}\right\|_{2} \leq \frac{1}{t}\left(\frac{\varepsilon_{v}}{36}\right)^{\frac{3}{2}}$. In $O\left(n d \log ^{2}\left(\frac{n}{\tilde{\varepsilon} \cdot \varepsilon \cdot \varepsilon_{v}}\right)\right)$ time, LineSearch $\left(y, t, t^{\prime}, u, \varepsilon\right)$ outputs $x^{\prime}$ such that $\left\|x^{\prime}-x_{t^{\prime}}\right\|_{2} \leq \frac{\varepsilon}{t^{\prime}}$ with high probability in $n$.

We also provide the following lemma useful for finding the first center.
Lemma 7.4.6. Let $\frac{1}{400 f\left(x_{*}\right)} \leq t \leq t^{\prime} \leq\left(1+\frac{1}{600}\right) t \leq \frac{2 n}{\tilde{\varepsilon}_{*} \cdot \hat{f}_{*}}$ and let $x \in \mathbb{R}^{d}$ satisfy $\left\|x-x_{t}\right\|_{2} \leq \frac{1}{100 t}$. Then, in $O\left(n d \log ^{2}\left(\frac{n}{\varepsilon \cdot \tilde{\varepsilon}_{*}^{*}}\right)\right)$ time, LineSearch $(x, t, t, u, \varepsilon)$ output $y$ such that $\left\|y-x_{t}\right\|_{2} \leq \frac{\varepsilon}{t}$ for any vector $u \in \mathbb{R}^{d}$.

## - 7.4.3 Putting It All Together

Combining the results of the previous sections to prove our main theorem.

Theorem 7.4.1. In $O\left(n d \log ^{3}\left(\frac{n}{\varepsilon}\right)\right)$ time, Algorithm 13 outputs an $(1+\varepsilon)$-approximate geometric median with constant probability.

## ■ 7.5 Analysis of the Central Path

Here we provide proofs of the claims in Section 7.3 as well as additional technical lemmas we use throughout the chapter.

## ■ 7.5.1 Basic Facts

Here we provide basic facts regarding the central path that we will use throughout our analysis. First we compute various derivatives of the penalized objective function.
Lemma 7.5.1 (Path Derivatives). We have

$$
\begin{gathered}
\nabla f_{t}(x)=\sum_{i \in[n]} \frac{t^{2}\left(x-a^{(i)}\right)}{1+g_{t}^{(i)}(x)} \quad, \quad \nabla^{2} f_{t}(x)=\sum_{i \in[n]} \frac{t^{2}}{1+g_{t}^{(i)}(x)}\left(\boldsymbol{I}-\frac{t^{2}\left(x-a^{(i)}\right)\left(x-a^{(i)}\right)^{\top}}{g_{t}^{(i)}(x)\left(1+g_{t}^{(i)}(x)\right)}\right) \quad, \text { and } \\
\frac{d}{d t} x_{t}=-\left(\nabla^{2} f_{t}\left(x_{t}\right)\right)^{-1} \sum_{i \in[n]} \frac{t\left(x_{t}-a^{(i)}\right)}{\left(1+g_{t}^{(i)}\left(x_{t}\right)\right) g_{t}^{(i)}\left(x_{t}\right)}
\end{gathered}
$$

Proof of Lemma 7.5.1. Direct calculation shows that

$$
\begin{aligned}
\nabla f_{t}^{(i)}(x) & =\frac{t^{2}\left(x-a^{(i)}\right)}{\sqrt{1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}}-\frac{1}{1+\sqrt{1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}}\left(\frac{t^{2}\left(x-a^{(i)}\right)}{\sqrt{1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}}\right) \\
& =\frac{t^{2}\left(x-a^{(i)}\right)}{1+\sqrt{1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}}=\frac{t^{2}\left(x-a^{(i)}\right)}{1+g_{t}^{(i)}(x)}
\end{aligned}
$$

and

$$
\begin{aligned}
\nabla^{2} f_{t}^{(i)}(x) & =\frac{t^{2}}{1+\sqrt{1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}} \boldsymbol{I}-\left(\frac{1}{1+\sqrt{1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}}\right)^{2} \frac{t^{4}\left(x-a^{(i)}\right)\left(x-a^{(i)}\right)^{\top}}{\sqrt{1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}} \\
& =\frac{t^{2}}{1+g_{t}^{(i)}(x)}\left(\boldsymbol{I}-\frac{t^{2}\left(x-a^{(i)}\right)\left(x-a^{(i)}\right)^{\top}}{g_{t}^{(i)}(x)\left(1+g_{t}^{(i)}(x)\right)}\right)
\end{aligned}
$$

and

$$
\begin{aligned}
\left(\frac{d}{d t} \nabla f_{t}^{(i)}\right)(x) & =\frac{2 t\left(x-a^{(i)}\right)}{1+\sqrt{1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}}-\frac{t^{2} \cdot\left(x-a^{(i)}\right) \cdot t\left\|x-a^{(i)}\right\|_{2}^{2}}{\left(1+\sqrt{1+t^{2}\left\|x-a^{(i)}\right\|}\right)^{2} \sqrt{1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}} \\
& =\frac{t \cdot\left(x-a^{(i)}\right)}{1+g_{t}^{(i)}(x)}\left(2-\frac{g_{t}^{(i)}(x)^{2}-1}{\left(1+g_{t}^{(i)}(x)\right) g_{t}^{(i)}(x)}\right) \\
& =\frac{t \cdot\left(x-a^{(i)}\right)}{1+g_{t}^{(i)}(x)}\left(\frac{2 g_{t}^{(i)}(x)-\left(g_{t}^{(i)}(x)-1\right)}{g_{t}^{(i)}(x)}\right)=\frac{t \cdot\left(x-a^{(i)}\right)}{g_{t}^{(i)}(x)}
\end{aligned}
$$

Finally, by the optimality of $x_{t}$ we have that $\nabla f_{t}\left(x_{t}\right)=0$. Consequently,

$$
\nabla^{2} f_{t}\left(x_{t}\right) \frac{d}{d t} x_{t}+\left(\frac{d}{d t} \nabla f_{t}\right)\left(x_{t}\right)=0
$$

and solving for $\frac{d}{d t} x_{t}$ then yields

$$
\begin{aligned}
\frac{d}{d t} x_{t} & =-\left(\nabla^{2} f_{t}\left(x_{t}\right)\right)^{-1}\left(\left(\frac{d}{d t} \nabla f_{t}\right)\left(x_{t}\right)\right) \\
& =-\left(\nabla^{2} f_{t}\left(x_{t}\right)\right)^{-1}\left(\left(\frac{d}{d t} \nabla f_{t}\right)\left(x_{t}\right)-\frac{1}{t} \nabla f_{t}\left(x_{t}\right)\right) \\
& =-\left(\nabla^{2} f_{t}\left(x_{t}\right)\right)^{-1}\left(\sum_{i \in[n]}\left[\frac{t}{g_{t}^{(i)}}-\frac{t}{1+g_{t}^{(i)}}\right]\left(x_{t}-a^{(i)}\right)\right) .
\end{aligned}
$$

Next, in we provide simple facts regarding the Hessian of the penalized objective function. Lemma 7.5.2. For all $t>0$ and $x \in \mathbb{R}^{d}$

$$
\nabla^{2} f_{t}(x)=\sum_{i \in[n]} \frac{t^{2}}{1+g_{t}^{(i)}(x)}\left(\boldsymbol{I}-\left(1-\frac{1}{g_{t}^{(i)}(x)}\right) u^{(i)}(x) u^{(i)}(x)^{\top}\right)
$$

and therefore

$$
\sum_{i \in[n]} \frac{t^{2}}{\left(1+g_{t}^{(i)}(x)\right) g_{t}^{(i)}(x)} \boldsymbol{I} \preceq \nabla^{2} f_{t}(x) \preceq \sum_{i \in[n]} \frac{t^{2}}{1+g_{t}^{(i)}(x)} \boldsymbol{I}
$$

Proof of Lemma 7.5.2. We have that

$$
\begin{aligned}
\nabla^{2} f_{t}(x) & =\sum_{i \in[n]} \frac{t^{2}}{1+g_{t}^{(i)}(x)}\left(\boldsymbol{I}-\frac{t^{2}\left(x-a^{(i)}\right)\left(x-a^{(i)}\right)^{\top}}{g_{t}^{(i)}(x)\left(1+g_{t}^{(i)}(x)\right)}\right) \\
& =\sum_{i \in[n]} \frac{t^{2}}{1+g_{t}^{(i)}(x)}\left(\boldsymbol{I}-\frac{t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}{\left(1+g_{t}^{(i)}(x)\right) g_{t}^{(i)}(x)} u^{(i)}(x) u^{(i)}(x)^{\top}\right) \\
& =\sum_{i \in[n]} \frac{t^{2}}{1+g_{t}^{(i)}(x)}\left(\boldsymbol{I}-u^{(i)}(x) u^{(i)}(x)^{\top}\right)
\end{aligned}
$$

Since

$$
\frac{g_{t}^{(i)}(x)^{2}-1}{g_{t}^{(i)}(x)\left(1+g_{t}^{(i)}(x)\right)}=\frac{\left(g_{t}^{(i)}(x)+1\right)\left(g_{t}^{(i)}(x)-1\right)}{\left(1+g_{t}^{(i)}(x)\right) g_{t}^{(i)}(x)}=1-\frac{1}{g_{t}^{(i)}(x)}
$$

the result follows.

## - 7.5.2 Stability of Hessian

Here we show that moving a point $x \in \mathbb{R}^{d}$ in $\ell_{2}$, does not change the Hessian, $\nabla^{2} f_{t}(x)$, too much spectrally. First we show that such changes do not change $g_{t}^{(i)}(x)$ by too much (Lemma 7.5.3) and then we use this to prove the claim, i.e. we prove Lemma 7.3.1.

Lemma 7.5.3 (Stability of $g$ ). For all $x, y \in \mathbb{R}^{d}$ and $t>0$, we have

$$
g_{t}^{(i)}(x)-t\|x-y\|_{2} \leq g_{t}^{(i)}(y) \leq g_{t}^{(i)}(x)+t\|x-y\|_{2}
$$

Proof of Lemma 7.5.3. Direct calculation reveals that

$$
\begin{aligned}
g_{t}^{(i)}(y)^{2} & =1+t^{2}\left\|x-a^{(i)}+y-x\right\|_{2}^{2} \\
& =1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}+2 t^{2}\left(x-a^{(i)}\right)^{\top}(y-x)+t^{2}\|y-x\|_{2}^{2} \\
& =g_{t}^{(i)}(x)^{2}+2 t^{2}\left(x-a^{(i)}\right)^{\top}(y-x)+t^{2}\|y-x\|_{2}^{2} .
\end{aligned}
$$

Consequently by Cauchy Schwarz

$$
g_{t}^{(i)}(y)^{2} \leq g_{t}^{(i)}(x)^{2}+2 t^{2}\left\|x-a^{(i)}\right\|_{2} \cdot\|y-x\|_{2}+t^{2}\|y-x\|_{2}^{2} \leq\left(g_{t}^{(i)}(x)+t\|y-x\|_{2}\right)^{2}
$$

and

$$
g_{t}^{(i)}(y)^{2} \geq g_{t}^{(i)}(x)^{2}-2 t^{2}\left\|x-a^{(i)}\right\|_{2} \cdot\|y-x\|_{2}+t^{2}\|y-x\|_{2}^{2} \geq\left(g_{t}^{(i)}(x)-t\|y-x\|_{2}\right)^{2}
$$

Lemma 7.3.1. Suppose that $\|x-y\|_{2} \leq \frac{\varepsilon}{t}$ with $\varepsilon \leq \frac{1}{20}$. Then, we have

$$
\left(1-6 \varepsilon^{2 / 3}\right) \nabla^{2} f_{t}(x) \preceq \nabla^{2} f_{t}(y) \preceq\left(1+6 \varepsilon^{2 / 3}\right) \nabla^{2} f_{t}(x) .
$$

Proof of Lemma 7.3.1. Here we prove the following stronger statement, for all $i \in[n]$

$$
\left(1-6 \varepsilon^{2 / 3}\right) \nabla^{2} f_{t}^{(i)}(x) \preceq \nabla^{2} f_{t}^{(i)}(y) \preceq\left(1+6 \varepsilon^{2 / 3}\right) \nabla^{2} f_{t}^{(i)}(x) .
$$

Without loss of generality let $y-x=\alpha v+\beta u^{(i)}(x)$ for some $v \perp u^{(i)}(x)$ with $\|v\|_{2}=1$. Since $\|x-y\|_{2}^{2} \leq \frac{\varepsilon^{2}}{t^{2}}$, we know that $\alpha^{2}, \beta^{2} \leq \frac{\varepsilon^{2}}{t^{2}}$. Also, let $\bar{x}=x+\beta u^{(i)}(x)$, so that clearly, $u^{(i)}(x)=u^{(i)}(\bar{x})$. Now some manipulation reveals that for all unit vectors $z \in \mathbb{R}^{d}$ the following holds (so long as $u^{(i)}(x) \neq 0$ and $\left.u^{(i)}(y) \neq 0\right)$

$$
\begin{aligned}
& \left|\left[u^{(i)}(x)^{\top} z\right]^{2}-\left[u^{(i)}(y) z\right]^{2}\right| \\
& =\left|\left[u^{(i)}(\bar{x})^{\top} z\right]^{2}-\left[u^{(i)}(y) z\right]^{2}\right| \\
& =\left|\left[\frac{\left(\bar{x}-a^{(i)}\right)^{\top} z}{\left\|\bar{x}-a^{(i)}\right\|_{2}}\right]^{2}-\left[\frac{\left(y-a^{(i)}\right)^{\top} z}{\left\|y-a^{(i)}\right\|_{2}}\right]^{2}\right| \\
& \leq\left|\left[\frac{\left(\bar{x}-a^{(i)}\right)^{\top} z}{\left\|\bar{x}-a^{(i)}\right\|_{2}}\right]^{2}-\left[\frac{\left(\bar{x}-a^{(i)}\right)^{\top} z}{\left\|y-a^{(i)}\right\|_{2}}\right]^{2}\right|+\left|\left[\frac{\left(\bar{x}-a^{(i)}\right)^{\top} z}{\left\|y-a^{(i)}\right\|_{2}}\right]^{2}-\left[\frac{\left(y-a^{(i)}\right)^{\top} z}{\left\|y-a^{(i)}\right\|_{2}}\right]^{2}\right| \\
& \leq\left|1-\frac{\left\|\bar{x}-a^{(i)}\right\|_{2}^{2}}{\left\|y-a^{(i)}\right\|_{2}^{2}}\right|+\frac{\mid\left[\left(\bar{x}-a^{(i)}+\alpha v^{\top} z\right]^{2}-\left[\left(\bar{x}-a^{(i)}\right)^{\top} z\right]^{2} \mid\right.}{\left\|y-a^{(i)}\right\|_{2}^{2}} \\
& =\frac{\alpha^{2}+\left|2\left[\left(\bar{x}-a^{(i)}\right)^{\top} z\right] \cdot\left[\alpha v^{\top} z\right]+\left[\alpha v^{\top} z\right]^{2}\right|}{\left\|\bar{x}-a^{(i)}\right\|_{2}^{2}+\alpha_{i}^{2}} .
\end{aligned}
$$

Where we used that $y=\bar{x}+\alpha v$ and $\left\|y-a^{(i)}\right\|_{2}^{2}=\alpha^{2}+\left\|\bar{x}-a^{(i)}\right\|_{2}^{2}\left(\right.$ since $\left.v \perp\left(\bar{x}-a^{(i)}\right)\right)$. Now we know that $\alpha^{2} \leq \frac{\varepsilon^{2}}{t^{2}}$ and and therefore, by Young's inequality and Cauchy Schwarz we have that for all $\gamma>0$

$$
\begin{align*}
\left|\left[u^{(i)}(x)^{\top} z\right]^{2}-\left[u^{(i)}(y) z\right]^{2}\right| & \leq \frac{2 \alpha^{2}+2\left|\left[\left(\bar{x}-a^{(i)}\right)^{\top} z\right] \cdot\left[\alpha v^{\top} z\right]\right|}{\left\|\bar{x}-a^{(i)}\right\|_{2}^{2}+\alpha^{2}} \\
& \leq \frac{2 \alpha^{2}+\gamma\left[\left(\bar{x}-a^{(i)}\right)^{\top} z\right]^{2}+\gamma^{-1} \alpha^{2}\left[v^{\top} z\right]^{2}}{\left\|\bar{x}-a^{(i)}\right\|_{2}^{2}+\alpha^{2}} \\
& \leq \frac{\alpha^{2}\left(2+\gamma^{-1}\left(v^{\top} z\right)^{2}\right)}{\left\|\bar{x}-a^{(i)}\right\|_{2}^{2}+\alpha^{2}}+\gamma\left[\left(u^{(i)}(x)\right)^{\top} z\right]^{2} \\
& \leq \frac{\varepsilon^{2}}{t^{2}\left\|\bar{x}-a^{(i)}\right\|_{2}^{2}+\varepsilon^{2}}\left(2+\frac{1}{\gamma}\left(v^{\top} z\right)^{2}\right)+\gamma\left[\left(u^{(i)}(x)\right)^{\top} z\right]^{2} \tag{7.6}
\end{align*}
$$

Note that

$$
\begin{aligned}
t^{2}\left\|\bar{x}-a^{(i)}\right\|_{2}^{2} & =t^{2}\left(\left\|x-a^{(i)}\right\|_{2}^{2}+2 \beta\left(x-a^{(i)}\right)^{\top} u^{(i)}(x)+\beta^{2}\right)=\left(t\left\|x-a^{(i)}\right\|_{2}+t \beta\right)^{2} \\
& \geq\left(\max \left\{t\left\|x-a^{(i)}\right\|_{2}-\varepsilon, 0\right\}\right)^{2}
\end{aligned}
$$

Consequently, if $t\left\|x-a^{(i)}\right\|_{2} \geq 2 \varepsilon^{1 / 3} \sqrt{g_{t}^{(i)}(x)}$ then since $\varepsilon \leq \frac{1}{20}$ we have that $\left\|x-a^{(i)}\right\|_{2} \geq 2 \varepsilon$ and $\left\|y-a^{(i)}\right\| \geq \varepsilon$, justifying our assumption that $u^{(i)}(x) \neq 0$ and $u^{(i)}(y) \neq 0$. Furthermore, this implies that $t^{2}\left\|\bar{x}-a^{(i)}\right\|_{2}^{2} \geq 2 \varepsilon^{2 / 3} g_{t}^{(i)}(x)$ and therefore letting $\gamma=\frac{\varepsilon^{2 / 3}}{g_{t}^{(i)}(x)}$ yields

$$
\begin{aligned}
\left|\left[u_{t}^{(i)}(x)^{\top} z\right]^{2}-\left[u_{t}^{(i)}(y) z\right]^{2}\right| & \leq \frac{\varepsilon^{4 / 3}}{2 g_{t}^{(i)}(x)}\left(2+\frac{g_{t}^{(i)}(x)}{\varepsilon^{2 / 3}}\left[v^{\top} z\right]^{2}\right)+\frac{\varepsilon^{2 / 3}}{g_{t}^{(i)}(x)}\left[\left(u^{(i)}(x)\right)^{\top} z\right]^{2} \\
& \leq 2 \varepsilon^{2 / 3}\left[v^{\top} z\right]^{2}+\frac{\varepsilon^{2 / 3}}{g_{t}^{(i)}(x)}\left[\left(u^{(i)}(x)\right)^{\top} z\right]^{2} \\
& \leq 2 \varepsilon^{2 / 3}\left(\frac{1+g_{t}^{(i)}(x)}{t^{2}}\right)\|z\|_{\nabla^{2} f_{t}^{(i)}(x)}^{2}
\end{aligned}
$$

and therefore if we let

$$
\boldsymbol{H} \stackrel{\text { def }}{=} \frac{t^{2}}{1+g_{t}^{(i)}(x)}\left(\boldsymbol{I}-\left(1-\frac{1}{g_{t}^{(i)}(x)}\right) u^{(i)}(y)\left(u^{(i)}(y)\right)^{\top}\right)
$$

we see that for unit vectors $z$,

$$
\left|z^{\top}\left(\boldsymbol{H}-\nabla^{2} f_{t}^{(i)}(x)\right) z\right| \leq 2 \varepsilon^{2 / 3}\|z\|_{\nabla^{2} f_{t}^{(i)}(x)}^{2}
$$

Otherwise, $t\left\|x-a^{(i)}\right\|_{2}<2 \varepsilon^{1 / 3} \sqrt{g_{t}^{(i)}(x)}$ and therefore

$$
g_{t}^{(i)}(x)^{2}=1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2} \leq 1+4 \varepsilon^{2 / 3} g_{t}^{(i)}(x)
$$

Therefore, we have

$$
g_{t}^{(i)}(x) \leq \frac{4 \varepsilon^{2 / 3}+\sqrt{\left(4 \varepsilon^{2 / 3}\right)^{2}+4}}{2} \leq 1+4 \varepsilon^{2 / 3}
$$

Therefore independent of (7.6) and the assumption that $u^{(i)}(x) \neq 0$ and $u^{(i)}(y) \neq 0$ we have

$$
\frac{1}{1+4 \varepsilon^{2 / 3}} \boldsymbol{H} \preceq \frac{t^{2}}{\left(1+g_{t}^{(i)}(x)\right) g_{t}^{(i)}(x)} \boldsymbol{I} \preceq \nabla^{2} f_{t}^{(i)}(x) \preceq \frac{t^{2}}{\left(1+g_{t}^{(i)}(x)\right)} \boldsymbol{I} \preceq\left(1+4 \varepsilon^{2 / 3}\right) \boldsymbol{H} .
$$

In either case, we have that

$$
\left|z^{\top}\left(\boldsymbol{H}-\nabla^{2} f_{t}^{(i)}(x)\right) z\right| \leq 4 \varepsilon^{2 / 3}\|z\|_{\nabla^{2} f_{t}^{(i)}(x)}^{2}
$$

Now, we note that $\|x-y\|_{2} \leq \frac{\varepsilon}{t} \leq \varepsilon \cdot \frac{g_{t}^{(i)}(x)}{t}$. Therefore, by Lemma 7.5 .3 we have that

$$
(1-\varepsilon) g_{t}^{(i)}(x) \leq g_{t}^{(i)}(y) \leq(1+\varepsilon) g_{t}^{(i)}(x)
$$

Therefore, we have

$$
\frac{1-4 \varepsilon^{2 / 3}}{(1+\varepsilon)^{2}} \nabla^{2} f_{t}^{(i)}(x) \preceq \frac{1}{(1+\varepsilon)^{2}} \boldsymbol{H} \preceq \nabla^{2} f_{t}^{(i)}(y) \preceq \frac{1}{(1-\varepsilon)^{2}} \boldsymbol{H} \preceq \frac{1+4 \varepsilon^{2 / 3}}{(1-\varepsilon)^{2}} \nabla^{2} f_{t}^{(i)}(x)
$$

Since $\varepsilon<\frac{1}{20}$, the result follows.
Consequently, so long as we have a point within a $O\left(\frac{1}{t}\right)$ sized Euclidean ball of some $x_{t}$ Newton's method (or an appropriately transformed first order method) within the ball will converge quickly.

## ■ 7.5.3 How Much Does the Hessian Change Along the Path?

Lemma 7.3.2. For all $t \geq 0$ and $i \in[n]$ the following hold

$$
\left\|\frac{d}{d t} x_{t}\right\|_{2} \leq \frac{1}{t^{2}} \bar{g}_{t}\left(x_{t}\right) \quad, \quad\left|\frac{d}{d t} g_{t}^{(i)}\left(x_{t}\right)\right| \leq \frac{1}{t}\left(g_{t}^{(i)}\left(x_{t}\right)+\bar{g}_{t}\right) \quad, \text { and } \quad\left|\frac{d}{d t} w_{t}\right| \leq \frac{2}{t} w_{t}
$$

Consequently, for all $t^{\prime} \geq t$ we have that $\left(\frac{t}{t^{\prime}}\right)^{2} w_{t} \leq w_{t^{\prime}} \leq\left(\frac{t^{\prime}}{t}\right)^{2} w_{t}$.
Proof of Lemma 7.3.2. From Lemma 7.5.1 we know that

$$
\frac{d}{d t} x_{t}=-\left(\nabla^{2} f_{t}\left(x_{t}\right)\right)^{-1} \sum_{i \in[n]} \frac{t\left(x_{t}-a^{(i)}\right)}{\left(1+g_{t}^{(i)}\left(x_{t}\right)\right) g_{t}^{(i)}\left(x_{t}\right)}
$$

and by Lemma 7.5.2 we know that

$$
\nabla^{2} f_{t}\left(x_{t}\right) \succeq \sum_{i \in[n]} \frac{t^{2}}{\left(1+g_{t}^{(i)}\left(x_{t}\right)\right) g_{t}^{(i)}\left(x_{t}\right)} \boldsymbol{I}=\frac{t^{2}}{\bar{g}_{t}\left(x_{t}\right)} \sum_{i \in[n]} \frac{1}{1+g_{t}^{(i)}\left(x_{t}\right)} \boldsymbol{I}
$$

Using this fact and the fact that $t\left\|x_{t}-a^{(i)}\right\|_{2} \leq g_{t}^{(i)}$ we have

$$
\begin{aligned}
\left\|\frac{d}{d t} x_{t}\right\|_{2} & =\left\|-\left(\nabla^{2} f_{t}\left(x_{t}\right)\right)^{-1} \frac{d}{d t} \nabla f_{t}\left(x_{t}\right)\right\|_{2} \\
& \leq\left(\frac{t^{2}}{\bar{g}_{t}\left(x_{t}\right)} \sum_{i \in[n]} \frac{1}{1+g_{t}^{(i)}\left(x_{t}\right)}\right)^{-1} \sum_{i \in[n]}\left\|\frac{t\left(x_{t}-a^{(i)}\right)}{g_{t}^{(i)}\left(x_{t}\right)\left(1+g_{t}^{(i)}\left(x_{t}\right)\right)}\right\|_{2} \leq \frac{\bar{g}_{t}\left(x_{t}\right)}{t^{2}} .
\end{aligned}
$$

Next, we have

$$
\begin{aligned}
\frac{d}{d t} g_{t}^{(i)}\left(x_{t}\right) & =\frac{d}{d t}\left(1+t^{2}\left\|x_{t}-a^{(i)}\right\|_{2}^{2}\right)^{\frac{1}{2}} \\
& =\frac{1}{2} \cdot g_{t}^{(i)}\left(x_{t}\right)^{-1}\left(2 t\left\|x_{t}-a^{(i)}\right\|_{2}^{2}+2 t^{2}\left(x_{t}-a^{(i)}\right)^{\top} \frac{d}{d t} x_{t}\right)
\end{aligned}
$$

which by Cauchy Schwarz and that $t\left\|x_{t}-a^{(i)}\right\|_{2} \leq g_{t}^{(i)}\left(x_{t}\right)$ yields the second equation. Furthermore,

$$
\begin{aligned}
\left|\frac{d}{d t} w_{t}\right| & =\left|\frac{d}{d t} \sum_{i \in[n]} \frac{1}{1+g_{t}^{(i)}\left(x_{t}\right)}\right| \leq \sum_{i \in[n]}\left|\frac{d}{d t} \frac{1}{1+g_{t}^{(i)}\left(x_{t}\right)}\right|=\sum_{i \in[n]}\left|\frac{1}{\left(1+g_{t}\left(x_{t}\right)\right)^{2}} \frac{d}{d t} g_{t}\left(x_{t}\right)\right| \\
& \leq \frac{1}{t} \sum_{i \in[n]} \frac{g_{t}^{(i)}\left(x_{t}\right)+\bar{g}_{t}}{\left(1+g_{t}^{(i)}\left(x_{t}\right)\right) g_{t}^{(i)}\left(x_{t}\right)} \leq 2 \frac{w_{t}}{t}
\end{aligned}
$$

which yields the third equation.
Finally, using our earlier results and Jensen's inequality yields that

$$
\left|\ln w_{t^{\prime}}-\ln w_{t}\right|=\left|\int_{t}^{t^{\prime}} \frac{\frac{d}{d \alpha} w_{\alpha}}{w_{\alpha}} d \alpha\right| \leq \int_{t}^{t^{\prime}} \frac{\left(2 \frac{w_{\alpha}}{\alpha}\right)}{w_{\alpha}} d \alpha=2 \int_{t}^{t^{\prime}} \frac{1}{\alpha} d \alpha=\ln \left(\frac{t^{\prime}}{t}\right)^{2}
$$

Exponentiating the above inequality yields the final inequality.

Lemma 7.3.3. For all $t \geq 0$ we have

$$
\begin{equation*}
-12 \cdot t \cdot w_{t} \boldsymbol{I} \preceq \frac{d}{d t}\left[\nabla^{2} f_{t}\left(x_{t}\right)\right] \preceq 12 \cdot t \cdot w_{t} \boldsymbol{I} \tag{7.3}
\end{equation*}
$$

and therefore for all $\beta \in\left[0, \frac{1}{8}\right]$

$$
\begin{equation*}
\nabla^{2} f\left(x_{t}\right)-15 \beta t^{2} w_{t} \boldsymbol{I} \preceq \nabla^{2} f\left(x_{t(1+\beta)}\right) \preceq \nabla^{2} f\left(x_{t}\right)+15 \beta t^{2} w_{t} \boldsymbol{I} \tag{7.4}
\end{equation*}
$$

Proof of Lemma 7.3.3. Let

$$
\boldsymbol{A}_{t}^{(i)} \stackrel{\text { def }}{=} \frac{t^{2}\left(x_{t}-a^{(i)}\right)\left(x_{t}-a^{(i)}\right)^{\top}}{\left(1+g_{t}^{(i)}\right) g_{t}^{(i)}}
$$

and recall that $\nabla^{2} f_{t}\left(x_{t}\right)=\sum_{i \in[n]} \frac{t^{2}}{1+g_{t}^{(i)}}\left(\boldsymbol{I}-\boldsymbol{A}_{t}^{(i)}\right)$. Consequently

$$
\begin{aligned}
\frac{d}{d t} \nabla^{2} f_{t}\left(x_{t}\right) & =\frac{d}{d t}\left(\sum_{i \in[n]} \frac{t^{2}}{1+g_{t}^{(i)}}\left(\boldsymbol{I}-\boldsymbol{A}_{t}^{(i)}\right)\right) \\
& =2 t\left(\frac{1}{t^{2}}\right) \nabla^{2} f_{t}\left(x_{t}\right)+t^{2} \sum_{i \in[n]} \frac{-\frac{d}{d t} g_{t}^{(i)}}{\left(1+g_{t}^{(i)}\right)^{2}}\left(\boldsymbol{I}-\boldsymbol{A}_{t}^{(i)}\right)-\sum_{i \in[n]} \frac{t^{2}}{1+g_{t}^{(i)}} \frac{d}{d t} \boldsymbol{A}_{t}^{(i)}
\end{aligned}
$$

Now, since $\mathbf{0} \preceq \boldsymbol{A}_{t}^{(i)} \preceq \boldsymbol{I}$ we have $0 \preceq \nabla^{2} f_{t}\left(x_{t}\right) \preceq t^{2} w_{t} \boldsymbol{I}$. For all unit vectors $v$, using Lemma 7.3.2
yields

$$
\begin{aligned}
\left|v^{\top}\left(\frac{d}{d t} \nabla^{2} f_{t}\left(x_{t}\right)\right) v\right| & \leq 2 t \cdot w_{t} \cdot\|v\|_{2}^{2}+t^{2} \sum_{i \in[n]} \frac{\left|\frac{d}{d t} g_{t}^{(i)}\right|}{\left(1+g_{t}^{(i)}\right)^{2}}\|v\|_{2}^{2}+\sum_{i \in[n]} \frac{t^{2}}{1+g_{t}^{(i)}}\left|v^{\top}\left(\frac{d}{d t} \boldsymbol{A}_{t}^{(i)}\right) v\right| \\
& \leq 4 t \cdot w_{t}+\sum_{i \in[n]} \frac{t^{2}}{1+g_{t}^{(i)}}\left|v^{\top}\left(\frac{d}{d t} \boldsymbol{A}_{t}^{(i)}\right) v\right|
\end{aligned}
$$

Next

$$
\begin{gathered}
\frac{d}{d t} \boldsymbol{A}_{t}^{(i)}=2 t\left(\frac{1}{t^{2}}\right) \boldsymbol{A}_{t}^{(i)}-\left(\frac{t}{\left(1+g_{t}^{(i)}\right) g_{t}^{(i)}}\right)^{2}\left[\left(1+g_{t}^{(i)}\right) \frac{d}{d t} g_{t}^{(i)}+g_{t}^{(i)} \frac{d}{d t} g_{t}^{(i)}\right]\left(x_{t}-a^{(i)}\right)\left(x_{t}-a^{(i)}\right)^{\top} \\
+\frac{t^{2}}{\left(1+g_{t}^{(i)}\right) g_{t}^{(i)}}\left[\left(x_{t}-a^{(i)}\right)\left(\frac{d}{d t} x_{t}\right)^{\top}+\left(\frac{d}{d t} x_{t}\right)\left(x_{t}-a^{(i)}\right)^{\top}\right],
\end{gathered}
$$

and therefore by Lemma 7.3.2 and the that $t\left\|x_{t}-a^{(i)}\right\|_{2} \leq g_{t}^{(i)}$ we have

$$
\begin{aligned}
\left|v^{\top}\left(\frac{d}{d t} \boldsymbol{A}_{t}^{(i)}\right) v\right| & \leq\left(\frac{2}{t}+\frac{2 t^{2}\left|\frac{d}{d t} g_{t}^{(i)}\right|}{\left(1+g_{t}^{(i)}\right)\left(g_{t}^{(i)}\right)^{2}}\left\|x_{t}-a^{(i)}\right\|_{2}^{2}+\frac{2 t^{2}\left\|x_{t}-a^{(i)}\right\|_{2}\left\|\frac{d}{d t} x_{t}\right\|_{2}}{\left(1+g_{t}^{(i)}\right) g_{t}^{(i)}}\right)\|v\|_{2}^{2} \\
& \leq \frac{2}{t}+\frac{2}{t} \cdot \frac{g_{t}^{(i)}+\bar{g}_{t}}{1+g_{t}^{(i)}}+\frac{2}{t} \cdot \frac{\bar{g}_{t}}{1+g_{t}^{(i)}} \leq \frac{4}{t}+\frac{4}{t} \frac{\bar{g}_{t}}{1+g_{t}^{(i)}}
\end{aligned}
$$

Consequently, we have

$$
\left|v^{\top}\left(\frac{d}{d t} \nabla^{2} f_{t}\left(x_{t}\right)\right) v\right| \leq 8 t \cdot w_{t}+4 t \sum_{i \in[n]} \frac{\bar{g}_{t}}{\left(1+g_{t}^{(i)}\right)^{2}} \leq 12 t \cdot w_{t}
$$

which completes the proof of (7.3). To prove (7.4), let $v$ be any unit vector and note that

$$
\begin{aligned}
\left|v^{\top}\left(\nabla^{2} f_{t(1+\beta)}(x)-\nabla^{2} f_{t}(x)\right) v\right| & =\left|\int_{t}^{t(1+\beta)} v^{\top} \frac{d}{d t}\left[\nabla^{2} f_{\alpha}\left(x_{\alpha}\right)\right] v \cdot d \alpha\right| \leq 12 \int_{t}^{t(1+\beta)} \alpha \cdot w_{\alpha} d \alpha \\
& \leq 12 \int_{t}^{t(1+\beta)} \alpha\left(\frac{\alpha}{t}\right)^{2} w_{t} d \alpha \leq \frac{12}{t^{2}}\left(\frac{1}{4}[t(1+\beta)]^{4}-\frac{1}{4} t^{4}\right) w_{t} \\
& =3 t^{2}\left[(1+\beta)^{4}-1\right] w_{t} \leq 15 t^{2} \beta w_{t}
\end{aligned}
$$

where we used Lemma 7.3.3 and $0 \leq \beta \leq \frac{1}{8}$ at the last line.

## - 7.5.4 Where is the next Optimal Point?

Lemma 7.5.4. For all $t$ we have $\frac{1}{2}\left[t^{2} \cdot w_{t} \boldsymbol{I}-\left(t^{2} \cdot w_{t}-\mu_{t}\right) v_{t} v_{t}^{\top}\right] \preceq \nabla^{2} f_{t}\left(x_{t}\right) \preceq t^{2} \cdot w_{t} \boldsymbol{I}-\left(t^{2} \cdot w_{t}-\mu_{t}\right) v_{t} v_{t}^{\top}$
Proof of Lemma 7.5.4. This follows immediately from Lemma 7.5.2, regarding the hessian of the penalized objective function, and Lemma 7.9.1, regarding the sum of PSD matrices expressed as the identity matrix minus a rank 1 matrix.

Lemma 7.3.5 (The Central Path is Almost Straight). For all $t \geq 0, \beta \in\left[0, \frac{1}{600}\right]$, and any unit vector $y$ with $\left|\left\langle y, v_{t}\right\rangle\right| \leq \frac{1}{t^{2} \cdot \kappa}$ where $\kappa=\max _{\delta \in[t,(1+\beta) t]} \frac{w_{\delta}}{\mu_{\delta}}$, we have $y^{\top}\left(x_{(1+\beta) t}-x_{t}\right) \leq \frac{6 \beta}{t}$.

Proof of Lemma 7.3.5. Clearly

$$
\begin{aligned}
y^{\top}\left(x_{(1+\beta) t}-x_{t}\right) & =\int_{t}^{(1+\beta) t} y^{\top} \frac{d}{d \alpha} x_{\alpha} d_{\alpha} \leq \int_{\beta}^{(1+\beta) t}\left|y^{\top} \frac{d}{d \alpha} x_{\alpha}\right| d_{\alpha} \\
& \leq \int_{t}^{(1+\beta) t}\left|y^{\top}\left(\nabla^{2} f_{\alpha}\left(x_{\alpha}\right)\right)^{-1} \sum_{i \in[n]} \frac{\alpha}{\left(1+g_{\alpha}^{(i)}\right) g_{\alpha}^{(i)}}\left(x_{\alpha}-a^{(i)}\right)\right| d_{\alpha} \\
& \leq \int_{t}^{(1+\beta) t}\left\|\left(\nabla^{2} f_{\alpha}\left(x_{\alpha}\right)\right)^{-1} y\right\|_{2} \cdot\left\|\sum_{i \in[n]} \frac{\alpha}{\left(1+g_{\alpha}^{(i)}\right) g_{\alpha}^{(i)}}\left(x_{\alpha}-a^{(i)}\right)\right\|_{2} d_{\alpha}
\end{aligned}
$$

Now since clearly $\alpha\left\|x_{\alpha}-a^{(i)}\right\|_{2} \leq g_{\alpha}^{(i)}$, invoking Lemma 7.3.2 yields that

$$
\left\|\sum_{i \in[n]} \frac{\alpha\left(x_{\alpha}-a^{(i)}\right)}{\left(1+g_{\alpha}^{(i)}\right) g_{\alpha}^{(i)}}\right\|_{2} \leq \sum_{i \in[n]} \frac{1}{1+g_{\alpha}^{(i)}}=w_{\alpha} \leq\left(\frac{\alpha}{t}\right)^{2} w_{t} .
$$

Now by invoking Lemma 7.3.3 and the Lemma 7.5.4, we have that

$$
\nabla^{2} f_{\alpha}\left(x_{\alpha}\right) \succeq \nabla^{2} f_{t}\left(x_{t}\right)-15 \beta t^{2} w_{t} \boldsymbol{I} \succeq \frac{1}{2}\left[t^{2} \cdot w_{t} \boldsymbol{I}-\left(t^{2} \cdot w_{t}-\mu_{t}\right) v_{t} v_{t}^{\top}\right]-15 \beta t^{2} w_{t} \boldsymbol{I} .
$$

For notational convenience let $\boldsymbol{H}_{t} \stackrel{\text { def }}{=} \nabla^{2} f_{t}\left(x_{t}\right)$ for all $t>0$. Then Lemma 7.3.3 shows that $\boldsymbol{H}_{\alpha}=$ $\boldsymbol{H}_{t}+\Delta_{\alpha}$ where $\left\|\Delta_{\alpha}\right\|_{2} \leq 15 \beta t^{2} w_{t}$. Using $\boldsymbol{H}_{\alpha}^{2}=\boldsymbol{H}_{t}^{2}+\Delta_{\alpha} \boldsymbol{H}_{t}+\boldsymbol{H}_{t} \Delta_{\alpha}+\Delta_{\alpha}^{2}$, we have

$$
\begin{aligned}
\left\|\boldsymbol{H}_{\alpha}^{2}-\boldsymbol{H}_{t}^{2}\right\|_{2} & \leq\left\|\Delta_{\alpha} \boldsymbol{H}_{t}\right\|_{2}+\left\|\boldsymbol{H}_{t} \Delta_{\alpha}\right\|_{2}+\left\|\Delta_{\alpha}^{2}\right\|_{2} \\
& \leq 2\|\Delta\|_{2}\left\|\boldsymbol{H}_{t}\right\|_{2}+\|\Delta\|_{2}^{2} \leq 40 \beta t^{4} w_{t}^{2}
\end{aligned}
$$

Let $S$ be the subspace orthogonal to $v_{t}$. Then, Lemma 7.5 .4 shows that $\boldsymbol{H}_{t} \succeq \frac{1}{2} t^{2} w_{t} \boldsymbol{I}$ on $S$ and hence $\boldsymbol{H}_{t}^{2} \succeq \frac{1}{4} t^{4} w_{t}^{2} \boldsymbol{I}$ on $S .{ }^{4}$ Since $\left\|\boldsymbol{H}_{\alpha}^{2}-\boldsymbol{H}_{t}^{2}\right\|_{2} \leq 40 \beta t^{4} w_{t}^{2}$, we have that

$$
\boldsymbol{H}_{\alpha}^{2} \succeq \frac{1}{4} t^{4} w_{t}^{2}-40 \beta t^{4} w_{t}^{2} \boldsymbol{I} \text { on } S
$$

and hence

$$
\boldsymbol{H}_{\alpha}^{-2} \preceq\left(\frac{1}{4} t^{4} w_{t}^{2}-40 \beta t^{4} w_{t}^{2}\right)^{-1} \quad \boldsymbol{I} \text { on } S .
$$

Therefore, for any $z \in S$, we have

$$
\left\|\left(\nabla^{2} f_{\alpha}\left(x_{\alpha}\right)\right)^{-1} z\right\|_{2}=\left\|\boldsymbol{H}_{\alpha}^{-1} z\right\|_{2} \leq \frac{\|z\|_{2}}{\sqrt{\frac{1}{4} t^{4} w_{t}^{2}-40 \beta t^{4} w_{t}^{2}}}
$$

Now, we split $y=z+\left\langle y, v_{t}\right\rangle v_{t}$ where $z \in S$. Then, we have that

$$
\begin{aligned}
\left\|\left(\nabla^{2} f_{\alpha}\left(x_{\alpha}\right)\right)^{-1} y\right\|_{2} & \leq\left\|\left(\nabla^{2} f_{\alpha}\left(x_{\alpha}\right)\right)^{-1} z\right\|_{2}+\left|\left\langle y, v_{t}\right\rangle\right|\left\|\left(\nabla^{2} f_{\alpha}\left(x_{\alpha}\right)\right)^{-1} v_{t}\right\|_{2} \\
& \leq \frac{1}{\sqrt{\frac{1}{4} t^{4} w_{t}^{2}-40 \beta t^{4} w_{t}^{2}}}+\frac{1}{t^{2} \cdot \kappa}\left\|\left(\nabla^{2} f_{\alpha}\left(x_{\alpha}\right)\right)^{-1} v_{t}\right\|_{2}
\end{aligned}
$$

Note that, we also know that $\lambda_{\min }\left(\nabla^{2} f_{\alpha}\left(x_{\alpha}\right)\right) \geq \mu_{\alpha}$ and hence $\lambda_{\max }\left(\nabla^{2} f_{\alpha}\left(x_{\alpha}\right)^{-2}\right) \leq \mu_{\alpha}^{-2}$. Therefore,

[^16]we have
$$
\left\|\left(\nabla^{2} f_{\alpha}\left(x_{\alpha}\right)\right)^{-1} y\right\|_{2} \leq \frac{1}{t^{2} w_{t} \sqrt{\frac{1}{4}-40 \beta}}+\frac{1}{t^{2}} \frac{\mu_{\alpha}}{w_{\alpha}} \frac{1}{\mu_{\alpha}} \leq \frac{1}{t^{2} w_{t}}\left(2+\frac{1}{\sqrt{\frac{1}{4}-40 \beta}}\right) \leq \frac{5}{t^{2} w_{t}}
$$

Combining these and using that $\beta \in[0,1 / 600]$ yields that

$$
\begin{aligned}
y^{\top}\left(x_{(1+\beta) t}-x_{t}\right) & \leq \int_{t}^{(1+\beta) t} \frac{5}{t^{2} w_{t}}\left(\frac{\alpha}{t}\right)^{2} w_{t} d_{\alpha} \leq \frac{5}{t^{4}}\left(\frac{1}{3}(1+\beta)^{3} t^{3}-\frac{1}{3} t^{3}\right) \\
& \leq \frac{5}{3 t}\left[(1+\beta)^{3}-1\right] \leq \frac{6 \beta}{t}
\end{aligned}
$$

### 7.5.5 Where is the End?

Lemma 7.3.6. $f\left(x_{t}\right)-f\left(x_{*}\right) \leq \frac{2 n}{t}$ for all $t>0$.
Proof of Lemma 7.3.6. Clearly, $\nabla f_{t}\left(x_{t}\right)=0$ by definition of $x_{t}$. Consequently $\frac{1}{t} \nabla f_{t}\left(x_{t}\right)^{\top}\left(x_{t}-x_{*}\right)=0$ and using Lemma 7.5 .1 to give the formula for $\nabla f_{t}\left(x_{t}\right)$ yields

$$
0=\sum_{i \in[n]} \frac{t\left(x_{t}-a^{(i)}\right)^{\top}\left(x_{t}-x_{*}\right)}{1+g_{t}^{(i)}(x)}=\sum_{i \in[n]} \frac{t\left\|x_{t}-a^{(i)}\right\|_{2}^{2}+t\left(x_{t}-a^{(i)}\right)^{\top}\left(a^{(i)}-x_{*}\right)}{1+g_{t}^{(i)}\left(x_{t}\right)}
$$

Therefore, by Cauchy Schwarz and the fact that $t\left\|x_{t}-a^{(i)}\right\|_{2} \leq g_{t}^{(i)}\left(x_{t}\right) \leq 1+g_{t}^{(i)}$

$$
\sum_{i \in[n]} \frac{t\left\|x_{t}-a^{(i)}\right\|_{2}^{2}}{1+g_{t}^{(i)}\left(x_{t}\right)} \leq \sum_{i \in[n]} \frac{t\left\|x_{t}-a^{(i)}\right\|_{2}\left\|a^{(i)}-x_{*}\right\|_{2}}{1+g_{t}^{(i)}\left(x_{t}\right)} \leq f\left(x_{*}\right)
$$

Furthermore, since $1+g_{t}^{(i)}\left(x_{t}\right) \leq 2+t\left\|x_{t}-a^{(i)}\right\|_{2}$ we have

$$
\sum_{i \in[n]} \frac{t\left\|x_{t}-a^{(i)}\right\|_{2}^{2}}{1+g_{t}^{(i)}\left(x_{t}\right)} \geq \sum_{i \in[n]}\left\|x_{t}-a^{(i)}\right\|_{2}-\sum_{i \in[n]} \frac{2\left\|x_{t}-a^{(i)}\right\|_{2}}{1+g_{t}^{(i)}\left(x_{t}\right)} \geq f\left(x_{t}\right)-\frac{2 n}{t}
$$

Combining yields the result.

## - 7.5.6 Simple Lemmas

Here we provide various small technical Lemmas that we will use to bound the accuracy with which we need to carry out various operations in our algorithm. Here we use some notation from Section 7.4 to simplify our bounds and make them more readily applied.
Lemma 7.5.5. For any $x$, we have that $\left\|x-x_{t}\right\|_{2} \leq f(x)$.

Proof of Lemma 7.5.5. Since $\sum_{i \in[n]}\left\|x-a^{(i)}\right\|_{2}=f(x)$, we have that $\left\|x-a^{(i)}\right\|_{2} \leq f(x)$ for all $i \in[n]$. Since $\nabla f\left(x_{t}\right)=0$ by Lemma 7.5 .1 we see that $x_{t}$ is a convex combination of the $a^{(i)}$ and therefore $\left\|x-x_{t}\right\|_{2} \leq f(x)$ by convexity.

Lemma 7.5.6. $x^{(0)}=\frac{1}{n} \sum_{i \in[n]} a^{(i)}$ is a 2-approximate geometric median, i.e. $\tilde{f}_{*} \leq 2 \cdot f\left(x_{*}\right)$.

Proof. For all $x \in \mathbb{R}^{d}$ we have

$$
\left\|x^{(0)}-x\right\|_{2}=\left\|\frac{1}{n} \sum_{i \in[n]} a^{(i)}-\frac{1}{n} \sum_{i \in[n]} x\right\|_{2} \leq \frac{1}{n} \sum_{i \in[n]}\left\|a^{(i)}-x\right\|_{2} \leq \frac{f(x)}{n} .
$$

Consequently,

$$
f\left(x^{(0)}\right) \leq \sum_{i \in[n]}\left\|x^{(0)}-a^{(i)}\right\|_{2} \leq \sum_{i \in[n]}\left(\left\|x^{(0)}-x_{*}\right\|_{2}+\left\|x_{*}-a^{(i)}\right\|_{2}\right) \leq 2 \cdot f\left(x_{*}\right)
$$

Lemma 7.5.7. For all $t \geq 0$, we have

$$
1 \leq \frac{t^{2} \cdot w_{t}(x)}{\mu_{t}(x)} \leq \bar{g}_{t}(x) \leq \max _{i \in[n]} g_{t}^{(i)}(x) \leq 1+t \cdot f(x)
$$

In particular, if $\left\|x-x_{t}\right\|_{2} \leq \frac{1}{t}+\alpha$ and $t \leq \frac{2 n}{\tilde{\varepsilon}_{*} \cdot \tilde{f}_{*}}$ then $g_{t}^{(i)}(x) \leq 6 n \varepsilon_{*}^{-1}+t \cdot \alpha$.
Proof of Lemma 7.5.7. The first claim $1 \leq \frac{t^{2} \cdot w_{t}(x)}{\mu_{t}(x)} \leq \bar{g}_{t}(x)$, follows from $\mu_{t}(x) \geq \sum_{i \in[n]} \frac{t^{2}}{g_{t}^{(i)}(x)\left(1+g_{t}^{(i)}(x)\right)}$ and the fact that the largest eigenvalue of $\nabla^{2} f_{t}(x)$ is at most $t^{2} \cdot w_{t}(x)$. The second follows from the fact that $\bar{g}_{t}(x)$ is a weighted harmonic mean of $g_{t}^{(i)}(x)$ and therefore

$$
\bar{g}_{t}(x) \leq \max _{i \in[n]} g_{t}^{(i)}(x) \leq 1+t \cdot \max _{i \in[n]}\left\|x-a^{(i)}\right\|_{2} \leq 1+t \cdot f(x) .
$$

The final inequality comes from the fact that $\left\|x-a^{(i)}\right\|_{2} \leq\left\|x-x_{t}\right\|_{2}+\left\|x_{t}-a^{(i)}\right\|_{2} \leq\left\|x-x_{t}\right\|_{2}+f\left(x_{t}\right)$ and the fact that $f\left(x_{t}\right) \leq f\left(x_{*}\right)+\frac{2 n}{t}$ by Lemma 7.3.6.

Lemma 7.5.8. For all $x \in \mathbb{R}^{d}$ and $t>0$ such that $\left\|x-x_{t}\right\|_{2} \leq \frac{1}{t}+\alpha$ and $t \leq \frac{2 n}{\tilde{\varepsilon}_{*} f_{*}}$ we have

$$
n\left(\frac{t}{7 n \tilde{\varepsilon}_{*}^{-1}+t \cdot \alpha}\right)^{2}\left\|x-x_{t}\right\|_{2}^{2} \leq f_{t}(x)-f_{t}\left(x_{t}\right) \leq \frac{n t^{2}}{2}\left\|x-x_{t}\right\|_{2}^{2}
$$

Proof of Lemma 7.5.8. For the first inequality, note that $\nabla^{2} f_{t}(x) \preceq \sum_{i \in[n]} \frac{t^{2}}{1+g_{t}^{(i)}(x)} \boldsymbol{I} \preceq n \cdot t^{2} \boldsymbol{I}$. Consequently, if we let $n \cdot t^{2} \boldsymbol{I}=\boldsymbol{H}$ in Lemma 7.9.5, we have that

$$
f_{t}(x)-f_{t}\left(x_{t}\right) \leq \frac{1}{2}\left\|x-x_{t}\right\|_{\boldsymbol{H}}^{2} \leq \frac{n t^{2}}{2}\left\|x-x_{t}\right\|_{2}^{2}
$$

For the second inequality, note that Lemma 7.5.2 and Lemma 7.5.7 yields that

$$
\nabla^{2} f_{t}(x) \succeq \sum_{i \in[n]} \frac{t^{2}}{\left(1+g_{t}^{(i)}(x)\right) g_{t}^{(i)}(x)} \boldsymbol{I} \succeq n\left(\frac{t}{7 n \tilde{\varepsilon}^{-1}+t \cdot \alpha}\right)^{2} \boldsymbol{I}
$$

Consequently, if we let $S$ denote the set of all $x \in \mathbb{R}^{n}$ with $\left\|x-x_{t}\right\|_{2} \leq \frac{1}{t}+\alpha$ then we see that $S$ is convex and contains the minimizer of $f_{t}$ and therefore by applying 7.9.5 again yields the lower bound.

## ■ 7.6 Analysis of the Algorithm

Here we provide proofs, algorithms, and technical lemmas from Section 7.4.

## ■ 7.6.1 Eigenvector Computation and Hessian Approximation

Below we prove that the power method can be used to compute an $\varepsilon$-approximate top eigenvector of a symmetric PSD matrix $\boldsymbol{A} \in \mathbb{R}^{d \times d}$ with a non-zero eigenvalue gap $g=\frac{\lambda_{1}(\boldsymbol{A})-\lambda_{2}(\boldsymbol{A})}{\lambda_{1}(\boldsymbol{A})}$. While it is well know that this can be by applying $\boldsymbol{A}$ to a random initial vector $O\left(\frac{\alpha}{g} \log \left(\frac{d}{\varepsilon}\right)\right)$ times in the following theorem we provide a slightly less known refinement that the dimension $d$ can be replaced with the stable rank of $\boldsymbol{A}, s=\sum_{i \in[d]} \frac{\lambda_{i}(\boldsymbol{A})}{\lambda_{1}(\boldsymbol{A})}$. We use this fact to avoid a dependence on $d$ in our logarithmic factors.

```
Algorithm 17: PowerMethod \((\boldsymbol{A}, k)\)
    Input: symmetric PSD matrix \(\boldsymbol{A} \in \mathbb{R}^{d \times d}\) and a number of iterations \(k \geq 1\).
    Let \(x \sim \mathcal{N}(0,1)\) be drawn from a \(d\) dimensional normal distribution.
    Let \(y=A^{k} z\)
    Output: \(u=y /\|y\|_{2}\)
```

Lemma 7.6.1 (Power Method). Let $\boldsymbol{A} \in \mathbb{R}^{d \times d}$ be a symmetric PSD matrix , let $g \stackrel{\text { def }}{=} \frac{\lambda_{1}(\boldsymbol{A})-\lambda_{2}(\boldsymbol{A})}{\lambda_{1}(\boldsymbol{A})}$, $s=\sum_{i \in d} \frac{\lambda_{i}}{\lambda_{1}}$, and let $\varepsilon>0$ and $k \geq \frac{\alpha}{g} \log \left(\frac{n}{\varepsilon}\right)$ for large enough constant $\alpha$. In time $O\left(\operatorname{nnz}(\boldsymbol{A}) \cdot \log \left(\frac{n}{\varepsilon}\right)\right)$, the algorithm PowerMethod $(\boldsymbol{A}, k)$ outputs a vector $u$ such that $\left\langle v_{1}(\boldsymbol{A}), u\right\rangle^{2} \geq 1-\varepsilon$ and $u^{\top} \boldsymbol{A} u \geq$ $(1-\varepsilon) \lambda_{1}(\boldsymbol{A})$ with high probability in $n / \varepsilon$.

Proof. We write $u=\sum_{i \in[d]} \alpha_{i} v_{i}(\boldsymbol{A})$. Then, we have

$$
\left\langle v_{1}(\boldsymbol{A}), u\right\rangle^{2}=\left\langle v_{1}(\boldsymbol{A}), \frac{\sum_{i \in[d]} \alpha_{i} \lambda_{i}(\boldsymbol{A})^{k} v_{i}(\boldsymbol{A})}{\sqrt{\sum_{i \in[d]} \alpha_{i}^{2} \lambda_{i}(\boldsymbol{A})^{2 k}}}\right\rangle^{2}=\frac{\alpha_{1}^{2}}{\alpha_{1}^{2}+\sum_{j \neq 1} \alpha_{j}^{2}\left(\frac{\lambda_{j}(\boldsymbol{A})}{\lambda_{1}(\boldsymbol{A})}\right)^{2 k}} \geq 1-\sum_{j \neq 1} \frac{\alpha_{j}^{2}}{\alpha_{1}^{2}}\left(\frac{\lambda_{j}(\boldsymbol{A})}{\lambda_{1}(\boldsymbol{A})}\right)^{2 k}
$$

Re arranging terms we have

$$
\begin{aligned}
1-\left\langle v_{1}(\boldsymbol{A}), u\right\rangle^{2} & \leq \sum_{j \neq 1} \frac{\alpha_{j}^{2}}{\alpha_{1}^{2}}\left(\frac{\lambda_{j}(\boldsymbol{A})}{\lambda_{1}(\boldsymbol{A})}\right)\left(\frac{\lambda_{j}(\boldsymbol{A})}{\lambda_{1}(\boldsymbol{A})}\right)^{2 k-1} \leq \sum_{j \neq 1} \frac{\alpha_{j}^{2}}{\alpha_{1}^{2}}\left(\frac{\lambda_{j}(\boldsymbol{A})}{\lambda_{1}(\boldsymbol{A})}\right)\left(\frac{\lambda_{2}(\boldsymbol{A})}{\lambda_{1}(\boldsymbol{A})}\right)^{2 k-1} \\
& =\sum_{j \neq 1,} \frac{\alpha_{j}^{2}}{\alpha_{1}^{2}} \cdot\left(\frac{\lambda_{j}(\boldsymbol{A})}{\lambda_{1}(\boldsymbol{A})}\right) \cdot(1-g)^{2 k-1} \leq \sum_{j \neq 1} \frac{\alpha_{j}^{2}}{\alpha_{1}^{2}} \cdot\left(\frac{\lambda_{j}(\boldsymbol{A})}{\lambda_{1}(\boldsymbol{A})}\right) \cdot \exp (-(2 k-1) g)
\end{aligned}
$$

where we used that $\frac{\lambda_{2}}{\lambda_{1}}=1-g \leq e^{-g}$.
Now with high probability in $n$ we have that $\alpha_{1}^{2} \geq \frac{1}{O(\operatorname{poly}(n / \varepsilon))}$ by known properties of the chisquared distribution. All that remains is to upper bound $\sum_{j \neq 1} \alpha_{j}^{2} \cdot\left(\frac{\lambda_{j}(\boldsymbol{A})}{\lambda_{1}(\boldsymbol{A})}\right)$. To bound this consider $h(\alpha) \stackrel{\text { def }}{=} \sqrt{\sum_{j \neq 1} \alpha_{j}^{2}\left(\lambda_{j}(\boldsymbol{A}) / \lambda_{1}(\boldsymbol{A})\right)}$. Note that
where $\overrightarrow{1}_{j}$ is the indicator vector for coordinate $j$. Consequently $h$ is 1 -Lipschitz and by Gaussian concentration for Lipschitz functions we know there are absolute constants $C$ and $c$ such that

$$
\operatorname{Pr}[h(\alpha) \geq \mathbf{E} h(\alpha)+\lambda] \leq C \exp \left(-c \lambda^{2}\right) .
$$

By the concavity of square root and the expected value of the chi-squared distribution we have

$$
\mathbf{E} h(\alpha) \leq \sqrt{\mathbf{E} \sum_{j \neq i} \alpha_{j}^{2} \cdot\left(\frac{\lambda_{j}(\boldsymbol{A})}{\lambda_{1}(\boldsymbol{A})}\right)}=\sqrt{\sum_{j \neq i}\left(\frac{\lambda_{j}(\boldsymbol{A})}{\lambda_{1}(\boldsymbol{A})}\right)} \leq \sqrt{s} .
$$

Consequently, since $s \geq 1$ we have that $\operatorname{Pr}\left[h(\alpha) \geq(1+\sqrt{t})^{2} \cdot s\right] \leq C \exp (-c \cdot t)$ for $t \geq 1$ and that $\sum_{j \neq 1} \alpha_{j} \cdot\left(\frac{\lambda_{j}(\boldsymbol{A})}{\lambda_{1}(\boldsymbol{A})}\right)=O(\operatorname{poly}(n / \varepsilon) \cdot s)$ with high probability in $n / \varepsilon$. Since $k=\Omega\left(\frac{1}{g} \log \left(\frac{n s}{\varepsilon}\right)\right)$, we have $\left\langle v_{1}(\boldsymbol{A}), u\right\rangle^{2} \geq 1-\varepsilon$ with high probability in $n$. Furthermore, this implies that

$$
u^{\top} \boldsymbol{A} u=u^{\top}\left(\sum_{i \in[d]} \lambda_{i}(\boldsymbol{A}) v_{i}(\boldsymbol{A}) v_{i}(\boldsymbol{A})^{\top}\right) u \geq \lambda_{1}(\boldsymbol{A})\left\langle v_{1}(\boldsymbol{A}), u\right\rangle^{2} \geq(1-\varepsilon) \lambda_{1}(\boldsymbol{A}) .
$$

Lemma 7.4.1 (Computing Hessian Approximation). Let $x \in \mathbb{R}^{d}, t>0$, and $\varepsilon \in\left(0, \frac{1}{4}\right)$. The algorithm ApproxMinEig $(x, t, \varepsilon)$ outputs $(\lambda, u)$ in $O\left(n d \log \frac{n}{\varepsilon}\right)$ time such that if $\mu_{t}(x) \leq \frac{1}{4} t^{2} w_{t}(x)$ then $\left\langle v_{t}(x), u\right\rangle^{2} \geq 1-\varepsilon$ with high probability in $n$. Furthermore, if $\varepsilon \leq \frac{\mu_{t}(x)}{8 t^{2} \cdot w_{t}(x)}$ then $\frac{1}{4} \boldsymbol{Q} \preceq \nabla^{2} f_{t}(x) \preceq 4 \boldsymbol{Q}$ with high probability in $n$ where $\boldsymbol{Q} \stackrel{\text { def }}{=} t^{2} \cdot w_{t}(x)-\left(t^{2} \cdot w_{t}(x)-\lambda\right) u u^{\top}$.

Proof of Lemma 7.4.1. By Lemma 7.5.4 we know that $\frac{1}{2} \boldsymbol{Z} \preceq \nabla^{2} f_{t}(x) \preceq \boldsymbol{Z}$ where

$$
\boldsymbol{Z}=t^{2} \cdot w_{t}(x)-\left(t^{2} \cdot w_{t}(x)-\mu_{t}(x)\right) v_{t}(x) v_{t}(x)^{\top} .
$$

Consequently, if $\mu_{t}(x) \leq \frac{1}{4} t^{2} w_{t}(x)$, then for all unit vectors $w \perp v_{t}(x)$, we have that

$$
w^{\top} \nabla^{2} f_{t}(x) w \geq \frac{1}{2} w^{\top} \boldsymbol{Z} w \geq \frac{1}{2} t^{2} w_{t}(x) .
$$

Since $\nabla^{2} f_{t}(x)=t^{2} \cdot w_{t}(x)-\boldsymbol{A}$, for $\boldsymbol{A}$ in the definition of ApproxMinEig (Algorithm 14) this implies that $v_{t}(x)^{\top} \boldsymbol{A} v_{t}(x) \geq \frac{3}{4} t^{2} \cdot w_{t}(x)$ and $w^{\top} \boldsymbol{A} w \leq \frac{1}{2} t^{2} w_{t}(x)$. Furthermore, we see that

$$
\sum_{i \in[d]} \lambda_{i}(\boldsymbol{A})=\operatorname{Tr}(\boldsymbol{A})=\sum_{i \in[n]} \frac{t^{4}\left\|x-a^{(i)}\right\|_{2}^{2}}{\left(1+g_{t}^{(i)}(x)\right)^{2} g_{t}^{(i)}(x)} \leq t^{2} \cdot w_{t}(x)
$$

Therefore, in this case, $\boldsymbol{A}$ has a constant multiplicative gap between its top two eigenvectors stable rank at most a constant (i.e. $g=O(1)$ and $s=O(1)$ in Theorem 7.6.1). Consequently, by Theorem 7.6.1 we have $\left\langle v_{t}(x), u\right\rangle^{2} \geq 1-\varepsilon$.

For the second claim, we note that

$$
t^{2} \cdot w_{t}(x)-\mu_{t}(x) \geq u^{\top} \boldsymbol{A} u \geq(1-\varepsilon) \lambda_{1}(\boldsymbol{A})=(1-\varepsilon)\left(t^{2} \cdot w_{t}(x)-\mu_{t}(x)\right)
$$

Therefore, since $\lambda=u^{\top} \nabla^{2} f_{t}(x) u=t^{2} \cdot w_{t}(x)-u^{\top} \boldsymbol{A} u$, we have

$$
\begin{equation*}
(1-\varepsilon) \mu_{t}(x)-\varepsilon \cdot t^{2} w_{t}(x) \leq \lambda \leq \mu_{t}(x) \tag{7.7}
\end{equation*}
$$

On the other hand, by Lemma 7.9.2, we have that

$$
\begin{equation*}
\sqrt{\varepsilon} \boldsymbol{I} \preceq v_{t}(x) v_{t}(x)^{\top}-u u^{\top} \preceq \sqrt{\varepsilon} \boldsymbol{I} . \tag{7.8}
\end{equation*}
$$

Combining (7.7) and (7.8), we have $\frac{1}{2} \boldsymbol{Z} \preceq \boldsymbol{Q} \preceq 2 \boldsymbol{Z}$ if $\varepsilon \leq \frac{\mu_{t}(x)}{8 t^{2} \cdot w_{t}(x)}$ and $\frac{1}{4} \boldsymbol{Q} \preceq \nabla^{2} f_{t}(x) \preceq 4 \boldsymbol{Q}$ follows. On the other hand, when $\mu_{t}(x)>\frac{1}{4} t^{2} w_{t}(x)$. t is the case that $\frac{1}{4} t^{2} \cdot w_{t}(x) \boldsymbol{I} \preceq \nabla^{2} f_{t}(x) \preceq t^{2} \cdot w_{t}(x) \boldsymbol{I}$ and $\frac{1}{4} t^{2} \cdot w_{t}(x) \boldsymbol{I} \preceq \boldsymbol{Q} \preceq t^{2} \cdot w_{t}(x) \boldsymbol{I}$ again yielding $\frac{1}{4} \boldsymbol{Q} \preceq \nabla^{2} f_{t}(x) \preceq 4 \boldsymbol{Q}$.

Lemma 7.4.2. Let $u=\operatorname{ApproxMinEig}\left(x, t, \varepsilon_{v}\right)$ for $\varepsilon_{v}<\frac{1}{8}$ and $x$ such that $\left\|x-x_{t}\right\|_{2} \leq \frac{\varepsilon_{c}}{t}$ for $\varepsilon_{c} \leq\left(\frac{\varepsilon_{v}}{36}\right)^{\frac{3}{2}}$. $\mu_{t} \leq \frac{1}{4} t^{2} \cdot w_{t}$. For all unit vectors $y \perp u$, we have $\left\langle y, v_{t}\right\rangle^{2} \leq 8 \varepsilon_{v}$.

Proof of Lemma 7.4.2. By Lemma 7.4.1 we know that $\left\langle v_{t}(x), u\right\rangle^{2} \geq 1-\varepsilon_{v}$. Since clearly $\left\|x-x_{t}\right\|_{2} \leq \frac{1}{20 t}$, by assumption, Lemma 7.3 .1 swhows

$$
\left(1-6 \varepsilon_{c}^{2 / 3}\right) \nabla^{2} f_{t}\left(x_{t}\right) \preceq \nabla^{2} f_{t}(x) \preceq\left(1+6 \varepsilon_{c}^{2 / 3}\right) \nabla^{2} f_{t}\left(x_{t}\right) .
$$

Furthermore, since $\mu_{t} \leq \frac{1}{4} t^{2} \cdot w_{t}$, as in Lemma 7.4.1 we know that the largest eigenvalue of $\boldsymbol{A}$ defined in ApproxMinEig $(x, t, \varepsilon)$ is at least $\frac{3}{4} t^{2} \cdot w_{t}$ while the second largest eigenvalue is at most $\frac{1}{2} t^{2} \cdot w_{t}$. Consequently, the eigenvalue gap, $g$, defined in Lemma 7.9.3 is at least $\frac{1}{3}$ and this lemma shows that $\left\langle v_{t}(x), v_{t}\right\rangle^{2} \geq 1-36 \varepsilon_{c}^{2 / 3} \geq 1-\varepsilon_{v}$. Consequently, by Lemma 7.9.4, we have that $\left\langle u, v_{t}\right\rangle^{2} \geq 1-4 \varepsilon_{v}$.

To prove the final claim, we write $u=\alpha v_{t}+\beta w$ for an unit vector $w \perp v_{t}$. Since $y \perp u$, we have that $0=\alpha\left\langle v_{t}, y\right\rangle+\beta\langle w, y\rangle$. Then, either $\left\langle v_{t}, y\right\rangle=0$ and the result follows or $\alpha^{2}\left\langle v_{t}, y\right\rangle^{2}=\beta^{2}\langle w, y\rangle^{2}$ and since $\alpha^{2}+\beta^{2}=1$, we have

$$
\left\langle v_{t}, y\right\rangle^{2} \leq \frac{\beta^{2}\langle w, y\rangle^{2}}{\alpha^{2}} \leq \frac{1-\alpha^{2}}{\alpha^{2}} \leq 2\left(1-\alpha^{2}\right) \leq 8 \varepsilon_{v}
$$

where in the last line we used that $\alpha^{2} \geq 1-4 \varepsilon_{v}>\frac{1}{2}$ since $\varepsilon_{v} \leq \frac{1}{8}$.

Lemma 7.4.3. Suppose $\mu_{t} \geq \frac{1}{4} t^{2} \cdot w_{t}$ and let $t^{\prime} \in\left[t,\left(1+\frac{1}{600}\right) t\right]$ then $\left\|x_{t^{\prime}}-x_{t}\right\|_{2} \leq \frac{1}{100 t}$.
Proof of Lemma 7.4.3. Note that $t^{\prime}=(1+\beta) t$ where $\beta \in\left[0, \frac{1}{600}\right]$. Since $\frac{1}{4} t^{2} \cdot w_{t} \boldsymbol{I} \preceq \mu_{t} \boldsymbol{I} \preceq \nabla^{2} f\left(x_{t}\right)$ applying Lemma 7.3.3 then yields that for all $s \in\left[t, t^{\prime}\right]$

$$
\nabla^{2} f\left(x_{s}\right) \succeq \nabla^{2} f\left(x_{t}\right)-15 \beta t^{2} w_{t} \boldsymbol{I} \succeq\left(\frac{1}{4}-15 \beta\right) t^{2} \cdot w_{t} \boldsymbol{I} \succeq \frac{t^{2} \cdot w_{t}}{6} \boldsymbol{I}
$$

Consequently, by Lemma 7.5.1, the fact that $t\left\|x_{t}-a^{(i)}\right\|_{2} \leq g_{t}^{(i)}$, and Lemma 7.3.2 we have

$$
\begin{aligned}
\left\|x_{t^{\prime}}-x_{t}\right\|_{2} & \leq \int_{t}^{t^{\prime}}\left\|\frac{d}{d s} x_{s}\right\|_{2} d_{s}=\int_{t}^{t^{\prime}}\left\|\left(\nabla^{2} f_{s}\left(x_{s}\right)\right)^{-1} \sum_{i \in[n]} \frac{s}{\left(1+g_{s}^{(i)}\right) g_{s}^{(i)}}\left(x_{s}-a^{(i)}\right)\right\|_{2} d_{s} \\
& \leq \int_{t}^{t^{\prime}} \frac{6}{t^{2} \cdot w_{t}} \sum_{i \in[n]} \frac{s\left\|x_{s}-a^{(i)}\right\|_{2}}{\left(1+g_{s}^{(i)}\right) g_{s}^{(i)}} d_{s} \leq \int_{t}^{t^{\prime}} \frac{6 w_{s}}{t^{2} \cdot w_{t}} d_{s} \leq \int_{t}^{t^{\prime}} \frac{6}{t^{2}} \cdot\left(\frac{s}{t}\right)^{2} d s \\
& =\frac{6}{3 t^{4}}\left[\left(t^{\prime}\right)^{2}-(t)^{3}\right]=\frac{6}{3 t}\left[(1+\beta)^{3}-1\right] \leq \frac{6 \beta}{t} \leq \frac{1}{100 t}
\end{aligned}
$$

where in the last line we used that for $\alpha \geq 0$ we have $1+3 \alpha \leq(1+\alpha)^{3}$.

## ■ 7.6.2 Line Searching

Here we prove the main results we use on centering, Lemma 7.4.4, and line searching Lemma 7.4.5. These results are our main tools for computing approximations to the central path. To prove Lemma 7.4.5 we also include here two preliminary lemmas, Lemma 7.6.2 and Lemma 7.6.3, on the structure of $g_{t, y, v}$ defined in (7.5).
Lemma 7.4.4. Given some $y \in \mathbb{R}^{d}, t>0$ and $\varepsilon \in\left(0, \min \left\{\frac{1}{4}, \frac{\mu_{t}(x)}{8 t^{2} \cdot w_{t}(x)}\right\}\right)$. In $O\left(n d \log \left(\frac{n}{\varepsilon}\right)\right)$ time LocalCenter $(y, t, \varepsilon)$ computes $x^{(k)}$ such that with high probability in $n / \varepsilon$.

$$
f_{t}\left(x^{(k)}\right)-\min _{\|x-y\|_{2} \leq \frac{1}{49 t}} f_{t}(x) \leq \varepsilon\left(f_{t}(y)-\min _{\|x-y\|_{2} \leq \frac{1}{49 t}} f_{t}(x)\right) .
$$

Proof of Lemma 7.4.4. By Lemma 7.4.1 we know that $\frac{1}{4} \boldsymbol{Q} \preceq \nabla^{2} f_{t}(y) \preceq 4 \boldsymbol{Q}$ with high probability in $n / \varepsilon$. Furthermore for $x$ such that $\|x-y\|_{2} \leq \frac{1}{50 t}$ Lemma 7.3.1 shows that $\frac{1}{2} \nabla^{2} f_{t}(x) \preceq \nabla^{2} f_{t}(y) \preceq$ $2 \nabla^{2} f_{t}(x)$. Combining these we have that $\frac{1}{8} \boldsymbol{Q} \preceq \nabla^{2} f_{t}(x) \preceq 8 \boldsymbol{Q}$ for all $x$ with $\|x-y\|_{2} \leq \frac{1}{50 t}$. Therefore, Lemma 7.9.5 shows that

$$
f_{t}\left(x^{(k)}\right)-\min _{\|x-y\|_{2} \leq \frac{1}{49 t}} f_{t}(x) \leq\left(1-\frac{1}{64}\right)^{k}\left(f_{t}\left(x^{(0)}\right)-\min _{\|x-y\|_{2} \leq \frac{1}{49 t}} f_{t}(x)\right) .
$$

The guarantee on $x^{(k)}$ then follows from our choice of $k$.
For the running time, Lemma 7.4 .1 showed the cost of ApproxMinEig is $O\left(n d \log \left(\frac{n}{\varepsilon}\right)\right)$. Using Lemma 7.9.7 we see that the cost per iteration is $O(n d)$ and therefore, the total cost of the $k$ iterations is $O\left(n d \log \left(\frac{1}{\varepsilon}\right)\right)$. Combining yields the running time.

Lemma 7.6.2. For $t>0, y \in \mathbb{R}^{d}$, and unit vector $v \in \mathbb{R}^{d}$, the function $g_{t, y, v}: \mathbb{R} \rightarrow \mathbb{R}$ defined by (7.5) is convex and $n t$-Lipschitz.

Proof. Changing variables yields $g_{t, y, v}(\alpha)=\min _{z \in S} f_{t}(z+\alpha v)$ for $S \stackrel{\text { def }}{=}\left\{z \in \mathbb{R}^{d}:\|z-y\|_{2} \leq \frac{1}{49 t}\right\}$. Since $f_{t}$ is convex and $S$ is a convex set, by Lemma 7.9 .6 we have that $g_{t, y, v}$ is convex.

Next, by Lemma 7.5.1, triangle inequality, and the fact that $t\left\|x-a^{(i)}\right\|_{2} \leq g_{t}^{(i)}(x)$ we have

$$
\begin{equation*}
\left\|\nabla f_{t}(x)\right\|_{2}=\left\|\sum_{i \in[n]} \frac{t^{2}\left(x-a^{(i)}\right)}{1+g_{t}^{(i)}(x)}\right\|_{2} \leq \sum_{i \in[n]} \frac{t^{2}\left\|x-a^{(i)}\right\|_{2}}{1+g_{t}^{(i)}(x)} \leq t n . \tag{7.9}
\end{equation*}
$$

Consequently, $f_{t}(x)$ is $n t$-Lipschitz and for all $x, y \in \mathbb{R}^{n}$ we have $\left|f_{t}(x)-f_{t}(y)\right| \leq n t\|x-y\|_{2}$. Now if we consider the set $S_{\alpha} \stackrel{\text { def }}{=}\left\{x \in \mathbb{R}^{d}:\|x-(y+\alpha v)\|_{2} \leq \frac{1}{49 t}\right\}$ then we see that for all $\alpha, \beta \in \mathbb{R}$ there is a bijection from $S_{\alpha}$ to $S_{\beta}$ were every point in the set moves by at most $\|(\alpha-\beta) v\|_{2} \leq|\alpha-\beta|$. Consequently, since $g_{t, y, v}(\alpha)$ simply minimizes $f_{t}$ over $S_{\alpha}$ we have that $\left|g_{t, y, v}(\alpha)-g_{t, y, v}(\beta)\right| \leq n t|\alpha-\beta|$ for all $\alpha, \beta \in \mathbb{R}$ we have that $g_{t, y v}$ is $n t$-Lipschitz as desired.

Lemma 7.6.3. Let $\frac{1}{400 f\left(x_{*}\right)} \leq t \leq t^{\prime} \leq\left(1+\frac{1}{600}\right) t \leq \frac{2 n}{\tilde{\varepsilon} \cdot \tilde{f}_{*}}$ and let $u=\operatorname{ApproxMinEig}\left(y, t, \varepsilon_{v}\right)$ for $\varepsilon_{v} \leq \frac{1}{8}\left(\frac{\tilde{\varepsilon}_{*}}{7 n}\right)^{2}$ and $y \in \mathbb{R}^{d}$ such that $\left\|y-x_{t}\right\|_{2} \leq \frac{1}{t}\left(\frac{\varepsilon_{v}}{36}\right)^{\frac{3}{2}}$. The function $g_{t, \prime}, y, v: \mathbb{R} \rightarrow \mathbb{R}$ defined in (7.5) satisfies $g_{t, y, v}\left(\alpha_{*}\right)=\min _{\alpha} g_{t, y, v}(\alpha)=f_{t^{\prime}}\left(x_{t^{\prime}}\right)$ for some $\alpha_{*} \in\left[-6 f\left(x_{*}\right), 6 f\left(x_{*}\right)\right]$.

Proof of Lemma 7.4.5. Let $z \in \mathbb{R}^{d}$ be an arbitrary unit vector. By Lemma 7.3.5 we know if $\left|\left\langle z, v_{t}\right\rangle\right| \leq$ $\min _{\delta \in[t,(1+\beta) t]} \frac{\mu_{\delta}}{t^{2} \cdot w_{\delta}}$ for $\beta \in\left[0, \frac{1}{600}\right]$ then $z^{\top}\left(x_{(1+\beta) t}-x_{t}\right) \leq \frac{6 \beta}{t}$. Now by Lemma 7.5.7 and our bound on $\beta$ we know that $\frac{\delta^{2} \cdot w_{\delta}}{\mu_{\delta}} \leq 7 n \tilde{\varepsilon}_{*}^{-1}$ and consequently $\frac{\tilde{\varepsilon}_{*}}{6 n} \leq \min _{\delta \in[t,(1+\beta) t]} \frac{w_{\delta}}{t^{2} \cdot \mu_{\delta}}$. Consequently, if $\mu_{t} \leq$
$\frac{1}{4} t^{2} \cdot w_{t}$ then by Lemma 7.4.2 and our choice of $\varepsilon_{v}$ we have that if $z \perp u$ then $\left|\left\langle z, v_{t}\right\rangle\right|^{2} \leq\left(\frac{\tilde{\varepsilon}_{*}}{6 n}\right)^{2}$ and $z^{\top}\left(x_{t^{\prime}}-x_{t}\right) \leq \frac{6}{600 t}=\frac{1}{100 t}$. Otherwise, $\mu_{t} \geq \frac{1}{4} t^{2} \cdot w_{t}$ and by Lemma 7.4.3 we have $\left\|x_{t^{\prime}}-x_{t}\right\|_{2} \leq \frac{1}{100 t}$. In either case, since $\left\|y-x_{t}\right\|_{2} \leq \frac{1}{100 t}$, we can reach $x_{t^{\prime}}$ from $y$ by first moving an Euclidean distance of $\frac{1}{100 t}$ to go from $x$ to $x_{t}$, then adding some multiple of $v$, then moving an Euclidean distance of $\frac{1}{100 t}$ in a direction perpendicular to $v$. Since the total movement perpendicular to $v$ is $\frac{1}{100 t}+\frac{1}{100 t} \leq \frac{1}{49 t^{\prime}}$ we have that $\min _{\alpha} g_{t^{\prime}, y, v}(\alpha)=f_{t^{\prime}}\left(x_{t^{\prime}}\right)$ as desired.

All that remains is to show that there is a minimizer of $g_{t, y, v}$ in the range $\left[-6 f\left(x_{*}\right), 6 f\left(x_{*}\right)\right]$. However, by Lemma 7.3.6 and Lemma 7.5.5 we know that

$$
\left\|y-x_{t^{\prime}}\right\|_{2} \leq\left\|y-x_{t}\right\|_{2}+\left\|x_{t}-x_{*}\right\|_{2}+\left\|x_{*}-x_{t^{\prime}}\right\|_{2} \leq \frac{1}{100 t}+f\left(x_{*}\right)+f\left(x_{*}\right) \leq 6 f\left(x_{*}\right) .
$$

Consequently, $\alpha_{*} \in\left[-6 f\left(x_{*}\right), 6 f\left(x_{*}\right)\right]$ as desired.
Lemma 7.4.5. Let $\frac{1}{400 f\left(x_{*}\right)} \leq t \leq t^{\prime} \leq\left(1+\frac{1}{600}\right) t \leq \frac{2 n}{\tilde{\varepsilon} \cdot f_{*}}$ and let $u=\operatorname{ApproxMinEig}\left(y, t, \varepsilon_{v}\right)$ for $\varepsilon_{v} \leq$ $\frac{1}{8}\left(\frac{\tilde{z}_{*}}{7 n}\right)^{2}$ and $y \in \mathbb{R}^{d}$ such that $\left\|y-x_{t}\right\|_{2} \leq \frac{1}{t}\left(\frac{\varepsilon_{v}}{36}\right)^{\frac{3}{2}}$. In $O\left(n d \log ^{2}\left(\frac{n}{\tilde{\varepsilon} \cdot \varepsilon \cdot \varepsilon \cdot \varepsilon_{v}}\right)\right)$ time, LineSearch $\left(y, t, t^{\prime}, u, \varepsilon\right)$ outputs $x^{\prime}$ such that $\left\|x^{\prime}-x_{t^{\prime}}\right\|_{2} \leq \frac{\varepsilon}{t^{\prime}}$ with high probability in $n$.

By (7.9) we know that $f_{t^{\prime}}$ is $n t^{\prime}$ Lipschitz and therefore

$$
f_{t}(y)-\min _{\|x-y\|_{2} \leq \frac{1}{49 t^{\prime}}} f_{t^{\prime}}(x) \leq \frac{n t^{\prime}}{49 t^{\prime}}=\frac{n}{49} .
$$

Furthermore, for $\alpha \in\left[-6 \tilde{f}_{*}, 6 \tilde{f}_{*}\right]$ we know that by Lemma 7.5.5

$$
\left\|y+\alpha u-x_{t^{\prime}}\right\|_{2} \leq\left\|y-x_{t}\right\|_{2}+|\alpha|+\left\|x_{t}-x_{*}\right\|_{2}+\left\|x_{t^{\prime}}-x_{*}\right\|_{2} \leq \frac{1}{t^{\prime}}+14 f\left(x_{*}\right)
$$

consequently by Lemma 7.5 .7 we have

$$
\frac{\left(t^{\prime}\right)^{2} \cdot w_{t^{\prime}}(y+\alpha u)}{\mu_{t^{\prime}}(y+\alpha u)} \leq 6 n \tilde{\varepsilon}^{-1}+t^{\prime} \cdot 14 f\left(x_{*}\right) \leq 34 n \tilde{\varepsilon}^{-1}
$$

Since $\varepsilon_{O} \leq \frac{1}{8} \cdot \frac{\tilde{\varepsilon}_{*}}{34 n}$ invoking Lemma 7.4.4 yields that $\left\lvert\, q(\alpha)-g_{t^{\prime}, y, u}(\alpha) \leq \frac{n \varepsilon_{O}}{49}\right.$ with high probability in $n / \tilde{\varepsilon}_{*}$ by and each call to LocalCenter takes $O\left(n d \log \frac{n}{\varepsilon_{O}}\right)$ time. Furthermore, by Lemma 7.6.2 we have that $g_{t^{\prime}, y, u}$ is a $n t^{\prime}$-Lipschitz convex function and by Lemma 7.6 .3 we have that the minimizer has value $f_{t^{\prime}}\left(x_{t^{\prime}}\right)$ and is achieved in the range $\left[-12 \tilde{f}_{*}, 12 \tilde{f}_{*}\right]$. Consequently, combining all these facts and invoking Lemma 7.9.3, i.e. our result on on one dimensional function minimization, we have $f_{t^{\prime}}\left(x^{\prime}\right)-f_{t^{\prime}}\left(x_{t^{\prime}}\right) \leq \frac{\varepsilon_{O}}{49 t^{\prime}}$ using only $O\left(\log \frac{n t^{\prime} f\left(x_{*}\right)}{\varepsilon_{O}}\right)$ calls to LocalCenter.

Finally, by Lemma 7.5.8 and Lemma 7.3.6 we have

$$
\left\|x-x_{t^{\prime}}\right\| \leq \frac{1}{\sqrt{n}} \cdot\left(\frac{7 n \tilde{\varepsilon}^{-1}+t^{\prime} \alpha}{t^{\prime}}\right) \cdot \sqrt{\frac{n \varepsilon_{O}}{49}} \leq \frac{34 n}{7 \cdot \tilde{\varepsilon}_{*} \cdot t^{\prime}} \sqrt{\varepsilon_{O}} \leq \frac{6 n}{\tilde{\varepsilon}_{*} \cdot t^{\prime}} \sqrt{\varepsilon_{O}} .
$$

Since $\varepsilon_{O} \leq \frac{\varepsilon^{2} \cdot \dot{\varepsilon}_{x}^{2}}{36 n^{2}}$ we have that $\left\|x-x_{t^{\prime}}\right\|_{2} \leq \frac{\varepsilon}{t^{\prime}}$ as desired.
Lemma 7.4.6. Let $\frac{1}{400 f\left(x_{*}\right)} \leq t \leq t^{\prime} \leq\left(1+\frac{1}{600}\right) t \leq \frac{2 n}{\tilde{\varepsilon}_{*} \cdot \tilde{f}_{*}}$ and let $x \in \mathbb{R}^{d}$ satisfy $\left\|x-x_{t}\right\|_{2} \leq \frac{1}{100 t}$. Then, in $O\left(n d \log ^{2}\left(\frac{n}{\varepsilon \cdot \tilde{\varepsilon}_{*}}\right)\right)$ time, LineSearch $(x, t, t, u, \varepsilon)$ output $y$ such that $\left\|y-x_{t}\right\|_{2} \leq \frac{\varepsilon}{t}$ for any vector $u \in \mathbb{R}^{d}$.

Proof of Lemma 7.4.6. The proof is strictly easier than the proof of Lemma 7.4.5 as $\left\|x-\alpha^{*} u-x_{t}\right\|_{2} \leq$ $\frac{1}{100 t}$ is satisfied automatically for $\alpha^{*}=0$. Note that this lemma assume less for the initial point.

## ■ 7.6.3 Putting It All Together

Theorem 7.4.1. In $O\left(n d \log ^{3}\left(\frac{n}{\varepsilon}\right)\right)$ time, Algorithm 13 outputs an $(1+\varepsilon)$-approximate geometric median with constant probability.

Proof of Theorem 7.4.1. By Lemma 7.5.6 we know that $x^{(0)}$ is a 2 -approximate geometric median and therefore $f\left(x^{(0)}\right)=\tilde{f}_{*} \leq 2 \cdot f\left(x_{*}\right)$. Furthermore, since $\left\|x^{(0)}-x_{t_{1}}\right\|_{2} \leq f\left(x^{(0)}\right)$ by Lemma 7.5.5 and $t_{1}=\frac{1}{400 \tilde{f}_{*}}$ we have $\left\|x^{(0)}-x_{t_{1}}\right\|_{2} \leq \frac{1}{400 t_{1}}$. Hence, by Lemma 7.4.6, we have $\left\|x^{(1)}-x_{t_{1}}\right\|_{2} \leq \frac{\varepsilon_{c}}{t_{1}}$ with high probability in $n / \varepsilon$. Consequently, by Lemma 7.4 .5 we have that $\left\|x^{(k)}-x_{t_{i}}\right\|_{2} \leq \frac{\varepsilon_{c}}{t_{i}}$ for all $i$ with high probability in $n / \varepsilon$.

Now, Lemma 7.3.6 shows that

$$
f\left(x_{t_{k}}\right)-f\left(x^{*}\right) \leq \frac{2 n}{t_{k}} \leq \frac{2 n}{\tilde{t}_{*}}\left(1+\frac{1}{600}\right) \leq \tilde{\varepsilon}_{*} \cdot \tilde{f}_{*}\left(1+\frac{1}{600}\right) \leq \frac{2}{3}\left(1+\frac{1}{600}\right) \varepsilon \cdot f\left(x_{*}\right) .
$$

Since $\left\|x^{(k)}-x_{t_{k}}\right\|_{2} \leq \frac{\varepsilon_{c}}{t_{k}} \leq 400 \cdot \tilde{f}_{*} \cdot \varepsilon_{c}$ we have that $f\left(x_{k}\right) \leq f\left(x_{t_{k}}\right)+400 n \cdot \tilde{f}_{*} \cdot \varepsilon_{c}$ by triangle inequality. Combining these facts and using that $\varepsilon_{c}$ is sufficiently small yields that $f\left(x^{(k)}\right) \leq(1+\varepsilon) f\left(x_{*}\right)$ as desired.

To bound the running time, Lemma 7.4.1 shows ApproxMinEvec takes $O\left(n d \log \left(\frac{n}{\varepsilon}\right)\right)$ per iteration and Lemma 7.4 .5 shows LineSearch takes $O\left(n d \log ^{2}\left(\frac{n}{\varepsilon}\right)\right)$ time per iteration, using that $\varepsilon_{v}$ and $\varepsilon_{c}$ are $O(\Omega(\varepsilon / n))$. Since for $l=\Omega\left(\log \frac{n}{\varepsilon}\right)$ we have that $t_{l}>\tilde{t}_{*}$ we have that $k=O\left(\log \frac{n}{\varepsilon}\right) \cdot t_{i+1} \leq \frac{1}{400}$. Since there are $O\left(\log \left(\frac{n}{\varepsilon}\right)\right)$ iterations taking time $O\left(n d \log ^{2}\left(\frac{n}{\varepsilon}\right)\right)$ the running time follows.

## ■ 7.7 Pseudo Polynomial Time Algorithm

Here we provide a self-contained result on computing a $1+\varepsilon$ approximate geometric median in $O\left(d \varepsilon^{-2}\right)$ time. Note that it is impossible to achieve such approximation for the mean, $\min _{x \in \mathbb{R}^{d}} \sum_{i \in[n]}\left\|x-a^{(i)}\right\|_{2}^{2}$, because the mean can be changed arbitrarily by changing only 1 point. However, [172] showed that the geometric median is far more stable. In Section 7.7.1, we show how this stability property allows us to get an constant approximate in $O(d)$ time. In Section 7.7.2, we show how to use stochastic subgradient descent to then improve the accuracy.

## ■ 7.7.1 A Constant Approximation of Geometric Median

We first prove that the geometric median is stable even if we are allowed to modify up to half of the points. The following lemma is a strengthening of the robustness result in [172].
Lemma 7.7.1. Let $x_{*}$ be a geometric median of $\left\{a^{(i)}\right\}_{i \in[n]}$ and let $S \subseteq[n]$ with $\|S\|<\frac{n}{2}$. For all $x$

$$
\left\|x_{*}-x\right\|_{2} \leq\left(\frac{2 n-2|S|}{n-2|S|}\right) \max _{i \notin S}\left\|a^{(i)}-x\right\|_{2} .
$$

Proof. For notational convenience let $r=\left\|x_{*}-x\right\|_{2}$ and let $M=\max _{i \notin S}\left\|a^{(i)}-x\right\|_{2}$.
For all $i \notin S$, we have that $\left\|x-a^{(i)}\right\|_{2} \leq M$, hence, we have

$$
\begin{aligned}
\left\|x_{*}-a^{(i)}\right\|_{2} & \geq r-\left\|x-a^{(i)}\right\|_{2} \\
& \geq r-2 M+\left\|x-a^{(i)}\right\|_{2} .
\end{aligned}
$$

Furthermore, by triangle inequality for all $i \in S$, we have

$$
\left\|x_{*}-a^{(i)}\right\|_{2} \geq\left\|x-a^{(i)}\right\|_{2}-r .
$$

Hence, we have that

$$
\sum_{i \in[n]}\left\|x_{*}-a^{(i)}\right\|_{2} \geq \sum_{i \in[n]}\left\|x-a^{(i)}\right\|_{2}+(n-|S|)(r-2 M)-|S| r .
$$

Since $x^{*}$ is a minimizer of $\sum_{i \in[n]}\left\|x_{*}-a^{(i)}\right\|_{2}$, we have that

$$
(n-|S|)(r-2 M)-|S| r \leq 0
$$

Hence, we have

$$
\left\|x^{*}-x\right\|_{2}=r \leq \frac{2 n-2|S|}{n-2|S|} M .
$$

Now, we use Lemma 7.7.1 to show that the algorithm CrudeApproximate outputs a constant approximation of the geometric median with high probability.

## Algorithm 18: CrudeApproximate ${ }_{K}$

Input: $a^{(1)}, a^{(2)}, \cdots, a^{(n)} \in \mathbb{R}^{d}$.
Sample two independent random subset of $[n]$ of size $K$. Call them $S_{1}$ and $S_{2}$.
Let $i^{*} \in \arg \min _{i \in S_{2}} \alpha_{i}$ where $\alpha_{i}$ is the 65 percentile of the numbers $\left\{\left\|a^{(i)}-a^{(j)}\right\|_{2}\right\}_{j \in S_{1}}$.
Output: Output $a^{\left(i^{*}\right)}$ and $\alpha_{i^{*}}$.

Lemma 7.7.2. Let $x_{*}$ be a geometric median of $\left\{a^{(i)}\right\}_{i \in[n]}$ and $(\widetilde{x}, \lambda)$ be the output of CrudeApproximate ${ }_{K}$. We define $d_{T}^{k}(x)$ be the $k$-percentile of $\left\{\left\|x-a^{(i)}\right\|\right\}_{i \in T}$. Then, we have that $\left\|x_{*}-\widetilde{x}\right\|_{2} \leq 6 d_{[n]}^{60}(\widetilde{x})$. Furthermore, with probability $1-e^{-\Theta(K)}$, we have

$$
d_{[n]}^{60}(\widetilde{x}) \leq \lambda=d_{S_{1}}^{60}(\widetilde{x}) \leq 2 d_{[n]}^{70}\left(x_{*}\right)
$$

Proof. Lemma 7.7.1 shows that for all $x$ and $T \subseteq[n]$ with $|T| \leq \frac{n}{2}$

$$
\left\|x_{*}-x\right\|_{2} \leq\left(\frac{2 n-2|T|}{n-2|T|}\right) \max _{i \notin T}\left\|a^{(i)}-x\right\|_{2} .
$$

Picking $T$ to be the indices of largest $40 \%$ of $\left\|a^{(i)}-\widetilde{x}\right\|_{2}$, we have

$$
\begin{equation*}
\left\|x_{*}-\widetilde{x}\right\|_{2} \leq\left(\frac{2 n-0.8 n}{n-0.8 n}\right) d_{[n]}^{60}(\widetilde{x})=6 d_{[n]}^{60}(\widetilde{x}) \tag{7.10}
\end{equation*}
$$

For any point $x$, we have that $d_{[n]}^{60}(x) \leq d_{S_{1}}^{65}(x)$ with probability $1-e^{-\Theta(K)}$ because $S_{1}$ is a random subset of $[n]$ with size $K$. Taking union bound over elements on $S_{2}$, with probability $1-K e^{-\Theta(K)}=$ $1-e^{-\Theta(K)}$, for all points $x \in S_{2}$

$$
\begin{equation*}
d_{[n]}^{60}(x) \leq d_{S_{1}}^{65}(x) \tag{7.11}
\end{equation*}
$$

yielding that $d_{[n]}^{60}(\widetilde{x}) \leq \lambda$.
Next, for any $i \in S_{2}$, we have

$$
\left\|a^{(i)}-a^{(j)}\right\|_{2} \leq\left\|a^{(i)}-x_{*}\right\|_{2}+\left\|x_{*}-a^{(j)}\right\|_{2} .
$$

and hence

$$
d_{[n]}^{70}\left(a^{(i)}\right) \leq\left\|a^{(i)}-x_{*}\right\|_{2}+d_{[n]}^{70}\left(x_{*}\right)
$$

Again, since $S_{1}$ is a random subset of $[n]$ with size $K$, we have that $d_{S_{1}}^{65}\left(a^{(i)}\right) \leq d_{[n]}^{70}\left(a^{(i)}\right)$ with probability $1-K e^{-\Theta(K)}=1-e^{-\Theta(K)}$. Therefore,

$$
d_{S_{1}}^{65}\left(a^{(i)}\right) \leq\left\|a^{(i)}-x_{*}\right\|_{2}+d_{[n]}^{70}\left(x_{*}\right) .
$$

Since $S_{2}$ is an independent random subset, with probability $1-e^{-\Theta(K)}$, there is $i \in S_{2}$ such that $\left\|a^{(i)}-x_{*}\right\|_{2} \leq d_{[n]}^{70}\left(x_{*}\right)$. In this case, we have

$$
d_{S_{1}}^{65}\left(a^{(i)}\right) \leq 2 d_{[n]}^{70}\left(x_{*}\right) .
$$

Since $i^{*}$ minimize $d_{S_{1}}^{65}\left(a^{(i)}\right)$ over all $i \in S_{2}$, we have that

$$
\lambda \stackrel{\text { def }}{=} d_{S_{1}}^{65}(\widetilde{x}) \stackrel{\text { def }}{=} d_{S_{1}}^{65}\left(a^{\left(i^{*}\right)}\right) \leq d_{S_{1}}^{65}\left(a^{(i)}\right) \leq 2 d_{[n]}^{70}\left(x_{*}\right) .
$$

## ■ 7.7.2 A $1+\varepsilon$ Approximation of Geometric Median

Here we show how to improve the constant approximation in the previous section to a $1+\varepsilon$ approximation. Our algorithm is essentially stochastic subgradient where we use the information from the previous section to bound the domain in which we need to search for a geometric median.

```
Algorithm 19: ApproximateMedian( \(\varepsilon\) )
    Input: \(a^{(1)}, a^{(2)}, \cdots, a^{(n)} \in \mathbb{R}^{d}\).
    Let \(T=(120 / \varepsilon)^{2}\) and let \(\eta=\frac{6 \lambda}{n} \sqrt{\frac{2}{T}}\).
    Let \(\left(x^{(1)}, \lambda\right)=\) CrudeApproximate \(_{\sqrt{T}}\left(a^{(1)}, a^{(2)}, \cdots, a^{(n)}\right)\).
    for \(k \leftarrow 1,2, \cdots, T\) do
        Sample \(i_{k}\) from [ \(n\) ] and let
            \(g^{(k)}= \begin{cases}n\left(x^{(k)}-a^{\left(i_{k}\right)}\right) /\left\|x^{(k)}-a^{\left(i_{k}\right)}\right\|_{2} & \text { if } x^{(i)} \neq a^{\left(i_{k}\right)} \\ 0 & \text { otherwise }\end{cases}\)
        Let \(x^{(k+1)}=\arg \min _{\left\|x-x^{(1)}\right\|_{2} \leq 6 \lambda} \eta\left\langle g^{(k)}, x-x^{(k)}\right\rangle+\frac{1}{2}\left\|x-x^{(k)}\right\|_{2}^{2}\).
    end
    Output: Output \(\frac{1}{T} \sum_{i=1}^{T} x^{(k)}\).
```

Theorem 7.7.1. Let $x$ be the output of ApproximateMedian( $\varepsilon$ ). With probability $1-e^{-\Theta(1 / \varepsilon)}$, we have

$$
\mathbb{E} f(x) \leq(1+\varepsilon) \min _{x \in \mathbb{R}^{d}} f(x) .
$$

Furthermore, the algorithm takes $O\left(d / \varepsilon^{2}\right)$ time.
Proof. After computing $x^{(1)}$ and $\lambda$ the remainder of our algorithm is the stocastic subgradient descent method applied to $f(x)$. It is routine to check that $\mathbf{E}_{i^{(k)}} g^{(k)}$ is a subgradient of $f$ at $x^{(k)}$. Furthermore, since the diameter of the domain, $\left\{x:\left\|x-x^{(1)}\right\|_{2} \leq 6 \lambda\right\}$, is clearly $\lambda$ and the norm of sampled gradient, $g^{(k)}$, is at most $n$, we have that

$$
\mathbb{E} f\left(\frac{1}{T} \sum_{i=1}^{T} x^{(k)}\right)-\min _{\left\|x-x^{(1)}\right\|_{2} \leq 6 \lambda} f(x) \leq 6 n \lambda \sqrt{\frac{2}{T}}
$$

(see [41, Thm 6.1]). Lemma 7.7.2 shows that $\left\|x^{*}-x^{(1)}\right\|_{2} \leq 6 \lambda$ and $\lambda \leq 2 d d_{[n]}^{70}\left(x^{*}\right)$ with probability $1-\sqrt{T} e^{-\Theta(\sqrt{T})}$. In this case, we have

$$
\mathbb{E} f\left(\frac{1}{T} \sum_{i=1}^{T} x^{(k)}\right)-f\left(x^{*}\right) \leq \frac{12 \sqrt{2} n d_{[n]}^{70}\left(x_{*}\right)}{\sqrt{T}}
$$

Since $d_{[n]}^{70}\left(x^{*}\right) \leq \frac{1}{0.3 n} f\left(x^{*}\right)$, we have

$$
\mathbb{E} f\left(\frac{1}{T} \sum_{i=1}^{T} x^{(k)}\right) \leq\left(1+\frac{60}{\sqrt{T}}\right) f\left(x_{*}\right) \leq\left(1+\frac{\varepsilon}{2}\right) f\left(x_{*}\right)
$$

## ■ 7.8 Derivation of Penalty Function

Here we derive our penalized objective function. Consider the following optimization problem:

$$
\min _{x \in \mathbb{R}^{d}, \alpha \geq 0 \in \mathbb{R}^{n}} f_{t}(x, \alpha) \quad \text { where } \quad t \cdot 1^{T} \alpha+\sum_{i \in[n]}-\ln \left(\alpha_{i}^{2}-\left\|x-a^{(i)}\right\|_{2}^{2}\right) .
$$

Since $p_{i}(\alpha, x) \stackrel{\text { def }}{=}-\ln \left(\alpha_{i}^{2}-\left\|x-a^{(i)}\right\|_{2}^{2}\right)$ is a barrier function for the set $\alpha_{i}^{2} \geq\left\|x-a^{(i)}\right\|_{2}^{2}$, i.e. as $\alpha_{i} \rightarrow\left\|x-a^{(i)}\right\|_{2}$ we have $p_{i}(\alpha, x) \rightarrow \infty$, we see that as we minimize $f_{t}(x, \alpha)$ for increasing values of $t$ the $x$ values converge to a solution to the geometric median problem. Our penalized objective function, $f_{t}(x)$, is obtain simply by minizing the $\alpha_{i}$ in the above formula and dropping terms that do not affect the minimizing $x$. In the remainder of this section we show this formally.

Fix some $x \in \mathbb{R}^{d}$ and $t>0$. Note that for all $j \in[n]$ we have

$$
\frac{\partial}{\partial \alpha_{j}} f_{t}(x, \alpha)=t-\left(\frac{1}{\alpha_{j}^{2}-\left\|x-a^{(i)}\right\|_{2}^{2}}\right) 2 \alpha_{j} .
$$

Since $f(x, \alpha)$ is convex in $\alpha$, the minimum $\alpha_{j}^{*}$ must satisfy

$$
\begin{equation*}
t\left(\left(\alpha_{j}^{*}\right)^{2}-\left\|x-a^{(i)}\right\|_{2}^{2}\right)-2 \alpha_{j}^{*}=0 \tag{7.12}
\end{equation*}
$$

Solving for such $\alpha_{j}^{*}$ under the restriction $\alpha_{j}^{*} \geq 0$ we obtain

$$
\begin{equation*}
\alpha_{j}^{*}=\frac{2+\sqrt{4+4 t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}}{2 t}=\frac{1}{t}\left[1+\sqrt{1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}\right] \tag{7.13}
\end{equation*}
$$

Using (7.12) and (7.13) we have that

$$
\min _{\alpha \geq 0 \in \mathbb{R}^{n}} f_{t}(x, \alpha)=\sum_{i \in[n]}\left[1+\sqrt{1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}-\ln \left[\frac{2}{t^{2}}\left(1+\sqrt{1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}\right)\right]\right] .
$$

If we drop the terms that do not affect the minimizing $x$ we obtain our penalty function $f_{t}$ :

$$
f_{t}(x)=\sum_{i \in[n]}\left[\sqrt{1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}-\ln \left(1+\sqrt{1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}\right)\right]
$$

## ■ 7.9 Technical Lemmas

Here we provide various technical lemmas we use through this chapter.

## ■ 7.9.1 Linear Algebra

First we provide the following lemma that shows that any matrix obtained as a non-negative linear combination of the identity minus a rank 1 matrix less than the identity results in a matrix that is well approximated spectrally by the identity minus a rank 1 matrix. We use this lemma to characterize the Hessian of our penalized objective function and thereby imply that it is possible to apply the inverse of the Hessian to a vector with high precision.
Lemma 7.9.1. Let $\boldsymbol{A}=\sum_{i}\left(\alpha_{i} \boldsymbol{I}-\beta_{i} a_{i} a_{i}^{\top}\right) \in \mathbb{R}^{d \times d}$ where the $a_{i}$ are unit vectors and $0 \leq \beta_{i} \leq \alpha_{i}$ for all $i$. Let $v$ denote a unit ector that is the maximum eigenvector of $\sum_{i} \beta_{i} a_{i} a_{i}^{\top}$ and let $\lambda$ denote the corresponding eigenvalue. Then,

$$
\frac{1}{2}\left(\sum_{i} \alpha_{i} \boldsymbol{I}-\lambda v v^{\top}\right) \preceq \boldsymbol{A} \preceq \sum_{i} \alpha_{i} \boldsymbol{I}-\lambda v v^{\top} .
$$

Proof. Let $\alpha \stackrel{\text { def }}{=} \sum_{i} \alpha_{i}$. Since clearly $v^{\top} \boldsymbol{A} v=v^{\top}\left(\sum_{i} \alpha_{i} \boldsymbol{I}-\lambda v v^{\top}\right) v$ it suffices to show that for $w \perp v$ it is the case that $\frac{1}{2} \alpha\|w\|_{2}^{2} \preceq w^{\top} \boldsymbol{A} w \preceq \alpha\|w\|_{2}^{2}$ or equivalently, that $\lambda_{i}(\boldsymbol{A}) \in\left[\frac{1}{2} \alpha, \alpha\right]$ for $i \neq d$. However we know that $\sum_{i \in[d]} \lambda_{i}(\boldsymbol{A})=\operatorname{Tr}(\boldsymbol{A})=d \alpha-\sum_{i} \beta_{i} \geq(d-1) \alpha$ and $\lambda_{i}(\boldsymbol{A}) \leq \alpha$ for all $i \in[d]$. Consequently, since $\lambda_{d}(\boldsymbol{A}) \in\left[0, \lambda_{d-1}\right]$ we have

$$
2 \cdot \lambda_{d-1}(\boldsymbol{A}) \geq(d-1) \alpha-\sum_{i=1}^{d-2} \lambda_{i}(\boldsymbol{A}) \geq(d-1) \alpha-(d-2) \alpha=\alpha .
$$

Consequently, $\lambda_{d-1}(\boldsymbol{A}) \in\left[\frac{\alpha}{2}, \alpha\right]$ and the result holds by the monotonicity of $\lambda_{i}$.
Next we bound the spectral difference between the outer product of two unit vectors by their inner product. We use this lemma to bound the amount of precision required in our eigenvector computations.

Lemma 7.9.2. For unit vectors $u_{1}$ and $u_{2}$ we have

$$
\begin{equation*}
\left\|u_{1} u_{1}^{\top}-u_{2} u_{2}^{\top}\right\|_{2}^{2}=1-\left(u_{1}^{\top} u_{2}\right)^{2} \tag{7.14}
\end{equation*}
$$

Consequently if $\left(u_{1}^{\top} u_{2}\right)^{2} \geq 1-\varepsilon$ for $\varepsilon \leq 1$ we have that

$$
-\sqrt{\varepsilon} \boldsymbol{I} \preceq u_{1} u_{1}^{\top}-u_{2} u_{2}^{\top} \preceq \sqrt{\varepsilon} \boldsymbol{I}
$$

Proof. Note that $u_{1} u_{1}^{\top}-u_{2} u_{2}^{\top}$ is a symmetric matrix and all eigenvectors are either orthogonal to both $u_{1}$ and $u_{2}$ (with eigenvalue 0 ) or are of the form $v=\alpha u_{1}+\beta u_{2}$ where $\alpha$ and $\beta$ are real numbers that are not both 0 . Thus, if $v$ is an eigenvector of eigenvalue $\lambda$ it must be that

$$
\begin{aligned}
\lambda\left(\alpha u_{1}+\beta u_{2}\right) & =\left(u_{1} u_{1}^{\top}-u_{2} u_{2}^{\top}\right)\left(\alpha u_{1}+\beta u_{2}\right) \\
& =\left(\alpha+\beta\left(u_{1}^{\top} u_{2}\right)\right) u_{1}-\left(\alpha\left(u_{1}^{\top} u_{2}\right)+\beta\right) u_{2}
\end{aligned}
$$

or equivalently

$$
\left(\begin{array}{cc}
(1-\lambda) & u_{1}^{\top} u_{2} \\
-\left(u_{1}^{\top} u_{2}\right) & -(1+\lambda)
\end{array}\right)\binom{\alpha}{\beta}=\binom{0}{0}
$$

By computing the determinant we see this has a solution only when

$$
-\left(1-\lambda^{2}\right)+\left(u_{1}^{\top} u_{2}\right)^{2}=0
$$

Solving for $\lambda$ then yields (7.14) and completes the proof.
Next we show how the top eigenvectors of two spectrally similar matrices are related. We use this to bound the amount of spectral approximation we need to obtain accuracate eigenvector approximations.
Lemma 7.9.3. Let $\boldsymbol{A}$ and $\boldsymbol{B}$ be symmetric PSD matrices such that $(1-\varepsilon) \boldsymbol{A} \preceq \boldsymbol{B} \preceq(1+\varepsilon) \boldsymbol{A}$. Then if $g \stackrel{\text { def }}{=} \frac{\lambda_{1}(\boldsymbol{A})-\lambda_{2}(\boldsymbol{A})}{\lambda_{1}(\boldsymbol{A})}$ satisfies $g>0$ we have $\left[v_{1}(\boldsymbol{A})^{\top} v_{1}(\boldsymbol{B})\right]^{2} \geq 1-2(\varepsilon / g)$.

Proof. Without loss of generality $v_{1}(\boldsymbol{B})=\alpha v_{1}(\boldsymbol{A})+\beta v$ for some unit vector $v \perp v_{1}(\boldsymbol{A})$ and $\alpha, \beta \in \mathbb{R}$ such that $\alpha^{2}+\beta^{2}=1$. Now we know that

$$
v_{1}(\boldsymbol{B})^{\top} \boldsymbol{B} v_{1}(\boldsymbol{B}) \leq(1+\varepsilon) v_{1}(\boldsymbol{B})^{\top} \boldsymbol{A} v_{1}(\boldsymbol{B}) \leq(1+\varepsilon)\left[\alpha^{2} \lambda_{1}(\boldsymbol{A})+\beta^{2} \lambda_{2}(\boldsymbol{A})\right]
$$

Furthermore, by the optimality of $v_{1}(\boldsymbol{B})$ we have that

$$
v_{1}(\boldsymbol{B})^{\top} \boldsymbol{B} v_{1}(\boldsymbol{B}) \geq(1-\varepsilon) v_{1}(\boldsymbol{A})^{\top} \boldsymbol{A} v_{1}(\boldsymbol{A}) \geq(1-\varepsilon) \lambda_{1}(\boldsymbol{A}) .
$$

Now since $\beta^{2}=1-\alpha^{2}$ combining these inequalities yields

$$
(1-\varepsilon) \lambda_{1}(\boldsymbol{A}) \leq(1+\varepsilon) \alpha^{2}\left(\lambda_{1}(\boldsymbol{A})-\lambda_{2}(\boldsymbol{A})\right)+(1+\varepsilon) \lambda_{2}(\boldsymbol{A}) .
$$

Rearranging terms, using the definition of $g$, and that $g \in(0,1]$ and $\varepsilon \geq 0$ yields

$$
\alpha^{2} \geq \frac{\lambda_{1}(\boldsymbol{A})-\lambda_{2}(\boldsymbol{A})-\varepsilon\left(\lambda_{1}(\boldsymbol{A})+\lambda_{2}(\boldsymbol{A})\right)}{(1+\varepsilon)\left(\lambda_{1}(\boldsymbol{A})-\lambda_{2}(\boldsymbol{A})\right)}=1-\frac{2 \varepsilon \lambda_{1}(\boldsymbol{A})}{(1+\varepsilon)\left(\lambda_{1}(\boldsymbol{A})-\lambda_{2}(\boldsymbol{A})\right)} \geq 1-2(\varepsilon / g)
$$

Here we prove a an approximate transitivity lemma for inner products of vectors. We use this to bound the accuracy need for certain eigenvector computations.
Lemma 7.9.4. Suppose that we have vectors $v_{1}, v_{2}, v_{3} \in \mathbb{R}^{n}$ such that $\left\langle v_{1}, v_{2}\right\rangle^{2} \geq 1-\varepsilon$ and $\left\langle v_{2}, v_{3}\right\rangle^{2} \geq$ $1-\varepsilon$ for $0<\varepsilon \leq \frac{1}{2}$ then $\left\langle v_{1}, v_{3}\right\rangle^{2} \geq 1-4 \varepsilon$.

Proof. Without loss of generality, we can write $v_{1}=\alpha_{1} v_{2}+\beta_{1} w_{1}$ for $\alpha_{1}^{2}+\beta_{1}^{2}=1$ and unit vector $w_{1} \perp v_{2}$. Similarly we can write $v_{3}=\alpha_{3} v_{2}+\beta_{3} w_{3}$ for $\alpha_{3}^{2}+\beta_{3}^{2}=1$ and unit vector $w_{3} \perp v_{2}$. Now, by the inner products we know that $\alpha_{1}^{2} \geq 1-\varepsilon$ and $\alpha_{3}^{2} \geq 1-\varepsilon$ and therefore $\left|\beta_{1}\right| \leq \sqrt{\varepsilon}$ and $\left|\beta_{3}\right| \leq \sqrt{\varepsilon}$. Consequently, since $\varepsilon \leq \frac{1}{2},\left|\beta_{1} \beta_{3}\right| \leq \varepsilon \leq 1-\varepsilon \leq\left|\alpha_{1} \alpha_{3}\right|$, and we have

$$
\begin{aligned}
\left\langle v_{1}, v_{3}\right\rangle^{2} & \geq\left\langle\alpha_{1} v_{2}+\beta_{1} w_{1}, \alpha_{3} v_{2}+\beta_{3} w_{3}\right\rangle^{2} \geq\left(\left|\alpha_{1} \alpha_{3}\right|-\left|\beta_{1} \beta_{3}\right|\right)^{2} \\
& \geq(1-\varepsilon-\varepsilon)^{2}=(1-2 \varepsilon)^{2} \geq 1-4 \varepsilon .
\end{aligned}
$$

## ■ 7.9.2 Convex Optimization

First we provide a single general lemma about about first order methods for convex optimization. We use this lemma for multiple purposes including bounding errors and quickly compute approximations to the central path.

Lemma 7.9.5. [206]Let $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ be a twice differentiable function, let $B \subseteq \mathbb{R}$ be a convex set, and let $x_{*}$ be a point that achieves the minimum value of $f$ restricted to $B$. Further suppose that for a symmetric positive definite matrix $\boldsymbol{H} \in \mathbb{R}^{n \times n}$ we have that $\mu \boldsymbol{H} \preceq \nabla^{2} f(y) \preceq L \boldsymbol{H}$ for all $y \in B$. Then for all $x \in B$ we have

$$
\frac{1}{2 L}\|\nabla f(x)\|_{\boldsymbol{H}^{-1}}^{2} \leq f(x)-f\left(x_{*}\right) \leq \frac{L}{2}\left\|x-x_{*}\right\|_{\boldsymbol{H}}^{2}
$$

and

$$
\frac{\mu}{2}\left\|x-x_{*}\right\|_{\boldsymbol{H}}^{2} \leq f(x)-f\left(x_{*}\right) \leq \frac{1}{2 \mu}\|\nabla f(x)\|_{\boldsymbol{H}^{-1}}^{2} .
$$

Furthermore, if

$$
x_{1}=\underset{x \in B}{\arg \min }\left[f\left(x_{0}\right)+\left\langle\nabla f\left(x_{0}\right), x-x_{0}\right\rangle+\frac{L}{2}\left\|x_{0}-x\right\|_{\boldsymbol{H}}^{2}\right]
$$

then

$$
\begin{equation*}
f\left(x_{1}\right)-f\left(x_{*}\right) \leq\left(1-\frac{\mu}{L}\right)\left(f\left(x_{0}\right)-f\left(x_{*}\right)\right) . \tag{7.15}
\end{equation*}
$$

Next we provide a short technical lemma about the convexity of functions that arises naturally in our line searching procedure.
Lemma 7.9.6. Let $f: \mathbb{R}^{n} \rightarrow \mathbb{R} \cup\{\infty\}$ be a convex function and and let $g(\alpha) \xlongequal{\text { def }} \min _{x \in S} f(x+\alpha d)$ for any convex set $S$ and $d \in \mathbb{R}^{n}$. Then $g$ is convex.

Proof. Let $\alpha, \beta \in \mathbb{R}$ and define $x_{\alpha}=\arg \min _{x \in S} f(x+\alpha d)$ and $x_{\beta}=\arg \min _{x \in S} f(x+\beta d)$. For any $t \in[0,1]$ we have

$$
\begin{aligned}
g(t \alpha+(1-t) \beta) & =\min _{x \in S} f(x+(t \alpha+(1-t) \beta) & & \\
& \leq f\left(t x_{\alpha}+(1-t) x_{\beta}+(t \alpha+(1-t) \beta) d\right) & & \text { (Convexity of } S \text { ) } \\
& \leq t \cdot f\left(x_{\alpha}+\alpha d\right)+(1-t) \cdot f\left(x_{\beta}+\beta \cdot d\right) & & (\text { Convexity of } f) \\
& =t \cdot g(\alpha)+(1-t) \cdot g(\beta) & &
\end{aligned}
$$

Lemma 7.9.7. For any vectors $y, z, v \in \mathbb{R}^{d}$ and scalar $\alpha$, we can minimize $\min _{\|x-y\|_{2}^{2} \leq \alpha}\|x-z\|_{I-v v^{\top}}^{2}$ exactly in time $O(d)$.

Proof. Let $x^{*}$ be the solution of this problem. If $\left\|x^{*}-y\right\|_{2}^{2}<\alpha$, then $x^{*}=z$. Otherwise, there is $\lambda>0$ such that $x^{*}$ is the minimizer of

$$
\min _{x \in \mathbb{R}^{d}}\|x-z\|_{I-v v^{\top}}^{2}+\lambda\|x-y\|_{2}^{2}
$$

Let $\boldsymbol{Q}=\boldsymbol{I}-v v^{\top}$. Then, the optimality condition of the above equation shows that

$$
\boldsymbol{Q}\left(x^{*}-z\right)+\lambda\left(x^{*}-y\right)=\overrightarrow{0} .
$$

Therefore,

$$
\begin{equation*}
x^{*}=(\boldsymbol{Q}+\lambda \boldsymbol{I})^{-1}(\boldsymbol{Q} z+\lambda y) . \tag{7.16}
\end{equation*}
$$

Hence,

$$
\alpha=\|x-y\|_{2}^{2}=(z-y)^{\top} \boldsymbol{Q}(\boldsymbol{Q}+\lambda \boldsymbol{I})^{-2} \boldsymbol{Q}(z-y) .
$$

Let $\eta=1+\lambda$, then we have $(\boldsymbol{Q}+\lambda \boldsymbol{I})=\eta \boldsymbol{I}-v v^{\top}$ and hence Sherman-Morrison formula shows that

$$
(\boldsymbol{Q}+\lambda \boldsymbol{I})^{-1}=\eta^{-1} \boldsymbol{I}+\frac{\eta^{-2} v v^{\top}}{1-\|v\|^{2} \eta^{-1}}=\eta^{-1}\left(\boldsymbol{I}+\frac{v v^{\top}}{\eta-\|v\|^{2}}\right) .
$$

Hence, we have

$$
(\boldsymbol{Q}+\lambda \boldsymbol{I})^{-2}=\eta^{-2}\left(I+\frac{2 v v^{\top}}{\eta-\|v\|^{2}}+\frac{v v^{\top}\|v\|^{2}}{\left(\eta-\|v\|^{2}\right)^{2}}\right)=\eta^{-2}\left(I+\frac{2 \eta-\|v\|^{2}}{\left(\eta-\|v\|^{2}\right)^{2}} v v^{\top}\right)
$$

Let $c_{1}=\|\boldsymbol{Q}(z-y)\|_{2}^{2}$ and $c_{2}=\left(v^{\top} \boldsymbol{Q}(z-y)\right)^{2}$, then we have

$$
\alpha=\eta^{-2}\left(c_{1}+\frac{2 \eta-\|v\|^{2}}{\left(\eta-\|v\|^{2}\right)^{2}} c_{2}\right) .
$$

Hence, we have

$$
\alpha \eta^{2}\left(\eta-\|v\|^{2}\right)^{2}=c_{1}\left(\eta-\|v\|^{2}\right)^{2}+c_{2}\left(2 \eta-\|v\|^{2}\right) .
$$

Note that this is a polynomial of degree 4 in $\eta$ and all coefficients can be computed in $O(d)$ time. Solving this by explicit formula, one can test all 4 possible $\eta$ 's into the formula (7.16) of $x$. Together with trivial case $x^{*}=z$, we simply need to check among 5 cases to check which is the solution.

## ■ 7.9.3 Noisy One Dimensional Convex Optimization

Here we show how to minimize a one dimensional convex function giving a noisy oracle for evaluating the function. While this could possibly be done using general results on convex optimization with a membership oracle, the proof in one dimension is much simpler and we provide it here for completeness.

```
Algorithm 20: OneDimMinimizer \((\ell, u, \varepsilon, g, L)\)
    Input: Interval \([\ell, u] \subseteq \mathbb{R}\) and target additive error \(\varepsilon \in \mathbb{R}\)
    Input: noisy additive evaluation oracle \(g: \mathbb{R} \rightarrow \mathbb{R}\) and Lipschitz bound \(L>0\)
    Let \(x=\ell, y_{\ell}^{(0)}=\ell, y_{u}^{(0)}=u\)
    for \(i=1, \ldots,\left\lceil\log _{3 / 2}\left(\frac{L(u-\ell)}{\varepsilon}\right)\right\rceil\) do
        Let \(z_{\ell}^{(i)}=\frac{2 y_{\ell}^{(i-1)}+y_{u}^{(i-1)}}{3}\) and \(z_{u}^{(i)}=\frac{y_{\ell}^{(i-1)}+2 y_{u}^{(i-1)}}{3}\)
        if \(g\left(z_{\ell}^{(i)}\right) \leq g\left(z_{u}^{(i)}\right)\) then
            Let \(\left(y_{\ell}^{(i)}, y_{u}^{(i)}\right)=\left(y_{\ell}^{(i-1)}, z_{u}^{(i)}\right)\).
            If \(g\left(z_{\ell}^{(i)}\right) \leq g\left(x^{(i-1)}\right)\) update \(x:=z_{\ell}^{(i)}\)..
        else if \(g\left(z_{\ell}^{(i)}\right)>g\left(z_{u}^{(i)}\right)\) then
            Let \(\left(y_{\ell}^{(i)}, y_{u}^{(i)}\right)=\left(z_{\ell}^{(i)}, y_{u}^{(i-1)}\right)\).
            If \(g\left(z_{u}^{(i)}\right) \leq g\left(x^{(i-1)}\right)\) update \(x:=z_{u}^{(i)}\).
        end
    end
    Output: \(x\)
```

Lemma 7.9.8. Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be an L-Lipschitz convex function defined on the $[\ell, u]$ interval and let $g: \mathbb{R} \rightarrow \mathbb{R}$ be an oracle such that $|g(y)-f(y)| \leq \varepsilon$ for all $y$. In $O\left(\log \left(\frac{L(u-\ell)}{\varepsilon}\right)\right)$ time and with $O\left(\log \left(\frac{L(u-\ell)}{\varepsilon}\right)\right)$ calls to $g$, the algorithm OneDimMinimizer $(\ell, u, \varepsilon, g, L)$ outputs a point $x$ such that

$$
f(x)-\min _{y \in[\ell, u]} f(y) \leq 4 \varepsilon .
$$

Proof. First, note that for any $y, y^{\prime} \in \mathbb{R}$ if $f(y)<f\left(y^{\prime}\right)-2 \varepsilon$ then $g(y)<g\left(y^{\prime}\right)$. This directly follows from our assumption on on $g$. Second, note that the output of the algorithm, $x$, is simply the point
queried by the algorithm (i.e. $\ell$ and the $z_{\ell}^{i}$ and $z_{u}^{i}$ ) with the smallest value of $g$. Combining these facts implies that $f(x)$ is within $2 \varepsilon$ of the minimum value of $f$ among the points queried. It thus suffices to show that the algorithm queries some point within $2 \varepsilon$ of optimal.

To do this, we break into two cases. First, consider the case where the intervals $\left[y_{\ell}^{(i)}, y_{u}^{(i)}\right]$ all contain a minimizer of $f$. In this case, the final interval contains an optimum, and is of size at most $\frac{\varepsilon}{L}$. Thus, by the Lipschitz property, all points in the interval are within $\varepsilon \leq 2 \varepsilon$ of optimal, and at least one endpoint of the interval must have been queried by the algorithm.

For the other case, consider the last $i$ for which this interval does contain an optimum of $f$. This means that $g\left(z_{\ell}^{(i)}\right) \leq g\left(z_{u}^{(i)}\right)$ while a minimizer $x^{*}$ is to the right of $z_{u}^{(i)}$, or the symmetric case with a minimizer is to the left of $z_{\ell}^{(i)}$. Without loss of generality, we assume the former. We then have $z_{\ell}^{(i)} \leq z_{u}^{(i)} \leq x^{*}$ and $x^{*}-z_{u}^{(i)} \leq z_{u}^{(i)}-z_{\ell}^{(i)}$. Consequently $z_{u}^{(i)}=\alpha z_{l}^{(i)}+(1-\alpha) x^{*}$ where $\alpha \in\left[0, \frac{1}{2}\right]$ and the convexity of $f$ implies $\left.f\left(z_{u}^{(i)}\right) \leq \frac{1}{2} f\left(z_{l}^{(i)}\right)+\frac{1}{2} f\left(x^{*}\right)\right)$ or equivalently $f\left(z_{u}^{(i)}\right)-f\left(x^{*}\right) \leq f\left(z_{\ell}^{(i)}\right)-f\left(z_{u}^{(i)}\right)$. But $f\left(z_{\ell}^{(i)}\right)-f\left(z_{u}^{(i)}\right) \leq 2 \varepsilon$ since $g\left(z_{\ell}^{(i)}\right) \leq g\left(z_{u}^{(i)}\right)$. Thus, $f\left(z_{u}^{(i)}\right)-f\left(x^{*}\right) \leq 2 \varepsilon$, and $z_{u}^{(i)}$ is queried by the algorithm, as desired.

## ■ 7.10 Weighted Geometric Median

In this section, we show how to extend our results to the weighted geometric median problem, also known as the Weber problem: given a set of $n$ points in $d$ dimensions, $a^{(1)}, \ldots, a^{(n)} \in \mathbb{R}^{d}$, with corresponding weights $w^{(1)}, \ldots, w^{(n)} \in \mathbb{R}_{>0}$, find a point $x_{*} \in \mathbb{R}^{d}$ that minimizes the weighted sum of Euclidean distances to them:

$$
x_{*} \in \underset{x \in \mathbb{R}^{d}}{\arg \min } f(x) \quad \text { where } \quad f(x) \stackrel{\text { def }}{=} \sum_{i \in[n]} w^{(i)}\left\|x-a^{(i)}\right\|_{2} .
$$

As in the unweighted problem, our goal is to compute $(1+\varepsilon)$-approximate solution, i.e. $x \in \mathbb{R}^{d}$ with $f(x) \leq(1+\varepsilon) f\left(x_{*}\right)$.

First, we show that it suffices to consider the case where the weights are integers with bounded sum (Lemma 7.10.1). Then, we show that such an instance of the weighted geometric median problem can be solved using the algorithms developed for the unweighted problem.
Lemma 7.10.1. Given points $a^{(1)}, a^{(2)}, \ldots, a^{(n)} \in \mathbb{R}^{d}$, non-negative weights $w^{(1)}, w^{(2)}, \ldots, w^{(n)} \in \mathbb{R}_{>0}$, and $\varepsilon \in(0,1)$, we can compute in linear time weights $w_{1}^{(1)}, w_{1}^{(2)}, \ldots, w_{1}^{(n)}$ such that:

- Any $(1+\varepsilon / 5)$-approximate weighted geometric median of $a^{(1)}, \ldots, a^{(n)}$ with the weights $w_{1}^{(1)}, \ldots, w_{1}^{(n)}$ is also a $(1+\varepsilon)$-approximate weighted geometric median of $a^{(1)}, \ldots, a^{(n)}$ with the weights $w^{(1)}, \ldots, w^{(n)}$, and
- $w_{1}^{(1)}, \ldots, w_{1}^{(n)}$ are nonnegative integers and $\sum_{i \in[n]} w_{1}^{(i)} \leq 5 n \varepsilon^{-1}$.

Proof. Let

$$
f(x)=\sum_{i \in[n]} w^{(i)}\left\|a^{(i)}-x\right\|
$$

and $W=\sum_{i \in[n]} w^{(i)}$. Furthermore, let $\varepsilon^{\prime}=\varepsilon / 5$ and for each $i \in[n]$, define $w_{0}^{(i)}=\frac{n}{\varepsilon^{\prime} W} w^{(i)}, w_{1}^{(i)}=$ $\left\lfloor w_{0}^{(i)}\right\rfloor$, and $w_{2}^{(i)}=w_{0}^{(i)}-w_{1}^{(i)}$. We also define $f_{0}, f_{1}, f_{2}, W_{0}, W_{1}, W_{2}$ analogously to $f$ and $W$.

Now, assume $f_{1}(x) \leq\left(1+\varepsilon^{\prime}\right) f_{1}\left(x_{*}\right)$, where $x_{*}$ is the minimizer of $f$ and $f_{0}$. Then:

$$
f_{0}(x)=f_{1}(x)+f_{2}(x) \leq f_{1}(x)+f_{2}\left(x_{*}\right)+W_{2}\left\|x-x_{*}\right\|_{2}
$$

and

$$
\begin{aligned}
W_{2}\left\|x-x_{*}\right\|_{2} & =\frac{W_{2}}{W_{1}} \sum_{i \in[n]} w_{1}^{(i)}\left\|x-x_{*}\right\|_{2} \leq \frac{W_{2}}{W_{1}} \sum_{i \in[n]} w_{1}^{(i)}\left(\left\|x-a^{(i)}\right\|_{2}+\left\|a^{(i)}-x_{*}\right\|\right) \\
& \leq \frac{W_{2}}{W_{1}}\left(f_{1}(x)+f_{1}\left(x_{*}\right)\right) .
\end{aligned}
$$

Now, since $W_{0}=\frac{n}{\varepsilon^{\prime}}$ and $W_{1} \geq W_{0}-n$ we have

$$
\frac{W_{2}}{W_{1}}=\frac{W_{0}-W_{1}}{W_{1}}=\frac{W_{0}}{W_{1}}-1 \leq \frac{W_{0}}{W_{0}-n}-\frac{W_{0}-n}{W_{0}-n}=\frac{n}{\frac{n}{\varepsilon^{\prime}}-n}=\frac{\varepsilon^{\prime}}{1-\varepsilon^{\prime}} .
$$

Combining these yields that

$$
\begin{aligned}
f_{0}(x) & \leq f_{1}(x)+f_{2}\left(x_{*}\right)+\frac{\varepsilon^{\prime}}{1-\varepsilon^{\prime}}\left(f_{1}(x)+f_{1}\left(x_{*}\right)\right) \\
& \leq\left(1+\frac{\varepsilon^{\prime}}{1-\varepsilon^{\prime}}\right)\left(1+\varepsilon^{\prime}\right) f_{1}\left(x^{*}\right)+\frac{\varepsilon^{\prime}}{1-\varepsilon^{\prime}} f_{1}\left(x_{*}\right)+f_{2}\left(x_{*}\right) \\
& \leq\left(1+5 \varepsilon^{\prime}\right) f_{0}\left(x_{*}\right)=(1+\varepsilon) f_{0}\left(x_{*}\right) .
\end{aligned}
$$

We now proceed to show the main result of this section.
Lemma 7.10.2. A $(1+\varepsilon)$-approximate weighted geometric median of $n$ points in $\mathbb{R}^{d}$ can be computed in $O\left(n d \log ^{3} \varepsilon^{-1}\right)$ time.

Proof. By applying Lemma 7.10.1, we can assume that the weights are integer and their sum does not exceed $n \varepsilon^{-1}$. Note that computing the weighted geometric median with such weights is equivalent to computing an unweighted geometric median of $O\left(n \varepsilon^{-1}\right)$ points (where each point of the original input is repeated with the appropriate multiplicity). We now show how to simulate the behavior of our unweighted geometric median algorithms on such a set of points without computing it explicitly.

If $\varepsilon>n^{-1 / 2}$, we will apply the algorithm ApproximateMedian $(\varepsilon)$, achieving a runtime of $O\left(d \varepsilon^{-2}\right)=$ $O(n d)$. It is only necessary to check that we can implement weighted sampling from our points with $O(n)$ preprocessing and $O(1)$ time per sample. This is achieved by the alias method [154].

Now assume $\varepsilon<n^{-1 / 2}$. We will employ the algorithm AccurateMedian $(\varepsilon)$. Note that the initialization using ApproximateMedian(2) can be performed as described in the previous paragraph. It remains to check that we can implement the subroutines LineSearch and ApproxMinEig on the implicitly represented multiset of $O\left(n \varepsilon^{-1}\right)$ points. It is enough to observe only $n$ of the points are distinct, and all computations performed by these subroutines are identical for identical points. The total runtime will thus be $O\left(n d \log ^{3}\left(n / \varepsilon^{2}\right)\right)=O\left(n d \log ^{3} \varepsilon^{-1}\right)$.

Part II

## Cutting



# Convex Minimization In Nearly-Cubic Time 

## ■ 8.1 Introduction

Throughout this chapter, we study the following feasibility problem:
Definition 8.1.1 (Feasibility Problem). Given a separation oracle (see Definition 2.3.5) for a set $K \subseteq \mathbb{R}^{n}$ contained in a box of radius $R$ either find a point $\boldsymbol{x} \in K$ or prove that $K$ does not contain a ball of radius $\varepsilon$.

This feasibility problem is one of the most fundamental and classic problems in optimization. Since the celebrated result of Shor [236], Yudin and Nemirovski [274] and Khachiyan [142] in 1970s essentially proving that it can be solved in time $O(\operatorname{poly}(n) \cdot \mathrm{SO} \cdot \log (R / \varepsilon))$, this problem has served as one of the key primitives for solving numerous problems in both combinatorial and convex optimization.

Despite the prevalence of this feasibility problem, the best known running time for solving this problem has not been improved in over 25 years. In a seminal paper of Vaidya in 1989 [253], he showed how to solve the problem in $\widetilde{O}\left(n \cdot \mathrm{SO} \cdot \log (n R / \varepsilon)+n^{\omega+1} \log (n R / \varepsilon)\right)$ time. Despite interesting generalizations and practical improvements [11, 226, 101, 20, 103, 204, 272, 102], the best theoretical guarantees for solving this problem have not been improved since.

In this chapter, we show how to improve upon Vaidya's running time in certain regimes. We provide a cutting plane algorithm which achieves an expected running time of $O(n \cdot \mathrm{SO} \cdot \log (n R / \varepsilon)+$ $\left.n^{3} \log ^{O(1)}(n R / \varepsilon)\right)$, improving upon the previous best known running time for the current known value of $\omega<2.373$ [269, 99] when $R / \varepsilon=O(\operatorname{poly}(n))$.

We achieve our results by the combination of multiple techniques. First we show how to use techniques from the work of Vaidya and Atkinson to modify Vaidya's scheme so that it is able to tolerate random noise in the computation in each iteration. We then show how to use known numerical machinery [257, 242, 164] in combination with some new techniques (Section 8.4.1 and Section 8.4.2) to implement each of these relaxed iterations efficiently. We hope that both these numerical techniques as well as our scheme for approximating complicated methods, such as Vaidya's, may find further applications.

## ■ 8.1.1 Previous Work

Throughout this chapter, we restrict our attention to algorithms for the feasibility problem that have a polynomial dependence on SO, $n$, and $\log (R / \varepsilon)$. Such "efficient" algorithms typically follow the following iterative framework. First, they compute some trivial region $\Omega$ that contains $K$. Then, they call the separation oracle at some point $\boldsymbol{x} \in \Omega$. If $\boldsymbol{x} \in K$, the algorithm successfully solved the problem. If $\boldsymbol{x} \notin K$ then the separation oracle must return a half-space containing $K$. The algorithm then uses this half-space to shrink the region $\Omega$ while maintaining the invariant that $K \subseteq \Omega$. The algorithm then repeats this process until it finds a point $\boldsymbol{x} \in K$ or the region $\Omega$ becomes too small to contain a ball with radius $\varepsilon$.

| Year | Algorithm | Complexity |
| :---: | :--- | :---: |
| 1979 | Ellipsoid Method [236, 274, 142] | $O\left(n^{2} \mathrm{SO} \log \kappa+n^{4} \log \kappa\right)$ |
| 1988 | Inscribed Ellipsoid [144, 213] | $O\left(n \mathrm{SOlog} \kappa+(n \log \kappa)^{4.5}\right)$ |
| 1989 | Volumetric Center [253] | $O\left(n \mathrm{SOlog} \kappa+n^{1+\omega} \log \kappa\right)$ |
| 1995 | Analytic Center [20] | $O\binom{n \mathrm{SO} \log ^{2} \kappa+n^{\omega+1} \log ^{2} \kappa}{+(n \log \kappa)^{2+\omega / 2}}$ |
| 2004 | Random Walk [36] | $\rightarrow O\left(n \mathrm{SOlog} \kappa+n^{7} \log \kappa\right)$ |
| - | This Chapter | $O\left(n \mathrm{SO} \log \kappa+n^{3} \log ^{O(1)} \kappa\right)$ |

Table 8.1: Algorithms for the Feasibility Problem. $\kappa$ indicates $n R / \varepsilon$. The arrow, $\rightarrow$, indicates that it solves a more general problem where only a membership oracle is given.

Previous works on efficient algorithms for the feasibility problem all follow this iterative framework. They vary in terms of what set $\Omega$ they maintain, how they compute the center to query the separation oracle, and how they update the set. In Table 8.1, we list the previous running times for solving the feasibility problem. As usual SO indicates the cost of the separation oracle.

The first efficient algorithm for the feasibility problem is the ellipsoid method, due to Shor [236], Nemirovksii and Yudin [274], and Khachiyan [142]. The ellipsoid method maintains an ellipsoid as $\Omega$ and uses the center of the ellipsoid as the next query point. It takes $\Theta\left(n^{2} \log (n R / \varepsilon)\right)$ calls of oracle which is far from the lower bound $\Omega(n \log (R / \varepsilon))$ calls [201].

To alleviate the problem, the algorithm could maintain all the information from the oracle, i.e., the polytope created from the intersection of all half-spaces obtained. The center of gravity method [168] achieves the optimal oracle complexity using this polytope and the center of gravity of this polytope as the next point. However, computing center of gravity is computationally expensive and hence we do not list its running time in Table 8.1. The Inscribed Ellipsoid Method [144] also achieved an optimal oracle complexity using this polytope as $\Omega$ but instead using the center of the maximal inscribed ellipsoid in the polytope to query the separation oracle. We listed it as occurring in year 1988 in Table 8.1 because it was [213] that yielded the first polynomial time algorithm to actually compute this maximal inscribed ellipsoid for polytope.

Vaidya [253] obtained a faster algorithm by maintaining an approximation of this polytope and using a different center, namely the volumetric center. Although the oracle complexity of this volumetric center method is very good, the algorithm is not extremely efficient as each iteration involves matrix inversion. Atkinson and Vaidya [20] showed how to avoid this computation in certain settings. However, they were unable to achieve the desired convergence rate from their method.

Bertsimas and Vempala [36] also gives an algorithm that avoids these expensive linear algebra operations while maintaining the optimal convergence rate by using techniques in sampling convex sets. Even better, this result works for a much weaker oracle, the membership oracle. However, the additional cost of this algorithm is relatively high. We remark that while there are considerable improvemenst on the sampling techniques [180,133,164], the additional cost is still quite high compared to standard linear algebra.

## ■ 8.1.2 Challenges in Improving Previous Work

Our algorithm builds upon the previous fastest algorithm of Vaidya [253]. Ignoring implementation details and analysis, Vaidya's algorithm is quite simple. This algorithm simply maintains a polytope $P^{(k)}=\left\{x \in \mathbb{R}^{n}: \boldsymbol{A x}-\boldsymbol{b}>\mathbf{0}\right\}$ as the current $\Omega$ and uses the volumetric center, the minimizer of the
following volumetric barrier function

$$
\begin{equation*}
\underset{\boldsymbol{x}}{\arg \min } \frac{1}{2} \log \operatorname{det}\left(\boldsymbol{A}^{T} \boldsymbol{S}_{\boldsymbol{x}}^{-2} \boldsymbol{A}\right) \quad \text { where } \quad \boldsymbol{S}_{\boldsymbol{x}} \stackrel{\text { def }}{=} \boldsymbol{D i a g}(\boldsymbol{A} \boldsymbol{x}-\boldsymbol{b}) \tag{8.1}
\end{equation*}
$$

as the point at which to query the separation oracle. The polytope is then updated by adding shifts of the half-spaces returned by the separation oracle and dropping unimportant constraints. By choosing the appropriate shift, picking the right rule for dropping constraints, and using Newton's method to compute the volumetric center he achieved a running time of $O\left(n \cdot S O \cdot \log \kappa+n^{1+\omega} \log \kappa\right)$.

While Vaidya's algorithm's dependence on SO is essentially optimal, the additional per-iteration costs of his algorithm could possibly be improved. The computational bottleneck in each iteration of Vaidya's algorithm is computing the gradient of log det which in turn involves computing the leverage scores $\boldsymbol{\sigma}(\boldsymbol{x}) \stackrel{\text { def }}{=} \boldsymbol{\operatorname { d i a g }}\left(\boldsymbol{S}_{x}^{-1} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{S}_{x}^{-2} \boldsymbol{A}\right)^{-1} \boldsymbol{A}^{T} \boldsymbol{S}_{x}^{-1}\right)$, a key concept in part I. As the best known algorithms for computing leverage scores exactly in this setting take time $O\left(n^{\omega}\right)$, directly improving the running time of Vaidya's algorithm seems challenging.

However, similar to previous chapters, these leverage scores can be computed up to a multiplicative $(1 \pm \varepsilon)$ error by solving $O\left(\varepsilon^{-2} \log n\right)$ linear systems involving $\boldsymbol{A}^{T} \boldsymbol{S}_{x}^{-2} \boldsymbol{A}$. While in general this still takes time $O\left(\varepsilon^{-2} n^{\omega}\right)$, there are known techniques for efficiently maintaining the inverse of a matrix so that solving linear systems take amortized $O\left(n^{2}\right)$ time [257, 164]. Consequently if it could be shown that computing approximate leverage scores sufficed, this would potentially decrease the amortized cost per iteration of Vaidya's method.

Unfortunately, Vaidya's method does not seem to tolerate this type of multiplicative error. If we compute compute gradients for (8.1) using this approximation of leverage scores, it seems that the point computed would be far from the true center. Moreover, without being fairly close to the true volumetric center, it is difficult to argue that such a cutting plane method would make sufficient progress.

To overcome this issue, it is tempting to directly use the linear program result in Chapter 6. In this chapter, we faced a similar issue where a volumetric, i.e. log det, potential function had the right analytic and geometric properties, however was computational expensive to minimize. To overcome this issue the authors instead computed a weighted analytic center:

$$
\underset{\boldsymbol{x}}{\arg \min }-\sum_{i \in[m]} w_{i} \log s_{i}(\boldsymbol{x}) \quad \text { where } \quad \boldsymbol{s}(\boldsymbol{x}) \stackrel{\text { def }}{=} \boldsymbol{A} \boldsymbol{x}-\boldsymbol{b}
$$

For carefully chosen weights this center provides the same convergence guarantees as the volumetric potential function, while each step can be computed by solving few linear systems (rather than forming the matrix inverse).

Unfortunately, it is unclear how to directly extend the work in Chapter 6 on solving an explicit linear program to the feasibility problem specified by a separation oracle. While it is possible to approximate the volumetric barrier by a weighted analytic center in many respects, proving that this approximation suffices for fast convergence remains open. In fact, the volumetric barrier function as used in Vaidya's algorithm is well approximated simply by the standard analytic center

$$
\underset{\boldsymbol{x}}{\arg \min }-\sum_{i \in[m]} \log s_{i}(\boldsymbol{x}) \quad \text { where } \quad \boldsymbol{s}(\boldsymbol{x}) \stackrel{\text { def }}{=} \boldsymbol{A} \boldsymbol{x}-\boldsymbol{b}
$$

as all the unimportant constraints are dropped during the algorithm. However, despite decades of research, the best running times known for solving the feasibility problem using the analytic center are Vaidya and Atkinson algorithm from 1995 [20]. While the running time of this algorithm could possibly be improved using approximate leverage score computations and amortized efficient linear
system solvers, unfortunately at best, without further insight this would yield an algorithm which requires a suboptimal $O\left(n \log ^{O(1)} \kappa\right)$ queries to the separation oracle.

As pointed out in [20], the primary difficulty in using any sort of analytic center is quantifying the amount of progress made in each step. We still believe providing direct near-optimal analysis of weighted analytic center is a tantalizing open question warranting further investigation. However, rather than directly address the question of the performance of weighted analytic centers for the feasibility problem, we take a slightly different approach that side-steps this issue. We provide a partial answer that still sheds some light on the performance of the weighted analytic center while still providing our desired running time improvements.

## ■ 8.1.3 Our Approach

To overcome the shortcoming of the volumetric and analytic centers we instead consider a hybrid barrier function

$$
\underset{\boldsymbol{x}}{\arg \min }-\sum_{i \in[m]} w_{i} \log s_{i}(\boldsymbol{x})+\log \operatorname{det}\left(\boldsymbol{A}^{T} \boldsymbol{S}_{x}^{-1} \boldsymbol{A}\right) \quad \text { where } \quad \boldsymbol{s}(\boldsymbol{x}) \stackrel{\text { def }}{=} \boldsymbol{A} \boldsymbol{x}-\boldsymbol{b} .
$$

for careful chosen weights. Our key observation is that for correct choice of weights, we can compute the gradient of this potential function. In particular if we let $\boldsymbol{w}=\boldsymbol{\tau}-\boldsymbol{\sigma}(\boldsymbol{x})$ then the gradient of this potential function is the same as the gradients of $\sum_{i \in[m]} \tau_{i} \log s_{i}(\boldsymbol{x})$, which we can compute efficiently. Moreover, since we are using log det, we can use analysis similar to Vaidya's algorithm [253] to analyze the convergence rate of this algorithm.

Unfortunately, this is a simple observation and does not immediately change the problem substantially. It simply pushes the problem of computing gradients of log det to computing $\boldsymbol{w}$. Therefore, for this scheme to work, we would need to ensure that the weights do not change too much and that when they change, they do not significantly hurt the progress of our algorithm. In other words, for this scheme to work, we would still need very precise estimates of leverage scores.

However, we note that the leverage scores $\boldsymbol{\sigma}(\boldsymbol{x})$ do not change too much between iterations. Moreover, we provide what we believe is an interesting technical result that an unbiased estimates to the changes in leverage scores can be computed using linear system solvers such that the total error of the estimate is bounded by the total change of the leverage scores (See Section 8.4.1). Using this result our scheme simply follows Vaidya's basic scheme in [253], however instead of minimizing the hybrid barrier function directly we alternate between taking Newton steps we can compute, changing the weights so that we can still compute Newton steps, and computing accurate unbiased estimates of the changes in the leverage scores so that the weights do not change adversarially by too much.

To make this scheme work, there are two additional details that need to be dealt with. First, we cannot let the weights vary too much as this might ultimately hurt the rate of progress of our algorithm. Therefore, in every iteration we compute a single leverage score to high precision to control the value of $w_{i}$ and we show that by careful choice of the index we can ensure that no weight gets too large (See Section 8.4.2).

Second, we need to show that changing weights does not affect our progress by much more than the progress we make with respect to log det. To do this, we need to show the slacks are bounded above and below. We enforce this by adding regularization terms and consider the potential function

$$
p_{\boldsymbol{e}}(\boldsymbol{x})=-\sum_{i \in[m]} w_{i} \log s_{i}(\boldsymbol{x})+\frac{1}{2} \log \operatorname{det}\left(\boldsymbol{A}^{T} \boldsymbol{S}_{x}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)+\frac{\lambda}{2}\|x\|_{2}^{2}
$$

This allows us to ensure that the entries of $\boldsymbol{s}(\boldsymbol{x})$ do not get too large or too small and therefore
changing the weighting of the analytic center cannot affect the function value too much.
Third, we need to make sure our potential function is convex. If we simply take $\boldsymbol{w}=\boldsymbol{\tau}-\boldsymbol{\sigma}(\boldsymbol{x})$ with $\boldsymbol{\tau}$ as an estimator of $\boldsymbol{\sigma}(\boldsymbol{x}), \boldsymbol{w}$ can be negative and the potential function could be non-convex. To circumvent this issue, we use $\boldsymbol{w}=c_{e}+\boldsymbol{\tau}-\boldsymbol{\sigma}(\boldsymbol{x})$ and make sure $\|\boldsymbol{\tau}-\boldsymbol{\sigma}(\boldsymbol{x})\|_{\infty}<c_{e}$.

Combining these insights, using efficient algorithms for solving a sequence of slowly changing linear systems [257, 164], and providing careful analysis ultimately allows us to achieve a running time of $O\left(n \mathrm{SO} \log \kappa+n^{3} \log ^{O(1)} \kappa\right)$ for the feasibility problem. Furthermore, in the case that $K$ does not contain a ball of radius $\varepsilon$, our algorithm provides a proof that the polytope does not contain a ball of radius $\varepsilon$. This proof ultimately allows us to achieve running time improvements for strongly polynomial submodular minimization in Chapter 10.

## ■ 8.1.4 Organization

The rest of Chapter 8 is organized as follows. In Section 8.2 we provide some preliminary information and notation we use throughout Chapter 8. In Section 8.3 we then provide and analyze our cutting plane method. In Section 8.4 we provide key technical tools which may be of independent interest.

## ■ 8.2 Preliminaries

Here we introduce some notation and concepts we use throughout this chapter.

## © 8.2.1 Hybrid Barrier Function

As explained in Section 8.1 .3 our cutting plane method maintains a polytope $P=\left\{\boldsymbol{x} \in \mathbb{R}^{n}: \boldsymbol{A} \boldsymbol{x} \geq \boldsymbol{b}\right\}$ for $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ and $\boldsymbol{b} \in \mathbb{R}^{n}$ that contains some target set $K$. We then maintain a minimizer of the following hybrid barrier function:

$$
p_{\boldsymbol{e}}(\boldsymbol{x}) \stackrel{\text { def }}{=}-\sum_{i \in[m]}\left(c_{e}+e_{i}\right) \log \left(s_{i}(\boldsymbol{x}) / R\right)+\frac{1}{2} \log \operatorname{det}\left(R^{2}\left(\boldsymbol{A}^{T} \boldsymbol{S}_{x}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)\right)+\frac{\lambda}{2}\|x\|_{2}^{2}
$$

where $\boldsymbol{e} \in \mathbb{R}^{m}$ is a variable we maintain, $c_{e} \geq 0$ and $\lambda \geq 0$ are constants we fix later, $\boldsymbol{s}(\boldsymbol{x}) \stackrel{\text { def }}{=} \boldsymbol{A} \boldsymbol{x}-\boldsymbol{b}$, $\boldsymbol{S}_{x}=\operatorname{Diag}(\boldsymbol{s}(\boldsymbol{x}))$, and $R$ is the radius of the box containing $P$. When the meaning is clear from context we often use the shorthand $\boldsymbol{A}_{x} \stackrel{\text { def }}{=} \boldsymbol{S}_{x}^{-1} \boldsymbol{A}$. The $R$ term is only used in the proof and can be ignored in the algorithm because it only affects a constant term in the potential function.

Rather than maintaining $\boldsymbol{e}$ explicitly, we instead maintain a vector $\boldsymbol{\tau} \in \mathbb{R}^{m}$ that approximates the leverage score

$$
\boldsymbol{\psi}(\boldsymbol{x}) \stackrel{\text { def }}{=} \boldsymbol{d i a g}\left(\boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{A}_{x}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{A}_{x}^{T}\right)
$$

and pick $\boldsymbol{e}$ using the function $\boldsymbol{e}_{P}(\boldsymbol{\tau}, \boldsymbol{x}) \stackrel{\text { def }}{=} \boldsymbol{\tau}-\boldsymbol{\psi}(\boldsymbol{x})$. Note that $\boldsymbol{\psi}(\boldsymbol{x})$ is simply the leverage scores of certain rows of the matrix

$$
\left[\begin{array}{c}
\boldsymbol{A}_{x} \\
\sqrt{\lambda} \boldsymbol{I}
\end{array}\right]
$$

and therefore the usual properties of leverage scores hold, i.e. $\psi_{i}(\boldsymbol{x}) \in(0,1)$ and $\left\|\psi_{i}(\boldsymbol{x})\right\|_{1} \leq n$. We write $\boldsymbol{\psi}(\boldsymbol{x})$ equivalently as $\boldsymbol{\psi}_{x}$ or $\boldsymbol{\psi}_{P}$ when we want the matrix to be clear. Furthermore, we let $\boldsymbol{\Psi}_{x} \stackrel{\text { def }}{=} \boldsymbol{D i a g}(\boldsymbol{\psi}(\boldsymbol{x}))$ and $\mu(\boldsymbol{x}) \stackrel{\text { def }}{=} \min _{i} \psi_{i}(\boldsymbol{x})$. Again, we use the subscripts of $x$ and $P$ interchangeably and often drop them when the meaning is clear from context.

We remark that the last term $\frac{\lambda}{2}\|x\|_{2}^{2}$ ensures that our point is always within a certain region (Lemma 8.3.18) and hence the term $\left(c_{e}+e_{i}\right) \log s_{i}(\boldsymbol{x})_{i}$ never gets too large. However, this $\ell^{2}$ term
changes the Hessian of the potential function and hence we need to put a $\lambda \boldsymbol{I}$ term inside both the $\log$ det and the leverage score to reflect this. This is the reason why we use $\boldsymbol{\psi}$ instead of the standard leverage score.

### 8.3 Our Cutting Plane Method

In this section we develop and prove the correctness of our cutting plane method. We break the presentation and proof of correctness of our cutting plane methods into multiple parts. First in Section 8.3.1 we describe how we maintain a center of the hybrid barrier function $p_{e}$ and analyze this procedure. Then, in Section 8.3.2 we carefully analyze the effect of changing constraints on the hybrid barrier function and in Section 8.3.3 we prove properties of an approximate center of hybrid barrier function, which we call a hybrid center. In Section 8.3 .4 we then provide our cutting plane method and in Section 8.3.5 we prove that the cutting plane method solves the feasibility problem as desired.

## - 8.3.1 Centering

In this section we show how to compute approximate centers or minimizers of the hybrid barrier function for the current polytope $P=\{\boldsymbol{x}: \boldsymbol{A} \boldsymbol{x} \geq \boldsymbol{b}\}$. We split this proof up into multiple parts. First we simply bound the gradient and Hessian of the hybrid barrier function, $p_{e}$, as follows.

Lemma 8.3.1. For $f(\boldsymbol{x}) \stackrel{\text { def }}{=} \frac{1}{2} \log \operatorname{det}\left(\boldsymbol{A}^{T} \boldsymbol{S}_{x}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)$, we have that

$$
\nabla f(\boldsymbol{x})=-\boldsymbol{A}_{x}^{T} \boldsymbol{\psi}(\boldsymbol{x}) \quad \text { and } \quad \boldsymbol{A}_{x}^{T} \boldsymbol{\Psi}(\boldsymbol{x}) \boldsymbol{A}_{x} \preceq \nabla^{2} f(\boldsymbol{x}) \preceq 3 \boldsymbol{A}_{x}^{T} \boldsymbol{\Psi}(\boldsymbol{x}) \boldsymbol{A}_{x} .
$$

Proof. Our proof is similar to [10, Appendix] which proved the statement when $\lambda=0$. This case does not change the derivation significantly, however for completeness we include the proof below.

We take derivatives on $\boldsymbol{s}$ first and then apply chain rule. Let $f(\boldsymbol{s})=\frac{1}{2} \log \operatorname{det}\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)$. We use the notation $D f(\boldsymbol{x})[\boldsymbol{h}]$ to denote the directional derivative of $f$ along the direction $\boldsymbol{h}$ at the point $\boldsymbol{x}$. Using the standard formula for the derivative of $\log$ det, i.e. $\frac{d}{d t} \log \operatorname{det} \boldsymbol{B}_{t}=\operatorname{Tr}\left(\left(\boldsymbol{B}_{t}\right)^{-1}\left(\frac{d \boldsymbol{B}_{t}}{d t}\right)\right)$, we have

$$
\begin{align*}
D f(\boldsymbol{s})[\boldsymbol{h}] & =\frac{1}{2} \operatorname{Tr}\left(\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1}\left(\boldsymbol{A}^{T}(-2) \boldsymbol{S}^{-3} \boldsymbol{H} \boldsymbol{A}\right)\right)  \tag{8.2}\\
& =-\sum_{i} \frac{h_{i}}{s_{i}} \mathbf{1}_{i}^{T} \boldsymbol{S}^{-1} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{A} \boldsymbol{S}^{-1} \mathbf{1}_{i}=-\sum_{i} \frac{\psi_{i} h_{i}}{s_{i}} .
\end{align*}
$$

Applying chain rules, we have $\nabla f(\boldsymbol{x})=-\boldsymbol{A}_{x}^{T} \boldsymbol{\psi}$. Now let $\boldsymbol{P} \stackrel{\text { def }}{=} \boldsymbol{S}^{-1} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{A}^{T} \boldsymbol{S}^{-1}$. Taking the derivative of (8.2) again and using the cyclic property of trace, we have

$$
\begin{aligned}
D^{2} f(\boldsymbol{s})\left[\boldsymbol{h}_{1}, \boldsymbol{h}_{2}\right]= & \operatorname{Tr}\left(\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1}\left(\boldsymbol{A}^{T}(-2) \boldsymbol{S}^{-3} \boldsymbol{H}_{2} \boldsymbol{A}\right)\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1}\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-3} \boldsymbol{H}_{1} \boldsymbol{A}\right)\right) \\
& -\operatorname{Tr}\left(\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1}\left(\boldsymbol{A}^{T}(-3) \boldsymbol{S}^{-4} \boldsymbol{H}_{2} \boldsymbol{H}_{1} \boldsymbol{A}\right)\right) \\
= & 3 \operatorname{Tr}\left(\boldsymbol{P} \boldsymbol{S}^{-2} \boldsymbol{H}_{2} \boldsymbol{H}_{1}\right)-2 \operatorname{Tr}\left(\boldsymbol{P} \boldsymbol{S}^{-1} \boldsymbol{H}_{2} \boldsymbol{P} \boldsymbol{S}^{-1} \boldsymbol{H}_{1}\right) \\
= & 3 \sum_{i} P_{i i} \frac{\boldsymbol{h}_{1}(i) \boldsymbol{h}_{2}(i)}{s_{i}^{2}}-2 \sum_{i j} P_{i j} \frac{\boldsymbol{h}_{2}(j)}{s_{j}} P_{j i} \frac{\boldsymbol{h}_{2}(i)}{s_{i}} \\
= & 3 \sum_{i} \psi_{i} \frac{\boldsymbol{h}_{1}(i) \boldsymbol{h}_{2}(i)}{s_{i}^{2}}-2 \sum_{i j} P_{i j}^{2} \frac{\boldsymbol{h}_{2}(j)}{s_{j}} \frac{\boldsymbol{h}_{2}(i)}{s_{i}} .
\end{aligned}
$$

Consequently, $D^{2} f(\boldsymbol{x})\left[\mathbf{1}_{i}, \mathbf{1}_{j}\right]=\left[\boldsymbol{S}^{-1}\left(3 \boldsymbol{\Psi}-2 \boldsymbol{P}^{(2)}\right) \boldsymbol{S}^{-1}\right]_{i j}$ where $\boldsymbol{P}^{(2)}$ is the Schur product of $\boldsymbol{P}$ with itself.

Now note that

$$
\begin{aligned}
\sum_{i} P_{i j}^{2} & =\mathbf{1}_{j} \boldsymbol{S}^{-1} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{A}^{T} \boldsymbol{S}^{-2} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{A}^{T} \boldsymbol{S}^{-1} \mathbf{1}_{j} \\
& \leq \mathbf{1}_{j} \boldsymbol{S}^{-1} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{S}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{A}^{T} \boldsymbol{S}^{-1} \mathbf{1}_{j}=P_{j j}=\boldsymbol{\Psi}_{j j}
\end{aligned}
$$

Hence, the Gershgorin circle theorem shows that the eigenvalues of $\boldsymbol{\Psi}-\boldsymbol{P}^{(2)}$ are lies in union of the interval $\left[0,2 \psi_{j}\right]$ over all $j$. Hence, $\boldsymbol{\Psi}-\boldsymbol{P}^{(2)} \succeq \mathbf{0}$. On the other hand, Schur product theorem shows that $\boldsymbol{P}^{(2)} \succeq \mathbf{0}$ as $\boldsymbol{P} \succeq \mathbf{0}$. Hence, the result follows by chain rule.

Lemma 8.3.1 immediately shows that under our choice of $\boldsymbol{e}=\boldsymbol{e}_{P}(\boldsymbol{x}, \boldsymbol{\tau})$ we can compute the gradient of the hybrid barrier function, $p_{\boldsymbol{e}}(\boldsymbol{x})$ efficiently. Formally, Lemma 8.3.1 immediately implies the following:
Lemma 8.3.2 (Gradient). For $\boldsymbol{x} \in P=\left\{\boldsymbol{y} \in \mathbb{R}^{n}: \boldsymbol{A} \boldsymbol{y} \geq \boldsymbol{b}\right\}$ and $\boldsymbol{e} \in \mathbb{R}^{m}$ we have

$$
\nabla p_{e}(\boldsymbol{x})=-\boldsymbol{A}_{x}^{T}\left(c_{e} \mathbf{1}+\boldsymbol{e}+\boldsymbol{\psi}_{P}(\boldsymbol{x})\right)+\lambda \boldsymbol{x}
$$

and therefore for all $\boldsymbol{\tau} \in \mathbb{R}^{m}$, we have

$$
\nabla p_{e(\boldsymbol{\tau}, \boldsymbol{x})}(\boldsymbol{x})=-\boldsymbol{A}_{x}^{T}\left(c_{e} \mathbf{1}+\boldsymbol{\tau}\right)+\lambda \boldsymbol{x} .
$$

Remark 8.3.3. To be clear, the vector $\nabla p_{e(\tau, x)}(\boldsymbol{x})$ is defined as the vector such that

$$
\left[\nabla p_{e(\tau, x)}(\boldsymbol{x})\right]_{i}=\lim _{t \rightarrow 0} \frac{1}{t}\left(p_{\boldsymbol{e}(\boldsymbol{\tau}, \boldsymbol{x})}\left(\boldsymbol{x}+t \mathbf{1}_{i}\right)-p_{\boldsymbol{e}(\tau, \boldsymbol{x})}(\boldsymbol{x})\right) .
$$

In other words, we treat the parameter $\boldsymbol{e}(\boldsymbol{\tau}, \boldsymbol{x})$ as fixed. This is the reason we denote it by subscript to emphasize that $p_{e}(\boldsymbol{x})$ is a family of functions, $p_{e(\tau, x)}$ is one particular function, and $\nabla p_{e(\tau, x)}$ means taking gradient on that particular function.

Consequently, we can always compute $\nabla p_{\boldsymbol{e}(\boldsymbol{\tau}, \boldsymbol{x})}(\boldsymbol{x})$ efficiently. Now, we measure centrality or how close we are to the hybrid center as follows.

Definition 8.3.4 (Centrality). For $\boldsymbol{x} \in P=\left\{\boldsymbol{y} \in \mathbb{R}^{n}: \boldsymbol{A} \boldsymbol{y} \geq \boldsymbol{b}\right\}$ and $\boldsymbol{e} \in \mathbb{R}^{m}$, we define the centrality of $\boldsymbol{x}$ by

$$
\delta_{e}(\boldsymbol{x}) \stackrel{\text { def }}{=}\left\|\nabla p_{e}(\boldsymbol{x})\right\|_{\boldsymbol{H}(\boldsymbol{x})^{-1}}
$$

where $\boldsymbol{H}(\boldsymbol{x}) \stackrel{\text { def }}{=} \boldsymbol{A}_{x}^{T}\left(c_{e} \boldsymbol{I}+\boldsymbol{\Psi}(\boldsymbol{x})\right) \boldsymbol{A}_{x}+\lambda \boldsymbol{I}$. Often, we use weights $\boldsymbol{w} \in \mathbb{R}_{>0}^{m}$ to approximate this Hessian and consider $\boldsymbol{Q}(\boldsymbol{x}, \boldsymbol{w}) \stackrel{\text { def }}{=} \boldsymbol{A}_{x}^{T}\left(c_{e} \boldsymbol{I}+\boldsymbol{W}\right) \boldsymbol{A}_{x}+\lambda \boldsymbol{I}$.

Next, we bound how much slacks can change in a region close to a nearly central point.
Lemma 8.3.5. Let $\boldsymbol{x} \in P=\left\{\boldsymbol{y} \in \mathbb{R}^{n}: \boldsymbol{A} \boldsymbol{y} \geq \boldsymbol{b}\right\}$ and $\boldsymbol{y} \in \mathbb{R}^{n}$ such that $\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})} \leq \varepsilon \sqrt{c_{e}+\mu(\boldsymbol{x})}$ for $\varepsilon<1$. Then $\boldsymbol{y} \in P$ and $(1-\varepsilon) \boldsymbol{S}_{x} \preceq \boldsymbol{S}_{y} \preceq(1+\varepsilon) \boldsymbol{S}_{x}$.
Proof. Direct calculation reveals the following:

$$
\begin{aligned}
\left\|\boldsymbol{S}_{x}^{-1}\left(s_{\boldsymbol{y}}-\boldsymbol{s}_{x}\right)\right\|_{\infty} & \leq\left\|\boldsymbol{A}_{x}(\boldsymbol{y}-\boldsymbol{x})\right\|_{2} \leq \frac{1}{\sqrt{c_{e}+\mu(\boldsymbol{x})}}\left\|\boldsymbol{A}_{x}(\boldsymbol{y}-\boldsymbol{x})\right\|_{c_{e} \boldsymbol{I}+\boldsymbol{\Psi}(\boldsymbol{x})} \\
& \leq \frac{1}{\sqrt{c_{e}+\mu(\boldsymbol{x})}}\|\boldsymbol{y}-\boldsymbol{x}\|_{\boldsymbol{H}(\boldsymbol{x})} \leq \varepsilon
\end{aligned}
$$

Consequently, $(1-\varepsilon) \boldsymbol{S}_{x} \preceq \boldsymbol{S}_{y} \preceq(1+\varepsilon) \boldsymbol{S}_{x}$. Since $y \in P$ if and only if $\boldsymbol{S}_{y} \succeq \mathbf{0}$ the result follows.

Combining the previous lemmas we obtain the following.
Lemma 8.3.6. Let $\boldsymbol{x} \in P=\left\{\boldsymbol{y} \in \mathbb{R}^{n}: \boldsymbol{A} \boldsymbol{y} \geq \boldsymbol{b}\right\}$ and $\boldsymbol{e}, \boldsymbol{w} \in \mathbb{R}^{m}$ such that $\|\boldsymbol{e}\|_{\infty} \leq \frac{1}{2} c_{e} \leq 1$ and $\boldsymbol{\Psi}(\boldsymbol{x}) \preceq \boldsymbol{W} \preceq \frac{4}{3} \boldsymbol{\Psi}(\boldsymbol{x})$. If $\boldsymbol{y} \in \mathbb{R}^{n}$ satisfies $\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{Q}(\boldsymbol{x}, \boldsymbol{w})} \leq \frac{1}{10} \sqrt{c_{e}+\mu(\boldsymbol{x})}$, then

$$
\frac{1}{4} \boldsymbol{Q}(\boldsymbol{x}, \boldsymbol{w}) \preceq \nabla^{2} p_{e}(\boldsymbol{y}) \preceq 8 \boldsymbol{Q}(\boldsymbol{x}, \boldsymbol{w}) \quad \text { and } \quad \frac{1}{2} \boldsymbol{H}(\boldsymbol{x}) \preceq \boldsymbol{H}(\boldsymbol{y}) \preceq 2 \boldsymbol{H}(\boldsymbol{x}) .
$$

Also, we have that $\boldsymbol{H}(\boldsymbol{x}) \preceq \boldsymbol{Q}(\boldsymbol{x}, \boldsymbol{w}) \preceq \frac{4}{3} \boldsymbol{H}(\boldsymbol{x})$.
Proof. Lemma 8.3.1 shows that

$$
\begin{equation*}
\boldsymbol{A}_{y}^{T}\left(c_{e} \boldsymbol{I}+\boldsymbol{E}+\boldsymbol{\Psi}(\boldsymbol{y})\right) \boldsymbol{A}_{y}+\lambda \boldsymbol{I} \preceq \nabla^{2} p_{e}(\boldsymbol{y}) \preceq \boldsymbol{A}_{y}^{T}\left(c_{e} \boldsymbol{I}+\boldsymbol{E}+3 \boldsymbol{\Psi}(\boldsymbol{y})\right) \boldsymbol{A}_{y}+\lambda \boldsymbol{I} . \tag{8.3}
\end{equation*}
$$

Since $\boldsymbol{W} \succeq \boldsymbol{\Psi}$, we have that $\boldsymbol{Q}(\boldsymbol{x}, \boldsymbol{w}) \succeq \boldsymbol{H}(\boldsymbol{x})$ and therefore $\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})} \leq \varepsilon \sqrt{c_{e}+\mu(\boldsymbol{x})}$ with $\varepsilon=0.1$. Consequently, by Lemma 8.3.5 we have $(1-\varepsilon) \boldsymbol{S}_{x} \preceq \boldsymbol{S}_{y} \preceq(1+\varepsilon) \boldsymbol{S}_{x}$ and therefore

$$
\frac{(1-\varepsilon)^{2}}{(1+\varepsilon)^{2}} \boldsymbol{\Psi}(\boldsymbol{x}) \preceq \boldsymbol{\Psi}(\boldsymbol{y}) \preceq \frac{(1+\varepsilon)^{2}}{(1-\varepsilon)^{2}} \boldsymbol{\Psi}(\boldsymbol{x})
$$

and

$$
\frac{1}{2} \boldsymbol{H}(\boldsymbol{x}) \preceq \frac{(1-\varepsilon)^{2}}{(1+\varepsilon)^{4}} \boldsymbol{H}(\boldsymbol{x}) \preceq \boldsymbol{H}(\boldsymbol{y}) \preceq \frac{(1+\varepsilon)^{2}}{(1-\varepsilon)^{4}} \boldsymbol{H}(\boldsymbol{x}) \preceq 2 \boldsymbol{H}(\boldsymbol{x})
$$

Furthermore, (8.3) shows that

$$
\begin{aligned}
\nabla^{2} p_{e}(\boldsymbol{y}) & \preceq \boldsymbol{A}_{y}^{T}\left(c_{e} \boldsymbol{I}+\boldsymbol{E}+3 \boldsymbol{\Psi}(\boldsymbol{y})\right) \boldsymbol{A}_{y}+\lambda \boldsymbol{I} \\
& \preceq \frac{(1+\varepsilon)^{2}}{(1-\varepsilon)^{4}} \boldsymbol{A}_{x}^{T}\left(c_{e} \boldsymbol{I}+\boldsymbol{E}+3 \boldsymbol{\Psi}(\boldsymbol{x})\right) \boldsymbol{A}_{x}+\lambda \boldsymbol{I} \\
& \preceq 2 \boldsymbol{A}_{x}^{T}\left(\frac{3}{2} c_{e} \boldsymbol{I}+3 \boldsymbol{W}\right) \boldsymbol{A}_{x}+\lambda \boldsymbol{I} \preceq 8 \boldsymbol{Q}(\boldsymbol{x}, \boldsymbol{w})
\end{aligned}
$$

and

$$
\begin{aligned}
\nabla^{2} p_{\boldsymbol{e}}(\boldsymbol{y}) & \succeq \boldsymbol{A}_{y}^{T}\left(c_{e} \boldsymbol{I}+\boldsymbol{E}+\boldsymbol{\Psi}(\boldsymbol{y})\right) \boldsymbol{A}_{y}+\lambda \boldsymbol{I} \\
& \succeq \frac{(1-\varepsilon)^{4}}{(1+\varepsilon)^{2}} \boldsymbol{A}_{x}^{T}\left(c_{e} \boldsymbol{I}+\boldsymbol{E}+\boldsymbol{\Psi}(\boldsymbol{x})\right) \boldsymbol{A}_{x}+\lambda \boldsymbol{I} \\
& \succeq \frac{1}{2} \boldsymbol{A}_{x}^{T}\left(\frac{1}{2} c_{e} \boldsymbol{I}+\frac{3}{4} \boldsymbol{W}\right) \boldsymbol{A}_{x}+\lambda \boldsymbol{I} \succeq \frac{1}{4} \boldsymbol{Q}(\boldsymbol{x}, \boldsymbol{w}) .
\end{aligned}
$$

The last inequality follows from the definition of $\boldsymbol{H}(\boldsymbol{x})$ and $\boldsymbol{Q}(\boldsymbol{x}, \boldsymbol{w})$ and the fact $\boldsymbol{\Psi}(\boldsymbol{x}) \preceq \boldsymbol{W} \preceq$ ${ }_{3}^{4} \boldsymbol{\Psi}(\boldsymbol{x})$.

To analyze our centering scheme we use standard facts about gradient descent we prove in Lemma 8.3.7.
Lemma 8.3.7 (Gradient Descent). Let $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ be twice differentiable and $\boldsymbol{Q} \in \mathbb{R}^{n \times n}$ be positive definite. Let $\boldsymbol{x}_{0} \in \mathbb{R}^{n}$ and $\boldsymbol{x}_{1} \stackrel{\text { def }}{=} \boldsymbol{x}_{0}-\frac{1}{L} \boldsymbol{Q}^{-1} \nabla f\left(\boldsymbol{x}_{0}\right)$. Furthermore, let $\boldsymbol{x}_{\alpha}=\boldsymbol{x}_{0}+\alpha\left(\boldsymbol{x}_{1}-\boldsymbol{x}\right)$ and suppose that $\mu \boldsymbol{Q} \preceq \nabla^{2} f\left(\boldsymbol{x}_{\alpha}\right) \preceq L \boldsymbol{Q}$ for all $\alpha \in[0,1]$. Then,

1. $\left\|\nabla f\left(\boldsymbol{x}_{1}\right)\right\|_{Q^{-1}} \leq\left(1-\frac{\mu}{L}\right)\left\|\nabla f\left(\boldsymbol{x}_{0}\right)\right\|_{Q^{-1}}$
2. $f\left(\boldsymbol{x}_{1}\right) \geq f\left(\boldsymbol{x}_{0}\right)-\frac{1}{L}\left\|\nabla f\left(\boldsymbol{x}_{0}\right)\right\|_{Q^{-1}}^{2}$

Proof. Integrating we have that

$$
\nabla f\left(\boldsymbol{x}_{1}\right)=\nabla f\left(\boldsymbol{x}_{0}\right)+\int_{0}^{1} \nabla^{2} f\left(\boldsymbol{x}_{\alpha}\right)\left(\boldsymbol{x}_{1}-\boldsymbol{x}_{0}\right) d \alpha=\int_{0}^{1}\left(\boldsymbol{Q}-\frac{1}{L} \nabla^{2} f\left(\boldsymbol{x}_{\alpha}\right)\right) \boldsymbol{Q}^{-1} \nabla f\left(\boldsymbol{x}_{0}\right) d \alpha
$$

Consequently, by applying Jensen's inequality we have

$$
\begin{aligned}
\left\|\nabla f\left(\boldsymbol{x}_{1}\right)\right\|_{\boldsymbol{Q}^{-1}} & =\left\|\int_{0}^{1}\left(\boldsymbol{Q}-\frac{1}{L} \nabla^{2} f\left(\boldsymbol{x}_{\alpha}\right)\right) \boldsymbol{Q}^{-1} \nabla f\left(\boldsymbol{x}_{0}\right) d \alpha\right\|_{\boldsymbol{Q}^{-1}} \\
& \leq \int_{0}^{1}\left\|\left(\boldsymbol{Q}-\frac{1}{L} \nabla^{2} f\left(\boldsymbol{x}_{\alpha}\right)\right) \boldsymbol{Q}^{-1} \nabla f\left(\boldsymbol{x}_{0}\right)\right\|_{\boldsymbol{Q}^{-1}} d \alpha \\
& \leq\left\|\boldsymbol{Q}^{-1 / 2} \nabla f\left(\boldsymbol{x}_{0}\right)\right\|_{\left[\boldsymbol{Q}^{-1 / 2}\left(\boldsymbol{Q}-\frac{1}{L} \nabla^{2} f\left(\boldsymbol{x}_{\alpha}\right)\right) \boldsymbol{Q}^{-1 / 2}\right]^{2}}
\end{aligned}
$$

Now we know that by assumption that

$$
\mathbf{0} \preceq \boldsymbol{Q}^{-1 / 2}\left(\boldsymbol{Q}-\frac{1}{L} \nabla^{2} f\left(\boldsymbol{x}_{\alpha}\right)\right) \boldsymbol{Q}^{-1 / 2} \preceq\left(1-\frac{\mu}{L}\right) \boldsymbol{I}
$$

and therefore combining these (1) holds.
Using the convexity of $f$, we have

$$
\begin{aligned}
f\left(\boldsymbol{x}_{1}\right) & \geq f\left(\boldsymbol{x}_{0}\right)+\left\langle\nabla f\left(\boldsymbol{x}_{0}\right), \boldsymbol{x}_{1}-\boldsymbol{x}_{0}\right\rangle \\
& \geq f\left(\boldsymbol{x}_{0}\right)-\left\|\nabla f\left(\boldsymbol{x}_{0}\right)\right\|_{Q^{-1}}\left\|\boldsymbol{x}_{1}-\boldsymbol{x}_{0}\right\|_{\boldsymbol{Q}}
\end{aligned}
$$

and since $\left\|\boldsymbol{x}_{1}-\boldsymbol{x}_{0}\right\|_{\boldsymbol{Q}}=\frac{1}{L}\left\|\nabla f\left(\boldsymbol{x}_{0}\right)\right\|_{Q^{-1}},(2)$ holds as well.
Next we bound the effect of changing $\boldsymbol{e}$ on the hybrid barrier function $p_{e}(\boldsymbol{x})$.
Lemma 8.3.8. For $\boldsymbol{x} \in P=\left\{\boldsymbol{y} \in \mathbb{R}^{n}: \boldsymbol{A} \boldsymbol{y} \geq \boldsymbol{b}\right\}, \boldsymbol{e}, \boldsymbol{f} \in \mathbb{R}^{m}$, and $\boldsymbol{w} \in \mathbb{R}_{>0}^{m}$ such that $\boldsymbol{W} \succeq \boldsymbol{\Psi}(\boldsymbol{x})$

$$
\left\|\nabla p_{\boldsymbol{f}}(\boldsymbol{x})\right\|_{\boldsymbol{Q}(\boldsymbol{x}, \boldsymbol{w})^{-1}} \leq\left\|\nabla p_{\boldsymbol{e}}(\boldsymbol{x})\right\|_{\boldsymbol{Q}(\boldsymbol{x}, \boldsymbol{w})^{-1}}+\frac{1}{\sqrt{c_{e}+\mu(\boldsymbol{x})}}\|\boldsymbol{f}-\boldsymbol{e}\|_{2}
$$

Proof. Direct calculation shows the following

$$
\begin{array}{rlr}
\left\|\nabla p_{\boldsymbol{f}}(\boldsymbol{x})\right\|_{\boldsymbol{Q}(\boldsymbol{x}, \boldsymbol{w})^{-1}} & =\left\|-\boldsymbol{A}_{x}^{T}\left(c_{e} \mathbf{1}+\boldsymbol{f}+\boldsymbol{\psi}_{P}(\boldsymbol{x})\right)+\lambda \boldsymbol{x}\right\|_{\boldsymbol{Q}(\boldsymbol{x}, \boldsymbol{w})^{-1}} & \text { (Formula for } \left.\nabla p_{\boldsymbol{f}}(\boldsymbol{x})\right) \\
& \leq\left\|\nabla p_{e}(\boldsymbol{x})\right\|_{\boldsymbol{Q}(\boldsymbol{x}, \boldsymbol{w})^{-1}}+\left\|\boldsymbol{A}_{x}^{T}(\boldsymbol{f}-\boldsymbol{e})\right\|_{\boldsymbol{Q}(\boldsymbol{x}, \boldsymbol{w})^{-1}} & \text { (Triangle inequality) } \\
& \leq\left\|\nabla p_{\boldsymbol{e}}(\boldsymbol{x})\right\|_{\boldsymbol{Q}(\boldsymbol{x}, \boldsymbol{w})^{-1}}+\frac{1}{\sqrt{c_{e}+\mu(\boldsymbol{x})}}\left\|\boldsymbol{A}_{x}^{T}(\boldsymbol{f}-\boldsymbol{e})\right\|_{\left(\boldsymbol{A}_{x}^{T} \boldsymbol{A}_{x}\right)^{-1}} & \text { (Bound on } \boldsymbol{Q}(\boldsymbol{x}, \boldsymbol{w})) \\
& \leq\left\|\nabla p_{\boldsymbol{e}}(\boldsymbol{x})\right\|_{\boldsymbol{Q}(\boldsymbol{x}, \boldsymbol{w})^{-1}}+\frac{1}{\sqrt{c_{e}+\mu(\boldsymbol{x})}}\|\boldsymbol{f}-\boldsymbol{e}\|_{2} & \text { (Property of projection matrix) }
\end{array}
$$

where in the second to third line we used $\boldsymbol{Q}(\boldsymbol{x}, \boldsymbol{w}) \succeq \boldsymbol{H}(\boldsymbol{x}) \succeq\left(c_{e}+\mu(\boldsymbol{x})\right) \boldsymbol{A}_{x}^{T} \boldsymbol{A}_{x}$.
We now have everything we need to analyze our centering algorithm.
Lemma 8.3.9. Let $\boldsymbol{x}^{(0)} \in P=\left\{\boldsymbol{y} \in \mathbb{R}^{n}: \boldsymbol{A} \boldsymbol{y} \geq \boldsymbol{b}\right\}$ and let $\boldsymbol{\tau}^{(0)} \in \mathbb{R}^{m}$ such that $\left\|\boldsymbol{e}\left(\boldsymbol{\tau}^{(0)}, \boldsymbol{x}^{(0)}\right)\right\|_{\infty} \leq$ $\frac{1}{3} c_{e} \leq \frac{1}{3}$. Assume that $r$ is a positive integer, $0 \leq c_{\Delta} \leq 0.01 c_{e}$ and $\delta_{\boldsymbol{e}^{(0)}}\left(\boldsymbol{x}^{(0)}\right) \leq \frac{1}{100} \sqrt{c_{e}+\mu\left(\boldsymbol{x}^{(0)}\right)}$. With high probability in $n$, the algorithm Centering $\left(\boldsymbol{x}^{(0)}, \boldsymbol{\tau}^{(0)}, r, c_{\Delta}\right)$ outputs $\left(\boldsymbol{x}^{(r)}, \boldsymbol{\tau}^{(r)}\right)$ such that

1. $\delta_{\boldsymbol{e}^{(r)}}\left(\boldsymbol{x}^{(r)}\right) \leq 2\left(1-\frac{1}{64}\right)^{r} \delta_{\boldsymbol{e}^{(0)}}\left(\boldsymbol{x}^{(0)}\right)$.
2. $\mathbf{E}\left[p_{\boldsymbol{e}^{(k)}}\left(\boldsymbol{x}^{(r)}\right)\right] \geq p_{\boldsymbol{e}^{(0)}}\left(\boldsymbol{x}^{(0)}\right)-8\left(\delta_{\boldsymbol{e}^{(0)}}\left(\boldsymbol{x}^{(0)}\right)\right)^{2}$.
```
Algorithm 21: \(\left(\boldsymbol{x}^{(r)}, \boldsymbol{\tau}^{(r)}\right)=\operatorname{Centering}\left(\boldsymbol{x}^{(0)}, \boldsymbol{\tau}^{(0)}, r, c_{\Delta}\right)\)
    Input: Initial point \(\boldsymbol{x}^{(0)} \in P=\left\{\boldsymbol{y} \in \mathbb{R}^{n}: \boldsymbol{A} \boldsymbol{y} \geq \boldsymbol{b}\right\}\), Estimator of leverage scores \(\boldsymbol{\tau}^{(0)} \in \mathbb{R}^{n}\)
    Input: Number of iterations \(r>0\), Accuracy of the estimator \(0 \leq c_{\Delta} \leq 0.01 c_{e}\).
    Given: \(\left\|\boldsymbol{e}^{(0)}\right\|_{\infty} \leq \frac{1}{3} c_{e} \leq \frac{1}{3}\) where \(\boldsymbol{e}^{(0)}=\boldsymbol{e}\left(\boldsymbol{\tau}^{(0)}, \boldsymbol{x}^{(0)}\right)\).
    Given: \(\delta_{\boldsymbol{e}^{(0)}}\left(\boldsymbol{x}^{(0)}\right)=\left\|\nabla p_{e^{(0)}}\left(\boldsymbol{x}^{(0)}\right)\right\|_{\boldsymbol{H}\left(\boldsymbol{x}^{(0)}\right)^{-1}} \leq \frac{1}{100} \sqrt{c_{e}+\mu\left(\boldsymbol{x}^{(0)}\right)}\).
    Compute \(\boldsymbol{w}\) such that \(\boldsymbol{\Psi}\left(\boldsymbol{x}^{(0)}\right) \preceq \boldsymbol{W} \preceq \frac{4}{3} \boldsymbol{\Psi}\left(\boldsymbol{x}^{(0)}\right)\) (Similar to Lemma 2.3.4)
    Let \(\boldsymbol{Q} \stackrel{\text { def }}{=} \boldsymbol{Q}\left(\boldsymbol{x}^{(0)}, \boldsymbol{w}\right)\).
    for \(k=1\) to \(r\) do
        \(\boldsymbol{x}^{(k)}:=\boldsymbol{x}^{(k-1)}-\frac{1}{8} \boldsymbol{Q}^{-1} \nabla p_{\boldsymbol{e}^{(k-1)}}\left(\boldsymbol{x}^{(k-1)}\right)\).
        Sample \(\boldsymbol{\Delta}^{(k)} \in \mathbb{R}^{n}\) s.t.
            \(\mathbf{E}\left[\boldsymbol{\Delta}^{(k)}\right]=\boldsymbol{\psi}\left(\boldsymbol{x}^{(k)}\right)-\boldsymbol{\psi}\left(\boldsymbol{x}^{(k-1)}\right)\) and
            with high probability in \(n\),
            \(\left\|\boldsymbol{\Delta}^{(k)}-\left(\boldsymbol{\psi}\left(\boldsymbol{x}^{(k)}\right)-\boldsymbol{\psi}\left(\boldsymbol{x}^{(k-1)}\right)\right)\right\|_{2} \leq c_{\Delta}\left\|\boldsymbol{S}_{\boldsymbol{x}^{(k-1)}}^{-1}\left(\boldsymbol{s}_{\boldsymbol{x}^{(k)}}-\boldsymbol{s}_{\boldsymbol{x}^{(k-1)}}\right)\right\|_{2}\) (See Section 8.4.1)
        \(\boldsymbol{\tau}^{(k)}:=\boldsymbol{\tau}^{(k-1)}+\boldsymbol{\Delta}^{(k)}\).
        \(\boldsymbol{e}^{(k)}:=\boldsymbol{e}\left(\boldsymbol{\tau}^{(k)}, \boldsymbol{x}^{(k)}\right)\).
    end
    Output: \(\left(\boldsymbol{x}^{(r)}, \boldsymbol{\tau}^{(r)}\right)\)
```

3. $\mathbf{E} \boldsymbol{e}^{(r)}=\boldsymbol{e}^{(0)}$ and $\left\|\boldsymbol{e}^{(r)}-\boldsymbol{e}^{(0)}\right\|_{2} \leq \frac{1}{10} c_{\Delta}$.
4. $\left\|\boldsymbol{S}_{\boldsymbol{x}^{(0)}}^{-1}\left(\boldsymbol{s}\left(\boldsymbol{x}^{(r)}\right)-\boldsymbol{s}\left(\boldsymbol{x}^{(0)}\right)\right)\right\|_{2} \leq \frac{1}{10}$.
where $\boldsymbol{e}^{(r)}=\boldsymbol{e}\left(\boldsymbol{\tau}^{(r)}, \boldsymbol{x}^{(r)}\right)$.

Proof. Let $\eta=\left\|\nabla p_{\boldsymbol{e}^{(0)}}\left(\boldsymbol{x}^{(0)}\right)\right\|_{Q^{-1}}$. First, we use induction to prove that $\left\|\boldsymbol{x}^{(r)}-\boldsymbol{x}^{(0)}\right\|_{Q} \leq 8 \eta$, $\left\|\nabla p_{\boldsymbol{e}^{(r)}}\left(\boldsymbol{x}^{(r)}\right)\right\|_{Q^{-1}} \leq\left(1-\frac{1}{64}\right)^{r} \eta$ and $\left\|\boldsymbol{e}^{(r)}-\boldsymbol{e}^{(0)}\right\|_{2} \leq \frac{1}{10} c_{\Delta}$ for all $r$.

Clearly the claims hold for $r=0$. We now suppose they hold for all $r \leq t$ and show that they hold for $r=t+1$. Now, since $\left\|\boldsymbol{x}^{(t)}-\boldsymbol{x}^{(0)}\right\|_{Q} \leq 8 \eta, \boldsymbol{x}^{(t+1)}=\boldsymbol{x}^{(t)}-\frac{1}{8} \boldsymbol{Q}^{-1} \nabla p_{e^{(t)}}\left(\boldsymbol{x}^{(t)}\right)$, and $\left\|\nabla p_{\boldsymbol{e}^{(t)}}\left(\boldsymbol{x}^{(t)}\right)\right\|_{Q^{-1}} \leq\left(1-\frac{1}{64}\right)^{t} \eta \leq \eta$, we have

$$
\left\|\boldsymbol{x}^{(t+1)}-\boldsymbol{x}^{(0)}\right\|_{Q} \leq\left\|\boldsymbol{x}^{(t)}-\boldsymbol{x}^{(0)}\right\|_{Q}+\frac{1}{8}\left\|\nabla p_{\boldsymbol{e}^{(t)}}\left(\boldsymbol{x}^{(t)}\right)\right\|_{Q^{-1}} \leq 9 \eta .
$$

We will improve this estimate later in the proof to finish the induction on $\left\|\boldsymbol{x}^{(t+1)}-\boldsymbol{x}^{(0)}\right\|_{\boldsymbol{Q}}$, but using this, $\eta \leq 0.01 \sqrt{c_{e}+\mu\left(\boldsymbol{x}^{(0)}\right)}$, and $\left\|\boldsymbol{e}^{(t)}\right\|_{\infty} \leq\left\|\boldsymbol{e}^{(t)}-\boldsymbol{e}^{(0)}\right\|_{\infty}+\left\|\boldsymbol{e}^{(0)}\right\|_{\infty} \leq \frac{c_{e}}{2}$, we can invoke Lemma 8.3.6 and Lemma 8.3.7 and therefore

$$
\left\|\nabla p_{\boldsymbol{e}^{(t)}}\left(\boldsymbol{x}^{(t+1)}\right)\right\|_{Q^{-1}} \leq\left(1-\frac{1}{32}\right)\left\|\nabla p_{\boldsymbol{e}^{(t)}}\left(\boldsymbol{x}^{(t)}\right)\right\|_{Q^{-1}}
$$

By Lemma 8.3.8 we have

$$
\begin{equation*}
\left\|\nabla p_{\boldsymbol{e}^{(t+1)}}\left(\boldsymbol{x}^{(t+1)}\right)\right\|_{Q^{-1}} \leq\left(1-\frac{1}{32}\right)\left\|\nabla p_{\boldsymbol{e}^{(t)}}\left(\boldsymbol{x}^{(t)}\right)\right\|_{Q^{-1}}+\frac{1}{\sqrt{c_{e}+\mu\left(\boldsymbol{x}^{(0)}\right)}}\left\|\boldsymbol{e}^{(t+1)}-\boldsymbol{e}^{(t)}\right\|_{2} . \tag{8.4}
\end{equation*}
$$

To bound $\left\|\boldsymbol{e}^{(t+1)}-\boldsymbol{e}^{(t)}\right\|_{2}$, we use the definition of $\boldsymbol{\Delta}$ to shows that

$$
\begin{aligned}
\left\|\boldsymbol{e}^{(t+1)}-\boldsymbol{e}^{(t)}\right\|_{2} & =\left\|\left(\boldsymbol{\tau}^{(t+1)}-\boldsymbol{\psi}\left(\boldsymbol{x}^{(t+1)}\right)\right)-\left(\boldsymbol{\tau}^{(t)}-\boldsymbol{\psi}\left(\boldsymbol{x}^{(t)}\right)\right)\right\|_{2} \\
& =\left\|\boldsymbol{\Delta}^{(t+1)}-\left(\boldsymbol{\psi}\left(\boldsymbol{x}^{(t+1)}\right)-\boldsymbol{\psi}\left(\boldsymbol{x}^{(t)}\right)\right)\right\|_{2} \\
& \leq c_{\Delta}\left\|\boldsymbol{S}_{x^{(t)}}^{-1}\left(\boldsymbol{s}_{x^{(t+1)}}-\boldsymbol{s}_{\left.x^{(t)}\right)}\right)\right\|_{2}
\end{aligned}
$$

with high probability in $n$. Now, we note that Lemma 8.3 .5 and the induction hypothesis $\| \boldsymbol{x}^{(t)}-$ $\boldsymbol{x}^{(0)}\left\|_{\boldsymbol{H}\left(\boldsymbol{x}^{(0)}\right)} \leq\right\| \boldsymbol{x}^{(t)}-\boldsymbol{x}^{(0)} \|_{Q} \leq 8 \eta$ shows that $(1-0.1) \boldsymbol{S}_{x^{(0)}} \preceq \boldsymbol{S}_{x^{(t)}} \preceq(1+0.1) \boldsymbol{S}_{x^{(0)}}$ and therefore

$$
\begin{align*}
\left\|\boldsymbol{e}^{(t+1)}-\boldsymbol{e}^{(t)}\right\|_{2} & \leq c_{\Delta}\left\|\boldsymbol{S}_{x^{(t)}}^{-1}\left(\boldsymbol{s}_{x^{(t)}}-\boldsymbol{s}_{x^{(t+1)}}\right)\right\|_{2} \\
& \leq \frac{c_{\Delta}}{1-0.1}\left\|\boldsymbol{S}_{x^{(0)}}^{-1} \boldsymbol{A}\left(\boldsymbol{x}^{(t)}-\boldsymbol{x}^{(t+1)}\right)\right\|_{2} \\
& =\frac{c_{\Delta}}{1-0.1}\left\|\frac{1}{8} \boldsymbol{Q}^{-1} \nabla p_{\boldsymbol{e}^{(t)}}\left(\boldsymbol{x}^{(t)}\right)\right\|_{\boldsymbol{A}^{T} \boldsymbol{S}_{x^{(0)}}^{-2} \boldsymbol{A}} \\
& \leq \frac{c_{\Delta}}{8(1-0.1) \sqrt{c_{e}+\mu\left(\boldsymbol{x}^{(0)}\right)}\left\|\nabla p_{\boldsymbol{e}^{(t)}}\left(\boldsymbol{x}^{(t)}\right)\right\|_{\boldsymbol{Q}^{-1}}} \tag{8.5}
\end{align*}
$$

where in the last line we used $\min _{i \in[m]} w_{i} \geq \mu\left(\boldsymbol{x}^{(0)}\right)$. Since $c_{\Delta}<0.01 c_{e} \leq 0.01$, by (8.4), we have

$$
\begin{aligned}
\left\|\nabla p_{\boldsymbol{e}^{(t+1)}}\left(\boldsymbol{x}^{(t+1)}\right)\right\|_{\boldsymbol{Q}^{-1}} & \leq\left(1-\frac{1}{32}\right)\left\|\nabla p_{\boldsymbol{e}^{(t)}}\left(\boldsymbol{x}^{(t)}\right)\right\|_{Q^{-1}}+\frac{0.01}{8(1-0.1)}\left\|\nabla p_{\boldsymbol{e}^{(t)}}\left(\boldsymbol{x}^{(t)}\right)\right\|_{Q^{-1}} \\
& \leq\left(1-\frac{1}{64}\right)\left\|\nabla p_{\boldsymbol{e}^{(t)}}\left(\boldsymbol{x}^{(t)}\right)\right\|_{Q^{-1}}
\end{aligned}
$$

Furthermore, this implies that

$$
\left\|\boldsymbol{x}^{(t+1)}-\boldsymbol{x}^{(0)}\right\|_{\boldsymbol{Q}} \leq\left\|\sum_{k=0}^{t} \frac{1}{8} \boldsymbol{Q}^{-1} \nabla p_{\boldsymbol{e}^{(k)}}\left(\boldsymbol{x}^{(k)}\right)\right\|_{Q^{-1}} \leq \frac{1}{8} \sum_{i=0}^{\infty}\left(1-\frac{1}{64}\right)^{k} \eta \leq \frac{64}{8} \eta=8 \eta .
$$

Similarly, we have that

$$
\begin{aligned}
\left\|\boldsymbol{e}^{(t+1)}-\boldsymbol{e}^{(0)}\right\|_{2} & \leq \sum_{k=0}^{t} \frac{c_{\Delta}}{8(1-0.1) \sqrt{c_{e}+\mu\left(\boldsymbol{x}^{(0)}\right)}}\left(1-\frac{1}{64}\right)^{k}\left\|\nabla p_{\boldsymbol{e}^{(0)}}\left(\boldsymbol{x}^{(0)}\right)\right\|_{Q^{-1}} \\
& \leq \frac{8 c_{\Delta \eta}}{(1-0.1) \sqrt{c_{e}+\mu\left(\boldsymbol{x}^{(0)}\right)}} \leq \frac{8 c_{\Delta}}{(1-0.1) \sqrt{c_{e}+\mu\left(\boldsymbol{x}^{(0)}\right)}} \delta_{\boldsymbol{e}^{(0)}}\left(\boldsymbol{x}^{(0)}\right) \leq \frac{1}{10} c_{\Delta}
\end{aligned}
$$

where we used $\eta=\left\|\nabla p_{\boldsymbol{e}^{(0)}}\left(\boldsymbol{x}^{(0)}\right)\right\|_{Q^{-1}} \leq\left\|\nabla p_{\boldsymbol{e}^{(0)}}\left(\boldsymbol{x}^{(0)}\right)\right\|_{\boldsymbol{H}^{-1}}=\delta_{\boldsymbol{e}^{(0)}}\left(\boldsymbol{x}^{(0)}\right)$ and this finishes the induction on $\left\|\nabla p_{\boldsymbol{e}^{(t)}}\left(\boldsymbol{x}^{(t)}\right)\right\|_{Q^{-1}},\left\|\boldsymbol{x}^{(t)}-\boldsymbol{x}^{(0)}\right\|_{Q}$ and $\left\|\boldsymbol{e}^{(t)}-\boldsymbol{e}^{(0)}\right\|_{2}$.

Hence, for all $r$, Lemma 8.3.6 shows that

$$
\begin{aligned}
\delta_{\boldsymbol{e}^{(r)}}\left(\boldsymbol{x}^{(r)}\right) & =\left\|\nabla p_{\boldsymbol{e}^{(r)}}\left(\boldsymbol{x}^{(r)}\right)\right\|_{\boldsymbol{H}\left(\boldsymbol{x}^{(r)}\right)^{-1}} \leq \sqrt{2}\left\|\nabla p_{\boldsymbol{e}^{(r)}}\left(\boldsymbol{x}^{(r)}\right)\right\|_{\boldsymbol{H}\left(\boldsymbol{x}^{(0)}\right)^{-1}} \\
& \leq \sqrt{\frac{8}{3}}\left\|\nabla p_{\boldsymbol{e}^{(r)}}\left(\boldsymbol{x}^{(r)}\right)\right\|_{Q^{-1}} \leq \sqrt{\frac{8}{3}}\left(1-\frac{1}{64}\right)^{r}\left\|\nabla p_{\boldsymbol{e}^{(0)}}\left(\boldsymbol{x}^{(0)}\right)\right\|_{Q^{-1}} \\
& \leq 2\left(1-\frac{1}{64}\right)^{r} \delta_{\boldsymbol{e}^{(0)}}\left(\boldsymbol{x}^{(0)}\right) .
\end{aligned}
$$

Using that $\mathbf{E} \boldsymbol{e}^{(t+1)}=\boldsymbol{e}^{(t)}$, we see that the expected change in function value is only due to the change
while taking centering steps and therefore Lemma 8.3.7 shows that

$$
\mathbf{E}\left[p_{\boldsymbol{e}^{(r)}}\left(\boldsymbol{x}^{(r)}\right)\right] \geq p_{\boldsymbol{e}^{(0)}}\left(\boldsymbol{x}^{(0)}\right)-\frac{1}{8} \sum_{k=0}^{\infty}\left(1-\frac{1}{64}\right)^{2 k}\left\|\nabla p_{\boldsymbol{e}^{(0)}}\left(\boldsymbol{x}^{(0)}\right)\right\|_{Q^{-1}}^{2} \geq p_{\boldsymbol{e}^{(0)}}\left(\boldsymbol{x}^{(0)}\right)-8\left(\delta_{\boldsymbol{e}^{(0)}}\left(\boldsymbol{x}^{(0)}\right)\right)^{2} .
$$

Finally, for (4), we note that

$$
\left\|\frac{s\left(\boldsymbol{x}^{(r)}\right)-s\left(\boldsymbol{x}^{(0)}\right)}{s\left(\boldsymbol{x}^{(0)}\right)}\right\|_{2}=\left\|\boldsymbol{x}^{(r)}-\boldsymbol{x}^{(0)}\right\|_{\boldsymbol{A}^{T} S_{x^{(0)}}^{-2} \boldsymbol{A}} \leq \frac{1}{\sqrt{\mu\left(\boldsymbol{x}^{(0)}\right)+c_{e}}}\left\|\boldsymbol{x}^{(r)}-\boldsymbol{x}^{(0)}\right\|_{Q^{-1}} \leq \frac{1}{10} .
$$

## ■ 8.3.2 Changing Constraints

Here we bound the effect that adding or a removing a constraint has on the hybrid barrier function. Much of the analysis in this section follows from the following lemma which follows easily from the Sherman Morrison Formula.

Lemma 8.3.10 (Sherman Morrison Formula Implications). Let $\boldsymbol{B} \in \mathbb{R}^{n \times n}$ be an invertible symmetric matrix and let $\boldsymbol{a} \in \mathbb{R}^{n}$ be arbitrary vector satisfying $\boldsymbol{a}^{T} \boldsymbol{B}^{-1} \boldsymbol{a}<1$. The following hold:

1. $\left(\boldsymbol{B} \pm \boldsymbol{a} \boldsymbol{a}^{T}\right)^{-1}=\boldsymbol{B}^{-1} \mp \frac{\boldsymbol{B}^{-1} \boldsymbol{a} \boldsymbol{a}^{T} B^{-1}}{1 \pm \boldsymbol{a}^{T} B^{-1} \boldsymbol{a}}$.
2. $0 \preceq \frac{B^{-1} a a^{T} B^{-1}}{1 \pm a^{T} B^{-1} a} \preceq \frac{a^{T} B^{-1} a}{1 \pm a^{T} B^{-1} a} B^{-1}$.
3. $\log \operatorname{det}\left(\boldsymbol{B} \pm \boldsymbol{a} \boldsymbol{a}^{T}\right)=\ln \operatorname{det} \boldsymbol{B}+\ln \left(1 \pm \boldsymbol{a}^{T} \boldsymbol{B}^{-1} \boldsymbol{a}\right)$.

Proof. (1) follows immediately from Sherman Morrison Formula (Lemma 2.3.7). (2) follows since $\boldsymbol{a}^{T}$ is PSD,

$$
\frac{\boldsymbol{B}^{-1} \boldsymbol{a} \boldsymbol{a}^{T} \boldsymbol{B}^{-1}}{1 \pm \boldsymbol{a}^{T} \boldsymbol{B}^{-1} \boldsymbol{a}}=\boldsymbol{B}^{-1 / 2}\left[\frac{\boldsymbol{B}^{-1 / 2} \boldsymbol{a} \boldsymbol{a}^{T} \boldsymbol{B}^{-1 / 2}}{1 \pm \boldsymbol{a}^{T} \boldsymbol{B}^{-1} \boldsymbol{a}}\right] \boldsymbol{B}^{-1 / 2}
$$

and $\boldsymbol{y} \boldsymbol{y}^{T} \preceq\|\boldsymbol{y}\|_{2}^{2} \boldsymbol{I}$ for any vector $\boldsymbol{y}$. (3) follows immediately from the Matrix Determinant Lemma.
We also make use of the following technical helper lemma.
Lemma 8.3.11. For $\boldsymbol{A} \in \mathbb{R}^{n \times m}$ and all $\boldsymbol{a} \in \mathbb{R}^{n}$ we have

$$
\sum_{i \in[m]} \frac{1}{\psi_{\boldsymbol{A}}[i]}\left(\boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{a}\right)_{i}^{4} \leq\left(\boldsymbol{a}^{T}\left(\boldsymbol{A}^{T} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{a}\right)^{2}
$$

Proof. We have by Cauchy Schwarz that

$$
\left(\mathbf{1}_{i}^{T} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{a}\right)^{2} \leq \psi_{\boldsymbol{A}}[i] \cdot \boldsymbol{a}^{T}\left(\boldsymbol{A}^{T} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{a}
$$

and consequently

$$
\sum_{i \in[m]} \frac{\left(\mathbf{1}_{i}^{T} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{a}\right)^{4}}{\psi_{\boldsymbol{A}}[i]} \leq\left(\boldsymbol{a}^{T}\left(\boldsymbol{A}^{T} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{a}\right) \sum_{i \in[m]}\left(\mathbf{1}_{i} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{a}\right)^{2}
$$

Since

$$
\begin{aligned}
\sum_{i \in[m]}\left(\mathbf{1}_{i}^{T} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{a}\right)^{2} & =\boldsymbol{a}^{T}\left(\boldsymbol{A}^{T} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{A}^{T} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{a} \\
& \leq \boldsymbol{a}^{T}\left(\boldsymbol{A}^{T} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{a}
\end{aligned}
$$

we have the desired result.

We now bound the effect of adding a constraint.
Lemma 8.3.12. Let $\boldsymbol{A} \in \mathbb{R}^{m \times n}, \boldsymbol{b} \in \mathbb{R}^{m}, \boldsymbol{\tau} \in \mathbb{R}^{m}$, and $\boldsymbol{x} \in P \stackrel{\text { def }}{=}\left\{\boldsymbol{y} \in \mathbb{R}^{n}: \boldsymbol{A} \boldsymbol{y} \geq \boldsymbol{b}\right\}$. Let $\overline{\boldsymbol{A}} \in$ $\mathbb{R}^{(m+1) \times n}$ be $\boldsymbol{A}$ with a row $\boldsymbol{a}_{m+1}$ added, let $\bar{b} \in \mathbb{R}^{m+1}$ be the vector $\boldsymbol{b}$ with an entry $b_{m+1}$ added, and let $\bar{P} \stackrel{\text { def }}{=}\left\{\boldsymbol{y} \in \mathbb{R}^{n}: \overline{\boldsymbol{A}} \boldsymbol{y} \geq \bar{b}\right\}$. Let $s_{m+1}=\boldsymbol{a}_{m+1}^{T} \boldsymbol{x}-b_{m+1}>0, \psi_{a}=\frac{\boldsymbol{a}_{m+1}^{T}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{A}_{\boldsymbol{x}}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{a}_{m+1}}{s_{m+1}^{2}}$.

Now, let $\boldsymbol{v} \in \mathbb{R}^{m+1}$ be defined so that $v_{m+1}=\frac{\psi_{a}}{1+\psi_{a}}$ and for all $i \in[m]$

$$
v_{i}=\tau_{i}-\frac{1}{1+\psi_{a}}\left[\boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{A}_{x}+\lambda \boldsymbol{I}\right)^{-1} \frac{\boldsymbol{a}_{m+1}}{s_{m+1}}\right]_{i}^{2}
$$

Then, the following hold

- [Leverage Score Estimation] $e_{\bar{P}}(\vec{v}, \boldsymbol{x})_{m+1}=0$ and $e_{\bar{P}}(\vec{v}, \boldsymbol{x})_{i}=e_{P}(\boldsymbol{\tau}, \boldsymbol{x})_{i}$ for all $i \in[m]$.
- [Function Value Increase] $p_{\boldsymbol{e}_{\bar{P}}(\boldsymbol{v}, \boldsymbol{x})}(\boldsymbol{x})=p_{\boldsymbol{e}_{P}(\tau, \boldsymbol{x})}(\boldsymbol{x})-c_{e} \ln \left(s(\boldsymbol{x})_{m+1} / R\right)+\ln \left(1+\psi_{a}\right)$.
- [Centrality Increase] $\delta_{e_{\bar{P}}(\boldsymbol{v}, \boldsymbol{x})}(\boldsymbol{x}) \leq \sqrt{1+\psi_{a}} \delta_{e_{P}(\boldsymbol{\tau}, \boldsymbol{x})}(\boldsymbol{x})+\left(c_{e} \sqrt{1+\psi_{a}}+\psi_{a}\right) \sqrt{\frac{\psi_{a}}{\mu(\boldsymbol{x})}}+\psi_{a}$.

Proof. By (1) in Lemma 8.3.10, we have that for all $i \in[m]$

$$
\psi_{\bar{P}}(\boldsymbol{x})_{i}=\psi_{P}(\boldsymbol{x})_{i}-\frac{1}{1+\psi_{a}}\left[\boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{A}_{x}+\lambda \boldsymbol{I}\right)^{-1} \frac{\boldsymbol{a}_{m+1}}{s_{m+1}}\right]_{i}^{2}
$$

and that

$$
\psi_{\bar{P}}(\boldsymbol{x})_{m+1}=\psi_{a}-\frac{\psi_{a}^{2}}{1+\psi_{a}}=\frac{\psi_{a}}{1+\psi_{a}} .
$$

Consequently [Leverage Score Estimation] holds.
By (3) in Lemma 8.3.10 we have that [Function Value Change] holds.
To bound the change in centrality, we first note that

$$
\begin{aligned}
\psi_{\bar{P}}(\boldsymbol{x})_{i} & \geq \psi_{P}(\boldsymbol{x})_{i}-\frac{1}{1+\psi_{a}}\left[\frac{\boldsymbol{a}_{i}^{T}}{s_{i}}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{A}_{x}+\lambda \boldsymbol{I}\right)^{-1} \frac{\boldsymbol{a}_{i}}{s_{i}}\right]\left[\frac{\boldsymbol{a}_{m+1}^{T}}{s_{m+1}}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{A}_{x}+\lambda \boldsymbol{I}\right)^{-1} \frac{\boldsymbol{a}_{m+1}}{s_{m+1}}\right] \\
& =\frac{\psi_{P}(\boldsymbol{x})_{i}}{1+\psi_{a}}
\end{aligned}
$$

Therefore, we have that the approximate Hessian for $\bar{P}$, denoted $\overline{\boldsymbol{H}}(\boldsymbol{x})$, is bounded by $\overline{\boldsymbol{H}}^{-1} \preceq(1+$ $\left.\psi_{a}\right) \boldsymbol{H}^{-1}$.

To bound the change in centrality note that by (2) in Lemma 8.3.10 we have that $\overline{\boldsymbol{H}}^{-1} \preceq \boldsymbol{H}^{-1}$. Therefore if let $\boldsymbol{v}^{\prime} \in \mathbb{R}^{m}$ be defined so that $\boldsymbol{v}_{i}^{\prime}=\boldsymbol{v}_{i}$ for all $i \in[m$ then by triangle inequality we have

$$
\begin{aligned}
\delta_{e_{\bar{p}}(\boldsymbol{v}, \boldsymbol{x})}(\boldsymbol{x}) & =\left\|\overline{\boldsymbol{A}}_{x}^{T}\left(c_{e} \mathbf{1}+\boldsymbol{v}\right)\right\|_{\overline{\boldsymbol{H}}^{-1}} \leq \sqrt{1+\psi_{a}}\left\|\overline{\boldsymbol{A}}_{x}^{T}\left(c_{e} \mathbf{1}+\boldsymbol{v}\right)\right\|_{\boldsymbol{H}^{-1}} \\
& \leq \sqrt{1+\psi_{a}}\left(\left\|\boldsymbol{A}_{x}^{T}\left(c_{e} \mathbf{1}+\boldsymbol{\tau}\right)\right\|_{\boldsymbol{H}^{-1}}+\left\|\frac{\boldsymbol{a}_{m+1}}{s_{m+1}}\left(c_{e}+v_{m+1}\right)\right\|_{\boldsymbol{H}^{-1}}+\left\|\boldsymbol{A}_{x}^{T}\left(\boldsymbol{v}^{\prime}-\boldsymbol{\tau}\right)\right\|_{\boldsymbol{H}^{-1}}\right) \\
& =\sqrt{1+\psi_{a}}\left(\delta_{e_{P}(\boldsymbol{\tau}, \boldsymbol{x})}(\boldsymbol{x})+\left(c_{e}+\frac{\psi_{a}}{1+\psi_{a}}\right)\left\|\frac{\boldsymbol{a}_{m+1}}{s_{m+1}}\right\|_{\boldsymbol{H}^{-1}}+\left\|\boldsymbol{A}_{x}^{T}\left(\boldsymbol{v}^{\prime}-\boldsymbol{\tau}\right)\right\|_{\boldsymbol{H}^{-1}}\right)
\end{aligned}
$$

Now, since $\boldsymbol{H}^{-1} \preceq \frac{1}{\mu(\boldsymbol{x})}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{A}_{x}+\lambda \boldsymbol{I}\right)^{-1}$, we have that

$$
\left\|\frac{\boldsymbol{a}_{m+1}}{s_{m+1}}\right\|_{\boldsymbol{H}^{-1}} \leq \frac{1}{\sqrt{\mu(\boldsymbol{x})}}\left\|\frac{\boldsymbol{a}_{m+1}}{s_{m+1}}\right\|_{\left(\boldsymbol{A}_{x}^{T} \boldsymbol{A}_{x}+\lambda \boldsymbol{I}\right)^{-1}}=\sqrt{\frac{\psi_{a}}{\mu(\boldsymbol{x})}}
$$

Since $\boldsymbol{\Psi}^{1 / 2} \boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{\Psi} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}_{x}^{T} \boldsymbol{\Psi}^{1 / 2}$ is a projection matrix, we have $\boldsymbol{\Psi}^{-1} \succeq \boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{\Psi} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}_{x}^{T} \succeq$ $\boldsymbol{A}_{x} \boldsymbol{H}^{-1} \boldsymbol{A}_{x}^{T}$. By Lemma 8.3.11, we have

$$
\begin{aligned}
\left\|\boldsymbol{A}_{x}^{T}\left(\boldsymbol{\tau}^{\prime}-\boldsymbol{v}\right)\right\|_{\boldsymbol{H}^{-1}}^{2} & \leq\left\|\boldsymbol{\tau}^{\prime}-\boldsymbol{v}\right\|_{\boldsymbol{\Psi}^{-1}}^{2} \\
& =\sum_{i \in[m]} \frac{1}{\psi(\boldsymbol{x})_{i}}\left(\frac{1}{1+\psi_{a}}\left(\mathbf{1}_{i} \boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{A}_{x}+\lambda \boldsymbol{I}\right)^{-1} \frac{\boldsymbol{a}_{m+1}}{s_{m+1}}\right)^{2}\right)^{2} \\
& \leq\left(\frac{1}{1+\psi_{a}}\right)^{2}\left(\frac{\boldsymbol{a}_{m+1}^{T}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{A}_{x}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{a}_{m+1}}{s_{m+1}^{2}}\right)^{2}=\left(\frac{\psi_{a}}{1+\psi_{a}}\right)^{2}
\end{aligned}
$$

Combining, we have that

$$
\begin{aligned}
\delta_{e_{\bar{P}}(v, \boldsymbol{x})}(\boldsymbol{x}) & \leq \sqrt{1+\psi_{a}}\left(\delta_{\boldsymbol{e}_{P}(\boldsymbol{\tau}, \boldsymbol{x})}(\boldsymbol{x})+\left(c_{e}+\frac{\psi_{a}}{1+\psi_{a}}\right) \sqrt{\frac{\psi_{a}}{\mu(\boldsymbol{x})}}+\frac{\psi_{a}}{1+\psi_{a}}\right) \\
& \leq \sqrt{1+\psi_{a}} \delta_{e_{P}(\boldsymbol{\tau}, \boldsymbol{x})}(\boldsymbol{x})+\left(c_{e} \sqrt{1+\psi_{a}}+\psi_{a}\right) \sqrt{\frac{\psi_{a}}{\mu(\boldsymbol{x})}}+\psi_{a} .
\end{aligned}
$$

We now bound the effect of removing a constraint.
Lemma 8.3.13 (Removing a Constraint). Let $\boldsymbol{A} \in \mathbb{R}^{m \times n}, \boldsymbol{b} \in \mathbb{R}^{m}, \boldsymbol{\tau} \in \mathbb{R}^{m}$, and $\boldsymbol{x} \in P \stackrel{\text { def }}{=}\{\boldsymbol{y} \in$ $\left.\mathbb{R}^{n}: \boldsymbol{A y} \geq \boldsymbol{b}\right\}$. Let $\overline{\boldsymbol{A}} \in \mathbb{R}^{(m-1) \times n}$ be $\boldsymbol{A}$ with row $m$ removed, let $\bar{b} \in \mathbb{R}^{m-1}$ denote the first $m-1$ coordinates of $\boldsymbol{b}$, and let $\bar{P} \xlongequal{\text { def }}\left\{\boldsymbol{y} \in \mathbb{R}^{n}: \overline{\boldsymbol{A}} \boldsymbol{y} \geq \bar{b}\right\}$. Let $\psi_{d}=\psi_{P}(\boldsymbol{x})_{m}$.

Now, let $\boldsymbol{v} \in \mathbb{R}^{m-1}$ be defined so that for all $i \in[m-1]$

$$
v_{i}=\tau_{i}+\frac{1}{1-\psi_{d}}\left(\boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{A}_{x}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{A}_{x}^{T} \mathbf{1}_{m}\right)_{i}^{2}
$$

Assume $\psi_{d} \leq 1.1 \mu(\boldsymbol{x}) \leq \frac{1}{10}$ and $\left\|\boldsymbol{e}_{P}(\boldsymbol{\tau}, \boldsymbol{x})\right\|_{\infty} \leq c_{e} \leq \frac{1}{2}$, we have the following:

- [Leverage Score Estimation] $e_{\bar{P}}(\vec{v}, \boldsymbol{x})_{i}=e_{P}(\boldsymbol{\tau}, \boldsymbol{x})_{i}$ for all $i \in[m-1]$.
- [Function Value Decrease] $p_{e_{\bar{p}}(\boldsymbol{v}, \boldsymbol{x})}(\boldsymbol{x})=p_{\boldsymbol{e}_{P}(\boldsymbol{\tau}, \boldsymbol{x})}(\boldsymbol{x})+\left[c_{e}+e_{P}(\boldsymbol{\tau}, \boldsymbol{x})_{m}\right] \ln \left(s(\boldsymbol{x})_{m} / R\right)+\ln \left(1-\psi_{d}\right)$
- [Centrality Increase] $\delta_{e_{\bar{p}}(\boldsymbol{v}, \boldsymbol{x})}(\boldsymbol{x}) \leq \frac{1}{\sqrt{1-2 \mu(\boldsymbol{x})}} \delta_{e_{P}(\boldsymbol{\tau}, \boldsymbol{x})}(\boldsymbol{x})+3\left(c_{e}+\mu(\boldsymbol{x})\right)$.

Proof. By (1) in Lemma 8.3.10, we have that for all $i \in[m-1]$

$$
\psi_{\bar{P}}(\boldsymbol{x})_{i}=\psi_{P}(\boldsymbol{x})_{i}+\frac{1}{1-\psi_{d}}\left(\mathbf{1}_{i}^{T} \boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{A}_{x}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{A}_{x}^{T} \mathbf{1}_{m}\right)^{2}
$$

Consequently, [Leverage Score Estimation] holds. Furthermore, by (3) in Lemma 8.3.10, this then implies that [Function Value Change] holds.

To bound the change in centrality, we first note that by (1) and (2) in Lemma 8.3.10 and the fact $\psi_{\bar{P}}(\boldsymbol{x})_{i} \geq \psi_{P}(\boldsymbol{x})_{i}$, we have that the approximate Hessian for $\bar{P}$, denoted $\overline{\boldsymbol{H}}(\boldsymbol{x})$, is bounded by

$$
\begin{aligned}
\overline{\boldsymbol{H}}(\boldsymbol{x})^{-1} & \preceq\left(\boldsymbol{H}(\boldsymbol{x})-\boldsymbol{A}_{x}^{T}\left(c_{e} \boldsymbol{I}+\mathbf{\Psi}_{x}\right)^{1 / 2} \mathbf{1}_{m} \mathbf{1}_{m}^{T}\left(c_{e} \boldsymbol{I}+\mathbf{\Psi}_{x}\right)^{1 / 2} \boldsymbol{A}_{x}\right)^{-1} \\
& \preceq\left(1+\frac{\alpha}{1-\alpha}\right) \boldsymbol{H}(\boldsymbol{x})^{-1}=\left(\frac{1}{1-\alpha}\right) \boldsymbol{H}(\boldsymbol{x})^{-1}
\end{aligned}
$$

where $\alpha \stackrel{\text { def }}{=} \mathbf{1}_{m}^{T}\left(c_{e} \boldsymbol{I}+\boldsymbol{\Psi}_{x}\right)^{1 / 2} \boldsymbol{A}_{x} \boldsymbol{H}(\boldsymbol{x})^{-1} \boldsymbol{A}_{x}^{T}\left(c_{e} \boldsymbol{I}+\mathbf{\Psi}_{x}\right)^{1 / 2} \mathbf{1}_{m}$. Using $c_{e}+\mu(\boldsymbol{x}) \leq \frac{1}{2}+\frac{1}{10} \leq 1$, we have

$$
\begin{equation*}
\boldsymbol{H}(\boldsymbol{x})^{-1} \preceq\left(\boldsymbol{A}_{x}^{T}\left(c_{e}+\mu(\boldsymbol{x})\right) \boldsymbol{A}_{x}+\lambda \boldsymbol{I}\right)^{-1} \preceq\left(c_{e}+\mu(\boldsymbol{x})\right)^{-1}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{A}_{x}+\lambda \boldsymbol{I}\right)^{-1} \tag{8.6}
\end{equation*}
$$

Using this, we have

$$
\begin{equation*}
\alpha \leq\left(\frac{c_{e}+\psi_{d}}{c_{e}+\mu(\boldsymbol{x})}\right) \mathbf{1}_{m}^{T} \boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{A}_{x}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{A}_{x}^{T} \mathbf{1}_{m}=\left(\frac{c_{e}+\psi_{d}}{c_{e}+\mu(\boldsymbol{x})}\right) \psi_{d} \tag{8.7}
\end{equation*}
$$

Now let $\boldsymbol{\tau}^{\prime} \in \mathbb{R}^{m-1}$ be defined so that $\tau_{i}^{\prime}=\tau_{i}$ for all $i \in[m-1]$. We have by above that

$$
\delta_{e_{\bar{p}}(\boldsymbol{v}, \boldsymbol{x})}(\boldsymbol{x})=\left\|\overline{\boldsymbol{A}}_{x}^{T}\left(c_{e} \mathbf{1}+\boldsymbol{v}\right)\right\|_{\overline{\boldsymbol{H}}^{-1}} \leq \frac{1}{\sqrt{1-\alpha}}\left\|\overline{\boldsymbol{A}}_{x}^{T}\left(c_{e} \mathbf{1}+\boldsymbol{v}\right)\right\|_{\boldsymbol{H}^{-1}}
$$

and therefore, by triangle inequality

$$
\begin{aligned}
\left\|\overline{\boldsymbol{A}}_{x}^{T}\left(c_{e} \mathbf{1}+\boldsymbol{v}\right)\right\|_{\boldsymbol{H}^{-1}} & \leq\left\|\boldsymbol{A}_{x}^{T}\left(c_{e} \mathbf{1}+\boldsymbol{\tau}\right)\right\|_{\boldsymbol{H}^{-1}}+\left\|\boldsymbol{A}_{x}^{T} \mathbf{1}_{m}\left(c_{e}+\tau_{m}\right)\right\|_{\boldsymbol{H}^{-1}}+\left\|\boldsymbol{A}_{x}^{T}\left(\boldsymbol{\tau}^{\prime}-\boldsymbol{v}\right)\right\|_{\boldsymbol{H}^{-1}} \\
& =\delta_{\boldsymbol{e}_{P}(\boldsymbol{\tau}, \boldsymbol{x})}(\boldsymbol{x})+\left(c_{e}+\tau_{m}\right)\left\|\boldsymbol{A}_{x}^{T} \mathbf{1}_{m}\right\|_{\boldsymbol{H}^{-1}}+\left\|\boldsymbol{A}_{x}^{T}\left(\boldsymbol{\tau}^{\prime}-\boldsymbol{v}\right)\right\|_{\boldsymbol{H}^{-1}}
\end{aligned}
$$

Now, (8.6) shows that

$$
\left\|\boldsymbol{A}_{x}^{T} \mathbf{1}_{m}\right\|_{\boldsymbol{H}^{-1}} \leq \frac{1}{\sqrt{c_{e}+\mu(\boldsymbol{x})}}\left\|\boldsymbol{A}_{x}^{T} \mathbf{1}_{m}\right\|_{\left(\boldsymbol{A}_{x}^{T} \boldsymbol{A}_{x}+\lambda \boldsymbol{I}\right)^{-1}} \leq \sqrt{\frac{\psi_{d}}{c_{e}+\mu(\boldsymbol{x})}}
$$

Furthermore, since $\boldsymbol{\Psi}^{-1} \succeq \boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{\Psi} \boldsymbol{A}_{x}\right)^{-1} \boldsymbol{A}_{x}^{T} \succeq \boldsymbol{A}_{x} \boldsymbol{H}^{-1} \boldsymbol{A}_{x}^{T}$, by Lemma 8.3.11 we have

$$
\begin{aligned}
\left\|\boldsymbol{A}_{x}^{T}\left(\boldsymbol{\tau}^{\prime}-\boldsymbol{v}\right)\right\|_{\boldsymbol{H}^{-1}}^{2} & \leq\left\|\boldsymbol{\tau}^{\prime}-\boldsymbol{v}\right\|_{\boldsymbol{\Psi}^{-1}}^{2} \\
& =\sum_{i \in[m]} \frac{1}{\psi(\boldsymbol{x})_{i}}\left(\frac{1}{1-\psi_{d}}\left(\mathbf{1}_{i}^{T} \boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{A}_{x}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{A}_{x}^{T} \mathbf{1}_{m}\right)^{2}\right)^{2} \\
& \leq\left(\frac{1}{1-\psi_{d}}\right)^{2}\left(\mathbf{1}_{m}^{T} \boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{A}_{x}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{A}_{x}^{T} \mathbf{1}_{m}\right)^{2}=\left(\frac{\psi_{d}}{1-\psi_{d}}\right)^{2}
\end{aligned}
$$

Combining, we have that

$$
\delta_{e_{\bar{p}}(\boldsymbol{v}, \boldsymbol{x})}(\boldsymbol{x}) \leq \frac{1}{\sqrt{1-\alpha}}\left[\delta_{\boldsymbol{e}_{P}(\boldsymbol{\tau}, \boldsymbol{x})}(\boldsymbol{x})+\left(c_{e}+\tau_{m}\right) \sqrt{\frac{\psi_{d}}{c_{e}+\mu(\boldsymbol{x})}}+\frac{\psi_{d}}{1-\psi_{d}}\right]
$$

Using the assumption $\psi_{d} \leq 1.1 \mu(\boldsymbol{x}) \leq \frac{1}{10},\left\|\boldsymbol{e}_{P}(\boldsymbol{\tau}, \boldsymbol{x})\right\|_{\infty} \leq c_{e}$ and (8.7), we have $\alpha \leq 1.1 \psi_{d} \leq 1.21 \mu(\boldsymbol{x})$ and $\tau_{m} \leq \psi_{d}+c_{e}$, and

$$
\begin{aligned}
\delta_{e_{\bar{p}}(\boldsymbol{v}, \boldsymbol{x})}(\boldsymbol{x}) & \leq \frac{1}{\sqrt{1-1.3 \mu(\boldsymbol{x})}}\left[\delta_{e_{P}(\tau, \boldsymbol{x})}(\boldsymbol{x})+\left(c_{e}+\tau_{m}\right) \sqrt{1.1}+1.2 \psi_{d}\right] \\
& \leq \frac{1}{\sqrt{1-2 \mu(\boldsymbol{x})}} \delta_{e_{P}(\tau, x)}(\boldsymbol{x})+\frac{1}{\sqrt{1-\frac{1.3}{10}}}\left(\sqrt{1.1} \cdot 2 c_{e}+1.1(\sqrt{1.1}+1.2) \mu(\boldsymbol{x})\right)
\end{aligned}
$$

## ■ 8.3.3 Hybrid Center Properties

Here we prove properties of points near the hybrid center. First we bound the distance between points in the $\boldsymbol{H}(\boldsymbol{x})$ norm in terms of the $\ell_{2}$ norm of the points.

Lemma 8.3.14. For $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ and $\boldsymbol{b} \in \mathbb{R}^{m}$ suppose that $\boldsymbol{x} \in P=\{\boldsymbol{y}: \boldsymbol{A} \boldsymbol{y} \geq \boldsymbol{b}\}$ and $\boldsymbol{e} \in \mathbb{R}^{m}$ such that $\|\boldsymbol{e}\|_{\infty} \leq \frac{1}{2} c_{e}<\frac{1}{20}$ and $\delta_{e} \leq 0.1 \sqrt{c_{e}+\mu(\boldsymbol{x})}$. Then for all $\boldsymbol{y} \in P$ we have

$$
\begin{equation*}
\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})} \leq \frac{12 m c_{e}+6 n+2 \lambda\|\boldsymbol{y}\|_{2}^{2}}{\sqrt{c_{e}+\mu(\boldsymbol{x})}} \tag{8.8}
\end{equation*}
$$

and

$$
\|\boldsymbol{x}\|_{2}^{2} \leq 4 \lambda^{-1}\left(m c_{e}+n\right)+2\|\boldsymbol{y}\|_{2}^{2} .
$$

Proof. For notational simplicity let $\boldsymbol{t} \stackrel{\text { def }}{=} c_{e} \mathbf{1}+\boldsymbol{e}+\boldsymbol{\psi}_{x}, \boldsymbol{T} \stackrel{\text { def }}{=} \boldsymbol{\operatorname { D i a g }}(\boldsymbol{t})$, and $\boldsymbol{M} \stackrel{\text { def }}{=} \boldsymbol{A}_{x}^{T}\left(c_{e} \boldsymbol{I}+\boldsymbol{\Psi}_{x}\right) \boldsymbol{A}_{x}$. We have

$$
\begin{equation*}
\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{A}_{x}^{T} \boldsymbol{T} \boldsymbol{A}_{x}}^{2}=\sum_{i \in[m]} t_{i} \frac{\left[\boldsymbol{s}_{x}-\boldsymbol{s}_{y}\right]_{i}^{2}}{\left[s_{x}\right]_{i}^{2}}=\sum_{i \in[m]} t_{i}\left(1-2 \frac{\left[\boldsymbol{s}_{y}\right]_{i}}{\left[\boldsymbol{s}_{x}\right]_{i}}+\frac{\left[\boldsymbol{s}_{y}\right]_{i}^{2}}{\left[s_{x}\right]_{i}^{2}}\right) \tag{8.9}
\end{equation*}
$$

and

$$
\begin{equation*}
\sum_{i \in[m]} t_{i} \frac{\left[\boldsymbol{s}_{y}\right]_{i}^{2}}{\left[\boldsymbol{s}_{x}\right]_{i}^{2}} \leq\left(\sum_{i \in[m]} t_{i} \frac{\left[\boldsymbol{s}_{y}\right]_{i}}{\left[\boldsymbol{s}_{x}\right]_{i}}\right) \max _{i \in[m]} \frac{\left[\boldsymbol{s}_{y}\right]_{i}}{\left[\boldsymbol{s}_{x}\right]_{i}} \leq\left(\sum_{i \in[m]} t_{i} \frac{\left[\boldsymbol{s}_{y}\right]_{i}}{\left[\boldsymbol{s}_{x}\right]_{i}}\right)\left(1+\left\|\boldsymbol{S}_{x}^{-1}\left(\boldsymbol{s}_{y}-\boldsymbol{s}_{x}\right)\right\|_{\infty}\right) \tag{8.10}
\end{equation*}
$$

and

$$
\begin{align*}
\left\|\boldsymbol{S}_{x}^{-1}\left(\boldsymbol{s}_{x}-\boldsymbol{s}_{y}\right)\right\|_{\infty} & =\max _{i \in[m]}\left|\mathbf{1}_{i} \boldsymbol{S}_{x}^{-1} \boldsymbol{A}(\boldsymbol{x}-\boldsymbol{y})\right| \leq\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})} \sqrt{\max _{i \in[m]}\left[\boldsymbol{S}_{x}^{-1} \boldsymbol{A H} \boldsymbol{H}(\boldsymbol{x})^{-1} \boldsymbol{A}^{T} \boldsymbol{S}_{x}^{-1}\right]_{i i}} \\
& \leq\left(c_{e}+\mu(\boldsymbol{x})\right)^{-1 / 2}\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})} \tag{8.11}
\end{align*}
$$

Now, clearly $\sum_{i \in[m]} t_{i}\left[\boldsymbol{s}_{y}\right]_{i} /\left[\boldsymbol{s}_{x}\right]_{i}$ is positive and since $\|\boldsymbol{e}\|_{\infty} \leq \frac{1}{2} c_{e}$ we know that $\frac{1}{2} \boldsymbol{M} \preceq \boldsymbol{A}_{x}^{T} \boldsymbol{T} \boldsymbol{A}_{x}$. Therefore, by combining, (8.9), (8.10), and (8.11) we have

$$
\begin{align*}
\frac{1}{2}\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{M}}^{2} & \leq\|\boldsymbol{t}\|_{1}-\sum_{i \in[m]} t_{i} \frac{\left[\boldsymbol{s}_{y}\right]_{i}}{\left[\boldsymbol{s}_{x}\right]_{i}}+\left(\sum_{i \in[m]} t_{i} \frac{\left[s_{y}\right]_{i}}{\left[\boldsymbol{s}_{x}\right]_{i}}\right) \frac{\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})}}{\sqrt{c_{e}+\mu(\boldsymbol{x})}} \\
& \leq\|\boldsymbol{t}\|_{1}+\left(\sum_{i \in[m]} t_{i} \frac{\left[\boldsymbol{s}_{y}\right]_{i}}{\left[s_{x}\right]_{i}}\right) \frac{\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})}}{\sqrt{c_{e}+\mu(\boldsymbol{x})}} \tag{8.12}
\end{align*}
$$

Now since $\nabla p_{e}(\boldsymbol{x})=-\boldsymbol{A}^{T} \boldsymbol{S}_{x}^{-1} \boldsymbol{T} \mathbf{1}+\lambda \boldsymbol{x}$ we have

$$
\left\langle\boldsymbol{x}-\boldsymbol{y}, \nabla p_{\boldsymbol{e}}(\boldsymbol{x})\right\rangle=-\sum_{i \in[m]} t_{i} \frac{\left[\boldsymbol{s}_{x}-\boldsymbol{s}_{y}\right]_{i}}{\left[\boldsymbol{s}_{x}\right]_{i}}+\lambda \boldsymbol{x}^{T}(\boldsymbol{x}-\boldsymbol{y})
$$

and therefore by Cauchy Schwarz and $\boldsymbol{x}^{T} \boldsymbol{y} \leq\|\boldsymbol{x}\|_{2}^{2}+\frac{1}{4}\|\boldsymbol{y}\|_{2}^{2}$,

$$
\begin{align*}
\sum_{i \in[m]} t_{i} \frac{\left[\boldsymbol{s}_{y}\right]_{i}}{\left[\boldsymbol{s}_{x}\right]_{i}} & =\|\boldsymbol{t}\|_{1}-\lambda\|\boldsymbol{x}\|_{2}^{2}+\lambda \boldsymbol{x}^{T} \boldsymbol{y}+\left\langle\boldsymbol{x}-\boldsymbol{y}, \nabla p_{\boldsymbol{e}}(\boldsymbol{x})\right\rangle  \tag{8.13}\\
& \leq\|t\|_{1}+\frac{\lambda}{4}\|\boldsymbol{y}\|_{2}^{2}+\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})} \delta_{\boldsymbol{e}}(\boldsymbol{x}) \tag{8.14}
\end{align*}
$$

Now, using (8.12), (8.14) and the definition of $\boldsymbol{H}(\boldsymbol{x})$, we have

$$
\begin{aligned}
\frac{1}{2}\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})}^{2} & =\frac{1}{2}\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{M}}^{2}+\frac{\lambda}{2}\|\boldsymbol{x}-\boldsymbol{y}\|_{2}^{2} \\
& \leq\|\boldsymbol{t}\|_{1}+\left(\sum_{i \in[m]} t_{i} \frac{\left[\boldsymbol{s}_{y}\right]_{i}}{\left[s_{x}\right]_{i}}\right) \frac{\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})}}{\sqrt{c_{e}+\mu(\boldsymbol{x})}}+\frac{\lambda}{2}\|\boldsymbol{x}-\boldsymbol{y}\|_{2}^{2} \\
& \leq\|\boldsymbol{t}\|_{1}+\frac{\|t\|_{1}+\frac{\lambda}{4}\|\boldsymbol{y}\|_{2}^{2}}{\sqrt{c_{e}+\mu(\boldsymbol{x})}}\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})}+\delta_{e}(\boldsymbol{x}) \frac{\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})}^{2}}{\sqrt{c_{e}+\mu(\boldsymbol{x})}}+\frac{\lambda}{2}\|\boldsymbol{x}-\boldsymbol{y}\|_{2}^{2} .
\end{aligned}
$$

Using the fact that $\delta_{e}(\boldsymbol{x}) \leq 0.1 \sqrt{c_{e}+\mu(\boldsymbol{x})}$, we have

$$
\begin{equation*}
\frac{1}{4}\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})}^{2} \leq\|\boldsymbol{t}\|_{1}+\frac{\lambda}{2}\|\boldsymbol{x}-\boldsymbol{y}\|_{2}^{2}+\frac{\|t\|_{1}+\frac{\lambda}{4}\|\boldsymbol{y}\|_{2}^{2}}{\sqrt{c_{e}+\mu(\boldsymbol{x})}}\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})} . \tag{8.15}
\end{equation*}
$$

Furthermore, since $\sum_{i \in[m]} t_{i}\left[s_{y}\right]_{i} /\left[s_{x}\right]_{i}$ is positive, (8.13) shows that

$$
\lambda \boldsymbol{x}^{T}(\boldsymbol{x}-\boldsymbol{y})=\lambda\|\boldsymbol{x}\|_{2}^{2}-\lambda \boldsymbol{x}^{T} \boldsymbol{y} \leq\|\boldsymbol{t}\|_{1}+\left\langle\boldsymbol{x}-\boldsymbol{y}, \nabla p_{\boldsymbol{e}}(\boldsymbol{x})\right\rangle \leq\|\boldsymbol{t}\|_{1}+\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})} \delta_{e}(\boldsymbol{x})
$$

and hence

$$
\begin{align*}
\frac{\lambda}{2}\|\boldsymbol{x}-\boldsymbol{y}\|_{2}^{2} & \leq \frac{\lambda}{2}\|\boldsymbol{x}-\boldsymbol{y}\|_{2}^{2}+\frac{\lambda}{2}\|\boldsymbol{x}\|_{2}^{2}=\lambda \boldsymbol{x}^{T}(\boldsymbol{x}-\boldsymbol{y})+\frac{\lambda}{2}\|\boldsymbol{y}\|_{2}^{2} \\
& \leq\|\boldsymbol{t}\|_{1}+\frac{\lambda}{2}\|\boldsymbol{y}\|_{2}^{2}+\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})} \delta_{e}(\boldsymbol{x}) \tag{8.16}
\end{align*}
$$

Putting (8.16) into (8.15) and using the fact that $\delta_{e}(\boldsymbol{x}) \leq 0.1 \sqrt{c_{e}+\mu(\boldsymbol{x})}$, we have

$$
\frac{1}{4}\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})}^{2} \leq 2\|\boldsymbol{t}\|_{1}+\frac{\lambda}{2}\|\boldsymbol{y}\|_{2}^{2}+\left(0.1+\frac{\|\boldsymbol{t}\|_{1}+\frac{\lambda}{4}\|\boldsymbol{y}\|_{2}^{2}}{\sqrt{c_{e}+\mu(\boldsymbol{x})}}\right)\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})} .
$$

Now, using $\|\boldsymbol{t}\|_{1} \leq 2 m c_{e}+n$ and $\sqrt{c_{e}+\mu(\boldsymbol{x})} \leq 1.05$, we have

$$
\frac{1}{4}\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})}^{2} \leq 2.2 \alpha+(0.1+\alpha)\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})} \quad \text { for } \quad \alpha=\frac{2 m c_{e}+n+\frac{\lambda}{4}\|\boldsymbol{y}\|_{2}^{2}}{\sqrt{c_{e}+\mu(\boldsymbol{x})}} .
$$

Since $\alpha \geq 1 / \sqrt{1.1}$, we have that

$$
\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})} \leq \frac{0.1+\alpha+\sqrt{(\alpha+0.1)^{2}+2.2 \alpha}}{2 \cdot \frac{1}{4}} \leq 6 \alpha
$$

yielding (8.8).
We also have by (8.13) and the fact that $\delta_{e}(\boldsymbol{x}) \leq 0.1 \sqrt{c_{e}+\mu(\boldsymbol{x})}$,

$$
\begin{aligned}
\lambda\|\boldsymbol{x}\|_{2}^{2} & =\|t\|_{1}+\lambda \boldsymbol{x}^{T} \boldsymbol{y}+\left\langle\boldsymbol{x}-\boldsymbol{y}, \nabla p_{\boldsymbol{e}}(\boldsymbol{x})\right\rangle-\sum_{i \in[m]} t_{i} \frac{\left[\boldsymbol{s}_{y}\right]_{i}}{\left[\boldsymbol{s}_{x}\right]_{i}} \\
& \leq\|t\|_{1}+\frac{\lambda}{2}\|\boldsymbol{x}\|_{2}^{2}+\frac{\lambda}{2}\|\boldsymbol{y}\|_{2}^{2}+\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})} \delta_{e}(\boldsymbol{x}) \\
& \leq\|t\|_{1}+\frac{\lambda}{2}\|\boldsymbol{x}\|_{2}^{2}+\frac{\lambda}{2}\|\boldsymbol{y}\|_{2}^{2}+0.1 \sqrt{c_{e}+\mu(\boldsymbol{x})}\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})}
\end{aligned}
$$

Hence, using $\|\boldsymbol{t}\|_{1} \leq 2 m c_{e}+n$ and $\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})} \leq 6 \alpha$, we have

$$
\begin{aligned}
\frac{\lambda}{2}\|\boldsymbol{x}\|_{2}^{2} & \leq\|t\|_{1}+\frac{\lambda}{2}\|\boldsymbol{y}\|_{2}^{2}+0.6\left(2 m c_{e}+n+\frac{\lambda}{4}\|\boldsymbol{y}\|_{2}^{2}\right) \\
& \leq \lambda\|\boldsymbol{y}\|_{2}^{2}+2\left(m c_{e}+n\right)
\end{aligned}
$$

In the following lemma we show how we can write one hyperplane in terms of the others provided that we are nearly centered and show there is a constraint that the central point is close to.

Lemma 8.3.15. Let $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ and $\boldsymbol{b} \in \mathbb{R}^{m}$ such that $\left\|a_{i}\right\|_{2}=1$ for all $i$. Suppose that $\boldsymbol{x} \in P=\{\boldsymbol{y}$ : $\boldsymbol{A} \boldsymbol{y} \geq \boldsymbol{b}\}$ and $\boldsymbol{e} \in \mathbb{R}^{m}$ such that $\|\boldsymbol{e}\|_{\infty} \leq \frac{1}{2} c_{e} \leq \frac{1}{2}$. Furthermore, let $\varepsilon=\min _{j \in[m]} s_{j}(\boldsymbol{x})$ and suppose that $i=\arg \min _{j \in[m]} s_{j}(\boldsymbol{x})$ then

$$
\left\|\boldsymbol{a}_{i}+\sum_{j \neq i}\left(\frac{s(x)_{i}}{s(x)_{j}}\right)\left(\frac{c_{e}+e_{j}+\psi_{j}(\boldsymbol{x})}{c_{e}+e_{i}+\psi_{i}(\boldsymbol{x})}\right) \boldsymbol{a}_{j}\right\|_{2} \leq \frac{2 \varepsilon}{\left(c_{e}+\mu(\boldsymbol{x})\right)}\left[\lambda\|\boldsymbol{x}\|_{2}+\delta_{e}(\boldsymbol{x}) \sqrt{\frac{m c_{e}+n}{\varepsilon^{2}}+\lambda}\right] .
$$

Proof. We know that

$$
\begin{aligned}
\nabla p_{e}(\boldsymbol{x}) & =-\boldsymbol{A}^{T} \boldsymbol{S}_{x}^{-1}\left(c_{e} \mathbf{1}+\boldsymbol{e}+\boldsymbol{\psi}_{x}\right)+\lambda \boldsymbol{x} \\
& =\lambda \boldsymbol{x}-\sum_{i \in[m]} \frac{\left(c_{e}+e_{i}+\psi_{i}\right)}{s(\boldsymbol{x})_{i}} \boldsymbol{a}_{i}
\end{aligned}
$$

Consequently, by $\|\boldsymbol{e}\|_{\infty} \leq \frac{1}{2} c_{e}$, and $\psi_{i}(\boldsymbol{x}) \geq \mu(\boldsymbol{x})$, we have

$$
\left\|\boldsymbol{a}_{i}+\sum_{j \neq i}\left(\frac{s(x)_{i}}{s(x)_{j}}\right)\left(\frac{c_{e}+e_{j}+\psi_{j}(\boldsymbol{x})}{c_{e}+e_{i}+\psi_{i}(\boldsymbol{x})}\right) \boldsymbol{a}_{j}\right\|_{2}=\frac{s_{i}(\boldsymbol{x})}{c_{e}+e_{i}+\psi_{i}(\boldsymbol{x})}\left\|\boldsymbol{A}^{T} \boldsymbol{S}_{x}^{-1}\left(c_{e} \mathbf{1}+\boldsymbol{e}+\boldsymbol{\psi}_{x}\right)\right\|_{2} .
$$

Using $\left\|\boldsymbol{a}_{i}\right\|=1, \sum_{i} \psi_{i} \leq n$, and $s_{i}(\boldsymbol{x}) \geq \varepsilon$, we have

$$
\begin{align*}
\operatorname{Tr}\left(\boldsymbol{A}_{x}^{T}\left(c_{e} \boldsymbol{I}+\boldsymbol{\Psi}_{x}\right) \boldsymbol{A}_{x}\right) & =\operatorname{Tr}\left(\boldsymbol{A}_{x} \boldsymbol{A}_{x}^{T}\left(c_{e} \boldsymbol{I}+\mathbf{\Psi}_{x}\right)\right) \\
& =\sum_{i}\left(c_{e}+\psi_{i}\right) \frac{\left\|a_{i}\right\|_{2}^{2}}{s_{i}^{2}(\boldsymbol{x})} \leq \frac{m c_{e}+n}{\varepsilon^{2}} \tag{8.17}
\end{align*}
$$

Hence, we have $\boldsymbol{H}(\boldsymbol{x}) \preceq\left(\frac{m c_{e}+n}{\varepsilon^{2}}+\lambda\right) \boldsymbol{I}$ and $\left\|\nabla p_{e}(\boldsymbol{x})\right\|_{2} \leq \delta_{e}(\boldsymbol{x}) \sqrt{\frac{m c_{e}+n}{\varepsilon^{2}}+\lambda}$ yielding the result.

### 8.3.4 The Algorithm

Here, we put all the results in the previous sections to get our ellipsoid algorithm. Below is a sketch of the pseudocode; we use $c_{a}, c_{d}, c_{e}, c_{\Delta}$ to denote parameters we decide later.

```
Algorithm 22: Our Cutting Plane Method
    Input: \(\boldsymbol{A}^{(0)} \in \mathbb{R}^{m \times n}, \boldsymbol{b}^{(0)} \in \mathbb{R}^{m}, \varepsilon>0\), and radius \(R>0\).
    Input: A separation oracle for a non-empty set \(K \subset B_{\infty}(R)\).
    Check: Throughout the algorithm, if \(s_{i}\left(\boldsymbol{x}^{(k)}\right)<\varepsilon\) output \(P^{(k)}\).
    Check: Throughout the algorithm, if \(\boldsymbol{x}^{(k)} \in K\), output \(\boldsymbol{x}^{(k)}\).
    Set \(P^{(0)}=B_{\infty}(R)\).
    Set \(\boldsymbol{x}^{(0)}:=\mathbf{0}\) and compute \(\tau_{i}^{(0)}=\psi_{P^{(0)}}\left(\boldsymbol{x}^{(0)}\right)_{i}\) for all \(i \in[m]\) exactly.
    for \(k=0\) to \(\infty\) do
        Let \(m^{(k)}\) be the number of constraints in \(P^{(k)}\).
        Compute \(\boldsymbol{w}^{(k)}\) such that \(\boldsymbol{\Psi}_{P^{(k)}}\left(\boldsymbol{x}^{(k)}\right) \preceq \boldsymbol{W}^{(k)} \preceq\left(1+c_{\Delta}\right) \boldsymbol{\Psi}_{P^{(k)}}\left(\boldsymbol{x}^{(k)}\right)\).
        Let \(i^{(k)} \in \arg \max _{i \in\left[m^{(k)}\right]}\left|w_{i}^{(k)}-\tau_{i}^{(k)}\right|\).
        Set \(\tau_{i^{(k)}}^{\left(k+\frac{1}{3}\right)}=\psi_{P^{(k)}}\left(\boldsymbol{x}^{(k)}\right)_{i^{(k)}}\) and \(\tau_{j}^{\left(k+\frac{1}{3}\right)}=\tau_{j}^{(k)}\) for all \(j \neq i^{(k)}\).
        if \(\min _{i \in\left[m^{(k)}\right]} w_{i}^{(k)} \leq c_{d}\) then
            Remove constraint with minimum \(w_{i}^{(k)}\) yielding polytope \(P^{(k+1)}\).
            Update \(\boldsymbol{\tau}\) according to Lemma 8.3.13 to get \(\tau_{j}^{\left(k+\frac{2}{3}\right)}\).
        else
            Use separation oracle at \(\boldsymbol{x}^{(k)}\) to get a constraint \(\left\{\boldsymbol{x}: \boldsymbol{a}^{T} \boldsymbol{x} \geq \boldsymbol{a}^{T} \boldsymbol{x}^{(k)}\right\}\) with \(\|\boldsymbol{a}\|_{2}=1\).
            Add constraint \(\left\{\boldsymbol{x}: \boldsymbol{a}^{T} \boldsymbol{x} \geq \boldsymbol{a}^{T} \boldsymbol{x}^{(k)}-c_{a}^{-1 / 2} \sqrt{\boldsymbol{a}^{T}\left(\boldsymbol{A}^{T} \boldsymbol{S}_{\boldsymbol{x}^{(k)}}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{a}}\right\}\) yielding polytope
                    \(P^{(k+1)}\).
                Update \(\boldsymbol{\tau}\) according to Lemma 8.3.12 to get \(\tau_{j}^{\left(k+\frac{2}{3}\right)}\).
        \(\left(\boldsymbol{x}^{(k+1)}, \boldsymbol{\tau}^{(k+1)}\right)=\) Centering \(\left(\boldsymbol{x}^{(k)}, \boldsymbol{\tau}^{\left(k+\frac{2}{3}\right)}, 200, c_{\Delta}\right)\).
    end
```

In the algorithm, there are two main invariants we maintain. First, we maintain that the centrality $\delta_{P, \boldsymbol{e}}(\boldsymbol{x})$, which indicates how close $\boldsymbol{x}$ is to the minimum point of $p_{\boldsymbol{e}}$, is small. Second, we maintain that $\|\boldsymbol{e}(\boldsymbol{\tau}, \boldsymbol{x})\|_{\infty}$, which indicates how accurate the leverage score estimate $\boldsymbol{\tau}$ is, is small. In the following lemma we show that we maintain both invariants throughout the algorithm.

Lemma 8.3.16. Assume that $c_{e} \leq c_{d} \leq \frac{1}{10^{6}}, c_{a} \sqrt{c_{a}} \leq \frac{c_{d}}{10^{3}}, c_{d} \leq c_{a}$, and $c_{\Delta} \leq C c_{e} / \log (n \log (R / \varepsilon))$ for some small enough universal constant $C$. During our cutting plane method, for all $k$, with high probability in $n$, we have

1. $\left\|\boldsymbol{e}\left(\boldsymbol{\tau}^{\left(k+\frac{1}{3}\right)}, \boldsymbol{x}^{(k)}\right)\right\|_{\infty} \leq \frac{1}{1000} c_{e},\left\|\boldsymbol{e}\left(\boldsymbol{\tau}^{\left(k+\frac{2}{3}\right)}, \boldsymbol{x}^{(k)}\right)\right\|_{\infty} \leq \frac{1}{1000} c_{e},\left\|\boldsymbol{e}\left(\boldsymbol{\tau}^{(k+1)}, \boldsymbol{x}^{(k+1)}\right)\right\|_{\infty} \leq \frac{1}{400} c_{e}$.
2. $\delta_{P^{(k)}, e\left(\tau^{\left(k+\frac{2}{3}\right)}, \boldsymbol{x}^{(k)}\right)}\left(\boldsymbol{x}^{(k)}\right) \leq \frac{1}{100} \sqrt{c_{e}+\min \left(\mu\left(\boldsymbol{x}^{(k)}\right), c_{d}\right)}$.
3. $\delta_{P^{(k+1)}, \boldsymbol{e}\left(\boldsymbol{\tau}^{(k+1)}, \boldsymbol{x}^{(k+1)}\right)}\left(\boldsymbol{x}^{(k+1)}\right) \leq \frac{1}{400} \sqrt{c_{e}+\min \left(\mu\left(\boldsymbol{x}^{(k+1)}\right), c_{d}\right)}$.

Proof. Some statements of the proof hold only with high probability in $n$; we omit mentioning this for simplicity.

We prove by induction on $k$. Note that the claims are written in order consistent with the algorithm and proving the statement for $k$ involves bounding centrality at the point $\boldsymbol{x}^{(k+1)}$. Trivially we define, $\boldsymbol{\tau}^{(-1)}=\boldsymbol{\tau}^{\left(-\frac{2}{3}\right)}=\boldsymbol{\tau}^{\left(-\frac{1}{3}\right)}=\boldsymbol{\tau}^{(0)}, \boldsymbol{x}^{(-1)}=\boldsymbol{x}^{(0)}$ and note that the claims then hold for $k=-1$ as we compute the initial leverage scores, $\boldsymbol{\tau}^{(0)}$, exactly and since the polytope is symmetric we have $\delta_{\boldsymbol{e}\left(\boldsymbol{\tau}^{(0)}, \boldsymbol{x}^{(0)}\right)}(\mathbf{0})=0$. We now suppose they hold for all $t<k$ and show that they hold for $t=k$.

We first bound $\delta$. For notational simplicity, let $\eta_{t} \stackrel{\text { def }}{=} \sqrt{c_{e}+\min \left(\mu\left(\boldsymbol{x}^{(t)}\right), c_{d}\right)}$. By the induction hypothesis we know that $\delta_{P^{(t)}, \boldsymbol{e}\left(\boldsymbol{\tau}^{(t)}, \boldsymbol{x}^{(t)}\right)}\left(\boldsymbol{x}^{(t)}\right) \leq \frac{1}{400} \eta_{t}$. Now, when we update $\tau^{(t)}$ to $\tau^{\left(t+\frac{1}{3}\right)}$, we set $\boldsymbol{e}_{i^{(t)}}$ to 0 . Consequently, Lemma 8.3.8 and the induction hypothesis $\left\|\boldsymbol{e}\left(\boldsymbol{\tau}^{(t)}, \boldsymbol{x}^{(t)}\right)\right\|_{\infty} \leq \frac{1}{400} c_{e}$ show that

$$
\begin{align*}
\delta_{P^{(t)}, \boldsymbol{e}\left(\boldsymbol{\tau}^{\left(t+\frac{1}{3}\right)}, \boldsymbol{x}^{(t)}\right)}\left(\boldsymbol{x}^{(t)}\right) & \leq \delta_{P^{(t)}, \boldsymbol{e}\left(\boldsymbol{\tau}^{(t)}, \boldsymbol{x}^{(t)}\right)}\left(\boldsymbol{x}^{(t)}\right)+\frac{1}{\sqrt{c_{e}+\mu\left(\boldsymbol{x}^{(t)}\right)}} e_{i^{(t)}}\left(\boldsymbol{\tau}^{(t)}, \boldsymbol{x}^{(t)}\right) \\
& \leq \frac{1}{400} \eta_{t}+\frac{\sqrt{c_{e}}}{400} \leq \frac{\eta_{t}}{200} \tag{8.18}
\end{align*}
$$

Next, we estimate the $\delta$ changes when we remove or add a constraint.
For the case of removal, we note that it happens only if $\mu\left(\boldsymbol{x}^{(t)}\right) \leq \min _{i} w_{i} \leq c_{d} \leq \frac{1}{10^{6}}$. Also, the row we remove has leverage score at most $1.1 \mu\left(\boldsymbol{x}^{(t)}\right)$ because we pick the row with minimum $w$. Hence, Lemma 8.3.13 show that

$$
\begin{aligned}
\delta_{P^{(t+1)}, \boldsymbol{e}\left(\boldsymbol{\tau}^{\left.\left(t+\frac{2}{3}\right), \boldsymbol{x}^{(t)}\right)}\right.}\left(\boldsymbol{x}^{(t)}\right) & \leq \frac{1}{\sqrt{1-2 \mu\left(\boldsymbol{x}^{(t)}\right)}} \delta_{P^{(t)}, \boldsymbol{e}\left(\boldsymbol{\tau}^{\left(t+\frac{1}{3}\right)}, \boldsymbol{x}^{(t)}\right)}\left(\boldsymbol{x}^{(t)}\right)+3\left(c_{e}+\mu\left(\boldsymbol{x}^{(t)}\right)\right) \\
& \leq \frac{1}{\sqrt{1-2 \cdot 10^{-6}}}\left(\frac{\eta_{t}}{200}\right)+3\left(c_{e}+\mu\left(\boldsymbol{x}^{(t)}\right)\right) \leq \frac{\eta_{t}}{100}
\end{aligned}
$$

where we used the fact $\mu\left(\boldsymbol{x}^{(t)}\right) \leq c_{d}$ and hence $c_{e}+\mu\left(\boldsymbol{x}^{(t)}\right) \leq \sqrt{2 c_{d}} \eta_{t} \leq \frac{\sqrt{2}}{1000} \eta_{t}$.
For the case of addition, we note that it happens only if $2 \mu\left(\boldsymbol{x}^{(t)}\right) \geq \min _{i} w_{i} \geq c_{d}$. Furthermore, in this case the hyperplane we add is chosen precisely so that $\psi_{a}=c_{a}$. Furthermore, since $c_{e} \leq c_{d} \leq$ $c_{a} \leq \frac{1}{100}$ by Lemma 8.3.12 we have that

$$
\begin{aligned}
\delta_{P^{(t+1)}, e\left(\boldsymbol{\tau}^{\left(t+\frac{2}{3}\right)}, \boldsymbol{x}^{(t)}\right)}\left(\boldsymbol{x}^{(t)}\right) & \leq \sqrt{1+c_{a}} \delta_{P^{(t)}, \boldsymbol{e}\left(\boldsymbol{\tau}^{\left(t+\frac{1}{3}\right)}, \boldsymbol{x}^{(t)}\right)}+\left(c_{e} \sqrt{1+c_{a}}+c_{a}\right) \sqrt{\frac{c_{a}}{\mu\left(\boldsymbol{x}^{(t)}\right)}}+c_{a} \\
& \leq \frac{\eta_{t}}{190}+4 c_{a} \sqrt{\frac{c_{a}}{c_{d}}}
\end{aligned}
$$

Furthermore, since $c_{a} \sqrt{c_{a}} \leq \frac{c_{d}}{1000}, \mu\left(\boldsymbol{x}^{(t)}\right) \geq c_{d} / 2$, and $c_{d} \leq 10^{-6}$ we know that $4 c_{a} \sqrt{c_{a} / c_{d}} \leq \frac{1}{250} \eta_{t}$ and consequently in both cases we have $\delta_{P^{(t+1)}, \boldsymbol{e}\left(\boldsymbol{\tau}^{\left(t+\frac{2}{3}\right)}, \boldsymbol{x}^{(t)}\right)}\left(\boldsymbol{x}^{(t)}\right) \leq \frac{1}{100} \eta_{t}$.

Now, note that Lemmas 8.3 .12 and 8.3 .13 show that $\boldsymbol{e}$ does not change during the addition or removal of an constraint. Hence, we have $\left\|\boldsymbol{e}\left(\boldsymbol{\tau}^{\left(t+\frac{2}{3}\right)}, \boldsymbol{x}^{(t)}\right)\right\|_{\infty} \leq\left\|\boldsymbol{e}\left(\boldsymbol{\tau}^{\left(t+\frac{1}{3}\right)}, \boldsymbol{x}^{(t)}\right)\right\|_{\infty}$. Furthermore, we know the step " $\boldsymbol{\tau}_{i^{(k)}}^{\left(k+\frac{1}{3}\right)}=\psi_{P^{(k)}}\left(\boldsymbol{x}^{(k)}\right)_{i^{(k)}}$ " only decreases $\|\boldsymbol{e}\|_{\infty}$ and hence we have $\left\|\boldsymbol{e}\left(\boldsymbol{\tau}^{\left(t+\frac{2}{3}\right)}, \boldsymbol{x}^{(t)}\right)\right\|_{\infty} \leq$ $\left\|\boldsymbol{e}\left(\boldsymbol{\tau}^{(t)}, \boldsymbol{x}^{(t)}\right)\right\|_{\infty} \leq \frac{c_{e}}{400}$. Thus, we have all the conditions needed for Lemma 8.3.9 and consequently

$$
\delta_{P^{(t+1)}, \boldsymbol{e}\left(\boldsymbol{\tau}^{(t+1)}, \boldsymbol{x}^{(t+1)}\right)}\left(\boldsymbol{x}^{(t+1)}\right) \leq 2\left(1-\frac{1}{64}\right)^{200} \delta_{P^{(t+1)}, \boldsymbol{e}\left(\boldsymbol{\tau}^{\left(t+\frac{2}{3}\right)}, \boldsymbol{x}^{(t)}\right)}\left(\boldsymbol{x}^{(t)}\right) \leq \frac{1}{1000} \eta_{t}
$$

Lemma 8.3 .9 also shows that that $\left\|\frac{s\left(\boldsymbol{x}^{(t+1)}\right)-s\left(\boldsymbol{x}^{(t)}\right)}{s\left(\boldsymbol{x}^{(t)}\right)}\right\|_{2} \leq \frac{1}{10}$ and hence $\psi_{i}\left(\boldsymbol{x}^{(t)}\right) \leq 2 \psi_{i}\left(\boldsymbol{x}^{(t+1)}\right)$ for all $i$.

Therefore, $\eta_{t} \leq 2 \eta_{t+1}$ and thus

$$
\delta_{P^{(t+1)}, e\left(\boldsymbol{\tau}^{(t+1)}, \boldsymbol{x}^{(t+1)}\right)}\left(\boldsymbol{x}^{(t+1)}\right) \leq \frac{\sqrt{c_{e}+\min \left(c_{d}, \mu\left(\boldsymbol{x}^{(t+1)}\right)\right)}}{400}
$$

completing the induction case for $\delta$.
Now, we bound $\|\boldsymbol{e}\|_{\infty}$. Lemma 8.3.12 and 8.3.13 show that $\boldsymbol{e}$ does not change during the addition or removal of an constraint. Hence, $\boldsymbol{e}$ is affected by only the update step " $\tau_{i^{(k)}}^{\left(k+\frac{1}{2}\right)}=\psi_{P^{(k)}}\left(\boldsymbol{x}^{(k)}\right)_{i^{(k)}}$ " and
 shows that $\operatorname{Ee}\left(\boldsymbol{\tau}^{(t+1)}, \boldsymbol{x}^{(t+1)}\right)=\boldsymbol{e}\left(\boldsymbol{\tau}^{\left(t+\frac{2}{3}\right)}, \boldsymbol{x}^{(t)}\right)$ and $\left\|\boldsymbol{e}\left(\boldsymbol{\tau}^{(t+1)}, \boldsymbol{x}^{(t+1)}\right)-\boldsymbol{e}\left(\boldsymbol{\tau}^{\left(t+\frac{2}{3}\right)}, \boldsymbol{x}^{(t)}\right)\right\|_{2} \leq \frac{1}{10} c_{\Delta}$. The goal for the update step is to decrease $\boldsymbol{e}$ by updating $\boldsymbol{\tau}$. In Section 8.4.2, we give a self-contained analysis of the effect of this step as a game. In each round, the vector $\boldsymbol{e}$ is corrupted by some mean 0 and bounded variance noise and the problem is how to update $\boldsymbol{e}$ such that $\|\boldsymbol{e}\|_{\infty}$ is bounded. Theorem 8.4.3 shows that we can do this by setting the $\boldsymbol{e}_{i}=0$ for the almost maximum coordinate in each iteration. This is exactly what the update step is doing. Since our algorithm would run only $O(n \log (n R / \varepsilon))$ many iterations, Theorem 8.4 .3 shows that this strategy guarantees that after the update step, we have

$$
\left\|\boldsymbol{e}\left(\tau^{\left(t+\frac{1}{3}\right)}, \boldsymbol{x}^{(t)}\right)\right\|_{\infty}=O\left(c_{\Delta} \log (n \log (R / \varepsilon))\right)
$$

Now, by our choice of $c_{\Delta}$, we have $\left\|\boldsymbol{e}\left(\boldsymbol{\tau}^{\left(t+\frac{1}{3}\right)}, \boldsymbol{x}^{(t)}\right)\right\|_{\infty} \leq \frac{1}{1000} c_{e}$. Lemma 8.3.12 and 8.3.13 show that $\boldsymbol{e}$ does not change during the addition or removal of an constraint. Hence, we have $\left\|\boldsymbol{e}\left(\boldsymbol{\tau}^{\left(t+\frac{2}{3}\right)}, \boldsymbol{x}^{(t)}\right)\right\|_{\infty} \leq$ $\frac{1}{1000} c_{e}$. Now, we note that again Lemma 8.3 .9 shows $\left\|\boldsymbol{e}\left(\boldsymbol{\tau}^{(t+1)}, \boldsymbol{x}^{(t+1)}\right)-\boldsymbol{e}\left(\boldsymbol{\tau}^{\left(t+\frac{2}{3}\right)}, \boldsymbol{x}^{(t)}\right)\right\|_{2} \leq \frac{1}{10} c_{\Delta} \leq$ $\frac{1}{1000} c_{e}$, and we have $\left\|\boldsymbol{e}\left(\boldsymbol{\tau}^{(t+1)}, \boldsymbol{x}^{(t+1)}\right)\right\|_{\infty} \leq \frac{c_{e}}{400}$. This finishes the induction case for $\|\boldsymbol{e}\|_{\infty}$ and proves this lemma.

Next, we show the number of constraints is always linear to $n$.
Lemma 8.3.17. Throughout our cutting plane method, there are at most $1+\frac{2 n}{c_{d}}$ constraints.
Proof. We only add a constraint if $\min _{i} w_{i} \geq c_{d}$. Since $2 \psi_{i} \geq w_{i}$, we have $\psi_{i} \geq \frac{c_{d}}{2}$ for all $i$. Letting $m$ denote the number of constraints after we add that row, we have $n \geq \sum_{i} \psi_{i} \geq(m-1)\left(c_{d} / 2\right)$.

Using $K \neq \emptyset$ and $K \subset B_{\infty}(R)$, here we show that the points are bounded.
Lemma 8.3.18. During our Cutting Plane Method, for all $k$, we have $\left\|\boldsymbol{x}^{(k)}\right\|_{2} \leq 6 \sqrt{n / \lambda}+2 \sqrt{n} R$.
Proof. By Lemma 8.3.16 and Lemma 8.3 .14 we know that $\left\|\boldsymbol{x}^{(k)}\right\|_{2}^{2} \leq 4 \lambda^{-1}\left(m c_{e}+n\right)+2\|\boldsymbol{y}\|_{2}^{2}$ for any $\boldsymbol{y} \in P^{(k)}$. Since our method never cuts out any point in $K$ and since $K$ is nonempty, there is some $\boldsymbol{y} \in K \subset P^{(k)}$. Since $K \subset B_{\infty}(R)$, we have $\|\boldsymbol{y}\|_{2}^{2} \leq n R^{2}$. Furthermore, by Lemma 8.3.17 we have that $m c_{e} \leq c_{e}+2 n \leq 3 n$ yielding the result.

Lemma 8.3.19. $s_{i}\left(\boldsymbol{x}^{(k)}\right) \leq 12 \sqrt{n / \lambda}+4 \sqrt{n} R+\sqrt{\frac{1}{c_{a} \lambda}}$ for all $i$ and $k$ in the our cutting plane method.
Proof. Let $\boldsymbol{x}^{(j)}$ be the current point at the time that the constraint corresponding to $s_{i}$, denoted $\left\{\boldsymbol{x}: \boldsymbol{a}_{i}^{T} \boldsymbol{x} \geq a_{i}^{T} \boldsymbol{x}^{(j)}-s_{i}\left(\boldsymbol{x}^{(j)}\right)\right\}$, was added. Clearly

$$
s_{i}\left(\boldsymbol{x}^{(k)}\right)=\boldsymbol{a}_{i}^{T} \boldsymbol{x}^{(k)}-a_{i}^{T} \boldsymbol{x}^{(j)}+s_{i}\left(\boldsymbol{x}^{(j)}\right) \leq\left\|\boldsymbol{a}_{i}\right\| \cdot\left\|\boldsymbol{x}^{(k)}\right\|+\left|\boldsymbol{a}_{i}^{T} \boldsymbol{x}^{(j)}-s_{i}\left(\boldsymbol{x}^{(j)}\right)\right|
$$

On the one hand, if the constraint for $s_{i}$ comes from the initial symmetric polytope $P^{(0)}=B_{\infty}(R)$, we know $\left|\boldsymbol{a}^{T} \boldsymbol{x}^{(j)}-\boldsymbol{s}_{i}\left(\boldsymbol{x}^{(j)}\right)\right| \leq R$. On the other hand, if the constraint was added later then we know that

$$
s_{i}\left(\boldsymbol{x}^{(j)}\right)=c_{a}^{-1 / 2} \sqrt{\boldsymbol{a}^{T}\left(\boldsymbol{A}^{T} \boldsymbol{S}_{\boldsymbol{x}^{(j)}}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{a}} \leq\left(c_{a} \lambda\right)^{-1 / 2}
$$

and $\left|\boldsymbol{a}^{T} \boldsymbol{x}^{(j)}-s_{i}\left(\boldsymbol{x}^{(j)}\right)\right| \leq\left\|\boldsymbol{a}_{i}\right\| \cdot\left\|\boldsymbol{x}^{(j)}\right\|+\left|s_{i}\left(\boldsymbol{x}^{(j)}\right)\right|$. Since $\left\|\boldsymbol{a}_{i}\right\|_{2}=1$ by design and $\left\|\boldsymbol{x}^{(j)}\right\|_{2}$ and $\left\|\boldsymbol{x}^{(k)}\right\|_{2}$ are upper bounded by $6 \sqrt{n / \lambda}+2 \sqrt{n} R$ by Lemma 8.3.18, in either case the result follows.

Now, we have everything we need to prove that the potential function is increasing in expectation.

Lemma 8.3.20. Under the assumptions of Lemma 8.3.16 if $\lambda=\frac{1}{c_{a} R^{2}}, c_{e}=\frac{c_{d}}{4 \ln (6 n R / \varepsilon)}$, and $24 c_{d} \leq$ $c_{a} \leq \frac{1}{1000}$ then for all $k$ we have

$$
\mathbb{E} p_{e\left(\boldsymbol{\tau}^{(k+1)}, \boldsymbol{x}^{(k+1)}\right)}\left(\boldsymbol{x}^{(k+1)}\right) \geq p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{(k)}, \boldsymbol{x}^{(k)}\right)}\left(\boldsymbol{x}^{(k)}\right)-c_{d}+\ln (1+\beta)
$$

where $\beta=c_{a}$ for the case of adding a constraint and $\beta=-c_{d}$ for the case of removal.

Proof. Note that there are three places which affect the function value, namely the update step for $\tau^{\left(k+\frac{1}{3}\right)}$, the addition/removal of constraints, and the centering step. We bound the effect of each separately.

First, for the update step, we have

$$
p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{\left(k+\frac{1}{3}\right)}, \boldsymbol{x}^{(k)}\right)}\left(\boldsymbol{x}^{(k)}\right)=\left(c+e_{i(k)}\right) \log \left(s_{i^{(k)}}\left(\boldsymbol{x}^{(k)}\right) / R\right)+p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{(k)}, \boldsymbol{x}^{(k)}\right)}\left(\boldsymbol{x}^{(k)}\right) .
$$

Lemma 8.3.19, the termination condition, $\lambda=\frac{1}{c_{a} R^{2}}$ and $c_{a}<\frac{1}{1000}$ ensure that

$$
\begin{equation*}
\varepsilon \leq s_{i(k)}\left(\boldsymbol{x}^{(k)}\right) \leq 12 \sqrt{n / \lambda}+4 \sqrt{n} R+\sqrt{\frac{1}{c_{a} \lambda}} \leq 6 \sqrt{n} R \tag{8.19}
\end{equation*}
$$

and Lemma 8.3.16 shows that $\left|e_{i(k)}\right| \leq c_{e}$. Hence, we have

$$
p_{e\left(\boldsymbol{\tau}^{\left(k+\frac{1}{3}\right)}, \boldsymbol{x}^{(k)}\right)}\left(\boldsymbol{x}^{(k)}\right) \geq p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{(k)}, \boldsymbol{x}^{(k)}\right)}\left(\boldsymbol{x}^{(k)}\right)+2 c_{e} \log (\varepsilon / R) .
$$

For the addition step, Lemma 8.3.12 shows that

$$
\begin{aligned}
p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{\left(k+\frac{2}{3}\right)}, \boldsymbol{x}^{(k)}\right)}\left(\boldsymbol{x}^{(k)}\right) & =p_{e\left(\boldsymbol{\tau}^{\left(k+\frac{1}{3}\right)}, \boldsymbol{x}^{(k)}\right)}\left(\boldsymbol{x}^{(k)}\right)-c_{e} \ln \left(s(\boldsymbol{x})_{m+1} / R\right)+\ln \left(1+c_{a}\right) \\
& \geq p_{e\left(\boldsymbol{\tau}^{\left(k+\frac{1}{3}\right)}, \boldsymbol{x}^{(k)}\right)}\left(\boldsymbol{x}^{(k)}\right)-c_{e} \log (6 \sqrt{n})+\ln \left(1+c_{a}\right)
\end{aligned}
$$

and for the removal step, Lemma 8.3.13 and $\left|e_{i}\right| \leq c_{e}$ shows that

$$
\begin{aligned}
p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{\left(k+\frac{2}{3}\right)}, \boldsymbol{x}^{(k))}\right)}\left(\boldsymbol{x}^{(k)}\right) & \geq p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{\left(k+\frac{1}{3}\right)}, \boldsymbol{x}^{(k)}\right)}\left(\boldsymbol{x}^{(k)}\right)+\left[c_{e}+e_{P}(\boldsymbol{\tau}, \boldsymbol{x})_{m}\right] \ln \left(s(\boldsymbol{x})_{m} / R\right)+\ln \left(1-c_{d}\right) \\
& \geq p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{\left(k+\frac{1}{3}\right)}, \boldsymbol{x}^{(k)}\right)}\left(\boldsymbol{x}^{(k)}\right)-2 c_{e} \log (R / \varepsilon)+\ln \left(1-c_{d}\right)
\end{aligned}
$$

After the addition or removal of a constraint, Lemma 8.3.16 shows that

$$
\delta_{P^{(k)}, \boldsymbol{e}\left(\boldsymbol{\tau}^{\left.\left(k+\frac{2}{3}\right), \boldsymbol{x}^{(k)}\right)}\right.}\left(\boldsymbol{x}^{(k)}\right) \leq \frac{1}{100} \sqrt{c_{e}+\min \left(\mu\left(\boldsymbol{x}^{(k)}\right), c_{d}\right)}
$$

and therefore Lemma 8.3.9 and $c_{e} \leq c_{d}$ show that

$$
\begin{aligned}
\mathbb{E} p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{(k+1)}, \boldsymbol{x}^{(k+1)}\right)}\left(\boldsymbol{x}^{(k+1)}\right) & \geq p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{\left(k+\frac{2}{3}\right)}, \boldsymbol{x}^{(k)}\right)}\left(\boldsymbol{x}^{(k)}\right)-8\left(\frac{\sqrt{c_{e}+\min \left(\mu\left(\boldsymbol{x}^{(k)}\right), c_{d}\right)}}{100}\right)^{2} \\
& \geq p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{\left(k+\frac{2}{3}\right)}, \boldsymbol{x}^{(k)}\right)}\left(\boldsymbol{x}^{(k)}\right)-\frac{c_{d}}{625}
\end{aligned}
$$

Combining them with $c_{e}=\frac{c_{d}}{4 \ln (6 n R / \varepsilon)}$, we have

$$
\begin{aligned}
\mathbb{E} p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{(k+1)}, \boldsymbol{x}^{(k+1)}\right)}\left(\boldsymbol{x}^{(k+1)}\right) & \geq p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{(k)}, \boldsymbol{x}^{(k)}\right)}\left(\boldsymbol{x}^{(k)}\right)-2 c_{e} \log (6 n R / \varepsilon)-\frac{c_{d}}{625}+\ln (1+\beta) \\
& \geq p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{(k)}, \boldsymbol{x}^{(k)}\right)}\left(\boldsymbol{x}^{(k)}\right)-c_{d}+\ln (1+\beta)
\end{aligned}
$$

where $\beta=c_{a}$ for the case of addition and $\beta=-c_{d}$ for the case of removal.

Theorem 8.3.21. For $c_{a}=\frac{1}{10^{10}}, c_{d}=\frac{1}{10^{12}}, c_{e}=\frac{c_{d}}{4 \ln (6 n R / \varepsilon)}, c_{\Delta}=\frac{C c_{e}}{\log (n \log (R / \varepsilon))}$ and $\lambda=\frac{1}{c_{a} R^{2}}$ for some small enough universal constant $C$, then we have

$$
\mathbb{E} p_{e\left(\boldsymbol{\tau}^{(k+1)}, \boldsymbol{x}^{(k+1)}\right)}\left(\boldsymbol{x}^{(k+1)}\right) \geq p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{(k)}, \boldsymbol{x}^{(k)}\right)}\left(\boldsymbol{x}^{(k)}\right)-\frac{1}{10^{11}}+\frac{9 \beta}{10^{11}}
$$

where $\beta=1$ for the case of addition and $\beta=0$ for the case of removal.
Proof. It is easy to see that these parameters satisfy the requirements of Lemma 8.3.20.

### 8.3.5 Guarantees of the Algorithm

In this section we put everything together to prove Theorem 8.3.26, the main result of this section, providing the guarantees of our cutting plane method.

For the remainder of this section we assume that $c_{a}=\frac{1}{10^{10}}, c_{d}=\frac{1}{10^{12}}, c_{e}=\frac{c_{d}}{4 \ln (6 n R / \varepsilon)}, c_{\Delta}=$ $\frac{C c_{e}}{\log (n \ln (R / \varepsilon))}$ and $\lambda=\frac{1}{c_{a} R^{2}}$. Consequently, throughout the algorithm we have

$$
\begin{equation*}
\|\boldsymbol{x}\|_{2} \leq 6 \sqrt{n / \lambda}+2 \sqrt{n} R \leq 3 \sqrt{n} R \tag{8.20}
\end{equation*}
$$

Lemma 8.3.22. If $s_{i}\left(\boldsymbol{x}^{(k)}\right)<\varepsilon$ for some $i$ and $k$ during our Cutting Plane Method then

$$
\max _{\boldsymbol{y} \in P^{(k)} \cap B_{\infty}(R)}\left\langle\boldsymbol{a}_{i}, \boldsymbol{y}\right\rangle-\min _{\boldsymbol{y} \in P^{(k)} \cap B_{\infty}(R)}\left\langle\boldsymbol{a}_{i}, \boldsymbol{y}\right\rangle \leq \frac{8 n \varepsilon}{c_{a} c_{e}} .
$$

Proof. Let $\boldsymbol{y} \in P^{(k)} \cap B_{\infty}(R)$ be arbitrary. Since $\boldsymbol{y} \in B_{\infty}(R)$ clearly $\|\boldsymbol{y}\|_{2}^{2} \leq n R^{2}$. Furthermore, by Lemma 8.3.17 and the choice of parameters $m c_{e}+n \leq 3 n$. Consequently, by Lemma 8.3.14 and the fact that $\lambda=\frac{1}{c_{a} R^{2}}$ and $c_{a}<1$ we have

$$
\|\boldsymbol{x}-\boldsymbol{y}\|_{\boldsymbol{H}(\boldsymbol{x})} \leq \frac{12 m c_{e}+6 n+2 \lambda\|\boldsymbol{y}\|_{2}^{2}}{\sqrt{c_{e}+\mu(\boldsymbol{x})}} \leq \frac{30 n+2 \frac{n}{c_{a}}}{\sqrt{c_{e}+\mu(\boldsymbol{x})}} \leq \frac{4 n}{c_{a} \sqrt{c_{e}}}
$$

and therefore

$$
\left\|\boldsymbol{S}_{x^{(k)}}^{-1}\left(s\left(\boldsymbol{x}^{(k)}\right)-s(\boldsymbol{y})\right)\right\|_{\infty} \leq \frac{1}{\sqrt{c_{e}}}\left\|\boldsymbol{S}_{x^{(k)}}^{-1}\left(s\left(\boldsymbol{x}^{(k)}\right)-s(\boldsymbol{y})\right)\right\|_{c_{e} \boldsymbol{I}+\boldsymbol{\Psi}} \leq \frac{4 n}{c_{a} c_{e}}
$$

Consequently, we have $\left(1-\frac{4 n}{c_{a} c_{e}}\right) s_{i}\left(\boldsymbol{x}^{(k)}\right) \leq s_{i}(\boldsymbol{y}) \leq\left(1+\frac{4 n}{c_{a} c_{e}}\right) s_{i}\left(\boldsymbol{x}^{(k)}\right)$ for all $\boldsymbol{y} \in P^{(k)} \cap B_{\infty}(R)$.

Now let us show how to compute a proof (or certificate) that the feasible region has small width on the direction $\boldsymbol{a}_{i}$.

Lemma 8.3.23. Suppose that during some iteration $k$ for $i=\arg \min _{j} s_{j}\left(\boldsymbol{x}^{(k)}\right)$ we have $s_{i}\left(\boldsymbol{x}^{(k)}\right) \leq \varepsilon$. Let $\left(\boldsymbol{x}_{*}, \boldsymbol{\tau}_{*}\right)=$ Centering $\left(\boldsymbol{x}^{(k)}, \boldsymbol{\tau}^{(k)}, 64 \log (2 R / \varepsilon), c_{\Delta}\right)$ where $\boldsymbol{\tau}^{(k)}$ is the $\tau$ at that point in the algorithm and let

$$
\boldsymbol{a}^{*}=\sum_{j \neq i} t_{j} \boldsymbol{a}_{j} \text { where } t_{j}=\left(\frac{s\left(\boldsymbol{x}_{*}\right)_{i}}{s\left(\boldsymbol{x}_{*}\right)_{j}}\right)\left(\frac{c_{e}+e_{j}\left(\boldsymbol{x}_{*}, \boldsymbol{\tau}_{*}\right)+\psi_{j}\left(\boldsymbol{x}_{*}\right)}{c_{e}+e_{i}\left(\boldsymbol{x}_{*}, \boldsymbol{\tau}_{*}\right)+\psi_{i}\left(\boldsymbol{x}_{*}\right)}\right) .
$$

Then, we have that $\left\|\boldsymbol{a}_{i}+\boldsymbol{a}^{*}\right\|_{2} \leq \frac{8 \sqrt{n} \varepsilon}{c_{a} c_{e} R}$ and $t_{j} \geq 0$ for all $j$. Furthermore, we have

$$
\left(\sum_{j \neq i}^{O(n)} t_{j} a_{j}\right)^{T} x_{*}-\sum_{j \neq i}^{O(n)} t_{j} b_{j} \leq \frac{4 n \varepsilon}{c_{e}} .
$$

Proof. By Lemma 8.3.9 and Lemma 8.3.16 we know that $\boldsymbol{e}\left(\boldsymbol{x}_{*}, \boldsymbol{\tau}_{*}\right) \leq \frac{1}{2} c_{e}$ and $\delta_{\boldsymbol{e}\left(\boldsymbol{x}_{*}, \boldsymbol{\tau}_{*}\right)} \leq \frac{\varepsilon}{R} \sqrt{c_{e}+\mu\left(\boldsymbol{x}_{*}\right)}$. Since $\boldsymbol{e}\left(\boldsymbol{x}_{*}, \boldsymbol{\tau}_{*}\right) \leq \frac{1}{2} c_{e}$, we have $t_{j} \geq 0$ for all $j$. Furthermore, by Lemma 8.3.15 and (8.20), we then have that with high probability in $n$

$$
\begin{aligned}
\left\|\boldsymbol{a}_{i}+\boldsymbol{a}^{*}\right\|_{2} & \leq \frac{2 \varepsilon}{\left(c_{e}+\mu\left(\boldsymbol{x}_{*}\right)\right)}\left[\lambda\left\|\boldsymbol{x}_{*}\right\|_{2}+\delta_{e}\left(\boldsymbol{x}_{*}\right) \sqrt{\frac{m c_{e}+n}{\varepsilon^{2}}+\lambda}\right] \\
& \leq \frac{2 \varepsilon}{c_{e}}\left[\frac{1}{c_{a} R^{2}}(3 \sqrt{n} R)+\frac{2 \varepsilon}{R} \sqrt{\frac{3 n}{\varepsilon^{2}}+\frac{1}{c_{a} R^{2}}}\right] \\
& \leq \frac{2 \varepsilon}{c_{e}}\left[\frac{3 \sqrt{n}}{c_{a} R}+\frac{2 \sqrt{3 n}}{R}+\frac{2}{\sqrt{c_{a}} R}\right] \leq \frac{8 \sqrt{n} \varepsilon}{c_{a} c_{e} R} .
\end{aligned}
$$

By Lemma 8.3.16 we know that $\boldsymbol{e}\left(\boldsymbol{x}_{*}, \boldsymbol{\tau}_{*}\right) \leq \frac{1}{2} c_{e}$ and hence

$$
\begin{aligned}
\left(\sum_{j \neq i}^{O(n)} t_{j} a_{j}\right)^{T} \boldsymbol{x}_{*}-\sum_{j \neq i}^{O(n)} t_{j} b_{j} & =\sum_{j \neq i}^{O(n)} t_{j} s\left(\boldsymbol{x}_{*}\right)_{j}=s_{i}\left(\boldsymbol{x}_{*}\right) \sum_{j \neq i}^{O(n)}\left(\frac{c_{e}+e_{j}\left(\boldsymbol{x}_{*}, \boldsymbol{\tau}_{*}\right)+\psi_{j}\left(\boldsymbol{x}_{*}\right)}{c_{e}+e_{i}\left(\boldsymbol{x}_{*}, \boldsymbol{\tau}_{*}\right)+\psi_{i}\left(\boldsymbol{x}_{*}\right)}\right) \\
& \leq s_{i}\left(\boldsymbol{x}_{*}\right) \sum_{j \neq i}^{O(n)}\left(\frac{\frac{3}{2} c_{e}+\psi_{j}\left(\boldsymbol{x}_{*}\right)}{\frac{1}{2} c_{e}+\psi_{i}\left(\boldsymbol{x}_{*}\right)}\right) \leq s_{i}\left(\boldsymbol{x}_{*}\right) \frac{3 m c_{e}+2 n}{c_{e}} \\
& \leq \frac{3 n}{c_{e}} s_{i}\left(\boldsymbol{x}_{*}\right) \leq \frac{4 n \varepsilon}{c_{e}}
\end{aligned}
$$

where the last line follows the fact that $s_{i}\left(\boldsymbol{x}_{*}\right) \leq 1.1 s_{i}\left(\boldsymbol{x}^{(k)}\right) \leq 1.1 \varepsilon$ (Lemma 8.3.9).

Lemma 8.3.24. During our Cutting Plane Method, if $p_{\boldsymbol{e}}\left(\boldsymbol{x}^{(k)}\right) \geq 2 n \log \left(\frac{n R}{c_{a} \varepsilon}\right)+\frac{6 n}{c_{a}}$, then we have $s_{i}\left(\boldsymbol{x}^{(k)}\right) \leq \varepsilon$ for some $i$.

Proof. Recall that

$$
p_{\boldsymbol{e}}\left(\boldsymbol{x}^{(k)}\right)=-\sum_{i \in[m]}\left(c_{e}+e_{i}\right) \log \left(s_{i}\left(\boldsymbol{x}^{(k)}\right)_{i} / R\right)+\frac{1}{2} \log \operatorname{det}\left(R^{2}\left(\boldsymbol{A}^{T} \boldsymbol{S}_{x^{(k)}}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)\right)+\frac{\lambda}{2}\left\|\boldsymbol{x}^{(k)}\right\|_{2}^{2}
$$

Using $\left\|\boldsymbol{x}^{(k)}\right\| \leq 3 \sqrt{n} R$ (8.20) and $\lambda=\frac{1}{c_{a} R^{2}}$, we have

$$
p_{\boldsymbol{e}}\left(\boldsymbol{x}^{(k)}\right) \leq-\sum_{i \in[m]}\left(c_{e}+e_{i}\right) \log \left(s\left(\boldsymbol{x}^{(k)}\right)_{i} / R\right)+\frac{1}{2} \log \operatorname{det}\left(R^{2}\left(\boldsymbol{A}^{T} \boldsymbol{S}_{x^{(k)}}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)\right)+\frac{5 n}{c_{a}}
$$

Next, we note that $\left\|e_{i}\right\|_{\infty} \leq c_{e} \leq \frac{1}{4 \ln (6 n R / \varepsilon)}$ and $s_{i}\left(\boldsymbol{x}^{(k)}\right) \leq 12 \sqrt{n / \lambda}+4 \sqrt{n} R+\sqrt{\frac{1}{c_{a} \lambda}} \leq 6 \sqrt{n} R$ (Lemma 8.3.19). Hence, we have

$$
p_{\boldsymbol{e}}\left(\boldsymbol{x}^{(k)}\right) \leq \frac{1}{2} \log \operatorname{det}\left(R^{2}\left(\boldsymbol{A}^{T} \boldsymbol{S}_{x^{(k)}}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)\right)+\frac{6 n}{c_{a}}-\frac{n}{2} \log \left(\frac{\min _{i} s_{i}}{R}\right)
$$

We prove $s_{i}\left(\boldsymbol{x}^{(k)}\right) \leq \varepsilon$ by contradiction. Since $p_{e}\left(\boldsymbol{x}^{(k)}\right) \geq 2 n \log \left(\frac{n R}{c_{a} \varepsilon}\right)+\frac{6 n}{c_{a}}$ and $s_{i}\left(\boldsymbol{x}^{(k)}\right)>\varepsilon$ for all $i$, we have that $\frac{1}{2} \log \operatorname{det}\left(R^{2}\left(\boldsymbol{A}^{T} \boldsymbol{S}_{x^{(k)}}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)\right) \geq n \log \left(\frac{n R}{c_{a} \varepsilon}\right)$. Using $\varepsilon<R$, we have that

$$
\sum_{i} \log \lambda_{i}\left(R^{2}\left(\boldsymbol{A}^{T} \boldsymbol{S}_{x^{(k)}}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)\right) \geq n \log \left(\frac{n^{2} R^{2}}{c_{a}^{2} \varepsilon^{2}}\right) \geq n \log \left(\frac{n R^{2}}{2 c_{a}^{2} \varepsilon^{2}}+R^{2} \lambda\right)
$$

Therefore, we have $\log \lambda_{\max }\left(R^{2}\left(\boldsymbol{A}^{T} \boldsymbol{S}_{x^{(k)}}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)\right) \geq \log \left(\frac{n R^{2}}{2 c_{a}^{2} \varepsilon^{2}}+R^{2} \lambda\right)$. Hence, we have some unit vector $\boldsymbol{v}$ such that $\boldsymbol{v} \boldsymbol{A}^{T} \boldsymbol{S}_{x}^{-2} \boldsymbol{A} \boldsymbol{v}+\lambda \boldsymbol{v}^{T} \boldsymbol{v} \geq \frac{n}{2 c_{a}^{2} \varepsilon^{2}}+\lambda$. Thus,

$$
\sum_{i} \frac{(\boldsymbol{A} \boldsymbol{v})_{i}^{2}}{s\left(\boldsymbol{x}^{(k)}\right)_{i}^{2}} \geq \frac{n}{2 c_{a}^{2} \varepsilon^{2}}
$$

Lemma 8.3 .17 shows that the number of constraints is bounded by $1+2 n / c_{d} \leq n /\left(2 c_{a}^{2}\right)$ and hence there is some $i$ such that $\frac{(\boldsymbol{A} \boldsymbol{v})_{i}^{2}}{s\left(\boldsymbol{x}^{(k)}\right)_{i}^{2}} \geq \frac{1}{\varepsilon^{2}}$. Since $\boldsymbol{a}_{i}$ and $\boldsymbol{v}$ are unit vectors, we have $1 \geq\left\langle\boldsymbol{a}_{i}, \boldsymbol{v}\right\rangle^{2} \geq s\left(\boldsymbol{x}^{(k)}\right)_{i}^{2} / \varepsilon^{2}$ and hence $s\left(\boldsymbol{x}^{(k)}\right)_{i} \leq \varepsilon($ contradiction $)$.

Lemma 8.3.25. With constant probability, the algorithm ends in $10^{24} n \log \left(\frac{n R}{\varepsilon}\right)$ iterations. ${ }^{1}$

Proof. Theorem 8.3.21 shows that for all $k$

$$
\begin{equation*}
\mathbb{E} p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{(k+1)}, \boldsymbol{x}^{(k+1)}\right)}\left(\boldsymbol{x}^{(k+1)}\right) \geq p_{e\left(\boldsymbol{\tau}^{(k)}, \boldsymbol{x}^{(k)}\right)}\left(\boldsymbol{x}^{(k)}\right)-\frac{1}{10^{11}}+\frac{9 \beta}{10^{11}} \tag{8.21}
\end{equation*}
$$

where $\beta=1$ for the case of adding a constraint and $\beta=0$ for the case of removing a constraint. Now, for all $t$ consider the random variable

$$
\boldsymbol{X}_{t}=p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{(t)}, \boldsymbol{x}^{(t)}\right)}\left(\boldsymbol{x}^{(t)}\right)-\frac{4.5 m^{(t)}}{10^{11}}-\frac{3.5 t}{10^{11}}
$$

where $m^{(t)}$ is the number of constraints in iteration $t$ of the algorithm. Then, since $m^{(t+1)}=m^{(t)}-$ $1+2 \beta$, (8.21) shows that

$$
\begin{aligned}
\mathbf{E} \boldsymbol{X}_{t+1} & \geq p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{(t)}, \boldsymbol{x}^{(t)}\right)}\left(\boldsymbol{x}^{(t)}\right)-\frac{1}{10^{11}}+\frac{9 \beta}{10^{11}}-\frac{4.5 m^{(t+1)}}{10^{11}}-\frac{3.5(t+1)}{10^{11}} \\
& =\boldsymbol{X}_{t}-\frac{1}{10^{11}}+\frac{9 \beta}{10^{11}}-\frac{4.5(-1+2 \beta)}{10^{11}}-\frac{3.5}{10^{11}}=\boldsymbol{X}_{t}
\end{aligned}
$$

[^17]Hence, it is a sub-martingale. Let $T$ be the iteration the algorithm outputs $\boldsymbol{x}^{(k)}$ or $P^{(k)}$. Optional stopping theorem shows that

$$
\begin{equation*}
\mathbb{E} \boldsymbol{X}_{\min (T, t)} \geq \mathbb{E} \boldsymbol{X}_{0} \tag{8.22}
\end{equation*}
$$

Since the polytope is $B_{\infty}(R)$, we have

$$
\begin{aligned}
p_{\mathbf{0}}(\mathbf{0}) & =-\sum_{i \in\left[m^{(0)}\right]} c_{e} \log \left(s_{i}(\mathbf{0}) / R\right)+\frac{1}{2} \log \operatorname{det}\left(R^{2}\left(\boldsymbol{A}^{T} \boldsymbol{S}_{0}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)\right)+\frac{\lambda}{2}\|\mathbf{0}\|_{2}^{2} \\
& \geq \frac{1}{2} \log \operatorname{det}\left(R^{2}\left(\frac{2}{R^{2}} \boldsymbol{I}+\frac{1}{c_{a} R^{2}} \boldsymbol{I}\right)\right) \\
& \geq-15 n
\end{aligned}
$$

where we used $c_{a}=1 / 10^{10}$ on the last line. Hence, we have

$$
\begin{aligned}
\boldsymbol{X}_{0} & \geq-15 n-\frac{4.5 m^{(0)}}{10^{11}} \\
& \geq-20 n
\end{aligned}
$$

Therefore, (8.22) shows that for all $t$ we have

$$
\begin{align*}
-20 n & \leq \mathbb{E} \boldsymbol{X}_{\min (T, t)} \\
& =p \mathbb{E}\left[\boldsymbol{X}_{\min (T, t)} \mid T<t\right]+(1-p) \mathbb{E}\left[\boldsymbol{X}_{\min (T, t)} \mid T \geq t\right] \tag{8.23}
\end{align*}
$$

where $p \stackrel{\text { def }}{=} \mathbb{P}(T<t)$.

Note that

$$
\begin{aligned}
\mathbb{E}\left[\boldsymbol{X}_{\min (T, t)} \mid T \geq t\right] & \leq \mathbb{E}\left[p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{(t)}, \boldsymbol{x}^{(t)}\right)}\left(\boldsymbol{x}^{(t)}\right) \mid T \geq t\right]-\frac{4.5 m^{(t)}}{10^{11}}-\frac{3.5 t}{10^{11}} \\
& \leq \mathbb{E}\left[p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{(t)}, \boldsymbol{x}^{(t)}\right)}\left(\boldsymbol{x}^{(t)}\right) \mid T \geq t\right]-\frac{3.5 t}{10^{11}}
\end{aligned}
$$

Furthermore, by Lemma 8.3.24 we know that when $p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{(t)}, \boldsymbol{x}^{(t)}\right)}\left(\boldsymbol{x}^{(t)}\right) \geq 2 n \log \left(\frac{n R}{c_{a} \varepsilon}\right)+\frac{6 n}{c_{a}}$, there is a slack that is too small and the algorithm terminates. Hence, we have

$$
\mathbf{E}\left[\boldsymbol{X}_{\min (T, t)} \mid T \geq t\right] \leq 2 n \log \left(\frac{n R}{c_{a} \varepsilon}\right)+\frac{6 n}{c_{a}}-\frac{3.5 t}{10^{11}}
$$

The proof of Lemma 8.3.16 shows that the function value does not change by more than 1 in one iteration by changing $\boldsymbol{x}$ and can change by at most $m c_{e} \log \left(\frac{3 n R}{\varepsilon}\right)$ by changing $\tau$. Since by Lemma 8.3.17 we know that $m \leq 1+\frac{2 n}{c_{a}}$ and $c_{e}=\frac{c_{d}}{4 \ln (6 n R / \varepsilon)}$, we have that $p_{e}(\boldsymbol{x}) \leq 2 n \log \left(\frac{n R}{c_{a} \varepsilon}\right)+\frac{7 n}{c_{a}}$ throughout the execution of the algorithm. Therefore, we have

$$
\mathbf{E}\left[\boldsymbol{X}_{\min (T, t)} \mid T \leq t\right] \leq \mathbb{E}_{T<t} p_{\boldsymbol{e}\left(\boldsymbol{\tau}^{(t)}, \boldsymbol{x}^{(t)}\right)}\left(\boldsymbol{x}^{(t)}\right) \leq 2 n \log \left(\frac{n R}{c_{a} \varepsilon}\right)+\frac{7 n}{c_{a}}
$$

Therefore, (8.23) shows that

$$
-20 n \leq 2 n \log \left(\frac{n R}{c_{a} \varepsilon}\right)+\frac{7 n}{c_{a}}-(1-p) \frac{3.5 t}{10^{11}}
$$

Hence, we have

$$
\begin{aligned}
(1-p) \frac{3.5 t}{10^{11}} & \leq 2 n \log \left(\frac{n R}{c_{a} \varepsilon}\right)+\frac{7 n}{c_{a}}+20 n \\
& \leq 2 n \log \left(\frac{n R}{c_{a} \varepsilon}\right)+\frac{7 n}{c_{a}}+20 n \\
& =2 n \log \left(\frac{R n}{c_{a} \varepsilon}\right)+8 \cdot 10^{10} n
\end{aligned}
$$

Thus, we have

$$
\mathbb{P}(T<t)=p \geq 1-\frac{1}{t}\left(10^{11} n \log \left(\frac{n R}{\varepsilon}\right)+10^{22} n\right)
$$

Now, we gather all the result as follows:
Theorem 8.3.26 (Our Cutting Plane Method). Let $K \subseteq \mathbb{R}^{n}$ be a non-empty set contained in a box of radius $R$, i.e. $K \subseteq B_{\infty}(R)$. For any $\varepsilon \in(0, R)$ in expected time $O\left(n \mathrm{SO}_{\Omega(\varepsilon / \sqrt{n})}(K) \log (n R / \varepsilon)+\right.$ $\left.n^{3} \log ^{O(1)}(n R / \varepsilon)\right)$ our cutting plane method either outputs $\boldsymbol{x} \in K$ or finds a polytope $P=\{\boldsymbol{x}: \boldsymbol{A} \boldsymbol{x} \geq$ $\boldsymbol{b}\} \supseteq K$ such that

1. Phas $O(n)$ many constraints (i.e. $\boldsymbol{A} \in \mathbb{R}^{O(n) \times n}$ and $\left.\boldsymbol{b} \in \mathbb{R}^{O(n)}\right)$.
2. Each constraint of $P$ is either an initial constraint from $B_{\infty}(R)$ or of the form $\langle\boldsymbol{a}, \boldsymbol{x}\rangle \geq b-\delta$ where $\langle\boldsymbol{a}, \boldsymbol{x}\rangle \geq b$ is a normalized hyperplane (i.e. $\|\boldsymbol{a}\|_{2}=1$ ) returned by the separation oracle and $\delta=\Omega\left(\frac{\varepsilon}{\sqrt{n}}\right)$.
3. The polytope $P$ has small width with respect to some direction $\boldsymbol{a}_{1}$ given by one of the constraints, i.e.

$$
\max _{\boldsymbol{y} \in P \cap B_{\infty}(R)}\left\langle\boldsymbol{a}_{1}, \boldsymbol{y}\right\rangle-\min _{\boldsymbol{y} \in P \cap B_{\infty}(R)}\left\langle\boldsymbol{a}_{1}, \boldsymbol{y}\right\rangle \leq O(n \varepsilon \ln (n R / \varepsilon))
$$

4. Furthermore, the algorithm produces a proof of the fact above involving convex combination of the constraints, namely, non-negatives $t_{2}, \ldots, t_{O(n)}$ and $\boldsymbol{x} \in P$ such that
(a) $\|\boldsymbol{x}\|_{2}=O(\sqrt{n} R)$,
(b) $\left\|\boldsymbol{a}_{1}+\sum_{i=2}^{O(n)} t_{i} \boldsymbol{a}_{i}\right\|_{2}=O\left(\frac{\varepsilon}{R} \sqrt{n} \log (n R / \varepsilon)\right)$,
(c) $\boldsymbol{a}_{1}^{T} \boldsymbol{x}-\boldsymbol{b}_{1} \leq \varepsilon$,
(d) $\left(\sum_{i=2}^{O(n)} t_{i} a_{i}\right)^{T} \boldsymbol{x}-\sum_{i=2}^{O(n)} t_{i} b_{i} \leq O(n \varepsilon \log (n R / \varepsilon))$.

Proof. Our algorithm either finds $\boldsymbol{x} \in K$ or we have $s_{i}\left(\boldsymbol{x}^{(k)}\right)<\varepsilon$. When $s_{i}\left(\boldsymbol{x}^{(k)}\right)<\varepsilon$, we apply Lemma 8.3.23 to construct the polytope $P$ and the linear combination $\sum_{i=2}^{O(n)} t_{i} \boldsymbol{a}_{i}$.

Notice that each iteration of our algorithm needs to solve constant number of linear systems and implements the sampling step to find $\boldsymbol{\Delta}^{(k)} \in \mathbb{R}^{n}$ s.t. $\mathbf{E}\left[\boldsymbol{\Delta}^{(k)}\right]=\boldsymbol{\psi}\left(\boldsymbol{x}^{(k)}\right)-\boldsymbol{\psi}\left(\boldsymbol{x}^{(k-1)}\right)$. Theorem 8.4.2 shows how to do the sampling in $\tilde{O}(1)$ many linear systems. Hence, in total, each iterations needs to solve $\tilde{O}(1)$ many linear systems plus nearly linear work. To output the proof for (4), we use Lemma 8.3.23.

Note that the linear systems the whole algorithm need to solve is of the form

$$
\left(\boldsymbol{A}^{T} \boldsymbol{S}_{x}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{x}=\boldsymbol{y} .
$$

where the matrix $\boldsymbol{A}^{T} \boldsymbol{S}_{x}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}$ can be written as $\overline{\boldsymbol{A}}^{T} \boldsymbol{D} \overline{\boldsymbol{A}}$ for the matrix $\overline{\boldsymbol{A}}=[\boldsymbol{A} \boldsymbol{I}]$ and diagonal matrix

$$
\boldsymbol{D}=\left[\begin{array}{cc}
\boldsymbol{S}^{-2} & \mathbf{0} \\
\mathbf{0} & \lambda \boldsymbol{I}
\end{array}\right]
$$

Note that Lemma 8.3.9 shows that $\left\|\left(\boldsymbol{S}^{(k)}\right)^{-1}\left(\boldsymbol{s}^{(k+1)}-\boldsymbol{s}^{(k)}\right)\right\|_{2} \leq \frac{1}{10}$ for the $k^{\text {th }}$ and $(k+1)^{\text {th }}$ linear systems we solved in the algorithm. Hence, we have $\left\|\left(\boldsymbol{D}^{(k)}\right)^{-1}\left(\boldsymbol{d}^{(k+1)}-\boldsymbol{d}^{(k)}\right)\right\|_{2} \leq \frac{1}{10}$. In [164], they showed how to solve such sequence of systems in $\tilde{O}\left(n^{2}\right)$ amortized cost. Moreover, since our algorithm always changes the constraints by $\delta$ amount where $\delta=\Omega\left(\frac{\varepsilon}{\sqrt{n}}\right)$ an inexact separation oracle $\mathrm{SO}_{\Omega(\varepsilon / \sqrt{n})}$ suffices. (see Def 2.3.5). Consequently, the total work $O\left(n \mathrm{SO}_{\Omega(\varepsilon / \sqrt{n})}(K) \log (n R / \varepsilon)+n^{3} \log ^{O(1)}(n R / \varepsilon)\right)$. Note that as the running time holds with only constant probability, we can restart the algorithm whenever the running time is too large.

To prove (2), we note that from the algorithm description, we know the constraints are either from $B_{\infty}(R)$ or of the form $\boldsymbol{a}^{T} \boldsymbol{x} \geq \boldsymbol{a}^{T} \boldsymbol{x}^{(k)}-\delta$ where

$$
\delta=\sqrt{\frac{\boldsymbol{a}^{T}\left(\boldsymbol{A}^{T} \boldsymbol{S}_{\boldsymbol{x}^{(k)}}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{a}}{c_{a}}} .
$$

From the proof of Lemma 8.3.24, we know that if $\lambda_{\max }\left(\boldsymbol{A}^{T} \boldsymbol{S}_{x}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right) \geq \frac{n}{c_{a}^{2} \varepsilon^{2}}$, then there is $s_{i}<\varepsilon$. Hence, we have $\lambda_{\min }\left(\left(\boldsymbol{A}^{T} \boldsymbol{S}_{x}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1}\right) \geq \frac{c_{a}^{2} \varepsilon^{2}}{n}$. Since $\boldsymbol{a}$ is a unit vector, we have

$$
\sqrt{\frac{\boldsymbol{a}^{T}\left(\boldsymbol{A}^{T} \boldsymbol{S}_{\boldsymbol{x}^{(k)}}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{a}}{c_{a}}} \geq \sqrt{\frac{c_{a} \varepsilon^{2}}{n}}
$$

## ■ 8.4 Technical Tools

In this section we provide stand-alone technical tools we use in our cutting plane method in Section 8.3. In Section 8.4 .1 we show how to efficiently compute accurate estimates of changes in leverage scores using access to a linear system solver. In Section 8.4 .2 we study what we call the "Stochastic Chasing 0 Game" and show how to maintain that a vector is small in $\ell_{\infty}$ norm by making small coordinate updates while the vector changes randomly in $\ell_{2}$.

## - 8.4.1 Estimating Changes in Leverage Scores

In previous sections, we needed to compute leverage scores accurately and efficiently for use in our cutting plane method. Note that the leverage score definition we used was

$$
\psi(\boldsymbol{w})_{i}=\mathbf{1}_{i}^{T} \sqrt{\boldsymbol{W}} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{A}^{T} \sqrt{\boldsymbol{W}} \mathbf{1}_{i}
$$

for some $\lambda>0$ which is different from the standard definition

$$
\sigma(\boldsymbol{w})_{i}=\mathbf{1}_{i}^{T} \sqrt{\boldsymbol{W}} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}\right)^{-1} \boldsymbol{A}^{T} \sqrt{\boldsymbol{W}} \mathbf{1}_{i} .
$$

However, note that the matrix $\boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}+\lambda \boldsymbol{I}$ can be written as $\overline{\boldsymbol{A}}^{T} \boldsymbol{D} \overline{\boldsymbol{A}}$ for the matrix $\overline{\boldsymbol{A}}^{T}=\left[\boldsymbol{A}^{T} \boldsymbol{I}\right]$ and diagonal matrix

$$
D=\left[\begin{array}{cc}
W & 0 \\
0 & \lambda I
\end{array}\right]
$$

and therefore computing $\psi$ is essentially same as computing typical leverage scores. Consequently, we use the standard definition $\sigma$ to simplify notation.

We explained in Section 2.3.1 that leverage scores can be approximated efficiently using dimension reduction. Unfortunately, the error incurred by this approximation is too large to use inside the cutting point method. In this section, we show how to efficiently approximate the change of leverage score more accurately.

In particular, we show how to approximate $\sigma(\boldsymbol{w})-\sigma(\boldsymbol{v})$ for any given $\boldsymbol{w}, \boldsymbol{v}$ with $\| \log (\boldsymbol{w})-$ $\log (\boldsymbol{v}) \|_{2} \ll 1$. Our algorithm breaks $\sigma(\boldsymbol{w})_{i}-\sigma(\boldsymbol{v})_{i}$ into the sum of the norm of small vectors and then uses the Johnson-Lindenstrauss dimension reduction to approximate the norm of each vector separately. Our algorithm makes use of the following version of Johnson-Lindenstrauss.
Lemma 8.4.1 ([2]). Let $0 \leq \varepsilon \leq \frac{1}{2}$ and let $\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{m} \in \mathbb{R}^{n}$ be arbitrary m points. For $k=$ $O\left(\varepsilon^{-2} \log (m)\right)$ let $\boldsymbol{Q}$ be a $k \times n$ random matrix with each entry sampled from $\left\{-\frac{1}{\sqrt{k}}, \frac{1}{\sqrt{k}}\right\}$ uniformly and independently. Then, $\mathbb{E}\left\|\boldsymbol{Q}_{i}\right\|^{2}=\left\|\boldsymbol{x}_{i}\right\|^{2}$ for all $i \in[m]$ and with high probability in $m$ we have that for all $i \in[m]$

$$
(1-\varepsilon)\left\|\boldsymbol{x}_{i}\right\|^{2} \leq\left\|\boldsymbol{Q} \boldsymbol{x}_{i}\right\|^{2} \leq(1+\varepsilon)\left\|\boldsymbol{x}_{i}\right\|^{2}
$$

```
Algorithm 23: \(\widehat{h}=\) LeverageChange \((\boldsymbol{A}, \boldsymbol{v}, \boldsymbol{w}, \varepsilon)\)
    Input: \(\boldsymbol{A} \in \mathbb{R}^{m \times n}, \boldsymbol{v}, \boldsymbol{w} \in \mathbb{R}_{>0}^{m}, \varepsilon \in(0,0.5)\).
    Given: \(\left\|\boldsymbol{V}^{-1}(\boldsymbol{v}-\boldsymbol{w})\right\|_{2} \leq \frac{1}{10} ; \boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\) and \(\boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}\) are invertible.
    Sample \(\boldsymbol{Q}_{d} \in \mathbb{R}^{O\left(\varepsilon^{-2} \log (m)\right) \times n}\) as in Lemma 8.4.1.
    Let \(\hat{d}_{i}=\left\|\boldsymbol{Q}_{d} \sqrt{\boldsymbol{W}} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}\right)^{-1} \boldsymbol{A}^{T} \mathbf{1}_{i}\right\|_{2}^{2}\) for all \(i \in[n]\).
    Let \(t=O\left(\log \left(\varepsilon^{-1}\right)\right)\).
    Sample \(\boldsymbol{Q}_{f} \in \mathbb{R}^{O\left(\varepsilon^{-2} \log (m t)\right) \times n}\) as in Lemma 8.4.1.
    Pick positive integer \(u\) randomly such that \(\operatorname{Pr}[u=i]=\left(\frac{1}{2}\right)^{i}\).
    for \(j \in\{1,2, \cdots, t\} \cup\{t+u\}\) do
        if \(j\) is even then
                Let \(\hat{f}_{i}^{(j)}=\left\|\boldsymbol{Q}_{f} \sqrt{\boldsymbol{V}} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1}\left(\boldsymbol{A}^{T}(\boldsymbol{V}-\boldsymbol{W}) \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1}\right)^{\frac{j}{2}} \boldsymbol{A}^{T} \mathbf{1}_{i}\right\|_{2}^{2}\).
        else
            Let \(\Delta^{+} \stackrel{\text { def }}{=}(\boldsymbol{V}-\boldsymbol{W})^{+}\), i.e. the matrix \(\boldsymbol{V}-\boldsymbol{W}\) with negative entries set to 0.
            Let \(\Delta^{-} \stackrel{\text { def }}{=}(\boldsymbol{W}-\boldsymbol{V})^{+}\), i.e. the matrix \(\boldsymbol{W}-\boldsymbol{V}\) with negative entries set to 0 .
                Let \(\hat{\alpha}_{i}^{(j)}=\left\|\boldsymbol{Q}_{f} \sqrt{\Delta^{+}} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1}\left(\boldsymbol{A}^{T}(\boldsymbol{V}-\boldsymbol{W}) \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1}\right)^{\frac{j-1}{2}} \boldsymbol{A}^{T} \mathbf{1}_{i}\right\|_{2}^{2}\).
                Let \(\hat{\beta}_{i}^{(j)}=\left\|\boldsymbol{Q}_{f} \sqrt{\Delta^{-}} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1}\left(\boldsymbol{A}^{T}(\boldsymbol{V}-\boldsymbol{W}) \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1}\right)^{\frac{j-1}{2}} \boldsymbol{A}^{T} \mathbf{1}_{i}\right\|_{2}^{2}\).
                Let \(\hat{f}_{i}^{(j)}=\hat{\alpha}_{i}^{(j)}-\hat{\beta}_{i}^{(j)}\).
        end
    end
    Let \(\hat{f}_{i}=2^{u} \hat{f}_{i}^{(t+u)}+\sum_{j=1}^{t} \hat{f}_{i}^{(j)}\).
    Output: \(\hat{h}_{i}=\left(w_{i}-v_{i}\right) \hat{d}_{i}+v_{i} \hat{f}_{i}\). for all \(i \in[m]\)
```

Theorem 8.4.2. Let $\boldsymbol{A} \in \mathbb{R}^{m \times n}$ and $\boldsymbol{v}, \boldsymbol{w} \in \mathbb{R}_{>0}^{m}$ be such that $\alpha \stackrel{\text { def }}{=}\left\|\boldsymbol{V}^{-1}(\boldsymbol{v}-\boldsymbol{w})\right\|_{2} \leq \frac{1}{10}$ and both $\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}$ and $\boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}$ are invertible. For any $\varepsilon \in(0,0.5)$, Algorithm 23 generates a random variable $\widehat{h}$
such that $\mathbf{E} \hat{h}=\sigma(\boldsymbol{w})-\sigma(\boldsymbol{v})$ and with high probability in m, we have $\|\widehat{h}-(\sigma(\boldsymbol{w})-\sigma(\boldsymbol{v}))\|_{2} \leq O(\alpha \varepsilon)$. Furthermore, the expected running time is $\widetilde{O}\left((\mathrm{nnz}(\boldsymbol{A})+\mathrm{LO}) / \varepsilon^{2}\right)$ where LO is the amount of time needed to apply $\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1}$ and $\left(\boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}\right)^{-1}$ to a vector.

Proof. First we bound the running time. To compute $\hat{d}_{i}, \hat{f}_{i}^{(j)}, \hat{\alpha}_{i}^{(j)}, \hat{\beta}_{i}^{(j)}$, we simply perform matrix multiplications from the left and then consider the dot products with each of the rows of $\boldsymbol{A}$. Naively this would take time $\widetilde{O}\left((t+u)^{2} \log (m t)(\operatorname{nnz}(\boldsymbol{A})+\mathrm{LO})\right)$. However, we can reuse the computation in computing high powers of $j$ to only take time $\widetilde{O}((t+u) \log (m t)(\mathrm{nnz}(\boldsymbol{A})+\mathrm{LO}))$. Now since $\mathbf{E}[u]$ is constant we see that the total running time is as desired. It only remains to prove the desired properties of $\hat{h}$.

First we note that we can re-write leverage score differences using

$$
\sigma(\boldsymbol{w})_{i}-\sigma(\boldsymbol{v})_{i}=\left(w_{i}-v_{i}\right)\left[\boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}\right)^{-1} \boldsymbol{A}^{T}\right]_{i i}+v_{i}\left[\boldsymbol{A}\left(\left(\boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}\right)^{-1}-\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1}\right) \boldsymbol{A}^{T}\right]_{i i}
$$

Consequently, for all $i \in[m]$, if we let

$$
\begin{aligned}
& d_{i} \stackrel{\text { def }}{=} \mathbf{1}_{i}^{T} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}\right)^{-1} \boldsymbol{A}^{T} \mathbf{1}_{i}, \\
& f_{i} \stackrel{\text { def }}{=} \\
& \mathbf{1}_{i}^{T} \boldsymbol{A}\left[\left(\boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}\right)^{-1}-\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1}\right] \boldsymbol{A}^{T} \mathbf{1}_{i} .
\end{aligned}
$$

then

$$
\begin{equation*}
\sigma(\boldsymbol{w})_{i}-\sigma(\boldsymbol{v})_{i}=\left(w_{i}-v_{i}\right) d_{i}+\left(v_{i}\right) f_{i} \tag{8.24}
\end{equation*}
$$

We show that $\hat{d}_{i}$ approximates $d$ and $\hat{f}_{i}$ approximate $\hat{f}$ well enough to satisfy the statements in the Theorem.

First we bound the quality of $\hat{d}_{i}$. Note that $d_{i}=\left\|\sqrt{\boldsymbol{W}} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}\right)^{-1} \boldsymbol{A}^{T} \mathbf{1}_{i}\right\|_{2}^{2}$. Consequently, Lemma 8.4 .1 shows that $\mathbf{E}\left[\hat{d}_{i}\right]=d_{i}$ and that with high probability in $m$ we have $(1-\varepsilon) d_{i} \leq \hat{d}_{i} \leq(1+\varepsilon) d_{i}$ for all $i \in[m]$. Therefore, with high probability in $m$, we have

$$
\begin{aligned}
\|(\boldsymbol{w}-\boldsymbol{v}) \hat{d}-(\boldsymbol{w}-\boldsymbol{v}) \boldsymbol{d}\|_{2}^{2} & =\sum_{i \in[m]}\left(w_{i}-v_{i}\right)^{2}\left(\widehat{d}_{i}-d_{i}\right)^{2} \leq \varepsilon^{2} \sum_{i \in[m]}\left(w_{i}-v_{i}\right)^{2} d_{i}^{2} \\
& =\varepsilon^{2} \sum_{i \in[m]}\left(w_{i}-v_{i}\right)^{2}\left(\frac{\sigma(\boldsymbol{w})_{i}}{\boldsymbol{w}_{i}}\right)^{2} \leq 2 \varepsilon^{2} \sum_{i \in[m]}\left(\frac{w_{i}-v_{i}}{v_{i}}\right)^{2}
\end{aligned}
$$

Next we show how to estimate $f$. Let $\boldsymbol{X} \stackrel{\text { def }}{=}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2} \boldsymbol{A}^{T}(\boldsymbol{V}-\boldsymbol{W}) \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2}$. By the assumption on $\alpha$ we know $-\frac{1}{2} \boldsymbol{V} \prec \boldsymbol{V}-\boldsymbol{W} \prec \frac{1}{2} \boldsymbol{V}$ and therefore $-\frac{1}{2} \boldsymbol{I} \prec \boldsymbol{X} \prec \frac{1}{2} \boldsymbol{I}$. Consequently we have that

$$
\begin{aligned}
\left(\boldsymbol{A}^{T} \boldsymbol{W} \boldsymbol{A}\right)^{-1} & =\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2}(\boldsymbol{I}-\boldsymbol{X})^{-1}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2} \\
& =\sum_{j=0}^{\infty}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2} \boldsymbol{X}^{j}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2}
\end{aligned}
$$

and therefore

$$
\begin{aligned}
f_{i} & =\mathbf{1}_{i}^{T} \boldsymbol{A}\left(\sum_{j=0}^{\infty}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2} \boldsymbol{X}^{j}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2}-\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1}\right) \boldsymbol{A}^{T} \mathbf{1}_{i} \\
& =\sum_{j=1}^{\infty} f_{i}^{(j)} \text { where } f_{i}^{(j)} \stackrel{\text { def }}{=} \mathbf{1}_{i}^{T} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2} \boldsymbol{X}^{j}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2} \boldsymbol{A}^{T} \mathbf{1}_{i}
\end{aligned}
$$

Furthermore, using the definition of $\boldsymbol{X}$ we have that for even j

$$
\begin{aligned}
f_{i}^{(j)} & =\left\|\boldsymbol{X}^{\frac{j}{2}}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2} \boldsymbol{A}^{T} \mathbf{1}_{i}\right\|_{2}^{2} \\
& =\left\|\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2}\left(\boldsymbol{A}^{T}(\boldsymbol{V}-\boldsymbol{W}) \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1}\right)^{\frac{j}{2}} \boldsymbol{A}^{T} \mathbf{1}_{i}\right\|_{2}^{2} \\
& =\left\|\sqrt{\boldsymbol{V}} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1}\left(\boldsymbol{A}^{T}(\boldsymbol{V}-\boldsymbol{W}) \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1}\right)^{\frac{j}{2}} \boldsymbol{A}^{T} \mathbf{1}_{i}\right\|_{2}^{2}
\end{aligned}
$$

For odd $j$, using our definition of $\boldsymbol{\Delta}^{+}$and $\boldsymbol{\Delta}^{-}$we have that

$$
\begin{aligned}
f_{i}^{(j)}= & \mathbf{1}_{i}^{T} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2} \boldsymbol{X}^{j}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2} \boldsymbol{A}^{T} \mathbf{1}_{i} \\
= & \mathbf{1}_{i}^{T} \boldsymbol{A}\left(\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1} \boldsymbol{A}^{T}(\boldsymbol{V}-\boldsymbol{W}) \boldsymbol{A}\right)^{\frac{j-1}{2}}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1} \boldsymbol{A}^{T}(\boldsymbol{V}-\boldsymbol{W}) \\
& \times \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1}\left(\boldsymbol{A}^{T}(\boldsymbol{V}-\boldsymbol{W}) \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1}\right)^{\frac{j-1}{2}} \boldsymbol{A}^{T} \mathbf{1}_{i} \\
= & \alpha_{i}^{(j)}-\beta_{i}^{(j)}
\end{aligned}
$$

where

$$
\begin{array}{ll}
\alpha_{i}^{(j)} \stackrel{\text { def }}{=}\left\|\sqrt{\Delta^{+}} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1}\left(\boldsymbol{A}^{T}(\boldsymbol{W}-\boldsymbol{V}) \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1}\right)^{\frac{j-1}{2}} \boldsymbol{A}^{T} \mathbf{1}_{i}\right\|_{2}^{2} \\
\beta_{i}^{(j)} \stackrel{\text { def }}{=}\left\|\sqrt{\Delta^{-}} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1}\left(\boldsymbol{A}^{T}(\boldsymbol{W}-\boldsymbol{V}) \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1}\right)^{\frac{j-1}{2}} \boldsymbol{A}^{T} \mathbf{1}_{i}\right\|_{2}^{2}
\end{array}
$$

Consequently, by Lemma 8.4.1 and the construction, we see that

$$
\mathbf{E} \hat{f}_{i}=\sum_{j=1}^{\infty} f_{i}^{(j)}=f_{i}
$$

and therefore $\mathbf{E} \hat{h}=\sigma(\boldsymbol{w})-\sigma(\boldsymbol{v})$ as desired. All that remains is to bound the variance of $\hat{f_{i}}$.

To bound the variance of $\hat{f}$, let $|\boldsymbol{X}|=\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2} \boldsymbol{A}^{T}|\boldsymbol{W}-\boldsymbol{V}| \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2}$. Note that $-\frac{1}{4} \boldsymbol{I} \preceq-|\boldsymbol{X}| \preceq \boldsymbol{X} \preceq|\boldsymbol{X}| \preceq \frac{1}{4} \boldsymbol{I}$ and consequently for all $j$

$$
\begin{aligned}
g_{i}^{(j)} & \stackrel{\text { def }}{=} \mathbf{1}_{i}^{T} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2}|\boldsymbol{X}|^{j}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2} \boldsymbol{A}^{T} \mathbf{1}_{i} \\
& \leq \frac{1}{4^{j-1}} \mathbf{1}_{i}^{T} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2}|\boldsymbol{X}|\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2} \boldsymbol{A}^{T} \mathbf{1}_{i} \\
& \stackrel{\text { def }}{=} \frac{1}{v_{i} 4^{j-1}} \mathbf{1}_{i}^{T} \boldsymbol{P}_{v} \boldsymbol{\Delta} \boldsymbol{P}_{v} \mathbf{1}_{i}
\end{aligned}
$$

where $\boldsymbol{P}_{v}=\sqrt{\boldsymbol{V}} \boldsymbol{A}\left(\boldsymbol{A}^{T} \boldsymbol{V} \boldsymbol{A}\right)^{-1 / 2} \boldsymbol{A}^{T} \sqrt{\boldsymbol{V}}$ and $\boldsymbol{\Delta}$ is a diagonal matrix with $\boldsymbol{\Delta}_{i i}=\left|\frac{w_{i}-v_{i}}{v_{i}}\right|$. Using that $\mathbf{0} \preceq \boldsymbol{P}_{v} \preceq \boldsymbol{I}$, we have that for all $j$

$$
\begin{aligned}
\left(4^{j-1}\right)^{2} \sum_{i=1}^{m}\left(v_{i} g_{i}^{(j)}\right)^{2} & =\sum_{i=1}^{m}\left(\mathbf{1}_{i}^{T} \boldsymbol{P}_{v} \boldsymbol{\Delta} \boldsymbol{P}_{v} \mathbf{1}_{i}\right)^{2}=\operatorname{Tr}\left(\boldsymbol{P}_{v} \boldsymbol{\Delta} \boldsymbol{P}_{v} \boldsymbol{P}_{v} \boldsymbol{\Delta} \boldsymbol{P}_{v}\right) \\
& \leq \operatorname{Tr}\left(\boldsymbol{P}_{v} \boldsymbol{\Delta} \boldsymbol{\Delta} \boldsymbol{P}_{v}\right)=\operatorname{Tr}\left(\boldsymbol{\Delta} \boldsymbol{P}_{v} \boldsymbol{P}_{v} \boldsymbol{\Delta}\right) \\
& \leq \operatorname{Tr}\left(\boldsymbol{\Delta}^{2}\right)=\sum_{i=1}^{m}\left(\frac{w_{i}-v_{i}}{v_{i}}\right)^{2} \leq \alpha^{2}
\end{aligned}
$$

and thus $\left\|\boldsymbol{V} \boldsymbol{g}^{(j)}\right\|_{2} \leq \frac{4 \alpha}{4 j}$. Furthermore, since $\boldsymbol{\Delta}^{+} \preceq|\boldsymbol{W}-\boldsymbol{V}|$ and $\boldsymbol{\Delta}^{-} \preceq|\boldsymbol{W}-\boldsymbol{V}|$ we have that $\left|\alpha_{i}^{(j)}\right| \leq g_{i}^{(j)}$ and $\left|\beta_{i}^{(j)}\right| \leq g_{i}^{(j)}$. Consequently, by Lemma 8.4.1 again, we have

$$
\begin{aligned}
\left\|\boldsymbol{V} \hat{f}^{(j)}-\boldsymbol{V} \boldsymbol{f}^{(j)}\right\|_{2}^{2} & =\sum_{i} v_{i}^{2}\left(\hat{f}_{i}^{(j)}-f_{i}^{(j)}\right)^{2} \\
& \leq 2 \sum_{i} v_{i}^{2}\left(\hat{\alpha}_{i}^{(j)}-\alpha_{i}^{(j)}\right)^{2}+2 \sum_{i} v_{i}^{2}\left(\hat{\beta}_{i}^{(j)}-\beta_{i}^{(j)}\right)^{2} \\
& \leq 2 \varepsilon^{2} \sum_{i} v_{i}^{2}\left(\left(\alpha_{i}^{(j)}\right)^{2}+\left(\beta_{i}^{(j)}\right)^{2}\right) \\
& \leq 4 \varepsilon^{2} \sum_{i}\left(v_{i} g_{i}^{(j)}\right)^{2} \leq \frac{4 \alpha^{2} \varepsilon^{2}}{\left(4^{j-1}\right)^{2}} .
\end{aligned}
$$

Putting this all together we have that

$$
\begin{aligned}
\|\boldsymbol{V} \hat{f}-\boldsymbol{V} \boldsymbol{f}\|_{2} & \leq\left\|2^{u} \boldsymbol{V} \hat{f}^{(t+u)}+\sum_{j=1}^{t} \boldsymbol{V} \hat{f}^{(j)}-\sum_{j=1}^{\infty} \boldsymbol{V} \boldsymbol{f}^{(j)}\right\|_{2} \\
& \leq 2^{u}\left\|\boldsymbol{V} \hat{f}^{(t+u)}\right\|_{2}+\sum_{j=1}^{t}\left\|\boldsymbol{V} \hat{f}^{(j)}-\boldsymbol{V} \boldsymbol{f}^{(j)}\right\|_{2}+\sum_{j=t+1}^{\infty}\left\|\boldsymbol{V} \boldsymbol{f}^{(j)}\right\|_{2} \\
& \leq 2^{u} \frac{4 \alpha}{4^{t+u-1}}+\sum_{j=1}^{t} \frac{2 \alpha \varepsilon}{4^{j-1}}+\sum_{j=t+1}^{\infty} \frac{2 \alpha}{4^{j-1}} \\
& =O\left(\alpha \varepsilon+\frac{\alpha}{4^{t}}\right)
\end{aligned}
$$

Consequently, since $t=O\left(\log \left(\varepsilon^{-1}\right)\right)$ we have the desired result.

## - 8.4.2 The Stochastic Chasing 0 Game

To avoid computing leverage scores exactly, in Section 8.4.1 we showed how to estimate the difference of leverage scores and use these to update the leverage scores. However, if we only applied this technique, the error of leverage scores would accumulate in the algorithm and we need to fix it. Naturally, one may wish to use dimension reduction to compute a multiplicative approximation to the leverage scores and update our computed value if the error is too large. However, this strategy would fail if there are too many rows with inaccurate leverage scores in the same iteration. In this case, we would change the central point too much that we are not able to recover. In this section, we present this update problem in a general form that we call Stochastic Chasing 0 game and provide an effective strategy for playing this game.

The Stochastic chasing 0 game is as follows. There is a player, a stochastic adversary, and a point
$\boldsymbol{x} \in \mathbb{R}^{m}$. The goal of the player is to keep the point close to $\mathbf{0} \in \mathbb{R}^{m}$ in $\ell_{\infty}$ norm and the goal of the stochastic adversary is to move $\boldsymbol{x}$ away from $\mathbf{0}$. The game proceeds for an infinite number of iterations where in each iteration the stochastic adversary moves the current point $\boldsymbol{x}^{(k)} \in \mathbb{R}^{m}$ to some new point $\boldsymbol{x}^{(k)}+\boldsymbol{\Delta}^{(k)} \in \mathbb{R}^{m}$ and the player needs to respond. The stochastic adversary cannot move the $\boldsymbol{\Delta}^{(k)}$ arbitrarily, instead he is only allowed to choose a probability distribution $\mathcal{D}^{(k)}$ and sample $\boldsymbol{\Delta}^{(k)}$ from it. Furthermore, it is required that $\mathbb{E}_{\mathcal{D}^{(k)}} \boldsymbol{\Delta}=\mathbf{0}$ and $\|\boldsymbol{\Delta}\|_{2}^{2} \leq c$ for some fixed $c$ and all $\boldsymbol{\Delta} \in \mathcal{D}^{(k)}$. The player does not know $\boldsymbol{x}^{(k)}$ or the distribution $\mathcal{D}^{(k)}$ or the move $\boldsymbol{\Delta}^{(k)}$ of the stochastic adversary. All the player knows is some $\boldsymbol{y}^{(k)} \in \mathbb{R}^{n}$ that is close to $\boldsymbol{x}^{(k)}$ in $\ell_{\infty}$ norm. With this information, the player is allowed to choose one coordinate $i$ and set $x_{i}^{(k+1)}$ to be zero and for other $j$, we have $x_{j}^{(k+1)}=x_{j}^{(k)}+\Delta_{j}^{(k)}$.

The question we would like to address is, what strategy the player should choose to keep $\boldsymbol{x}^{(k)}$ close to $\mathbf{0}$ in $\ell_{\infty}$ norm? We show that there is a trivial strategy that performs well: simply pick the largest coordinate and set it to 0 .

```
Algorithm 24: Stochastic chasing 0 game
    Constant: \(c>0, R>0\).
    Let \(\boldsymbol{x}^{(1)}=\mathbf{0} \in \mathbb{R}^{m}\).
    for \(k=1\) to \(\infty\) do
        Stochastic Adversary: Pick \(\mathcal{D}^{(k)}\) such that \(\mathbb{E}_{\mathcal{D}^{(k)}} \boldsymbol{\Delta}=\mathbf{0}\) and \(\|\boldsymbol{\Delta}\|_{2} \leq c\) all \(\boldsymbol{\Delta} \in \mathcal{D}^{(k)}\).
        Stochastic Adversary: Pick \(\boldsymbol{y}^{(k)} \in \mathbb{R}^{m}\) such that \(\left\|\boldsymbol{y}^{(k)}-\boldsymbol{x}^{(k)}\right\|_{\infty} \leq R\).
        Player: Pick a coordinate \(i^{(k)}\) using only \(\boldsymbol{y}^{(k)}\).
        Sample \(\boldsymbol{\Delta}^{(k)}\) from \(\mathcal{D}^{(k)}\).
        Set \(x_{i(k)}^{(k+1)}=0\) and \(x_{j}^{(k+1)}=x_{j}^{(k)}+\Delta_{j}^{(k)}\) for all \(j \neq i^{(k)}\).
    end
```

Theorem 8.4.3. Using the strategy $i^{(k)}=\arg \max _{i}\left|y_{i}^{(k)}\right|$, with probability at least $1-p$, we have

$$
\left\|\boldsymbol{x}^{(k)}\right\|_{\infty} \leq 2(c+R) \log \left(4 m k^{2} / p\right)
$$

for all $k$ in the Stochastic Chasing $\mathbf{0}$ Game.

Proof. Consider the potential function $\Phi(\boldsymbol{x})=\sum_{i} e^{\alpha x_{i}}+\sum_{i} e^{-\alpha x_{i}}$ where $\alpha$ is to be determined. Now for all $x$ we know that $e^{x} \leq 1+x+\frac{x^{2}}{2} e^{|x|}$ and therefore for all $|\delta| \leq c, x$ and $\alpha$, we have

$$
e^{\alpha x+\alpha \delta} \leq e^{\alpha x}+\alpha \delta e^{\alpha x}+\frac{1}{2} \alpha^{2} \delta^{2} e^{\alpha x+|\alpha| c}
$$

Consequently,

$$
\begin{aligned}
& \mathbb{E}_{\boldsymbol{\Delta} \in \mathcal{D}^{(k)}} \Phi\left(\boldsymbol{x}^{(k)}+\boldsymbol{\Delta}\right) \leq \Phi\left(\boldsymbol{x}^{(k)}\right)+\alpha \mathbb{E}_{\boldsymbol{\Delta} \in \mathcal{D}^{(k)}}\left(\sum_{i \in[m]} e^{\alpha x_{i}^{(k)}} \Delta_{i}-\sum_{i \in[m]} e^{-\alpha x_{i}^{(k)}} \Delta_{i}\right) \\
&+\frac{\alpha^{2}}{2} e^{\alpha\|\boldsymbol{\Delta}\|_{\infty} \mathbb{E}_{\boldsymbol{\Delta} \in \mathcal{D}^{(k)}}\left(\sum_{i \in[m]} e^{\alpha x_{i}^{(k)}} \Delta_{i}^{2}+\sum_{i \in[m]} e^{-\alpha x_{i}^{(k)}} \Delta_{i}^{2}\right)} .
\end{aligned}
$$

Since $\mathbb{E}_{\mathcal{D}^{(k)}} \boldsymbol{\Delta}=\mathbf{0}$ and $\|\boldsymbol{\Delta}\|_{2} \leq c$, we have $\mathbb{E}_{\boldsymbol{\Delta} \in \mathcal{D}^{(k)}}\left(\sum_{i} e^{\alpha x_{i}^{(k)}} \Delta_{i}-\sum_{i} e^{-\alpha x_{i}^{(k)}} \Delta_{i}\right)=0$ and

$$
\begin{aligned}
\mathbb{E}_{\boldsymbol{\Delta} \in \mathcal{D}^{(k)}}\left(\sum_{i} e^{\alpha x_{i}^{(k)}} \Delta_{i}^{2}+\sum_{i} e^{-\alpha x_{i}^{(k)}} \Delta_{i}^{2}\right) & \leq \mathbb{E}_{\boldsymbol{\Delta} \in \mathcal{D}^{(k)}}\left(\sum_{i} \Delta_{i}^{2}\right)\left(\max _{i} e^{\alpha x_{i}^{(k)}}+\max _{i} e^{-\alpha x_{i}^{(k)}}\right) \\
& \leq 2 c^{2} \max _{i} e^{\alpha\left|x_{i}^{(k)}\right|}
\end{aligned}
$$

Letting $\eta^{(k)}=\max _{i} e^{\alpha \mid x_{i}^{(k)}}$, we then have

$$
\mathbb{E}_{\boldsymbol{\Delta} \in \mathcal{D}^{(k)}} \Phi\left(\boldsymbol{x}^{(k)}+\boldsymbol{\Delta}\right) \leq \Phi\left(\boldsymbol{x}^{(k)}\right)+\alpha^{2} e^{\alpha c} c^{2} \eta^{(k)} .
$$

Since $i^{(k)}=\arg \max _{i}\left|y_{i}^{(k)}\right|$ and $\left\|\boldsymbol{y}^{(k)}-\boldsymbol{x}^{(k)}\right\|_{\infty} \leq R$, the player setting $x_{i^{(k)}}^{(k+1)}=0$ decreases $\Phi$ by at least $e^{-\alpha(R+c)} \eta^{(k)}$. Hence, we have

$$
\mathbb{E}_{\boldsymbol{\Delta} \in \mathcal{D}^{(k)}} \Phi\left(\boldsymbol{x}^{(k+1)}\right) \leq \Phi\left(\boldsymbol{x}^{(k)}\right)+\alpha^{2} e^{\alpha c} c^{2} \eta^{(k)}-e^{-\alpha(R+c)} \eta^{(k)} .
$$

Picking $\alpha=\frac{1}{2(c+R)}$, we have $e^{2 \alpha(c+R)}(\alpha(c+R))^{2} \leq 1$ and hence $\alpha^{2} e^{\alpha c} c^{2} \leq e^{-\alpha(R+c)}$. Therefore, we have that

$$
\mathbb{E}_{\boldsymbol{\Delta} \in \mathcal{D}^{(k)}} \Phi\left(\boldsymbol{x}^{(k+1)}\right) \leq \mathbf{E} \Phi\left(\boldsymbol{x}^{(k)}\right) \leq \ldots \leq \Phi\left(\boldsymbol{x}^{(1)}\right)=2 m
$$

Consequently, by Markov's inequality we have that $\operatorname{Pr}\left[\Phi\left(\boldsymbol{x}^{(k)}\right) \geq \lambda_{k}\right] \leq \frac{2 m}{\lambda_{k}}$ for any $\lambda_{k}$. Furthermore, since clearly $\Phi(\boldsymbol{x}) \geq e^{\alpha\|\boldsymbol{x}\|_{\infty}}$ we have that $\operatorname{Pr}\left[\left\|\boldsymbol{x}^{(k)}\right\|_{\infty} \geq \log \left(\lambda_{k}\right) / \alpha\right] \leq \frac{2 m}{\lambda_{k}}$ for all $k$. Choosing $\lambda_{k}=\frac{4 m k^{2}}{p}$ and taking a union bound over all $k$, we have that

$$
\left\|\boldsymbol{x}^{(k)}\right\|_{\infty} \leq 2(c+R) \log \left(4 m k^{2} / p\right)
$$

for all $k$ with probability at least

$$
1-\sum_{i=1}^{\infty} \frac{2 m}{\lambda_{k}}=1-\sum_{k=1}^{\infty} \frac{p}{2 k^{2}} \geq 1-p
$$

## ■ 8.5 Glossary

Here we summarize problem specific notation we use throughout this section. For many quantities we included the typical order of magnitude as they appear during our algorithms.

- Our algorithm maintains a polytope $\{\boldsymbol{A} x \geq b\}$ which contains the set of solution $K$.
- Slacks $\boldsymbol{s}(\boldsymbol{x})=\boldsymbol{A} \boldsymbol{x}-\boldsymbol{b}, \boldsymbol{S}(\boldsymbol{x})$ is the diagonal matrix corresponds to $\boldsymbol{s}(\boldsymbol{x})$. Rescaled constraint matrix $\boldsymbol{A}_{s}=\boldsymbol{S}^{-1} \boldsymbol{A}$.
- Leverage Score: $\boldsymbol{\psi}(\boldsymbol{x})=\boldsymbol{\operatorname { d i a g }}\left(\boldsymbol{A}_{x}\left(\boldsymbol{A}_{x}^{T} \boldsymbol{A}_{x}+\lambda \boldsymbol{I}\right)^{-1} \boldsymbol{A}_{x}^{T}\right), \mu(\boldsymbol{x})=\min \psi_{i}(\boldsymbol{x}), \lambda=\frac{1}{c_{a} R^{2}}$ and $c_{a}=\frac{1}{10^{10}}$.
- Each iteration, if there is a constraint with leverage score less than $c_{d}=\frac{1}{10^{12}}$, then we delete that constraint. Otherwise, we add a constraint and put the leverage score $c_{a}$.
- Each iteration, our algorithm computes the change of leverage score with accuracy $c_{\Delta}=\frac{C c_{e}}{\log (n \log (R / \varepsilon))}$ where $c_{e}=\frac{c_{d}}{4 \ln (6 n R / \varepsilon)}$ and $R$ is the diameter of the box containing $K$.
- Potential function $p_{\boldsymbol{e}}(\boldsymbol{x})=-\sum_{i \in[m]}\left(c_{e}+e_{i}\right) \log \left(s_{i}(\boldsymbol{x}) / R\right)+\frac{1}{2} \log \operatorname{det}\left(R^{2}\left(\boldsymbol{A}^{T} \boldsymbol{S}_{x}^{-2} \boldsymbol{A}+\lambda \boldsymbol{I}\right)\right)+$ $\frac{\lambda}{2}\|x\|_{2}^{2}$.
- The algorithm starts with $p_{e}(\boldsymbol{x}) \sim-\Theta(1) n$ and ends when $p_{\boldsymbol{e}}(\boldsymbol{x}) \sim \Theta(1) n \log (n R / \varepsilon)$.
- Centrality $\delta_{e}(\boldsymbol{x})=\left\|\nabla p_{e}(\boldsymbol{x})\right\|_{\boldsymbol{H}(\boldsymbol{x})^{-1}}$ where $\boldsymbol{H}(\boldsymbol{x})=\boldsymbol{A}_{x}^{T}\left(c_{e} \boldsymbol{I}+\boldsymbol{\Psi}(\boldsymbol{x})\right) \boldsymbol{A}_{x}+\lambda \boldsymbol{I} . \quad \delta_{e}(\boldsymbol{x}) \approx$ $\Theta\left(\sqrt{c_{e}+\mu(\boldsymbol{x})}\right)$.
- Approximate Hessian $\boldsymbol{Q}(\boldsymbol{x}, \boldsymbol{w})=\boldsymbol{A}_{x}^{T}\left(c_{e} \boldsymbol{I}+\boldsymbol{W}\right) \boldsymbol{A}_{x}+\lambda \boldsymbol{I}$.


# Effective Use of Cutting Plane Method 

## ■ 9.1 Introduction

Cutting plane methods have long been employed to obtain polynomial time algorithms for solving optimization problems. However, for many problems cutting plane methods are often regarded as inefficient both in theory and in practice. Here, in this chapter, we provide several techniques for applying cutting plane methods efficiently. Moreover, we illustrate the efficacy and versatility of these techniques by applying them to achieve improved running times for solving multiple problems including semidefinite programming, matroid intersection, and submodular flow.

We hope these results revive interest in ellipsoid and cutting plane methods. We believe these results demonstrate how cutting plan methods are often useful not just for showing that a problem is solvable in polynomial time, but in many yield substantial running time improvements. We stress that while some results in this chapter are problem-specific, the techniques introduced here are quite general and are applicable to a wide range of problems.

In the remainder of this introduction we survey the key techniques we use to apply our cutting plane method (Section 9.1.1) and the key results we obtain on improving the running time for solving various optimization problems (Section 9.1.2). We conclude in Section 9.1 .3 by providing an overview of where to find additional technical result in this chapter.

## ■ 9.1.1 Techniques

Although cutting plane methods are typically introduced as algorithms for finding a point in a convex set (as we did with the feasibility problem in Chapter 8), this is often not the easiest way to apply the methods. Moreover, improperly applying results on the feasibility problem to solve convex optimization problems can lead to vastly sub-optimal running times. Our central goal, here, in this chapter is to provide tools that allow cutting plane methods to be efficiently applied to solve complex optimization problems. Some of these tools are new and some are extensions of previously known techniques. Here we briefly survey the techniques we cover in Section 9.3 and Section 9.4.

## Technique 0: From Feasibility to Optimization

In Section 9.3.1, we explain how to use our cutting plane method to solve convex optimization problems using an approximate subgradient oracle. Our result is based on a result of Nemirovski [202] in which he showed how to use a cutting plane method to solve convex optimization problems without smoothness assumptions on the function and with minimal assumptions on the size of the function's domain. We adopt his proof to accommodate for an approximate separation oracle, an extension which is essential for our applications. We use this result as the starting point for two new techniques we discuss below.

## Technique 1: Dimension Reduction through Duality

In Section 9.3.2, we discuss how cutting plane methods can be applied to obtain both primal and dual solutions to convex optimization problems. Moreover, we show how this can be achieved while only applying the cutting plane method in the space, primal or dual, which has a fewer number of variables. Thus we show how to use duality to improve the convergence of cutting plane methods while still solving the original problem.

To illustrate this idea consider the following very simple linear program (LP)

$$
\min _{x_{i} \geq 0, \sum x_{i}=1} \sum_{i=1}^{n} w_{i} x_{i}
$$

where $\boldsymbol{x} \in \mathbb{R}^{n}$ and $\boldsymbol{w} \in \mathbb{R}^{n}$. Although this LP has $n$ variables, it should to be easy to solve purely on the grounds that it only has one equality constraint and thus dual linear program is simply

$$
\max _{y \leq w_{i} \forall i} y
$$

i.e. a LP with only one variable. Consequently, we can apply our cutting plane method to solve it efficiently.

However, while this simple example demonstrates how we can use duality to decrease dimensions, it is not always obvious how to recover the optimal primal solution $x$ variable given the optimal dual solution $y$. Indeed, for many problems their dual is significantly simpler than itself (primal), so some work is required to show that working in the space suffices to require a primal solution.

One such recent example of this approach proving successful is the linear program result explained in Chapter 6. In this result, the authors show how to take advantage of this observation and get a faster LP solver and maximum flow algorithm. It is interesting to study how far this technique can extend, that is, in what settings can one recover the solution to a more difficult dual problem from the solution to its easier primal problem?

There is in fact another precedent for such an approach. Grötschel, Lovász and Schrijver[109] showed how to obtain the primal solution for linear program by using a cutting plane method to solve the linear program exactly. This is based on the observation that cutting plane methods are able to find the active constraints of the optimal solution and hence one can take dual of the linear program to get the dual solution. This idea was further extended in [151] which also observed that cutting plane methods are incrementally building up a LP relaxation of the optimization problem. Hence, one can find a dual solution by taking the dual of that relaxation.

In Section 9.3.2, we provide a fairly general technique to recover a dual optimal solution from an approximately optimal primal solution. Unfortunately, the performance of this technique seems quite problem-dependent. We therefore only analyze this technique for semidefinite programming (SDP), a classic and popular convex optimization problem. As a result, we obtain a faster SDP solver in both the primal and dual formulations of the problem.

## Technique 2: Using Optimization Oracles Directly

In the seminal works of Grötschel, Lovász, Schrijver and independently Karp and Papadimitriou [110, 138], they showed the equivalence between optimization oracles and separation oracles, and gave a general method to construct a separation oracle for a convex set given an optimization oracle for that set, that is an oracle for minimizing linear functionals over the set. This seminal result led to the first weakly polynomial time algorithm for many algorithms such as submodular function minimization. Since then, this idea has been used extensively in various settings [127, 44, 45, 65].

Unfortunately, while this equivalence of separation and optimization is a beautiful and powerful tool for polynomial time solvability of problems, in many case it may lead to inefficient algorithms. In order to use this reduction to get a separation oracle, the optimization oracle may need to be called multiple times - essentially the number of times needed to run a cutting plane method and hence may be detrimental to obtaining small asymptotic running times. Therefore, it is an interesting question of whether there is a way of using an optimization oracle more directly.

In Section 9.4 we provide a partial answer to this question for the case of a broad class of problems, that we call the intersection problem. For these problems we demonstrate how to achieve running time improvements by using optimization oracles directly. The problem we consider is as follows. We wish to solve the problem for some cost vector $\boldsymbol{c} \in \mathbb{R}^{n}$ and convex set $K$. We assume that the convex set $K$ can be decomposed as $K=K_{1} \cap K_{2}$ such that $\max _{\boldsymbol{x} \in K_{1}}\langle\boldsymbol{c}, \boldsymbol{x}\rangle$ and $\max _{\boldsymbol{x} \in K_{2}}\langle\boldsymbol{c}, \boldsymbol{x}\rangle$ can each be solved efficiently. Our goal is to obtain a running time for this problem comparable to that of minimizing $K$ given only a separation oracle for it.

We show that by considering a carefully regularized variant, we obtain a problem such that optimization oracles for $K_{1}$ and $K_{2}$ immediately yield a separation oracle for this regularized problem. By analyzing the regularizer and bounding the domains of the problem we are able to show that this allows us to efficiently compute highly accurate solutions to the intersection problem by applying our cutting plane method once. In other words, we do not need to use a complicated iterative scheme or directly invoke the equivalence between separation and optimization and thereby save $O(\operatorname{poly}(n))$ factors in our running times.

We note that this intersection problem can be viewed as a generalization of the matroid intersection problem and in Section 9.4.2, we show our reduction gives a faster algorithm in certain parameter regimes. As another example, in Section 9.4 .3 we show our reduction gives a substantial polynomial improvement for the submodular flow problem. Furthermore, in Section 9.4.4 we show how our techniques allow us to minimize a linear function over the intersection of a convex set and an affine subspace in a number of iterations that depends only on the co-dimension of the affine space.

## ■ 9.1.2 Applications

Our main goal in this chapter is to provide general techniques for efficiently using cutting plane methods for various problems. Hence, in this chapter, we use minimally problem-specific techniques to achieve the best possible running time. However, we also demonstrate the efficacy of our approach by showing how techniques improve upon the previous best known running times for solve several classic problems in combinatorial and continuous optimization. Here we provide a brief overview of these applications, previous work on these problems, and our results.

In order to avoid deviating from our main discussion, our coverage of previous methods and techniques is brief. Given the large body of prior works on SDP, matroid intersection and submodular flow, it would be impossible to have an in-depth discussion on all of them. Therefore, this section focuses on running time comparisons and explanations of relevant preivous techniques.

## Semidefinite Programming

In Section 9.3.2 we consider the classic semidefinite programming (SDP) problem:

$$
\max _{\boldsymbol{X} \succeq \mathbf{0}} \boldsymbol{C} \bullet \boldsymbol{X} \text { s.t. } \boldsymbol{A}_{i} \bullet \boldsymbol{X}=b_{i} \text { (primal) } \quad \min _{\boldsymbol{y}} \boldsymbol{b}^{T} \boldsymbol{y} \text { s.t. } \sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i} \succeq \boldsymbol{C} \text { (dual) }
$$

where $\boldsymbol{X}, \boldsymbol{C}, \boldsymbol{A}_{i}$ are $m \times m$ symmetric matrices, $\boldsymbol{b}, \boldsymbol{y} \in \mathbb{R}^{n}$, and $\boldsymbol{A} \bullet \boldsymbol{B} \stackrel{\text { def }}{=} \operatorname{Tr}\left(\mathbf{A}^{T} \mathbf{B}\right)$. For many problems, $n \ll m^{2}$ and hence the dual problem has fewer variables than the primal. There are many

| Years | Authors | Running times |
| :---: | :---: | :---: |
| 1992 | Nesterov, Nemirovsky[210] | $\tilde{O}\left(\sqrt{n}\left(n m^{\omega}+n^{\omega-1} m^{2}\right)\right)$ |
| 2000 | Anstreicher [13] | $\tilde{O}\left((m n)^{1 / 4}\left(n m^{\omega}+n^{\omega-1} m^{2}\right)\right)$ |
| 2003 | Krishnan, Mitchell [152] | $\tilde{O}\left(m\left(n^{\omega}+m^{\omega}+S\right)(\right.$ dual SDP $)$ |
| - | This Chapter | $\tilde{O}\left(m\left(n^{2}+m^{\omega}+S\right)\right)$ |

Table 9.1: Previous algorithms for solving a $n \times n$ SDP with $m$ constraints and $S$ non-zeros entries
results and applications of SDP; see [259, 248, 196] for a survey on this topic. Since our focus is on polynomial time algorithms, we do not discuss pseudo-polynomial algorithms such as the spectral bundle method [115], multiplicative weight update methods [17, 19, 126, 8], etc.

Currently, there are two competing approaches for solving SDP problems, namely interior point methods (IPM) and cutting plane methods. Typically, IPMs require fewer iterations than the cutting plane methods, however each iteration of these methods is more complicated and possibly more computationally expensive. For SDP problems, interior point methods require the computations of the Hessian of the function $-\log \operatorname{det}\left(\boldsymbol{C}-\sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i}\right)$ whereas cutting plane methods usually only need to compute minimum eigenvectors of the slack matrix $\boldsymbol{C}-\sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i}$.

In [13], Anstreicher provided the current fastest IPM for solving the dual SDP problem using a method based on the volumetric barrier function. This method takes $O\left((m n)^{1 / 4}\right)$ iterations and each iteration is as cheap as usual IPMs. For general matrices $\boldsymbol{C}, \boldsymbol{X}, \boldsymbol{A}_{i}$, each iteration takes $O\left(n m^{\omega}+\right.$ $n^{\omega-1} m^{2}$ ) time where $\omega$ is the fast matrix multiplication exponent. If the constraint matrices $\boldsymbol{A}_{i}$ are rank one matrices, the iteration cost can be improved to $O\left(m^{\omega}+n m^{2}+n^{2} m\right)$ [153]. If the matrices are sparse, then [97, 197] show how to use matrix completion inside the IPM. However, the running time depends on the extended sparsity patterns which can be much larger than the total number of non-zeros.

In [152], Krishnan and Mitchell observed that the separation oracle for dual SDP takes only $O\left(m^{\omega}+S\right)$ time, where $S=\sum_{i=1}^{n} \mathrm{nnz}\left(\boldsymbol{A}_{i}\right)$ be the total number of non-zeros in the constant matrix. Hence, the cutting plane method by [253] gives a faster algorithm for SDP for many regimes. For $\omega=2.38$, the cutting plane method is faster when $\boldsymbol{A}_{i}$ is not rank 1 and the problem is not too dense, i.e. $\sum_{i=1}^{n} \mathrm{nnz}\left(\boldsymbol{A}_{i}\right)<n^{0.63} \mathrm{~m}^{2.25}$. While there are previous methods for using cutting plane methods to obtain primal solutions[151], to the best of our knowledge, there are no worst case running time analysis for these techniques.

In Section 9.3.2, show how to alleviate this issue. We provide an improved algorithm for finding the dual solution and prove carefully how to obtain a comparable primal solution as well. See Figure 9.1 for a summary of the algorithms for SDP and their running times.

## Matroid Intersection

In Section 9.4 .2 we show how our optimization oracle technique can be used to improve upon the previous best known running times for matroid intersection. Matroid intersection is one of the most fundamental problems in combinatorial optimization. The first algorithm for matroid intersection is due to the seminal paper by Edmonds [81]. In Figures 9.2 and 9.3 we provide a summary of the previous algorithms for unweighted and weighted matroid intersection as well as the new running times we obtain in this chapter. While there is no total ordering on the running times of these algorithms due to the different dependence on various parameters, we would like to point out that our algorithms outperform the previous ones in regimes where $r$ is close to $n$ and/or the oracle query costs are relatively expensive. In particular, in terms of oracle query complexity our algorithms are the first to

| Years | Authors | Running times |
| :---: | :---: | :---: |
| 1968 | Edmonds [81] | not stated |
| 1971 | Aigner, Dowling [6] | $O\left(n r^{2} \mathcal{T}_{\text {ind }}\right)$ |
| 1974 | Tomizawa, Iri [249] | not stated |
| 1975 | Lawler [157] | $O\left(n r^{2} \mathcal{T}_{\text {ind }}\right)$ |
| 1979 | Edmonds [80] | not stated |
| 1986 | Cunningham [60] | $O\left(n r^{1.5} \mathcal{T}_{\text {ind }}\right)$ |
| - | This Chapter | $O\left(n^{2} \log n \mathcal{T}_{\text {ind }}+n^{3} \log ^{O(1)} n\right)$ |
| $O\left(n r \log ^{2} n \mathcal{T}_{\text {rank }}+n^{3} \log ^{O(1)} n\right)$ |  |  |

Table 9.2: Previous algorithms for (unweighted) matroid intersection. Here $n$ is the size of the ground set, $r=\max \left\{r_{1}, r_{2}\right\}$ is the maximum rank of the two matroids, $\mathcal{T}_{\text {ind }}$ is the time needed to check if a set is independent (independence oracle), and $\mathcal{T}_{\text {rank }}$ is the time needed to compute the rank of a given set (rank oracle).

| Years | Authors | Running times |
| :---: | :---: | :---: |
| 1968 | Edmonds [81] | not stated |
| 1974 | Tomizawa, Iri [249] | not stated |
| 1975 | Lawler [157] | $O\left(n r^{2} \mathcal{T}_{\text {ind }}+n r^{3}\right)$ |
| 1979 | Edmonds [80] | not stated |
| 1981 | Frank [90] | $O\left(n^{2} r\left(\mathcal{T}_{\text {circuit }}+n\right)\right)$ |
| 1983 | Orlin, Ahuja [219] | not stated |
| 1986 | Brezovec, Cornuéjols, Glover[39] | $O\left(n r\left(\mathcal{T}_{\text {circuit }}+r+\log n\right)\right)$ |
| 1995 | Fujishige, Zhang [96] | $O\left(n^{2} r^{0.5} \log r M \cdot \mathcal{T}_{\text {ind }}\right)$ |
| 1995 | Shigeno, Iwata [234] | $O\left(\left(n+\mathcal{T}_{\text {circuit }}\right) n r^{0.5} \log r M\right)$ |
| - | This Chapter | $O\left(\left(n^{2} \log n \mathcal{T}_{\text {ind }}+n^{3} \log { }^{O(1)} n\right) \log n M\right)$ |
|  |  | $O\left(\left(n r \log ^{2} n \mathcal{T}_{\text {rank }}+n^{3} \log O(1) n\right) \log n M\right)$ |

Table 9.3: Previous algorithms for weighted matroid intersection. In additions to the notations used in the unweighted table, $\mathcal{T}_{\text {circuit }}$ is the time needed to find a fundamental circuit and $M$ is the bit complexity of the weights.
achieve the quadratic bounds of $\tilde{O}\left(n^{2}\right)$ and $\tilde{O}(n r)$ for independence and rank oracles. We hope our work will revive the interest in the problem of which progress has been mostly stagnated for the past 20-30 years.

## Minimum-Cost Submodular Flow

In Section 9.4 .3 we show how our optimization oracle technique can be used to improve upon the previous best known running times for (Minimum-cost) Submodular Flow. Submodular flow is a very general problem in combinatorial optimization which generalizes many problems such as minimum cost flow, the graph orientation, polymatroid intersection, directed cut covering [94]. In Figure 9.4 we provide an overview of the previous algorithms for submodular flow as well as the new running times we obtain in this chapter.

Many of the running times are in terms of a parameter $h$, which is the time required for computing an "exchange capacity". To the best of our knowledge, the most efficient way of computing an exchange capacity is to solve an instance of submodular minimization which previously took time $\tilde{O}\left(n^{4} \mathrm{EO}+n^{5}\right)$ (and now takes $\tilde{O}\left(n^{2} \mathrm{EO}+n^{3}\right)$ time using our result in Chapter 10). Readers may wish to substitute $h=\tilde{O}\left(n^{2} \mathrm{EO}+n^{3}\right)$ when reading the table.

| Years | Authors | Running times |
| :---: | :---: | :---: |
| 1978 | Fujishige [92] | not stated |
| 1981 | Grötschel, Lovász, Schrijver[110] | weakly polynomial |
| 1982 | Zimmermann [276] | not stated |
| 1984 | Barahona, Cunningham [30] | not stated |
| 1985 | Cunningham, Frank [62] | $\rightarrow O\left(n^{4} h \log C\right)$ |
| 1987 | Fujishige [93] | not stated |
| 1987 | Frank, Tardos [91] | strongly polynomial |
| 1988 | Cui, Fujishige [265] | not stated |
| 1989 | Fujishige, Röck, Zimmermann[95] | $\rightarrow O\left(n^{6} h \log n\right)$ |
| 1991 | Chung, Tcha [52] | not stated |
| 1992 | Zimmermann [277] | not stated |
| 1993 | McCormick, Ervolina [190] | $O\left(n^{7} h^{*} \log n C U\right)$ |
| 1994 | Wallacher, Zimmermann [266] | $O\left(n^{8} h \log n C U\right)$ |
| 1997 | Iwata [117] | $O\left(n^{7} h \log U\right)$ |
| 1998 | Iwata, McCormick, Shigeno [123] | $O\left(n^{4} h \min \left\{\log n C, n^{2} \log n\right\}\right)$ |
| 1999 | Iwata, McCormick, Shigeno [124] | $O\left(n^{6} h \min \left\{\log n U, n^{2} \log n\right\}\right)$ |
| 1999 | Fleischer, Iwata, McCormick[89] | $O\left(n^{4} h \min \left\{\log U, n^{2} \log n\right\}\right)$ |
| 1999 | Iwata, McCormick, Shigeno [122] | $O\left(n^{4} h \min \left\{\log C, n^{2} \log n\right\}\right)$ |
| 2000 | Fleischer, Iwata [88] | $O\left(m n^{5} \log n U \cdot \mathrm{EO}\right)$ |
| - | This Chapter | $O\left(n^{2} \log n C U \cdot \operatorname{EO}+n^{3} \log { }^{0(1)} n C U\right)$ |

Figure 9-1: Previous algorithms for Submodular Flow with $n$ vertices, maximum $\operatorname{cost} C$ and maximum capacity $U$. The factor $h$ is the time for an exchange capacity oracle, $h^{*}$ is the time for a "more complicated exchange capacity oracle" and EO is the time for evaluation oracle of the submodular function. The arrow, $\rightarrow$, indicates that it used currently best maximum submodular flow algorithm as subroutine which was non-existent at the time of the publication.

The previous fastest weakly polynomial algorithms for submodular flow are by [122, 88, 89], which take time $\tilde{O}\left(n^{6} \mathrm{EO}+n^{7}\right)$ and $O\left(m n^{5} \log n U \cdot \mathrm{EO}\right)$, assuming $h=\tilde{O}\left(n^{2} \mathrm{EO}+n^{3}\right)$. Our algorithm for submodular flow has a running time of $\tilde{O}\left(n^{2} \mathrm{EO}+n^{3}\right)$, which is significantly faster by roughly a factor of $\tilde{O}\left(n^{4}\right)$.

For strongly polynomial algorithms, our results do not yield a speedup but we remark that our faster strongly polynomial algorithm for submodular minimization in Chapter 10 improves the previous algorithms by a factor of $\tilde{O}\left(n^{2}\right)$ as a corollary (because $h$ requires solving an instance of submodular minimization).

## ■ 9.1.3 Overview

After providing covering some preliminaries on convex analysis in Section 9.2 we split the remainder of Chapter 9 into Section 9.3 and Section 9.4. In Section 9.3 we cover our algorithm for convex optimization using an approximate subgradient oracle (Section 9.3.1) as well as our technique on using duality to decrease dimensions and improve the running time of semidefinite programming (Section 9.3.2). In Section 9.4 we provide our technique for using minimization oracles to minimize functions over the intersection of convex sets and provide several applications including matroid intersection (Section 9.4.2), submodular flow (Section 9.4.3), and minimizing a linear function over the intersection of an affine subspace and a convex set (Section 9.4.4).

## ■ 9.2 Preliminaries

In this section we review basic facts about convex functions that we use throughout Chapter 9. We also introduce two oracles that we use throughout Chapter 9, i.e. subgradient and optimization oracles, and provide some basic reductions between them. Note that we have slightly extended some definitions and facts to accommodate for the noisy separation oracles used in this chapter.

First we recall the definition of strong convexity
Definition 9.2.1 (Strong Convexity ). A real valued function $f$ on a convex set $\Omega$ is $\alpha$-strongly convex if for any $\boldsymbol{x}, \boldsymbol{y} \in \Omega$ and $t \in[0,1]$, we have

$$
f(t \boldsymbol{x}+(1-t) \boldsymbol{y})+\frac{1}{2} \alpha t(1-t)\|\boldsymbol{x}-\boldsymbol{y}\|^{2} \leq t f(\boldsymbol{x})+(1-t) f(\boldsymbol{y}) .
$$

Next we define an approximate subgradient.
Definition 9.2.2 (Subgradient). For any convex function $f$ on a convex set $\Omega$, the $\delta$-subgradients of $f$ at $x$ are defined to be

$$
\partial_{\delta} f(\boldsymbol{x}) \stackrel{\text { def }}{=}\{\boldsymbol{g} \in \Omega: f(\boldsymbol{y})+\delta \geq f(\boldsymbol{x})+\langle\boldsymbol{g}, \boldsymbol{y}-\boldsymbol{x}\rangle \text { for all } \boldsymbol{y} \in \Omega\} .
$$

Here we provide some basic facts regarding convexity and subgradients. These statements are natural extensions of well known facts regarding convex functions and their proof can be found in any standard textbook on convex optimization.

Fact 9.2.3. For any convex set $\Omega$ and $\boldsymbol{x}$ be a point in the interior of $\Omega$, we have the following:

1. If $f$ is convex on $\Omega$, then $\partial_{0} f(\boldsymbol{x}) \neq \emptyset$ and $\partial_{s} f(\boldsymbol{x}) \subseteq \partial_{t} f(\boldsymbol{x})$ for all $0 \leq s \leq t$. Otherwise, we have $\|\boldsymbol{g}\|_{2}>\frac{1}{2} \sqrt{\frac{\delta}{D}}$. For any $f(\boldsymbol{y}) \leq f(\boldsymbol{x})$, we have $\delta \geq\langle\boldsymbol{g}, \boldsymbol{y}-\boldsymbol{x}\rangle$ and hence
2. If $f$ is a differential convex function on $\Omega$, then $\nabla f(\boldsymbol{x}) \in \partial_{0} f(\boldsymbol{x})$.
3. If $f_{1}$ and $f_{2}$ are convex function on $\Omega$, $\boldsymbol{g}_{1} \in \partial_{\delta_{1}} f_{1}(\boldsymbol{x})$ and $\boldsymbol{g}_{2} \in \partial_{\delta_{2}} f_{1}(\boldsymbol{x})$, then $\alpha \boldsymbol{g}_{1}+\beta \boldsymbol{g}_{2} \in$ $\partial_{\alpha \delta_{1}+\beta \delta_{2}}\left(\boldsymbol{g}_{1}+\boldsymbol{g}_{2}\right)(\boldsymbol{x})$.
4. If $f$ is $\alpha$-strongly convex on $\Omega$ with minimizer $x^{*}$, then for any $\boldsymbol{y}$ with $f(\boldsymbol{y}) \leq f\left(\boldsymbol{x}^{*}\right)+\varepsilon$, we have $\frac{1}{2} \alpha\left\|\boldsymbol{x}^{*}-\boldsymbol{y}\right\|^{2} \leq \varepsilon$.

Next we provide a reduction from subgradients to separation oracles. We will use this reduction several times in Chapter 9 to simplify our construction of separation oracles.

Lemma 9.2.4. Let $f$ be a convex function. Suppose we have $\boldsymbol{x}$ and $\boldsymbol{g} \in \partial_{\delta} f(\boldsymbol{x})$ with $\|\boldsymbol{x}\|_{2} \leq 1 \leq D$ and $\delta \leq 1$. If $\|\boldsymbol{g}\|_{2} \leq \frac{1}{2} \sqrt{\frac{\delta}{D}}$, then $f(\boldsymbol{x}) \leq \min _{\left\|\boldsymbol{y}_{2}\right\|_{2} \leq D} f(\boldsymbol{y})+2 \sqrt{\delta D}$ and if $\|\boldsymbol{g}\|_{2} \leq \frac{1}{2} \sqrt{\frac{\delta}{D}}$ then

$$
\left\{\|\boldsymbol{y}\|_{2} \leq D: f(\boldsymbol{y}) \leq f(\boldsymbol{x})\right\} \subset\left\{\boldsymbol{y}: \boldsymbol{d}^{T} \boldsymbol{y} \leq \boldsymbol{d}^{T} \boldsymbol{x}+2 \sqrt{\delta D}\right\}
$$

with $\boldsymbol{d}=\boldsymbol{g} /\|\boldsymbol{g}\|_{2}$. Hence, this gives a $(2 \sqrt{\delta D}, 2 \sqrt{\delta D})$-separation oracle on the set $\left\{\|\boldsymbol{x}\|_{2} \leq D\right\}$.
Proof. Let $\boldsymbol{y}$ such that $\|\boldsymbol{y}\|_{2} \leq D$. By the definition of $\delta$-subgradient, we have

$$
f(\boldsymbol{y})+\delta \geq f(\boldsymbol{x})+\langle\boldsymbol{g}, \boldsymbol{y}-\boldsymbol{x}\rangle .
$$

If $\|\boldsymbol{g}\| \leq \frac{1}{2} \sqrt{\frac{\delta}{D}}$, then, we have $|\langle\boldsymbol{g}, \boldsymbol{y}-\boldsymbol{x}\rangle| \leq \sqrt{\delta D}$ because $\|\boldsymbol{x}\| \leq D$ and $\|\boldsymbol{y}\|_{2} \leq D$. Therefore,

$$
\min _{\|\boldsymbol{y}\|_{2} \leq D} f(\boldsymbol{y})+2 \sqrt{\delta D} \geq f(\boldsymbol{x})
$$

Otherwise, we have $\|\boldsymbol{g}\|_{2}>\frac{1}{2} \sqrt{\frac{\delta}{D}}$. For any $f(\boldsymbol{y}) \leq f(\boldsymbol{x})$, we have $\delta \geq\langle\boldsymbol{g}, \boldsymbol{y}-\boldsymbol{x}\rangle$ and hence

$$
2 \sqrt{\delta D} \geq\left\langle\frac{\boldsymbol{g}}{\|\boldsymbol{g}\|}, \boldsymbol{y}-\boldsymbol{x}\right\rangle
$$

At several times in Chapter 9 we will wish to construct subgradient oracles or separation oracles given only the ability to approximately maximize a linear function over a convex set. In the remainder of this section we formally define such a optimization oracle and prove this equivalence.

Definition 9.2.5 (Optimization Oracle). Given a convex set $K$ and $\delta>0$ a $\delta$-optimization oracle for $K$ is a function on $\mathbb{R}^{n}$ such that for any input $\boldsymbol{c} \in \mathbb{R}^{n}$, it outputs $\boldsymbol{y}$ such that

$$
\max _{\boldsymbol{x} \in K}\langle\boldsymbol{c}, \boldsymbol{x}\rangle \leq\langle\boldsymbol{c}, \boldsymbol{y}\rangle+\delta .
$$

We denote by $\mathrm{OO}_{\delta}(K)$ the time complexity of this oracle.
Lemma 9.2.6. Given a convex set $K$, any $\varepsilon$-optimization oracle for $K$ is a $\varepsilon$-subgradient oracle for $f(\boldsymbol{c}) \stackrel{\text { def }}{=} \max _{\boldsymbol{x} \in K}\langle\boldsymbol{c}, \boldsymbol{x}\rangle$.

Proof. Let $\boldsymbol{x}_{c}$ be the output of $\varepsilon$-optimization oracle on the cost vector $\boldsymbol{c}$. We have

$$
\max _{x \in K}\langle\boldsymbol{c}, \boldsymbol{x}\rangle \leq\left\langle\boldsymbol{c}, \boldsymbol{x}_{c}\right\rangle+\varepsilon .
$$

Hence, for all $\boldsymbol{d}$, we have and therefore

$$
\left\langle\boldsymbol{x}_{c}, \boldsymbol{d}-\boldsymbol{c}\right\rangle+f(\boldsymbol{c}) \leq f(\boldsymbol{d})+\varepsilon .
$$

Hence, $\boldsymbol{x}_{c} \in \partial_{\delta} f(\boldsymbol{c})$.
Combining these lemmas shows that having an $\varepsilon$-optimization oracle for a convex set $K$ contained in a ball of radius $D$ yields a $O(\sqrt{D \varepsilon}, \sqrt{D \varepsilon})$ separation oracle for $\max _{x \in K}\langle\boldsymbol{c}, \boldsymbol{x}\rangle$. We use these ideas to construction separation oracles throughout Chapter 9.

## ■ 9.3 Convex Optimization

In this section we show how to apply our cutting plane method to efficiently solve problems in convex optimization. First, in Section 9.3.1 we show how to use our result to minimize a convex function given an approximate subgradient oracle. Then, in Section 9.3 .2 we illustrate how this result can be used to obtain both primal and dual solutions for a standard convex optimization problems. In particular, we show how our result can be used to obtain improved running times for semidefinite programming across a range of parameters.

## - 9.3.1 From Feasibility to Optimization

In this section we consider the following standard optimization problem. We are given a convex function $f: \mathbb{R}^{n} \rightarrow \mathbb{R} \cup\{+\infty\}$ and we want to find a point $\boldsymbol{x}$ that approximately solves the minimization
problem

$$
\min _{\boldsymbol{x} \in \mathbb{R}^{n}} f(\boldsymbol{x})
$$

given only a subgradient oracle for $f$.
Here we show how to apply the cutting plane method from Chapter 8 turning the small width guarantee of the output of that algorithm into a tool to find an approximate minimizer of $f$. Our result is applicable to any convex optimization problem armed with a separation or subgradient oracle. This result will serve as the foundation for many of our applications in this chapter.

Our approach is an adaptation of Nemiroski's method [202] which applies the cutting plane method to solve convex optimiziation problems, with only minimal assumption on the cutting plane method. The proof here is a generalization that accommodates for the noisy separation oracle. In the remainder of this subsection we provide a key definition we will use in our algorithm (Defintion 9.3.1), provide our main result (Theorem 9.3.2), and conclude with a brief discussion of this result.

Definition 9.3.1. For any compact set $K$, we define the minimum width by $\operatorname{MinWidth}(K) \stackrel{\text { def }}{=}$ $\min _{\|\boldsymbol{a}\|_{2}=1} \max _{\boldsymbol{x}, \boldsymbol{y} \in K}\langle\boldsymbol{a}, \boldsymbol{x}-\boldsymbol{y}\rangle$.

Theorem 9.3.2. Let $f$ be a convex function on $\mathbb{R}^{n}$ and $\Omega$ be a convex set that contains a minimizer of $f$. Suppose we have a $(\eta, \delta)$-separation oracle for $f$ and $\Omega$ is contained inside $B_{\infty}(R)$. Using $B_{\infty}(R)$ as the initial polytope for our Cutting Plane Method, for any $0<\alpha<1$, we can compute $\boldsymbol{x} \in \mathbb{R}^{n}$ such that

$$
\begin{equation*}
f(\boldsymbol{x})-\min _{\boldsymbol{y} \in \Omega} f(\boldsymbol{y}) \leq \eta+\alpha\left(\max _{\boldsymbol{y} \in \Omega} f(\boldsymbol{y})-\min _{\boldsymbol{y} \in \Omega} f(\boldsymbol{y})\right) . \tag{9.1}
\end{equation*}
$$

with an expected running time of

$$
O\left(n \mathrm{SO}_{\eta, \delta}(f) \log \left(\frac{n \kappa}{\alpha}\right)+n^{3} \log ^{O(1)}\left(\frac{n \kappa}{\alpha}\right)\right)
$$

where $\delta=\Theta\left(\frac{\alpha \operatorname{MinWidth}(\Omega)}{n^{3 / 2} \ln (\kappa)}\right)$ and $\kappa=\frac{R}{\operatorname{MinWidth}(\Omega)}$. Furthermore, we only need the oracle defined on the set $B_{\infty}(R)$.

Proof. Let $\boldsymbol{x}^{*} \in \arg \min _{\boldsymbol{x} \in \Omega} f(\boldsymbol{x})$. Since $B_{\infty}(R) \supset \Omega$ contains a minimizer of $f$, by the definition of $(\eta, \delta)$-separation oracles, our Cutting Plane Method (Theorem 8.3.26) either returns a point $\boldsymbol{x}$ that is almost optimal or returns a polytope $P$ of small width. In the former case we have a point $\boldsymbol{x}$ such that $f(\boldsymbol{x}) \leq \min _{\boldsymbol{y}} f(\boldsymbol{y})+\eta$. Hence, the error is clearly at most $\eta+\alpha\left(\max _{\boldsymbol{z} \in \Omega} f(\boldsymbol{z})-\min _{\boldsymbol{x} \in \Omega} f(\boldsymbol{x})\right)$ as desired. Consequently, we assume the latter case.

Theorem 8.3.26 shows MinWidth $(P)<C n \varepsilon \ln (R / \varepsilon)$ for some universal constant $C$. Picking

$$
\begin{equation*}
\varepsilon=C^{\prime} \frac{\alpha \operatorname{MinWidth}(\Omega)}{n \ln \left(\frac{n \kappa}{\alpha}\right)} \tag{9.2}
\end{equation*}
$$

for small enough constant $C^{\prime}$, we have $\operatorname{MinWidth}\left(P^{(i)}\right)<\alpha \operatorname{MinWidth}(\Omega)$. Let $\Omega^{\alpha}=\boldsymbol{x}^{*}+\alpha\left(\Omega-\boldsymbol{x}^{*}\right)$, namely, $\Omega^{\alpha}=\left\{\boldsymbol{x}^{*}+\alpha\left(\boldsymbol{z}-\boldsymbol{x}^{*}\right): \boldsymbol{z} \in \Omega\right\}$. Then, we have

$$
\operatorname{MinWidth}\left(\Omega^{\alpha}\right)=\alpha \operatorname{MinWidth}(\Omega)>\operatorname{MinWidth}(P)
$$

Therefore, $\Omega^{\alpha}$ is not a subset of $P^{(i)}$ and hence there is some point $\boldsymbol{y} \in \Omega^{\alpha} \backslash P$. Since $\Omega^{\alpha} \subseteq \Omega \subseteq B_{\infty}(R)$, we know that $\boldsymbol{y}$ does not violate any of the constraints of $P^{(0)}$ and therefore must violate one of the constraints added by querying the separation oracle. Therefore, for some $j \leq i$, we have

$$
\left\langle\boldsymbol{c}^{(j-1)}, \boldsymbol{y}\right\rangle>\left\langle\boldsymbol{c}^{(j-1)}, \boldsymbol{x}^{(j-1)}\right\rangle+c_{s} \varepsilon / \sqrt{n} .
$$

By the definition of $\left(\eta, c_{s} \varepsilon / \sqrt{n}\right)$-separation oracle (Definition 2.3.6), we have $f(\boldsymbol{y})>f\left(\boldsymbol{x}^{(j-1)}\right)$. Since $\boldsymbol{y} \in \Omega^{\alpha}$, we have $\boldsymbol{y}=(1-\alpha) \boldsymbol{x}^{*}+\alpha \boldsymbol{z}$ for some $\boldsymbol{z} \in \Omega$. Thus, the convexity of $f$ implies that

$$
f(\boldsymbol{y}) \leq(1-\alpha) f\left(\boldsymbol{x}^{*}\right)+\alpha f(\boldsymbol{z}) .
$$

Therefore, we have

$$
\min _{1 \leq k \leq i} f\left(\boldsymbol{x}^{(k)}\right)-\min _{\boldsymbol{x} \in \Omega} f(\boldsymbol{x})<f(\boldsymbol{y})-f\left(\boldsymbol{x}^{*}\right) \leq \alpha\left(\max _{\boldsymbol{z} \in \Omega} f(\boldsymbol{z})-\min _{\boldsymbol{x} \in \Omega} f(\boldsymbol{x})\right) .
$$

Hence, we can simply output the best $\boldsymbol{x}$ among all $\boldsymbol{x}^{(j)}$ and in either case $\boldsymbol{x}$ satisfies (9.1).
Note that we need to call $(\eta, \delta)$-separation oracle with $\delta=\Omega(\varepsilon / \sqrt{n})$ to ensure we do not cut out $\boldsymbol{x}^{*}$. Theorem 8.3.26 shows that the algorithm takes $O\left(n \mathrm{SO}_{\eta, \delta}(f) \log (n R / \varepsilon)+n^{3} \log ^{O(1)}(n R / \varepsilon)\right)$ expected time, as promised. Furthermore, the oracle needs only be defined on $B_{\infty}(R)$ as our cutting plane method guarantees $\boldsymbol{x}^{(k)} \in B_{\infty}(R)$ for all $k$ (although if needed, an obvious separating hyperplane can be returned for a query point outside $\left.B_{\infty}(R)\right)$.

Observe that this algorithm requires no information about $\Omega$ (other than that $\Omega \subseteq B_{\infty}(R)$ ) and does not guarantee that the output is in $\Omega$. Hence, even though $\Omega$ can be complicated to describe, the algorithm still gives a guarantee related to the gap $\max _{\boldsymbol{x} \in \Omega} f(\boldsymbol{x})-\min _{\boldsymbol{x} \in \Omega} f(\boldsymbol{x})$. For specific applications, it is therefore advantageous to pick a $\Omega$ as large as possible while the bound on function value is as small as possible.

Before indulging into specific applications, we remark on the dependence on $\kappa$. Using John's ellipsoid, it can be shown that any convex set $\Omega$ can be transformed linearly such that (1) $B_{\infty}$ (1) contains $\Omega$ and, (2) $\operatorname{MinWidth}(\Omega)=\Omega\left(n^{-3 / 2}\right)$. In other words, $\kappa$ can be effectively chosen as $O\left(n^{3 / 2}\right)$. Therefore if we are able to find such a linear transformation, the running time is simply $O\left(n \mathrm{SO}(f) \log (n / \alpha)+n^{3} \log ^{O(1)}(n / \alpha)\right)$. Often this can be done easily using the structure of the particular problem and the running time does not depend on the size of domain at all.

## ■ 9.3.2 Duality and Semidefinite Programming

In this section we illustrate how our result in Section 9.3.1 can be used to obtain both primal and dual solutions for standard problems in convex optimization. In particular we show how to obtain improved running times for semidefinite programming.

To explain our approach, consider the following minimax problem

$$
\begin{equation*}
\min _{\boldsymbol{y} \in Y} \max _{\boldsymbol{x} \in X}\langle\boldsymbol{A} \boldsymbol{x}, \boldsymbol{y}\rangle+\langle\boldsymbol{c}, \boldsymbol{x}\rangle+\langle\boldsymbol{d}, \boldsymbol{y}\rangle \tag{9.3}
\end{equation*}
$$

where $\boldsymbol{x} \in \mathbb{R}^{m}$ and $\boldsymbol{y} \in \mathbb{R}^{n}$. When $m \gg n$, solving this problem by directly using Chapter 8 could lead to an inefficient algorithm with running time at least $m^{3}$. In many situations, for any fixed $\boldsymbol{y}$, the problem $\max _{\boldsymbol{x} \in X}\langle\boldsymbol{A} \boldsymbol{x}, \boldsymbol{y}\rangle$ is very easy and hence one can use it as a separation oracle and apply Chapter 8 and this would gives a running time almost independent of $m$. However, this would only give us the $\boldsymbol{y}$ variable and it is not clear how to recover $\boldsymbol{x}$ variable from it.

In this section we show how to alleviate this issue and give semidefinite programming (SDP) as a concrete example of how to apply this general technique. We do not write down the general version as the running time of the technique seems to be problem specific and faster SDP is already an interesting application.

For the remainder of this section we focus on the semidefinite programming (SDP) problem:

$$
\begin{equation*}
\max _{\boldsymbol{X} \succeq \mathbf{0}} \boldsymbol{C} \bullet \boldsymbol{X} \text { s.t. } \boldsymbol{A}_{i} \bullet \boldsymbol{X}=b_{i} \tag{9.4}
\end{equation*}
$$

and its dual

$$
\begin{equation*}
\min _{\boldsymbol{y}} \boldsymbol{b}^{T} \boldsymbol{y} \text { s.t. } \sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i} \succeq \boldsymbol{C} \tag{9.5}
\end{equation*}
$$

where $\boldsymbol{X}, \boldsymbol{C}, \boldsymbol{A}_{i}$ are $m \times m$ symmetric matrices and $\boldsymbol{b}, \boldsymbol{y} \in \mathbb{R}^{n}$. Our approach is partially inspired by one of the key ideas of $[115,152]$. These results write down the dual SDP in the form

$$
\begin{equation*}
\min _{y} \boldsymbol{b}^{T} \boldsymbol{y}-K \min \left(\lambda_{\min }\left(\sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i}-\boldsymbol{C}\right), 0\right) \tag{9.6}
\end{equation*}
$$

for some large number $K$ and use non-smooth optimization techniques to solve the dual SDP problem. Here, we follow the same approach but instead write it as a max-min problem $\min _{y} f_{K}(\boldsymbol{y})$ where

$$
\begin{equation*}
f_{K}(\boldsymbol{y})=\max _{\operatorname{Tr} \boldsymbol{X} \leq K, \boldsymbol{X} \succeq \mathbf{0}}\left(\boldsymbol{b}^{T} \boldsymbol{y}+\left\langle\boldsymbol{X}, \boldsymbol{C}-\sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i}\right\rangle\right) . \tag{9.7}
\end{equation*}
$$

Thus the SDP problem in fact assumes the form (9.3) and many ideas in this section can be generalized to the minimax problem (9.3).

To get a dual solution, we notice that the cutting plane method maintains a subset of the primal feasible solution $\operatorname{conv}\left(\boldsymbol{X}_{i}\right)$ such that

$$
\min _{\boldsymbol{y}} \boldsymbol{b}^{T} \boldsymbol{y}+\max _{\operatorname{Tr} \boldsymbol{X} \leq K, \boldsymbol{X} \succeq \mathbf{0}}\left\langle\boldsymbol{X}, \boldsymbol{C}-\sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i}\right\rangle \sim \min _{\boldsymbol{y}} \boldsymbol{b}^{T} \boldsymbol{y}+\max _{\boldsymbol{X} \in \operatorname{conv}\left(\boldsymbol{X}_{i}\right)}\left\langle\boldsymbol{X}, \boldsymbol{C}-\sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i}\right\rangle .
$$

Applying minimax theorem, this shows that there exists an approximation solution $\boldsymbol{X}$ in $\operatorname{conv}\left(\boldsymbol{X}_{i}\right)$ for the primal problem. Hence, we can restrict the primal SDP on the polytope $\operatorname{conv}\left(\boldsymbol{X}_{i}\right)$, this reduces the primal SDP into a linear program which can be solved very efficiently. This idea of getting primal/dual solution from the cutting plane method is quite general and is the main purpose of this example. As a by-product, we have a faster SDP solver in both primal and dual! We remark that this idea has been used as a heuristic to obtain [151] for getting the primal SDP solution and our contribution here is mainly the asymptotic time analysis.

We first show how to construct the separation oracle for SDP. For that we need to compute smallest eigenvector of a matrix. Below, for completeness we provide a folklore result showing we can do this using fast matrix multiplication.

Lemma 9.3.3. Given a $n \times n$ symmetric matrix $\boldsymbol{Y}$ such that $-R \boldsymbol{I} \preceq \boldsymbol{Y} \preceq R \boldsymbol{I}$, for any $\varepsilon>0$, with high probability in $n$ in time $O\left(n^{\omega+o(1)} \log ^{O(1)}(R / \varepsilon)\right)$ we can find a unit vector $\boldsymbol{u}$ such that $\boldsymbol{u}^{T} \boldsymbol{Y} \boldsymbol{u} \geq$ $\lambda_{\max }(\boldsymbol{Y})-\varepsilon$.

Proof. Let $\boldsymbol{B} \stackrel{\text { def }}{=} \frac{1}{R} \boldsymbol{Y}+\boldsymbol{I}$. Note that $\boldsymbol{B} \succeq \mathbf{0}$. Now, we consider the repeated squaring $\boldsymbol{B}_{0}=\boldsymbol{B}$ and $\boldsymbol{B}_{k+1}=\frac{\boldsymbol{B}_{k}^{2}}{\operatorname{Tr} \boldsymbol{B}_{k}^{2}}$. Let $0 \leq \lambda_{1} \leq \lambda_{2} \leq \cdots \leq \lambda_{n}$ be the eigenvalues of $\boldsymbol{B}$ and $\boldsymbol{v}_{i}$ be the corresponding eigenvectors. Then, it is easy to see the the eigenvalues of $\boldsymbol{B}_{k}$ are $\frac{\lambda_{i}^{2^{k}}}{\sum_{i=1}^{n} \lambda_{i}^{\lambda^{k}}}$.

Let $\boldsymbol{q}$ be a random unit vector and $\boldsymbol{r} \stackrel{\text { def }}{=} \boldsymbol{B}_{k} \boldsymbol{q}$. Now $\boldsymbol{q}=\sum \alpha_{i} \boldsymbol{v}_{i}$ for some $\alpha_{i}$ such that $\sum \alpha_{i}^{2}=1$. Letting

$$
\boldsymbol{p}=\frac{\sum_{\lambda_{i}>(1-\delta) \lambda_{n}} \alpha_{i} \lambda_{i}^{2^{k}} \boldsymbol{v}_{i}}{\sum_{i=1}^{n} \lambda_{i}^{2^{k}}}
$$

we have

$$
\|\boldsymbol{r}-\boldsymbol{p}\|_{2}=\left\|\frac{\sum_{\lambda_{i} \leq(1-\delta) \lambda_{n}} \alpha_{i} \lambda_{i}^{2^{k}} \boldsymbol{v}_{i}}{\sum_{i=1}^{n} \lambda_{i}^{2^{k}}}\right\|_{2} \leq \frac{\sum_{\lambda_{i} \leq(1-\delta) \lambda_{n}} \lambda_{i}^{2^{k}}}{\sum_{i=1}^{n} \lambda_{i}^{2^{k}}} \leq(1-\delta)^{2^{k}}{ }_{n}
$$

Letting $k=\log _{2}\left(\frac{\log \left(n^{3 / 2} / \delta\right)}{\delta}\right)$, we have $\|\boldsymbol{r}-\boldsymbol{p}\|_{2} \leq \delta / \sqrt{n}$. Since $\mathbf{0} \preceq \boldsymbol{B} \preceq 2 \boldsymbol{I}$, we have

$$
\begin{aligned}
\sqrt{\boldsymbol{r}^{T} \boldsymbol{B r}} & \geq \sqrt{\boldsymbol{p}^{T} \boldsymbol{B} \boldsymbol{p}}-\sqrt{(\boldsymbol{r}-\boldsymbol{p})^{T} \boldsymbol{B}(\boldsymbol{r}-\boldsymbol{p})} \\
& \geq \sqrt{\boldsymbol{p}^{T} \boldsymbol{B} \boldsymbol{p}}-2 \delta / \sqrt{n} .
\end{aligned}
$$

Note that $\boldsymbol{p}$ involves only eigenvectors between $(1-\delta) \lambda_{n}$ to $\lambda_{n}$. Hence, we have

$$
\sqrt{\boldsymbol{r}^{T} \boldsymbol{B} \boldsymbol{r}} \geq \sqrt{(1-\delta) \lambda_{n}}\|\boldsymbol{p}\|_{2}-2 \delta / \sqrt{n}
$$

With constant probability, we have $\alpha_{n}=\Omega(1 / \sqrt{n})$. Hence, we have $\|\boldsymbol{p}\|_{2}=\Omega(1 / \sqrt{n})$. Using $\boldsymbol{B} \preceq 2 \boldsymbol{I}$ and $\|\boldsymbol{p}\|_{2} \geq\|\boldsymbol{r}\|_{2}-\delta / \sqrt{n}$ we have that so long as $\delta$ is a small enough universal constant

$$
\begin{aligned}
\frac{\sqrt{\boldsymbol{r}^{T} \boldsymbol{B r}}}{\|\boldsymbol{r}\|_{2}} & \geq \frac{\sqrt{(1-\delta) \lambda_{n}}\|\boldsymbol{p}\|_{2}-2 \delta / \sqrt{n}}{\|\boldsymbol{p}\|_{2}+\delta / \sqrt{n}} \\
& =(1-O(\delta)) \sqrt{\lambda_{n}}-O(\delta) \\
& =\sqrt{\lambda_{n}}-O(\delta \sqrt{R}) .
\end{aligned}
$$

Therefore, we have $\frac{r^{T} \boldsymbol{Y} r}{\|\boldsymbol{r}\|^{2}} \geq \lambda_{\max }(\boldsymbol{Y})-O(R \delta)$. Hence, we can find vector $\boldsymbol{r}$ by computing $k$ matrix multiplications. [66] showed that fast matrix multiplication is stable under Frobenius norm, i.e., for any $\eta>0$, using $O(\log (n / b))$ bits, we can find $\boldsymbol{C}$ such that $\|\boldsymbol{C}-\boldsymbol{A} \boldsymbol{B}\|_{F} \leq \frac{1}{b}\|\boldsymbol{A}\|\|\boldsymbol{B}\|$ in time $O\left(n^{\omega+\eta}\right)$ where $\omega$ is the matrix multiplicative constant. Hence, this algorithm takes only $O\left(n^{\omega+o(1)} \log ^{O(1)}\left(\delta^{-1}\right)\right)$ time. The result follows from renormalizing the vector $r$, repeating the algorithm $O(\log n)$ times to boost the probability and taking $\delta=\Omega(\varepsilon / R)$.

The following lemma shows how to compute a separation for $f_{K}$ defined in (9.7).
Lemma 9.3.4. Suppose that $\left\|\boldsymbol{A}_{i}\right\|_{F} \leq M$ and $\|\boldsymbol{C}\|_{F} \leq M$. For any $0<\varepsilon<1$ and $\boldsymbol{y}$ with $\|\boldsymbol{y}\|_{2}=$ $O(L)$, with high probability in $m$, we can compute a $(\varepsilon, \varepsilon)$-separation of $f_{K}$ on $\left\{\|\boldsymbol{x}\|_{2} \leq L\right\}$ at $\boldsymbol{y}$ in time $O\left(S+m^{\omega+o(1)} \log ^{O(1)}(n K M L / \varepsilon)\right)$ where where $S$ is the sparsity of the problem defined as $\mathrm{nnz}(\boldsymbol{C})+\sum_{i=1}^{n} \mathrm{nnz}\left(\boldsymbol{A}_{i}\right)$.

Proof. Note that $-O(n M L) \boldsymbol{I} \preceq \boldsymbol{C}-\sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i} \preceq O(n M L) \boldsymbol{I}$. Using Lemma 9.3.3, we can find a vector $\boldsymbol{v}$ with $\|\boldsymbol{v}\|_{2}=K$ in time $O\left(m^{\omega+o(1)} \log ^{O(1)}(n K M L / \delta)\right)$ such that

$$
\begin{equation*}
\boldsymbol{v}^{T}\left(\boldsymbol{C}-\sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i}\right) \boldsymbol{v} \geq \max _{\operatorname{Tr} \boldsymbol{X} \leq K, \boldsymbol{X} \succeq \mathbf{0}}\left\langle\boldsymbol{X}, \boldsymbol{C}-\sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i}\right\rangle-\delta . \tag{9.8}
\end{equation*}
$$

In other words, we have a $\delta$-optimization oracle for the function $f_{K}$. Lemma 9.2 .6 shows this yields a $\delta$-subgradient oracle and Lemma 9.2.4 then shows this yields a $(O(\sqrt{\delta L}), O(\sqrt{\delta L})$-separation oracle on the set $\left\{\|x\|_{2} \leq L\right\}$. By picking $\delta=\varepsilon^{2} / L$, we have the promised oracle.

With the separation oracle in hand, we are ready to give the algorithm for SDP:
Theorem 9.3.5. Given a primal-dual semidefinite programming problem in the form (9.4) and (9.5), suppose that for some $M \geq 1$ we have

1. $\|b\|_{2} \leq M,\|\boldsymbol{C}\|_{F} \leq M$ and $\left\|\boldsymbol{A}_{i}\right\|_{F} \leq M$ for all $i$.
2. The primal feasible set lies inside the region $\operatorname{Tr} \boldsymbol{X} \leq M$.
3. The dual feasible set lies inside the region $\|\boldsymbol{y}\|_{\infty} \leq M$.

Let OPT be the optimum solution of (9.4) and (9.5). Then, with high probability, we can find $\boldsymbol{X}$ and $\boldsymbol{y}$ such that

1. $\boldsymbol{X} \succeq \mathbf{0}, \operatorname{Tr} \boldsymbol{X}=O(M), \sum_{i}\left|b_{i}-\left\langle\boldsymbol{X}, \boldsymbol{A}_{i}\right\rangle\right| \leq \varepsilon$ for all $i$ and $\boldsymbol{C} \bullet \boldsymbol{X} \geq \mathrm{OPT}-\varepsilon$.
2. $\|\boldsymbol{y}\|_{\infty}=O(M), \sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i} \succeq \boldsymbol{C}-\varepsilon \boldsymbol{I}$ and $\boldsymbol{b}^{T} \boldsymbol{y} \leq \mathrm{OPT}+\varepsilon$.
in expected time $O\left(\left(n S+n^{3}+n m^{\omega+o(1)}\right) \log ^{O(1)}\left(\frac{n M}{\varepsilon}\right)\right)$ where $S$ is the sparsity of the problem defined as $\mathrm{nnz}(\boldsymbol{C})+\sum_{i=1}^{n} \mathrm{nnz}\left(\boldsymbol{A}_{i}\right)$ and $\omega$ is the fast matrix multiplication constant.

Proof. Let $K \geq M$ be some parameter to be determined. Since the primal feasible set is lies inside the region $\operatorname{Tr} \boldsymbol{X} \leq M \leq K$, we have

$$
\begin{aligned}
\min _{\sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i} \succeq \boldsymbol{C}} \boldsymbol{b}^{T} \boldsymbol{y} & =\underset{\boldsymbol{X} \succeq \mathbf{0}, \operatorname{Tr} \boldsymbol{X} \leq K, \boldsymbol{A}_{i} \bullet \boldsymbol{X}=b_{i}}{ } \boldsymbol{C} \bullet \boldsymbol{X} \\
& =\max _{\boldsymbol{X} \succeq \mathbf{0}, \operatorname{Tr} \boldsymbol{X} \leq K} \min _{\boldsymbol{y}} \boldsymbol{C} \bullet \boldsymbol{X}-\sum_{i} y_{i}\left(\boldsymbol{A}_{i} \bullet \boldsymbol{X}-b_{i}\right) \\
& =\min _{\boldsymbol{y}} \max _{\boldsymbol{X} \succeq \mathbf{0}, \operatorname{Tr} \boldsymbol{X} \leq K}\left(\boldsymbol{b}^{T} \boldsymbol{y}+\left(\boldsymbol{C}-\sum_{i} y_{i} \boldsymbol{A}_{i}\right) \bullet \boldsymbol{X}\right) \\
& =\min _{\boldsymbol{y}} f_{K}(\boldsymbol{y}) .
\end{aligned}
$$

Lemma 9.3.4 shows that it takes $\mathrm{SO}_{\delta, \delta}\left(f_{K}\right)=O\left(S+m^{\omega+o(1)} \log (n K M L / \delta)\right)$ time to compute a $(\delta, \delta)$-separation oracle of $f_{K}$ for any point $\boldsymbol{y}$ with $\|\boldsymbol{y}\|_{\infty}=O(L)$ where $L$ is some parameter with $L \geq$ $M$. Taking the radius $R=L$, Theorem 9.3.2 shows that it takes $O\left(n \mathrm{SO}_{\delta, \delta}\left(f_{K}\right) \log \left(\frac{n}{\alpha}\right)+n^{3} \log ^{O(1)}\left(\frac{n}{\alpha}\right)\right)$ expected time with $\delta=\Theta\left(\alpha n^{-3 / 2} L\right)$ to find $\boldsymbol{y}$ such that

$$
f_{K}(\boldsymbol{y})-\min _{\|\boldsymbol{y}\|_{\infty} \leq L} f_{K}(\boldsymbol{y}) \leq \delta+\alpha\left(\max _{\|\boldsymbol{y}\|_{\infty} \leq L} f_{K}(\boldsymbol{y})-\min _{\|\boldsymbol{y}\|_{\infty} \leq L} f_{K}(\boldsymbol{y})\right) \leq \delta+2 \alpha(n M L+2 n K M L) .
$$

Picking $\alpha=\frac{\varepsilon}{7 n M K L}$, we have $f_{K}(\boldsymbol{y}) \leq \min _{\boldsymbol{y}} f_{K}(\boldsymbol{y})+\varepsilon$. Therefore,

$$
\boldsymbol{b}^{T} \boldsymbol{y}+K \max \left(\lambda_{\max }\left(\boldsymbol{C}-\sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i}\right), 0\right) \leq \mathrm{OPT}+\varepsilon
$$

Let $\beta=\max \left(\lambda_{\max }\left(\boldsymbol{C}-\sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i}\right), 0\right)$. Then, we have that $\sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i} \succeq \boldsymbol{C}-\beta \boldsymbol{I}$ and

$$
\begin{aligned}
\boldsymbol{b}^{T} \boldsymbol{y} & \geq \min _{\sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i} \succeq \boldsymbol{C}-\beta \boldsymbol{I}} \boldsymbol{b}^{T} \boldsymbol{y} \\
& =\max _{\boldsymbol{X} \succeq \mathbf{0} \boldsymbol{A}_{i} \boldsymbol{X} \boldsymbol{X}=b_{i}}(\boldsymbol{C}-\beta \boldsymbol{I}) \bullet \boldsymbol{X} \\
& \geq \mathrm{OPT}-\beta M
\end{aligned}
$$

because $\operatorname{Tr} \boldsymbol{X} \leq M$. Hence, we have

$$
\mathrm{OPT}-\beta M+\beta K \leq \boldsymbol{b}^{T} \boldsymbol{y}+K \max \left(\lambda_{\max }\left(\boldsymbol{C}-\sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i}\right), 0\right) \leq \mathrm{OPT}+\varepsilon
$$

Putting $K=M+1$, we have $\beta \leq \varepsilon$. Thus,

$$
\sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i} \succeq \boldsymbol{C}-\varepsilon \boldsymbol{I} .
$$

This gives the result for the dual with the running time $O\left(\left(n S+n^{3}+n m^{\omega+o(1)}\right) \log ^{O(1)}\left(\frac{n M L}{\varepsilon}\right)\right)$.
Our Cutting Plane Method accesses the sub-problem

$$
\max _{\boldsymbol{X} \succeq \mathbf{0}, \operatorname{Tr} \boldsymbol{X} \leq K}\left(\boldsymbol{C}-\sum_{i} y_{i} \boldsymbol{A}_{i}\right) \bullet \boldsymbol{X}
$$

only through the separation oracle. Let $\boldsymbol{z}$ be the output of our Cutting Plane Method and $\left\{\boldsymbol{v}_{i} \boldsymbol{v}_{i}^{T}\right\}_{i=1}^{O(n)}$ be the matrices used to construct the separation for the $O(n)$ hyperplanes the algorithm maintains at the end. Let $\boldsymbol{u}$ be the maximum eigenvector of $\boldsymbol{C}-\sum_{i=1}^{n} z_{i} \boldsymbol{A}_{i}$. Now, we consider a realization of $f_{K}$

$$
\tilde{f}_{K}(\boldsymbol{y})=\boldsymbol{b}^{T} \boldsymbol{y}+\max _{\boldsymbol{X} \in \operatorname{conv}\left(K \boldsymbol{u} \boldsymbol{u}^{T}, \boldsymbol{v}_{i} \boldsymbol{v}_{i}^{T}\right)}\left\langle\boldsymbol{X}, \boldsymbol{C}-\sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i}\right\rangle .
$$

Since applying our Cutting Plane Method to either $f_{K}$ or $\tilde{f}_{K}$ gives the same result, the correctness of the our Cutting Plane Method shows that

$$
\tilde{f}_{K}(\boldsymbol{z}) \leq \min _{\|\boldsymbol{y}\|_{\infty} \leq L} \tilde{f}_{K}(\boldsymbol{y})+\varepsilon
$$

Note that the function $\tilde{f}_{K}$ is defined such that $\tilde{f}_{K}(\boldsymbol{z})=f_{K}(\boldsymbol{z})$. Hence, we have

$$
\min _{\|\boldsymbol{y}\|_{\infty} \leq L} f_{K}(\boldsymbol{y}) \leq f_{K}(\boldsymbol{z}) \leq \tilde{f}_{K}(\boldsymbol{z}) \leq \min _{\|\boldsymbol{y}\|_{\infty} \leq L} \tilde{f}_{K}(\boldsymbol{y})+\varepsilon .
$$

Also, note that $\tilde{f}_{K}(\boldsymbol{x}) \leq f_{K}(\boldsymbol{x})$ for all $\boldsymbol{x}$. Hence, we have

$$
\min _{\|\boldsymbol{y}\|_{\infty} \leq L} f_{K}(\boldsymbol{y})-\varepsilon \leq \min _{\|\boldsymbol{y}\|_{\infty} \leq L} \tilde{f}(\boldsymbol{y}) \leq \min _{\|\boldsymbol{y}\|_{\infty} \leq L} f_{K}(\boldsymbol{y}) .
$$

Now, we consider the primal version of $\tilde{f}$, namely

$$
g(\boldsymbol{X}) \stackrel{\text { def }}{=} \min _{\|\boldsymbol{y}\|_{\infty} \leq L} \boldsymbol{b}^{T} \boldsymbol{y}+\left\langle\boldsymbol{X}, \boldsymbol{C}-\sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i}\right\rangle .
$$

Sion's minimax theorem [239] shows that

$$
\mathrm{OPT} \geq \max _{\boldsymbol{x} \in \operatorname{conv}\left(K u u^{T}, \boldsymbol{v}_{i} \boldsymbol{v}_{i}^{T}\right)} g(\boldsymbol{X})=\min _{\|\boldsymbol{y}\|_{\infty} \leq L} \tilde{f}(\boldsymbol{y}) \geq \mathrm{OPT}-\varepsilon .
$$

Therefore, to get the primal solution, we only need to find $\boldsymbol{u}$ by Lemma 9.3.3 and solve the maximization problem on $g$. Note that

$$
\begin{aligned}
g(\boldsymbol{X}) & =\min _{\|\boldsymbol{y}\|_{\infty} \leq L} \sum_{i=1}^{n} y_{i}\left(b_{i}-\left\langle\boldsymbol{X}, \boldsymbol{A}_{i}\right\rangle\right)+\langle\boldsymbol{X}, \boldsymbol{C}\rangle \\
& =-L \sum_{i}\left|b_{i}-\left\langle\boldsymbol{X}, \boldsymbol{A}_{i}\right\rangle\right|+\langle\boldsymbol{X}, \boldsymbol{C}\rangle
\end{aligned}
$$

For notation simplicity, we write $K \boldsymbol{u} \boldsymbol{u}^{T}=\boldsymbol{v}_{0} \boldsymbol{v}_{0}^{T}$. Then, $\boldsymbol{X}=\sum_{j=0}^{O(n)} \alpha_{j} \boldsymbol{v}_{j} \boldsymbol{v}_{j}^{T}$ for some $\sum \alpha_{j}=1$ and
$\alpha_{j} \geq 0$. Substituting this into the function $g$, we have

$$
g(\boldsymbol{\alpha})=-L \sum_{j}\left|b_{i}-\sum_{j} \alpha_{j} \boldsymbol{v}_{j}^{T} \boldsymbol{A}_{i} \boldsymbol{v}_{j}\right|+\sum_{j} \alpha_{j} \boldsymbol{v}_{j}^{T} \boldsymbol{C} \boldsymbol{v}_{j}
$$

Hence, this can be easily written as a linear program with $O(n)$ variables and $O(n)$ constraints in time $O(n S)$. Now, we can apply interior point method to find $\boldsymbol{\alpha}$ such that

$$
g(\boldsymbol{\alpha}) \geq \max _{\boldsymbol{X} \in \operatorname{conv}\left(K \boldsymbol{u} \boldsymbol{u}^{T}, \boldsymbol{v}_{i} \boldsymbol{v}_{i}^{T}\right)} g(\boldsymbol{X})-\varepsilon \geq \mathrm{OPT}-2 \varepsilon
$$

Let the corresponding approximate solution be $\widetilde{\boldsymbol{X}}=\sum \alpha_{j} \boldsymbol{v}_{j} \boldsymbol{v}_{j}^{T}$. Then, we have

$$
\langle\widetilde{\boldsymbol{X}}, \boldsymbol{C}\rangle-L \sum_{i}\left|b_{i}-\left\langle\boldsymbol{X}, \boldsymbol{A}_{i}\right\rangle\right| \geq \mathrm{OPT}-2 \varepsilon
$$

Now, we let $\tilde{b}_{i}=\left\langle\widetilde{\boldsymbol{X}}, \boldsymbol{A}_{i}\right\rangle$. Then, we note that

$$
\begin{aligned}
\langle\widetilde{\boldsymbol{X}}, \boldsymbol{C}\rangle & \leq \max _{\boldsymbol{X} \succeq \mathbf{0} \boldsymbol{A}_{i} \bullet \boldsymbol{X}=\tilde{b}_{i}} \boldsymbol{C} \bullet \boldsymbol{X} \\
& =\min _{\sum_{i=1}^{n} y_{i} \boldsymbol{A}_{i} \succeq \boldsymbol{C}} \tilde{b}_{i}^{T} \boldsymbol{y} \\
& \leq \mathrm{OPT}+M \sum_{i}\left|b_{i}-\left\langle\widetilde{\boldsymbol{X}}, \boldsymbol{A}_{i}\right\rangle\right|
\end{aligned}
$$

because $\|\boldsymbol{y}\|_{\infty} \leq M$. Hence, we have

$$
\mathrm{OPT}+(M-L) \sum_{i}\left|b_{i}-\left\langle\widetilde{\boldsymbol{X}}, \boldsymbol{A}_{i}\right\rangle\right| \geq\langle\widetilde{\boldsymbol{X}}, \boldsymbol{C}\rangle-L \sum_{i}\left|b_{i}-\left\langle\widetilde{\boldsymbol{X}}, \boldsymbol{A}_{i}\right\rangle\right| \geq \mathrm{OPT}-2 \varepsilon
$$

Now, we put $L=M+2$, we have

$$
\sum_{i}\left|b_{i}-\left\langle\widetilde{\boldsymbol{X}}, \boldsymbol{A}_{i}\right\rangle\right| \leq \varepsilon
$$

This gives the result for the primal. Note that it only takes $O\left(n^{5 / 2} \log O(1)(n M / \varepsilon)\right)$ to solve a linear program with $O(n)$ variables and $O(n)$ constraints because we have an explicit interior point deep inside the feasible set, i.e. $\alpha_{i}=\frac{1}{m}$ for some parameter $m$ [164]. ${ }^{1}$ Hence, the running time is dominated by the cost of cutting plane method which is $O\left(\left(n S+n^{3}+n m^{\omega+o(1)}\right) \log ^{O(1)}\left(\frac{n M}{\varepsilon}\right)\right)$ by putting $L=M+2$.

We leave it as an open problem if it is possible to improve this result by reusing the computation in the separation oracle and achieve a running time of $O\left(\left(n S+n^{3}+n m^{2}\right) \log ^{O(1)}\left(\frac{n M}{\varepsilon}\right)\right)$.

## ■ 9.4 Intersection of Convex Sets

In this section we introduce a general technique to optimize a linear function over the intersection of two convex sets, whenever the linear optimization problem on each of them can be done efficiently. At the very high level, this is accomplished by applying cutting plane to a suitably regularized version of the problem. In Section 9.4 .1 we present the technique and in the remaining sections we provide

[^18]several applications including, matroid intersection (Section 9.4.2), submodular flow (Section 9.4.3), and minimizing over the intersection of an affine subspace and a convex set (Section 9.4.4).

## ■ 9.4.1 The Technique

Throughout this section we consider variants of the following general optimization problem

$$
\begin{equation*}
\max _{\boldsymbol{x} \in K_{1} \cap K_{2}}\langle\boldsymbol{c}, \boldsymbol{x}\rangle \tag{9.9}
\end{equation*}
$$

where $\boldsymbol{x}, \boldsymbol{c} \in \mathbb{R}^{n}, K_{1}$ and $K_{2}$ are convex subsets of $\mathbb{R}^{n}$. We assume that

$$
\begin{equation*}
\max _{\boldsymbol{x} \in K_{1}}\|\boldsymbol{x}\|_{2}<M, \max _{\boldsymbol{x} \in K_{2}}\|\boldsymbol{x}\|_{2}<M,\|\boldsymbol{c}\|_{2} \leq M \tag{9.10}
\end{equation*}
$$

for some constant $M \geq 1$ and we assume that

$$
\begin{equation*}
K_{1} \cap K_{2} \neq \emptyset . \tag{9.11}
\end{equation*}
$$

Instead of a separation oracle, we assume that $K_{1}$ and $K_{2}$ each have optimization oracles (see Section 9.2).

To solve this problem we first introduce a relaxation for the problem (9.9) that we can optimize efficiently. Because we have only the optimization oracles for $K_{1}$ and $K_{2}$, we simply have variables $\boldsymbol{x}$ and $\boldsymbol{y}$ for each of them in the objective. Since the output should (approximately) be in the intersection of $K_{1}$ and $K_{2}$, a regularization term $-\frac{\lambda}{2}\|\boldsymbol{x}-\boldsymbol{y}\|_{2}^{2}$ is added to force $\boldsymbol{x} \approx \boldsymbol{y}$ where $\lambda$ is a large number to be determined later. Furthermore, we add terms to make the problem strong concave.

Lemma 9.4.1. Assume (9.10) and (9.11). For $\lambda \geq 1$, let

$$
\begin{equation*}
f_{\lambda}(\boldsymbol{x}, \boldsymbol{y}) \stackrel{\text { def }}{=} \frac{1}{2}\langle\boldsymbol{c}, \boldsymbol{x}\rangle+\frac{1}{2}\langle\boldsymbol{c}, \boldsymbol{y}\rangle-\frac{\lambda}{2}\|\boldsymbol{x}-\boldsymbol{y}\|_{2}^{2}-\frac{1}{2 \lambda}\|\boldsymbol{x}\|_{2}^{2}-\frac{1}{2 \lambda}\|\boldsymbol{y}\|_{2}^{2} . \tag{9.12}
\end{equation*}
$$

There is an unique maximizer $\left(\boldsymbol{x}_{\lambda}, \boldsymbol{y}_{\lambda}\right)$ for the problem $\max _{\boldsymbol{x} \in K_{1}, \boldsymbol{y} \in K_{2}} f_{\lambda}(\boldsymbol{x}, \boldsymbol{y})$. The maximizer $\left(\boldsymbol{x}_{\lambda}, \boldsymbol{y}_{\lambda}\right)$ is a good approximation of the solution of (9.9), i.e. $\left\|\boldsymbol{x}_{\lambda}-\boldsymbol{y}_{\lambda}\right\|_{2}^{2} \leq \frac{6 M^{2}}{\lambda}$ and

$$
\begin{equation*}
\max _{\boldsymbol{x} \in K_{1} \cap K_{2}}\langle\boldsymbol{c}, \boldsymbol{x}\rangle \leq f_{\lambda}\left(\boldsymbol{x}_{\lambda}, \boldsymbol{y}_{\lambda}\right)+\frac{M^{2}}{\lambda} \tag{9.13}
\end{equation*}
$$

Proof. Let $\boldsymbol{x}^{*}$ be a maximizer of $\max _{\boldsymbol{x} \in K_{1} \cap K_{2}}\langle\boldsymbol{c}, \boldsymbol{x}\rangle$. By assumption (9.10), $\left\|\boldsymbol{x}^{*}\right\|_{2} \leq M$, and therefore

$$
\begin{equation*}
f_{\lambda}\left(\boldsymbol{x}^{*}, \boldsymbol{x}^{*}\right)=\left\langle\boldsymbol{c}, \boldsymbol{x}^{*}\right\rangle-\frac{\left\|\boldsymbol{x}^{*}\right\|_{2}^{2}}{\lambda} \geq \max _{\boldsymbol{x} \in K_{1} \cap K_{2}}\langle\boldsymbol{c}, \boldsymbol{x}\rangle-\frac{M^{2}}{\lambda} . \tag{9.14}
\end{equation*}
$$

This shows (9.13). Since $f_{\lambda}$ is strongly concave in $\boldsymbol{x}$ and $\boldsymbol{y}$, there is a unique maximizer ( $\boldsymbol{x}_{\lambda}, \boldsymbol{y}_{\lambda}$ ). Let $\mathrm{OPT}_{\lambda}=f_{\lambda}\left(\boldsymbol{x}_{\lambda}, \boldsymbol{y}_{\lambda}\right)$. Then, we have

$$
\begin{aligned}
\mathrm{OPT}_{\lambda} & \leq \frac{1}{2}\|\boldsymbol{c}\|_{2}\left\|\boldsymbol{x}_{\lambda}\right\|_{2}+\frac{1}{2}\|\boldsymbol{c}\|_{2}\left\|\boldsymbol{y}_{\lambda}\right\|_{2}-\frac{\lambda}{2}\left\|\boldsymbol{x}_{\lambda}-\boldsymbol{y}_{\lambda}\right\|_{2}^{2} \\
& \leq \frac{M^{2}}{2}+\frac{M^{2}}{2}-\frac{\lambda}{2}\left\|\boldsymbol{x}_{\lambda}-\boldsymbol{y}_{\lambda}\right\|_{2}^{2} .
\end{aligned}
$$

On the other hand, using $\lambda \geq 1$, (9.14) shows that

$$
\mathrm{OPT}_{\lambda} \geq f_{\lambda}\left(\boldsymbol{x}^{*}, \boldsymbol{x}^{*}\right) \geq \max _{x \in K_{1} \cap K_{2}}\langle\boldsymbol{c}, \boldsymbol{x}\rangle-\frac{M^{2}}{\lambda} \geq-2 M^{2}
$$

Hence, we have

$$
\begin{equation*}
\left\|\boldsymbol{x}_{\lambda}-\boldsymbol{y}_{\lambda}\right\|_{2}^{2} \leq \frac{2\left(M^{2}-\mathrm{OPT}_{\lambda}\right)}{\lambda} \leq \frac{6 M^{2}}{\lambda} \tag{9.15}
\end{equation*}
$$

Now we write $\max f_{\lambda}(\boldsymbol{x}, \boldsymbol{y})$ as a max-min problem. The reason for doing this is that the dual approximate solution is much easier to obtain and there is a way to read off a primal approximate solution from a dual approximate solution. This is analogous to the idea in [161] which showed how to convert a cut solution to a flow solution by adding regularization terms into the problem.

Lemma 9.4.2. Assume (9.10) and (9.11). Let $\lambda \geq 2$. For any $\boldsymbol{x} \in K_{1}$ and $\boldsymbol{y} \in K_{2}$, the function $f_{\lambda}$ can be represented as

$$
\begin{equation*}
f_{\lambda}(\boldsymbol{x}, \boldsymbol{y})=\min _{\left(\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right) \in \Omega} g_{\lambda}\left(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right) \tag{9.16}
\end{equation*}
$$

where $\Omega=\left\{\left(\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right):\left\|\boldsymbol{\theta}_{1}\right\|_{2} \leq 2 M,\left\|\boldsymbol{\theta}_{2}\right\|_{2} \leq M,\left\|\boldsymbol{\theta}_{3}\right\|_{2} \leq M\right\}$ and

$$
\begin{equation*}
g_{\lambda}\left(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right)=\left\langle\frac{\boldsymbol{c}}{2}+\lambda \boldsymbol{\theta}_{1}+\frac{\boldsymbol{\theta}_{2}}{\lambda}, \boldsymbol{x}\right\rangle+\left\langle\frac{\boldsymbol{c}}{2}-\lambda \boldsymbol{\theta}_{1}+\frac{\boldsymbol{\theta}_{3}}{\lambda}, \boldsymbol{y}\right\rangle+\frac{\lambda}{2}\left\|\boldsymbol{\theta}_{1}\right\|_{2}^{2}+\frac{1}{2 \lambda}\left\|\boldsymbol{\theta}_{2}\right\|_{2}^{2}+\frac{1}{2 \lambda}\left\|\boldsymbol{\theta}_{3}\right\|_{2}^{2} \tag{9.17}
\end{equation*}
$$

$\operatorname{Let}_{\lambda}\left(\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right)=\max _{\boldsymbol{x} \in K_{1}, \boldsymbol{y} \in K_{2}} g_{\lambda}\left(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right)$. For any $\left(\boldsymbol{\theta}_{1}^{\prime}, \boldsymbol{\theta}_{2}^{\prime}, \boldsymbol{\theta}_{3}^{\prime}\right)$ such that $h_{\lambda}\left(\boldsymbol{\theta}_{1}^{\prime}, \boldsymbol{\theta}_{2}^{\prime}, \boldsymbol{\theta}_{3}^{\prime}\right) \leq$ $\min _{\left(\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right) \in \Omega} h_{\lambda}\left(\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right)+\varepsilon$, we know $\boldsymbol{z}=\frac{1}{2}\left(\boldsymbol{\theta}_{2}^{\prime}+\boldsymbol{\theta}_{3}^{\prime}\right)$ satisfies

$$
\max _{\boldsymbol{x} \in K_{1} \cap K_{2}}\langle\boldsymbol{c}, \boldsymbol{x}\rangle \leq\langle\boldsymbol{c}, \boldsymbol{z}\rangle+\frac{20 M^{2}}{\lambda}+20 \lambda^{3} \varepsilon
$$

and $\left\|\boldsymbol{z}-\boldsymbol{x}_{\lambda}\right\|_{2}+\left\|\boldsymbol{z}-\boldsymbol{y}_{\lambda}\right\|_{2} \leq 4 \sqrt{2 \lambda \varepsilon}+\sqrt{\frac{6 M^{2}}{\lambda}}$ where $\left(\boldsymbol{x}_{\lambda}, \boldsymbol{y}_{\lambda}\right)$ is the unique maximizer for the problem $\max _{\boldsymbol{x} \in K_{1}, \boldsymbol{y} \in K_{2}} f_{\lambda}(\boldsymbol{x}, \boldsymbol{y})$.

Proof. Note that for any $\|\boldsymbol{\zeta}\|_{2} \leq \alpha$, we have

$$
-\frac{1}{2}\|\boldsymbol{\zeta}\|_{2}^{2}=\min _{\|\boldsymbol{\theta}\|_{2} \leq \alpha}\langle\boldsymbol{\theta}, \boldsymbol{\zeta}\rangle+\frac{1}{2}\|\boldsymbol{\theta}\|_{2}^{2}
$$

Using this and (9.10), we have (9.16) for all $\boldsymbol{x} \in K_{1}$ and $\boldsymbol{y} \in K_{2}$ as desired. Since $\Omega$ is closed and bounded set and the function $g_{\lambda}$ is concave in $(\boldsymbol{x}, \boldsymbol{y})$ and convex in $\left(\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right)$, Sion's minimax theorem [239] shows that

$$
\begin{equation*}
\max _{\boldsymbol{x} \in K_{1}, \boldsymbol{y} \in K_{2}} f_{\lambda}(\boldsymbol{x}, \boldsymbol{y})=\min _{\left(\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right) \in \Omega} h_{\lambda}\left(\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right) \tag{9.18}
\end{equation*}
$$

Since $f_{\lambda}$ is strongly concave, there is an unique maximizer $\left(\boldsymbol{x}_{\lambda}, \boldsymbol{y}_{\lambda}\right)$ of $f_{\lambda}$. Since $h_{\lambda}$ is strongly convex, there is a unique minimizer $\left(\boldsymbol{\theta}_{1}^{*}, \boldsymbol{\theta}_{2}^{*}, \boldsymbol{\theta}_{3}^{*}\right)$. By the definition of $f_{\lambda}$ and $h_{\lambda}$, we have

$$
h_{\lambda}\left(\boldsymbol{\theta}_{1}^{*}, \boldsymbol{\theta}_{2}^{*}, \boldsymbol{\theta}_{3}^{*}\right) \geq g_{\lambda}\left(\boldsymbol{x}_{\lambda}, \boldsymbol{y}_{\lambda}, \boldsymbol{\theta}_{1}^{*}, \boldsymbol{\theta}_{2}^{*}, \boldsymbol{\theta}_{3}^{*}\right) \geq f_{\lambda}\left(\boldsymbol{x}_{\lambda}, \boldsymbol{y}_{\lambda}\right)
$$

Using (9.18), the equality above holds and hence $\left(\boldsymbol{\theta}_{1}^{*}, \boldsymbol{\theta}_{2}^{*}, \boldsymbol{\theta}_{3}^{*}\right)$ is the minimizer of $g_{\lambda}\left(\boldsymbol{x}_{\lambda}, \boldsymbol{y}_{\lambda}, \boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right)$ over $\left(\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right)$. Since the domain $\Omega$ is large enough that $\left(\boldsymbol{\theta}_{1}^{*}, \boldsymbol{\theta}_{2}^{*}, \boldsymbol{\theta}_{3}^{*}\right)$ is an interior point in $\Omega$, the optimality condition of $g_{\lambda}$ shows that we have $\boldsymbol{\theta}_{2}^{*}=\boldsymbol{x}_{\lambda}$ and $\boldsymbol{\theta}_{3}^{*}=\boldsymbol{y}_{\lambda}$.

Since $h_{\lambda}$ is $\frac{1}{\lambda}$ strongly convex, we have $\left\|\boldsymbol{\theta}_{1}^{\prime}-\boldsymbol{\theta}_{1}^{*}\right\|_{2}^{2}+\left\|\boldsymbol{\theta}_{2}^{\prime}-\boldsymbol{\theta}_{2}^{*}\right\|_{2}^{2}+\left\|\boldsymbol{\theta}_{3}^{\prime}-\boldsymbol{\theta}_{3}^{*}\right\|_{2}^{2} \leq 2 \lambda \varepsilon$ (Fact 9.2.3). Since $\boldsymbol{\theta}_{2}^{*}=\boldsymbol{x}_{\lambda}$ and $\boldsymbol{\theta}_{3}^{*}=\boldsymbol{y}_{\lambda}$, we have

$$
\begin{equation*}
\left\|\boldsymbol{\theta}_{2}^{\prime}-\boldsymbol{x}_{\lambda}\right\|_{2}^{2}+\left\|\boldsymbol{\theta}_{3}^{\prime}-\boldsymbol{y}_{\lambda}\right\|_{2}^{2} \leq 2 \lambda \varepsilon \tag{9.19}
\end{equation*}
$$

Therefore, we have $\left\|\boldsymbol{x}_{\lambda}-\boldsymbol{y}_{\lambda}\right\|_{2} \geq\left\|\boldsymbol{\theta}_{2}^{\prime}-\boldsymbol{\theta}_{3}^{\prime}\right\|_{2}-2 \sqrt{2 \lambda \varepsilon},\left\|\boldsymbol{x}_{\lambda}\right\|_{2} \geq\left\|\boldsymbol{\theta}_{2}^{\prime}\right\|_{2}-\sqrt{2 \lambda \varepsilon}$ and $\left\|\boldsymbol{y}_{\lambda}\right\|_{2} \geq\left\|\boldsymbol{\theta}_{3}^{\prime}\right\|_{2}-\sqrt{2 \lambda \varepsilon}$. Using these, $\left\|\boldsymbol{x}_{\lambda}\right\|_{2} \leq M$ and $\left\|\boldsymbol{y}_{\lambda}\right\|_{2} \leq M$, we have

$$
\begin{aligned}
f_{\lambda}\left(\boldsymbol{\theta}_{2}^{\prime}, \boldsymbol{\theta}_{3}^{\prime}\right)= & \frac{1}{2}\left\langle\boldsymbol{c}, \boldsymbol{\theta}_{2}^{\prime}\right\rangle+\frac{1}{2}\left\langle\boldsymbol{c}, \boldsymbol{\theta}_{3}^{\prime}\right\rangle-\frac{\lambda}{2}\left\|\boldsymbol{\theta}_{2}^{\prime}-\boldsymbol{\theta}_{3}^{\prime}\right\|_{2}^{2}-\frac{1}{2 \lambda}\left\|\boldsymbol{\theta}_{2}^{\prime}\right\|_{2}^{2}-\frac{1}{2 \lambda}\left\|\boldsymbol{\theta}_{3}^{\prime}\right\|_{2}^{2} \\
\geq & \frac{1}{2}\left\langle\boldsymbol{c}, \boldsymbol{x}_{\lambda}\right\rangle+\frac{1}{2}\left\langle\boldsymbol{c}, \boldsymbol{y}_{\lambda}\right\rangle-M \sqrt{2 \lambda \varepsilon} \\
& -\frac{\lambda}{2}\left(\left\|\boldsymbol{x}_{\lambda}-\boldsymbol{y}_{\lambda}\right\|_{2}+2 \sqrt{2 \lambda \varepsilon}\right)^{2} \\
& -\frac{1}{2 \lambda}\left(\left\|\boldsymbol{x}_{\lambda}\right\|_{2}+\sqrt{2 \lambda \varepsilon}\right)^{2}-\frac{1}{2 \lambda}\left(\left\|\boldsymbol{y}_{\lambda}\right\|_{2}+\sqrt{2 \lambda \varepsilon}\right)^{2} \\
= & \frac{1}{2}\left\langle\boldsymbol{c}, \boldsymbol{x}_{\lambda}\right\rangle+\frac{1}{2}\left\langle\boldsymbol{c}, \boldsymbol{y}_{\lambda}\right\rangle-\frac{\lambda}{2}\left\|\boldsymbol{x}_{\lambda}-\boldsymbol{y}_{\lambda}\right\|_{2}^{2}-\frac{1}{2 \lambda}\left\|\boldsymbol{x}_{\lambda}\right\|_{2}^{2}-\frac{1}{2 \lambda}\left\|\boldsymbol{y}_{\lambda}\right\|_{2}^{2} \\
& -M \sqrt{2 \lambda \varepsilon}-2 \lambda \sqrt{2 \lambda \varepsilon}\left\|\boldsymbol{x}_{\lambda}-\boldsymbol{y}_{\lambda}\right\|_{2}-4 \lambda^{2} \varepsilon \\
& -\frac{1}{\lambda}\left\|\boldsymbol{x}_{\lambda}\right\|_{2} \sqrt{2 \lambda \varepsilon}-\varepsilon-\frac{1}{\lambda}\left\|\boldsymbol{y}_{\lambda}\right\|_{2} \sqrt{2 \lambda \varepsilon}-\varepsilon .
\end{aligned}
$$

Using $\left\|\boldsymbol{x}_{\lambda}-\boldsymbol{y}_{\lambda}\right\|_{2} \leq \sqrt{\frac{6 M^{2}}{\lambda}}$ (Lemma 9.4.1), $\left\|\boldsymbol{x}_{\lambda}\right\|_{2}<M$ and $\left\|\boldsymbol{y}_{\lambda}\right\|_{2}<M$, we have

$$
\begin{aligned}
f_{\lambda}\left(\boldsymbol{\theta}_{2}^{\prime}, \boldsymbol{\theta}_{3}^{\prime}\right) \geq & f_{\lambda}\left(\boldsymbol{x}_{\lambda}, \boldsymbol{y}_{\lambda}\right) \\
& -M \sqrt{2 \lambda \varepsilon}-2 \lambda \sqrt{2 \lambda \varepsilon}\left\|\boldsymbol{x}_{\lambda}-\boldsymbol{y}_{\lambda}\right\|_{2}-4 \lambda^{2} \varepsilon \\
& -\frac{1}{\lambda}\left\|\boldsymbol{x}_{\lambda}\right\|_{2} \sqrt{2 \lambda \varepsilon}-\varepsilon-\frac{1}{\lambda}\left\|\boldsymbol{y}_{\lambda}\right\|_{2} \sqrt{2 \lambda \varepsilon}-\varepsilon . \\
\geq & f_{\lambda}\left(\boldsymbol{x}_{\lambda}, \boldsymbol{y}_{\lambda}\right) \\
& -M \sqrt{2 \lambda \varepsilon}-2 \lambda \sqrt{12 \varepsilon} M-4 \lambda^{2} \varepsilon \\
& -2 M \sqrt{2 \frac{\varepsilon}{\lambda}}-2 \varepsilon .
\end{aligned}
$$

Since $\lambda \geq 2$, we have

$$
f_{\lambda}\left(\boldsymbol{\theta}_{2}^{\prime}, \boldsymbol{\theta}_{3}^{\prime}\right) \geq f_{\lambda}\left(\boldsymbol{x}_{\lambda}, \boldsymbol{y}_{\lambda}\right)-20 M \lambda \sqrt{\varepsilon}-10 \lambda^{2} \varepsilon .
$$

Let $\boldsymbol{z}=\frac{\boldsymbol{\theta}_{2}^{\prime}+\boldsymbol{\theta}_{3}^{\prime}}{2}$. Lemma 9.4.1 shows that

$$
\begin{aligned}
\max _{\boldsymbol{x} \in K_{1} \cap K_{2}}\langle\boldsymbol{c}, \boldsymbol{x}\rangle & \leq \max _{\boldsymbol{x} \in K_{1}, \boldsymbol{y} \in K_{2}} f_{\lambda}(\boldsymbol{x}, \boldsymbol{y})+\frac{M^{2}}{\lambda} \\
& \leq f_{\lambda}\left(\boldsymbol{\theta}_{2}^{\prime}, \boldsymbol{\theta}_{3}^{\prime}\right)+\frac{M^{2}}{\lambda}+20 M \lambda \sqrt{\varepsilon}+10 \lambda^{2} \varepsilon \\
& \leq\langle\boldsymbol{c}, \boldsymbol{z}\rangle+\frac{20 M^{2}}{\lambda}+20 \lambda^{3} \varepsilon
\end{aligned}
$$

because $20 M \lambda \sqrt{\varepsilon} \leq 10 \frac{M^{2}}{\lambda}+10 \lambda^{3} \varepsilon$. Furthermore, we have

$$
\begin{aligned}
\left\|\boldsymbol{z}-\boldsymbol{x}_{\lambda}\right\|_{2}+\left\|\boldsymbol{z}-\boldsymbol{y}_{\lambda}\right\|_{2} & \leq\left\|\boldsymbol{\theta}_{2}^{\prime}-\boldsymbol{x}_{\lambda}\right\|_{2}+\| \boldsymbol{\theta}_{3}^{\prime} \\
& \leq 4 \sqrt{2 \lambda \varepsilon}+\sqrt{\frac{6 M^{2}}{\lambda}}\left\|_{2}+\right\| \boldsymbol{\theta}_{2}^{\prime}-\boldsymbol{\theta}_{3}^{\prime} \|_{2}
\end{aligned}
$$

We now apply our cutting plane method to solve the optimization problem (9.9). First we show how to transform the optimization oracles for $K_{1}$ and $K_{2}$ to get a separation oracle for $h_{\lambda}$, with the
appropriate parameters.
Lemma 9.4.3. Suppose we have a ع-optimization oracle for $K_{1}$ and $K_{2}$ for some $0<\varepsilon<1$. Then on the set $\left\{\|\boldsymbol{\theta}\|_{2} \leq D\right\}$, we have a $(O(\sqrt{\varepsilon \lambda D}), O(\sqrt{\varepsilon \lambda D}))$-separation oracle for $h_{\lambda}$ with time complexity $\mathrm{OO}_{\varepsilon}\left(K_{1}\right)+\mathrm{OO}_{\varepsilon}\left(K_{2}\right)$.

Proof. Recall that the function $h_{\lambda}$ is defined by

$$
\begin{aligned}
& h_{\lambda}\left(\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right) \\
= & \max _{\boldsymbol{x} \in K_{1}, \boldsymbol{y} \in K_{2}}\left(\left\langle\frac{\boldsymbol{c}}{2}+\lambda \boldsymbol{\theta}_{1}+\frac{\boldsymbol{\theta}_{2}}{\lambda}, \boldsymbol{x}\right\rangle+\left\langle\frac{\boldsymbol{c}}{2}-\lambda \boldsymbol{\theta}_{1}+\frac{\boldsymbol{\theta}_{3}}{\lambda}, \boldsymbol{y}\right\rangle+\frac{\lambda}{2}\left\|\boldsymbol{\theta}_{1}\right\|_{2}^{2}+\frac{1}{2 \lambda}\left\|\boldsymbol{\theta}_{2}\right\|_{2}^{2}+\frac{1}{2 \lambda}\left\|\boldsymbol{\theta}_{3}\right\|_{2}^{2}\right) \\
= & \max _{\boldsymbol{x} \in K_{1}}\left\langle\frac{\boldsymbol{c}}{2}+\lambda \boldsymbol{\theta}_{1}+\frac{\boldsymbol{\theta}_{2}}{\lambda}, \boldsymbol{x}\right\rangle+\max _{\boldsymbol{y} \in K_{2}}\left\langle\frac{\boldsymbol{c}}{2}-\lambda \boldsymbol{\theta}_{1}+\frac{\boldsymbol{\theta}_{3}}{\lambda}, \boldsymbol{y}\right\rangle+\frac{\lambda}{2}\left\|\boldsymbol{\theta}_{1}\right\|_{2}^{2}+\frac{1}{2 \lambda}\left\|\boldsymbol{\theta}_{2}\right\|_{2}^{2}+\frac{1}{2 \lambda}\left\|\boldsymbol{\theta}_{3}\right\|_{2}^{2} .
\end{aligned}
$$

Lemma 9.2 .6 shows how to compute the subgradient of functions of the form $f(\boldsymbol{c})=\max _{\boldsymbol{x} \in K}\langle\boldsymbol{c}, \boldsymbol{x}\rangle$ using the optimization oracle for $K$. The rest of the term are differentiable so its subgradient is just the gradient. Hence, by addition rule for subgradients (Fact 9.2.3), we have a $O(\varepsilon \lambda)$-subgradient oracle for $f_{\lambda}$ using a $O(\varepsilon)$-optimization oracle for $K_{1}$ and $K_{2}$. The result then follows from Lemma 9.2.4.

Theorem 9.4.4. Assume (9.10) and (9.11). Suppose that we have $\varepsilon$-optimization oracle for every $\varepsilon>0$. For $0<\delta<1$, we can find $\boldsymbol{z} \in \mathbb{R}^{n}$ such that

$$
\max _{\boldsymbol{x} \in K_{1} \cap K_{2}}\langle\boldsymbol{c}, \boldsymbol{x}\rangle \leq \delta+\langle\boldsymbol{c}, \boldsymbol{z}\rangle
$$

and $\|\boldsymbol{z}-\boldsymbol{x}\|_{2}+\|\boldsymbol{z}-\boldsymbol{y}\|_{2} \leq \delta$ for some $\boldsymbol{x} \in K_{1}$ and $\boldsymbol{y} \in K_{2}$ in time

$$
O\left(n\left(\mathrm{OO}_{\eta}\left(K_{1}\right)+\mathrm{OO}_{\eta}\left(K_{2}\right)\right) \log \left(\frac{n M}{\delta}\right)+n^{3} \log ^{O(1)}\left(\frac{n M}{\delta}\right)\right)
$$

where $\eta=\Omega\left(\left(\frac{\delta}{n M}\right)^{O(1)}\right)$.
Proof. Setting $\lambda=\frac{40 M^{2}}{\delta^{2}}$ and $\varepsilon=\frac{\delta^{7}}{10^{7} M^{6}}$ in Lemma 9.4 .2 we see that so long as we obtain any approximate solution $\left(\boldsymbol{\theta}_{1}^{\prime}, \boldsymbol{\theta}_{2}^{\prime}, \boldsymbol{\theta}_{3}^{\prime}\right)$ such that

$$
h_{\lambda}\left(\boldsymbol{\theta}_{1}^{\prime}, \boldsymbol{\theta}_{2}^{\prime}, \boldsymbol{\theta}_{3}^{\prime}\right) \leq \min _{\left(\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right) \in \Omega} h_{\lambda}\left(\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right)+\varepsilon
$$

then we obtain the point we want. To apply Theorem 9.3.2, we use

$$
\tilde{h}\left(\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right)= \begin{cases}h_{\lambda}\left(\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right) & \text { if }\left(\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right) \in \Omega \\ +\infty & \text { else }\end{cases}
$$

Lemma 9.4 .3 shows that for any $\gamma>0$ we can obtain a $(\gamma, \gamma)$-separation oracle of $h_{\lambda}(\boldsymbol{\theta})$ by using sufficiently accurate optimization oracles. Since $\Omega$ is just a product of $\ell^{2}$ balls, we can produce a separating hyperplane easily when $\left(\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right) \notin \Omega$. Hence, we can obtain a $(\gamma, \gamma)$-separation oracle of $\tilde{h}(\boldsymbol{\theta})$. For simplicity, we use $\boldsymbol{\theta}$ to represent $\left(\boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \boldsymbol{\theta}_{3}\right)$. Note that $B_{\infty}(2 M) \supseteq \Omega$ and therefore we can apply Theorem 9.3 .2 with $R=2 M$ to compute $\boldsymbol{\theta}^{\prime}$ such

$$
\tilde{h}\left(\boldsymbol{\theta}^{\prime}\right)-\min _{\boldsymbol{\theta} \in \Omega} \tilde{h}(\boldsymbol{\theta}) \leq \gamma+\alpha\left(\max _{\boldsymbol{\theta} \in \Omega} \tilde{h}(\boldsymbol{\theta})-\min _{\boldsymbol{\theta} \in \Omega} \tilde{h}(\boldsymbol{\theta})\right)
$$

in time $O\left(n \mathrm{SO}_{\gamma, \gamma} \log \left(\frac{n \kappa}{\alpha}\right)+n^{3} \log ^{O(1)}\left(\frac{n \kappa}{\alpha}\right)\right)$ where $\gamma=\Omega\left(\alpha \operatorname{MinWidth}(\Omega) / n^{O(1)}\right)=\Omega\left(\alpha M / n^{O(1)}\right)$
and $\kappa=\frac{2 M}{\operatorname{MinWidth}(\Omega)}=O(1)$. Using $\lambda \geq 1$ and $M \geq 1$, we have

$$
\max _{\boldsymbol{\theta} \in \Omega} \tilde{h}(\boldsymbol{\theta})-\min _{\boldsymbol{\theta} \in \Omega} \tilde{h}(\boldsymbol{\theta}) \leq O\left(\lambda M^{2}\right) \leq O\left(\frac{M^{4}}{\delta^{2}}\right) .
$$

Setting $\alpha=\Theta\left(\frac{\delta^{9}}{M^{10}}\right)$ with some small enough constant, we have that we can find $\boldsymbol{\theta}^{\prime}$ such that

$$
\begin{aligned}
h_{\lambda}\left(\boldsymbol{\theta}^{\prime}\right) & \leq \min _{\boldsymbol{\theta} \in P} h_{\lambda}(\boldsymbol{\theta})+\gamma+\alpha O\left(\frac{M^{4}}{\delta^{2}}\right) \\
& =\min _{\boldsymbol{\theta} \in P} h_{\lambda}(\boldsymbol{\theta})+O\left(\frac{\delta^{7}}{M^{6}}\right) \\
& =\min _{\boldsymbol{\theta} \in P} h_{\lambda}(\boldsymbol{\theta})+\varepsilon
\end{aligned}
$$

in time $O\left(n \mathrm{SO}_{\gamma, \gamma} \log \left(\frac{n M}{\delta}\right)+n^{3} \log ^{O(1)}\left(\frac{n M}{\delta}\right)\right)$ where $\gamma=\Omega\left(\left(\frac{\delta}{n M}\right)^{O(1)}\right)$. Lemma 9.4.3 shows that the cost of $(\gamma, \gamma)$-separation oracle is just $O\left(\mathrm{OO}_{\eta}\left(K_{1}\right)+\mathrm{OO}_{\eta}\left(K_{2}\right)\right)$ where $\eta=\Omega\left(\left(\frac{\delta}{n M}\right)^{O(1)}\right)$.

Remark 9.4.5. Note that the algorithm does not promise that we obtain a point close to $K_{1} \cap K_{2}$. It only promises to give a point that is close to both some point in $K_{1}$ and some point in $K_{2}$. It appears to the authors that a further assumption is needed to get a point close to $K_{1} \cap K_{2}$. For example, if $K_{1}$ and $K_{2}$ are two almost parallel lines, it would be difficult to get an algorithm that does not depend on the angle. However, as far as we know, most algorithms tackling this problem are pseudo-polynomial and have polynomial dependence on the angle. Our algorithm depends on the logarithmic of the angle which is useful for combinatorial problems.

This reduction is very useful for problems in many areas including linear programming, semidefinite programming and algorithmic game theory. In the remainder of this section we demonstrate its power by applying it to classical combinatorial problems.

There is however one issue with applying our cutting plane algorithm to these problems. As with other convex optimization methods, only an approximately optimal solution is found. On the other hand, typically an exact solution is insisted in combinatorial optimization. To overcome this gap, we introduce the following lemma which (1) transforms the objective function so that there is only one optimal solution and (2) shows that an approximate solution is close to the optimal solution whenever it is unique. As we shall see in the next two subsections, this allows us to round an approximate solution to an optimal one.
Lemma 9.4.6. Given a linear program $\min _{\boldsymbol{A} x \geq b} \boldsymbol{c}^{T} \boldsymbol{x}$ where $\boldsymbol{x}, \boldsymbol{c} \in \mathbb{Z}^{n}, \boldsymbol{b} \in \mathbb{Z}^{m}$ and $\boldsymbol{A} \in \mathbb{Z}^{m \times n}$. Suppose $\{\boldsymbol{A} \boldsymbol{x} \geq \boldsymbol{b}\}$ is an integral polytope (i.e. all extreme points are integral) contained in the set $\left\{\|\boldsymbol{x}\|_{\infty} \leq M\right\}$. Then we can find a random cost vector $\boldsymbol{z} \in \mathbb{Z}^{n}$ with $\|\boldsymbol{z}\|_{\infty} \leq O\left(n^{2} M^{2}\|\boldsymbol{c}\|_{\infty}\right)$ such that with constant probability, $\min _{\boldsymbol{A} \boldsymbol{x} \geq \boldsymbol{b}} \boldsymbol{z}^{T} \boldsymbol{x}$ has an unique minimizer $\boldsymbol{x}^{*}$ and this minimizer is one of the minimizer(s) of $\min _{\boldsymbol{A x \geq b}} \boldsymbol{c}^{T} \boldsymbol{x}$. Furthermore, if there is an interior point $\boldsymbol{y}$ such that $\boldsymbol{z}^{T} \boldsymbol{y}<\min _{\boldsymbol{A} \boldsymbol{x} \geq \boldsymbol{b}} \boldsymbol{z}^{T} \boldsymbol{x}+\delta$, then $\left\|\boldsymbol{y}-\boldsymbol{x}^{*}\right\|_{\infty} \leq 2 n M \delta$.

Proof. The first part of the lemma follows by randomly perturbing the cost vector $\boldsymbol{c}$. We consider a new cost vector $\boldsymbol{z}=100 n^{2} M^{2} \boldsymbol{c}+\boldsymbol{r}$ where each coordinate of $\boldsymbol{r}$ is sampled randomly from $\{0,1, \cdots, 10 n M\}$. [145, Lem 4] shows that the linear program $\min _{\boldsymbol{A} \boldsymbol{x} \geq \boldsymbol{b}} \boldsymbol{z}^{T} \boldsymbol{x}$ has a unique minimizer with constant probability. Furthermore, it is clear that the minimizer of $\min _{\boldsymbol{A} \boldsymbol{x} \geq \boldsymbol{b}} \boldsymbol{z}^{T} \boldsymbol{x}$ is a minimizer of $\min _{\boldsymbol{A} \boldsymbol{x} \geq \boldsymbol{b}} \boldsymbol{c}^{T} \boldsymbol{x}$ (as $\boldsymbol{r}_{i} \ll 100 n^{2} M^{2}\left|\boldsymbol{c}_{i}\right|$ ).

Now we show the second part of the lemma. Given an interior point $\boldsymbol{y}$ of the polytope $\{\boldsymbol{A} \boldsymbol{x} \geq \boldsymbol{b}\}$, we can write $\boldsymbol{y}$ as a convex combination of the vertices of $\{\boldsymbol{A} \boldsymbol{x} \geq \boldsymbol{b}\}$, i.e. $\boldsymbol{y}=\sum t_{i} \boldsymbol{v}_{i}$. Note that
$\boldsymbol{z}^{T} \boldsymbol{y}=\sum t_{i} \boldsymbol{z}^{T} \boldsymbol{v}_{i}$. If all $\boldsymbol{v}_{i}$ are not the minimizer, then $\boldsymbol{z}^{T} \boldsymbol{v}_{i} \geq$ OPT +1 and hence $\boldsymbol{z}^{T} \boldsymbol{y} \geq$ OPT +1 which is impossible. Hence, we can assume that $\boldsymbol{v}_{1}$ is the minimizer. Hence, $\boldsymbol{z}^{T} \boldsymbol{v}_{i}=$ OPT if $i=1$ and $\boldsymbol{z}^{T} \boldsymbol{v}_{i} \geq \mathrm{OPT}+1$ otherwise. We then have $\boldsymbol{z}^{T} \boldsymbol{y} \geq \mathrm{OPT}+\left(1-t_{1}\right)$ which gives $1-t_{1}<\delta$. Finally, the claim follows from $\left\|\boldsymbol{y}-\boldsymbol{v}_{1}\right\|_{\infty} \leq \sum_{i \neq 1} t_{i}\left\|\boldsymbol{v}_{i}-\boldsymbol{v}_{1}\right\|_{\infty} \leq 2 n M \delta$.

## - 9.4.2 Matroid Intersection

Let $M_{1}=\left(E, \mathcal{I}_{1}\right)$ and $M_{2}=\left(E, \mathcal{I}_{2}\right)$ be two matroids sharing the same ground set. In this section we consider the weighted matroid intersection problem

$$
\min _{S \in \mathcal{I}_{1} \cap \mathcal{I}_{2}} \boldsymbol{w}(S) .
$$

where $\boldsymbol{w} \in \mathbb{R}^{E}$ and $w(S) \stackrel{\text { def }}{=} \sum_{e \in S} w_{e}$.
For any matroid $M=(E, \mathcal{I})$, it is well known that the polytope of all independent sets has the following description [80]:

$$
\begin{equation*}
\operatorname{conv}\left(\mathcal{I}_{1}\right)=\left\{\boldsymbol{x} \in \mathbb{R}^{E} \text { s.t. } 0 \leq x(S) \leq r(S) \text { for all } S \subseteq E\right\} \tag{9.20}
\end{equation*}
$$

where $r$ is the rank function for $M$, i.e. $r(S)$ is the size of the largest independent set that is a subset of $S$. Furthermore, the polytope of the matroid intersection satisfies $\operatorname{conv}\left(\mathcal{I}_{1} \cap \mathcal{I}_{2}\right)=\operatorname{conv}\left(\mathcal{I}_{1}\right) \cap \operatorname{conv}\left(\mathcal{I}_{2}\right)$.

It is well known that the optimization problem

$$
\min _{S \in \mathcal{I}_{1}} w(S) \text { and } \min _{S \in \mathcal{I}_{2}} w(S)
$$

can be solved efficiently by the greedy method. Given a matroid (polytope), the greedy method finds a maximum weight independent subset by maintaining a candidate independent subset $S$ and iteratively attempts to add new element to $S$ in descending weight. A element $i$ is added to $S$ if $S \cup\{i\}$ is still independent. A proof of this algorithm is well-known and can be found in any standard textbook on combinatorial optimization.

Clearly, the greedy method can be implemented by $O(n)$ calls to the independence oracle (also called membership oracle). For rank oracle, it requires $O(r \log n)$ calls by finding the next element to add via binary search. Therefore, we can apply Theorem 9.4.4 to get the following result (note that this algorithm is the fastest if $r$ is close to $n$ for the independence oracle).

Theorem 9.4.7. Suppose that the weights $\boldsymbol{w}$ are integer with $\|w\|_{\infty} \leq M$. Then, we can find

$$
S \in \underset{S \in \mathcal{I}_{1} \cap \mathcal{I}_{2}}{\arg \min } w(S)
$$

in time $O\left(n \mathrm{GO} \log (n M)+n^{3} \log ^{O(1)}(n M)\right)$ where GO is the cost of greedy method for $\mathcal{I}_{1}$ and $\mathcal{I}_{2}$.
Proof. Applying Lemma 9.4.6, we can find a new cost $\boldsymbol{z}$ such that

$$
\min _{x \in \operatorname{conv}\left(\mathcal{I}_{1}\right) \cap \operatorname{conv}\left(\mathcal{I}_{2}\right)} z^{T} \boldsymbol{x}
$$

has an unique solution. Note that for any $\boldsymbol{x} \in \operatorname{conv}\left(\mathcal{I}_{1}\right)$, we have $\|\boldsymbol{x}\|_{\infty} \leq 1$. Hence, applying theorem 9.4.4, we can find $\boldsymbol{q}$ such that $\boldsymbol{q}^{T} \boldsymbol{z} \leq \mathrm{OPT}+\varepsilon$ and $\|\boldsymbol{q}-\boldsymbol{x}\|_{2}+\|\boldsymbol{q}-\boldsymbol{y}\|_{2} \leq \varepsilon$ for some $\boldsymbol{x} \in \operatorname{conv}\left(\mathcal{I}_{1}\right)$ and $\boldsymbol{y} \in \operatorname{conv}\left(\mathcal{I}_{2}\right)$. Using (9.20), we have the coordinate wise minimum of $\boldsymbol{x}, \boldsymbol{y}$, i.e. $\min \{\boldsymbol{x}, \boldsymbol{y}\}$, is in $\operatorname{conv}\left(\mathcal{I}_{1}\right) \cap \operatorname{conv}\left(\mathcal{I}_{2}\right)$. Since $\|\boldsymbol{q}-\min \{\boldsymbol{x}, \boldsymbol{y}\}\|_{2} \leq\|\boldsymbol{q}-\boldsymbol{x}\|_{2}+\|\boldsymbol{q}-\boldsymbol{y}\|_{2} \leq \varepsilon$, we have

$$
(\min \{\boldsymbol{x}, \boldsymbol{y}\})^{T} \boldsymbol{z} \leq \mathrm{OPT}+n M \varepsilon
$$

Hence, we have a feasible point $\min \{\boldsymbol{x}, \boldsymbol{y}\}$ which has value close to optimal and Lemma 9.4.6 shows that $\|\min (\boldsymbol{x}, \boldsymbol{y})-s\|_{\infty} \leq 2 n^{2} M^{2} \varepsilon$ where $s$ is the optimal solution. Hence, we have $\|\boldsymbol{q}-s\|_{\infty} \leq 2 n^{2} M^{2} \varepsilon+\varepsilon$. Picking $\varepsilon=\frac{1}{6 n^{2} M^{2}}$, we have $\|\boldsymbol{q}-\boldsymbol{s}\|_{\infty}<\frac{1}{2}$ and hence, we can get the optimal solution by rounding to the nearest integer.

Since optimization over $\mathcal{I}_{1}$ and $\mathcal{I}_{2}$ involves applying greedy method on certain vectors, it takes only $O$ (GO) time. Theorem 9.4 .4 shows it only takes $O\left(n \mathrm{GO} \log (n M)+n^{3} \log ^{O(1)}(n M)\right)$ in finding $\operatorname{such} \boldsymbol{q}$.

This gives the following corollary.
Corollary 9.4.8. We have $O\left(n \min \left(n \mathcal{T}_{\text {ind }}, r \mathcal{T}_{\text {rank }}\right) \log (n M)+n^{3} \log ^{O(1)} n M\right)$ time algorithm for weighted matroid intersection. Here $\mathcal{T}_{\text {ind }}$ is the time needed to check if a subset is independent, and $\mathcal{T}_{\text {rank }}$ is the time needed to compute the rank of a given subset.

Proof. By Theorem 9.4.7, it suffices to show that the optimization oracle for the matroid polytope can be implemented in $O\left(n \mathcal{T}_{\text {ind }}\right)$ and $O\left(r \mathcal{T}_{\text {rank }} \log n\right)$ time. This is simply attained by the well-known greedy algorithm which iterates through all the positively-weighted elements in decreasing order, and adds an element to our candidate independent set whenever possible.

For the independence oracle, this involves one oracle call for each element. On the other hand, for the rank oracle, we can find the next element to add by binary search which takes time $O\left(\mathcal{T}_{\text {rank }} \log n\right)$. Since there are at most $r$ elements to add, we have the desired running time.

### 9.4.3 Submodular Flow

Let $G=(V, E)$ be a directed graphwith $|E|=m$, let $f$ be a submodular function on $\mathbb{R}^{V}$ with $|V|=n$, $f(\emptyset)=0$ and $f(V)=0$, and let $A$ be the incidence matrix of $G$. In this section we consider the submodular flow problem

$$
\begin{align*}
\text { Minimize } & \langle c, \varphi\rangle  \tag{9.21}\\
\text { subject to } & l(e) \leq \varphi(e) \leq u(e) \quad \forall e \in E \\
& x(v)=(A \varphi)(v) \quad \forall v \in V \\
& \sum_{v \in S} x(v) \leq f(S) \quad \forall S \subseteq V
\end{align*}
$$

where $c \in \mathbb{Z}^{E}, l \in \mathbb{Z}^{E}, u \in \mathbb{Z}^{E}$ where $C=\|\boldsymbol{c}\|_{\infty}$ and $U=\max \left(\|u\|_{\infty},\|l\|_{\infty}, \max _{S \subset V}|f(S)|\right)$. Here $c$ is the cost on edges, $\varphi$ is the flow on edges, $l$ and $u$ are lower and upper bounds on the amount of flow on the edges, and $x(v)$ is the net flow out of vertex $v$. The submodular function $f$ upper bounds the total net flow out of any subset $S$ of vertices by $f(S)$.

Theorem 9.4.9. Suppose that the cost vector $\boldsymbol{c}$ is integer weight with $\|\boldsymbol{c}\|_{\infty} \leq C$ and the capacity vector and the submodular function satisfy $U=\max \left(\|u\|_{\infty},\|l\|_{\infty}, \max _{S \subset V}|f(S)|\right)$. Then, we can solve the submodular flow problem (9.21) in time $O\left(n^{2} \mathrm{EO} \log (m C U)+n^{3} \log ^{O(1)}(m C U)\right)$ where EO is the cost of function evaluation.

Proof. First, we can assume $l(e) \leq u(e)$ for every edge $e$, otherwise, the problem is infeasible. Now, we apply a similar transformation in [110] to modify the graph. We create a new vertex $v_{0}$. For every vertex $v$ in $V$, we create a edge from $v_{0}$ to $v$ with capacity lower bounded by 0 , upper bounded by $4 n U$, and with cost $2 m C U$. Edmonds and Giles showed that the submodular flow polytope is integral [83]. Hence, there is an integral optimal flow on this new graph. If the optimal flow passes through the
newly created edge, then it has cost at least $2 m C U-m C U$ because the cost of all other edges in total has at least $-m C U$. That means the optimal flow has the cost larger than $m C U$ which is impossible. So the optimal flow does not use the newly created edges and vertex and hence the optimal flow in the new problem gives the optimal solution of the original problem. Next, we note that for any $\varphi$ on the original graph such that $l(e) \leq \varphi(e) \leq u(e)$, we can send suitable amount of flow from $v_{0}$ to $v$ to make $\varphi$ feasible. Hence, this modification makes the feasibility problem trivial.

Lemma 9.4.6 shows that we can assume the new problem has an unique solution and it only blows up $C$ by a $(m U)^{O(1)}$ factors.

Note that the optimal value is an integer and its absolute value at most $m C U$. By binary search, we can assume we know the optimal value OPT. Now, we reduce the problem to finding a feasible $\varphi$ with $\{\langle d, \varphi\rangle \leq \mathrm{OPT}+\varepsilon\}$ with $\varepsilon$ determined later. Let $P_{\varepsilon}$ be the set of such $\varphi$. Note that $P_{\varepsilon}=K_{1, \varepsilon} \cap K_{2, \varepsilon}$ where

$$
\begin{aligned}
& K_{1, \varepsilon}=\left\{\begin{array}{c}
l(e) \leq \varphi(e) \leq u(e) \quad \forall e \in E \\
\left.x \in \mathbb{R}^{V} \text { such that } \begin{array}{c}
x(v)=(A \varphi)(v) \quad \forall v \in V \quad \text { for some } \varphi \\
\langle d, \varphi\rangle \leq \mathrm{OPT}+\varepsilon
\end{array}\right\} \\
K_{2, \varepsilon}=\left\{y \in \mathbb{R}^{V} \text { such that } \sum_{v \in S} y(v) \leq f(S) \quad \forall S \subseteq V, \sum_{v \in V} y(v)=f(V)\right\}
\end{array} .\right.
\end{aligned}
$$

Note that the extra condition $\sum_{v} y(v)=f(V)$ is valid because $\sum_{v} y(v)=\sum_{v}(A \varphi)(v)=0$ and $f(V)=0$, and $K_{1, \varepsilon}$ has radius bounded by $O\left((m C U)^{O(1)}\right)$ and $K_{2, \varepsilon}$ has radius bounded by $O(n U)$. Furthermore, for any vector $c \in \mathbb{R}^{V}$, we note that

$$
\begin{aligned}
\max _{x \in K_{1, \varepsilon}}\langle c, x\rangle & =\max _{l \leq \varphi \leq u,\langle d, \varphi\rangle \leq \mathrm{OPT}+\varepsilon, x=A \varphi}\langle c, x\rangle \\
& =\max _{l \leq \varphi \leq u,\langle d, \varphi\rangle \leq \mathrm{OPT}+\varepsilon}\langle c, A \varphi\rangle \\
& =\max _{l \leq \varphi \leq u,\langle d, \varphi\rangle \leq \mathrm{OPT}+\varepsilon}\left\langle A^{T} c, \varphi\right\rangle
\end{aligned}
$$

To solve this problem, again we can do a binary search on $\langle d, \varphi\rangle$ and reduce the problem to

$$
\max _{l \leq \varphi \leq u,\langle d, \varphi\rangle=K}\left\langle A^{T} c, \varphi\right\rangle
$$

for some value of $K$. Since $A^{T} c$ is fixed, this is a linear program with only the box constraints and an extra equality constraint. Hence, it can be solved in nearly linear time [164, Thm 17, ArXiv v1]. As the optimization oracle for $K_{1, \varepsilon}$ involves only computing $A^{T} c$ and solving this simple linear program, it takes only $O\left(n^{2} \log ^{O(1)}(m C U / \varepsilon)\right)$ time. On the other hand, since $K_{2, \varepsilon}$ is just a base polyhedron, the optimization oracle for $K_{2, \varepsilon}$ can be done by greedy method and only takes $O(n \mathrm{EO})$ time.

Applying Theorem 9.4.4, we can find $q$ such that $\|q-x\|_{2}+\|q-y\|_{2} \leq \delta$ for some $x \in K_{1, \varepsilon}, y \in K_{2, \varepsilon}$ and $\delta$ to be chosen later. According to the definition of $K_{1, \varepsilon}$, there is $\varphi$ such that $l(e) \leq \varphi(e) \leq u(e)$ and $x(v)=(A \varphi)(v)$ for all $v$ and $\langle d, \varphi\rangle \leq$ OPT $+\varepsilon$. Since $\|y-x\|_{2} \leq 2 \delta$, that means $|y(v)-(A \varphi)(v)| \leq 2 \delta$ for all $v$.

- Case 1) If $y(v) \geq(A \varphi)(v)$, then we can replace $y(v)$ by $(A \varphi)(v)$, note that $y$ is still in $K_{2, \varepsilon}$ because of the submodular constraints.
- Case 2) If $y(v) \leq(A \varphi)(v)$, then we can send a suitable amount of flow from $v_{0}$ to $v$ to make $\varphi$ feasible $y(v) \leq(A \varphi)(v)$.

Note that under this modification, we increased the objective value by $(\delta n)(2 m C U)$ because the new edge cost $2 m C U$ per unit of flow. Hence, we find a flow $\varphi$ which is feasible in new graph with objective
value $\varepsilon+(\delta n)(2 m C U)$ far from optimum value. By picking $\delta=\frac{1}{2 m n C U}$, we have the value $2 \varepsilon$ far from OPT. Now, we use Lemma 9.4.6 to shows that when $\varepsilon$ is small enough, i.e, $\frac{1}{(m C U)^{c}}$ for some constant $c$, then we can guarantee that $\left\|y-x^{*}\right\|_{\infty} \leq \frac{1}{4}$ where $x^{*}$ is the optimal demand. Now, we note that $\|q-y\|_{2} \leq \delta$ and we note that we only modify $y$ by a small amount, we in fact have $\left\|q-x^{*}\right\|_{\infty}<\frac{1}{2}$. Hence, we can read off the solution $x^{*}$ by rounding $q$ to the nearest integer. Note that we only need to solve the problem $K_{1, \varepsilon} \cap K_{2, \varepsilon}$ to $\frac{1}{(m C U)^{\Theta(1)}}$ accuracy and the optimization oracle for $K_{1, \varepsilon}$ and $K_{2, \varepsilon}$ takes time $O\left(n^{2} \log ^{O(1)}(m C U)\right)$ and $O(n \mathrm{EO})$ respectively. Hence, Theorem 9.4 .4 shows that it takes $O\left(n^{2} \mathrm{EO} \log (m C U)+n^{3} \log ^{O(1)}(m C U)\right)$ time to find $x^{*}$ exactly.

After getting $x^{*}$, one can find $\varphi^{*}$ by solving a min cost flow problem using interior point method (Chapter 6), which takes $O\left(m \sqrt{n} \log ^{O(1)}(m C U)\right)$ time.

## - 9.4.4 Affine Subspace of Convex Set

In this section, we give another example about using optimization oracle directly via regularization. We consider the following optimization problem

$$
\begin{equation*}
\max _{\boldsymbol{x} \in K \text { and } \boldsymbol{A} \boldsymbol{x}=\boldsymbol{b}}\langle\boldsymbol{c}, \boldsymbol{x}\rangle \tag{9.22}
\end{equation*}
$$

where $\boldsymbol{x}, \boldsymbol{c} \in \mathbb{R}^{n}, K$ is a convex subset of $\mathbb{R}^{n}, \boldsymbol{A} \in \mathbb{R}^{r \times n}$ and $\boldsymbol{b} \in \mathbb{R}^{m}$. We suppose that $r \ll n$ and thus, the goal of this subsection is to show how to obtain an algorithm takes only $\tilde{O}(r)$ many iterations. To do this, we assume a slightly stronger optimization oracle for $K$ :

Definition 9.4.10. Given a convex set $K$ and $\delta>0$. A $\delta$-2nd-order-optimization oracle for $K$ is a function on $\mathbb{R}^{n}$ such that for any input $\boldsymbol{c} \in \mathbb{R}^{n}$ and $\lambda>0$, it outputs $\boldsymbol{y}$ such that

$$
\max _{\boldsymbol{x} \in K}\left(\langle\boldsymbol{c}, \boldsymbol{x}\rangle-\lambda\|\boldsymbol{x}\|^{2}\right) \leq \delta+\langle\boldsymbol{c}, \boldsymbol{y}\rangle-\lambda\|\boldsymbol{y}\|^{2}
$$

We denote by $\mathrm{OO}_{\delta, \lambda}^{(2)}(K)$ the time complexity of this oracle.
The strategy for solving this problem is very similar to the intersection problem and hence some details are omitted.

Theorem 9.4.11. Assume that $\max _{\boldsymbol{x} \in K}\|\boldsymbol{x}\|_{2}<M,\|\boldsymbol{b}\|_{2}<M,\|\boldsymbol{c}\|_{2}<M,\|\boldsymbol{A}\|_{2}<M$ and $\lambda_{\min }(\boldsymbol{A})>$ $1 / M$. Assume that $K \cap\{\boldsymbol{A x}=\boldsymbol{b}\} \neq \emptyset$ and we have $\varepsilon$-2nd-order-optimization oracle for every $\varepsilon>0$. For $0<\delta<1$, we can find $\boldsymbol{z} \in K$ such that

$$
\max _{\boldsymbol{x} \in K \text { and } \boldsymbol{A x}=\boldsymbol{b}}\langle\boldsymbol{c}, \boldsymbol{x}\rangle \leq \delta+\langle\boldsymbol{c}, \boldsymbol{z}\rangle
$$

and $\|\boldsymbol{A} \boldsymbol{z}-\boldsymbol{b}\|_{2} \leq \delta$. This algorithm takes time

$$
O\left(r \mathrm{OO}_{\eta, \lambda}^{(2)}(K) \log \left(\frac{n M}{\delta}\right)+r^{3} \log ^{O(1)}\left(\frac{n M}{\delta}\right)\right)
$$

where $r$ is the number of rows in $\boldsymbol{A}, \eta=\left(\frac{\delta}{n M}\right)^{\Theta(1)}$ and $\lambda=\left(\frac{\delta}{n M}\right)^{\Theta(1)}$.
Proof. The proof is based on the minimax problem

$$
\mathrm{OPT}_{\lambda} \stackrel{\text { def }}{=} \min _{\|\boldsymbol{\eta}\|_{2} \leq 1 / \lambda} \max _{\boldsymbol{x} \in K}\langle\boldsymbol{c}, \boldsymbol{x}\rangle+\langle\boldsymbol{\eta}, \boldsymbol{A} \boldsymbol{x}-\boldsymbol{b}\rangle-\frac{1}{\lambda}\|\boldsymbol{x}\|_{2}^{2}
$$

where $\lambda=\left(\frac{\delta}{n M}\right)^{c}$ for some large constant $c$. We note that

$$
\begin{aligned}
\mathrm{OPT}_{\lambda} & =\max _{\boldsymbol{x} \in K} \min _{\|\boldsymbol{\eta}\|_{2} \leq 1 / \lambda}\langle\boldsymbol{c}, \boldsymbol{x}\rangle+\langle\boldsymbol{\eta}, \boldsymbol{A} \boldsymbol{x}-\boldsymbol{b}\rangle-\frac{1}{\lambda}\|\boldsymbol{x}\|_{2}^{2} \\
& =\max _{\boldsymbol{x} \in K}\langle\boldsymbol{c}, \boldsymbol{x}\rangle-\frac{1}{\lambda}\|\boldsymbol{A} \boldsymbol{x}-\boldsymbol{b}\|_{2}-\frac{1}{\lambda}\|\boldsymbol{x}\|_{2}^{2} .
\end{aligned}
$$

Since $\lambda_{\min }(\boldsymbol{A})>1 / M$ and the set $K$ is bounded by $M$, one can show that the saddle point $\left(\boldsymbol{x}^{*}, \boldsymbol{\eta}^{*}\right)$ of the minimax problem gives a good enough solution $\boldsymbol{x}$ for the original problem for large enough constant $c$.

For any $\boldsymbol{\eta}$, we define

$$
\boldsymbol{x}_{\boldsymbol{\eta}}=\underset{\boldsymbol{x} \in K}{\arg \max }\langle\boldsymbol{c}, \boldsymbol{x}\rangle+\langle\boldsymbol{\eta}, \boldsymbol{A} \boldsymbol{x}-\boldsymbol{b}\rangle-\frac{1}{\lambda}\|\boldsymbol{x}\|_{2}^{2}
$$

Since the problem is strongly concave in $\boldsymbol{x}$, one can prove that

$$
\left\|\boldsymbol{x}_{\boldsymbol{\eta}}-\boldsymbol{x}^{*}\right\|_{2} \leq\left(\frac{n M}{\delta}\right)^{O(1)}\left\|\boldsymbol{\eta}-\boldsymbol{\eta}^{*}\right\|_{2}
$$

Hence, we can first find an approximate minimizer of the function $f(\boldsymbol{\eta})=\max _{\boldsymbol{x} \in K}\langle\boldsymbol{c}, \boldsymbol{x}\rangle+\langle\boldsymbol{\eta}, \boldsymbol{A} \boldsymbol{x}-\boldsymbol{b}\rangle-$ $\frac{1}{\lambda}\|\boldsymbol{x}\|_{2}^{2}$ and use the oracle to find $\boldsymbol{x}_{\boldsymbol{\eta}}$.

To find an approximate minimizer of $f$, we note that the subgradient of $f$ can be found using the optimization oracle similar to Theorem 9.4.4. Hence, the result follows from our cutting plane method and the fact that $\boldsymbol{\eta} \in \mathbb{R}^{r}$.

Remark 9.4.12. In Chapter 6, we considered the special case $K=\left\{\boldsymbol{x}: 0 \leq x_{i} \leq 1\right\}$ and showed that it can be solved in $\tilde{O}(\sqrt{r})$ iterations using interior point methods. This gives the current fastest algorithm for the maximum flow problem on directed weighted graphs. This result generalizes that result to any convex set $K$ but with $\tilde{O}(r)$ iterations. This suggests the following open problem: under what condition on $K$ can one optimize linear functions over affine subspaces of $K$ with $r$ constraints in $\tilde{O}(\sqrt{r})$ iterations?

# Submodular Minimization In Nearly-Cubic Time 

## ■ 10.1 Introduction

Submodular functions and submodular function minimization (SFM) are fundamental to the field of combinatorial optimization. Examples of submodular functions include graph cut functions, set coverage function, and utility functions from economics. Since the seminal work by Edmonds in 1970 [82], submodular functions and the problem of minimizing such functions (i.e. submodular function minimization) have served as a popular modeling and optimization tool in various fields such as theoretical computer science, operations research, game theory, and most recently, machine learning. Given its prevalence, fast algorithms for SFM are of immense interest both in theory and in practice.

Throughout this chapter, we consider the standard formulation of SFM: we are given a submodular function $f$ defined over the subsets of a $n$-element ground set. The values of $f$ are integers, have absolute value at most $M$, and are evaluated by querying an oracle that takes time EO. Our goal is to produce an algorithm that solves this SFM problem, i.e. finds a minimizer of $f$, while minimizing both the number of oracle calls made and the total running time.

We provide new $O\left(n^{2} \log n M \cdot \mathrm{EO}+n^{3} \log ^{O(1)} n M\right)$ and $O\left(n^{3} \log ^{2} n \cdot \mathrm{EO}+n^{4} \log ^{O(1)} n\right)$ time algorithms for SFM. These algorithms improve upon the previous fastest weakly and strongly polynomial time algorithms for SFM which had a a running time of $O\left(\left(n^{4} \cdot \mathrm{EO}+n^{5}\right) \log M\right)[118]$ and $O\left(n^{5} \cdot \mathrm{EO}+n^{6}\right)$ [216] respectively. Consequently, we improve the running times in both regimes by roughly a factor of $O\left(n^{2}\right)$.

Both of our algorithms bear resemblance to the classic approach of Grötschel, Lovász and Schrijver [110, 109] using the Lovász extension. In fact our weakly polynomial time algorithm directly uses the Lovász extension as well as the results of Chapter 9 to achieve these results. Our strongly polynomial time algorithm also uses the Lovász extension, along with more modern tools from the past 15 years.

At a high level, our strongly polynomial algorithms apply our cutting plane method in conjunction with techniques originally developed by Iwata, Fleischer, and Fujishige (IFF) [121]. Our cutting plane method is performed for enough iterations to sandwich the feasible region in a narrow strip from which useful structural information about the minimizers can be deduced. Our ability to derive the new information hinges on a significant extension of IFF techniques.

Over the past few decades, SFM has drawn considerable attention from various research communities, most recently in machine learning [24, 149]. Given this abundant interest in SFM, we hope that our ideas will be of value in various practical applications. Indeed, one of the critiques against existing theoretical algorithms is that their running time is too slow to be practical. Our contribution, on the contrary, shows that this school of algorithms can actually be made fast theoretically and we hope it may potentially be competitive against heuristics which are more commonly used.

## ■ 10.1.1 Previous Work

Here we provide a brief survey of the history of algorithms for SFM. For a more comprehensive account of the rich history of SFM, we refer the readers to recent surveys [189, 120].

The first weakly and strongly polynomial time algorithms for SFM were based on the ellipsoid method [142] and were established in the foundational work of Grötschel, Lovász and Schrijver in 1980's [110, 109]. Their work was complemented by a landmark paper by Cunningham in 1985 which provided a pseudopolynomial algorithm that followed a flow-style algorithmic framework [61]. His tools foreshadowed much of the development in SFM that would take place 15 years later. Indeed, modern algorithms synthesize his framework with inspirations from various max flow algorithms.

The first such "flow style" strongly polynomial algorithms for SFM were discovered independently in the breakthrough papers by Schrijver [231] and Iwata, Fleischer, and Fujishige (IFF) [121]. Schrijver's algorithm has a running of $O\left(n^{8} \cdot \mathrm{EO}+n^{9}\right)$ and borrows ideas from the push-relabel algorithms [105, 68] for the maximum flow problem. On the other hand, IFF's algorithm runs in time $O\left(n^{7} \log n \cdot \mathrm{EO}\right)$ and $O\left(n^{5} \cdot\right.$ EO $\left.\log M\right)$, and applies a flow-scaling scheme with the aid of certain proximity-type lemmas as in the work of Tardos [247]. Their method has roots in flow algorithms such as [117, 106].

Subsequent work on SFM provided algorithms with considerably faster running time by extending the ideas in these two "genesis" papers [231, 121] in various novel directions [264, 87, 118, 216, 125]. Currently, the fastest weakly and strongly polynomial time algorithms for SFM have a running time of $O\left(\left(n^{4} \cdot \mathrm{EO}+n^{5}\right) \log M\right)[118]$ and $O\left(n^{5} \cdot \mathrm{EO}+n^{6}\right)$ [216] respectively. Despite this impressive track record, the running time has not been improved in the last eight years.

We remark that all of the previous algorithms for SFM proceed by maintaining a convex combination of $O(n)$ BFS's of the base polyhedron, and incrementally improving it in a relatively local manner. As we shall discuss in Section 10.1.2, our algorithms do not explicitly maintain a convex combination. This may be one of the fundamental reasons why our algorithms achieve a faster running time.

Finally, beyond the distinction between weakly and strongly polynomial time algorithms for SFM, there has been interest in another type of SFM algorithm, known as fully combinatorial algorithms in which only additions and subtractions are permitted. Previous such algorithms include [125, 118, 119]. We do not consider such algorithms in the remainder of this chapter and leave it as an open question if it is possible to turn our algorithms into fully combinatorial ones.

## ■ 10.1.2 Our Results and Techniques

In this chapter, we show how to improve upon the previous best known running times for SFM by a factor of $O\left(n^{2}\right)$ in both the strongly and weakly polynomial regimes. In Table 10.1 summarizes the running time of the previous algorithms as well as the running times of the fastest algorithms presented in this chapter.

Both our weakly and strongly polynomial algorithms for SFM utilize a convex relaxation of the submodular function, called the Lovász extension. Our algorithms apply our cutting plane method from Chapter 8 using a separation oracle given by the subgradient of the Lovász extension. To the best of the author's knowledge, Grötschel, Lovász and Schrijver were the first to formulate this convex optimization framework for SFM [110, 109].

For weakly polynomial algorithms, our contribution is two-fold. First, we show that cutting plane methods such as Vaidya's [253] can be applied to SFM to yield faster algorithms. Second, as our cutting plane method, Theorem 9.3.2, improves upon previous cutting plane algorithms and consequently the running time for SFM as well. This gives a running time of $O\left(n^{2} \log n M \cdot \mathrm{EO}+n^{3} \log O(1) n M\right)$, an improvement over the previous best algorithm by Iwata [118] by a factor of almost $O\left(n^{2}\right)$.

Our strongly polynomial algorithms, on the other hand, require substantially more innovation. We

| Years | Authors | Running times | Remarks |
| :---: | :---: | :---: | :---: |
| 1981,1988 | Grötschel, Lovasz, | $\widetilde{O}\left(n^{5} \cdot \mathrm{EO}+n^{7}\right)[189]$ | first weakly <br> and strongly |
| 1985 | Cunningham [61] | $O\left(M n^{6} \log n M \cdot \mathrm{EO}\right)$ | first pseudopoly |
| 2000 | Schrijver [231] | $O\left(n^{8} \cdot \mathrm{EO}+n^{9}\right)$ | first combin. strongly |
| 2000 | Iwata, Fleischer, | $O\left(n^{5} \cdot \mathrm{EO} \log M\right)$ <br> $O\left(n^{7} \log n \cdot \mathrm{EO}\right)$ | first combin. strongly |
| 2000 | Iwata, Fleischer [87] | $O\left(n^{7} \cdot \mathrm{EO}+n^{8}\right)$ |  |
| 2003 | Iwata [118] | $O\left(\left(n^{4} \cdot \mathrm{EO}+n^{5}\right) \log M\right)$ <br> $O\left(\left(n^{6} \cdot \mathrm{EO}+n^{7}\right) \log n\right)$ | current best weakly |
| 2003 | Vygen [264] | $O\left(n^{7} \cdot \mathrm{EO}+n^{8}\right)$ |  |
| 2007 | Orlin [216] | $O\left(n^{5} \cdot \mathrm{EO}+n^{6}\right)$ | current best strongly |
| 2009 | Iwata, Orlin [125] | $O\left(\left(n^{4} \cdot \mathrm{EO}+n^{5}\right) \log n M\right)$ <br> $O\left(\left(n^{5} \cdot \mathrm{EO}+n^{6}\right) \log n\right)$ |  |
| - | This Chapter | $\left.\begin{array}{c}O\left(n^{2} \log ^{3} M \cdot \mathrm{EO}+n^{3} \log O(1)\right. \\ O\left(n^{3} \log ^{2} n \cdot \mathrm{EO}+n^{4} \log O(1)\right.\end{array}\right)$ |  |

Table 10.1: Algorithms for submodular function minimization. Note that some of these algorithms were published in both conferences and journals, in which case the year we provided is the earlier one.
first begin with a very simple geometric argument that SFM can be solved in $O\left(n^{3} \log n \cdot\right.$ EO $)$ oracle calls (but in exponential time). This proof only uses Grunbaum's Theorem from convex geometry and is completely independent from the rest of this chapter. It was the starting point of our method and suggests that a running time of $\widetilde{O}\left(n^{3} \cdot \mathrm{EO}+n^{O(1)}\right)$ for submodular minimization is in principle achievable.

To make this existence result algorithmic, we first run cutting plane, Theorem 8.3.26, for enough iterations such that we compute either a minimizer or a set $P$ containing the minimizers that fits within in a narrow strip. This narrow strip consists of the intersection of two approximately parallel hyperplanes. If our narrow strip separates $P$ from one of the faces $x_{i}=0, x_{i}=1$, we can effectively eliminate the element $i$ from our consideration and reduce the dimension of our problem by 1 . Otherwise a pair of elements $p, q$ can be identified for which $q$ is guaranteed to be in any minimizer containing $p$ (but $p$ may not be contained in a minimizer). Our first algorithm deduces only one such pair at a time. This technique immediately suffices to achieve a $\widetilde{O}\left(n^{4} \cdot \mathrm{EO}+n^{5}\right)$ time algorithm for SFM (See Section 10.4.3). We then improve the running time to $\widetilde{O}\left(n^{3} \cdot \mathrm{EO}+n^{4}\right)$ by showing how to deduce many such pairs simultaneously. Similar to past algorithms, this structural information is deduced from a point in the so-called base polyhedron (See Section 10.2).

Readers well-versed in SFM literature may recognize that our strongly polynomial algorithms are reminiscent of the scaling-based approach first used by IFF [121] and later in [118, 125]. While both approaches share the same skeleton, there are differences as to how structural information about minimizers is deduced. A comparison of our algorithms and previous ones are presented in Section 10.5.

Finally, there is one more crucial difference between these algorithms which we believe is responsible for much of our speedup. One common feature shared by all the previous algorithms is that they maintain a convex combination of $O(n)$ BFS's of the base polyhedron, and incrementally improve on it by introducing new BFS's by making local changes to existing ones. Our algorithms, on the other hand, choose new BFS's by the cutting plane method. Because of this, our algorithm considers the geometry of the existing BFS's where each of them has influences over the choice of the next BFS. In some sense, our next BFS is chosen in a more "global" manner.

## - 10.1.3 Organization

The rest of this chapter is organized as follows. We first begin with a gentle introduction to submodular functions in Section 10.2. In Section 10.3, we apply our cutting plane method to SFM to obtain a faster weakly polynomial algorithms. In Section 10.4 we then present our results for achieving better strongly polynomial algorithms, where a warm-up $\widetilde{O}\left(n^{4} \cdot \mathrm{EO}+n^{5}\right)$ algorithm is given before the fullfledged $\widetilde{O}\left(n^{3} \cdot \mathrm{EO}+n^{4}\right)$ algorithm. Finally, we end this chaper with a discussion and comparison between our algorithms and previous ones in Section 10.5.

We note that there are a few results in this chapter that can be read fairly independently of the rest of this chapter. In Theorem 10.3.2 we show how Vaidya's algorithm can be applied to SFM to obtain a faster weakly polynomial running time. Also in Theorem 10.4.1 we present a simple geometric argument that SFM can be solved with $O\left(n^{3} \log n \cdot E O\right)$ oracle calls but with exponential time. These results can be read with only a working knowledge of the Lovasz extension of submodular functions.

## ■ 10.2 Preliminaries

Here we introduce background on submodular function minimization (SFM) and notation that we use throughout this chapter. Our exposition is kept to a minimal amount sufficient for our purposes. We refer interested readers to the extensive survey by McCormick [189] for further intuition.

## ■ 10.2.1 Submodular Function Minimization

Throughout the rest of this chapter, let $V=\{1, \ldots, n\}=[n]$ denote a ground set and let $f: 2^{V} \longrightarrow \mathbb{Z}$ denote a submodular function defined on subsets of this ground set. We use $V$ and $[n]$ interchangeably and let $[0] \stackrel{\text { def }}{=} \emptyset$. We abuse notation by letting $S+i \stackrel{\text { def }}{=} S \cup\{i\}$ and $S-i \stackrel{\text { def }}{=} S \backslash\{i\}$ for an element $i \in V$ and a set $S \subseteq 2^{V}$. Formally, we call a function submodular if it obeys the following property of diminishing marginal differences:

Definition 10.2.1 (Submodularity). A function $f: 2^{V} \longrightarrow \mathbb{Z}$ is submodular if $f(T+i)-f(T) \leq$ $f(S+i)-f(S)$ for any $S \subseteq T$ and $i \in V \backslash T$.

For convenience we assume without loss of generality that $f(\emptyset)=0$ by replacing $f(S)$ by $f(S)-f(\emptyset)$ for all $S$. We also let $M \xlongequal{\text { def }} \max _{S \in 2^{V}}|f(S)|$.

The central goal of this chapter is to design algorithms for SFM, i.e. computing the minimizer of $f$. We call such an algorithm strongly polynomial if its running time depends only polynomially on $n$ and EO, the time needed to compute $f(S)$ for a set $S$, and we call such an algorithm weakly polynomial if it also depends polylogarithmically on $M$.

## ■ 10.2.2 Lovász Extension

Our new algorithms for SFM all consider a convex relaxation of a submodular function, known as the Lovász extension, and then carefully apply our cutting plane methods to it. Here we formally introduce the Lovász extension and present basic facts that we use throughout this chapter.

The Lovász extension of $\hat{f}:[0,1]^{n} \longrightarrow \mathbb{R}$ of our submodular function $f$ is defined for all $\boldsymbol{x}$ by

$$
\hat{f}(\boldsymbol{x}) \stackrel{\text { def }}{=} \mathbb{E}_{t \sim[0,1]}\left[f\left(\left\{i: x_{i} \geq t\right\}\right)\right],
$$

where $t \sim[0,1]$ is drawn uniformly at random from $[0,1]$. The Lovász extension allows us to reduce SFM to minimizing a convex function defined over the interior of the hypercube. Below we state that the Lovász extension is a convex relaxation of $f$ and that it can be evaluated efficiently.

Theorem 10.2.2. The Lovász extension $\hat{f}$ satisfies the following properties:

1. $\hat{f}$ is convex and $\min _{x \in[0,1]^{n}} \hat{f}(\boldsymbol{x})=\min _{S \subset[n]} f(S)$;
2. $f(S)=\hat{f}\left(I_{S}\right)$, where $I_{S}$ is the characteristic vector for $S$, i.e. $I_{S}(i)=\left\{\begin{array}{ll}1 & \text { if } i \in S \\ 0 & \text { if } i \notin S\end{array}\right.$;
3. If $S$ is a minimizer of $f$, then $I_{S}$ is a minimizer of $\hat{f}$;
4. Suppose $x_{1} \geq \cdots \geq x_{n} \geq x_{n+1} \stackrel{\text { def }}{=} 0$, then

$$
\hat{f}(\boldsymbol{x})=\sum_{i=1}^{n} f([i])\left(x_{i}-x_{i+1}\right)=\sum_{i=1}^{n}(f([i])-f([i-1])) x_{i}
$$

Proof. See [109] or any standard textbook on combinatorial optimization, e.g. [232].

Next we show that we can efficiently compute a subgradient of the Lovász or alternatively, a separating hyperplane for the set of minimizers of our submdoular function $f$. First we remind the reader of the definition of a separation oracle, and then we prove the necessary properties of the hyperplane, Theorem 10.2.4.

Definition 10.2.3 (separation oracle, Defintion 2.3.5 restated for Lovász extension). Given a point $\bar{x}$ and a convex function $\hat{f}$ over a convex set $P, \boldsymbol{a}^{T} \boldsymbol{x} \leq b$ is a separating hyperplane if $\boldsymbol{a}^{T} \bar{x} \geq b$ and any minimizer $x^{*}$ of $\hat{f}$ over $P$ satisfies $\boldsymbol{a}^{T} x^{*} \leq b$.

Theorem 10.2.4. Given a point $\bar{x} \in[0,1]^{n}$ assume without loss of generality (by re-indexing the coordinates) that $\bar{x}_{1} \geq \cdots \geq \bar{x}_{n}$. Then the following inequality is a valid separating hyperplane for $\boldsymbol{x}$ and $f$ :

$$
\sum_{i=1}^{n}(f([i])-f([i-1])) x_{i} \leq \hat{f}(\bar{x})
$$

i.e., it satisfies the following:

1. (separating) $\bar{x}$ lies on $\sum_{i=1}^{n}(f([i])-f([i-1])) x_{i} \leq \hat{f}(\bar{x})$.
2. (valid) For any $\boldsymbol{x}$, we have $\sum_{i=1}^{n}(f([i])-f([i-1])) x_{i} \leq \hat{f}(\boldsymbol{x})$. In particular, $\sum_{i=1}^{n}(f([i])-$ $f([i-1])) x_{i}^{*} \leq \hat{f}(\bar{x})$ for any minimizer $\boldsymbol{x}^{*}$, i.e. the separating hyperplane does not cut out any minimizer.

Moreover, such a hyperplane can be computed with $n$ oracle calls to $f$ and in time $O\left(n \cdot E O+n^{2}\right)$.
Proof. Note that by Theorem 10.2.2 we have that $\sum_{i \in[n]}(f([i])-f([i-1])) x_{i}=\hat{f}(\bar{x})$ and thus the hyperplane satisfies the separating condition. Moreover, clearly computing it only takes time $O(n$. $\mathrm{EO}+n^{2}$ ) as we simply need to sort the coordinates and evaluate $f$ at $n$ points, i.e. each of the [i]. All that remains is to show that the hyperplane satisfies the valid condition.

Let $L^{(t)} \stackrel{\text { def }}{=}\left\{i: x_{i} \geq t\right\}$. Recall that $\hat{f}(\boldsymbol{x})=\mathbb{E}_{t \sim[0,1]}\left[f\left(L_{t}\right)\right]$. Thus $\hat{f}(\boldsymbol{x})$ can be written as a convex combination $\hat{f}(\boldsymbol{x})=\sum_{t} \alpha_{t} f\left(L^{(t)}\right)$, where $\alpha_{t} \geq 0$ and $\sum_{t} \alpha_{t}=1$. However, by diminishing marginal
differences we see that for all $t$

$$
\begin{aligned}
\sum_{i \in[n]}(f([i])-f([i-1]))\left(I_{L^{(t)}}\right)_{i} & =\sum_{i \in L^{(t)}}(f([i])-f([i-1])) \\
& \leq \sum_{i \in L^{(t)}}\left(f\left([i] \cap L^{(t)}\right)-f\left([i-1] \cap L^{(t)}\right)\right) \\
& =f\left(L^{(t)}\right)-f(\emptyset)=f\left(L^{(t)}\right)
\end{aligned}
$$

and therefore since $\sum_{t} \alpha_{t} I_{L^{(t)}}=\boldsymbol{x}$ we have

$$
\sum_{i \in[n]}\left(f[i]-f([i-1]) x_{i}=\sum_{t} \alpha_{t} \sum_{i=1}^{n}(f([i])-f([i-1]))\left(I_{L^{(t)}}\right)_{i} \leq \sum_{t} \alpha_{t} f\left(L^{(t)}\right)=\hat{f}(\boldsymbol{x}) .\right.
$$

## ■ 10.2.3 Polyhedral Aspects of SFM

Here we provide a natural primal dual view of SFM that we use throughout the analysis. We provide a dual convex optimization program to minimizing the Lovász extension and provide several properties of these programs. We believe the material in this section helps crystallize some of the intuition behind our algorithm and we make heavy use of the notation presented in this section. However, we will not need to appeal to the strong duality of these programs in our proofs.

Consider the following primal and dual programs, where we use the shorthands $y(S)=\sum_{i \in S} y_{i}$ and $y_{i}^{-}=\min \left\{0, y_{i}\right\}$. Here the primal constraints are often called the base polyhedron $\mathcal{B}(f) \stackrel{\text { def }}{=}\{\boldsymbol{y} \in$ $\left.\mathbb{R}^{n}: y(S) \leq f(S) \forall S \subsetneq V, y(V)=f(V)\right\}$ and the dual program directly corresponds to minimizing the Lovász extension and thus $f$.

| Primal | Dual |
| :---: | :---: |
| $\max y^{-}(V)$ | $\min \hat{f}(\boldsymbol{x})$ |
| $y(S) \leq f(S) \forall S \subsetneq V$ | $0 \leq \boldsymbol{x} \leq 1$ |
| $y(V)=f(V)$ |  |

Theorem 10.2.5. $\boldsymbol{h}$ is a basic feasible solution (BFS) of the base polyhedron $\mathcal{B}(f)$ if and only if

$$
h_{i}=f\left(\left\{v_{1}, \ldots, v_{i}\right\}\right)-f\left(\left\{v_{1}, \ldots, v_{i-1}\right\}\right)
$$

for some permutation $v_{1}, \ldots, v_{n}$ of the ground set $V$. We call $v_{1}, \ldots, v_{n}$ the defining permutation of $\boldsymbol{h}$. We call $v_{i}$ precedes $v_{j}$ for $i<j$.
This theorem gives a nice characterization of the BFS's of $\mathcal{B}(f)$. It also gives the key observation underpinning our approach: the coefficients of each separating hyperplane in Theorem 10.2.4 precisely corresponds to a primal BFS (Theorem 10.2.5). Our analysis relies heavily on this connection. We re-state Theorem 10.2.4 in the language of BFS.

Lemma 10.2.6. We have $\boldsymbol{h}^{T} \boldsymbol{x} \leq \hat{f}(\boldsymbol{x})$ for any $\boldsymbol{x} \in[0,1]^{n}$ and BFS $\boldsymbol{h}$.
Proof. Any BFS is given by some permutation. Thus this is just Theorem 10.2.4 in disguise.
We also note that since the objective function of the primal program is non-linear, we cannot say that the optimal solution to the primal program is a BFS. Instead we only know that it is a convex combination of the BFS's that satisfy the following property. A proof can be found in any standard textbook on combinatorial optimization.

Theorem 10.2.7. The above primal and dual programs have no duality gap. Moreover, there always exists a primal optimal solution $\boldsymbol{y}=\sum_{k} \lambda^{(k)} \boldsymbol{h}^{(k)}$ with $\sum_{k} \lambda^{(k)}=1$ (a convex combination of BFS $\left.\boldsymbol{h}^{(k)}\right)$ s.t. any $i$ with $y_{i}<0$ precedes any $j$ with $y_{j}>0$ in the defining permutation for each BFS $\boldsymbol{h}^{(k)}$.

Our algorithms will maintain collections of BFS and use properties of $\boldsymbol{h} \in \mathcal{B}(f)$, i.e. convex combination of BFS. To simplify our analysis at several points we will want to assume that such a vector $\boldsymbol{h} \in \mathcal{B}(f)$ is non-degenerate, meaning it has both positive and negative entries. Below, we prove that such degenerate points in the base polytope immediately allow us to trivially solve the SFM problem.

Lemma 10.2.8 (Degenerate Hyperplanes). If $\boldsymbol{h} \in \mathcal{B}(f)$ is non-negative then $\emptyset$ is a minimizer of $f$ and if $\boldsymbol{h}$ is non-positive then $V$ is a minimizer of $f$.

Proof. While this follows immediately from Theorem 10.2.7, for completeness we prove this directly. Let $S \in 2^{V}$ be arbitrary. If $\boldsymbol{h} \in \mathcal{B}(f)$ is non-negative then by the we have

$$
f(S) \geq \boldsymbol{h}(S)=\sum_{i \in S} h_{i} \geq 0=f(\emptyset) .
$$

On the other hand if $\boldsymbol{h}$ is non-positive then by definition we have

$$
f(S) \geq \boldsymbol{h}(S)=\sum_{i \in S} h_{i} \geq \sum_{i \in V} h_{i}=h(V)=f(V) .
$$

## ■ 10.3 Improved Weakly Polynomial Algorithms for SFM

In this section we show how our cutting plane method can be used to obtain a $O\left(n^{2} \log n M \cdot \mathrm{EO}+\right.$ $\left.n^{3} \log ^{O(1)} n M\right)$ time algorithm for SFM. Our main result in this section is the following theorem, which shows how directly applying our results from earlier chapters to minimize the Lovász extension yields the desired running time.

Theorem 10.3.1. We have an $O\left(n^{2} \log n M \cdot E O+n^{3} \log ^{O(1)} n M\right)$ time algorithm for submodular function minimization.

Proof. We apply Theorem 9.3 .2 to the Lovász extension $\hat{f}:[0,1]^{n} \longrightarrow \mathbb{R}$ with the separation oracle given by Theorem 10.2.4. $\hat{f}$ fulfills the requirement on the domain as its domain $\Omega=[0,1]^{n}$ is symmetric about the point $(1 / 2, \ldots, 1 / 2)$ and has exactly $2 n$ constraints.

In the language of Theorem 9.3.2, our separation oracle is a ( 0,0 )-separation oracle with $\eta=0$ and $\delta=0$.

We first show that $\delta=0$. Firstly, our separating hyperplane can be written as

$$
\sum_{i=1}^{n}(f([i])-f([i-1])) x_{i} \leq \hat{f}(\bar{x})=\sum_{i=1}^{n}(f([i])-f([i-1])) \bar{x}_{i},
$$

where the equality follows from Theorem 10.2.2. Secondly, for any $\boldsymbol{x}$ with $\hat{f}(\boldsymbol{x}) \leq \hat{f}(\bar{x})$ we have by Theorem 10.2.4 that

$$
\sum_{i=1}^{n}(f([i])-f([i-1])) x_{i} \leq \hat{f}(\boldsymbol{x}) \leq \hat{f}(\bar{x})
$$

which implies that $\boldsymbol{x}$ is not cut away by the hyperplane.

Next we show that $\eta=0$. Our separating hyperplane induces a valid halfspace whenever it is not nonzero, i.e. $f([i]) \neq f([i-1])$ for some $i$. In the case that it is zero $f([i])=f([i-1]) \forall i$, by the same argument above, we have $\hat{f}(\bar{x})=\sum_{i=1}^{n}(f([i])-f([i-1])) \bar{x}_{i}=0$ and

$$
\hat{f}(\boldsymbol{x}) \geq \sum_{i=1}^{n}(f([i])-f([i-1])) x_{i}=0=\hat{f}(\bar{x}) .
$$

In other words, $\bar{x}$ is an exact minimizer, i.e. $\eta=0$.
Note that $|\hat{f}(\boldsymbol{x})|=\left|\mathbb{E}_{t \sim[0,1]}\left[f\left(\left\{i: x_{i} \geq t\right\}\right)\right]\right| \leq M$ as $M=\max _{S}|f(S)|$. Now plugging in $\alpha=\frac{1}{4 M}$ in the guarantee of Theorem 8.3.26, we can find a point $x^{*}$ such that

$$
\begin{aligned}
\hat{f}\left(x^{*}\right)-\min _{x \in[0,1]^{n}} \hat{f}(\boldsymbol{x}) & \leq \frac{1}{4 M}\left(\max _{\boldsymbol{x} \in[0,1]^{n}} \hat{f}(\boldsymbol{x})-\min _{x \in[0,1]^{n}} \hat{f}(\boldsymbol{x})\right) \\
& \leq \frac{1}{4 M}(2 M) \\
& <1
\end{aligned}
$$

We claim that $\min _{t \in[0,1]} f\left(\left\{i: x_{i}^{*} \geq t\right\}\right)$ is minimum. To see this, recall from 10.2.2 that $\hat{f}$ has an integer minimizer and hence $\min _{x \in[0,1]^{n}} \hat{f}(\boldsymbol{x})=\min _{S} f(S)$. Moreover, $\hat{f}\left(x^{*}\right)$ is a convex combination of $f\left(\left\{i: x_{i}^{*} \geq t\right\}\right)$ which gives

$$
1>\hat{f}\left(x^{*}\right)-\min _{x \in[0,1]^{n}} \hat{f}(\boldsymbol{x})=\hat{f}\left(x^{*}\right)-\min _{S} f(S) \geq \min _{t \in[0,1]} f\left(\left\{i: x_{i}^{*} \geq t\right\}\right)-\min _{S} f(S) .
$$

Since $f$ is integer-valued, we must then have $\min _{t \in[0,1]} f\left(\left\{i: x_{i}^{*} \geq t\right\}\right)=\min _{S} f(S)$ as desired. Since our separation oracle can be computed by $n$ oracle calls and runs in time $O\left(n \cdot \mathrm{EO}+n^{2}\right)$, by Theorem 9.3.2 the overall running time is then $O\left(n^{2} \log n M \cdot \mathrm{EO}+n^{3} \log ^{O(1)} n M\right)$ as claimed.

Needless to say the proof above completely depends on Theorem 9.3.2. We remark that one can use the Vaidya's cutting plane instead of ours to get a time complexity $O\left(n^{2} \log n M \cdot \mathrm{EO}+n^{\omega+1} \log ^{O(1)} n\right.$. $\log M)$. There is actually an alternate argument that gives a time complexity of $O\left(n^{2} \log M \cdot \mathrm{EO}+\right.$ $\left.n^{O(1)} \cdot \log M\right)$. Thus it requires slightly fewer oracle calls at the expense of slower running time. A proof is offered in this section, which can be skipped without any risk of discontinuation. This proof relies the following cutting plane method.

Theorem 10.3.2 ([36]). Given any convex set $K \subset[0,1]^{n}$ with a separation oracle of cost $S O$, in time $O\left(k S O+k n^{O(1)}\right)$ one can find either find a point $\boldsymbol{x} \in K$ or find a polytope $P$ such that $K \subset P$ and the volume of $K$ is at most $\left(\frac{2}{3}\right)^{k}$.

The Theorem allows us to decrease the volume of the feasible region by a factor of $\left(\frac{2}{3}\right)^{k}$ after $k$ iterations. Similar to above, we apply cutting plane to minimize $\hat{f}$ over the hypercube $[0,1]^{n}$ for $O(n \log M)$ iterations, and outputs any integral point in the remaining feasible region $P$.

Lemma 10.3.3. Let $x^{*}$ achieve the minimum function value $\hat{f}\left(x^{*}\right)$ among the points used to query the separation oracle. Then

1. $x^{*} \in P^{(k)}$, the current feasible region.
2. Any $\boldsymbol{x}$ with $\hat{f}(\boldsymbol{x}) \leq \hat{f}\left(x^{*}\right)$ belongs to $P^{(k)}$.
3. suppose $x_{i_{1}}^{*} \geq \cdots \geq x_{i_{n}}^{*}$ and let $S_{j}=\left\{i_{1}, \ldots, i_{j}\right\}$. Then $S_{l} \in \arg \min _{S_{j}} f\left(S_{j}\right)$ also belongs to $P^{(k)}$.

Proof. For any separating hyperplane $\boldsymbol{h}^{T} x \leq \hat{f}(\bar{x})$ given by $\bar{x}$, we have by Lemma 10.2 .6 that $\boldsymbol{h}^{T} x^{*} \leq$ $\hat{f}\left(x^{*}\right)$. Since $\hat{f}\left(x^{*}\right)$ is the minimum among all $\hat{f}(\bar{x}), \boldsymbol{h}^{T} x^{*} \leq \hat{f}(\bar{x})$ and hence $x^{*}$ is not removed by any new separating hyperplane. In other words, $x^{*} \in P^{(k)}$. The argument for (2) is analogous.

For (3), recall that by the definition of Lovász extension $\hat{f}\left(x^{*}\right)$ is a convex combination of $f\left(S_{j}\right)$ and thus the indicator variable $I_{S_{l}}$ for $S_{l}$ satisfies $f\left(I_{S_{l}}\right) \leq \hat{f}\left(x^{*}\right)$. By Lemma 10.2.6 again, this implies $\boldsymbol{h}^{T} I_{S_{l}} \leq f\left(I_{S_{l}}\right) \leq \hat{f}\left(x^{*}\right) \leq \hat{f}(\bar{x})$ for any separating hyperplane $\boldsymbol{h}^{T} x \leq \hat{f}(\bar{x})$.

Theorem 10.3.4. Suppose that we run Cutting Plane in Theorem 10.3.2 for $O(n \log M)$ iterations. Then $S_{l}$ from the last lemma also minimizes $f$.

Proof. We use the notations from the last lemma. After $k=K n \log _{2 / 3} M$ iterations, the volume of the feasible region $P^{(k)}$ is at most $1 / M^{K n}$. By the last lemma, $I_{S_{l}} \in P^{(k)}$.

Suppose for the sake of contradiction that $S$ minimizes $f$ but $f(S)<f\left(S_{l}\right)$. Since $f$ is integervalued, $f(S)+1 \leq f\left(S_{l}\right)$. Let $r \stackrel{\text { def }}{=} 1 / 6 M$. Consider the set $B \stackrel{\text { def }}{=}\left\{\boldsymbol{x}: 0 \leq x_{i} \leq r \forall \mathrm{i} \notin S, 1-r \leq x_{i} \leq\right.$ $1 \forall i \in S\}$. We claim that for $\boldsymbol{x} \in B$,

$$
\hat{f}(\boldsymbol{x}) \leq f(S)+1
$$

To show this, note that $f\left(\left\{i: x_{i} \geq t\right\}\right)=f(S)$ for $r<t \leq 1-r$ as $x_{i} \leq r$ for $i \notin S$ and $x_{i} \geq 1-r$ for $i \in S$. Now using conditional probability and $|f(T)| \leq M$ for any $T$,

$$
\begin{aligned}
\hat{f}(\boldsymbol{x})= & \mathbb{E}_{t \sim[0,1]}\left[f\left(\left\{i: x_{i} \geq t\right\}\right)\right] \\
= & (1-2 r) \mathbb{E}\left[f\left(\left\{i: x_{i} \geq t\right\}\right) \mid r<t \leq 1-r\right]+ \\
& \left.r\left(\mathbb{E}\left[f\left(\left\{i: x_{i} \geq t\right\}\right) \mid 0 \leq t \leq r\right]+\mathbb{E}\left[f\left(\left\{i: x_{i} \geq t\right\}\right) \mid 1-r \leq t \leq 1\right]\right]\right) \\
= & (1-r) f(S)+r\left(\mathbb{E}\left[f\left(\left\{i: x_{i} \geq t\right\}\right) \mid 0 \leq t \leq r+\mathbb{E}\left[f\left(\left\{i: x_{i} \geq t\right\}\right) \mid 1-r \leq t \leq 1\right]\right]\right) \\
\leq & (1-2 r) f(S)+2 r M \\
\leq & f(S)+4 r M \\
\leq & f(S)+1
\end{aligned}
$$

But now $B \subseteq P^{(k)}$ as $\hat{f}(\boldsymbol{x}) \leq f(S)+1 \leq f\left(S_{l}\right)$ and by (2) of the last lemma. This would lead to a contradiction since

$$
\operatorname{vol}(B)=\frac{1}{(6 M)^{n}}>\frac{1}{M^{K n}} \geq \operatorname{vol}\left(P^{(k)}\right)
$$

for sufficiently large $K$.
Corollary 10.3.5. There is an $O\left(n^{2} \log M \cdot E O+n^{O(1)} \log M\right)$ time algorithm for submodular function minimization.

Proof. This simply follows from the last lemma, Theorem 10.3.2, and the fact that our separation oracle runs in time $O\left(n \cdot \mathrm{EO}+n^{2}\right)$.

Curiously, we obtained $O(\log M)$ rather than $O(\log n M)$ as in our algorithm. We leave it as an open problem whether one can slightly improve our running time to $O\left(n^{2} \log M \cdot \mathrm{EO}+n^{3} \log ^{O(1)} n \cdot \log M\right)$. The rest of this chapter is devoted to obtaining better strongly polynomial running time.

## ■ 10.4 Improved Strongly Polynomial Algorithms for SFM

In this section we show how our cutting plane method can be used to obtain a $\widetilde{O}\left(n^{3} \cdot \mathrm{EO}+n^{4}\right)$ time algorithm for SFM, which improves over the currently fastest $O\left(n^{5} \cdot \mathrm{EO}+n^{6}\right)$ time algorithm by Orlin.

## ■ 10.4.1 Improved Oracle Complexity

We first present a simple geometric argument that $f$ can be minimized with just $O\left(n^{3} \log n \cdot\right.$ EO $)$ oracle calls. While this is our desired query complexity (and it improves upon the previous best known bounds by a factor of $O\left(n^{2}\right)$ unfortunately the algorithm runs in exponential time. Nevertheless, it does provide some insight into how our more efficient algorithms should proceed and it alone, does suggests that information theoretically, $O\left(n^{3} \log n \cdot \mathrm{EO}\right)$ calls suffice to solve SFM. In the rest of this chapter, we combine this insight with some of the existing SFM tools developed over the last decade to get improved polynomial time algorithms.
Theorem 10.4.1. Submodular functions can be minimized with $O\left(n^{3} \log n \cdot E O\right)$ oracle calls.
Proof. We use the cutting plane method in Theorem 10.3.2 with the separation oracle given by Theorem 10.2.4. This method reduce the volume of the feasible region by a factor of $\left(\frac{2}{3}\right)^{k}$ after $k$ iterations if the optimal has not found yet.

Now, we argue that after $O(n \log n)$ iterations of this procedure we have either found a minimizer of $f$ or we have enough information to reduce the dimension of the problem by 1 . To see this, first note that if the separation oracle ever returns a degenerate hyperplane, then by Lemma 10.2.8 then either $\emptyset$ or $V$ is the minimizer, which we can determine in time $O(\mathrm{EO}+n)$. Otherwise, after $100 n \log n$ iterations, our feasible region $P$ must have a volume of at most $1 / n^{10 n}$. In this case, we claim that the remaining integer points in $P$ all lie on a hyperplane. This holds, as if this was not the case, then there is a simplex $\triangle$, with integral vertices $v_{0}, v_{1}, \ldots, v_{n}$, contained in $P$. But then

$$
\operatorname{vol}(P) \geq \operatorname{vol}(\triangle)=\frac{1}{n!}\left|\operatorname{det}\left(v_{1}-v_{0} v_{2}-v_{0} \ldots v_{n}-v_{0}\right)\right| \geq \frac{1}{n!}
$$

where the last inequality holds since the determinant of an integral matrix is integral, yielding a contradiction.

In other words after $O(n \log n)$ iterations, we have reduced the dimension of all viable solutions by at least 1 . Thus, we can recurse by applying the cutting plane method to the lower dimensional feasible region, i.e. $P$ is (replaced by) the convex combination of all the remaining integer points. There is a minor technical issue we need to address as our space is now lower dimensional and the starting region is not necessarily the hypercube anymore and the starting volume is not necessarily equal to 1 .

We argue that the starting volume is bounded by $n^{O(n)}$. If this is indeed the case, then our previous argument still works as the volume goes down by a factor of $1 / n^{O(n)}$ in $O(n \log n)$ iterations.

Let $v \in P$ be an integer point. Now the $\operatorname{dim}(() P)$-dimensional ball of radius $\sqrt{n}$ centered at $v$ must contain all the other integer points in $P$ as any two points of $\{0,1\}^{n}$ are at most $\sqrt{n}$ apart. Thus the volume of $P$ is bounded by the volume of the ball which is $n^{O(n)}$. Now to get the volume down to $1 / n^{10 n}$, the number of iterations is still $O(n \log n)$.

In summary, we have reduced our dimension by 1 using $O(n \log n)$ iterations which requires $O\left(n^{2} \log n \cdot \mathrm{EO}\right)$ oracle calls (as each separating hyperplane is computed with $n \cdot \mathrm{EO}$ oracle calls). This can happen at most $n$ times. The overall query complexity is then $O\left(n^{3} \log n \cdot \mathrm{EO}\right)$.

Note that the minimizer $\boldsymbol{x}$ obtained may not be integral. This is not a problem as the definition of Lovász extension implies that if $\hat{f}(\boldsymbol{x})$ is minimal, then $f\left(\left\{i: x_{i} \geq t\right\}\right)$ is minimal for any $t \in[0,1]$.

We remark that this algorithm does not have a polynomial runtime. Even though all the integral vertices of $P$ lie on a hyperplane, the best way we know of that identifies it takes exponential time by checking for all the integer points $\{0,1\}^{n}$.

Remark 10.4.2. Note that this algorithm works for minimizing any convex function over the hypercube that obtains its optimal value at a vertex of the hypercube. Formally, our proof of Theorem 10.4.1
holds whenever a function $f: 2^{V} \longrightarrow \mathbb{R}^{n}$ admits a convex relaxation $\hat{f}$ with the following properties:

1. For every $S \subseteq V, \hat{f}\left(I_{S}\right)=f(S)$.
2. Every $\hat{f}(\boldsymbol{x})$ can be written as a convex combination $\sum_{S \in \mathcal{S}} \alpha_{S} f(S)$, where $\sum \alpha_{S}=1,|\mathcal{S}|=O(n)$, and $\mathcal{S}$ can be computed without any oracle call.
3. A subgradient $\partial \hat{f}(\boldsymbol{x})$ of $\hat{f}$ at any point $\boldsymbol{x} \in[0,1]^{n}$ can be computed with $O(n \cdot$ EO) oracle calls.

In this case, the proof of Theorem 10.4.1, implies that $\hat{f}$ and $f$ can be minimized with $O\left(n^{3} \log n \cdot \mathrm{EO}\right)$ oracle calls by using the separating hyperplane $\partial \hat{f}(\bar{x})^{T}(\boldsymbol{x}-\bar{x}) \leq 0$.

## ■ 10.4.2 Technical Tools

To improve upon the running time of the algorithm in the previous section, we use more structure of our submodular function $f$. Rather than merely showing that we can decrease the dimension of our SFM problem by 1 we show how we can reduce the degrees of freedom of our problem in a more principled way. In Section 10.4.2.1 we formally define the abstraction we use for this and discuss how to change our separation oracle to accommodate this abstraction, and in Section 10.4.2.2 we show how we can deduce these constraints. These tools serve as the foundation for the faster strongly polynomial time SFM algorithms we present in Section 10.4.3 and Section 10.4.4.

### 10.4.2.1 SFM over Ring Family

For the remainder of this chapter we consider a more general problem than SFM in which we wish to compute a minimizer of our submodular function $f$ over a ring family of the ground set $V=[n]$. A ring family $\mathcal{F}$ is a collection of subsets of $V$ such that for any $S_{1}, S_{2} \in \mathcal{F}$, we have $S_{1} \cup S_{2}, S_{1} \cap S_{2} \in \mathcal{F}$. Thus SFM corresponds to the special case where $\mathcal{F}$ consists of every subset of $V$. This generalization has been considered before in the literature and was essential to the IFF algorithm.

It is well known that any ring family $\mathcal{F}$ over $V$ can be represented by a directed graph $D=(V, A)$ where $S \in \mathcal{F}$ iff $S$ contains all of the descendants of any $i \in S$. An equivalent definition is that for any arc $(i, j) \in A, i \in S$ implies $j \in S$. It is customary to assume that $A$ is acyclic as any (directed) cycle of $A$ can be contracted (see section 10.4.3.1).

We denote by $R(i)$ the set of descendants of $i$ (including $i$ itself) and $Q(i)$ the set of ancestors of $i$ (including $i$ itself). Polyhedrally, an arc $(i, j) \in A$ can be encoded as the constraint $x_{i} \leq x_{j}$ as shown by the next lemma.

Lemma 10.4.3. Let $\mathcal{F}$ be a ring family over $V$ and $D=(V, A)$ be its directed acyclic graph representation. Suppose $f: V \longrightarrow \mathbb{R}$ is submodular with Lovász extension $\hat{f}$. Then the characteristic vector $I_{S}$ of any minimizer $S=\arg \min _{S \in \mathcal{F}} f(S)$ over $\mathcal{F}$ is also the solution to

$$
\begin{align*}
& \min \hat{f}(\boldsymbol{x}) \\
& \quad x_{i} \leq x_{j} \forall(i, j) \in A  \tag{10.1}\\
& \quad 0 \leq \boldsymbol{x} \leq 1
\end{align*}
$$

Proof. Let $x^{*}$ be a minimizer, and $L^{(t)}=\left\{i: x_{i}^{*} \geq t\right\}$. It is easy to check that the indicator variable $I_{L^{(t)}}$ satisfies (10.1) since $x^{*}$ does. Moreover, recall that $\hat{f}\left(x^{*}\right)=\mathbb{E}_{t \sim[0,1]}\left[f\left(L_{t}\right)\right]$. Thus $\hat{f}\left(x^{*}\right)$ can be written as a convex combination $\hat{f}\left(x^{*}\right)=\sum_{t} \alpha_{t} f\left(L^{(t)}\right)=\sum_{t} \alpha_{t} \hat{f}\left(I_{L^{(t)}}\right)$, where $\alpha_{t}>0$ and $\sum_{t} \alpha_{t}=1$. Thus all such $\hat{f}\left(I_{L^{(t)}}\right)$ are minimal, i.e. (10.1) has no "integrality gap".

We also modify our separation oracle to accommodate for this generalization as follows. Before doing so we need a definition which relates our BFS to the ring family formalism.

Definition 10.4.4. A permutation $\left(v_{1}, \ldots, v_{n}\right)$ of $V$ is said to be consistent with an arc $(i, j)$ if $j$ precedes $i$ in $\left(v_{1}, \ldots, v_{n}\right)$. Similarly, a BFS of the base polyhedron is consistent with $(i, j)$ if $j$ precedes $i$ in its defining permutation. $\left(v_{1}, \ldots, v_{n}\right)$ (or a BFS) is consistent with $A$ if it is consistent with every $(i, j) \in A$.

Readers may find it helpful to keep in mind the following picture which depicts the relative positions between $R(i), i, Q(i)$ in the defining permutation of $\boldsymbol{h}$ that is consistent with $A$ :

$$
\cdots \cdots(i) \backslash\{i\} \cdots \cdots(i) \backslash\{i\} \cdots \cdots \cdot
$$

In Theorem 10.2.4, given $\bar{x} \in[0,1]^{n}$ our separating hyperplane is constructed by sorting the entries of $\bar{x}$. This hyperplane is associated with some BFS $\boldsymbol{h}$ of the base polyhedron. As we shall see towards the end of the section, we would like $\boldsymbol{h}$ to be consistent with every arc $(i, j) \in A$.

This task is easy initially as $\bar{x}$ satisfies $x_{i} \leq x_{j}$ for $(i, j) \in A$ for the starting polytope of (10.1). If $x_{i}<x_{j}$, nothing special has to be done as $j$ must precede $i$ in the ordering. On the other hand, whenever $x_{i}=x_{j}$, we can always break ties by ranking $j$ ahead of $i$.

However, a technical issue arises due to the fact that our cutting plane algorithm may drop constraints from the current feasible region $P$. In other words, $\bar{x}$ may violate $x_{i} \geq 0, x_{j} \leq 1$ or $x_{i} \leq x_{j}$ if it is ever dropped. Fortunately this can be fixed by reintroducing the constraint. We summarize the modification needed in the pseudocode below and formally show that it fulfills our requirement.

```
Algorithm 25: Modified Separation Oracle
    Input: \(\bar{x} \in \mathbb{R}^{n}\) and the set of arcs \(A\)
    if \(\bar{x}_{i}<0\) for some \(i\) then
        Output: \(x_{i} \geq 0\)
    else if \(\bar{x}_{j}>1\) for some \(j\) then
        Output: \(x_{j} \leq 1\)
    else if \(\bar{x}_{i}>\bar{x}_{j}\) for some \((i, j) \in A\) then
        Output: \(x_{i} \leq x_{j}\)
    else
        Let \(i_{1}, \ldots, i_{n}\) be a permutation of \(V\) such that \(\bar{x}_{i_{1}} \geq \ldots \geq \bar{x}_{i_{n}}\) and for all \((i, j) \in A, j\)
        precedes \(i\) in \(i_{1}, \ldots, i_{n}\).
        Output: \(\boldsymbol{h}^{T} \boldsymbol{x} \leq \hat{f}(\bar{x})\), where \(\boldsymbol{h}\) is the BFS defined by the permutation \(i_{1}, \ldots, i_{n}\).
```

Lemma 10.4.5. Our modified separation oracle returns either some BFS $\boldsymbol{h}=0$ or a valid separating hyperplane, i.e.

1. $\bar{x}$ either lies on the separating hyperplane or is cut away by it.
2. Any minimizer of (10.1) is not cut away by the separating hyperplane.

Such a hyperplane can be computed with $n$ oracle calls to $f$ and in time $O\left(n \cdot E O+n^{2}\right)$.
Proof. If we get $x_{i} \geq 0, x_{j} \leq 1$ or $x_{i} \leq x_{j}$ (if loop or the first two else loops), then clearly $\bar{x}$ is cut away by it and any minimizer must of course satisfy $x_{i} \geq 0, x_{j} \leq 1$ and $x_{i} \leq x_{j}$ as they are the constraints in (10.1). This proves (1) and (2) for the case of getting $x_{i} \geq 0, x_{j} \leq 1$ or $x_{i} \leq x_{j}$.

Thus it remains to consider the case $\boldsymbol{h}^{T} \boldsymbol{x} \leq \hat{f}(\bar{x})$ (last else loop). First of all, $\bar{x}$ lies on it as $\hat{f}(\bar{x})=\boldsymbol{h}^{T} \bar{x}$. This proves (1). For (2), we have from Lemma 10.2.6 that $\boldsymbol{h}^{T} \boldsymbol{x} \leq \hat{f}(\boldsymbol{x})$. If $x^{*}$ is a minimizer of (10.1), we must then have $\boldsymbol{h}^{T} x^{*} \leq \hat{f}\left(x^{*}\right) \leq \hat{f}(\bar{x})$ as $\bar{x}$ is also feasible for (10.1).

Finally we note that the running time is self-evident.
We stress again that the main purpose of modifying our separation oracle is to ensure that any BFS $\boldsymbol{h}$ used to define a new separating hyperplane must be consistent with every $(i, j) \in A$.

### 10.4.2.2 Identifying New Valid Arcs

The reason for considering the ring family generalization of SFM is that our algorithms (and some previous algorithms too) work by adding new arcs to our digraph $D$. This operation yields a strongly polynomial algorithm since there are only $2 \cdot\binom{n}{2}$ possible arcs to add. Of course, a new arc $(i, j)$ is valid only if $i \in S_{\min } \Longrightarrow j \in S_{\min }$ for some minimizer $S_{\text {min }}$. Here we show how to identify such valid arcs by extracting information from certain nice elements of the base polyhedron.

This is guaranteed by the next four lemmas, which are stated in a way different from previous works e.g. our version is extended to the ring family setting. This is necessary as our algorithms require a more general formulation. We also give a new polyhedral proof, which is mildly simpler than the previous combinatorial proof. On the other hand, Lemma 10.4.10 is new and unique to our work. It is an important ingredient of our $\widetilde{O}\left(n^{3} \cdot \mathrm{EO}+n^{4}\right)$ time algorithm.

Recall that each BFS of the base polyhedron is defined by some permutation of the ground set elements.

First, we prove the following two lemmas which show that should we ever encounter a nondegenerate point in the base polytope with a coordinate of very large value, then we can immediately conclude that that coordinate must be or must not be in solution to SFM over the ring family.

Lemma 10.4.6. If $\boldsymbol{y} \in \mathcal{B}(f)$ is non-degenerate and satisfies $y_{i}>-(n-1) \min _{j} y_{j}$, then $i$ is not in any minimizer of $f$ (over the ring family $A$ ).

Proof. We proceed by contradiction and suppose that $S$ is a minimizer of $f$ that contains $i$. Now since $\boldsymbol{y}$ is non-degenerate we know that $\min _{j} y_{j} \leq 0$ and by the definition of $\boldsymbol{y}$ we have the following contradiction

$$
0<y_{i}+(n-1) \min _{j} y_{j} \leq \sum_{j \in S} y_{j}=\boldsymbol{y}(S) \leq f(S) \leq f(\emptyset)=0 .
$$

Lemma 10.4.7. If $\boldsymbol{y} \in \mathcal{B}(f)$ is non-degenerate and satisfies $y_{i}<-(n-1) \max _{j} y_{j}$, then $i$ is in every minimizer of $f$ (over the ring family $A$ ).

Proof. We proceed by contradiction and suppose that $S$ is a minimizer of $f$ that does not contain $i$. Now since $\boldsymbol{y}$ is non-degenerate we know that $\max _{j} y_{j} \geq 0$ and therefore

$$
\sum_{j \in[n]} y_{j}=y_{i}+\sum_{j \in S} y_{j}+\sum_{j \in V-(S+i)} y_{j}<-(n-1) \max _{j} y_{j}+\sum_{j \in S} y_{j}+(|V|-|S|-1) \max _{j} y_{j} \leq \sum_{j \in S} y_{j} .
$$

However by the definition of $\boldsymbol{y}$ we have

$$
\sum_{j \in S} y_{j}=\boldsymbol{y}(S) \leq f(S) \leq f(V)=\sum_{j \in[n]} y_{j} .
$$

Thus we have a contradiction and the result follows.

Now we are ready to present conditions under which a new valid arc can be added. We begin with a simple observation. Let upper $(i) \stackrel{\text { def }}{=} f(R(i))-f(R(i)-i)$ and lower $(i) \stackrel{\text { def }}{=} f(V \backslash Q(i)+i)-f(V \backslash Q(i))$. As the names suggest, they bound the value of $h_{i}$ for any BFS used.

Lemma 10.4.8. For any BFS $\boldsymbol{h}$ used to construct a separating hyperplane given by our modified separation oracle, we have lower $(i) \leq h_{i} \leq \operatorname{upper}(i)$.

Proof. Note that by Lemma $10.4 .5, \boldsymbol{h}$ is consistent with every $\left(j_{1}, j_{2}\right) \in A$ and hence $i$ must precede $Q(i)$ and be preceded by $R(i)$. Let $S$ be the set of elements preceding $i$ in the defining permutation of $\boldsymbol{h}$. Then $h_{i}=f(S+i)-f(S) \leq f(R(i))-f(R(i)-i)$ because of diminishing return and $R(i)-i \subseteq S$. The lower bound follows from the same argument as $Q(i)-i$ comes after $i$, and so $Q(i) \subseteq V \backslash S$.

In the following two lemmas, we show that if upper $(i)$ is ever sufficiently positive or lower $(i)$ is sufficiently negative, then we find a new arc.

While these lemmas may appear somewhat technical but actually has an intuitive interpretation. Suppose an element $p$ is in a minimizer $S_{\min }$ of $f$ over the ring family $D$. Then $R(p)$ must also be part of $S_{\text {min }}$. Now if $f(R(p))$ is very large relative to $f(R(p)-p)$, there should be some element $q \in S_{\min } \backslash R(p)$ compensating for the discrepancy. The lemma says that such an element $q$ can in fact be found efficiently.

Lemma 10.4.9 (new arc). Let $\boldsymbol{y}=\sum_{k} \lambda^{(k)} \boldsymbol{y}^{(k)}$ be a non-degenerate convex combination of $O(n)$ base polyhedron BFS's $\boldsymbol{y}^{(k)}$ which are consistent with every arc $(i, j) \in A$. If some element $p$ satisfies $\operatorname{upper}(p)>n^{4} \max y_{j}$, then we can find, using $O(n \cdot E O)$ oracle calls and $O\left(n^{2}\right)$ time, some $q \notin R(p)$ such that the arc $(p, q)$ is valid, i.e. if $p$ is in a minimizer, then so is $q$.

Proof. If max $y_{j}<0$ then we are immediately done by Lemma 10.2 .8 . We assume max $y_{j} \geq 0$ in the proof. For all $k$ let $\boldsymbol{y}^{\prime(k)}$ be the BFS obtained by taking the defining permutation of $\boldsymbol{y}^{(k)}$ and moving $R(p)$ to the front while preserving the relative ordering of $R(p)$ within each permutation). Furthermore, let $\boldsymbol{y}^{\prime} \stackrel{\text { def }}{=} \sum_{k} \lambda^{(k)} \boldsymbol{y}^{\prime(k)}$. Then since ${y^{\prime}}_{p}^{(k)}=f(R(p))-f(R(p)-p)=\operatorname{upper}(p)$ we have $\operatorname{upper}(p)=y_{p}^{\prime}=f(R(p))-f(R(p)-p)$. Moreover,

$$
\begin{equation*}
y_{j}^{\prime} \geq y_{j} \quad \forall j \in R(p) \text { and } y_{j}^{\prime} \leq y_{j} \quad \forall j \notin R(p) \tag{10.2}
\end{equation*}
$$

by diminishing marginal return.
Now, suppose $p$ is in a minimizer $S_{\text {min }}$. Then $R(p) \subseteq S_{\text {min }}$ by definition. We then define $f^{\prime}(S)=$ $f(S \cup R(p))$ for $S \subseteq V \backslash R(p)$. It can be checked readily that $f^{\prime}$ is submodular and $S_{\min } \backslash R(p)$ is a minimizer of $f^{\prime}$ (over the corresponding ring family). Note that now $\boldsymbol{y}_{V \backslash R(p)}^{\prime}$ (the restriction of $\boldsymbol{y}^{\prime}$ to $V \backslash R(p))$ is a convex combination of the BFS's of the base polyhedron $\mathcal{B}\left(f^{\prime}\right)$ of $f^{\prime}$. We shall show that $\boldsymbol{y}_{V \backslash R(p)}^{\prime}$ has the desired property in Lemma 10.4.7.

Note that $y^{\prime}(V \backslash R(p)+p) \leq y(V \backslash R(p)+p)$ since

$$
y^{\prime}(V \backslash R(p)+p)=y^{\prime}(V)-y^{\prime}(R(p)-p)=y(V)-y^{\prime}(R(p)-p) \leq y(V)-y(R(p)-p)=y(V \backslash R(p)+p) .
$$

But now since $\boldsymbol{y}$ is non-degenerate $\max _{j} y_{j} \geq 0$ and therefore

$$
\begin{align*}
y^{\prime}(V \backslash R(p)) & \leq y(V \backslash R(p)+p)-y_{p}^{\prime} \\
& =y(V \backslash R(p)+p)-(f(R(p))-f(R(p)-p))  \tag{10.3}\\
& \leq n \max y_{j}-(f(R(p))-f(R(p)-p)) \\
& <\left(n-n^{4}\right) \max y_{j}
\end{align*}
$$

Therefore by the Pigeonhole Principle some $q \notin R(p)$ must satisfy

$$
\begin{aligned}
y_{q}^{\prime} & <\left(\left(n-n^{4}\right) \max y_{j}\right) /(n-1) \\
& =-\left(n^{3}+n^{2}+n\right) \max y_{j} \\
& \leq-\left(n^{3}+n^{2}+n\right) \max _{j \notin R(p)} y_{j} \\
& \leq-\left(n^{3}+n^{2}+n\right) \max _{j \notin R(p)} y_{j}^{\prime} \quad \text { by }(10.2)
\end{aligned}
$$

By Lemma 10.4.7, this $q$ must be in any minimizer of $f^{\prime}$. In other words, whenever $p$ is in a minimizer of $f$, then so is $q$.

Note however that computing all $\boldsymbol{y}^{\prime}$ would take $O\left(n^{2}\right)$ oracle calls in the worst case as there are $O(n) \boldsymbol{y}^{\prime(k)}$ s. We use the following trick to identify some $q$ with $y_{q}^{\prime}<-(n-1) \max y_{j}$ using just $O(n)$ calls. The idea is that we actually only want to have sufficient decreases in $y^{\prime}(V \backslash R(p))$ which can be accomplished by having a large corresponding decrease in some $\boldsymbol{y}^{\prime(k)}$.

For each $k$, by the same argument above (see (10.3))

$$
\begin{equation*}
y^{\prime(k)}(V \backslash R(p))-y^{(k)}(V \backslash R(p)) \leq y_{p}^{(k)}-(f(R(p))-f(R(p)-p)) \tag{10.4}
\end{equation*}
$$

The "weighted decrease" $\lambda^{(k)}\left(y_{p}^{(k)}-(f(R(p))-f(R(p)-p))\right)$ for $\boldsymbol{y}^{\prime(k)}$ sum up to

$$
\sum \lambda^{(k)}\left(y_{p}^{(k)}-(f(R(p))-f(R(p)-p))\right)=y_{p}-(f(R(p))-f(R(p)-p))<\left(1-n^{4}\right) \max y_{j}
$$

Thus by the Pigeonhole Principle, some $l$ will have

$$
\lambda^{(l)}\left(y_{p}^{(l)}-(f(R(p))-f(R(p)-p))\right)<\left(\left(1-n^{4}\right) \max y_{j}\right) / O(n)<-n^{2} \max y_{j}
$$

For this $\boldsymbol{y}^{(l)}$ we compute $\boldsymbol{y}^{\prime(l)}$. We show that $\boldsymbol{y}^{\prime \prime}=\lambda^{(l)} \boldsymbol{y}^{\prime(l)}+\sum_{k \neq l} \lambda^{(k)} \boldsymbol{y}^{(k)}$ has the same property as $\boldsymbol{y}^{\prime}$ above.

$$
\begin{aligned}
y^{\prime \prime}(V \backslash R(p)) & =\lambda^{(l)} y^{\prime(l)}(V \backslash R(p))+\sum_{k \neq l} \lambda^{(k)} y^{(k)}(V \backslash R(p)) \\
& =y(V \backslash R(p))+\lambda^{(l)}\left(y^{\prime(l)}(V \backslash R(p))-y^{(l)}(V \backslash R(p))\right) \\
& \leq y(V \backslash R(p))+\lambda^{(l)}\left(y_{p}^{(l)}-(f(R(p))-f(R(p)-p))\right) \quad \text { by }(10.4) \\
& <(n-1) \max y_{j}-n^{2} \max y_{j} \\
& <\left(n-n^{2}\right) \max y_{j}
\end{aligned}
$$

Then some $q \in V \backslash R(p)$ must satisfy

$$
y_{q}^{\prime \prime}<\frac{n-n^{2}}{n-1} \max y_{j}=-n \max y_{j}
$$

That is, the $\operatorname{arc}(p, q)$ is valid. This takes $O(n)$ oracle calls as given $\boldsymbol{y}=\sum_{k} \lambda^{(k)} \boldsymbol{y}^{(k)}$, computing $\boldsymbol{y}^{\prime \prime}$ requires knowing only $f(R(p)), f(R(p)-p)$, and $\boldsymbol{y}^{(l)}$ which can be computed from $\boldsymbol{y}^{(l)}$ with $n$ oracle calls. The runtime is $O\left(n^{2}\right)$ which is needed for computing $\boldsymbol{y}^{\prime \prime}$.

Lemma 10.4.10. Let $\boldsymbol{y}=\sum_{k} \lambda^{(k)} \boldsymbol{y}^{(k)}$ be a non-degenerate convex combination of base polyhedron BFS $\boldsymbol{y}^{(k)}$ which is consistent with every arc $(i, j) \in A$. If lower $(p)<n^{4} \min y_{j}$, then we can find, using $O(n \cdot E O)$ oracle calls and $O\left(n^{2}\right)$ time, some $q \notin Q(p)$ such that the arc $(q, p)$ is valid, i.e. if $p$ is not in a minimizer, then $q$ is not either.

Proof. It is possible to follow the same recipe in the proof of Lemma 10.4.9 but using Lemma 10.4.6 instead of Lemma 10.4.7. Here we offer a proof which directly invokes Lemma 10.4.7 on a different submodular function.

Let $g$ be defined by $g(S) \stackrel{\text { def }}{=} f(V \backslash S)$ for any $S$, and $A_{g}$ be the set of arcs obtained by reversing the directions of the arcs of $A$. Consider the problem of minimizing $g$ over the ring family $A_{g}$. Using subscripts to avoid confusion with $f$ and $g$, e.g. $R_{g}(i)$ is the set of descendants of $i$ w.r.t. $A_{g}$, it is not hard to verify the following:

- $g$ is submodular
- $R_{g}(i)=Q_{f}(i)$
- $g\left(R_{g}(p)\right)-g\left(R_{g}(p)-p\right)=-\left(f\left(V \backslash Q_{f}(p)+p\right)-f\left(V \backslash Q_{f}(p)\right)\right)$
- $-\boldsymbol{y}^{(k)}$ is a BFS of $\mathcal{B}(g)$ if and only if $\boldsymbol{y}^{(k)}$ is a BFS of $\mathcal{B}(f)$
- $\max \left(-y_{j}\right)=-\min y_{j}$

By using the above correspondence and applying Lemma 10.4 .9 to $g$ and $A_{g}$, we can find, using $O(n)$ oracle calls and $O\left(n^{2}\right)$ time, some $q \notin R_{g}(p)=Q(p)$ such that the arc $(p, q)$ is valid for $g$ and $A_{g}$. In other words, the reverse $(q, p)$ will be valid for $f$ and $A$.

These lemmas lay the foundation of our algorithm. They suggests that if the positive entries of a point in the base polyhedron are small relative to some upper $(p)=f(R(p))-f(R(p)-p)$, a new arc $(p, q)$ can be added to $A$. This can be seen as a robust version of Lemma 10.2.8.

Finally, we end the section with a technical lemma that will be used crucially for both of our algorithms. The importance of it would become obvious when it is invoked in our analyses.

Lemma 10.4.11. Let $\boldsymbol{h}^{\prime \prime}$ denote a convex combination of two vectors $\boldsymbol{h}$ and $\boldsymbol{h}^{\prime}$ in the base polyhedron, i.e. $\boldsymbol{h}^{\prime \prime}=\lambda \boldsymbol{h}+(1-\lambda) \boldsymbol{h}^{\prime}$ for some $\lambda \in[0,1]$. Further suppose that

$$
\left\|\boldsymbol{h}^{\prime \prime}\right\|_{2} \leq \alpha \min \left\{\lambda\|\boldsymbol{h}\|_{2},(1-\lambda)\left\|\boldsymbol{h}^{\prime}\right\|_{2}\right\}
$$

for some $\alpha \leq \frac{1}{2 \sqrt{n}}$. Then for $p=\arg \max _{j}\left(\max \left\{\lambda\left|h_{j}\right|,(1-\lambda)\left|h_{j}^{\prime}\right|\right\}\right)$ we have

$$
\text { lower }(p) \leq-\frac{1}{2 \alpha \sqrt{n}} \cdot\left\|\boldsymbol{h}^{\prime \prime}\right\|_{\infty} \quad \text { and } \quad \operatorname{upper}(p) \geq \frac{1}{2 \alpha \sqrt{n}} \cdot\left\|\boldsymbol{h}^{\prime \prime}\right\|_{\infty}
$$

Proof. Suppose without loss of generality that $\lambda\left|h_{p}\right| \geq(1-\lambda)\left|h_{p}^{\prime}\right|$. Then by assumptions we have

$$
\left\|\boldsymbol{h}^{\prime \prime}\right\|_{\infty} \leq\left\|\boldsymbol{h}^{\prime \prime}\right\|_{2} \leq \alpha \cdot \min \left\{\lambda\|\boldsymbol{h}\|_{2},(1-\lambda)\left\|\boldsymbol{h}^{\prime}\right\|_{2}\right\} \leq \alpha \sqrt{n}\left|\lambda h_{p}\right|
$$

However, since $\alpha \leq \frac{1}{2 \sqrt{n}}$ we see that

$$
\left|\lambda h_{p}+(1-\lambda) h_{p}^{\prime}\right| \leq\left\|h^{\prime \prime}\right\|_{\infty} \leq \alpha \sqrt{n}\left|\lambda h_{p}\right| \leq \frac{1}{2}\left|\lambda h_{p}\right|
$$

Consequently, $\lambda h_{p}$ and $(1-\lambda) h_{p}^{\prime}$ have opposite signs and $\left|(1-\lambda) h_{p}^{\prime}\right| \geq \frac{1}{2}\left|\lambda h_{p}^{\prime}\right|$. We then have,

$$
\operatorname{lower}(p) \leq \min \left\{h_{p}, h_{p}^{\prime}\right\} \leq \min \left\{\lambda h_{p},(1-\lambda) h_{p}^{\prime}\right\} \leq-\frac{1}{2}\left|\lambda h_{p}\right| \leq-\frac{1}{2 \alpha \sqrt{n}}\left\|h^{\prime \prime}\right\|_{\infty}
$$

and

$$
\operatorname{upper}(p) \geq \max \left\{h_{p}, h_{p}^{\prime}\right\} \geq \max \left\{\lambda h_{p},(1-\lambda) h_{p}^{\prime}\right\} \geq \frac{1}{2}\left|\lambda h_{p}\right| \geq \frac{1}{2 \alpha \sqrt{n}}\left\|h^{\prime \prime}\right\|_{\infty}
$$

## ■ 10.4.3 $\widetilde{O}\left(n^{4} \cdot \mathbf{E O}+n^{5}\right)$ Time Algorithm

Here we present a $\widetilde{O}\left(n^{4} \cdot \mathrm{EO}+n^{5}\right)$ time, i.e. strongly polynomial time algorithm, for SFM. We build upon the algorithm achieved in the section to achieve a faster running time in Section 10.4.4.

Our new algorithm combines the existing tools for SFM developed over the last decade with our cutting plane method. While there are certain similarities with previous algorithms (especially [118, $125,121]$ ), our approach significantly departs from all the old approaches in one important aspect.

All of the previous algorithms actively maintain a point in the base polyhedron and represent it as a convex combination of BFS's. At each step, a new BFS may enter the convex combination and an old BFS may exit. Our algorithm, on the other hand, maintains only a collection of BFS's (corresponding to our separating hyperplanes), rather than an explicit convex combination. A "good" convex combination is computed from the collection of BFS's only after running Cutting Plane for enough iterations. We believe that this crucial difference is the fundamental reason which offers the speedup. This is achieved by the Cutting Plane method which considers the geometry of the collection of BFS's. On the other hand, considering only a convex combination of BFS's effectively narrows our sight to only one point in the base polyhedron.

## Overview

Now we are ready to describe our strongly polynomial time algorithm. Similar to the weakly polynomial algorithm, we first run our cutting plane for enough iterations on the initial feasible region $\left\{\boldsymbol{x} \in[0,1]^{n}: x_{i} \leq x_{j} \forall(i, j) \in A\right\}$, after which a pair of approximately parallel supporting hyperplanes $F_{1}, F_{2}$ of width $1 / n^{\Theta(1)}$ can be found. Our strategy is to write $F_{1}$ and $F_{2}$ as a nonnegative combination of the facets of remaining feasible region $P$. This combination is made up of newly added separating hyperplanes as well as the inequalities $x_{i} \geq 0, x_{j} \leq 1$ and $x_{i} \leq x_{j}$. We then argue that one of the following updates can be done:

- Collapsing: $x_{i}=0, x_{j}=1$ or $x_{i}=x_{j}$
- Adding a new arc $(i, j): x_{i} \leq x_{j}$ for some $(i, j) \notin A$

The former case is easy to handle by elimination or contraction. If $x_{i}=0$, we simply eliminate $i$ from the ground set $V$; and if $x_{i}=1$, we redefine $f$ so that $f(S)=f(S+i)$ for any $S \subseteq V-i$. $x_{i}=x_{j}$ can be handled in a similar fashion. In the latter case, we simply add the $\operatorname{arc}(i, j)$ to $A$. We then repeat the same procedure on the new problem.

Roughly speaking, our strongly polynomial time guarantee follows as eliminations and contractions can happen at most $n$ times and at most $2 \cdot\binom{n}{2}$ new arcs can be added. While the whole picture is simple, numerous technical details come into play in the execution. We advise readers to keep this overview in mind when reading the subsequent sections.

## Algorithm

Our algorithm is summarized below. Again, we remark that our algorithm simply uses Theorem 10.4.12 regarding our cutting plane and is agnostic as to how the cutting plane works, thus it could be replaced with other methods, albeit at the expense of slower runtime.

1. Run cutting plane on (10.1) (Theorem 10.4.12 with $\tau=\Theta(1)$ ) using our modified separation oracle (Section 10.4.2.1).
2. Identify a pair of "narrow" approximately parallel supporting hyperplanes or get some BFS $\boldsymbol{h}=0$ (in which case both $\emptyset$ and $V$ are minimizers).
3. Deduce from the hyperplanes some new constraint of the forms $x_{i}=0, x_{j}=1, x_{i}=x_{j}$ or $x_{i} \leq x_{j}$ (Section 10.4.3.2).
4. Consolidate $A$ and $f$ (Section 10.4.3.1).
5. Repeat by running our cutting plane method on (10.1) with updated $A$ and $f$. (Note that Any previously found separating hyperplanes are discarded.)

We call step (1) a phase of cutting plane. The minimizer can be constructed by unraveling the recursion.

### 10.4.3.1 Consolidating $A$ and $f$

Here we detail how the set of valid $\operatorname{arcs} A$ and submodular function $f$ should be updated once we deduce new information $x_{i}=0, x_{i}=1, x_{i}=x_{j}$ or $x_{i} \leq x_{j}$. Recall that $R(i)$ and $Q(i)$ are the sets of descendants and ancestors of $i$ respectively (including $i$ itself). The changes below are somewhat self-evident, and are actually used in some of the previous algorithms so we only sketch how they are done without a detailed justification.

Changes to the digraph representation $D$ of our ring family include:

- $x_{i}=0$ : remove $Q(i)$ from the ground set and all the arcs incident to $Q(i)$
- $x_{i}=1$ : remove $R(i)$ from the ground set and all the arcs incident to $R(i)$
- $x_{i}=x_{j}$ : contract $i$ and $j$ in $D$ and remove any duplicate arcs
- $x_{i} \leq x_{j}$ : insert the arc $(i, j)$ to $A$
- For the last two cases, we also contract the vertices on a directed cycle of $A$ until there is no more. Remove any duplicate arcs.

Here we can contract any cycle $\left(i_{1}, \ldots, i_{k}\right)$ because the inequalities $x_{i_{1}} \leq x_{i_{2}}, \ldots, x_{i_{k-1}} \leq x_{i_{k}}, x_{i_{k}} \leq x_{i_{1}}$ imply $x_{i_{1}}=\ldots=x_{i_{k}}$.

Changes to $f$ :

- $x_{i}=0$ : replace $f$ by $f^{\prime}: 2^{V \backslash Q(i)} \longrightarrow \mathbb{R}, f^{\prime}(S)=f(S)$ for $S \subseteq V \backslash Q(i)$
- $x_{i}=1$ : replace $f$ by $f^{\prime}: 2^{V \backslash R(i)} \longrightarrow \mathbb{R}, f^{\prime}(S)=f(S \cup R(i))$ for $S \subseteq V \backslash R(i)$
- $x_{i}=x_{j}$ : see below
- $x_{i} \leq x_{j}$ : no changes to $f$ needed if it does not create a cycle in $A$; otherwise see below
- Contraction of $C=\left\{i_{1}, \ldots, i_{k}\right\}$ : replace $f$ by $f^{\prime}: 2^{V \backslash C+l} \longrightarrow \mathbb{R}, f^{\prime}(S)=f(S)$ for $S \subseteq V \backslash C$ and $f^{\prime}(S)=f((S-l) \cup C)$ for $S \ni l$

Strictly speaking, these changes are in fact not needed as they will automatically be taken care of by our cutting plane method. Nevertheless, performing them lends a more natural formulation of the algorithm and simplifies its description.

### 10.4.3.2 Deducing New Constraints $x_{i}=0, x_{j}=1, x_{i}=x_{j}$ or $x_{i} \leq x_{j}$

Here we show how to deduce new constraints through the result of our cutting plane method. This is the most important ingredient of our algorithm. As mentioned before, similar arguments were used first by IFF [121] and later in [118, 125]. There are however two important differences for our method:

- We maintain a collection of BFS's rather a convex combination; a convex combination is computed and needed only after each phase of cutting plane.
- As a result, our results are proved mostly geometrically whereas the previous ones were proved mostly combinatorially.

Our ability to deduce such information hinges on the power of the cutting plane method in Chapter 8. We re-state our main result Theorem 8.3.26 in the language of SFM. Note that Theorem 10.4.12 is formulated in a fairly general manner in order to accommodate for the next section. Readers may wish to think $\tau=\Theta(1)$ for now.

Theorem 10.4.12 (Theorem 8.3.26 restated for SFM). For any $\tau \geq 100$, applying our cutting plane method, Theorem 10.4.12, to (10.1) with our modified separation oracle (or its variant in Section 10.4.4) with high probability in $n$ either

1. Finds a degenerate $B F S \boldsymbol{h} \geq \overrightarrow{0}$ or $\boldsymbol{h} \leq \mathbf{0}$.
2. Finds a polytope $P$ consisting of $O(n)$ constraints which are our separating hyperplanes or the constraints in (10.1). Moreover, $P$ satisfies the following inequalities

$$
\boldsymbol{c}^{T} \boldsymbol{x} \leq M \quad \text { and } \quad \boldsymbol{c}^{T} \boldsymbol{x} \leq M^{\prime}
$$

both of which are nonnegative combinations of the constraints of $P$, where

$$
\left\|\boldsymbol{c}+\boldsymbol{c}^{\prime}\right\|_{2} \leq \min \left\{\|\boldsymbol{c}\|_{2},\left\|\boldsymbol{c}^{\prime}\right\|_{2}\right\} / n^{\Theta(\tau)} \quad \text { and } \quad\left|M+M^{\prime}\right| \leq \min \left\{\|\boldsymbol{c}\|_{2},\left\|\boldsymbol{c}^{\prime}\right\|_{2}\right\} / n^{\Theta(\tau)} .
$$

Furthermore, the algorithm runs in expected time $O\left(n^{2} \tau \log n \cdot E O+n^{3} \tau^{O(1)} \log { }^{O(1)} n\right)$.
Proof. In applying Theorem 10.4 .12 we let $K$ be the set of minimizers of $f$ over the ring family and the box is the hypercube with $R=1$. We run cutting plane with our modified separation oracle (Lemma 10.4.5). The initial polytope $P^{(0)}$ can be chosen to be, say, the hypercube. If some separating hyperplane is degenerate, then we have the desired result (and know that either $\emptyset$ or $V$ is optimal). Otherwise let $P$ be the current feasible region. Note that $P \neq \emptyset$, because our minimizers of $\hat{f}$ are all in $P^{(0)}$ and $P^{(k)}$ as they are never cut away by the separating hyperplanes.

Let $\mathcal{S}$ be the collection of inequalities (10.1) as well as the separating hyperplanes $\boldsymbol{h}^{T} \boldsymbol{x} \leq \hat{f}\left(\bar{x}_{h}\right)=$ $\boldsymbol{h}^{T} \bar{x}_{h}$ used. By Theorem 8.3.26, all of our minimizers will be contained in $P$, consisting of $O(n)$ constraints $\boldsymbol{A} \boldsymbol{x} \geq \boldsymbol{b}$. Each such constraint $\boldsymbol{a}_{i}^{T} \boldsymbol{x} \geq b_{i}$ is a scaling and shifting of some inequality $\boldsymbol{p}_{i}^{T} \boldsymbol{x} \geq q_{i}$ in $\mathcal{S}$, i.e. $\boldsymbol{a}_{i}=\boldsymbol{p}_{i} /\left\|\boldsymbol{p}_{i}\right\|_{2}$ and $b_{i} \leq q_{i} /\left\|\boldsymbol{p}_{i}\right\|_{2}$.

By taking $\varepsilon=1 / n^{\Theta(\tau)}$ with sufficiently large constant in $\Theta$, our theorem certifies that $P$ has a narrow width by $\boldsymbol{a}_{1}$, some nonnegative combination $\sum_{i=2}^{O(n)} t_{i} \boldsymbol{a}_{i}$ and point $\boldsymbol{x}_{o} \in P$ with $\left\|\boldsymbol{x}_{o}\right\|_{\infty} \leq$ $3 \sqrt{n} R=3 \sqrt{n}$ satisying the following:

$$
\begin{aligned}
& \left\|\boldsymbol{a}_{1}+\sum_{i=2}^{O(n)} t_{i} \boldsymbol{a}_{i}\right\|_{2} \leq 1 / n^{\Theta(\tau)} \\
& 0 \leq \boldsymbol{a}_{1}^{T} \boldsymbol{x}_{o}-\boldsymbol{b}_{1} \leq 1 / n^{\Theta(\tau)}
\end{aligned}
$$

$$
0 \leq\left(\sum_{i=2}^{O(n)} t_{i} a_{i}\right)^{T} \boldsymbol{x}_{o}-\sum_{i=2}^{O(n)} t_{i} b_{i} \leq 1 / n^{\Theta(\tau)}
$$

We convert these inequalities to $\boldsymbol{p}$ and $q$. Let $t_{i}^{\prime} \stackrel{\text { def }}{=} t_{i} \cdot\left\|\boldsymbol{p}_{1}\right\|_{2} /\left\|\boldsymbol{p}_{i}\right\|_{2} \geq 0$.

$$
\begin{gathered}
\left\|\boldsymbol{p}_{1}+\sum_{i=2}^{O(n)} t_{i}^{\prime} \boldsymbol{p}_{i}\right\|_{2} \leq\left\|\boldsymbol{p}_{1}\right\|_{2} / n^{\Theta(\tau)} \\
0 \leq \boldsymbol{p}_{1}^{T} \boldsymbol{x}_{o}-q_{1} \leq\left\|\boldsymbol{p}_{1}\right\|_{2} / n^{\Theta(\tau)} \\
0 \leq\left(\sum_{i=2}^{O(n)} t_{i}^{\prime} \boldsymbol{p}_{i}\right)^{T} \boldsymbol{x}_{o}-\sum_{i=2}^{O(n)} t_{i}^{\prime} q_{i} \leq\left\|\boldsymbol{p}_{1}\right\|_{2} / n^{\Theta(\tau)}
\end{gathered}
$$

We claim that ${ }^{1} \boldsymbol{c}=-\boldsymbol{p}_{1}, M=-q_{1}, \boldsymbol{c}^{\prime}=-\sum_{i=2}^{O(n)} t_{i}^{\prime} \boldsymbol{p}_{i}, M^{\prime}=-\sum_{i=2}^{O(n)} t_{i}^{\prime} q_{i}$ satisfy our requirement.
We first show that $\left\|\boldsymbol{c}+\boldsymbol{c}^{\prime}\right\|_{2} \leq \min \left\{\|\boldsymbol{c}\|_{2},\left\|\boldsymbol{c}^{\prime}\right\|_{2}\right\} / n^{\Theta(\tau)}$. We have $\left\|\boldsymbol{c}+\boldsymbol{c}^{\prime}\right\|_{2} \leq\|\boldsymbol{c}\|_{2} / n^{\Theta(\tau)}$ from the first inequality. If $\|\boldsymbol{c}\|_{2} \leq\left\|\boldsymbol{c}^{\prime}\right\|_{2}$ we are done. Otherwise, by triangle inequality

$$
\left\|\boldsymbol{c}^{\prime}\right\|_{2}-\|\boldsymbol{c}\|_{2} \leq\left\|\boldsymbol{c}+\boldsymbol{c}^{\prime}\right\|_{2} \leq\|\boldsymbol{c}\|_{2} / n^{\Theta(\tau)} \Longrightarrow 2\|\boldsymbol{c}\|_{2} \geq\left\|\boldsymbol{c}^{\prime}\right\|_{2}
$$

and hence $\left\|\boldsymbol{c}+\boldsymbol{c}^{\prime}\right\|_{2} \leq\|\boldsymbol{c}\|_{2} / n^{\Theta(\tau)} \leq\left\|\boldsymbol{c}^{\prime}\right\|_{2} / 2 n^{\Theta(\tau)}=\left\|\boldsymbol{c}^{\prime}\right\|_{2} / n^{\Theta(\tau)}$.
We also need to prove $\left|M+M^{\prime}\right| \leq \min \left\{\|\boldsymbol{c}\|_{2},\left\|\boldsymbol{c}^{\prime}\right\|_{2}\right\} / n^{\Theta(\tau)}$. Summing the second and third inequalities,

$$
-\|\boldsymbol{c}\|_{2} / n^{\Theta(\tau)} \leq\left(\boldsymbol{c}+\boldsymbol{c}^{\prime}\right)^{T} \boldsymbol{x}_{o}-\left(M+M^{\prime}\right) \leq 0
$$

Recall that we have $\left\|\boldsymbol{x}_{o}\right\|_{\infty} \leq 3 \sqrt{n}$. Then

$$
\begin{aligned}
\left|M+M^{\prime}\right| & \leq\left|\left(\boldsymbol{c}+\boldsymbol{c}^{\prime}\right)^{T} \boldsymbol{x}_{o}-\left(M+M^{\prime}\right)\right|+\left|\left(\boldsymbol{c}+\boldsymbol{c}^{\prime}\right)^{T} \boldsymbol{x}_{o}\right| \\
& \leq\|\boldsymbol{c}\|_{2} / n^{\Theta(\tau)}+3 \sqrt{n}\left\|\boldsymbol{c}+\boldsymbol{c}^{\prime}\right\|_{2} \\
& \leq\|\boldsymbol{c}\|_{2} / n^{\Theta(\tau)}+3 \sqrt{n}\|\boldsymbol{c}\|_{2} / n^{\Theta(\tau)} \\
& =\|\boldsymbol{c}\|_{2} / n^{\Theta(\tau)}
\end{aligned}
$$

as desired. Our result then follows as we proved $2\left\|\boldsymbol{c}^{\prime}\right\|_{2} \geq\|\boldsymbol{c}\|_{2}$.
Finally, we have the desired runtime as our modified separation oracle runs in time $O(n \cdot \mathrm{EO}+$ $\left.n^{2} \log ^{O(1)} n\right)$.

Informally, the theorem above simply states that after $O(n \tau \log n)$ iterations of cutting plane, the remaining feasible region $P$ can be sandwiched between two approximately parallel supporting hyperplanes of width $1 / n^{O(\tau)}$. A good intuition to keep in mind is that every $O(n)$ iterations of cutting plane reduces the minimum width by a constant factor.

Remark 10.4.13. As shown in the proof of Theorem 10.4.12, one of the two approximately parallel hyperplanes can actually be chosen to be a constraint of our feasible region $P$. However we do not exploit this property as it does not seem to help us and would break the notational symmetry in $\boldsymbol{c}$ and $\boldsymbol{c}^{\prime}$.

## Setup

[^19]In each phase, we run cutting plane using Theorem 10.4.12 with $\tau=\Theta(1)$. If some separating hyperplane used is degenerate, we have found the minimizer by Lemma 10.2.8.

Now assume none of the separating hyperplanes is degenerate. By Theorem 10.4.12, $P$ is sandwiched by a pair of approximately parallel supporting hyperplanes $F, F^{\prime}$ which are of width $1 / 10 n^{10}$ apart. The width here can actually be $1 / n^{c}$ for any constant $c$ by taking a sufficiently large constant in Theta.

Here, we show how to deduce from $F$ and $F^{\prime}$ some $x_{i}=0, x_{j}=1, x_{i}=x_{j}$, or $x_{i} \leq x_{j}$ constraint on the minimizers of $f$ over the ring family. Let

$$
\boldsymbol{c}^{T} \boldsymbol{x}=\sum c_{i} x_{i} \leq M \quad \text { and } \quad \boldsymbol{c}^{\prime T} \boldsymbol{x}=\sum c_{i}^{\prime} x_{i} \leq M^{\prime}
$$

be the inequality for $F$ and $F^{\prime}$ such that

$$
\left|M+M^{\prime}\right|,\left\|\boldsymbol{c}+\boldsymbol{c}^{\prime}\right\|_{2} \leq \text { gap }, \quad \text { where gap } \stackrel{\text { def }}{=} \frac{1}{10 n^{10}} \min \left\{\|\boldsymbol{c}\|_{2},\left\|\boldsymbol{c}^{\prime}\right\|_{2}\right\} .
$$

By the same theorem we can write $\boldsymbol{c}^{T} \boldsymbol{x} \leq M$ as a nonnegative combination of the constraints for $P$. Recall that the constraints for $P$ take on four different forms: (1) $-x_{i} \leq 0$; (2) $x_{j} \leq 1$; (3) $-\left(x_{j}-x_{i}\right) \leq 0$; (4) $\boldsymbol{h}^{T} \boldsymbol{x}=\sum h_{i} x_{i} \leq \hat{f}\left(\bar{x}_{h}\right)$. Here the first three types are present initially whereas the last type is the separating hyperplane added. As alleged previously, the coefficient vector $\boldsymbol{h}$ corresponds to a BFS of the base polyhedron for $f$. Our analysis crucially exploits this property.

Thus suppose $\boldsymbol{c}^{T} \boldsymbol{x}=\sum_{i} c_{i} x_{i} \leq M$ is a nonnegative combination of our constraints with weights $\alpha_{i}, \beta_{j}, \gamma_{i j}, \lambda_{h} \geq 0$. The number of (positive) $\alpha_{i}, \beta_{j}, \gamma_{i j}, \lambda_{h}$ is at most $O(n)$. Here we denote separating hyperplanes by $\boldsymbol{h}^{T} \boldsymbol{x} \leq \hat{f}\left(\bar{x}_{h}\right)$. Let $H$ be the set of BFS's used to construct separating hyperplanes.

$$
\begin{equation*}
\boldsymbol{c}^{T} \boldsymbol{x}=-\sum_{i} \alpha_{i} x_{i}+\sum_{j} \beta_{j} x_{j}+\sum_{(i, j) \in A} \gamma_{i j}\left(x_{i}-x_{j}\right)+\sum_{h \in H} \lambda_{h} \boldsymbol{h}^{T} \boldsymbol{x} \quad \text { and } \quad M=\sum_{j} \beta_{j}+\sum_{h \in H} \lambda_{h} \hat{f}\left(\bar{x}_{h}\right) . \tag{10.5}
\end{equation*}
$$

Similarly, we write the inequality for $F^{\prime}$ as a nonnegative combination of the constraints for $P$ and the number of (positive) $\alpha_{i}^{\prime}, \beta_{j}^{\prime}, \gamma_{i j}^{\prime}, \lambda_{h}^{\prime}$ is $O(n)$ :
$\boldsymbol{c}^{\prime T} \boldsymbol{x}=-\sum \alpha_{i}^{\prime} x_{i}+\sum \beta_{j}^{\prime} x_{j}+\sum_{(i, j) \in A} \gamma_{i j}^{\prime}\left(x_{i}-x_{j}\right)+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}^{T} \boldsymbol{x} \quad$ and $\quad M^{\prime}=\sum \beta_{j}^{\prime}+\sum_{h \in H} \lambda_{h}^{\prime} \hat{f}\left(\bar{x}_{h}\right)$.
We also scale $\boldsymbol{c}, \boldsymbol{c}^{\prime}, \alpha, \alpha^{\prime}, \beta, \beta^{\prime}, \gamma, \gamma^{\prime}, \lambda, \lambda^{\prime}$ so that

$$
\sum_{h \in H}\left(\lambda_{h}+\lambda_{h}^{\prime}\right)=1
$$

as this does not change any of our preceding inequalities regarding $F$ and $F^{\prime}$.
Now that $F, F^{\prime}$ have been written as combinations of our constraints, we have gathered the necessary ingredients to derive our new arc. We first give a geometric intuition why we would expect to be able to derive a new constraint. Consider the nonnegative combination making up $F$. We think of the coefficient $\beta_{j}$ as the contribution of $x_{j} \leq 1$ to $F$. Now if $\beta_{j}$ is very large, $F$ is "very parallel" to $x_{j} \leq 1$ and consequently $F^{\prime}$ would miss $x_{j}=0$ as the gap between $F$ and $F^{\prime}$ is small. $P$ would then miss $x_{j}=0$ too as it is sandwiched between $F$ and $F^{\prime}$. Similarly, a large $\alpha_{i}$ and a large $\gamma_{i j}$ would respectively imply that $x_{i}=1$ and ( $x_{i}=0, x_{j}=1$ ) would be missed. The same argument works for $F^{\prime}$ as well.

But on the other hand, if the contributions from $x_{i} \geq 0, x_{j} \leq 1, x_{i} \leq x_{j}$ to both $F$ and $F^{\prime}$ are small, then the supporting hyperplanes $\boldsymbol{c}^{T} \boldsymbol{x} \leq \ldots$ and $\boldsymbol{c}^{\prime T} \boldsymbol{x} \leq \ldots$ would be mostly made up of
separating hyperplanes $\boldsymbol{h}^{T} \boldsymbol{x} \leq \hat{f}\left(\bar{x}_{h}\right)$. By summing up these separating hyperplanes (whose coefficients form BFS's), we would then get a point in the base polyhedron which is very close to the origin 0 . Moreover, by Lemma 10.4.11 and Lemma 10.4 .9 we should then be able to deduce some interesting information about the minimizer of $f$ over $D$.

The rest of this section is devoted to realizing the vision sketched above. We stress that while the algebraic manipulations may be long, they are simply the execution of this elegant geometric picture.

Now, consider the following weighted sum of $\boldsymbol{h}^{T} \boldsymbol{x} \leq \hat{f}\left(\bar{x}_{h}\right)$ :

$$
\left(\sum_{h \in H} \lambda_{h} \boldsymbol{h}^{T}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}^{T}\right) \boldsymbol{x}=\sum_{h \in H} \lambda_{h} \boldsymbol{h}^{T} \boldsymbol{x}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}^{T} \boldsymbol{x} \leq \sum_{h \in H} \lambda_{h} \hat{f}\left(\bar{x}_{h}\right)+\sum_{h \in H} \lambda_{h}^{\prime} \hat{f}\left(\bar{x}_{h}\right) .
$$

Observe that $\sum_{h \in H} \lambda_{h} \boldsymbol{h}^{T}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}^{T}$ is in the base polyhedron since it is a convex combination of BFS $\boldsymbol{h}$. Furthermore, using (10.5) and (10.6) this can also be written as

$$
\begin{align*}
\left(\sum_{h \in H} \lambda_{h} \boldsymbol{h}^{T}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}^{T}\right) \boldsymbol{x}=\left(\boldsymbol{c}^{T} \boldsymbol{x}\right. & \left.+\sum \alpha_{i} x_{i}-\sum \beta_{j} x_{j}+\sum_{(i, j) \in A} \gamma_{i j}\left(x_{j}-x_{i}\right)\right) \\
& +\left(\boldsymbol{c}^{\prime T} \boldsymbol{x}+\sum \alpha_{i}^{\prime} x_{i}-\sum \beta_{j}^{\prime} x_{j}+\sum_{(i, j) \in A} \gamma_{i j}^{\prime}\left(x_{j}-x_{i}\right)\right) \tag{10.7}
\end{align*}
$$

and

$$
\begin{aligned}
\sum_{h \in H} \lambda_{h} \hat{f}\left(\bar{x}_{h}\right)+\sum_{h \in H} \lambda_{h}^{\prime} \hat{f}\left(\bar{x}_{h}\right) & =\left(M-\sum \beta_{j}\right)+\left(M^{\prime}-\sum \beta_{j}^{\prime}\right) \\
& =\left(M+M^{\prime}\right)-\sum \beta_{j}-\sum \beta_{j}^{\prime}
\end{aligned}
$$

Furthermore, we can bound $\boldsymbol{c}^{T} \boldsymbol{x}+\boldsymbol{c}^{\prime T} \boldsymbol{x}$ by $\boldsymbol{c}^{T} \boldsymbol{x}+\boldsymbol{c}^{T} \boldsymbol{x} \geq-\left\|\boldsymbol{c}+\boldsymbol{c}^{\prime}\right\|_{1} \geq-\sqrt{n}\left\|\boldsymbol{c}+\boldsymbol{c}^{\prime}\right\|_{2} \geq-\sqrt{n}$ gap as $\boldsymbol{x} \leq 1$. Since $M+M^{\prime} \leq$ gap, we obtain

$$
\begin{gathered}
L H S \stackrel{\text { def }}{=} \sum \alpha_{i} x_{i}+\sum \alpha_{i}^{\prime} x_{i}-\sum \beta_{j} x_{j}-\sum \beta_{j}^{\prime} x_{j}+\sum_{(i, j) \in A} \gamma_{i j}\left(x_{j}-x_{i}\right)+\sum_{(i, j) \in A} \gamma_{i j}^{\prime}\left(x_{j}-x_{i}\right) \\
\leq 2 \sqrt{n} \text { gap }-\sum \beta_{j}-\sum \beta_{j}^{\prime}
\end{gathered}
$$

Geometrically, the next lemma states that if the contribution from, say $x_{i} \geq 0$, to $F$ is too large, then $F^{\prime}$ would be forced to miss $x_{i}=1$ because they are close to one another.

Lemma 10.4.14. Suppose $\boldsymbol{x}$ satisfies (10.1), LHS $\leq 2 \sqrt{n}$ gap $-\sum \beta_{j}-\sum \beta_{j}^{\prime}$ and $\alpha_{i}, \beta_{j}, \gamma_{i j}, \alpha_{i}^{\prime}, \beta_{j}^{\prime}, \gamma_{i j}^{\prime} \geq$ 0.

1. If $\alpha_{i}>2 \sqrt{n}$ gap or $\alpha_{i}^{\prime}>2 \sqrt{n}$ gap, then $x_{i}<1$.
2. If $\beta_{j}>2 \sqrt{n}$ gap or $\beta_{j}^{\prime}>2 \sqrt{n}$ gap, then $x_{j}>0$.
3. If $\gamma_{i j}>2 \sqrt{n}$ gap or $\gamma_{i j}^{\prime}>2 \sqrt{n}$ gap, then $0 \leq x_{j}-x_{i}<1$.

Proof. We only prove it for $\alpha_{i}, \beta_{j}, \gamma_{i j}$ as the other case follows by symmetry.
Using $0 \leq x \leq 1$ and $x_{i} \leq x_{j}$ for $(i, j) \in A$, we have LHS $\geq \alpha_{i} x_{i}-\sum \beta_{j}-\sum \beta_{j}^{\prime}$. Hence $\alpha_{i} x_{i} \leq 2 \sqrt{n}$ gap and we get $x_{i}<1$ if $\alpha_{i}>2 \sqrt{n}$ gap.

Similarly, LHS $\geq-\beta_{k} x_{k}-\sum_{j \neq k} \beta_{j}-\sum \beta_{j}^{\prime}$ which gives $-\beta_{k} x_{k} \leq 2 \sqrt{n}$ gap $-\beta_{k}$. Then $x_{k}>0$ if $\beta_{k}>2 \sqrt{n}$ gap.

Finally, LHS $\geq \gamma_{i j}\left(x_{j}-x_{i}\right)-\sum \beta_{j}-\sum \beta_{j}^{\prime}$ which gives $\gamma_{i j}\left(x_{j}-x_{i}\right) \leq 2 \sqrt{n}$ gap. Then $x_{j}-x_{i}<1$ if $\gamma_{i j}>2 \sqrt{n}$ gap. We have $x_{i} \leq x_{j}$ since $(i, j) \in A$.

So if either condition of Lemma 10.4.14 holds, we can set $x_{i}=0$ or $x_{j}=1$ or $x_{i}=x_{j}$ since our problem (10.1) has an integral minimizer and any minimizer of $\hat{f}$ is never cut away by Lemma 10.4.5. Consequently, in this case we can reduce the dimension by at least 1 . From now on we may assume that

$$
\begin{equation*}
\max \left\{\alpha_{i}, \alpha_{i}^{\prime}, \beta_{j}, \beta_{j}^{\prime}, \gamma_{i j}, \gamma_{i j}^{\prime}\right\} \leq 2 \sqrt{n} \text { gap. } \tag{10.8}
\end{equation*}
$$

Geometrically, (10.8) says that if the supporting hyperplanes are both mostly made up of the separating hyperplanes, then their aggregate contributions to $F$ and $F^{\prime}$ should be small in absolute value.

The next lemma identifies some $p \in V$ for which $f(R(p))-f(R(p)-p)$ is "big". This prepares for the final step of our approach which invokes Lemma 10.4.9.
Lemma 10.4.15. Let $\boldsymbol{y} \stackrel{\text { def }}{=} \sum_{h \in H} \lambda_{h} \boldsymbol{h}$ and $\boldsymbol{y}^{\prime} \stackrel{\text { def }}{=} \sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}$ and let $p \in \arg \max _{l}\left\{\max \left\{\left|y_{l}\right|,\left|y_{l}^{\prime}\right|\right\}\right\}$ then

$$
\operatorname{upper}(p) \geq n^{7}\left\|\boldsymbol{y}+\boldsymbol{y}^{\prime}\right\|_{\infty}
$$

assuming (10.8).
Proof. Recall that $\left\|\boldsymbol{c}+\boldsymbol{c}^{\prime}\right\|_{2} \leq$ gap where gap $=\frac{1}{10 n^{10}} \min \left\{\|\boldsymbol{c}\|_{2},\left\|\boldsymbol{c}^{\prime}\right\|_{2}\right\}$,
$\boldsymbol{c}=\boldsymbol{y}-\sum_{i} \alpha_{i} \mathbf{1}_{i}+\sum_{j} \beta_{j} \mathbf{1}_{j}+\sum_{(i, j)} \gamma_{i j}\left(\mathbf{1}_{i}-\mathbf{1}_{j}\right) \quad$ and $\boldsymbol{c}^{\prime}=\boldsymbol{y}^{\prime}-\sum_{i} \alpha_{i}^{\prime} \mathbf{1}_{i}+\sum_{j} \beta_{j}^{\prime} \mathbf{1}_{j}+\sum_{(i, j)} \gamma_{i j}^{\prime}\left(\mathbf{1}_{i}-\mathbf{1}_{j}\right)$.
By (10.8) we know that $\|\boldsymbol{c}-\boldsymbol{y}\|_{2} \leq 4 n^{2}$ gap $\leq \frac{4}{10 n^{8}}\|\boldsymbol{c}\|_{2}$ and $\left\|\boldsymbol{c}^{\prime}-\boldsymbol{y}^{\prime}\right\|_{2} \leq 4 n^{2}$ gap $\leq \frac{4}{10 n^{8}}\left\|\boldsymbol{c}^{\prime}\right\|_{2}$. Consequently, by the triangle inequality we have that

$$
\left\|\boldsymbol{y}+\boldsymbol{y}^{\prime}\right\|_{2} \leq\left\|\boldsymbol{c}+\boldsymbol{c}^{\prime}\right\|_{2}+\|\boldsymbol{c}-\boldsymbol{y}\|_{2}+\left\|\boldsymbol{c}^{\prime}-\boldsymbol{y}^{\prime}\right\|_{2} \leq 9 n^{2} \text { gap }
$$

and

$$
\|\boldsymbol{c}\|_{2} \leq\|\boldsymbol{c}-\boldsymbol{y}\|_{2}+\|\boldsymbol{y}\|_{2} \leq \frac{4}{10 n^{8}}\|\boldsymbol{c}\|_{2}+\|\boldsymbol{y}\|_{2} \quad \Rightarrow \quad\|\boldsymbol{c}\|_{2} \leq 2\|\boldsymbol{y}\|_{2}
$$

Similarly, we have that $\left\|\boldsymbol{c}^{\prime}\right\|_{2} \leq 2\left\|\boldsymbol{y}^{\prime}\right\|_{2}$. Consequently since gap $\leq \frac{1}{10 n^{10}} \min \left\{\|\boldsymbol{c}\|_{2},\left\|\boldsymbol{c}^{\prime}\right\|_{2}\right\}$, we have that

$$
\left\|\boldsymbol{y}+\boldsymbol{y}^{\prime}\right\|_{2} \leq \frac{2}{n^{8}} \min \left\{\|\boldsymbol{y}\|_{2},\left\|\boldsymbol{y}^{\prime}\right\|_{2}\right\}
$$

and thus, invoking Lemma 10.4.11 yields the result.
We summarize the results in the lemma below.
Corollary 10.4.16. Let $P$ be the feasible region after running cutting plane on (10.1). Then one of the following holds:

1. We found a degenerate BFS and hence either $\emptyset$ or $V$ is a minimizer.
2. The integral points of $P$ all lie on some hyperplane $x_{i}=0, x_{j}=1$ or $x_{i}=x_{j}$ which we can find.
3. Let $H$ be the collection of BFS's $\boldsymbol{h}$ used to construct our separating hyperplanes for $P$. Then there is a convex combination $\boldsymbol{y}$ of $H$ such that $n^{4}\left|y_{i}\right|<\max _{p} \operatorname{upper}(p)$ for all $i$.

Proof. As mentioned before, (1) happens if some separating hyperplane is degenerate. We have (2) if one of the conditions in Lemma 10.4.14 holds. Otherwise, $y=\sum_{h \in H} \lambda_{h} \boldsymbol{h}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}$ is a candidate for Case 3 by Lemma 10.4.15.

Let us revisit the conditions of Lemma 10.4.9 and explain that they are satisfied by Case 3 of the last lemma.

- $\boldsymbol{y}$ is a convex combination of at most $O(n)$ BFS's. This holds in Case 3 since our current feasible region consists of only $O(n)$ constraints thanks to the Cutting Plane method.
- Those BFS's must be consistent with every arc of $A$. This holds because Case 3 uses the BFS's for constructing our separating hyperplane. Our modified separation oracle guarantees that they are consistent with $A$.

Thus in Case 3 of the last corollary, Lemma 10.4.9 allows us to deduce a new constraint $x_{p} \leq x_{q}$ for some $q \notin R(p)$.

### 10.4.3.3 Running Time

Here we bound the total running time of our algorithm and prove the following.
Theorem 10.4.17. Our algorithm runs in time $O\left(n^{4} \log n \cdot E O+n^{5} \log O(1) n\right)$.
Proof. To avoid being repetitive, we appeal to Corollary 10.4.16. Each phase of cutting plane takes time $O\left(n^{2} \log n \cdot \mathrm{EO}+n^{3} \log { }^{O(1)} n\right)$ (Theorem 10.4.12 with $\tau$ being a big constant. Given $F$ and $F^{\prime}$ represented as a nonnegative combination of facets, we can check for the conditions in Lemma 10.4.14 in $O(n)$ time as there are only this many facets of $P$. This settles Case 2 of Corollary 10.4.16. Finally, Lemma 10.4.9 tells us that we can find a new arc in $O\left(n \cdot \mathrm{EO}+n^{2}\right)$ time for Case 3 of Corollary 10.4.16. Our conclusion follows from the fact that we can get $x_{i}=0, x_{i}=1, x_{i}=x_{j}$ at most $n$ times and $x_{i} \leq x_{j}$ at most $O\left(n^{2}\right)$ times.

## - 10.4.4 $\widetilde{O}\left(n^{3} \cdot \mathbf{E O}+n^{4}\right)$ Time Algorithm

Here we show how to improve our running time for strongly polynomial SFM to $\widetilde{O}\left(n^{3} \cdot \mathrm{EO}+n^{4}\right)$. Our algorithm can be viewed as an extension of the algorithm we presented in the previous Section 10.4.3. The main bottleneck of our previous algorithm was the time needed to identify a new arc, which cost us $\widetilde{O}\left(n^{2} \cdot \mathrm{EO}+n^{3}\right)$. Here we show how to reduce our amortized cost for identifying a valid arc down to $\widetilde{O}\left(n \cdot \mathrm{EO}+n^{2}\right)$ and thereby achieve our result.

The key observation we make to improve this running time is that our choice of $p$ for adding an arc in the previous lemma can be relaxed. $p$ actually need not be arg $\max _{i}$ upper $(i)$; instead it is enough to have upper $(p)>n^{4} \max \left\{\alpha_{i}, \alpha_{i}^{\prime}, \beta_{j}, \beta_{j}^{\prime}, \gamma_{i j}, \gamma_{i j}^{\prime}\right\}$. For each such $p$ a new constraint $x_{p} \leq x_{q}$ can be identified via Lemma 10.4.9. So if there are many $p$ 's satisfying this we will be able to obtain many new constraints and hence new valid arcs $(p, q)$.

On the other hand, the bound in Lemma 10.4.15 says that our point in the base polyhedron is small in absolute value. This is actually stronger than what we need in Lemma 10.4.9 which requires only its positive entries to be "small". However as we saw in Lemma 10.4.10 we can generate a constraint of the form $x_{q} \leq x_{p}$ whenever lower $(p)$ is sufficiently negative.

Using this idea, we divide $V$ into different buckets according to $\operatorname{upper}(p)$ and $\operatorname{lower}(p)$. This will allow us to get a speedup for two reasons.

First, bucketing allows us to disregard unimportant elements of $V$ during certain executions of our cutting plane method. If both upper $(i)$ and lower $(i)$ are small in absolute value, then $i$ is essentially negligible because for a separating hyperplane $\boldsymbol{h}^{T} \boldsymbol{x} \leq \hat{f}(\bar{x})$, any $h_{i} \in[$ lower $(i)$, upper $(i)]$ small in absolute value would not really make a difference. We can then run our cutting plane algorithm only on those non-negligible $i$ 's, thereby reducing our time complexity. Of course, whether $h_{i}$ is small is something relative. This suggests that partitioning the ground set by the relative size of upper $(i)$ and lower $(i)$ is a good idea.

Second, bucketing allows us to ensure that we can always add an arc for many edges simultaneously. Recall that we remarked that all we want is $n^{O(1)}\left|y_{i}\right| \leq \operatorname{upper}(p)$ for some $\boldsymbol{y}$ in the base polyhedron. This would be sufficient to identify a new valid arc $(p, q)$. Now if the marginal differences upper $(p)$ and $\operatorname{upper}\left(p^{\prime}\right)$ are close in value, knowing $n^{O(1)}\left|y_{i}\right| \leq \operatorname{upper}(p)$ would effectively give us the same for $p^{\prime}$ for free. This suggests that elements with similar marginal differences should be grouped together.

The remainder of this section simply formalizes these ideas. In Section 10.4.4.1 we discuss how we partition the ground set $V$. In Section 10.4.4.2, we present our cutting plane method on a subset of the coordinates. Then in Section 10.4.4.3 we show how we find new arcs. Finally, in Section 10.4.4.4 we put all of this together to achieve our desired running time.

### 10.4.4.1 Partitioning Ground Set into Buckets

We partition the ground set $V$ into different buckets according to the values of upper $(i)$ and lower $(i)$. This is reminiscent to Iwata-Orlin's algorithm [125] which considers elements with big upper $(i)$. However they did not need to do bucketing by size or to consider lower $(i)$, whereas these seem necessary for our algorithm.

Let $N=\max _{i}\{\max \{\operatorname{upper}(i),-\operatorname{lower}(i)\}\}$ be the largest marginal difference in absolute value. By Lemma (10.4.8), $N \geq 0$. We partition our ground set $V$ as follows:

$$
\begin{gathered}
B_{1}=\left\{i: \operatorname{upper}(i) \geq N / n^{10} \text { or lower }(i) \leq-N / n^{10}\right\} \\
B_{k}=\left\{i \notin B_{1} \cup \ldots \cup B_{k-1}: \quad N / n^{10 k} \leq \operatorname{upper}(i)<N / n^{10(k-1)}\right. \\
\\
\text { or } \left.-N / n^{10(k-1)}<\operatorname{lower}(i) \leq-N / n^{10 k}\right\}, \quad k \geq 2
\end{gathered}
$$

We call $B_{k}$ buckets. Our buckets group elements by the values of upper $(i)$ and lower $(i)$ at $1 / n^{10}$ "precision". There are two cases.

- Case 1: the number of buckets is at most $\log n^{2}$, in which case $\operatorname{upper}(i)>N / n^{O(\log n)}$ or lower $(i)<-N / n^{O(\log n)}$ for all $i$.
- Case 2: there is some $k$ for which $\left|B_{1} \cup \ldots \cup B_{k}\right| \geq\left|B_{k+1}\right|$.

This is because if there is no such $k$ in Case 2, then by induction each bucket $B_{k+1}$ has at least $2^{k}\left|B_{1}\right| \geq 2^{k}$ elements and hence $k \leq \log n$.

Case 1 is easier to handle, and is in fact a special case of Case 2 . We first informally sketch the treatment for Case 1 which should shed some light into how we deal with Case 2.

We run Cutting Plane for $O\left(n \log ^{2} n\right)$ iterations (i.e. $\tau=\Theta(\log n)$ ). By Theorem 10.4.12, our feasible region $P$ would be sandwiched by a pair of approximately parallel supporting hyperplanes of width at most $1 / n^{\Theta(\log n)}$. Now proceeding as in the last section, we would be able to find some $\boldsymbol{y}$ in the base polyhedron and some element $p$ such that $n^{\Theta(\log n)}\left|y_{i}\right| \leq \operatorname{upper}(p)$. This gives

$$
n^{\Theta(\log n)}\left|y_{i}\right| \leq \frac{\operatorname{upper}(p)}{n^{\Theta(\log n)}} \leq \frac{N}{n^{\Theta(\log n)}}
$$

Since upper $(i)>N / n^{\Theta(\log n)}$ or lower $(i)<-N / n^{\Theta(\log n)}$ for all $i$ in Case 1, we can then conclude that some valid arc $(i, q)$ or $(q, i)$ can be added for every $i$. Thus we add $n / 2$ arcs simultaneously in one phase of the algorithm at the expense of blowing up the runtime by $O(\log n)$. This saves a factor of $n / \log n$ from our runtime in the last section, and the amortized cost for an arc would then be $\widetilde{O}\left(n \cdot \mathrm{EO}+n^{2}\right)$.

[^20]On the other hand, in Case 2 we have a "trough" at $B_{k+1}$. Roughly speaking, this trough is useful for acting as a soft boundary between $B_{1} \cup \ldots \cup B_{k}$ and $\bigcup_{l \geq k+2} B_{l}$. Recall that we are able to "ignore" $\bigcup_{l \geq k+2} B_{l}$ because their $h_{i}$ is relatively small in absolute value. In particular, we know that for any $p \in B_{1} \cup \ldots \cup B_{k}$ and $i \in B_{l}$, where $l \geq k+2$,

$$
\max \{\operatorname{upper}(p),- \text { lower }(p)\} \geq n^{10} \max \{\operatorname{upper}(i),- \text { lower }(i)\} .
$$

This is possible because $B_{k+1}$, which is sandwiched in between, acts like a shield preventing $B_{l}$ to "mess with" $B_{1} \cup \ldots \cup B_{k}$. This property comes at the expense of sacrificing $B_{k+1}$ which must confront $B_{l}$.

Furthermore, we require that $\left|B_{1} \cup \ldots \cup B_{k}\right| \geq\left|B_{k+1}\right|$, and run Cutting Plane on $B=\left(B_{1} \cup \ldots \cup\right.$ $\left.B_{k}\right) \cup B_{k+1}$. If $\left|B_{k+1}\right| \gg\left|B_{1} \cup \ldots \cup B_{k}\right|$, our effort would mostly be wasted on $B_{k+1}$ which is sacrificed, and the amortized time complexity for $B_{1} \cup \ldots \cup B_{k}$ would then be large.

Before discussing the algorithm for Case 2, we need some preparatory work.

### 10.4.4.2 Separating Hyperplane: Project and Lift

Our speedup is achieved by running our cutting plane method on the projection of our feasible region onto $B:=\left(B_{1} \cup \cdots \cup B_{k}\right) \cup B_{k+1}$. More precisely, we start by running our cutting plane on $P^{B}=$ $\left\{\boldsymbol{x} \in \mathbb{R}^{B}: \exists \boldsymbol{x}^{\prime} \in \mathbb{R}^{\bar{B}}\right.$ s.t. $\left(\boldsymbol{x}, \boldsymbol{x}^{\prime}\right)$ satisfies (10.1) $\}$, which has a lower dimension. However, to do this, we need to specify a separation oracle for $P^{B}$. Here we make one of the most natural choices.

We begin by making an apparently immaterial change to our set of arcs $A$. Let us take the transitive closure of $A$ by adding the arc $(i, j)$ whenever there is a path from $i$ to $j$. Clearly this would not change our ring family as a path from $i$ to $j$ implies $j \in R(i)$. Roughly speaking, we do this to handle pathological cases such as $(i, k),(k, j) \in A,(i, j) \notin A$ and $i, j \in B, k \notin B$. Without introducing the arc $(i, j)$, we risk confusing a solution containing $i$ but not $j$ as feasible since we are restricting our attention to $B$ and ignoring $k \notin B$.

Definition 10.4.18. Given a digraph $D=(V, A)$, the transitive closure of $A$ is the set of arcs $(i, j)$ for which there is a directed path from $i$ to $j$. We say that $A$ is complete if it is equal to its transitive closure.

Given $\bar{x} \in[0,1]^{B}$, we define the completion of $\bar{x}$ with respect to $A$ as follows.
Definition 10.4.19. Given $\bar{x} \in[0,1]^{B}$ and a set of $\operatorname{arcs} A, x^{\mathcal{C}} \in[0,1]^{n}$ is a completion of $\bar{x}$ if $x_{B}^{\mathcal{C}}=\bar{x}$ and $x_{i}^{\mathcal{C}} \leq x_{j}^{\mathcal{C}}$ for every $(i, j) \in A$. Here $x_{B}^{\mathcal{C}}$ denotes the restriction of $x^{\mathcal{C}}$ to $B$.

Lemma 10.4.20. Given $\bar{x} \in[0,1]^{B}$ and a complete set of arcs $A$, there is a completion of $\bar{x}$ if $\bar{x}_{i} \leq \bar{x}_{j}$ for every $(i, j) \in A \cap(B \times B)$. Moreover, it can be computed in $O\left(n^{2}\right)$ time.

Proof. We set $x_{B}^{\mathcal{C}}=\bar{x}$. For $i \notin B$, we set

$$
x_{i}^{\mathcal{C}}= \begin{cases}1 & \text { if } \nexists j \in B \text { s.t. }(i, j) \in A \\ \min _{(i, j) \in A, j \in B} x_{j}^{\mathcal{C}} & \text { otherwise }\end{cases}
$$

One may verify that $x^{\mathcal{C}}$ satisfies our requirement as $A$ is complete. Computing each $x_{i}^{\mathcal{C}}$ takes $O(n)$ time. Since $|V \backslash B|=|\bar{B}| \leq n$, computing the whole $x^{\mathcal{C}}$ takes $O\left(n^{2}\right)$ time.

This notion of completion is needed since our original separation oracle requires a full dimensional input $\bar{x}$. Now that $\bar{x} \in \mathbb{R}^{B}$, we need a way of extending it to $\mathbb{R}^{n}$ while retaining the crucial property that $\boldsymbol{h}$ is consistent with every arc in $A$.

```
Algorithm 26: Projected Separation Oracle
    Input: \(\bar{x} \in \mathbb{R}^{B}\) and a complete set of arcs \(A\)
    if \(\bar{x}_{i}<0\) for some \(i \in B\) then
        Output: \(x_{i} \geq 0\)
    else if \(\bar{x}_{j}>1\) for some \(j \in B\) then
        Output: \(x_{j} \leq 1\)
    else if \(\bar{x}_{i}>\bar{x}_{j}\) for some \((i, j) \in A \cap B^{2}\) then
        Output: \(x_{i} \leq x_{j}\)
    else
        Let \(x^{\mathcal{C}} \in \mathbb{R}^{n}\) be a completion of \(\bar{x}\)
        Let \(i_{1}, \ldots, i_{n}\) be a permutation of \(V\) such that \(x_{i_{1}}^{\mathcal{C}} \geq \ldots \geq x_{i_{n}}^{\mathcal{C}}\) and for all \((i, j) \in A, j\)
        precedes \(i\) in \(i_{1}, \ldots, i_{n}\).
        Output: \(\boldsymbol{h}_{B}^{T} \boldsymbol{x}_{B}=\sum_{i \in B} h_{i} x_{i} \leq \sum_{i \in B} h_{i} \bar{x}_{i}\), where \(\boldsymbol{h}\) is the BFS defined by the permutation
        \(i_{1}, \ldots, i_{n}\).
```

Note that the runtime is still $O\left(n \cdot \mathrm{EO}+n^{2} \log ^{O(1)} n\right)$ as $x^{\mathcal{C}}$ can be computed in $O\left(n^{2}\right)$ time by the last lemma.

We reckon that the hyperplane $\boldsymbol{h}_{B}^{T} \boldsymbol{x}_{B} \leq \sum_{i \in B} h_{i} \bar{x}_{i}$ returned by the oracle is not a valid separating hyperplane (i.e. it may cut out the minimizers). Nevertheless, we will show that it is a decent "proxy" to the true separating hyperplane $\boldsymbol{h}^{T} \boldsymbol{x} \leq \hat{f}\left(x^{\mathcal{C}}\right)=\sum_{i \in V} h_{i} x_{i}^{\mathcal{C}}$ and is good enough to serve our purpose of sandwiching the remaining feasible region in a small strip. To get a glimpse, note that the terms missing $\boldsymbol{h}_{B}^{T} \boldsymbol{x}_{B} \leq \sum_{i \in B} h_{i} \bar{x}_{i}$ all involve $h_{i}$ for $i \notin B$, which is "negligible" compared to $B_{1} \cup \cdots \cup B_{k}$.

One may try to make $\boldsymbol{h}_{B}^{T} \boldsymbol{x}_{B} \leq \sum_{i \in B} h_{i} \bar{x}_{i}$ valid, say, by $\boldsymbol{h}_{B}^{T} \boldsymbol{x}_{B} \leq \sum_{i \in B} h_{i} \bar{x}_{i}+\sum_{i \notin B}\left|h_{i}\right|$. The problem is that such hyperplanes would not be separating for $\bar{x}$ anymore as $\boldsymbol{h}_{B}^{T} \bar{x}=\sum_{i \in B} h_{i} \bar{x}_{i}<$ $\sum_{i \in B} h_{i} \bar{x}_{i}+\sum_{i \notin B}\left|h_{i}\right|$. Consequently, we lose the width (or volume) guarantee of our cutting plane algorithm. Although this seems problematic, it is actually still possible to show a guarantee sufficient for our purpose as $\sum_{i \notin B}\left|h_{i}\right|$ is relatively small. We leave it as a nontrivial exercise to interested readers.

In conclusion, it seems that one cannot have the best of both worlds: the hyperplane returned by the oracle cannot be simultaneously valid and separating.

## Algorithm

We take $k$ to be the first for which $\left|B_{1} \cup \ldots \cup B_{k}\right| \geq\left|B_{k+1}\right|$, i.e. $\left|B_{1} \cup \ldots \cup B_{l}\right|<\left|B_{l+1}\right|$ for $l \leq k-1$. Thus $k \leq \log n$. Let $b=|B|$, and so $\left|B_{1} \cup \cdots \cup B_{k}\right| \geq b / 2$. Case 1 is a special case by taking $B=V$.

Our algorithm is summarized below. Here $A$ is always complete as $A$ is replaced its transitive closure whenever a new valid arc is added.

1. Run Cutting Plane on $P^{B}=\left\{x \in \mathbb{R}^{B}: \exists x^{\prime} \in \mathbb{R}^{\bar{B}}\right.$ s.t. $\left(x, x^{\prime}\right)$ satisfies (10.1) $\}$ with the new projected separation oracle.
2. Identify a pair of "narrow" approximately parallel supporting hyperplanes.
3. Deduce from the hyperplanes certain new constraints of the forms $x_{i}=0, x_{j}=1, x_{i}=x_{j}$ or $x_{i} \leq x_{j}$ by lifting separating hyperplanes back to $\mathbb{R}^{n}$
4. Consolidate $A$ and $f$. If some $x_{i} \leq x_{j}$ added, replace $A$ by its transitive closure.
5. Repeat Step 1 with updated $A$ and $f$. (Any previously found separating hyperplanes are discarded.)

The minimizer can be constructed by unraveling the recursion.
First of all, to be able to run Cutting Plane on $P^{B}$ we must come up with a polyhedral description of $P^{B}$ which consists of just the constraints involving $B$. This is shown in the next lemma.

Lemma 10.4.21. Let $P^{B}=\left\{\boldsymbol{x} \in \mathbb{R}^{B}: \exists \boldsymbol{x}^{\prime} \in \mathbb{R}^{\bar{B}}\right.$ s.t. $\left(\boldsymbol{x}, \boldsymbol{x}^{\prime}\right)$ satisfies (10.1) $\}$. Then

$$
P^{B}=\left\{\boldsymbol{x} \in \mathbb{R}^{B}: 0 \leq \boldsymbol{x} \leq 1, x_{i} \leq x_{j} \forall(i, j) \in A \cap(B \times B)\right\}
$$

Proof. It is clear that $P^{B} \subseteq\left\{\boldsymbol{x} \in \mathbb{R}^{B}: 0 \leq \boldsymbol{x} \leq 1, x_{i} \leq x_{j} \forall(i, j) \in A \cap(B \times B)\right\}$ as the constraints $0 \leq x \leq 1, x_{i} \leq x_{j} \forall(i, j) \in A \cap(B \times B)$ all appear in (10.1).

Conversely, for any $\boldsymbol{x} \in \mathbb{R}^{B}$ satisfying $0 \leq \boldsymbol{x} \leq 1, x_{i} \leq x_{j} \forall(i, j) \in A \cap(B \times B)$, we know there is some completion $x^{\mathcal{C}}$ of $\boldsymbol{x}$ by Lemma 10.4.20 as $A$ is complete. Now $x^{\mathcal{C}}$ satisfies (10.1) by definition, and hence $\boldsymbol{x} \in P^{B}$.

The only place where we have really changed the algorithm is Step (3).

### 10.4.4.3 Deducing New Constraints $x_{i}=0, x_{j}=1, x_{i}=x_{j}$ or $x_{i} \leq x_{j}$

Our method will deduce one of the following:

- $x_{i}=0, x_{j}=1$ or $x_{i}=x_{j}$
- for each $p \in B_{1} \cup \cdots \cup B_{k}, x_{p} \leq x_{q}$ for some $q \notin R(p)$ or $x_{p} \geq x_{q}$ for some $q \notin Q(p)$

Our argument is very similar to the last section's. Roughly speaking, it is the same argument but with "noise" introduced by $i \notin B$. We use extensively the notations from the last section.

Our main tool is again Theorem 10.4.12. Note that $n$ should be replaced by $b$ in the Theorem statement. We invoke it with $\tau=k \log _{b} n=O\left(\log ^{2} n\right)$ (using $k \leq \log n$ ) to get a width of $1 / b^{\Theta(\tau)}=$ $1 / n^{\Theta(k)}$. This takes time at most $O\left(b n \log ^{2} n \cdot \mathrm{EO}+b n^{2} \log ^{O(1)} n\right)$. Again, this is intuitively clear as we run it for $O(k b \log n)$ iterations, each of which takes time $O\left(n \cdot \mathrm{EO}+n^{2} \log ^{O(1)} n\right)$.

After each phase of (roughly $O(k b \log n)$ iterations) of Cutting Plane, $P^{B}$ is sandwiched between a pair of approximately parallel supporting hyperplanes $F$ and $F^{\prime}$ which have width $1 / n^{20 k}$. Let $F$ and $F^{\prime}$ be

$$
\boldsymbol{c}^{T} \boldsymbol{x}_{B}=\sum_{i \in B} c_{i} x_{i} \leq M, \quad \boldsymbol{c}^{\prime T} \boldsymbol{x}_{B}=\sum_{i \in B} c_{i}^{\prime} x_{i} \leq M^{\prime}
$$

such that

$$
\left|M+M^{\prime}\right|,\left\|\boldsymbol{c}+\boldsymbol{c}^{\prime}\right\|_{2} \leq \text { gap, } \quad \text { where gap }=\frac{1}{n^{20 k}} \min \left\{\|\boldsymbol{c}\|_{2},\left\|\boldsymbol{c}^{\prime}\right\|_{2}\right\} .
$$

The rest of this section presents an execution of the ideas discussed above. All of our work is basically geared towards bringing the amortized cost for identifying a valid arc down to $\widetilde{O}\left(n \cdot \mathrm{EO}+n^{2}\right)$. Again, we can write these two constraints as a nonnegative combination. Here $\bar{x}_{h}^{\mathcal{C}}$ is the completion of the point $\bar{x}_{h}$ used to construct $\boldsymbol{h}_{B}^{T} \boldsymbol{x}_{B} \leq \boldsymbol{h}_{B}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{B}$. (Recall that $\left(\bar{x}_{h}^{\mathcal{C}}\right)_{B}$ is the restriction of $\bar{x}_{h}^{\mathcal{C}}$ to B.)

$$
\begin{aligned}
& \boldsymbol{c}^{T} \boldsymbol{x}_{B}=-\sum_{i \in B} \alpha_{i} x_{i}+\sum_{j \in B} \beta_{j} x_{j}+\sum_{(i, j) \in A \cap B^{2}} \gamma_{i j}\left(x_{i}-x_{j}\right)+\sum_{h \in H} \lambda_{h} \boldsymbol{h}_{B}^{T} \boldsymbol{x}_{B} \quad \text { and } \quad M=\sum_{j \in B} \beta_{j}+\sum_{h \in H} \lambda_{h} \boldsymbol{h}_{B}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{B} . \\
& \boldsymbol{c}^{\prime T} \boldsymbol{x}_{B}=-\sum_{i \in B} \alpha_{i}^{\prime} x_{i}+\sum_{j \in B} \beta_{j}^{\prime} x_{j}+\sum_{(i, j) \in A \cap B^{2}} \gamma_{i j}^{\prime}\left(x_{i}-x_{j}\right)+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}_{B}^{T} \boldsymbol{x}_{B} \quad \text { and } \quad M^{\prime}=\sum_{j \in B} \beta_{j}^{\prime}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}_{B}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{B} .
\end{aligned}
$$

As we have discussed, the problem is that the separating hyperplanes $\boldsymbol{h}_{B}^{T} \boldsymbol{x}_{B} \leq \boldsymbol{h}_{B}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{B}$ are not actually valid. We can, however, recover their valid counterpart by lifting them back to $\boldsymbol{h}^{T} \boldsymbol{x} \leq \boldsymbol{h}^{T} \bar{x}_{h}^{\mathcal{C}}$. The hope is that $\boldsymbol{h}_{B}^{T} \boldsymbol{x}_{B} \leq \boldsymbol{h}_{B}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{B}$ and $\boldsymbol{h}^{T} \boldsymbol{x} \leq \boldsymbol{h}^{T} \bar{x}_{h}^{\mathcal{C}}$ are not too different so that the arguments will still go through. We show that this is indeed the case.

Again, we scale $c, c^{\prime}, \alpha, \alpha^{\prime}, \beta, \beta^{\prime}, \gamma, \gamma^{\prime}, \lambda, \lambda^{\prime}$ so that

$$
\sum_{h \in H}\left(\lambda_{h}+\lambda_{h}^{\prime}\right)=1 .
$$

By adding all the constituent separating hyperplane inequalities, we get

$$
\sum_{h \in H} \lambda_{h} \boldsymbol{h}^{T} \boldsymbol{x}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}^{T} \boldsymbol{x} \leq \sum_{h \in H} \lambda_{h} \boldsymbol{h}^{T} \bar{x}_{h}^{\mathcal{C}}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}^{T} \bar{x}_{h}^{\mathcal{C}}
$$

Let

$$
L H S \stackrel{\text { def }}{=} \sum \alpha_{i} x_{i}+\sum \alpha_{i}^{\prime} x_{i}-\sum \beta_{j} x_{j}-\sum \beta_{j}^{\prime} x_{j}+\sum \gamma_{i j}\left(x_{j}-x_{i}\right)+\sum \gamma_{i j}^{\prime}\left(x_{j}-x_{i}\right) .
$$

Here we know that

$$
\begin{gathered}
\sum_{h \in H} \lambda_{h} \boldsymbol{h}^{T} \boldsymbol{x}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}^{T} \boldsymbol{x}=L H S+\left(\boldsymbol{c}+\boldsymbol{c}^{\prime}\right)^{T} \boldsymbol{x}_{B}+\sum_{h \in H} \lambda_{h} \boldsymbol{h}_{\bar{B}}^{T} \boldsymbol{x}_{\bar{B}}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}_{\bar{B}}^{T} \boldsymbol{x}_{\bar{B}} \\
\sum_{h \in H} \lambda_{h} \boldsymbol{h}^{T} \bar{x}_{h}^{\mathcal{C}}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}^{T} \bar{x}_{h}^{\mathcal{C}}=\left(M+M^{\prime}\right)+\sum_{h \in H} \lambda_{h} \boldsymbol{h}_{\bar{B}}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{\bar{B}}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}_{\bar{B}}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{\bar{B}}-\sum \beta_{j}-\sum \beta_{j}^{\prime}
\end{gathered}
$$

Combining all yields

$$
\begin{aligned}
& L H S+\left(\boldsymbol{c}+\boldsymbol{c}^{\prime}\right)^{T} \boldsymbol{x}_{B}+\sum_{h \in H} \lambda_{h} \boldsymbol{h}_{\bar{B}}^{T} \boldsymbol{x}_{\bar{B}}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}_{\bar{B}}^{T} \boldsymbol{x}_{\bar{B}} \\
\leq & \left(M+M^{\prime}\right)+\sum_{h \in H} \lambda_{h} \boldsymbol{h}_{\bar{B}}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{\bar{B}}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}_{\bar{B}}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{\bar{B}}-\sum \beta_{j}-\sum \beta_{j}^{\prime} .
\end{aligned}
$$

Here $\left(\boldsymbol{c}+\boldsymbol{c}^{\prime}\right)^{T} \boldsymbol{x}_{B}$ can be bounded as before: $\left(\boldsymbol{c}+\boldsymbol{c}^{\prime}\right)^{T} \boldsymbol{x}_{B} \geq-\sqrt{n}\left\|\boldsymbol{c}+\boldsymbol{c}^{\prime}\right\|_{2} \geq-\sqrt{n}$ gap. Since $M+M^{\prime} \leq$ gap, We then obtain

$$
L H S+\sum_{h \in H} \lambda_{h} \boldsymbol{h}_{\bar{B}}^{T} \boldsymbol{x}_{\bar{B}}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}_{\bar{B}}^{T} \boldsymbol{x}_{\bar{B}} \leq 2 \sqrt{n} \operatorname{gap}+\sum_{h \in H} \lambda_{h} \boldsymbol{h}_{\bar{B}}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{\bar{B}}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}_{\bar{B}}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{\bar{B}}-\sum \beta_{j}-\sum \beta_{j}^{\prime}
$$

We should expect the contribution from $\boldsymbol{h}_{\bar{B}}$ to be small as $h_{i}$ for $i \notin B$ is small compared to $B_{1} \cup \ldots \cup B_{k}$. We formalize our argument in the next two lemmas.

Lemma 10.4.22. We have $\sum_{h \in H} \lambda_{h} \boldsymbol{h}_{\bar{B}}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{\bar{B}}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}_{\bar{B}}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{\bar{B}} \leq N / n^{10(k+1)-1}$.
Proof. We bound each component of $\sum_{h \in H} \lambda_{h} \boldsymbol{h}_{\bar{B}}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{\bar{B}}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}_{\bar{B}}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{\bar{B}}$. For $i \in \bar{B}$, we have $\operatorname{upper}(i) \leq N / n^{10(k+1)}$. By Lemma 10.4.8 $h_{i} \leq \operatorname{upper}(i)$. Therefore,

$$
\sum_{h \in H} \lambda_{h} \boldsymbol{h}_{i}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{i}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}_{i}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{i} \leq\left(\sum_{h \in H} \lambda_{h}+\sum_{h \in H} \lambda_{h}^{\prime}\right) N / n^{10(k+1)}=N / n^{10(k+1)}
$$

Our result then follows since

$$
\sum_{h \in H} \lambda_{h} \boldsymbol{h}_{\bar{B}}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{\bar{B}}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}_{\bar{B}}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{\bar{B}}=\sum_{i \in \bar{B}}\left(\sum_{h \in H} \lambda_{h} \boldsymbol{h}_{i}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{i}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}_{i}^{T}\left(\bar{x}_{h}^{\mathcal{C}}\right)_{i}\right) .
$$

Lemma 10.4.23. We have $\sum_{h \in H} \lambda_{h} \boldsymbol{h}_{\bar{B}}^{T} \boldsymbol{x}_{\bar{B}}+\sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}_{\bar{B}}^{T} \boldsymbol{x}_{\bar{B}} \geq-N / n^{10(k+1)-1}$.
Proof. The proof is almost identical to the last lemma except that we use $h_{i} \geq \operatorname{lower}(i)$ instead of $h_{i} \leq \operatorname{upper}(i)$, and lower $(i) \geq-N / n^{10(k+1)}$.

The two lemmas above imply that

$$
L H S \leq 2 \sqrt{n} \mathrm{gap}-\sum \beta_{j}-\sum \beta_{j}+2 N / n^{10(k+1)-1}=\mathrm{gap}^{\prime}-\sum \beta_{j}-\sum \beta_{j}
$$

where gap $^{\prime}=2 \sqrt{n}$ gap $+2 N / n^{10(k+1)-1}$.
Lemma 10.4.24. Suppose $x$ satisfies (10.1) and LHS $\leq \operatorname{gap}^{\prime}-\sum \beta_{j}-\sum \beta_{j}^{\prime}$ with $\alpha_{i}, \beta_{j}, \gamma_{i j}, \alpha_{i}^{\prime}, \beta_{j}^{\prime}, \gamma_{i j}^{\prime} \geq$ 0 .

1. If $\alpha_{i}>\mathrm{gap}^{\prime}$ or $\alpha_{i}^{\prime}>\mathrm{gap}^{\prime}$, then $x_{i}<1$.
2. If $\beta_{j}>$ gap $^{\prime}$ or $\beta_{j}^{\prime}>$ gap $^{\prime}$, then $x_{j}>0$.
3. If $\gamma_{i j}>$ gap $^{\prime}$ or $\gamma_{i j}^{\prime}>$ gap $^{\prime}$, then $0 \leq x_{j}-x_{i}<1$.

Proof. The proof is exactly the same as Lemma 10.4 .14 with $2 \sqrt{n}$ gap replaced by gap' .
From now on we may assume that

$$
\begin{equation*}
\max \left\{\alpha_{i}, \alpha_{i}^{\prime}, \beta_{j}, \beta_{j}^{\prime}, \gamma_{i j}, \gamma_{i j}^{\prime}\right\} \leq \operatorname{gap}^{\prime} \tag{10.9}
\end{equation*}
$$

Lemma 10.4.25. Let $\boldsymbol{y} \stackrel{\text { def }}{=} \sum_{h \in H} \lambda_{h} \boldsymbol{h}$ and $\boldsymbol{y}^{\prime} \stackrel{\text { def }}{=} \sum_{h \in H} \lambda_{h}^{\prime} \boldsymbol{h}$ and let $p \in \arg \max _{l \in B}\left\{\max \left\{\left|y_{l}\right|,\left|y_{l}^{\prime}\right|\right\}\right.$ then

$$
N \geq n^{10 k+6}\left\|\boldsymbol{y}_{B}+\boldsymbol{y}_{B}^{\prime}\right\|_{\infty}
$$

assuming (10.9).
Proof. Recall that $\left\|\boldsymbol{c}+\boldsymbol{c}^{\prime}\right\|_{2} \leq$ gap $<$ gap $^{\prime}$ where gap $=\frac{1}{n^{20 k}} \min \left\{\|\boldsymbol{c}\|_{2},\left\|\boldsymbol{c}^{\prime}\right\|_{2}\right\}$ and gap ${ }^{\prime}=2 \sqrt{n}$ gap + $2 N / n^{10(k+1)-1}$. Now there are two cases.

Case 1: $2 \sqrt{n}$ gap $\geq 2 N / n^{10(k+1)-1}$. Then gap ${ }^{\prime} \leq 4 \sqrt{n}$ gap and we follow the same proof of Lemma 10.4.15. We have
$\boldsymbol{c}=\boldsymbol{y}_{B}-\sum_{i} \alpha_{i} \mathbf{1}_{i}+\sum_{j} \beta_{j} \mathbf{1}_{j}+\sum_{(i, j)} \gamma_{i j}\left(\mathbf{1}_{i}-\mathbf{1}_{j}\right) \quad$ and $\boldsymbol{c}^{\prime}=\boldsymbol{y}_{B}^{\prime}-\sum_{i} \alpha_{i}^{\prime} \mathbf{1}_{i}+\sum_{j} \beta_{j}^{\prime} \mathbf{1}_{j}+\sum_{(i, j)} \gamma_{i j}^{\prime}\left(\mathbf{1}_{i}-\mathbf{1}_{j}\right)$.
By (10.9) we know that $\left\|\boldsymbol{c}-\boldsymbol{y}_{B}\right\|_{2} \leq 4 n^{2} \operatorname{gap}^{\prime} \leq \frac{1}{n^{17 k}}\|\boldsymbol{c}\|_{2}$ and $\left\|\boldsymbol{c}^{\prime}-\boldsymbol{y}_{B}^{\prime}\right\|_{2} \leq 4 n^{2}$ gap $^{\prime} \leq \frac{1}{n^{17 k}}\|\boldsymbol{c}\|_{2}$. Consequently, by the triangle inequality we have that

$$
\left\|\boldsymbol{y}_{B}+\boldsymbol{y}_{B}^{\prime}\right\|_{2} \leq\left\|\boldsymbol{c}+\boldsymbol{c}^{\prime}\right\|_{2}+\left\|\boldsymbol{c}-\boldsymbol{y}_{B}\right\|_{2}+\left\|\boldsymbol{c}^{\prime}-\boldsymbol{y}_{B}^{\prime}\right\|_{2} \leq 9 n^{2} \text { gap }^{\prime}
$$

and

$$
\|\boldsymbol{c}\|_{2} \leq\left\|\boldsymbol{c}-\boldsymbol{y}_{B}\right\|_{2}+\left\|\boldsymbol{y}_{B}\right\|_{2} \leq \frac{1}{n^{17 k}}\|\boldsymbol{c}\|_{2}+\left\|\boldsymbol{y}_{B}\right\|_{2} \quad \Rightarrow \quad\|\boldsymbol{c}\|_{2} \leq 2\left\|\boldsymbol{y}_{B}\right\|_{2}
$$

Similarly, we have that $\left\|\boldsymbol{c}^{\prime}\right\|_{2} \leq 2\left\|\boldsymbol{y}_{B}^{\prime}\right\|_{2}$. Consequently since gap ${ }^{\prime} \leq \frac{1}{n^{19 k}} \min \left\{\|\boldsymbol{c}\|_{2},\left\|\boldsymbol{c}^{\prime}\right\|_{2}\right\}$, we have that

$$
\left\|\boldsymbol{y}_{B}+\boldsymbol{y}_{B}^{\prime}\right\|_{2} \leq \frac{18}{n^{17 k}} \min \left\{\left\|\boldsymbol{y}_{B}\right\|_{2},\left\|\boldsymbol{y}_{B}^{\prime}\right\|_{2}\right\}
$$

and thus, invoking Lemma 10.4 .11 yields $N \geq \operatorname{upper}(p) \geq n^{16 k}\left\|\boldsymbol{y}_{B}+\boldsymbol{y}_{B}^{\prime}\right\|_{\infty}$, as desired.
Case 2: $2 \sqrt{n}$ gap $<2 N / n^{10(k+1)-1}$. Then for any $i \in B,\left|c_{i}+c_{i}^{\prime}\right| \leq\left\|\boldsymbol{c}+\boldsymbol{c}^{\prime}\right\|_{2} \leq \operatorname{gap}<$ $2 N / n^{10(k+1)-1}$. Since

$$
\boldsymbol{y}_{B}+\boldsymbol{y}_{B}^{\prime}=\left(\boldsymbol{c}+\boldsymbol{c}^{\prime}\right)+\sum_{i} \alpha_{i} \mathbf{1}_{i}-\sum_{j} \beta_{j} \mathbf{1}_{j}-\sum_{(i, j)} \gamma_{i j}\left(\mathbf{1}_{i}-\mathbf{1}_{j}\right)+\sum_{i} \alpha_{i}^{\prime} \mathbf{1}_{i}-\sum_{j} \beta_{j}^{\prime} \mathbf{1}_{j}-\sum_{(i, j)} \gamma_{i j}^{\prime}\left(\mathbf{1}_{i}-\mathbf{1}_{j}\right)
$$

we have

$$
\left\|\boldsymbol{y}_{B}+\boldsymbol{y}_{B}^{\prime}\right\|_{\infty} \leq 2 N / n^{10(k+1)-1}+2 n^{1.5} \operatorname{gap}^{\prime} \leq N / n^{10 k+7} .
$$

Corollary 10.4.26. Let $P$ be the feasible region after running Cutting Plane on (10.1) with the projected separation oracle. Then one of the following holds:

1. We found a BFS $\boldsymbol{h}$ with $\boldsymbol{h}_{B}=0$.
2. The integral points of $P$ all lie on some hyperplane $x_{i}=0, x_{j}=1$ or $x_{i}=x_{j}$.
3. Let $H$ be the collection of BFS's $\boldsymbol{h}$ used to construct our separating hyperplanes for $P$. Then there is a convex combination $\boldsymbol{y}$ of $H$ such that for $p \in B_{1} \cup \cdots \cup B_{k}$, we have $n^{4}\left|y_{i}\right|<\operatorname{upper}(p)$ or lower $(p)<-n^{4}\left|y_{i}\right|$ for all $i$.

Proof. As mentioned before, (1) happens if some separating hyperplane satisfies $\boldsymbol{h}_{B}=0$ when running cutting plane on the non-negligible coordinates. We have (2) if some condition in Lemma 10.4.24 holds. Otherwise, we claim $y=\sum_{h} \lambda_{h} h+\sum_{h} \lambda_{h}^{\prime} h$ is a candidate for Case 3. $y$ is a convex combination of BFS and by Lemma 10.4.25, for the big elements $i \in B$ we have

$$
\left|y_{i}\right| \leq N / n^{10 k+6} \leq \frac{1}{n^{4}} \max \{\operatorname{upper}(p),-\operatorname{lower}(p)\} .
$$

where the last inequality holds since for $p \in B_{1} \cup \cdots \cup B_{k}, \max \{\operatorname{upper}(p),-\operatorname{lower}(p)\} \geq N / n^{10 k}$.
On the other hand, for the small elements $i \notin B,\left|y_{i}\right| \leq N / n^{10(k+1)} \leq \frac{1}{n^{4}} \max \{\operatorname{upper}(p),-\operatorname{lower}(p)\}$ as desired.

The gap is then smaller enough to add an arc for each $p \in B_{1} \cup \cdots \cup B_{k}$ by Lemmas 10.4.9 and 10.4.10. Therefore we can add a total of $\left|B_{1} \cup \cdots \cup B_{k}\right| / 2 \geq b / 4$ arcs with roughly $O(k b \log n)=\widetilde{O}(b)$ iterations of Cutting Plane, each of which takes $\widetilde{O}\left(n \cdot \mathrm{EO}+n^{2}\right)$. That is, the amortized cost for each arc is $\widetilde{O}\left(n \cdot \mathrm{EO}+n^{2}\right)$. We give a more formal time analysis in below but it should be somewhat clear why we have the desired time complexity.

Lemma 10.4.27. Suppose there is a convex combination $\boldsymbol{y}$ of $H$ such that for $p \in B_{1} \cup \cdots \cup B_{k}$, we have $n^{4}\left|y_{i}\right|<\operatorname{upper}(p)$ or $\operatorname{lower}(p)<-n^{4}\left|y_{i}\right|$ for all $i$. Then we can identify at least $b / 4$ new valid arcs.

Proof. We have $|H|=O(n)$ since $H$ is the set of BFS's used for the constraints of $P$ which has $O(n)$ constraints. By Lemmas 10.4.9 and 10.4.10, for $p \in B_{1} \cup \cdots \cup B_{k}$ we can add a new valid arc $(p, q)$ or $(q, p)$. However note that a new arc $\left(p_{1}, p_{2}\right)$ may added twice by both $p_{1}$ and $p_{2}$. Therefore the total number of new arcs is only at least $\left|B_{1} \cup \cdots \cup B_{k}\right| / 2 \geq b / 4$.

### 10.4.4.4 Running Time

Not much changes to the previous runtime analysis are needed. To avoid repetition, various details already present in the corresponding part of the last section are omitted. Recall $k \leq \log n$, and of course, $b \leq n$.

For each (roughly) $O(k b \log n)$ iterations of Cutting Plane we either get $x_{i}=0, x_{i}=1, x_{i}=x_{j}$ or $b / 4 x_{i} \leq x_{j}$ 's. The former can happen at most $n$ times while in the latter case, the amortized cost of each arc is $O(k \log n)$ iterations of Cutting Plane. In the worst case the overall number of iterations required is $\widetilde{O}\left(n^{2}\right)$. Thus our algorithm has a runtime of $\widetilde{O}\left(n^{3} \cdot \mathrm{EO}+n^{4}\right)$ since each iteration is $\widetilde{O}\left(n \cdot \mathrm{EO}+n^{2}\right)$ as shown below.
Theorem 10.4.28. Our algorithm runs in time $O\left(n^{3} \log ^{2} n \cdot E O+n^{4} \log O(1) n\right)$.
Proof. We use Corollary 10.4.26. First we note that Case 1 can actually be integrated into Case 3 since $\max \{\operatorname{upper}(p),-$ lower $(p)\} \geq N / n^{10 k}=n^{10} N / n^{10(k+1)} \geq h_{i}$ for $i \notin B$.

As we have argued in the beginning of the last section, Theorem 10.4.12 with $\tau=k \log _{b} n$ implies that the runtime for each phase is $O\left(b n \log ^{2} n \cdot \mathrm{EO}+b n^{2} \log ^{O(1)} n\right)$. In each phase we either get $x_{i}=0$, $x_{i}=1, x_{i}=x_{j}$ (Case 2) or $b / 4 x_{i} \leq x_{j}$ 's (Case 3), the latter of which follows from Corollary 10.4.26 and Lemma 10.4.27.

Case 2 can only happen $n$ times. Thus the total cost is at most $O\left(n^{3} \log ^{2} n \cdot \mathrm{EO}+n^{4} \log ^{O(1)} n\right)$. The overhead cost is also small. Similar to before, given $F$ and $F^{\prime}$ represented as a nonnegative combination of facets, we can check for the conditions in Lemma 10.4.24 in $O(n)$ time as there are only this many facets of $P$. This settles Case 2 .

For case 3 the amortized cost for each arc is $O\left(n \log ^{2} n \cdot \mathrm{EO}+n^{2} \log ^{O(1)} n\right)$. Our desired runtime follows since there are only $O\left(n^{2}\right)$ arcs to add. Unlike Case 2 some extra care is needed to handle the overhead cost. The time needed to deduce a new arc (applying Lemmas 10.4.9 and 10.4.10 to $\boldsymbol{y}$ and $\left.p \in B_{1} \cup \cdots \cup B_{k}\right)$ is still $O\left(n \cdot \mathrm{EO}+n^{2}\right)$. But as soon as we get a new arc, we must update $A$ to be its transitive closure so that it is still complete. Given $A$ complete and a new arc $(p, q) \notin A$, we can simply add the arcs from the ancestors of $p$ to $q$ and from $p$ to the descendants of $q$. There are at most $O(n)$ arcs to add so this takes time $O\left(n^{2}\right)$ per arc, which is okay.

## ■ 10.5 Discussion and Comparison with Previous Algorithms

We compare and contrast our algorithms with the previous ones. We focus primarily on strongly polynomial time algorithms.

## Convex combination of BFS's

All of the previous algorithms maintain a convex combination of BFS's and iteratively improve over it to get a better primal solution. In particular, the new BFS's used are typically obtained by making local changes to existing ones. Our algorithms, on the other hand, considers the geometry of the existing BFS's. The weighted "influences" ${ }^{3}$ then aggregately govern the choice of the next BFS. We believe that this is the main driving force for the speedup of our algorithms.

## Scaling schemes

Many algorithms for combinatorial problems are explicitly or implicitly scaling a potential function or a parameter. In this chapter, our algorithms in some sense aim to minimize the volume of the feasible region. Scaling schemes for different potential functions and parameters were also designed in previous works [121, 118, 125, 119]. All of these functions and parameters have an explict form. On the contrary, our potential function is somewhat unusual in the sense that it has no closed form.

[^21]
## Deducing new constraints

As mentioned in the main text, our algorithms share the same skeleton and tools for deducing new constraints with $[121,118,125,119]$. Nevertheless, there are differences in the way these tools are employed. Our algorithms proceed by invoking them in a geometric manner, whereas previous algorithms were mostly combinatorial.

## Big elements and bucketing

Our bucketing idea has roots in Iwata-Orlin's algorithm [125] but is much more sophisticated. For instance, it is sufficient for their algorithm to consider only big elements, i.e. upper $(i) \geq N / n^{O(1)}$. Our algorithm, on the other hand, must carefully group elements by the size of both upper $(i)$ and lower $(i)$. The speedup appears impossible without these new ideas. We do however note that it is unfair to expect such a sophisticated scheme in Iwata-Orlin's algorithm as it would not lead to a speedup. In other words, their method is fully sufficient for their purposes, and the simplicity in their case is a virtue rather than a shortcoming.

## ■ 10.5.1 Open Problems

One natural open problem is improving our weakly polynomial algorithm to $O\left(n^{2} \log M \cdot \mathrm{EO}+\right.$ $\left.n^{3} \log ^{O(1)} n \cdot \log M\right)$ time. Our application of center of mass to SFM demonstrates that it should be possible.

For strongly polynomial algorithms, the existential result of Theorem 10.4 . 1 shows that SFM can be solved with $O\left(n^{3} \log n \cdot E O\right)$ oracle calls. Unfortunately, our algorithm incurs an overhead of $\log n$ as there can be as many as $\log n$ buckets each time. One may try to remove this $\log n$ overhead by designing a better bucketing scheme or arguing that more arcs can be added.

The other $\log n$ overhead seem much trickier to remove. Our method currently makes crucial use of the tools developed by [121], where the $\log n$ factors in the runtime seem inevitable. We suspect that our algorithm may have an analogue similar to [231, 216], which do not carry any $\log n$ overhead in the running time.

Perhaps an even more interesting open problem is whether our algorithm is optimal (up to polylogarithmic factors). There are grounds for optimism. So far the best way of certifying the optimality of a given solution $S \subseteq V$ is to employ duality and express some optimal solution to the base polyhedron as a convex combination of $n+1$ BFS's. This already takes $n^{2}$ oracle calls as each BFS requires $n$. Thus one would expect the optimal number of oracle calls needed for SFM to be at least $n^{2}$. Our bound is not too far off from it, and anything strictly between $n^{2}$ and $n^{3}$ seems instinctively unnatural.

## First Order Method vs Cutting Plane Method

## ■ 11.1 Introduction

Let $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ be a $\beta$-smooth and $\alpha$-strongly convex function. Thus, for any $x, y \in \mathbb{R}^{n}$, we have

$$
f(x)+\nabla f(x)^{\top}(y-x)+\frac{\alpha}{2}|y-x|^{2} \leq f(y) \leq f(x)+\nabla f(x)^{\top}(y-x)+\frac{\beta}{2}|y-x|^{2} .
$$

Let $\kappa=\beta / \alpha$ be its condition number. It is a one line calculation to verify that a step of gradient descent on $f$ will decrease (multiplicatively) the squared distance to the optimum by $1-1 / \kappa$.

In Section 11.3, we propose a new method, which can be viewed as some combination of gradient descent and the ellipsoid method, for which the squared distance to the optimum decreases at a rate of $(1-1 / \sqrt{\kappa})$ (and each iteration requires one gradient evaluation and two line-searches). This matches the optimal rate of convergence among the class of first-order methods, [203, 206].

In Section 11.4, we generalize this idea and show how to automatically combine different optimization algorithms. In particular, we show how to combine gradient descent and cutting plane methods to get an algorithm with convergence rate $\min (1-1 / \kappa, 1-1 / O(n))$. This is the first algorithm we know that can naturally achieve both $\kappa$ dependence and $n$ dependence guarantee.

In Section 11.5, we provide some numerical evidence that the new methods can be superior to Nesterov's accelerated gradient descent.

## ■ 11.1.1 Preliminaries

We write $|\cdot|$ for the Euclidean norm in $\mathbb{R}^{n}$, and $\mathrm{B}\left(x, r^{2}\right):=\left\{y \in \mathbb{R}^{n}:|y-x|^{2} \leq r^{2}\right\}$ (note that the second argument is the radius squared). We define the map line_search : $\mathbb{R}^{n} \times \mathbb{R}^{n} \rightarrow \mathbb{R}^{n}$ by

$$
\text { line_search }(x, y)=\underset{t \in \mathbb{R}}{\arg \min } f(x+t(y-x)),
$$

and we denote

$$
x^{+}=x-\frac{1}{\beta} \nabla f(x), \text { and } x^{++}=x-\frac{1}{\alpha} \nabla f(x) .
$$

Recall that by strong convexity one has

$$
\forall y \in \mathbb{R}^{n}, f(y) \geq f(x)+\nabla f(x)^{\top}(y-x)+\frac{\alpha}{2}|y-x|^{2},
$$

which implies in particular:

$$
x^{*} \in \mathrm{~B}\left(x^{++}, \frac{|\nabla f(x)|^{2}}{\alpha^{2}}-\frac{2}{\alpha}\left(f(x)-f\left(x^{*}\right)\right)\right) .
$$

Furthermore recall that by smoothness one has $f\left(x^{+}\right) \leq f(x)-\frac{1}{2 \beta}|\nabla f(x)|^{2}$ which allows to shrink the above ball by a factor of $1-\frac{1}{\kappa}$ and obtain the following:

$$
\begin{equation*}
x^{*} \in \mathrm{~B}\left(x^{++}, \frac{|\nabla f(x)|^{2}}{\alpha^{2}}\left(1-\frac{1}{\kappa}\right)-\frac{2}{\alpha}\left(f\left(x^{+}\right)-f\left(x^{*}\right)\right)\right) \tag{11.1}
\end{equation*}
$$

## ■ 11.2 Intuition



Figure 11-1: One ball shrinks.


Figure 11-2: Two balls shrink.

The left diagram shows the intersection shrinks at the same rate if only one of the ball shrinks; the right diagram shows the intersection shrinks much faster if two balls shrinks at the same absolute amount.

In Section 11.2.1 we describe a geometric alternative to gradient descent (with the same convergence rate) which gives the core of our new optimal method. Then in Section 11.2 .2 we explain why one can expect to accelerate this geometric algorithm.

## ■ 11.2.1 A Suboptimal Algorithm

Assume that we are given a guarantee $R_{0}>0$ on the distance from some point $x_{0}$ to the optimum, that is $x^{*} \in \mathrm{~B}\left(x_{0}, R_{0}^{2}\right)$. Combining this original enclosing ball for $x^{*}$ and the one obtained by (11.1) (with $f\left(x^{*}\right) \leq f\left(x_{0}^{+}\right)$) one obtains

$$
x^{*} \in \mathrm{~B}\left(x_{0}, R_{0}^{2}\right) \cap \mathrm{B}\left(x_{0}-\frac{1}{\alpha} \nabla f\left(x_{0}\right), \frac{\left|\nabla f\left(x_{0}\right)\right|^{2}}{\alpha^{2}}\left(1-\frac{1}{\kappa}\right)\right) .
$$

If $\frac{\left|\nabla f\left(x_{0}\right)\right|^{2}}{\alpha^{2}} \leq R_{0}^{2}\left(1-\frac{1}{\kappa}\right)$ then the second ball already shrinks by a factor of $\left(1-\frac{1}{\kappa}\right)$. In the other case when $\frac{\left|\nabla f\left(x_{0}\right)\right|^{2}}{\alpha^{2}}>R_{0}^{2}\left(1-\frac{1}{\kappa}\right)$, the center of the two balls are far apart and therefore there is a much smaller ball containing the intersection of two balls. Formally, it is an easy calculation to see that for any $g \in \mathbb{R}^{n}, \varepsilon \in(0,1)$, there exists $x \in \mathbb{R}^{n}$ such that

$$
\begin{equation*}
\mathrm{B}(0,1) \cap \mathrm{B}\left(g,|g|^{2}(1-\varepsilon)\right) \subset \mathrm{B}(x, 1-\varepsilon) . \tag{Figure11-1}
\end{equation*}
$$

In particular the two above display implies that there exists $x_{1} \in \mathbb{R}^{n}$ such that

$$
x^{*} \in \mathrm{~B}\left(x_{1}, R_{0}^{2}\left(1-\frac{1}{\kappa}\right)\right) .
$$

Denote by $T$ the map from $x_{0}$ to $x_{1}$ defined implicitely above, and let $\left(x_{k}\right)$ be defined by $x_{k+1}=T\left(x_{k}\right)$. Then we just proved

$$
\left|x^{*}-x_{k}\right|^{2} \leq\left(1-\frac{1}{\kappa}\right)^{k} R_{0}^{2} .
$$

In other words, after $2 \kappa \log \left(R_{0} / \varepsilon\right)$ iterations where each iteration cost one call to the gradient oracle) one obtains a point $\varepsilon$-close to the minimizer of $f$.

## ■ 11.2.2 Why One Can Accelerate

Assume now that we are give a guarantee $R_{0}>0$ such that $x^{*} \in \mathrm{~B}\left(x_{0}, R_{0}^{2}-\frac{2}{\alpha}\left(f(y)-f\left(x^{*}\right)\right)\right)$ where $f\left(x_{0}\right) \leq f(y)$ (say by choosing $y=x_{0}$ ). Using the fact that $f\left(x_{0}^{+}\right) \leq f\left(x_{0}\right)-\frac{1}{2 \beta}\left|\nabla f\left(x_{0}\right)\right|^{2} \leq$ $f(y)-\frac{1}{2 \alpha \kappa}\left|\nabla f\left(x_{0}\right)\right|^{2}$, we obtain that

$$
x^{*} \in \mathrm{~B}\left(x_{0}, R_{0}^{2}-\frac{\left|\nabla f\left(x_{0}\right)\right|^{2}}{\alpha^{2} \kappa}-\frac{2}{\alpha}\left(f\left(x_{0}^{+}\right)-f\left(x^{*}\right)\right)\right)
$$

which, intuitively, allows us the shrink the radius squared from $R_{0}^{2}$ to $R_{0}^{2}-\frac{\left|\nabla f\left(x_{0}\right)\right|^{2}}{\alpha^{2} \kappa}$ using the local information at $x_{0}$. From (11.1), we have

$$
x^{*} \in \mathrm{~B}\left(x_{0}^{++}, \frac{\left|\nabla f\left(x_{0}\right)\right|^{2}}{\alpha^{2}}\left(1-\frac{1}{\kappa}\right)-\frac{2}{\alpha}\left(f\left(x_{0}^{+}\right)-f\left(x^{*}\right)\right)\right) .
$$

Now, intersecting the above two shrunk balls and using Lemma 11.3.2 (see below and also see Figure 11-2), we obtain that there is an $x_{1}^{\prime}$ such that

$$
x^{*} \in \mathrm{~B}\left(x_{1}^{\prime}, R_{0}^{2}\left(1-\frac{1}{\sqrt{\kappa}}\right)-\frac{2}{\alpha}\left(f\left(x_{0}^{+}\right)-f\left(x^{*}\right)\right)\right)
$$

giving us an acceleration in shrinking of the radius. To carry the argument for the next iteration, we would have required that $f\left(x_{1}^{\prime}\right) \leq f\left(x_{0}^{+}\right)$but it may not hold. Thus, we choose $x_{1}$ by a line search

$$
x_{1}=\text { line_search }\left(x_{1}^{\prime}, x_{0}^{+}\right)
$$

which ensures that $f\left(x_{1}\right) \leq f\left(x_{0}^{+}\right)$. To remedy the fact that the ball for the next iteration is centered at $x_{1}^{\prime}$ and not $x_{1}$, we observe that the line search also ensures that $\nabla f\left(x_{1}\right)$ is perpendicular to the line going through $x_{1}$ and $x_{1}^{\prime}$. This geometric fact is enough for the algorithm to work at the next iteration as well. In the next section we describe precisely our proposed algorithm which is based on the above insights.

## ■11.3 An Optimal Algorithm

Let $x_{0} \in \mathbb{R}^{n}, c_{0}=x_{0}^{++}$, and $R_{0}^{2}=\left(1-\frac{1}{\kappa}\right) \frac{\left|\nabla f\left(x_{0}\right)\right|^{2}}{\alpha^{2}}$. For any $k \geq 0$ let

$$
x_{k+1}=\text { line_search }\left(c_{k}, x_{k}^{+}\right),
$$

and $c_{k+1}$ (respectively $R_{k+1}^{2}$ ) be the center (respectively the squared radius) of the ball given by (the proof of) Lemma 11.3.2 which contains

$$
\mathrm{B}\left(c_{k}, R_{k}^{2}-\frac{\left|\nabla f\left(x_{k+1}\right)\right|^{2}}{\alpha^{2} \kappa}\right) \cap \mathrm{B}\left(x_{k+1}^{++}, \frac{\left|\nabla f\left(x_{k+1}\right)\right|^{2}}{\alpha^{2}}\left(1-\frac{1}{\kappa}\right)\right) .
$$

```
Algorithm 27: Minimum Enclosing Ball of the Intersection to Two Balls
    Input: a ball centered at \(x_{A}\) with radius \(R_{A}\) and a ball centered at \(x_{B}\) with radius \(R_{B}\).
    if \(\left|x_{A}-x_{B}\right|^{2} \geq\left|R_{A}^{2}-R_{B}^{2}\right|\) then
        \(c=\frac{1}{2}\left(x_{A}+x_{B}\right)-\frac{R_{A}^{2}-R_{B}^{2}}{2\left|x_{A}-x_{B}\right|^{2}}\left(x_{A}-X_{B}\right)\).
        \(R^{2}=R_{B}^{2}-\frac{\left(\left|x_{A}-x_{B}\right|^{2}+R_{B}^{2}-R_{A}^{2}\right)^{2}}{4\left|x_{A}-x_{B}\right|^{2}}\).
    else if \(\left|x_{A}-x_{B}\right|^{2}<R_{A}^{2}-R_{B}^{2}\) then
        \(c=x_{B} . R=R_{B}\).
    else
        \(c=x_{A} . \quad R=R_{A}\).
    end
    Output: a ball centered at \(c\) with radius \(R\).
```

```
Algorithm 28: Geometric Descent Method (GeoD)
    Input: parameters \(\alpha\) and initial points \(x_{0}\).
    \(x_{0}^{+}=\)line_search \(\left(x_{0}, x_{0}-\nabla f\left(x_{0}\right)\right)\).
    \(c_{0}=x_{0}-\bar{\alpha}^{-1} \nabla f\left(x_{0}\right)\).
    \(R_{0}^{2}=\frac{\left|\nabla f\left(x_{0}\right)\right|^{2}}{\alpha^{2}}-\frac{2}{\alpha}\left(f\left(x_{0}\right)-f\left(x_{0}^{+}\right)\right)\).
    for \(i \leftarrow 1,2, \cdots\) do
        Combining Step:
        \(x_{k}=\) line_search \(\left(x_{k-1}^{+}, c_{k-1}\right)\).
        Gradient Step:
        \(x_{k}^{+}=\)line_search \(\left(x_{k}, x_{k}-\nabla f\left(x_{k}\right)\right)\).
        Ellipsoid Step:
        \(x_{A}=x_{k}-\alpha^{-1} \nabla f\left(x_{k}\right) . R_{A}^{2}=\frac{\left|\nabla f\left(x_{k}\right)\right|^{2}}{\alpha^{2}}-\frac{2}{\alpha}\left(f\left(x_{k}\right)-f\left(x_{k}^{+}\right)\right)\).
        \(x_{B}=c_{k-1} . R_{B}^{2}=R_{k-1}^{2}-\frac{2}{\alpha}\left(f\left(x_{k-1}^{+}\right)-f\left(x_{k}^{+}\right)\right)\).
        Let \(B\left(c_{k}, R_{k}^{2}\right)\) is the minimum enclosing ball of \(B\left(x_{A}, R_{A}^{2}\right) \cap B\left(x_{B}, R_{B}^{2}\right)\).
    end
    Output: \(x_{T}\).
```

The formulas for $c_{k+1}$ and $R_{k+1}^{2}$ are given in Algorithm 27.
Theorem 11.3.1. For any $k \geq 0$, one has $x^{*} \in \mathrm{~B}\left(c_{k}, R_{k}^{2}\right), R_{k+1}^{2} \leq\left(1-\frac{1}{\sqrt{\kappa}}\right) R_{k}^{2}$, and thus

$$
\left|x^{*}-c_{k}\right|^{2} \leq\left(1-\frac{1}{\sqrt{\kappa}}\right)^{k} R_{0}^{2}
$$

Proof. We will prove a stronger claim by induction that for each $k \geq 0$, one has

$$
x^{*} \in \mathrm{~B}\left(c_{k}, R_{k}^{2}-\frac{2}{\alpha}\left(f\left(x_{k}^{+}\right)-f\left(x^{*}\right)\right)\right) .
$$

The case $k=0$ follows immediately by (11.1). Let us assume that the above display is true for some $k \geq 0$. Then using $f\left(x^{*}\right) \leq f\left(x_{k+1}^{+}\right) \leq f\left(x_{k+1}\right)-\frac{1}{2 \beta}\left|\nabla f\left(x_{k+1}\right)\right|^{2} \leq f\left(x_{k}^{+}\right)-\frac{1}{2 \beta}\left|\nabla f\left(x_{k+1}\right)\right|^{2}$, one gets

$$
x^{*} \in \mathrm{~B}\left(c_{k}, R_{k}^{2}-\frac{\left|\nabla f\left(x_{k+1}\right)\right|^{2}}{\alpha^{2} \kappa}-\frac{2}{\alpha}\left(f\left(x_{k+1}^{+}\right)-f\left(x^{*}\right)\right)\right) .
$$

Furthermore by (11.1) one also has

$$
\text { В }\left(x_{k+1}^{++}, \frac{\left|\nabla f\left(x_{k+1}\right)\right|^{2}}{\alpha^{2}}\left(1-\frac{1}{\kappa}\right)-\frac{2}{\alpha}\left(f\left(x_{k+1}^{+}\right)-f\left(x^{*}\right)\right)\right) .
$$

Thus it only remains to observe that the squared radius of the ball given by Lemma 11.3.2 which encloses the intersection of the two above balls is smaller than $\left(1-\frac{1}{\sqrt{\kappa}}\right) R_{k}^{2}-\frac{2}{\alpha}\left(f\left(x_{k+1}^{+}\right)-f\left(x^{*}\right)\right)$. We apply Lemma 11.3 .2 after moving $c_{k}$ to the origin and scaling distances by $R_{k}$. We set $\varepsilon=\frac{1}{\kappa}$, $g=\frac{\left|\nabla f\left(x_{k+1}\right)\right|}{\alpha}, \delta=\frac{2}{\alpha}\left(f\left(x_{k+1}^{+}\right)-f\left(x^{*}\right)\right)$ and $a=x_{k+1}^{++}-c_{k}$. The line search step of the algorithm implies that $\nabla f\left(x_{k+1}\right)^{\top}\left(x_{k+1}-c_{k}\right)=0$ and therefore, $|a|=\left|x_{k+1}^{++}-c_{k}\right| \geq\left|\nabla f\left(x_{k+1}\right)\right| / \alpha=g$ and Lemma 11.3.2 applies to give the result.

Lemma 11.3.2. Let $a \in \mathbb{R}^{n}$ and $\varepsilon \in(0,1), g \in \mathbb{R}_{+}$. Assume that $|a| \geq g$. Then there exists $c \in \mathbb{R}^{n}$ such that for any $\delta \geq 0$,

$$
\mathrm{B}\left(0,1-\varepsilon g^{2}-\delta\right) \cap \mathrm{B}\left(a, g^{2}(1-\varepsilon)-\delta\right) \subset \mathrm{B}(c, 1-\sqrt{\varepsilon}-\delta) .
$$

Proof. First observe that if $g^{2} \leq 1 / 2$ then one can take $c=a$ since $\frac{1}{2}(1-\varepsilon) \leq 1-\sqrt{\varepsilon}$. Thus we assume now that $g^{2}>1 / 2$, and note that we can also clearly assume that $n=2$. Consider the segment joining the two points at the intersection of the two balls under consideration. We denote $c$ for the point at the intersection of this segment and $[0, a]$, and $x=|c|$ (that is $c=x \frac{a}{|a|}$ ). A simple picture reveals that $x$ satisfies

$$
1-\varepsilon g^{2}-\delta-x^{2}=g^{2}(1-\varepsilon)-\delta-(|a|-x)^{2} \Leftrightarrow x=\frac{1+|a|^{2}-g^{2}}{2|a|} .
$$

When $x \leq|a|$, neither of the balls covers more than half of the other ball and hence the intersection of the two balls is contained in the ball $\mathrm{B}\left(x \frac{a}{|a|}, 1-\varepsilon g^{2}-\delta-x^{2}\right)$ (See figure 11-2). Thus it only remains to show that $x \leq|a|$ and that $1-\varepsilon g^{2}-\delta-x^{2} \leq 1-\sqrt{\varepsilon}-\delta$. The first inequality is equivalent to $|a|^{2}+g^{2} \geq 1$ which follows from $|a|^{2} \geq g^{2} \geq 1 / 2$. The second inequality to prove can be written as

$$
\varepsilon g^{2}+\frac{\left(1+|a|^{2}-g^{2}\right)^{2}}{4|a|^{2}} \geq \sqrt{\varepsilon}
$$

which is straightforward to verify (recall that $|a|^{2} \geq g^{2} \geq 1 / 2$ ).
Algorithm 28 we give is more agressive than Theorem 11.3.1, for instance, using line search instead of fixed step size. The correctness of this version follows from a similar proof as Theorem 11.3.1.

This algorithm does not require the smoothness parameter and the number of iterations; and it guarantees the function value is strictly decreasing. They are useful properties for machine learning applications because the only required parameter $\alpha$ is usually given. Furthermore, we believe that the integration of zeroth and first order information about the function makes our new method particularly well-suited in practice.

## ■ 11.4 Politician

In this section, we discuss how to in general combine different optimization algorithms. In standard black-box convex optimization [201, 206, 40] first-order methods interact with an oracle: given a query point $x$, the oracle reports the value and gradient of the underlying objective function $f$ at $x$. Here, we propose to replace the oracle by a politician. Instead of answering the original query $x$ the politician changes the question and answers a new query $y$ which is guaranteed to be better than the original
query $x$ in the sense that $f(y) \leq f(x)$. The newly selected query $y$ also depends on the history of queries that were made to the politician. Formally we introduce the following definition (for sake of simplicty we write $\nabla f(x)$ for either a gradient or a subgradient of $f$ at $x)$.

Definition 11.4.1. Let $\mathcal{X} \subset \mathbb{R}^{n}$ and $f: \mathcal{X} \rightarrow \mathbb{R}$. A politician $\Phi$ for $f$ is a mapping from $\mathcal{X} \times$ $\cup_{k=0}^{\infty}\left(\mathcal{X} \times \mathbb{R} \times \mathbb{R}^{n}\right)^{k}$ to $\mathcal{X}$ such that for any $k \geq 0, x \in \mathcal{X}, h \in\left(\mathcal{X} \times \mathbb{R} \times \mathbb{R}^{n}\right)^{k}$ one has $f(\Phi(x, h)) \leq f(x)$. Furthermore when queried at $x$ with history $h$ a politician for $f$ also output $f(\Phi(x, h))$ and $\nabla f(\Phi(x, h))$ (in order to not overload notation we do not include these outputs in the range of $\Phi$ ).

Let us clarify the interaction of a first-order method with a politician. Note that we refer to the couple (first-order method, politician) as the algorithm. Let $M: \cup_{k=0}^{\infty}\left(\mathcal{X} \times \mathbb{R} \times \mathbb{R}^{n}\right)^{k} \rightarrow \mathcal{X}$ be a firstorder method and $\Phi$ a politician for some function $f: \mathcal{X} \rightarrow \mathbb{R}$. The course of the algorithm $(M, \Phi)$ then goes as follows: at iteration $k+1$ one first calculates the method's query point $x_{k+1}=M\left(h_{k}\right)$ (with $h_{0}=\emptyset$ ), then one calculates the politician's new query point $y_{k+1}=\Phi\left(x_{k+1}, h_{k}\right)$ and the first order information at this point $\left(f\left(y_{k+1}\right), \nabla f\left(y_{k+1}\right)\right)$, and finally one updates the history with this new information $h_{k+1}=\left(h_{k},\left(y_{k+1}, f\left(y_{k+1}\right), \nabla f\left(y_{k+1}\right)\right)\right)$. Note that a standard oracle simply corresponds to a politician $\mathcal{O}$ for $f$ such that $\mathcal{O}(x, h)=(x, f(x), \nabla f(x))$ (in particular the algorithm $(M, \mathcal{O})$ is the usual algorithm corresponding to the first-order method $M)$.

The philosophy of the above definition is that it gives in some sense an automatic way to combine different optimization algorithms. Say for example that we wish to combine the ellipsoid method with gradient descent. One way to do so is to design an "ellipsoidal politician": the politician keeps track of a feasible ellipsoidal region based on the previously computed gradients, and when asked with the query $x$ the politician chooses as a new query $y$ the result of a line-search on the line between $x$ and the center of current ellipsoid. Gradient descent with this ellipsoidal politician would then replace the step $x \leftarrow x-\eta \nabla f(x)$ by $x \leftarrow y-\eta \nabla f(y)$. The hope is that in practice such a combination would integrate the fast incremental progress of gradient descent with the geometrical progress of the ellipsoid method.

In this section, we focus on unconstrained convex optimization. We are particularly interested in situations where calculating a (sub)gradient has superlinear complexity (i.e., $\gg n$ ) such as in logistic regression and semidefinite programming. In such cases it is natural to try to make the most out of the computed gradients by incorporating geometric reasoning (such as in the ellipsoid method). We do so by introducing the geometric politician (Section 11.4.2), which is based on a combination of the ideas in previous section with standard cutting plane machinery (through the notion of a "center" of a set, see Section 11.4.3). For a given first order method $M$, we denote by $M+$ the algorithm obtained by running $M$ with the geometric politician. We demonstrate empirically (Section 11.5) the effectiveness of the geometric politician on various standard first-order methods for convex optimization (gradient descent, Nesterov's accelerated gradient descent, non-linear conjugate gradient, BFGS). In particular we show that BFGS + is a surprisingly robust and parameter-free algorithm with state of the art performance across a wide range of problems (both smooth and non-smooth).

## ■ 11.4.1 Affine Invariant Politician

As mentioned above we assume that the complexity of computing the map $x \mapsto \nabla f(x)$ is superlinear. This implies that we can afford to have a politician such that the complexity of computing the map $(x, h) \mapsto \Phi(x, h)$ is $O(n \times \operatorname{poly}(k))$ (we think of the number of iterations $k$ as typically much smaller than the dimension $n$ ). We show in this section that this condition is (essentially) automatically satisfied as long as the politician is affine invariant in the following sense (we use a slight abuse of language and refer to a map $f \mapsto \Phi_{f}$, where $\Phi_{f}$ is a politician for $f$, as a politician):

Definition 11.4.2. A politician $f \mapsto \Phi_{f}$ is called affine invariant if for any function $f$ and any affine $\operatorname{map} T: \mathbb{R}^{m} \rightarrow \mathbb{R}^{n}$ such that $T(x)=z+L x$ for some matrix $L, k \geq 0, x \in \mathbb{R}^{m},\left(y_{i}, v_{i}, g_{i}\right) \in \mathbb{R}^{m} \times \mathbb{R} \times \mathbb{R}^{n}$, one has

$$
T\left(\Phi_{f \circ T}\left(x,\left(y_{i}, v_{i}, L^{\top} g_{i}\right)_{i \in[k]}\right)\right)=\Phi_{f}\left(T(x),\left(T\left(y_{i}\right), v_{i}, g_{i}\right)_{i \in[k]}\right) .
$$

We say that an affine invariant politician has cost $\psi: \mathbb{N} \rightarrow \mathbb{N}$ if for any $f: \mathbb{R}^{k} \rightarrow \mathbb{R}$ the map $(x, h) \in \mathbb{R}^{k} \times\left(\mathbb{R}^{k} \times \mathbb{R} \times \mathbb{R}^{k}\right)^{k} \mapsto \Phi_{f}(x, h)$ can be computed in time $\psi(k)$.
Proposition 11.4.3. Let $\Phi$ be an affine invariant politician with cost $\psi$. Then for any $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$, $\left(y_{i}, v_{i}, g_{i}\right) \in \mathbb{R}^{n} \times \mathbb{R} \times \mathbb{R}^{n}, i \in[k]$ and $x, y_{i} \in y_{1}+\operatorname{Span}\left(g_{1}, \cdots, g_{k}\right)$ one can compute $\Phi_{f}\left(x,\left(y_{i}, v_{i}, g_{i}\right)_{i \in[k]}\right) \in$ $\mathbb{R}^{n}$ in time $\psi(k)+O\left(n k^{2}\right)$.
Proof. Let $G$ be the $n \times k$ matrix with $i^{\text {th }}$ column given by $g_{i}$. We consider the $Q R$ decomposition of $G$ which can be computed in time $O\left(n k^{2}\right)$, that is $Q$ is an $n \times k$ matrix and $R$ a $k \times k$ matrix such that $G=Q R$ and $Q^{\top} Q=\mathrm{I}_{k}$. Let $T$ be the affine map defined by $T=y_{1}+Q$. Note that since $x \in y_{1}+\operatorname{Span}\left(g_{1}, \cdots, g_{k}\right)$ one has $x=T\left(Q^{\top}\left(x-y_{1}\right)\right)$ (and similarly for $\left.y_{i}\right)$. Thus by affine invariance one has

$$
\begin{aligned}
\Phi_{f}\left(x,\left(y_{i}, v_{i}, g_{i}\right)\right) & =\Phi_{f}\left(T\left(Q^{\top}\left(x-y_{1}\right)\right),\left(T\left(Q^{\top}\left(y_{i}-y_{1}\right)\right), v_{i}, g_{i}\right)\right) \\
& =y_{1}+Q \Phi_{f \circ T}\left(Q^{\top}\left(x-y_{1}\right),\left(Q^{\top}\left(y_{i}-y_{1}\right), v_{i}, R_{i}\right)\right),
\end{aligned}
$$

where $R_{i}$ is the $i^{\text {th }}$ column of $R$. Furthermore by definition of the cost $\psi$ and since $f \circ T$ is defined on $\mathbb{R}^{k}$ we see that this last quantity can be computed in time $\psi(k)+O\left(n k^{2}\right)$, thus concluding the proof.

The above proposition shows that with an affine invariant politician and a first order method $M$ verifying for any $\left(y_{i}, v_{i}, g_{i}\right)_{i \in[k]} \in\left(\mathbb{R}^{n} \times \mathbb{R} \times \mathbb{R}^{n}\right)^{k}$,

$$
M\left(\left(y_{i}, v_{i}, g_{i}\right)_{i \in[k]}\right) \in y_{1}+\operatorname{Span}\left(y_{1}, \cdots, y_{k}, g_{1}, \cdots, g_{k}\right)
$$

one can run $k$ steps of the corresponding algorithm in time $O\left(n k^{2}+k \psi(k)\right)$ plus the time to compute the $k$ function values and gradients of the underlying function $f$ to be optimized. Note that one gets a time of $O\left(n k^{2}\right)$ instead of $O\left(n k^{3}\right)$ as one can store the $Q R$ decomposition from one step to the next, and updating the decomposition only cost $O(n k)$.

## ■ 11.4.2 Geometric Politician

We describe in this section the geometric politician which is based on ideas developed in Section 11.3. A key observation in that Section is that if $f$ is a $\alpha$-strongly convex function minimized at $x^{*}$ then one has for any $x$,

$$
\left\|x^{*}-x-\frac{1}{\alpha} \nabla f(x)\right\|^{2} \leq \frac{\|\nabla f(x)\|^{2}}{\alpha^{2}}-\frac{2}{\alpha}\left(f(x)-f\left(x^{*}\right)\right) .
$$

This motivates the following definition:

$$
\mathrm{B}(x, \alpha, \text { fval }):=\left\{z \in \mathbb{R}^{n}:\left\|z-x-\frac{1}{\alpha} \nabla f(x)\right\|^{2} \leq \frac{\|\nabla f(x)\|^{2}}{\alpha^{2}}-\frac{2}{\alpha}(f(x)-\mathrm{fval})\right\} .
$$

In particular given the first order information at $y_{1}, \cdots, y_{k}$ one knows that the optimum $x^{*}$ lies in the region $R_{k} \subset \mathbb{R}^{n}$ defined by

$$
\begin{equation*}
R_{k}=\bigcap_{i \in[k]} B\left(y_{i}, \alpha, \text { fval }\right) \text { where fval }=\min _{i \in[k]} f\left(y_{i}\right) . \tag{11.2}
\end{equation*}
$$

```
Algorithm 29: Geometric Politician
    Parameter: An upper bound on the strong convexity parameter \(\alpha\). (Can be \(+\infty\).)
    Input: Query \(x\), past queries and the corresponding first order information
    \(\left(y_{i}, f\left(y_{i}\right), \nabla f\left(y_{i}\right)\right)_{i \in[k]}\).
    Let fval \(=\min _{i \in[k]} f\left(y_{i}\right)\) and the feasible region \(R_{k}(\alpha)=\bigcap_{i \in[k]} \mathrm{B}\left(y_{i}, \alpha\right.\), fval \()\).
    if \(R_{k}(\alpha)=\emptyset\) then
        Let \(\alpha\) be the largest number such that \(R_{k}(\alpha) \neq \emptyset\).
        \(\alpha \leftarrow \alpha / 4\).
    end
    Let \(y_{k+1}=\arg \min _{y \in\left\{t x+(1-t) c\left(R_{k}(\alpha)\right), t \in \mathbb{R}\right\}} f(y)\)
    where \(c\left(R_{k}(\alpha)\right)\) is the volumetric center of \(R_{k}(\alpha)\) (see Section 11.4.3).
    Output: \(y_{k+1}, f\left(y_{k+1}\right)\) and \(\nabla f\left(y_{k+1}\right)\).
```

Now suppose that given this first order information at $y_{1}, \cdots, y_{k}$ the first order method asks to query $x$. How should we modify this query in order to take into account the geometric information that $x^{*} \in R_{k}$ ? First observe that for any $z, \mathrm{~B}(z, \alpha, \mathrm{fval})$ is contained in a halfspace that has $z$ on its boundary (in the limiting case $\alpha \rightarrow 0$ the set $\mathrm{B}(z, \alpha, f(z))$ is exactly a halfspace). In particular if the next query point $y_{k+1}$ is the center of gravity of $R_{k}$ then we have that the volume of $R_{k+1}$ is at most $1-1 / e$ times the volume of $R_{k}$ (see [111]), thus leading to an exponential convergence rate. However the region $R_{k}$ can be very large initially, and the center of gravity might have a large function value and gradient, which means that $R_{k}$ would be intersected with a large sphere (possibly so large that it is close to a halfspace). On the other hand the first order method recommends to query $x$, which we can think of as a local improvement of $y_{k}$, which should lead to a much smaller sphere. The issue is that the position of this sphere might be such that the intersection with $R_{k}$ is almost as large as the sphere itself. In order to balance between the geometric and function value/gradient considerations we propose for the new query to do a line search between the center of $R_{k}$ and the recommended query $x$. The geometric politician follows this recipe with two important modifications: (i) there are many choices of centers that would guarantee an exponential convergence rate while being much easier to compute than the center of gravity, and we choose here to consider the volumetric center, see Section 11.4.3 for the definition and more details about this notion; (ii) we use a simple heuristic to adapt online the strong convexity parameter $\alpha$, namely we start with some large value for $\alpha$ and if it happens that the feasible region $R_{k}$ is empty then we know that $\alpha$ was too large, in which case we reduce it. We can now describe formally the geometric politician, see Algorithm 29. Importantly one can verify that the geometric politician is affine invariant and thus can be implemented efficiently (see the proof of Proposition 11.4.3).

## ■ 11.4.3 Volumetric Center

The volumetric barrier for a polytope was introduced in [253] to construct an algorithm with both the oracle complexity of the center of gravity method and the computational complexity of the ellipsoid method (see [Section 2.3, [40]] for more details and Chapter 8 for recent advances on this construction). Recalling that the standard logarithmic barrier $F_{P}$ for the polytope $P=\left\{x \in \mathbb{R}^{n}: a_{i}^{\top} x<b_{i}, i \in[m]\right\}$ is defined by

$$
F_{P}(x)=-\sum_{i=1}^{m} \log \left(b_{i}-a_{i}^{\top} x\right),
$$

one defines the volumetric barrier $v_{P}$ for $P$ by

$$
v_{P}(x)=\log \operatorname{det}\left(\nabla^{2} F_{P}(x)\right) .
$$

The volumetric center $c(P)$ is then defined as the minimizer of $v_{P}$. In the context of the geometric politician (see Algorithm 29) we are dealing with an intersection of balls rather than an intersection of halfspaces. More precisely the region of interest is of the form:

$$
R=\bigcap_{i=1}^{k}\left\{x \in \mathbb{R}^{n}:\left\|x-c_{i}\right\| \leq r_{i}\right\} .
$$

For such a domain the natural self-concordant barrier to consider is:

$$
F_{R}(x)=-\frac{1}{2} \sum_{i=1}^{k} \log \left(r_{i}^{2}-\left\|x-c_{i}\right\|^{2}\right) .
$$

The volumetric barrier is defined as before by

$$
v_{R}(x)=\log \operatorname{det}\left(\nabla^{2} F_{R}(x)\right),
$$

and the volumetric center of $R$ is the minimizer of $v_{R}$. It is shown in [12] that $v_{R}$ is a self-concordant barrier which means that the center can be updated (when a new ball is added to $R$ ) via few iterations of Newton's method. Often in practice, it takes less than 5 iterations to update the minimizer of a self-concordant barrier $[103,26]$ when we add a new constraint. Hence, the complexity merely depends on how fast we can compute the gradient and Hessian of $F_{R}$ and $v_{R}$.
Proposition 11.4.4. For the analytic barrier $F_{R}$, we have that

$$
\begin{aligned}
\nabla F_{R}(x) & =A^{\top} 1_{k \times 1}, \\
\nabla^{2} F_{R}(x) & =2 A^{\top} A+\lambda^{(1)} I
\end{aligned}
$$

where $d$ is a vector defined by $\left(r_{i}^{2}-\left\|x-c_{i}\right\|^{2}\right)^{-1}$, $A$ is a $k \times n$ matrix with $i^{\text {th }}$ row given by $d_{i}(x)\left(x-c_{i}\right)$, $\lambda^{(p)}=\sum_{i \in[k]} d_{i}^{p}(x)$ and $1_{k \times 1}$ is a $k \times 1$ matrix with all entries being 1.

For the volumetric center, we have that

$$
\begin{aligned}
\nabla v_{R}(x) & =\left(\left(2 \operatorname{Tr} H^{-1}\right) \mathrm{I}+4 H^{-1}\right) A^{\top} d+8 A^{\top} \sigma, \\
\nabla^{2} v_{R}(x) & =48 A^{\top} \Sigma A-64 A^{\top}\left(A H^{-1} A^{\top}\right)^{(2)} A \\
& +\left(8 \operatorname{Tr}(D \Sigma)+2 \lambda^{(2)} \operatorname{Tr}\left(H^{-1}\right)\right) I+4 \lambda^{(2)} H^{-1} \\
& +8 \operatorname{Tr}\left(H^{-1}\right) A^{\top} D A+16 \operatorname{sym}\left(\left(A^{\top} D A H^{-1}\right)\right) \\
& -4 \operatorname{Tr}\left(H^{-2}\right) A^{\top} D \mathrm{~J} D A-8 H^{-1} A^{\top} D \mathrm{~J} D A H^{-1} \\
& -8 \operatorname{sym}\left(() A^{\top} D \mathrm{~J} D A H^{-2}\right)-8\left(d^{\top} A H^{-1} A^{\top} d\right) H^{-1} \\
& -16 \operatorname{sym}\left(\left(A^{\top} \boldsymbol{D i a g}\left(A H^{-2} A^{\top}\right) \mathrm{J} D A\right)\right) \\
& -32 \operatorname{sym}\left(\left(A^{\top} \boldsymbol{D i a g}\left(A H^{-1} A^{\top} d\right) A H^{-1}\right)\right)
\end{aligned}
$$

where $H=\nabla^{2} F_{R}(x), \sigma_{i}=e_{i}^{\top} A H^{-1} A^{\top} e_{i}, e_{i}$ is the indicator vector with $i^{\text {th }}$ coordinate, J is a $k \times k$ matrix with all entries being 1, $\operatorname{sym}(() B)=B+B^{\top}, \boldsymbol{\operatorname { D i a g }}(v)$ is a diagonal matrix with $\boldsymbol{\operatorname { D i a g }}(v)_{i i}=v_{i}$, $D=\boldsymbol{\operatorname { D i a g }}(d), \Sigma=\boldsymbol{\operatorname { D i a g }}(\sigma)$, and $B^{(2)}$ is the Schur square of $B$ defined by $B_{i j}^{(2)}=B_{i j}^{2}$.

The above proposition shows that one step of Newton method for analytic center requires 1 dense matrix multiplication and solving 1 linear system; and for volumetric center, it requires 5 dense matrix multiplications, 1 matrix inversion and solving 1 linear system if implemented correctly. Although the analytic center is a more popular choice for "geometrical" algorithms, we choose volumetric center here because it gives a better convergence rate $[253,20]$ and the extra cost $\psi(k)$ is negligible to the cost of updating QR decomposition $n k$.

## ■ 11.5 Experiments

In this section, we compare the geometric politician against two libraries for first order methods, minFunc [230] and TFOCS [32]. Both are popular MATLAB libraries for minimizing general smooth convex functions. Since the focus of this chapter is all about how to find a good step direction, we use the exact line search (up to machine accuracy) whenever possible. This eliminates the effect of different line searches and reduces the number of algorithms we need to test. TFOCS is the only algorithm we use which does not use line search because they do not provide such option. To compensate on the unfairness to TFOCS, we note that the algorithm TFOCS uses is accelerated gradient descent and hence we implement the Gonzaga-Karas's accelerated gradient descent [108], which is specifically designed to be used with exact line search. Another reason we pick this variant of accelerated gradient descent is because we found it to be the fastest variant of accelerated gradient descent (excluding the geometric descent in Section 11.3) for our tested data (Gonzaga and Karas also observed that on their own dataset).

The algorithms to be tested are the following:

- [SD] Steepest descent algorithm in minFunc.
- [Nes] Accelerated gradient descent, General Scheme 2.2.6 in [206].
- [TFOCS] Accelerated gradient descent in TFOCS.
- [GK] Gonzaga-Karas's of Accelerated Gradient Descent (Section 11.5.1).
- [Geo] Geometric Descent (Section 11.3).
- [CG] Non-Linear Conjugate Gradient in minFunc.
- [BFGS] Broyden-Fletcher-Goldfarb-Shanno algorithm in minFunc.
- [PCG] Preconditioned Non-Linear Conjugate Gradient in minFunc.
- [Ø+] Geometric Politician itself (Section 11.5.1).
- [GK+] Using GK with Geometric Oracle (Section 11.5.1).
- [BFGS+] Using BFGS with Geometric Oracle (Section 11.5.1).

We only tested the geometric oracle on GK and BFGS because they are respectively the best algorithms in theory and practice on our tested data. The $\emptyset+$ algorithm is used as the control group to test if the geometric politician by itself is sufficient to achieve good convergence rate. We note that all algorithms except Nes are parameter free; each step of SD, Nes, TFOCS, GK, Geo, CG takes $O(n)$ time and each step of BFGS, PCG, $\emptyset+$, GK+ and BFGS+ takes roughly $O(n k)$ time for $k^{\text {th }}$ iteration.

```
Algorithm 30: \(\emptyset+\)
    Input: \(x_{0}\).
    for \(k \leftarrow 1,2, \cdots\) do
        Set \(x_{k+1} \leftarrow \Phi_{f}\left(x_{k},\left(x_{i}, f\left(x_{i}\right), \nabla f\left(x_{i}\right)\right)_{i \in[k]}\right)\).
    end
```

```
Algorithm 31: Gonzaga-Karas's variant of Accelerated Gradient Descent
    Input: \(x_{1}\).
    \(\gamma=2 \alpha, v_{0}=x_{0}\) and \(y_{0}=x_{0}\).
    for \(k \leftarrow 1,2, \cdots\) do
        \(y_{k} \leftarrow \Phi_{f}\left(y_{k-1}\right)\).
        \(x_{k+1}=\) line_search \(\left(y_{k},-\nabla f\left(y_{k}\right)\right)\).
        if \(\alpha \geq \gamma / 1.02\) and we are using first order oracle then \(\alpha=\gamma / 2 .\left(^{*}\right)\);
        if \(\alpha \geq \frac{\left\|\nabla f\left(y_{k}\right)\right\|^{2}}{2\left(f\left(y_{k}\right)-f\left(x_{k+1}\right)\right)}\) then \(\alpha=\frac{\left\|\nabla f\left(y_{k}\right)\right\|^{2}}{20\left(f\left(y_{k}\right)-f\left(x_{k+1}\right)\right)}\);
        \(G=\gamma\left(\frac{\alpha}{2}\left\|v_{k}-y_{k}\right\|^{2}+\left\langle\nabla f\left(y_{k}\right), v_{k}-y_{k}\right\rangle\right)\).
        \(A=G+\frac{1}{2}\left\|\nabla f\left(y_{k}\right)\right\|^{2}+(\alpha-\gamma)\left(f\left(x_{k}\right)-f\left(y_{k}\right)\right)\).
        \(B=(\alpha-\gamma)\left(f\left(x_{k+1}\right)-f\left(x_{k}\right)\right)-\gamma\left(f\left(y_{k}\right)-f\left(x_{k}\right)\right)-G\).
        \(C=\gamma\left(f\left(x_{k+1}\right)-f\left(x_{k}\right)\right)\).
        \(\beta=\frac{-B+\sqrt{B^{2}-4 A C}}{2 A}, \gamma=(1-\beta) \gamma+\beta \alpha\).
        \(v_{k+1}=\frac{1}{\gamma}\left((1-\beta) \gamma v_{k}+\beta\left(\alpha y_{k}-\nabla f\left(y_{k}\right)\right)\right.\).
    end
```


## ■ 11.5.1 Details of Implementations

The first algorithm we implement is the $\emptyset+$ algorithm which simply repeatedly call the politician. As we will see, this algorithm is great for non smooth problems but not competitive for smooth problems.

The second algorithm we implement is the accelerated gradient descent proposed by Gonzaga and Karas [108]. This algorithm uses line search to learn the the smoothness parameter and strong convexity parameter, see Algorithm 31. We disable the line $\left({ }^{*}\right)$ in the algorithm if $\Phi_{f}$ is a politician instead of an oracle because $\gamma \geq \alpha$ does not hold for the strong convexity parameter $\alpha$ if $\Phi_{f}$ is not an oracle.

The third algorithm we implemented is the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm. This algorithm uses the gradients to reconstruct the Hessian and use it to approximate Newton's method, see Algorithm 32. We note that another natural way to employ the politician with BFGS is to set $x_{k+1}=$ line_search $\left(\Phi_{f}\left(x_{k}\right), p\right)$ and this runs faster in practice; however, this algorithm computes two gradients per iteration (namely $\nabla f\left(x_{k}\right)$ and $\nabla f\left(\Phi_{f}\left(x_{k}\right)\right)$ ) while we restrict ourselves to algorithms which compute one gradient per iteration.

## $\square$ 11.5.2 Quadratic Function

We consider the function

$$
\begin{equation*}
f(x)=(x-c)^{\top} D(x-c) \tag{11.3}
\end{equation*}
$$

where $D$ is a diagonal matrix with entries uniformly sampled from $[0,1]$ and $c$ is a random vector with entries uniformly sampled from the normal distribution $N(0,1)$. Since this is a quadratic function, CG, BFGS and BFGS + are equivalent and optimal, namely, they output the minimum point in the

```
Algorithm 32: BFGS
    Input: \(x_{1}\).
    for \(k \leftarrow 1,2, \cdots\) do
        \(p=-\nabla f\left(x_{k}\right)\).
        for \(i \leftarrow k-1, \cdots, 1\) do
            \(\alpha_{i} \leftarrow\left\langle s_{i}, p\right\rangle /\left\langle s_{i}, y_{i}\right\rangle\).
            \(p=p-\alpha_{i} y_{i}\).
        end
        \(p=\left\langle s_{k-1}, y_{k-1}\right\rangle /\left\langle y_{k-1}, y_{k-1}\right\rangle p\).
        for \(i \leftarrow 1, \cdots, k-1\) do
            \(\beta_{i} \leftarrow\left\langle y_{i}, p\right\rangle /\left\langle s_{i}, y_{i}\right\rangle\).
            \(p=p+\left(\alpha_{i}-\beta_{i}\right) y_{i}\).
        end
        \(x_{k+1}=\Phi_{f}\left(\right.\) line_search \(\left.\left(x_{k}, p\right)\right)\).
        \(s_{k}=x_{k+1}-x_{k}, y_{k}=\nabla f\left(x_{k+1}\right)-\nabla f\left(x_{k}\right)\).
    end
```

span of all previous gradients.

## ■ 11.5.3 Variant of Nesterov's Worst Function

[206] introduced the function

$$
f(x)=(1-x[1])^{2}+\sum_{k=1}^{n-1}(x[k]-x[k+1])^{2}
$$

and used it to give a lower bound for all first-order methods. To distinguish the performance between CG, BFGS and BFGS+, we consider the following non-quadratic variant

$$
\begin{equation*}
f(x)=g(1-x[1])+\sum_{k=1}^{n-1} g(x[k]-x[k+1]) \tag{11.4}
\end{equation*}
$$

for some function $g$ to be defined. If we pick $g(x)=|x|$ then all first order methods takes at least $n$ iterations to minimize $f$ exactly. On the other hand with $g(x)=\max (|x|-0.1,0)$ one of the minimizer of $f$ is $\left(1, \frac{9}{10}, \frac{8}{10}, \cdots, \frac{1}{10}, 0,0, \cdots, 0\right)$, and thus it takes at least 11 iterations for first order methods to minimize $f$ in this case. We "regularize" the situation a bit and consider the function

$$
g(x)= \begin{cases}\sqrt{(x-0.1)^{2}+0.001^{2}}-0.001 & \text { if } x \geq 0.1 \\ \sqrt{(x+0.1)^{2}+0.001^{2}}-0.001 & \text { if } x \leq-0.1 \\ 0 & \text { otherwise }\end{cases}
$$

Since this function is far from quadratic, our algorithms ( $\emptyset+$, GK+, BFGS+) converge much faster. This is thus a nice example where the geometric politician helps a lot because the underlying dimension of the problem is small.


Figure 11-3: Comparison of first-order methods for the function (11.3) with $n=10000$.

## ■ 11.5.4 Binary Regression with Smoothed Hinge Loss

We consider the binary classification problem on the datasets from [48]. The problem is to minimize the regularized empirical risk:

$$
\begin{equation*}
f_{t}(x)=\frac{1}{n} \sum_{i=1}^{n} \varphi_{t}\left(b_{i} a_{i}^{T} x\right)+\frac{\lambda}{2}|x|^{2} \tag{11.5}
\end{equation*}
$$

where $a_{i} \in \mathbb{R}^{d}, b_{i} \in \mathbb{R}$ are given by the datasets, $\lambda$ is the regularization coefficient, $\varphi_{t}$ is the smoothed hinge loss defined by

$$
\varphi_{t}(z)= \begin{cases}0 & \text { if } z \leq-1 \\ z+1-\frac{t}{2} & \text { if } z \geq-1+t \\ \frac{1}{2 t}(z+1)^{2} & \text { otherwise }\end{cases}
$$

and $t$ is the smoothness parameter. The usual choice for $t$ is 1 , here we test both $t=1$ and $t=10^{-4}$. The latter case is to test how well the algorithms perform when the function is non-smooth.

We note that for this problem it would be natural to compare ourselves with SGD (stochastic gradient descent) or more refined stochastic algorithms such as SAG [159] or SVRG [129]. However since the focus of this chapter is on general black-box optimization we stick to comparing only to general methods. It is an interesting open problem to extend our algorithms to the stochastic setting, see Section 11.6.

In figures 11-5 and 11-6, we show the performance profile for problems in the LIBSVM datasets (and with different values for the regularization parameter $\lambda$ ). More precisely for a given algorithm we plot $x \in[1,10]$ versus the fraction of datasets that the algorithm can solve (up to a certain prespecified accuracy) in a number of iterations which is at most $x$ times the number of iterations of the best algorithm for this dataset. Figure 11-5 shows the case $t=1$ with the targeted accuracy $10^{-6}$; Figure 11-6 shows the case $t=10^{-4}$ with the targeted accuracy $10^{-3}$. We see that TFOCS is slower than SD for many problems, this is simply because SD uses the line search while TFOCS does not, and this makes a huge difference for simple problems. Among algorithms taking $O(n)$ time per iteration, CG and Geo perform the best, while for the $O(n k)$ algorithms we see that BFGS, BFGS+


Figure 11-4: Comparison of first-order methods for the function (11.4) with $n=10000$.
and GK+ perform the best. The gap in performance is particularly striking in the non-smooth case where BFGS+ is the fastest algorithm on almost all problems and all other methods (except GK+) are lagging far behind (for $20 \%$ of the problems all other methods take 10 times more iterations than BFGS+ and GK+).

Finally in figures 11-7 and 11-8 we test five algorithms on three specific datasets (respectively in the smooth and non-smooth case). In both figures we see that BFGS+ performs the best for all three datasets. BFGS performs second for smooth problems while GK+ performs second for nonsmooth problems.

## - 11.5.5 Summary

The experiments show that BFGS+ and BFGS perform the best among all methods for smooth test problems while BFGS+ and GK+ perform the best for nonsmooth test problems. The first phenomenon is due to the optimality of these algorithm for quadratic problems. We leave the explanation for the second phenomenon as an open problem. At least, the experiments show that this is not due to the geometric oracle itself since $\emptyset+$ is much slower, and this is not due to the original algorithm since GK performs much worse than GK+ for those problems. Overall these experiments are very promising for the geometric oracle as a replacement of quasi Newton method for non-smooth problems and as a general purpose solver due to its robustness.

## ■ 11.6 Discussion

First order methods generally involve only very basic operations at each step (addition, scalar multiplication). In last section, we formalize each step's operations (besides the gradient calculation) as


Figure 11-5: Performance profile on problem (11.5) with $t=1$ and $\lambda=10^{-4}, 10^{-5}, 10^{-6}, 10^{-7}, 10^{-8}$.
the work of the politician. We showed that the cost per step of an affine invariant politician $\psi(k)$ is negligible compared to the gradient calculation (which is $\Omega(n)$ ). This opens up a lot of possibilities: instead of basic addition or scalar multiplication one can imagine computing a center of gravity, solving a linear program, or even searching over an exponential space (indeed, say $k<30$ and $n>10^{10}$, then $\left.2^{k}<n\right)$. Our experiments demonstrate the effectiveness of this strategy. On the other hand from a theoretical point of view a lot remains to be done. For example one can prove results of the following flavor:

Theorem 11.6.1. Let $f$ such that $\alpha \mathrm{I} \preceq \nabla^{2} f(x) \preceq \beta \mathrm{I}, \forall x \in \mathbb{R}^{n}$ and let $\kappa=\beta / \alpha$. Suppose that in the Geometric Politician we replace the volumetric center by the center of gravity or the center of the John ellipsoid. Let $y_{k}$ be the output of the $k^{\text {th }}$ step of $S D+$ with some initial point $x_{0}$. Then, we have that

$$
f\left(y_{k}\right)-f\left(x^{*}\right) \leq \kappa\left(1-\frac{1}{\Theta(\min (n \log (\kappa), \kappa))}\right)^{k}\left(f\left(x_{0}\right)-f\left(x^{*}\right)\right)
$$

and

$$
f\left(y_{k}\right)-f\left(x^{*}\right) \leq \frac{2 \beta R^{2}}{k+4}
$$

where $R=\max _{f(x) \leq f\left(x_{0}\right)}\left\|x-x^{*}\right\|$.
This claim says that, up to a logarithmic factor, $\mathrm{SD}+$ enjoys simultaneously the incremental progress of gradient descent and the geometrical progress of cutting plane methods. There are three caveats in this claim:

- We use the center of gravity or the center of the John ellipsoid instead of the volumetric center. Note however that it is well-known that the volumetric center is usually more difficult to analyze, [253, 20].


Figure 11-6: Performance profile on problem (11.5) with $t=10^{-4}$ and $\lambda=10^{-4}, 10^{-5}, 10^{-6}, 10^{-7}, 10^{-8}$.

- The extraneous $\log (\kappa)$ comes from the number of potential restart when we decrease $\alpha$. Is there a better way to learn $\alpha$ that would not incur this additional logarithmic term?
- Theorem 11.3.1 shows essentially that one can combine the ellipsoid method with gradient descent to achieve the optimal $1-\sqrt{1 / \kappa}$ rate. Can we prove such a result for $\mathrm{SD}+$ ?

The geometric politician could be refined in many ways. Here are two simple questions that we leave for future work:

- One can think that gradient descent stores 1 gradient information, accelerated gradient descent stores 2 gradient information, and our method stores all past gradient information. We believe that neither 1,2 nor all is the correct answer. Instead, the algorithm should dynamically decide the number of gradients to store based on the size of its memory, the cost of computing gradients, and the information each gradient reveals.
- Is there a stochastic version of our algorithm? How well would such a method compare with state of the art stochastic algorithms such as SAG [159] and SVRG [129]?


## ■ 11.7 Appendix: Convergence of SD+

Let $f$ such that $\alpha \mathrm{I} \preceq \nabla^{2} f(x) \preceq \beta \mathrm{I}$ for all $x \in \mathbb{R}^{n}$ and let $\kappa=\beta / \alpha$. Let $y_{k}$ be the output of the $k^{t h}$ step of SD+ (where the volumetric center is replaced by the center of gravity or the center of the John ellipsoid) with some initial point $x_{0}$. We prove two rates of convergence for $\mathrm{SD}+$, one with the condition number $\kappa$, and one with the ambient dimension $n$. We start by the former.


Figure 11-7: Comparison between Geo, CG, BFGS, GK+, BFGS+ on problem (11.5) with $t=1$ and $\lambda=$ $10^{-4}, 10^{-6}, 10^{-8}$.

Theorem 11.7.1. One has

$$
f\left(y_{k}\right)-f\left(x^{*}\right) \leq\left(1-\frac{1}{\kappa}\right)^{k}\left(f\left(x_{0}\right)-f\left(x^{*}\right)\right)
$$

and

$$
f\left(y_{k}\right)-f\left(x^{*}\right) \leq \frac{2 \beta r^{2}}{k+4}
$$

where $r=\max _{f(x) \leq f\left(x_{0}\right)}\left\|x-x^{*}\right\|$.

Proof. Let $\delta_{k}=f\left(y_{k}\right)-f\left(x^{*}\right)$. Since $f$ is $\alpha$-strongly convex we have that

$$
\delta_{k} \leq \frac{1}{2 \alpha}\left\|\nabla f\left(y_{k}\right)\right\|^{2}
$$

Due to the decrease guarantee of politicians and the line search in steepest descent, we have that $f\left(y_{k+1}\right) \leq f\left(x_{k+1}\right) \leq f\left(y_{k}\right)-\frac{1}{2 \beta}\left\|\nabla f\left(y_{k}\right)\right\|^{2}$ and hence

$$
\begin{equation*}
\delta_{k}-\delta_{k+1} \geq \frac{1}{2 \beta}\left\|\nabla f\left(y_{k}\right)\right\|^{2} \geq \frac{\delta_{k}}{\kappa} . \tag{11.6}
\end{equation*}
$$

Hence, we have $\delta_{k+1} \leq\left(1-\frac{1}{\kappa}\right) \delta_{k}$ and this gives the first inequality.
To obtain a rate independent of $\alpha$ we instead use the following estimate

$$
\delta_{k} \leq\left\langle\nabla f\left(y_{k}\right), y_{k}-x^{*}\right\rangle \leq\left\|\nabla f\left(y_{k}\right)\right\| \cdot\left\|y_{k}-x^{*}\right\| .
$$

Using the decrease guarantee of politicians and line search we have that $f\left(y_{k}\right) \leq f\left(x_{k}\right) \leq f\left(y_{k-1}\right) \leq$


Figure 11-8: Comparison between Geo, CG, BFGS, GK+, BFGS+ on problem (11.5) with $t=10^{-4}$ and $\lambda=10^{-4}, 10^{-6}, 10^{-8}$.
$\cdots \leq f\left(x_{0}\right)$, and thus by definition of $R$ :

$$
\left\|y_{k}-x^{*}\right\| \leq r
$$

Due to the line search in steepest descent again, we have that

$$
\delta_{k}-\delta_{k+1} \geq \frac{1}{2 \beta}\left\|\nabla f\left(y_{k}\right)\right\|^{2} \geq \frac{1}{2 \beta}\left(\frac{\delta_{k}}{r}\right)^{2}
$$

Since $\delta_{k} \geq \delta_{k+1}$, we have

$$
\frac{1}{\delta_{k+1}}-\frac{1}{\delta_{k}}=\frac{\delta_{k}-\delta_{k+1}}{\delta_{k} \delta_{k+1}} \geq \frac{\delta_{k}-\delta_{k+1}}{\delta_{k}^{2}} \geq \frac{1}{2 \beta r^{2}}
$$

So, by induction, we have that $\frac{1}{\delta_{k}} \geq \frac{1}{\delta_{0}}+\frac{k}{2 \beta r^{2}}$. Now, we note that

$$
\delta_{0} \leq\left\langle\nabla f\left(x^{*}\right), x_{0}-x^{*}\right\rangle+\frac{\beta}{2}\left\|x_{0}-x^{*}\right\|^{2} \leq \frac{\beta r^{2}}{2}
$$

Thus, we have that

$$
\delta_{k} \leq \frac{2 \beta r^{2}}{k+4}
$$

We now turn to the dimension dependent analysis of SD+. We first show a simple geometric result, namely that if an intersection of spheres has a "small" volume then the intersection must lie close close to the boundary of one of the spheres.

Lemma 11.7.2. Let $R=\cap_{i=1}^{k}\left\{x \in \mathbb{R}^{n}:\left\|x-c_{i}\right\| \leq r_{i}\right\}, D=\max _{i \in[k]} r_{i}$, and $\omega_{n}$ the volume of the unit ball in $\mathbb{R}^{n}$. Then, there exists $i \in[k]$ such that for all $x \in R$,

$$
\left\|x-c_{i}\right\|^{2} \geq r_{i}^{2}-24 k^{2}\left(\frac{\mathrm{vol} R}{D^{n} \omega_{n}}\right)^{1 / n} D^{2}
$$

Proof. Since $-\log \left(1-\|x\|^{2}\right)$ is a 1 -self concordant barrier function, $2 F_{R}$ is a $k$-self concordant function. Let $y$ be the minimizer of $F_{R}$. Let $E=\left\{x \in \mathbb{R}^{n}: x^{\top} \nabla^{2}\left(2 F_{R}\right)(y) x \leq 1\right\}$. Theorem 4.2.6 in [206] shows that

$$
\begin{equation*}
y+E \subset R \subset y+(k+2 \sqrt{k}) E \tag{11.7}
\end{equation*}
$$

In particular, this shows that $\operatorname{vol} E \leq \operatorname{vol} R$. We have that

$$
\left(\operatorname{det} \nabla^{2} F_{R}(y)\right)^{1 / 2}=\frac{\omega_{n}}{2^{n / 2}} \frac{1}{\operatorname{vol} E} \geq \frac{\omega_{n}}{2^{n / 2}} \frac{1}{\operatorname{vol} R}
$$

By the AM-GM inequality, we have that

$$
\begin{equation*}
\frac{\operatorname{Tr} \nabla^{2} F_{R}(y)}{n} \geq \frac{1}{2}\left(\frac{\omega_{n}}{\operatorname{vol} R}\right)^{2 / n} \tag{11.8}
\end{equation*}
$$

By Proposition 11.4.3, we have that $\nabla^{2} F_{R}(y)=2 A^{\top} A+\lambda^{(1)} \mathrm{I}$ and hence,

$$
\begin{aligned}
\operatorname{Tr} \nabla^{2} F_{R}(y) & =2 \operatorname{Tr} A^{\top} A+n \lambda^{(1)} \\
& =2 \sum_{i=1}^{k} \frac{\left\|y-c_{i}\right\|^{2}}{\left(r_{i}^{2}-\left\|y-c_{i}\right\|^{2}\right)^{2}}+n \sum_{i=1}^{k} \frac{1}{r_{i}^{2}-\left\|y-c_{i}\right\|^{2}}
\end{aligned}
$$

Applying (11.8), we have that

$$
\frac{2}{n} \sum_{i=1}^{k} \frac{\left\|y-c_{i}\right\|^{2}}{\left(r_{i}^{2}-\left\|y-c_{i}\right\|^{2}\right)^{2}}+\sum_{i=1}^{k} \frac{1}{r_{i}^{2}-\left\|y-c_{i}\right\|^{2}} \geq \frac{1}{2}\left(\frac{\omega_{n}}{\operatorname{vol} R}\right)^{2 / n}
$$

So, there exists $i$ such that

$$
\frac{\left\|y-c_{i}\right\|^{2}}{\left(r_{i}^{2}-\left\|y-c_{i}\right\|^{2}\right)^{2}} \geq \frac{n}{8 k}\left(\frac{\omega_{n}}{\operatorname{vol} R}\right)^{2 / n} \text { or } \frac{1}{r_{i}^{2}-\left\|y-c_{i}\right\|^{2}} \geq \frac{1}{4 k}\left(\frac{\omega_{n}}{\operatorname{vol} R}\right)^{2 / n} .
$$

Using $\left\|y-c_{i}\right\| \leq r_{i} \leq D$ and $\operatorname{vol} R \leq D^{n} \omega_{n}$, we have that

$$
\begin{aligned}
r_{i}^{2}-\left\|y-c_{i}\right\|^{2} & \leq \max \left(\sqrt{\frac{8 k}{n}}\left(\frac{\mathrm{vol} R}{\omega_{n}}\right)^{1 / n}\left\|y-c_{i}\right\|, 4 k\left(\frac{\mathrm{vol} R}{\omega_{n}}\right)^{2 / n}\right) \\
& \leq 4 k\left(\frac{\mathrm{vol} R}{D^{n} \omega_{n}}\right)^{1 / n} D^{2}
\end{aligned}
$$

Therefore, the width of the ellipsoid $E$ in the direction $y-c_{i}$ is at most

$$
r_{i}-\left\|y-c_{i}\right\| \leq r_{i}-\sqrt{r_{i}^{2}-4 k\left(\frac{\operatorname{vol} R}{D^{n} \omega_{n}}\right)^{1 / n} D^{2}}
$$

The right hand side of (11.7) shows that, for all $x \in R$, we have

$$
\begin{aligned}
\left\|x-c_{i}\right\| & \geq r_{i}-(1+k+2 \sqrt{k})\left(r_{i}-\sqrt{r_{i}^{2}-4 k\left(\frac{\mathrm{vol} R}{D^{n} \omega_{n}}\right)^{1 / n} D^{2}}\right) \\
& \geq r_{i}-(1+k+2 \sqrt{k})\left(r_{i}-r_{i}\left(1-4 k\left(\frac{\mathrm{vol} R}{D^{n} \omega_{n}}\right)^{1 / n} \frac{D^{2}}{r_{i}^{2}}\right)\right) \\
& \geq\left[1-12 k^{2}\left(\frac{\mathrm{vol} R}{D^{n} \omega_{n}}\right)^{1 / n} \frac{D^{2}}{r_{i}^{2}}\right] r_{i} .
\end{aligned}
$$

Hence, we have

$$
\left\|x-c_{i}\right\|^{2} \geq\left[1-24 k^{2}\left(\frac{\mathrm{vol} R}{D^{n} \omega_{n}}\right)^{1 / n} \frac{D^{2}}{r_{i}^{2}}\right] r_{i}^{2} .
$$

Finally, equipped with the above geometrical result, we can bound the convergence of SD+ using the dimension $n$. We start with a lemma taking care of the adaptivity to the strong convexity in the geometric politician.
Lemma 11.7.3. In the first $k=\Theta\left(n \log \left(\frac{\kappa n}{\varepsilon}\right)\right)$ iterations, either $S D+$ restarts the estimate of the strong convexity or

$$
f\left(y_{k}\right)-f\left(x^{*}\right) \leq \varepsilon\left(f\left(x_{0}\right)-f\left(x^{*}\right)\right) .
$$

Proof. The decrease guarantee and the smoothness imply that

$$
\frac{\left\|\nabla f\left(y_{k}\right)\right\|^{2}}{2 \beta} \leq f\left(y_{k}\right)-f\left(x^{*}\right) \leq f\left(x_{0}\right)-f\left(x^{*}\right) .
$$

Therefore, all the spheres found by the geometric politician have radius squared at most $D^{2}$ where, denoting $\bar{\alpha}$ for the convexity upper bound the algorithm is currently using,

$$
D^{2}=\max _{k \geq 1} \frac{\left\|\nabla f\left(y_{k}\right)\right\|^{2}}{\bar{\alpha}^{2}} \leq \frac{2 \beta\left(f\left(x_{0}\right)-f\left(x^{*}\right)\right)}{\bar{\alpha}^{2}} .
$$

Lemma 11.7.2 shows that for any step $k$, there is $i \in[k]$ such that for all $x \in R_{k}$,

$$
\left\|x-c_{i}\right\|^{2} \geq r_{i}^{2}-\frac{48 \beta k^{2}}{\bar{\alpha}^{2}}\left(\frac{\operatorname{vol} R_{k}}{D^{n} \omega_{n}}\right)^{1 / n}\left(f\left(x_{0}\right)-f\left(x^{*}\right)\right) .
$$

Let $k=\Theta\left(n \log \left(\frac{\kappa n}{\varepsilon}\right)\right)$ and recall the discussion in Section 11.4.2 about the volume decrease of the geometric politician with the center of gravity (the same discussion applies to the John ellipsoid). We see that if the algorithm does not restart $\bar{\alpha}$ within the first $k$ iterations then we have

$$
\frac{\operatorname{vol} R_{k}}{D^{n} \omega_{n}}=\left(O\left(\frac{\varepsilon}{\kappa^{2} k^{2}}\right)\right)^{n},
$$

and hence (for an appropriate numerical constant in $k$ )

$$
\begin{equation*}
\left\|x-c_{i}\right\|^{2} \geq r_{i}^{2}-\frac{\varepsilon\left(f\left(x_{0}\right)-f\left(x^{*}\right)\right)}{\bar{\alpha} \kappa} . \tag{11.9}
\end{equation*}
$$

Recall from (11.6) that

$$
f\left(y_{k+1}\right) \leq f\left(y_{k}\right)-\frac{f\left(y_{k}\right)-f\left(x^{*}\right)}{\kappa}
$$

and therefore we have (by the improvement of the previous balls):

$$
R_{k+1} \subset\left\{\left\|x-c_{i}\right\|^{2} \leq r_{i}^{2}-\frac{2\left(f\left(y_{k}\right)-f\left(x^{*}\right)\right)}{\bar{\alpha} \kappa}\right\} \cap R_{k}
$$

However, from (11.9), we know that either the above intersection is empty or $f\left(y_{k}\right)-f\left(x^{*}\right)<\varepsilon\left(f\left(x_{0}\right)-\right.$ $f\left(x^{*}\right)$ ). This proves the statement.

Theorem 11.7.4. We have that

$$
f\left(y_{k}\right)-f\left(x^{*}\right) \leq \kappa\left(1-\frac{1}{\Theta(n \log (\kappa))}\right)^{k}\left(f\left(x_{0}\right)-f\left(x^{*}\right)\right)
$$

Proof. If $\kappa<n$, the statements follows from Theorem 11.7.1. Hence, we can assume $\kappa \geq n$.
Set $T=\Theta\left(n \log \left(\frac{n \kappa}{\varepsilon}\right) \log (\kappa)\right)$, Lemma 11.7.3 shows that for every $\Theta\left(n \log \left(\frac{n \kappa}{\varepsilon}\right)\right)$ iteration, the algorithm either finds $y$ such that

$$
f(y)-f\left(x^{*}\right) \leq \varepsilon\left(f\left(x_{0}\right)-f\left(x^{*}\right)\right)
$$

or decreases $\alpha_{k}$ by a constant where $\alpha_{k}$ is the convexity upper bound the algorithm is using at $k^{t h}$ iteration. Note that $\alpha_{1} \leq \beta$ because of the line search, and thus the algorithm can restart $\alpha_{k}$ at most $\log (\kappa)$ many times. Hence, after $T$ iterations, we must have

$$
f\left(y_{T}\right)-f\left(x^{*}\right) \leq \varepsilon\left(f\left(x_{0}\right)-f\left(x^{*}\right)\right)
$$

thus concluding the proof.

Part III

## Collapsing



# Undirected MaxFlow In Almost-Linear Time 

## ■ 12.1 Introduction

In this chapter, we study the maximum flow problem. Given a graph $G=(V, E)$ in which each edge $e \in E$ is assigned a nonnegative capacity $\boldsymbol{\mu}_{e}$, the maximum $s-t$ flow problem asks us to find a flow $\boldsymbol{f}$ that routes as much flow as possible from a source vertex $s$ to a sink vertex $t$ while sending at most $\boldsymbol{\mu}_{e}$ units of flow over each edge $e$. Its generalization, the maximum concurrent multicommodity flow problem, supplies $k$ source-sink pairs $\left(s_{i}, t_{i}\right)$ and asks for the maximum $\alpha$ such that we may simultaneously route $\alpha$ units of flow between each source-sink pair. That is, it asks us to find flows $\boldsymbol{f}_{1}, \ldots, \boldsymbol{f}_{k}$ (which we think of as corresponding to $k$ different commodities) such that $\boldsymbol{f}_{i}$ sends $\alpha$ units of flow from $s_{i}$ to $t_{i}$, and $\sum_{i}\left|\boldsymbol{f}_{i}(e)\right| \leq \boldsymbol{\mu}_{e}$ for all $e \in E$.

These problems lie at the core of graph algorithms and combinatorial optimization, and they have been extensively studied over the past 60 years [232, 4]. They have found a wide range of theoretical and practical applications [5], and they are widely used as key subroutines in other algorithms (see [18, 233]).

In this chapter, we introduce a new framework for approximately solving flow problems in capacitated, undirected graphs and apply it to provide asymptotically faster algorithms for the maximum $s$ - $t$ flow and maximum concurrent multicommodity flow problems. For graphs with $n$ vertices and $m$ edges, it allows us to find an $\varepsilon$-approximately maximum $s$ - $t$ flows in time $O\left(m^{1+o(1)} \varepsilon^{-2}\right)$, improving on the previous best bound of $\widetilde{O}\left(m n^{1 / 3} \operatorname{poly}(1 / \varepsilon)\right)[51]$. Applying the same framework in the multicommodity setting solves a maximum concurrent multicommodity flow problem with $k$ commodities in $O\left(m^{1+o(1)} \varepsilon^{-2} k^{2}\right)$ time, improving on the existing bound of $\widetilde{O}\left(m^{4 / 3} \operatorname{poly}\left(k, \varepsilon^{-1}\right)\right)$ [140].

We believe that both our general framework and several of the pieces necessary for its present instantiation are of independent interest, and we hope that they will find other applications. These include:

- a non-Euclidean generalization of gradient descent, bounds on its performance, and a way to use this to reduce approximate maximum flow and maximum concurrent flow to oblivious routing;
- the definition and efficient construction of flow sparsifiers; and
- the construction of a new oblivious routing scheme that can be implemented extremely efficiently.

We have aimed to make our algorithm fairly modular and have thus occasionally worked in slightly more generality than is strictly necessary for the problem at hand. This has slightly increased the length of the exposition, but we believe that it clarifies the high-level structure of the argument, and it will hopefully facilitate the application of these tools in other settings.

## ■ 12.1.1 Related Work

For the first several decades of its study, the fastest algorithms for the maximum flow problem were essentially all deterministic algorithms based on combinatorial techniques, such as augmenting paths, blocking flows, preflows, and the push-relabel method. These culminated in the work of Goldberg and Rao [104], which computes exact maximum flows in time $O\left(\min \left(n^{2 / 3}, m^{1 / 2}\right) \log \left(n^{2} / m\right) \log U\right)$ on graphs with edge weights in $\{0, \ldots, U\}$. We refer the reader to [104] for a survey of these results.

More recently, a collection of new techniques based on randomization, spectral graph theory and numerical linear algebra, graph decompositions and embeddings, and iterative methods for convex optimization have emerged. These have allowed researchers to provide better provable algorithms for a wide range of flow and cut problems, particularly when one aims to obtain approximately optimal solutions on undirected graphs.

Our algorithm draws extensively on the intellectual heritage established by these works. In this section, we will briefly review some of the previous advances that inform our algorithm. We do not give a comprehensive review of the literature, but instead aim to provide a high-level view of the main tools that motivated the present work, along with the limitations of these tools that had to be overcome. For simplicity of exposition, we primarily focus on the maximum $s$ - $t$ flow problem for the remainder of the introduction.

Sparsification In [33], Benczur and Karger showed how to efficiently approximate any graph $G$ with a sparse graph $G^{\prime}$ on the same vertex set. To do this, they compute a carefully chosen probability $p_{e}$ for each $e \in E$, sample each edge $e$ with probability $p_{e}$, and include $e$ in $G^{\prime}$ with its weight increased by a factor of $1 / p_{e}$ if it is sampled. Using this, they obtain, in nearly linear time, a graph $G^{\prime}$ with $O\left(n \log n / \varepsilon^{2}\right)$ edges such that the total weight of the edges crossing any cut in $G^{\prime}$ is within a multiplicative factor of $1 \pm \varepsilon$ of the weight crossing the corresponding cut in $G$. In particular, the Max-Flow Min-Cut Theorem implies that the value of the maximum flow on $G^{\prime}$ is within a factor of $1 \pm \varepsilon$ of that of $G$.

This is an extremely effective tool for approximately solving cut problems on a dense graph $G$, since one can simply solve the corresponding problem on the sparsified graph $G^{\prime}$. However, while this means that one can approximately compute the value of the maximum $s$ - $t$ flow on $G$ by solving the problem on $G^{\prime}$, it is not known how to use the maximum $s$ - $t$ flow on $G^{\prime}$ to obtain an actual approximately maximum flow on $G$. Intuitively, this is because the weights of edges included in $G^{\prime}$ are larger than they were in $G$, and the sampling argument does not provide any guidance about how to route flows over these edges in the original graph $G$.

Iterative algorithms based on linear systems and electrical flows In 2010, Christiano et al.[51] described a new linear algebraic approach to the problem that found $\varepsilon$-approximately maximum $s$ - $t$ flows in time $\widetilde{O}\left(m n^{1 / 3} \operatorname{poly}(1 / \varepsilon)\right)$. They treated the edges of $G$ as electrical resistors and then computed the electrical flow that would result from sending electrical current from $s$ to $t$ in the corresponding circuit. They showed that these flows can be computed in nearly-linear time using fast Laplacian linear system solvers [147, 146, 141, 163], which we further discuss below. The electrical flow obeys the flow conservation constraints, but it could violate the capacity constraints. They then adjusted the resistances of edges to penalize the edges that were flowing too much current and repeated the process. Kelner, Miller, and Peng [140] later showed how to use more general objects that they called quadratically coupled flows to use a similar approach to solve the maximum concurrent multicommodity flow problem in time $\widetilde{O}\left(m^{4 / 3} \operatorname{poly}(k, 1 / \varepsilon)\right)$.

Following this, Lee, Rao, and Srivastava [161] proposed another iterative algorithm that uses electrical flows, but in a way that was substantially different than in [51]. Instead of adjusting the
resistances of the edges in each iteration to correct overflowing edges, they keep the resistances the same but compute a new electrical flow to reroute the excess current. They explain how to interpret this as gradient descent in a certain space, from which a standard analysis would give an algorithm that runs in time $\widetilde{O}\left(m^{3 / 2} \operatorname{poly}(1 / \varepsilon)\right)$. By replacing the standard gradient descent step with Nesterov's accelerated gradient descent method [206] and using a regularizer to make the penalty function smoother, they obtain an algorithm that runs in time $\widetilde{O}\left(m n^{1 / 3} \operatorname{poly}(1 / \varepsilon)\right)$ in unweighted graphs.

In all of these algorithms, the superlinear running times arise from an intrinsic $\Theta(\sqrt{m})$ factor introduced by using electrical flows, which minimize an $\ell_{2}$ objective function, to approximate the maximum congestion, which is an $\ell_{\infty}$ quantity.

Fast solvers for Laplacian linear systems In their breakthrough paper [243], Spielman and Teng showed how to solve Laplacian systems in nearly-linear time. (This was later sped up and simplified by Koutis, Miller, and Peng [147, 146] and Kelner, Orecchia, Sidford, and Zhu [141].) Their algorithm worked by showing how to approximate the Laplacian $\mathcal{L}_{G}$ of a graph $G$ with the Laplacian $\mathcal{L}_{H}$ of a much simpler graph $H$ such that one could use the ability to solve linear systems in $\mathcal{L}_{H}$ to accelerate the solution of a linear system in $\mathcal{L}_{G}$. They then applied this recursively to solve the linear systems in $\mathcal{L}_{H}$. In addition to providing the electrical flow primitive used by the algorithms described above, the structure of their recursive sequence of graph simplifications provides the motivating framework for much of the technical content of our oblivious routing construction.

Oblivious routing In an oblivious routing scheme, one specifies a linear operator taking any demand vector to a flow routing these demands over the edges of $G$. Given a collection of demand vectors, one can produce a multicommodity flow meeting these demands by routing each demand vector using this pre-specified operator, independently of the others. The competitive ratio of such an operator is the worst possible ratio between the congestion incurred by a set of demands in this scheme and the congestion of the best multicommodity flow routing these demands.

In [225], Räcke showed how to construct an oblivious routing scheme with a competitive ratio of $O(\log n)$. His construction worked by providing a probability distribution over trees $T_{i}$ such that $G$ embeds into each $T_{i}$ with congestion at most 1 , and such that the corresponding convex combination of trees embeds into $G$ with congestion $O(\log n)$. In a sense, one can view this as showing how to approximate $G$ by a probability distribution over trees. Using this, he was able to show how to obtain polylogarithmic approximations for a variety of cut and flow problems, given only the ability to solve these problems on trees.

We note that such an oblivious routing scheme clearly yields a logarithmic approximation to the maximum flow and maximum concurrent multicommodity flow problems. However, Räcke's construction took time substantially superlinear time, making it too slow to be useful for computing approximately maximum flows. Furthermore, it only gives a logarithmic approximation, and it is not clear how to use this a small number of times to reduce the error to a multiplicative $\varepsilon$.

In a later paper [182], Madry applied a recursive technique similar to the one employed by Spielman and Teng in their Laplacian solver to accelerate many of the applications of Räcke's construction at the cost of a worse approximation ratio. Using this, he obtained almost-linear-time polylogarithmic approximation algorithms for a wide variety of cut problems.

Unfortunately, his algorithm made extensive use of sparsification, which, for the previously mentioned reasons, made it unable to solve the corresponding flow problems. This meant that, while it could use flow-cut duality to find a polylogarithmic approximation of the value of a maximum flow, it could not construct a corresponding flow or repeatedly apply such a procedure a small number of times to decrease the error to a multiplicative $\varepsilon$.

In simultaneous, independent work [233], Jonah Sherman used somewhat different techniques to find another almost-linear-time algorithm for the (single-commodity) maximum flow problem. His approach is essentially dual to ours: Our algorithm maintains a flow that routes the given demands throughout its execution and iteratively works to improve its congestion. Our main technical tools thus consist of efficient methods for finding ways to route flow in the graph while maintaining flow conservation. Sherman, on the other hand, maintains a flow that does not route the given demands, along with a bound on the congestion required to route the excess flow at the vertices. He then uses this to iteratively work towards achieving flow conservation. (In a sense, our algorithm is more in the spirit of augmenting paths, whereas his is more like preflow-push.) As such, his main technical tools are efficient methods for producing dual objects that give congestion bounds. Objects meeting many of his requirements were given in the work of Madry [182] (whereas there were no previous constructions of flow-based analogues, requiring us to start from scratch); leveraging these allows him to avoid some of the technical complexity required by our approach. We believe that these papers nicely complement each other, and we enthusiastically refer the reader to Sherman's paper.

## ■ 12.1.2 Our Approach

In this section, we give a high-level description of how we overcome the obstacles described in the previous section. For simplicity, we suppose for the remainder of this introduction that all edges have capacity 1.

The problem is thus to send as many units of flow as possible from $s$ to $t$ without sending more than one unit over any edge. It will be more convenient for us to work with an equivalent congestion minimization problem, where we try to find the unit $s$ - $t$ flow $\boldsymbol{f}$ (i.e., a flow sending one unit from $s$ to $t$ ) that minimizes $\|f\|_{\infty}=\max _{e}\left|\boldsymbol{f}_{e}\right|$. If we begin with some initial unit $s-t$ flow $\boldsymbol{f}_{0}$, the goal will be thus be to find the circulation $\boldsymbol{c}$ to add to $\boldsymbol{f}_{0}$ that minimizes $\left\|\boldsymbol{f}_{0}+\boldsymbol{c}\right\|_{\infty}$.

We give an iterative algorithm to approximately find such a c. There are $2^{O(\sqrt{\log n \log \log n})} / \varepsilon^{2}$ iterations, each of which adds a circulation to the present flow and runs in $m \cdot 2^{O(\sqrt{\log n \log \log n})}$ time. Constructing this scheme consists of two main parts: an iterative scheme that reduces the problem to the construction of a projection matrix with certain properties; and the construction of such an operator.

## The iterative scheme: Non-Euclidean gradient descent

The simplest way to improve the flow would be to just perform gradient descent on the maximum congestion of an edge. There are two problems with this:

The first problem is that gradient descent depends on having a smoothly varying gradient, but the infinity norm is very far from smooth. This is easily remedied by a standard technique: we replace the infinity norm with a smoother "soft max" function. Doing this would lead to an update that would be a linear projection onto the space of circulations. This could be computed using an electrical flow, and the resulting algorithm would be very similar to the unaccelerated gradient descent algorithm in [161].

The more serious problem is the one discussed in the previous section: the difference between $\ell_{2}$ and $\ell_{\infty}$. Gradient steps choose a direction by optimizing a local approximation of the objective function over a sphere, whereas the $\ell_{\infty}$ constraint asks us to optimize over a cube. The difference between the size of the largest sphere inside a cube and the smallest sphere containing it gives rise to an inherent $O(\sqrt{m})$ in the number of iterations, unless one can somehow exploit additional structure in the problem.

To deal with this, we introduce and analyze a non-Euclidean variant of gradient descent that
operates with respect to an arbitrary norm. ${ }^{1}$ Rather than choosing the direction by optimizing a local linearization of the objective function over the sphere, it performs an optimization over the unit ball in the given norm. By taking this norm to be $\ell_{\infty}$ instead of $\ell_{2}$, we are able to obtain a much smaller bound on the number of iterations, albeit at the expense of having to solve a nonlinear minimization problem at every step. The number of iterations required by the gradient descent method depends on how quickly the gradient can change over balls in the norm we are using, which we express in terms of the Lipschitz constant of the gradient in the chosen norm.

To apply this to our problem, we write flows meeting our demands as $f_{0}+\boldsymbol{c}$, as described above. We then need a parametrization of the space of circulations so that the objective function (after being smoothed using soft max) has a good bound on its Lipschitz constant. Similarly to what occurs in [141], this comes down to finding a good linear representation of the space of circulations, which we show amounts in the present setting to finding a matrix that projects into the space of circulations while meetings certain norm bounds.

## Constructing a projection matrix

This reduces our problem to the construction of such a projection matrix. A simple calculation shows that any linear oblivious routing scheme $A$ with a good competitive ratio gives rise to a projection matrix with the desired properties, and thus leads to an iterative algorithm that converges in a small number of iterations. Each of these iterations performs a matrix-vector multiplication with both $\boldsymbol{A}$ and $\boldsymbol{A}^{T}$.

Intuitively, this is letting us replace the electrical flows used in previous algorithms with the flows given by an oblivious routing scheme. Since the oblivious routing scheme was constructed to meet $\ell_{\infty}$ guarantees, while the electrical flow could only obtain such guarantees by relating $\ell_{2}$ to $\ell_{\infty}$, it is quite reasonable that we should expect this to lead to a better iterative algorithm.

However, the computation involved in existing oblivious routing schemes is not fast enough to be used in this setting. Our task thus becomes constructing an oblivious routing scheme that we can compute and work with very efficiently. We do this with a recursive construction that reduces oblivious routing in a graph to oblivious routing in various successively simpler graphs.

To this end, we show that if $G$ can be embedded with low congestion into $H$ (existentially), and $H$ can be embedded with low congestion into $G$ efficiently, one can use an oblivious routing on $H$ to obtain an oblivious routing on $G$. The crucial difference between the simplification operations we perform here and those in previous papers (e.g., in the work of Benczur-Karger [33] and Madry [182]) is that ours are accompanied by such embeddings, which enables us to transfer flows from the simpler graphs to the more complicated ones.

We construct our routing scheme by recursively composing two types of reductions, each of which we show how to implement without incurring a large increase in the competitive ratio:

- Vertex elimination This shows how to efficiently reduce oblivious routing on a graph $G=$ $(V, E)$ to routing on $t$ graphs with roughly $\widetilde{O}(|E| / t)$ vertices. To do this, we show how to efficiently embed $G$ into $t$ simpler graphs, each consisting of a tree plus a subgraph supported on roughly $\widetilde{O}(|E| / t)$ vertices. This follows easily from a careful reading of Madry's paper [182]. We then show that routing on such a graph can be reduced to routing on a graph with at most $\widetilde{O}(|E| / t)$ vertices by collapsing paths and eliminating leaves.

[^22]- Flow sparsification This allows us to efficiently reduce oblivious routing on an arbitrary graph to oblivious routing on a graph with $\widetilde{O}(|V|)$ edges, which we call a flow sparsifier.
To construct flow sparsifiers, we use local partitioning to decompose the graph into well-connected clusters that contain many of the original edges. (These clusters are not quite expanders, but they are contained in graphs with good expansion in a manner that is sufficient for our purposes.) We then sparsify these clusters using standard techniques and then show that we can embed the sparse graph back into the original graph using electrical flows. If the graph was originally dense, this results in a sparser graph, and we can recurse on the result. While the implementation of these steps is somewhat different, the outline of this construction parallels Spielman and Teng's approach to the construction of spectral sparsifiers [243, 244].

Combining these two reductions recursively yields an efficient oblivious routing scheme, and thus an algorithm for the maximum flow problem.

Finally, we show that the same framework can be applied to the maximum concurrent multicommodity flow problem. While the norm and regularization change, the structure of the argument and the construction of the oblivious routing scheme go through without requiring substantial modification.

## ■ 12.2 Preliminaries

Here, we define some notations that is used throughout in this chapter.
General Notation: We let $\boldsymbol{I}$ be the identity matrix and $\boldsymbol{I}_{a \rightarrow b} \in \mathbb{R}^{b \times a}$ denote the matrix such that for all $i \leq \min \{a, b\}$ we have $\boldsymbol{I}_{i i}=1$ and $\boldsymbol{I}_{i j}=0$ otherwise.

Graphs: Throughout this chapter we let $G=(V, E, \boldsymbol{\mu})$ denote an undirected capacitated graph with $n=|V|$ vertices, $m=|E|$ edges, and non-negative capacities $\boldsymbol{\mu} \in \mathbb{R}^{E}$. We let $w_{e} \geq 0$ denote the weight of an edge and let $r_{e} \stackrel{\text { def }}{=} 1 / w_{e}$ denote the resistance of an edge. Here we make no connection between $\mu_{e}$ and $r_{e}$; we fix their relationship later.

Sizes: For all $a \in V$ we let $d_{a} \stackrel{\text { def }}{=} \sum_{\{a, b\}} w_{a, b}$ denote the (weighted) degree of vertex $a$ and we let $\operatorname{deg}(a) \stackrel{\text { def }}{=} \mid\{e \in E \mid e=\{a, b\}$ for some $b \in V\} \mid$ denote its (combinatorial) degree. We let $\boldsymbol{D} \in \mathbb{R}^{V \times V}$ be the diagonal matrix where $\boldsymbol{D}_{a, a}=d_{a}$. Furthermore, for any vertex subset $S \subseteq V$ we define its volume by $\operatorname{vol}(S) \stackrel{\text { def }}{=} \sum_{a \in V} d_{a}$.

Cuts: For any vertex subset $S \subseteq V$ we denote the cut induced by $S$ by the edge subset

$$
\partial(S) \stackrel{\text { def }}{=}\{e \in E \mid e \nsubseteq S \text { and } e \nsubseteq E \backslash S\}
$$

and we denote the cost of $F \subseteq E$ by $w(F) \stackrel{\text { def }}{=} \sum_{e \in F} w_{e}$. We denote the conductance of $S \subseteq V$ by

$$
\Phi(S) \stackrel{\text { def }}{=} \frac{w(\partial(S))}{\min \{\operatorname{vol}(S), \operatorname{vol}(V-S)\}}
$$

and we denote the conductance of a graph by

$$
\Phi(G) \stackrel{\text { def }}{=} \min _{S \subseteq V: S \notin\{0, V\}} \phi(S)
$$

Subgraphs: For a graph $G=(V, E)$ and a vertex subset $S \subseteq V$ let $G(S)$ denote the subgraph of $G$ consisting of vertex set $S$ and all the edges of $E$ with both endpoints in $S$, i.e. $\{(a, b) \in E \mid a, b \in S\}$. When we we refer to a graph property, such as vol or $\Phi$, we use subscripts to specify which graph we are considering. For example $\operatorname{vol}_{G(S)}(A)$ denotes the volume of vertex set $A$ in the subgraph of $G$ induced by $S$.

Congestion: Thinking of edge vectors, $\boldsymbol{f} \in \mathbb{R}^{E}$, as flows we let the congestion ${ }^{2}$ of $\boldsymbol{f}$ be given by $\operatorname{cong}(\boldsymbol{f}) \stackrel{\text { def }}{=}\left\|\boldsymbol{U}^{-1} \boldsymbol{f}\right\|_{\infty}$. For any collection of flows $\left\{\boldsymbol{f}_{i}\right\}=\left\{\boldsymbol{f}_{1}, \ldots, \boldsymbol{f}_{k}\right\}$ we overload notation and let their total congestion be given by

$$
\operatorname{cong}\left(\left\{\boldsymbol{f}_{i}\right\}\right) \stackrel{\text { def }}{=}\left\|\boldsymbol{U}^{-1} \sum_{i}\left|\boldsymbol{f}_{i}\right|\right\|_{\infty}
$$

Demands and Multicommodity Flow: We call a vector $\chi \in \mathbb{R}^{V}$ a demand vector if it is the case that $\sum_{a \in V} \boldsymbol{\chi}(a)=0$ and we say $\boldsymbol{f} \in \mathbb{R}^{E}$ meets the demands if $\boldsymbol{B}^{T} \boldsymbol{f}=\boldsymbol{\chi}$. Given a set of demands $D=\left\{\chi_{1}, \ldots, \chi_{k}\right\}$, i.e. $\forall i \in[k], \sum_{a \in V} \chi_{i}(a)=0$, we denote the optimal low congestion routing of these demands as follows

$$
\operatorname{OPT}(D) \stackrel{\text { def }}{=} \min _{\left\{\boldsymbol{f}_{i}\right\} \in \mathbb{R}^{E}}:\left\{\boldsymbol{B}^{T} \boldsymbol{f}_{i}\right\}=\left\{\chi_{i}\right\}<10 \operatorname{cong}\left(\left\{\boldsymbol{f}_{i}\right\}\right)
$$

We call a set of flows $\left\{\boldsymbol{f}_{i}\right\}$ that meet demands $\left\{\boldsymbol{\chi}_{i}\right\}$, i.e. $\forall i, \boldsymbol{B}^{T} \boldsymbol{f}_{i}=\chi_{i}$, a multicommodity flow meeting the demands.

Running Time: For matrix $\boldsymbol{A}$, we let $\mathcal{T}(\boldsymbol{A})$ denote the maximum amount of time needed to apply $\boldsymbol{A}$ or $\boldsymbol{A}^{T}$ to a vector.

## ■ 12.3 Solving Max-Flow Using a Circulation Projection

## $\square$ 12.3.1 Gradient Descent

In this section, we discuss the gradient descent method for general norms. Let $\|\cdot\|: \mathbb{R}^{n} \rightarrow \mathbb{R}$ be an arbitrary norm on $\mathbb{R}^{n}$ and recall that the gradient of $f$ at $\boldsymbol{x}$ is defined to be the vector $\nabla f(\boldsymbol{x}) \in \mathbb{R}^{n}$ such that

$$
\begin{equation*}
f(\boldsymbol{y})=f(\boldsymbol{x})+\langle\nabla f(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\rangle+o(\|\boldsymbol{y}-\boldsymbol{x}\|) . \tag{12.1}
\end{equation*}
$$

The gradient descent method is a greedy minimization method that updates the current vector, $\boldsymbol{x}$, using the direction which minimizes $\langle\nabla f(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\rangle$. To analyze this method's performance, we need a tool to compare the improvement $\langle\nabla f(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\rangle$ with the step size, $\|\boldsymbol{y}-\boldsymbol{x}\|$, and the quantity, $\|\nabla f(\boldsymbol{x})\|$. For $\ell_{2}$ norm, this can be done by Cauchy Schwarz inequality and in general, we can define a new norm for $\nabla f(\boldsymbol{x})$ to make this happens. We call this the dual norm $\|\cdot\|^{*}$ defined as follows

$$
\|\boldsymbol{x}\|^{*} \stackrel{\text { def }}{=} \max _{\boldsymbol{y} \in \mathbb{R}^{n}:\|\boldsymbol{y}\| \leq 1}\langle\boldsymbol{y}, \boldsymbol{x}\rangle .
$$

Fact 12.8.4 shows that this definition indeed yields that $\langle\boldsymbol{y}, \boldsymbol{x}\rangle \leq\|\boldsymbol{y}\|^{*}\|\boldsymbol{x}\|$. Next, we define the fastest increasing direction $x^{\#}$, which is an arbitrary point satisfying the following

$$
\boldsymbol{x}^{\#} \stackrel{\text { def }}{=} \underset{\boldsymbol{s} \in \mathbb{R}}{\arg \max }\langle\boldsymbol{x}, \boldsymbol{s}\rangle-\frac{1}{2}\|\boldsymbol{s}\|^{2} .
$$

In the appendix, we provide some facts about $\|\cdot\|^{*}$ and $\boldsymbol{x}^{\#}$ that we will use in this section. Using the notations defined, the gradient descent method simply produces a sequence of $\boldsymbol{x}_{k}$ such that

$$
\boldsymbol{x}_{k+1}:=\boldsymbol{x}_{k}-t_{k}\left(\nabla f\left(\boldsymbol{x}_{k}\right)\right)^{\#}
$$

[^23]where $t_{k}$ is some chosen step size for iteration $k$. To determine what these step sizes should be we need some information about the smoothness of the function, in particular, the magnitude of the second order term in (12.1). The natural notion of smoothness for gradient descent is the Lipschitz constant of the gradient of $f$, that is the smallest constant $L$ such that
$$
\forall \boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{n} \quad: \quad\|\nabla f(\boldsymbol{x})-\nabla f(\boldsymbol{y})\|^{*} \leq L \cdot\|\boldsymbol{x}-\boldsymbol{y}\|
$$

In the appendix we provide an equivalent definition and a way to compute $L$, which is useful later.
Let $X^{*} \subseteq \mathbb{R}^{n}$ denote the set of optimal solutions to the unconstrained minimization problem $\min _{x \in \mathbb{R}^{n}} f$ and let $f^{*}$ denote the optimal value of this minimization problem, i.e.

$$
\forall \boldsymbol{x} \in X^{*}: f(\boldsymbol{x})=f^{*}=\min _{\boldsymbol{y} \in \mathbb{R}^{n}} f(y) \quad \text { and } \quad \forall \boldsymbol{x} \notin X^{*}: f(\boldsymbol{x})>f^{*}
$$

We assume that $X^{*}$ is non-empty. Now, we are ready to estimate the convergence rate of the gradient descent method.

Theorem 12.3.1. Let $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ be a convex continuously differentiable function and let $L$ be the Lipschitz constant of $\nabla f$. For initial point $\boldsymbol{x}_{0} \in \mathbb{R}^{n}$ we define a sequence of $\boldsymbol{x}_{k}$ by the update rule

$$
\boldsymbol{x}_{k+1}:=\boldsymbol{x}_{k}-\frac{1}{L}\left(\nabla f\left(\boldsymbol{x}_{k}\right)\right)^{\#}
$$

For all $k \geq 0$, we have

$$
f\left(\boldsymbol{x}_{k}\right)-f^{*} \leq \frac{2 \cdot L \cdot R^{2}}{k+4} \quad \text { where } \quad R \stackrel{\text { def }}{=} \max _{\boldsymbol{x} \in \mathbb{R}^{n}: f(\boldsymbol{x}) \leq f\left(\boldsymbol{x}_{0}\right)} \min _{\boldsymbol{x}^{*} \in X^{*}}\left\|\boldsymbol{x}-\boldsymbol{x}^{*}\right\| .
$$

Proof. ${ }^{3}$ By the Lipschitz continuity of the gradient of $f$ and Lemma 12.8.5 we have

$$
f\left(\boldsymbol{x}_{k+1}\right) \leq f\left(\boldsymbol{x}_{k}\right)-\frac{1}{2 L}\left(\left\|\nabla f\left(\boldsymbol{x}_{k}\right)\right\|^{*}\right)^{2}
$$

Furthermore, by the convexity of $f$, we know that

$$
\forall \boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{n} \quad: \quad f(\boldsymbol{y}) \geq f(\boldsymbol{x})+\langle\nabla f(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\rangle
$$

Using this and the fact that $f\left(\boldsymbol{x}_{k}\right)$ decreases monotonically with $k$, we get

$$
f\left(\boldsymbol{x}_{k}\right)-f^{*} \leq \min _{\boldsymbol{x}^{*} \in X^{*}}\left\langle\nabla f\left(\boldsymbol{x}_{k}\right), \boldsymbol{x}_{k}-\boldsymbol{x}^{*}\right\rangle \leq \min _{\boldsymbol{x}^{*} \in X^{*}}\left\|\nabla f\left(\boldsymbol{x}_{k}\right)\right\|^{*}\left\|\boldsymbol{x}_{k}-\boldsymbol{x}^{*}\right\| \leq R\left\|\nabla f\left(\boldsymbol{x}_{k}\right)\right\|^{*}
$$

Therefore, letting $\phi_{k} \stackrel{\text { def }}{=} f\left(\boldsymbol{x}_{k}\right)-f^{*}$, we have

$$
\phi_{k}-\phi_{k+1} \geq \frac{1}{2 L}\left(\left\|\nabla f\left(\boldsymbol{x}_{k}\right)\right\|^{*}\right)^{2} \geq \frac{\phi_{k}^{2}}{2 \cdot L \cdot R^{2}}
$$

Furthermore, since $\phi_{k} \geq \phi_{k+1}$, we have

$$
\frac{1}{\phi_{k+1}}-\frac{1}{\phi_{k}}=\frac{\phi_{k}-\phi_{k+1}}{\phi_{k} \phi_{k+1}} \geq \frac{\phi_{k}-\phi_{k+1}}{\phi_{k}^{2}} \geq \frac{1}{2 \cdot L \cdot R^{2}}
$$

So, by induction, we have that

$$
\frac{1}{\phi_{k}}-\frac{1}{\phi_{0}} \geq \frac{k}{2 \cdot L \cdot R^{2}}
$$

[^24]Now, note that since $\nabla f\left(\boldsymbol{x}^{*}\right)=0$, we have that

$$
f\left(\boldsymbol{x}_{0}\right) \leq f\left(\boldsymbol{x}^{*}\right)+\left\langle\nabla f\left(\boldsymbol{x}^{*}\right), \boldsymbol{x}_{0}-\boldsymbol{x}^{*}\right\rangle+\frac{L}{2}\left\|\boldsymbol{x}_{0}-\boldsymbol{x}^{*}\right\|^{2} \leq f\left(\boldsymbol{x}^{*}\right)+\frac{L}{2} R^{2} .
$$

So, we have that $\phi_{0} \leq \frac{L}{2} R^{2}$ and putting this all together yields that

$$
\frac{1}{\phi_{k}} \geq \frac{1}{\phi_{0}}+\frac{k}{2 \cdot L \cdot R^{2}} \geq \frac{4}{2 \cdot L \cdot R^{2}}+\frac{k}{2 \cdot L \cdot R^{2}} .
$$

## ■ 12.3.2 Maximum Flow Formulation

For an arbitrary set of demands $\boldsymbol{\chi} \in \mathbb{R}^{V}$ we wish to solve the following maximum flow problem

$$
\max _{\alpha \in R, \boldsymbol{f} \in \mathbb{R}^{E}} \alpha \quad \text { subject to } \quad \boldsymbol{B}^{T} \boldsymbol{f}=\alpha \boldsymbol{\chi} \quad \text { and } \quad\left\|\boldsymbol{U}^{-1} \boldsymbol{f}\right\|_{\infty} \leq 1 .
$$

Equivalently, we want to compute a minimum congestion flow

$$
\min _{\boldsymbol{f} \in \mathbb{R}^{E}: \boldsymbol{B}^{T} \boldsymbol{f}=\chi}\left\|\boldsymbol{U}^{-1} \boldsymbol{f}\right\|_{\infty} .
$$

where we call $\left\|\boldsymbol{U}^{-1} \boldsymbol{f}\right\|_{\infty}$ the congestion of $\boldsymbol{f}$.
Letting $\boldsymbol{f}_{0} \in \mathbb{R}^{E}$ be some initial feasible flow, i.e. $\boldsymbol{B}^{T} \boldsymbol{f}_{0} \stackrel{\text { def }}{=} \boldsymbol{\chi}$, we write the problem equivalently as

$$
\min _{\boldsymbol{c} \in \mathbb{R}^{E}: \boldsymbol{B}^{T} \boldsymbol{c}=0}\left\|\boldsymbol{U}^{-1}\left(\boldsymbol{f}_{0}+\boldsymbol{c}\right)\right\|_{\infty}
$$

where the output flow is $\boldsymbol{f}=\boldsymbol{f}_{0}+\boldsymbol{c}$. Although the gradient descent method is applicable to constrained optimization problems and has a similar convergence guarantee, the sub-problem involved in each iteration is a constrained optimization problem, which is quite complicated in this case. Since the domain is a linear subspace, the constraints can be avoided by projecting the variables onto this subspace.

Formally, we define a circulation projection matrix as follows.
Definition 12.3.2. A matrix $\tilde{\boldsymbol{P}} \in \mathbb{R}^{E \times E}$ is a circulation projection matrix if it is a projection matrix onto the circulation space, i.e. it satisfies the following

- $\forall \boldsymbol{x} \in \mathbb{R}^{E}$ we have $\boldsymbol{B}^{T} \tilde{\boldsymbol{P}} \boldsymbol{x}=\mathbf{0}$.
- $\forall \boldsymbol{x} \in \mathbb{R}^{E}$ with $\boldsymbol{B}^{T} \boldsymbol{x}=\mathbf{0}$ we have $\boldsymbol{P} \boldsymbol{x}=\boldsymbol{x}$.

Then, the problem becomes

$$
\min _{c \in \mathbb{R}^{E}}\left\|\boldsymbol{U}^{-1}\left(\boldsymbol{f}_{0}+\tilde{\boldsymbol{P}} \boldsymbol{c}\right)\right\|_{\infty}
$$

Applying gradient descent on this problem is similar to applying projected gradient method on the original problem. But, instead of using the orthogonal projection that is not suitable for $\|\cdot\|_{\infty}$, we will pick a better projection matrix.

Applying the change of basis $\boldsymbol{x}=\boldsymbol{U}^{-1} \boldsymbol{c}$ and letting $\boldsymbol{\alpha}_{\boldsymbol{0}}=\boldsymbol{U}^{-1} \boldsymbol{f}_{0}$ and $\boldsymbol{P}=\boldsymbol{U}^{-1} \tilde{\boldsymbol{P}} \boldsymbol{U}$, we write the problem equivalently as

$$
\min _{x \in \mathbb{R}^{E}}\left\|\alpha_{0}+\boldsymbol{P} \boldsymbol{x}\right\|_{\infty}
$$

where the output maximum flow is

$$
\boldsymbol{f}(\boldsymbol{x})=\boldsymbol{U}\left(\boldsymbol{\alpha}_{0}+\boldsymbol{P} \boldsymbol{x}\right) /\left\|\boldsymbol{U}\left(\boldsymbol{\alpha}_{0}+\boldsymbol{P} \boldsymbol{x}\right)\right\|_{\infty} .
$$

### 12.3.3 An Approximate Maximum Flow Algorithm

Since the gradient descent method requires the objective function to be differentiable, we introduce a smooth version of $\|\cdot\|_{\infty}$ which we call $\operatorname{smax}_{t}$. In next section, we prove that there is a convex differentiable function $\operatorname{smax}_{t}$ such that $\nabla$ smax $_{t}$ is Lipschitz continuous with Lipschitz constant $\frac{1}{t}$ and such that

$$
\forall \boldsymbol{x} \in \mathbb{R}^{E} \quad: \quad\|\boldsymbol{x}\|_{\infty}-t \ln (2 m) \leq \operatorname{smax}_{t}(\boldsymbol{x}) \leq\|\boldsymbol{x}\|_{\infty} .
$$

Now we consider the following regularized optimization problem

$$
\min _{\boldsymbol{x} \in \mathbb{R}^{E}} g_{t}(\boldsymbol{x}) \quad \text { where } \quad g_{t}(\boldsymbol{x})=\operatorname{smax}_{t}\left(\boldsymbol{\alpha}_{0}+\boldsymbol{P} \boldsymbol{x}\right) .
$$

For the rest of this section, we consider solving this optimization problem using gradient descent under $\|\cdot\|_{\infty}$.

First, we bound the Lipschitz constant of the gradient of $g_{t}$.
Lemma 12.3.3. The gradient of $g_{t}$ is Lipschitz continuous with Lipschitz constant $L=\frac{\|P\|_{\infty}^{2}}{t}$.
Proof. By Lemma 12.8.5 and the Lipschitz continuity of $\nabla \operatorname{smax}_{t}$, we have

$$
\operatorname{smax}_{t}(\boldsymbol{y}) \leq \operatorname{smax}_{t}(\boldsymbol{x})+\left\langle\nabla \operatorname{smax}_{t}(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\right\rangle+\frac{1}{2 t}\|\boldsymbol{y}-\boldsymbol{x}\|_{\infty} .
$$

Seting $\boldsymbol{x} \leftarrow \boldsymbol{\alpha}_{0}+\boldsymbol{P} \boldsymbol{x}$ and $\boldsymbol{y} \leftarrow \boldsymbol{\alpha}_{0}+\boldsymbol{P} \boldsymbol{y}$, we have

$$
\begin{aligned}
g_{t}(\boldsymbol{y}) & \leq g_{t}(\boldsymbol{y})+\left\langle\nabla \operatorname{smax}_{t}\left(\boldsymbol{\alpha}_{0}+\boldsymbol{P} \boldsymbol{x}\right), \boldsymbol{P} \boldsymbol{y}-\boldsymbol{P} \boldsymbol{x}\right\rangle+\frac{1}{2 t}\|\boldsymbol{P} \boldsymbol{y}-\boldsymbol{P} \boldsymbol{x}\|_{\infty}^{2} \\
& \leq g_{t}(\boldsymbol{y})+\left\langle\boldsymbol{P}^{T} \nabla \operatorname{smax}_{t}\left(\boldsymbol{\alpha}_{0}+\boldsymbol{P} \boldsymbol{x}\right), \boldsymbol{y}-\boldsymbol{x}\right\rangle+\frac{1}{2 t}\|\boldsymbol{P}\|_{\infty}^{2}\|\boldsymbol{y}-\boldsymbol{x}\|_{\infty}^{2} \\
& =g_{t}(\boldsymbol{y})+\left\langle\nabla g_{t}(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\right\rangle+\frac{1}{2 t}\|\boldsymbol{P}\|_{\infty}^{2}\|\boldsymbol{y}-\boldsymbol{x}\|_{\infty}^{2} .
\end{aligned}
$$

Hence, the result follows from Lemma 12.8.5.

Now, we apply gradient descent to find an approximate max flow as follows.

```
Algorithm 33: MaxFlow
    Input: any initial feasible flow \(\boldsymbol{f}_{0}\) and OPT \(=\min _{\boldsymbol{x}}\left\|\boldsymbol{U}^{-1} \boldsymbol{f}_{0}+\boldsymbol{P} \boldsymbol{x}\right\|_{\infty}\).
    Let \(\boldsymbol{\alpha}_{0}=(\boldsymbol{I}-\boldsymbol{P}) \boldsymbol{U}^{-1} \boldsymbol{f}_{0}\) and \(\boldsymbol{x}_{0}=0\).
    Let \(t=\varepsilon \mathrm{OPT} / 2 \ln (2 m)\) and \(k=300\|\boldsymbol{P}\|_{\infty}^{4} \ln (2 m) / \varepsilon^{2}\).
    Let \(g_{t}=\operatorname{smax}_{t}\left(\boldsymbol{\alpha}_{0}+\boldsymbol{P} \boldsymbol{x}\right)\).
    for \(i=1,2, \cdots, k\) do
        \(\boldsymbol{x}_{i+1}=\boldsymbol{x}_{i}-\frac{t}{\|\boldsymbol{P}\|_{\infty}^{2}}\left(\nabla g_{t}\left(\boldsymbol{x}_{i}\right)\right)^{\#}\). (See Lemma 12.3.5)
    end
    Output: \(\boldsymbol{U}\left(\overrightarrow{\alpha_{0}}+\boldsymbol{P} \boldsymbol{x}_{k}\right) /\left\|\overrightarrow{\alpha_{0}}+\boldsymbol{P} \boldsymbol{x}_{k}\right\|_{\infty}\).
```

We remark that the initial flow can be obtained by BFS and the OPT value can be approximted using binary search. In Section 12.7, we will give an algorithm with better dependence on $\|\boldsymbol{P}\|$.

Theorem 12.3.4. Let $\widetilde{\boldsymbol{P}}$ be a cycle projection matrix, let $\boldsymbol{P}=\boldsymbol{U}^{-1} \widetilde{\boldsymbol{P}} \boldsymbol{U}$, and let $\varepsilon<1$. MaxFlow outputs an $(1-\varepsilon)$-approximate maximum flow in time

$$
O\left(\frac{\|\boldsymbol{P}\|_{\infty}^{4} \ln (m)(\mathcal{T}(\boldsymbol{P})+m)}{\varepsilon^{2}}\right)
$$

Proof. First, we bound $\left\|\overrightarrow{\alpha_{0}}\right\|_{\infty}$. Let $\boldsymbol{x}^{*}$ be a minimizer of $\min _{x}\left\|\boldsymbol{U}^{-1} \boldsymbol{f}_{0}+\boldsymbol{P} \boldsymbol{x}\right\|_{\infty}$ such that $\boldsymbol{P} \vec{x}^{*}=\vec{x}^{*}$. Then, we have

$$
\begin{aligned}
\left\|\boldsymbol{\alpha}_{0}\right\|_{\infty} & =\left\|\boldsymbol{U}^{-1} \boldsymbol{f}_{0}-\boldsymbol{P} \boldsymbol{U}^{-1} \boldsymbol{f}_{0}\right\|_{\infty} \\
& \leq\left\|\boldsymbol{U}^{-1} \boldsymbol{f}_{0}+\boldsymbol{x}^{*}\right\|_{\infty}+\left\|\boldsymbol{x}^{*}+\boldsymbol{P} \boldsymbol{U}^{-1} \boldsymbol{f}_{0}\right\|_{\infty} \\
& =\left\|\boldsymbol{U}^{-1} \boldsymbol{f}_{0}+\boldsymbol{x}^{*}\right\|_{\infty}+\left\|\boldsymbol{P} \boldsymbol{x}^{*}+\boldsymbol{P} \boldsymbol{U}^{-1} \boldsymbol{f}_{0}\right\|_{\infty} \\
& \leq\left(1+\|\boldsymbol{P}\|_{\infty}\right)\left\|\boldsymbol{U}^{-1} \boldsymbol{f}_{0}+\boldsymbol{x}^{*}\right\|_{\infty} \\
& =\left(1+\|\boldsymbol{P}\|_{\infty}\right) \text { OPT. }
\end{aligned}
$$

Second, we bound $R$ in Theorem 12.3.1. Note that

$$
g_{t}\left(\boldsymbol{x}_{0}\right)=\operatorname{smax}_{t}\left(\boldsymbol{\alpha}_{0}\right) \leq\left\|\boldsymbol{\alpha}_{0}\right\|_{\infty} \leq\left(1+\|\boldsymbol{P}\|_{\infty}\right) \mathrm{OPT} .
$$

Hence, the condition $g_{t}(\boldsymbol{x}) \leq g_{t}\left(\boldsymbol{x}_{0}\right)$ implies that

$$
\left\|\boldsymbol{\alpha}_{0}+\boldsymbol{P} \boldsymbol{x}\right\|_{\infty} \leq\left(1+\|\boldsymbol{P}\|_{\infty}\right) \mathrm{OPT}+t \ln (2 m) .
$$

For any $\boldsymbol{y} \in X^{*}$ let $\boldsymbol{c}=\boldsymbol{x}-\boldsymbol{P} \boldsymbol{x}+\boldsymbol{y}$ and note that $\boldsymbol{P} \boldsymbol{c}=\boldsymbol{P} \boldsymbol{y}$ and therefore $\boldsymbol{c} \in X^{*}$. Using these facts, we can bound $R$ as follows

$$
\begin{aligned}
R & =\max _{\boldsymbol{x} \in \mathbb{R}^{E}: g_{t}(\boldsymbol{x}) \leq g_{t}\left(\boldsymbol{x}_{0}\right)}\left\{\min _{\boldsymbol{x}^{*} \in X^{*}}\left\|\boldsymbol{x}-\boldsymbol{x}^{*}\right\|_{\infty}\right\} \\
& \leq \max _{x \in \mathbb{R}^{E}: g_{t}(\boldsymbol{x}) \leq g_{t}\left(\boldsymbol{x}_{0}\right)}\|\boldsymbol{x}-\boldsymbol{c}\|_{\infty} \\
& \leq \max _{\boldsymbol{x} \in \mathbb{R}^{E}: g_{t}(\boldsymbol{x}) \leq g_{t}\left(\boldsymbol{x}_{0}\right)}\|\boldsymbol{P} \boldsymbol{x}-\boldsymbol{P} \boldsymbol{y}\|_{\infty} \\
& \leq \max _{x \in \mathbb{R}^{E}: g_{t}(\boldsymbol{x}) \leq g_{t}\left(\boldsymbol{x}_{0}\right)}\|\boldsymbol{P} \boldsymbol{x}\|_{\infty}+\|\boldsymbol{P} \boldsymbol{y}\|_{\infty} \\
& \leq 2\left\|\boldsymbol{\alpha}_{0}\right\|_{\infty}+\left\|\boldsymbol{\alpha}_{0}+\boldsymbol{P} \boldsymbol{x}\right\|_{\infty}+\left\|\boldsymbol{\alpha}_{0}+\boldsymbol{P} \boldsymbol{y}\right\|_{\infty} \\
& \leq 2\left\|\boldsymbol{\alpha}_{0}\right\|_{\infty}+2\left\|\boldsymbol{\alpha}_{0}+\boldsymbol{P} \boldsymbol{x}\right\|_{\infty} \\
& \leq 4\left(1+\|\boldsymbol{P}\|_{\infty}\right) \mathrm{OPT}+2 t \ln (2 m) .
\end{aligned}
$$

From Lemma 12.3.3, we know that the Lipschitz constant of $\nabla g_{t}$ is $\|\boldsymbol{P}\|_{\infty}^{2} / t$. Hence, Theorem 12.3.1 shows that

$$
\begin{aligned}
g_{t}\left(\boldsymbol{x}_{k}\right) & \leq \min _{x} g_{t}(\boldsymbol{x})+\frac{2 \cdot L \cdot R^{2}}{k+4} \\
& \leq \mathrm{OPT}+\frac{2 \cdot L \cdot R^{2}}{k+4}
\end{aligned}
$$

So, we have

$$
\begin{aligned}
\left\|\boldsymbol{\alpha}_{0}+\boldsymbol{P} \boldsymbol{x}_{k}\right\|_{\infty} & \leq g_{t}\left(\boldsymbol{x}_{k}\right)+t \ln (2 m) \\
& \leq \mathrm{OPT}+t \ln (2 m)+\frac{2\|\boldsymbol{P}\|_{\infty}^{2}}{t(k+4)}\left(4\left(1+\|\boldsymbol{P}\|_{\infty}\right) \mathrm{OPT}+2 t \ln (2 m)\right)^{2}
\end{aligned}
$$

Using $t=\varepsilon \mathrm{OPT} / 2 \ln (2 m)$ and $k=300\|\boldsymbol{P}\|_{\infty}^{4} \ln (2 m) / \varepsilon^{2}$, we have

$$
\left\|\boldsymbol{\alpha}_{0}+\boldsymbol{P} \boldsymbol{x}_{k}\right\|_{\infty} \leq(1+\varepsilon) \mathrm{OPT}
$$

Therefore, $\boldsymbol{\alpha}_{0}+\boldsymbol{P} \boldsymbol{x}_{k}$ is an $(1-\varepsilon)$ approximate maximum flow.
Now, we estimate the running time. In each step 5 , we are required to compute $\left(\nabla g\left(\boldsymbol{x}_{k}\right)\right)^{\#}$. The gradient

$$
\nabla g(\boldsymbol{x})=\boldsymbol{P}^{T} \nabla \operatorname{smax}_{t}\left(\boldsymbol{\alpha}_{0}+\boldsymbol{P} \boldsymbol{x}\right)
$$

can be computed in $O(\mathcal{T}(\boldsymbol{P})+m)$ using the formula of the gradient of smax $_{t}$, applications of $\boldsymbol{P}$ and $\boldsymbol{P}^{T}$. Lemma 12.3 .5 shows that the \# operator can be computed in $O(\mathrm{~m})$.

Lemma 12.3.5. In $\|\cdot\|_{\infty}$, the $\#$ operator is given by the explicit formula

$$
\left(\boldsymbol{x}^{\#}\right)_{e}=\operatorname{sign}\left(x_{e}\right)\|\boldsymbol{x}\|_{1} \quad \text { for } e \in E
$$

Proof. Recall that

$$
\boldsymbol{x}^{\#}=\underset{s \in \mathbb{R}}{\arg \max }\langle\boldsymbol{x}, \boldsymbol{s}\rangle-\frac{1}{2}\|\vec{s}\|_{\infty}^{2}
$$

It is easy to see that for all $e \in E,\left\|\boldsymbol{x}^{\#}\right\|_{\infty}=\left|\left(\boldsymbol{x}^{\#}\right)_{e}\right|$. In particular, we have

$$
\left(\boldsymbol{x}^{\#}\right)_{e}=\operatorname{sign}\left(x_{e}\right)\left\|\boldsymbol{x}^{\#}\right\|_{\infty}
$$

Fact 12.8 .3 shows that $\left\|\boldsymbol{x}^{\#}\right\|_{\infty}=\|\boldsymbol{x}\|_{1}$ and the result follows.

## ■ 12.3.4 Properties of soft max

In this section, we define $\operatorname{smax}_{t}$ and discuss its properties. Formally, the regularized convex function can be found by smoothing technique using convex conjugate [208] [35, Sec 5.4]. For simplicity and completeness, we define it explicitly and prove its properties directly. Formally, we define

$$
\forall \boldsymbol{x} \in \mathbb{R}^{E}, \forall t \in \mathbb{R}_{>0} \quad: \quad \operatorname{smax}_{t}(\boldsymbol{x}) \stackrel{\text { def }}{=} t \ln \left(\frac{\sum_{e \in E} \exp \left(\frac{x_{e}}{t}\right)+\exp \left(-\frac{x_{e}}{t}\right)}{2 m}\right)
$$

For notational simplicity, for all $\boldsymbol{x}$ where this vector is clear from context, we define $\boldsymbol{c}$ and $\boldsymbol{s}$ as follows

$$
\forall e \in E \quad: \quad \boldsymbol{c}_{e} \stackrel{\text { def }}{=} \exp \left(\frac{x_{e}}{t}\right)+\exp \left(-\frac{x_{e}}{t}\right) \quad \text { and } \quad \boldsymbol{s}_{e} \stackrel{\text { def }}{=} \exp \left(\frac{x_{e}}{t}\right)-\exp \left(-\frac{x_{e}}{t}\right)
$$

where the letters are chosen due to the very close resemblance to hyperbolic sine and hyperbolic cosine.

## Lemma 12.3.6.

$$
\begin{gathered}
\forall \boldsymbol{x} \in \mathbb{R}^{n} \quad: \quad \nabla \operatorname{smax}_{t}(\boldsymbol{x})=\frac{1}{\mathbf{1}^{T} \boldsymbol{c}} \boldsymbol{s} \\
\forall \boldsymbol{x} \in \mathbb{R}^{n} \quad: \quad \nabla^{2} \operatorname{smax}_{t}(\boldsymbol{x})=\frac{1}{t\left(\mathbf{1}^{T} \boldsymbol{c}\right)}\left[\boldsymbol{\operatorname { c i a }} \boldsymbol{g}(\boldsymbol{c})-\frac{\boldsymbol{s} \boldsymbol{s}^{T}}{\mathbf{1}^{T} \boldsymbol{c}}\right]
\end{gathered}
$$

Proof. For all $i \in E$ and $\boldsymbol{x} \in \mathbb{R}^{E}$, we have

$$
\begin{aligned}
\frac{\partial}{\partial x_{i}} \operatorname{smax}_{t}(\boldsymbol{x}) & =\frac{\partial}{\partial x_{i}}\left(t \ln \left(\frac{\sum_{e \in E} \exp \left(\frac{x_{e}}{t}\right)+\exp \left(-\frac{x_{e}}{t}\right)}{2 m}\right)\right) \\
& =\frac{\exp \left(\frac{x_{i}}{t}\right)-\exp \left(-\frac{x_{i}}{t}\right)}{\sum_{e \in E} \exp \left(\frac{x_{e}}{t}\right)+\exp \left(-\frac{x_{e}}{t}\right)} .
\end{aligned}
$$

For all $i, j \in E$ and $\boldsymbol{x} \in \mathbb{R}^{E}$, we have

$$
\begin{aligned}
\frac{\partial^{2}}{\partial x_{i} \partial x_{j}} \operatorname{smax}_{t}(\boldsymbol{x}) & =\frac{\partial^{2}}{\partial x_{i} \partial x_{j}}\left(t \ln \left(\frac{\sum_{e \in E} \exp \left(\frac{x_{e}}{t}\right)+\exp \left(-\frac{x_{e}}{t}\right)}{2 m}\right)\right) \\
& =\frac{\partial}{\partial j}\left[\frac{\exp \left(\frac{x_{i}}{t}\right)-\exp \left(-\frac{x_{i}}{t}\right)}{\sum_{e \in E} \exp \left(\frac{x_{e}}{t}\right)+\exp \left(-\frac{x_{e}}{t}\right)}\right] \\
& =\frac{1}{t} \frac{\left(\mathbf{1}^{T} \boldsymbol{c}\right) \mathbf{1}_{i=j}\left(\boldsymbol{c}_{i}\right)-\boldsymbol{s}_{i} \boldsymbol{s}_{j}}{\left(\mathbf{1}^{T} \boldsymbol{c}\right)^{2}} .
\end{aligned}
$$

Lemma 12.3.7. The function $\operatorname{smax}_{t}$ is a convex continuously differentiable function and it has Lipschitz continuous gradient with Lipschitz constant $1 / t$ and

$$
\|\boldsymbol{x}\|_{\infty}-t \ln (2 m) \leq \operatorname{smax}_{t}(\boldsymbol{x}) \leq\|\boldsymbol{x}\|_{\infty}
$$

for $\boldsymbol{x} \in \mathbb{R}^{E}$.
Proof. By the formulation of the Hessian, for all $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{E}$, we have

$$
\boldsymbol{y}^{T}\left(\nabla^{2} \operatorname{smax}_{t}(\boldsymbol{x})\right) \boldsymbol{y} \leq \frac{\sum_{i} c_{i} \boldsymbol{y}_{i}^{2}}{t\left(\mathbf{1}^{T} \boldsymbol{c}\right)} \leq \frac{\sum_{i} c_{i}\left(\max _{j} \boldsymbol{y}_{j}^{2}\right)}{t\left(\mathbf{1}^{T} \boldsymbol{c}\right)} \leq \frac{1}{t}\|y\|_{\infty}^{2}
$$

On the other side, for all $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{E}$, we have by $s_{i} \leq\left|s_{i}\right| \leq c_{i}$ and Cauchy Schwarz shows that

$$
\boldsymbol{y}^{T} s^{T} \boldsymbol{y} \leq\left(\mathbf{1}^{T}|s|\right)\left(\boldsymbol{y}^{T} \operatorname{diag}(|s|) \boldsymbol{y}\right) . \leq\left(\mathbf{1}^{T} \boldsymbol{c}\right)\left(\boldsymbol{y}^{T} \operatorname{diag}(\boldsymbol{c}) \boldsymbol{y}\right) .
$$

and hence

$$
0 \leq \boldsymbol{y}^{T}\left(\nabla^{2} \operatorname{smax}_{t}(\boldsymbol{x})\right) \boldsymbol{y}
$$

Thus, the first part follows from Lemma 12.8.6. For the later part, we have

$$
\|\boldsymbol{x}\|_{\infty} \geq t \ln \left(\frac{\sum_{e \in E} \exp \left(\frac{x_{e}}{t}\right)+\exp \left(-\frac{x_{e}}{t}\right)}{2 m}\right) \geq t \ln \left(\frac{\exp \left(\frac{\|x\|_{\infty}}{t}\right)}{2 m}\right)=\|\boldsymbol{x}\|_{\infty}-\ln (2 m) .
$$

## ■ 12.4 Oblivious Routing

In the previous sections, we saw how a circulation projection matrix can be used to solve max flow. In the next few sections, we show how to efficiently construct a circulation projection matrix to obtain an almost linear time algorithm for solving max flow.

Our proof focuses on the notion of (linear) oblivious routings. Rather than constructing the circulation projection matrix directly, we show how the efficient construction of an oblivious routing
algorithm with a good competitive ratio immediately allows us to produce a circulation projection matrix.

In the remainder of this section, we formally define oblivious routings and prove the relationship between oblivious routing and circulation projection matrices (Section 12.4.1), provide a high level overview of our recursive approach and state the main theorems we will prove in later sections (Section 12.4.2). Finally, we prove the main theorem about our almost-linear-time construction of circulation projection with norm $2^{O(\sqrt{\log (n) \log \log (n)})}$ assuming the proofs in the later sections (Section 12.4.3).

## ■ 12.4.1 From Oblivious Routing to Circulation Projection

Here we provide definitions and prove basic properties of oblivious routings, that is, fixed mappings from demands to flows that meet the input demands. While non-linear algorithms could be considered, we restrict our attention to linear oblivious routing strategies and use the term oblivious routing to refer to the linear subclass for the remainder of this chapter. ${ }^{4}$

Definition 12.4.1 (Oblivious Routing). An oblivious routing on graph $G=(V, E)$ is a linear operator $\boldsymbol{A} \in \mathbb{R}^{E \times V}$ such that for all demands $\boldsymbol{\chi} \in \mathbb{R}^{V}, \boldsymbol{B}^{T} \boldsymbol{A} \boldsymbol{\chi}=\boldsymbol{\chi}$. We call $\boldsymbol{A} \boldsymbol{\chi}$ the routing of $\boldsymbol{\chi}$ by $\boldsymbol{A}$.

Given an oblivious routing strategy $\boldsymbol{A}$ and a set of demands $D=\left\{\chi_{1}, \ldots, \chi_{k}\right\}$, one can construct a multicommodity flow satisfying all the demands in $D$ by using $\boldsymbol{A}$ to route each demand individually, obliviously to the existence of the other demands. We measure the competitive ratio of such an oblivious routing strategy to be the ratio of the worst relative congestion of such a routing to the minimal-congestion routing of the demands.

Definition 12.4.2 (Competitive Ratio). The competitive ratio of oblivious routing $\boldsymbol{A} \in \mathbb{R}^{E \times V}$, denoted $\rho(\boldsymbol{A})$, is given by

$$
\rho(\boldsymbol{A}) \stackrel{\text { def }}{=} \max _{\left\{\boldsymbol{\chi}_{i}\right\}: \forall i} \boldsymbol{\chi}_{\chi_{i} \perp \mathbf{1}} \frac{\operatorname{cong}\left(\left\{\boldsymbol{A}_{\boldsymbol{1}}\right\}\right)}{\operatorname{OPT}\left(\left\{\boldsymbol{\chi}_{i}\right\}\right)}
$$

At times, it will be more convenient to analyze an oblivious routing as a linear algebraic object rather a combinatorial algorithm; towards this end, we note that the competitive ratio of a linear oblivious routing strategy can be gleaned from the operator norm of a related matrix (see also [158] and [139]). Below, we state and prove a generalization of this result to weighted graphs that will be vital to relating $\boldsymbol{A}$ to $\widetilde{\boldsymbol{P}}$.

Lemma 12.4.3. For any oblivious routing $\boldsymbol{A}$, we have $\rho(\boldsymbol{A})=\left\|\boldsymbol{U}^{-1} \boldsymbol{A} \boldsymbol{B}^{T} \boldsymbol{U}\right\|_{\infty}$
Proof. For a set of demands $D$, let $D_{\infty}$ be the set of demands that results by taking the routing of every demand in $D$ by $\operatorname{OPT}(D)$ and splitting it up into demands on every edge corresponding to the flow sent by $\operatorname{OPT}(D)$. Now, clearly $\operatorname{OPT}(D)=\operatorname{OPT}\left(D_{\infty}\right)$ since routing $D$ can be used to route $D_{\infty}$ and vice versa, and clearly cong $(\boldsymbol{A} D) \leq \operatorname{cong}\left(\boldsymbol{A} D_{\infty}\right)$ by the linearity of $\boldsymbol{A}$ (routing $D_{\infty}$ simply doesn't reward $\boldsymbol{A}$ routing for cancellations). Using this and the fact that $D_{\infty} \subseteq D$, we have:

$$
\begin{aligned}
\rho_{p}(\boldsymbol{A}) & =\max _{D} \frac{\operatorname{cong}(\{\boldsymbol{A} D\})}{\operatorname{OPT}(D)}=\max _{D_{\infty}} \frac{\operatorname{cong}\left(\boldsymbol{A} D_{\infty}\right)}{\operatorname{OPT}\left(D_{\infty}\right)}=\max _{x \in \mathbb{R}^{E}} \frac{\left\|\sum_{e \in E} \boldsymbol{x}_{e}\left|\boldsymbol{U}^{-1} \boldsymbol{A} \boldsymbol{\chi}_{e}\right|\right\|_{\infty}}{\left\|\boldsymbol{U}^{-1} \boldsymbol{x}\right\|_{\infty}} \\
& =\max _{\boldsymbol{x} \in \mathbb{R}^{E}} \frac{\left\|\left|\boldsymbol{U}^{-1} \boldsymbol{A} \boldsymbol{B}^{T}\right| \boldsymbol{x}\right\|_{\infty}}{\left\|\boldsymbol{U}^{-1} \boldsymbol{x}\right\|_{\infty}}=\max _{\boldsymbol{x} \in \mathbb{R}^{E}} \frac{\left\|\left|\boldsymbol{U}^{-1} \boldsymbol{A} \boldsymbol{B}^{T} \boldsymbol{U}\right| \boldsymbol{x}\right\|_{\infty}}{\|\boldsymbol{x}\|_{\infty}} .
\end{aligned}
$$

[^25]To make this lemma easily applicable in a variety of settings, we make use of the following easy-to-prove lemma.

Lemma 12.4.4 (Operator Norm Bounds). For all $\boldsymbol{A} \in \mathbb{R}^{n \times m}$, we have that

$$
\|\boldsymbol{A}\|_{\infty}=\||\boldsymbol{A}|\|_{\infty}=\||\boldsymbol{A}| \mathbf{1}\|_{\infty}=\max _{i \in n}\left\|\boldsymbol{A}^{T} \mathbf{1}_{i}\right\|_{1}
$$

The previous two lemmas make the connection between oblivious routings and circulation projection matrices clear. Below, we prove it formally.

Lemma 12.4.5 (Oblivious Routing to Circulation Projection). For oblivious routing $\boldsymbol{A} \in \mathbb{R}^{E \times V}$ the matrix $\widetilde{\boldsymbol{P}} \stackrel{\text { def }}{=} \boldsymbol{I}-\boldsymbol{A} \boldsymbol{B}^{T}$ is a circulation projection matrix such that $\left\|\boldsymbol{U}^{-1} \widetilde{\boldsymbol{P}} \boldsymbol{U}\right\|_{\infty} \leq 1+\rho(\boldsymbol{A})$.

Proof. First, we verify that $\operatorname{im}(\widetilde{\boldsymbol{P}})$ is contained in cycle space:

$$
\forall \boldsymbol{x} \in \mathbb{R}^{E} \quad: \quad \boldsymbol{B}^{T} \widetilde{\boldsymbol{P}} \boldsymbol{x}=\boldsymbol{B}^{T} \boldsymbol{x}-\boldsymbol{B}^{T} \boldsymbol{A} \boldsymbol{B}^{T} \boldsymbol{x}=\mathbf{0}
$$

Next, we check that $\widetilde{\boldsymbol{P}}$ is the identity on cycle space

$$
\forall \boldsymbol{x} \in \mathbb{R}^{E} \text { s.t. } \boldsymbol{B}^{T} \boldsymbol{x}=\mathbf{0} \quad: \quad \widetilde{\boldsymbol{P}} \boldsymbol{x}=\boldsymbol{x}-\boldsymbol{A} \boldsymbol{B}^{T} \boldsymbol{x}=\boldsymbol{x}
$$

Finally, we bound the $\ell_{\infty}$-norm of the scaled projection matrix:

$$
\left\|\boldsymbol{U}^{-1} \widetilde{\boldsymbol{P}} \boldsymbol{U}\right\|_{\infty}=\left\|\boldsymbol{I}-\boldsymbol{U}^{-1} \boldsymbol{A} \boldsymbol{B}^{T} \boldsymbol{U}\right\|_{\infty} \leq 1+\rho(\boldsymbol{A})
$$

## - 12.4.2 A Recursive Approach by Embeddings

We construct an oblivious routing for a graph recursively. Given a generic, possibly complicated, graph, we show how to reduce computing an oblivious routing on this graph to computing an oblivious routing on a simpler graph on the same vertex set. A crucial concept in these constructions will be the notion of an embedding, which will allow us to relate the competitive ratios of an oblivious routing algorithms over graphs on the same vertex sets but different edge sets.

Definition 12.4.6 (Embedding). Let $G=(V, E, \boldsymbol{\mu})$ and $G^{\prime}=\left(V, E^{\prime}, \boldsymbol{\mu}^{\prime}\right)$ denote two undirected capacitated graphs on the same vertex set with incidence matrices $\boldsymbol{B} \in \mathbb{R}^{E \times V}$ and $\boldsymbol{B}^{\prime} \in \mathbb{R}^{E^{\prime} \times V}$ respectively. An embedding from $G$ to $G^{\prime}$ is a matrix $\boldsymbol{M} \in \mathbb{R}^{E^{\prime} \times E}$ such that $\boldsymbol{B}^{\prime T} \boldsymbol{M}=\boldsymbol{B}^{T}$.

In other words, an embedding is a map from flows in one graph $G$ to flows in another graph $G^{\prime}$ that preserves the demands met by the flow. We can think of an embedding as a way of routing any flow in graph $G$ into graph $G^{\prime}$ that has the same vertex set, but different edges. We will be particularly interested in embeddings that increase the congestion of the flow by a small amount going from $G$ to $G^{\prime}$ 。

Definition 12.4.7 (Embedding Congestion). Let $\boldsymbol{M} \in \mathbb{R}^{E^{\prime} \times E}$ be an embedding from $G=(V, E, \boldsymbol{\mu})$ to $G^{\prime}=\left(V, E^{\prime}, \boldsymbol{\mu}^{\prime}\right)$ and let $\boldsymbol{U} \in \mathbb{R}^{E \times E}$ and $\boldsymbol{U}^{\prime} \in \mathbb{R}^{E^{\prime} \times E^{\prime}}$ denote the capacity matrices of $G$ and $G^{\prime}$ respectively. The congestion of embedding $\boldsymbol{M}$ is given by

$$
\operatorname{cong}(\boldsymbol{M}) \stackrel{\text { def }}{=} \max _{\boldsymbol{x} \in \mathbb{R}^{E}} \frac{\left\|\boldsymbol{U}^{\prime-1} \boldsymbol{M} \boldsymbol{x}\right\|_{\infty}}{\left\|\boldsymbol{U}^{-1} \boldsymbol{x}\right\|_{\infty}}=\left\|\boldsymbol{U}^{\prime-1}|\boldsymbol{M}| \boldsymbol{U} \mathbf{1}\right\|_{\infty}
$$

We say $G$ embeds into $G^{\prime}$ with congestion $\alpha$ if there exists an embedding $\boldsymbol{M}$ from $G$ to $G^{\prime}$ such that $\operatorname{cong}(\boldsymbol{M}) \leq \alpha$.

Embeddings potentially allow us to reduce computing an oblivious routing in a complicated graph to computing an oblivious routing in a simpler graph. Specifically, if we can embed a complicated graph in a simpler graph and we can efficiently embed the simple graph in the original graph, both with low congestion, then we can just focus on constructing oblivious routings in the simpler graph. We prove this formally as follows.

Lemma 12.4.8 (Embedding Lemma). Let $G=(V, E, \boldsymbol{\mu})$ and $G^{\prime}=\left(V, E^{\prime}, \boldsymbol{\mu}^{\prime}\right)$ denote two undirected capacitated graphs on the same vertex sets, let $M \in \mathbb{R}^{E^{\prime} \times E}$ denote an embedding from $G$ into $G^{\prime}$, let $\boldsymbol{M}^{\prime} \in \mathbb{R}^{E \times E^{\prime}}$ denote an embeding from $G^{\prime}$ into $G$, and let $\boldsymbol{A}^{\prime} \in \mathbb{R}^{E^{\prime} \times V}$ denote an oblivious routing algorithm on $G^{\prime}$. Then $\boldsymbol{A} \stackrel{\text { def }}{=} \boldsymbol{M}^{\prime} \boldsymbol{A}^{\prime}$ is an oblivious routing algorithm on $G$ and

$$
\rho(\boldsymbol{A}) \leq \operatorname{cong}(\boldsymbol{M}) \cdot \operatorname{cong}\left(\boldsymbol{M}^{\prime}\right) \cdot \rho\left(\boldsymbol{A}^{\prime}\right)
$$

Proof. For all $\boldsymbol{x} \in \mathbb{R}^{V}$ we have by definition of embeddings and oblivious routings that

$$
\boldsymbol{B}^{T} \boldsymbol{A} \boldsymbol{x}=\boldsymbol{B}^{T} \boldsymbol{M}^{\prime} \boldsymbol{A}^{\prime} \boldsymbol{x}=\boldsymbol{B}^{T} \boldsymbol{x}
$$

To bound $\rho(A)$, we let $\boldsymbol{U}$ denote the capacity matrix of $G$ and $\boldsymbol{U}^{\prime}$ denote the capacity matrix of $G^{\prime}$. Using Lemma 12.4.3, we get

$$
\rho(\boldsymbol{A})=\left\|\boldsymbol{U}^{-1} \boldsymbol{A} \boldsymbol{B}^{T} \boldsymbol{U}\right\|_{\infty}=\left\|\boldsymbol{U}^{-1} \boldsymbol{M}^{\prime} \boldsymbol{A}^{\prime} \boldsymbol{B}^{T} \boldsymbol{U}\right\|_{\infty}
$$

Using that $\boldsymbol{M}$ is an embedding and therefore $\boldsymbol{B}^{\prime T} \boldsymbol{M}=\boldsymbol{B}^{T}$, we get

$$
\rho(\boldsymbol{A})=\left\|\boldsymbol{U}^{-1} \boldsymbol{M}^{\prime} \boldsymbol{A}^{\prime} \boldsymbol{B}^{\prime T} \boldsymbol{M} \boldsymbol{U}\right\|_{\infty} \leq\left\|\boldsymbol{U}^{-1} \boldsymbol{M}^{\prime} \boldsymbol{U}^{\prime}\right\|_{\infty} \cdot\left\|\boldsymbol{U}^{\prime-1} \boldsymbol{A}^{\prime} \boldsymbol{B}^{\prime T} \boldsymbol{U}^{\prime}\right\|_{\infty} \cdot\left\|\boldsymbol{U}^{\prime-1} \boldsymbol{M} \boldsymbol{U}\right\|_{\infty}
$$

By the definition of competitive ratio and congestion, we obtain the result.
Note how in this lemma we only use the embedding from $G$ to $G^{\prime}$ to certify the quality of flows in $G^{\prime}$, we do not actually need to apply this embedding in the reduction.

Using this concept, we construct oblivious routings via recursive application of two techniques. First, in Section 12.5, we show how to take an arbitrary graph $G=(V, E)$ and approximate it by a sparse graph $G^{\prime}=\left(V, E^{\prime}\right)$ (i.e. one in which $\left.\left|E^{\prime}\right|=\widetilde{O}(|V|)\right)$ such that flows in $G$ can be routed in $G^{\prime}$ with low congestion and that there is an $\widetilde{O}(1)$ embedding from $G^{\prime}$ to $G$ that can be applied in $\widetilde{O}(|E|)$ time. We call such a construction a flow sparsifiers and prove the following theorem.

Theorem 12.4.9. Let $G=(V, E, \boldsymbol{\mu})$ be an undirected capacitated graph. In $\widetilde{O}(|E| \log U)$ time we can construct a graph $G^{\prime}$ on the same vertex set with at most $\widetilde{O}(|V| \log \underset{\widetilde{O}}{U})$ edges and capacity ratio at most $U \cdot \operatorname{poly}(|V|)$. Moreover, given an oblivious routing $\boldsymbol{A}^{\prime}$ on $G^{\prime}$, in $\widetilde{O}(|E| \log U)$ time we can construct an oblivious routing $\boldsymbol{A}$ on $G$ such that

$$
\mathcal{T}(\boldsymbol{A})=\widetilde{O}\left(|E| \log U+\mathcal{T}\left(\boldsymbol{A}^{\prime}\right)\right) \quad \text { and } \quad \rho(\boldsymbol{A})=\widetilde{O}\left(\rho\left(\boldsymbol{A}^{\prime}\right)\right)
$$

Next, in Section 12.6 we show how to embed a graph into a collection of graphs consisting of trees plus extra edges. Then, we will show how to embed these graphs into better structured graphs consisting of trees plus edges so that by simply removing degree 1 and degree 2 vertices we are left with graphs with fewer vertices. Formally, we prove the following.

Theorem 12.4.10. Let $G=(V, E, \boldsymbol{\mu})$ be an undirected capacitated graph with capacity ratio $U$. For all $t>0$ in $\widetilde{O}(t \cdot|E|)$ time we can compute graphs $G_{1}, \ldots, G_{t}$ each with at most $\widetilde{O}\left(\frac{|E| \log (U)}{t}\right)$ vertices, at most $|E|$ edges, and capacity ratio at most $|V| \cdot U$. Moreover, given oblivious routings $\boldsymbol{A}_{i}$ for each
$G_{i}$, in $\widetilde{O}(t \cdot|E|)$ time we can compute an oblivious routing $\boldsymbol{A}$ on $G$ such that

$$
\mathcal{T}(\boldsymbol{A})=\widetilde{O}\left(t \cdot|E|+\sum_{i=1}^{t} \mathcal{T}\left(\boldsymbol{A}_{i}\right)\right) \quad \text { and } \quad \rho(\boldsymbol{A})=\widetilde{O}\left(\max _{i} \rho\left(\boldsymbol{A}_{i}\right)\right)
$$

In the next section we show that the careful application of these two ideas along with a powerful primitive for routing on constant sized graphs suffices to produce an oblivious routing with the desired properties.

## ■ 12.4.3 Efficient Oblivious Routing Construction Proof

First, we provide the lemma that will serve as the base case of our recursion. In particular, we show that electric routing can be used to obtain a routing algorithm with constant competitive ratio for constant-size graphs.

Lemma 12.4.11 (Base Case). Let $G=(V, E, \boldsymbol{\mu})$ be an undirected capacitated graph and let us assign weights to edges so that $\boldsymbol{W}=\boldsymbol{U}^{2}$. For $\mathcal{L} \stackrel{\text { def }}{=}{\underset{\sim}{\boldsymbol{O}}}^{T} \boldsymbol{W} \boldsymbol{B}$ we have that $\boldsymbol{A} \stackrel{\text { def }}{=} \boldsymbol{W} \boldsymbol{B L}^{+}$is an oblivious routing on $G$ with $\rho(\boldsymbol{A}) \leq \sqrt{|E|}$ and $\mathcal{T}\left(\mathcal{L}^{+}\right)=\widetilde{O}(|E|)$.

Proof. To see that $\boldsymbol{A}$ is an oblivious routing strategy we note that for any demands $\boldsymbol{\chi} \in \mathbb{R}^{V}$ we have $\boldsymbol{B}^{T} \boldsymbol{A}=\mathcal{L} \mathcal{L}^{+}=\boldsymbol{I}$. To see bound $\rho(\boldsymbol{A})$ we note that by Lemma 12.4.3 and standard norm inequalities we have

The result follows from the fact in [242] that $\boldsymbol{\Pi} \xlongequal{\text { def }} \boldsymbol{U} \boldsymbol{B} \mathcal{L}^{+} \boldsymbol{B}^{T} \boldsymbol{U}$ is an orthogonal projection, and therefore $\|\boldsymbol{\Pi}\|_{2} \leq 1$, and the fact in $[243,147,146,141]$ that $\mathcal{T}\left(\mathcal{L}^{+}\right)=\widetilde{O}(|E|)$.

Assuming Theorem 12.4.9 and Theorem 12.4.10, which we prove in the next two sections, we prove that low-congestion oblivious routings can be constructed efficiently.

Theorem 12.4.12. Given an undirected capacitated graph $G=(V, E, \boldsymbol{\mu})$ with capacity ratio $U$. Assume $U=\operatorname{poly}(|V|)$. We can construct an oblivious routing algorithm $\boldsymbol{A}$ on $G$ in time

$$
O\left(|E| 2^{O(\sqrt{\log |V| \log \log |V|})}\right)
$$

such that

$$
\mathcal{T}(\boldsymbol{A})=|E| 2^{O(\sqrt{\log |V| \log \log |V|})} \quad \text { and } \quad \rho(\boldsymbol{A})=2^{O(\sqrt{\log |V| \log \log |V|})} .
$$

Proof. Let $c$ be the constant hidden in the exponent terms, including $\tilde{O}(\cdot)$ and poly $(\cdot)$ in Theorem 12.4.9 and Theorem 12.4.10. Apply Theorem 12.4 .9 to construct a sparse graph $G^{(1)}$, then apply Theorem 12.4.10 with $t=\left\lceil 2^{\sqrt{\log |V| \log \log |V|}}\right\rceil$ to get $t$ graphs $G_{1}^{(1)}, \cdots G_{t}^{(1)}$ such that each graphs have at most $O\left(\frac{1}{t}|E| \log ^{2 c}|V| \log ^{2 c} U\right)$ vertices and at most $U \cdot|V|^{2 c}$ capacity ratio.

Repeat this process on each $G_{i}^{(1)}$, it produces $t^{2}$ graphs $G_{1}^{(2)}, \cdots, G_{t^{2}}^{(2)}$. Keep doing this until all graphs $G_{i}$ produced have $O(1)$ vertices. Let $k$ be the highest level we go through in this process. Since at the $k$-th level the number of vertices of each graph is at most $O\left(\frac{c^{k}}{t^{k}}|E| \log ^{2 k c}|V| \log ^{2 k}\left(U|V|^{2 c k}\right)\right)$ vertices, we have $k=O\left(\sqrt{\frac{\log |V|}{\log \log |V|}}\right)$.

On each graph $G_{i}$, we use Theorem 12.4.11 to get an oblivious routing algorithm $\boldsymbol{A}_{i}$ for each $G_{i}$ with

$$
\mathcal{T}\left(\boldsymbol{A}_{i}\right)=\tilde{O}(1) \quad \text { and } \quad \rho\left(\boldsymbol{A}_{i}\right)=\tilde{O}(1)
$$

Then, Theorem 12.4.10 and 12.4.9 shows that we have an oblivious routing algorithm $\boldsymbol{A}$ for $G$ with

$$
\mathcal{T}(\boldsymbol{A})=O\left(t k|E| \log ^{c k}(|V|) \log ^{2 k}\left(U|V|^{2 c k}\right)\right) \quad \text { and } \quad \rho(\boldsymbol{A})=O\left(\log ^{2 k c}|V| \log ^{k}\left(U|V|^{2 c k}\right)\right)
$$

The result follows from $k=O\left(\sqrt{\frac{\log |V|}{\log \log |V|}}\right)$ and $t=\left\lceil 2^{\sqrt{\log |V| \log \log |V|}}\right]$.
Using Theorem 12.4.12, Lemma 12.4.5 and Theorem 12.3.4, we have the following almost linear time max flow algorithm on undirected graph.

Theorem 12.4.13. Given an undirected capacitated graph $G=(V, E, \boldsymbol{\mu})$ with capacity ratio $U$. Assume $U=\operatorname{poly}(|V|)$. There is an algorithm finds an $(1-\varepsilon)$ approximate maximum flow in time

$$
O\left(\frac{|E| 2^{O(\sqrt{\log |V| \log \log |V|})}}{\varepsilon^{2}}\right) .
$$

## - 12.5 Flow Sparsifiers

In order to prove Theorem 12.4.9, i.e. reduce the problem of efficiently computing a competitive oblivious routing on a dense graph to the same problem on a sparse graph, we introduce a new algorithmic tool called flow sparsifiers. ${ }^{5}$ A flow sparsifier is an efficient cut-sparsification algorithm that also produces an efficiently-computable low-congestion embedding mapping the sparsified graph back to the original graph.

Definition 12.5.1 (Flow Sparsifier). An algorithm is a ( $h, \varepsilon, \alpha$ )-flow sparsifier if on input graph $G=$ $(V, E, \mu)$ with capacity ratio $U$ it outputs a graph $G^{\prime}=\left(V, E^{\prime}, \mu^{\prime}\right)$ with capacity ratio $U^{\prime} \leq U \cdot \operatorname{poly}(|V|)$ and an embedding $M: \mathbb{R}^{E^{\prime}} \rightarrow \mathbb{R}^{E}$ of $G^{\prime}$ into $G$ with the following properties:

- Sparsity: $G^{\prime}$ is $h$-sparse, i.e.

$$
\left|E^{\prime}\right| \leq h
$$

- Cut Approximation: $G^{\prime}$ is an $\varepsilon$-cut approximation of $G$, i.e.

$$
\forall S \subseteq V \quad: \quad(1-\varepsilon) \mu\left(\partial_{G}(S)\right) \leq \mu^{\prime}\left(\partial_{G^{\prime}}(S)\right) \leq(1+\varepsilon) \mu\left(\partial_{G}(S)\right)
$$

- Flow Approximation: $\boldsymbol{M}$ has congestion at most $\alpha$, i.e.

$$
\operatorname{cong}(\boldsymbol{M}) \leq \alpha
$$

- Efficiency: The algorithm runs in $\tilde{O}(m \log U)$ time and $\mathcal{T}(\boldsymbol{M})$ is also $\tilde{O}(m \log U)$.

Flow sparsifiers allow us to solve a multi-commodity flow problem on a possibly dense graph $G$ by converting $G$ into a sparse graph $G^{\prime}$ and solving the flow problem on $G^{\prime}$, while suffering a loss of a factor of at most $\alpha$ in the congestion when mapping the solution back to $G$ using $M$.

[^26]Theorem 12.5.2. Consider a graph $G=(V, E, \mu)$ and let $G^{\prime}=\left(V, E^{\prime}, \mu^{\prime}\right)$ be given by an $(h, \varepsilon, \alpha)$-flow sparsifier of $G$. Then, for any set of $k$ demands $D=\left\{\chi_{1}, \chi_{2}, \ldots, \chi_{k}\right\}$ between vertex pairs of $V$, we have:

$$
\begin{equation*}
\operatorname{OPT}_{G^{\prime}}(D) \leq \frac{O(\log k)}{1-\varepsilon} \cdot \operatorname{OPT}_{G}(D) \tag{12.2}
\end{equation*}
$$

Given the optimum flow $\left\{f_{i}^{\star}\right\}$ over $G^{\prime}$, we have

$$
\operatorname{cong}_{G}\left(\left\{\boldsymbol{M} f_{i}^{\star}\right\}\right) \leq \alpha \cdot \operatorname{OPT}_{G^{\prime}}(D) \leq \frac{O(\alpha \log k)}{1-\varepsilon} \cdot \operatorname{OPT}_{G}(D)
$$

Proof. By the flow-cut gap theorem of Aumann and Rabani [21], we have that, for any set of $k$ demands $D$ on $V$ we have:

$$
\operatorname{OPT}_{G}(D) \geq O\left(\frac{1}{\log k}\right) \cdot \max _{S \subset V} \frac{D(\partial(S))}{\mu\left(\partial_{G}(S)\right)}
$$

where $D(\partial(S))$ denotes the total amount of demand separated by the cut between $S$ and $\bar{S}$. As any cut $S \subseteq V$ in $G^{\prime}$ has capacity $\mu^{\prime}\left(\partial_{G^{\prime}}(S)\right) \geq(1-\varepsilon) \mu\left(\partial_{G}(S)\right)$, we have:

$$
\operatorname{OPT}_{G^{\prime}}(D) \leq \max _{S \subset V} \frac{D(\partial(S))}{\mu^{\prime}\left(\partial_{G^{\prime}}(S)\right)} \leq \frac{1}{1-\varepsilon} \cdot \max _{S \subset V} \frac{D(\partial(S))}{\mu\left(\partial_{G}(S)\right)} \leq \frac{O(\log k)}{1-\varepsilon} \cdot \operatorname{OPT}_{G}(D)
$$

The second part of the theorem follows as a consequence of the definition of the congestion of the embedding $\boldsymbol{M}$.

Our flow sparsifiers should be compared with the cut-based decompositions of Räcke [225]. Räcke constructs a probability distribution over trees and gives explicit embeddings from $G$ to this distribution and backwards, achieving a congestion of $O(\log n)$. However, this distribution over tree can include up to $O(m \log n)$ trees and it is not clear how to use it to obtain an almost linear time algorithm. Flow sparsifiers answer this problem by embedding $G$ into a single graph $G^{\prime}$, which is larger than a tree, but still sparse. Moreover, they provide an explicit efficient embedding of $G^{\prime}$ into $G$. Interestingly, the embedding from $G$ to $G^{\prime}$ is not necessary for our notion of flow sparsifier, and is replaced by the cut-approximation guarantee. This requirement, together with the application of the flow-cut gap [21], lets us argue that the optimal congestion of a $k$-commodity flow problem can change at most by a factor of $O(\log k)$ between $G$ and $G^{\prime}$.

## Main Theorem on Flow Sparsifiers and Proof of Theorem 12.4.9

The main goal of this section will be to prove the following theorem:
Theorem 12.5.3. For any constant $\varepsilon \in(0,1)$, there is an $(\tilde{O}(n \log U), \varepsilon, \tilde{O}(1))$-flow sparsifier.
Assuming Theorem 12.5.3, we can now prove Theorem 12.4.9, the main theorem necessary for edge reduction in our construction of low-congestion projections.

Proof of Theorem 12.4.9. We apply the flow sparsifier of Theorem 12.5.3 to $G=(V, E, \boldsymbol{\mu})$ and obtain output $G^{\prime}=\left(V, E^{\prime}, \boldsymbol{\mu}^{\prime}\right)$ with embedding $\boldsymbol{M}$. By the definition of flow sparsifier, we know that the capacity ratio $U^{\prime}$ of $G^{\prime}$ is at most $U \cdot \operatorname{poly}(|V|)$, as required. Moreover, again by Theorem 12.5.3, $G^{\prime}$ has at most $\tilde{O}(|V| \log U)$ edges. Given an oblivious routing $\boldsymbol{A}^{\prime}$ on $G^{\prime}$ consider the oblivious routing $\boldsymbol{A} \stackrel{\text { def }}{=} \boldsymbol{M} \boldsymbol{A}^{\prime}$. By the definition of flow sparsifier, we have that $\mathcal{T}(\boldsymbol{M})=\widetilde{O}(|E| \log U)$. Hence $\mathcal{T}(\boldsymbol{A})=$ $\mathcal{T}(\boldsymbol{M})+\mathcal{T}\left(\boldsymbol{A}^{\prime}\right)=\widetilde{O}(|E| \log U)+\mathcal{T}\left(\boldsymbol{A}^{\prime}\right)$. To complete the proof, we bound the competivite ratio $\rho(\boldsymbol{A})$. Using the same argument as in Lemma 12.4.3, we can write $\rho(\boldsymbol{A})$ as

$$
\rho(\boldsymbol{A})=\max _{D} \frac{\operatorname{cong}_{G}(\{\boldsymbol{A} D\})}{\operatorname{OPT}_{G}(D)} \leq \max _{D_{\infty}} \frac{\operatorname{cong}_{G}\left(\boldsymbol{A} D_{\infty}\right)}{\operatorname{OPT}_{G}\left(D_{\infty}\right)}
$$

where $D_{\infty}$ is the set of demands that result by taking the routing of every demand in $D$ by $\operatorname{OPT}(D)$ and splitting it up into demands on every edge corresponding to the flow sent by $\operatorname{OPT}(D)$. Notice that $D_{\infty}$ has at most $|E|$ demands that are routed between pairs of vertices in $V$. Then, because $G^{\prime}$ is an $\varepsilon$-cut approximation to $G$, the flow-cut gap of Aumann and Rabani [21] guarantees that

$$
\operatorname{OPT}_{G}\left(D_{\infty}\right) \geq \frac{1}{O(\log n)} \operatorname{OPT}_{G^{\prime}}\left(D_{\infty}\right)
$$

As a result, we obtain:

$$
\begin{aligned}
\rho(\boldsymbol{A}) & \leq O(\log n) \cdot \max _{D_{\infty}} \frac{\operatorname{cong}_{G}\left(\boldsymbol{A} D_{\infty}\right)}{\operatorname{OPT}_{G^{\prime}}\left(D_{\infty}\right)}=O(\log n) \cdot \max _{D_{\infty}} \frac{\operatorname{cong}_{G}\left(\boldsymbol{M} \boldsymbol{A}^{\prime} D_{\infty}\right)}{\operatorname{OPT}_{G^{\prime}}\left(D_{\infty}\right)} \\
& \leq O(\log n) \cdot \operatorname{cong}(\boldsymbol{M}) \cdot \max _{D_{\infty}} \frac{\operatorname{cong}_{G^{\prime}}\left(\boldsymbol{A}^{\prime} D_{\infty}\right)}{\operatorname{OPT}_{G^{\prime}}\left(D_{\infty}\right)} \leq \widetilde{O}\left(\rho\left(\boldsymbol{A}^{\prime}\right)\right) .
\end{aligned}
$$

## Techniques

We will construct flow sparsifiers by taking as a starting point the construction of spectral sparsifiers of Spielman and Teng [244]. Their construction achieves a sparsity of $\widetilde{O}\left(\frac{n}{\varepsilon^{2}}\right)$ edges, while guaranteeing an $\varepsilon$-spectral approximation. As the spectral approximation implies the cut approximation, the construction in [244] suffices to meet the first two conditions in Definition 12.5.1. Moreover, their algorithm also runs in time $\tilde{O}(m)$, meeting the fourth condition. Hence, to complete the proof of Theorem 12.5.3, we will modify the construction of Spielman and Teng to endow their sparsifier $G^{\prime}$ with an embedding $\boldsymbol{M}$ onto $G$ of low congestion that can be both computed and invoked efficiently. The main tool we use in constructing $\boldsymbol{M}$ is the notion of electrical-flow routing and the fact that electrical-flow routing schemes achieve a low competitive ratio on near-expanders and subsets thereof [139, 158].

To exploit this fact and construct a flow sparsifier, we follow Spielman and Teng [244] and partition the input graph into vertex sets, where each sets induces a near-expanders and most edges of the graph do not cross set boundaries. We then sparsify these induced subgraphs using standard sparsification techniques and iterate on the edges not in the subgraphs. As each iteration removes a constant fraction of the edges, by using standard sparsification techniques, we immediately obtain the sparsity and cut approximation properties. To obtain the embedding $\boldsymbol{M}$ with $\operatorname{cong}(\boldsymbol{M})=\widetilde{O}(1)$, we prove a generalization of results in $[139,158]$ and show that the electrical-flow routing achieves a low competitive ratio on near-expanders and subsets thereof.

In the next two subsections, we introduce the necessary concept about electrical-flow routing and prove that it achieves low competitive ratio over near-expanders (and subsets of near-expanders).

## ■ 12.5.1 Subgraph Routing

Given an oblivious routing strategy $\boldsymbol{A}$, we may be interested only in routing demands coming from a subset of edge $F \subseteq E$. In this setting, given a set of demands $D$ routable in $F$, we let $\operatorname{OPT}^{F}(D)$ denote the minimal congestion achieved by any routing restricted to only sending flow on edges in $F$ and we measure the $F$-competitive ratio of $\boldsymbol{A}$ by

$$
\rho^{F}(\boldsymbol{A}) \stackrel{\text { def }}{=} \max _{D \text { routable in } F} \frac{\operatorname{cong}(\boldsymbol{A} D)}{\operatorname{OPT}^{F}(D)}
$$

Note that $\boldsymbol{A}$ may use all the edges in $G$ but $\rho^{F}(\boldsymbol{A})$ compares it only against routings that are restricted to use only edges in $F$. As before, we can upper bound the $F$-competitive ratio $\rho^{F}(\boldsymbol{A})$ by operator
norms.
Lemma 12.5.4. Let $\mathbf{1}_{F} \in \mathbb{R}^{E}$ denote the indicator vector for set $F$ (i.e. $\mathbf{1}_{F}(e)=1$ if $e \in F$ and $\mathbf{1}_{F}(e)=0$ ) and let $\boldsymbol{I}_{F} \stackrel{\text { def }}{=} \operatorname{diag}\left(\mathbf{1}_{F}\right)$. For any $F \subseteq E$ we have

$$
\rho^{F}(\boldsymbol{A})=\left\|\boldsymbol{U}^{-1} \boldsymbol{A} \boldsymbol{B}^{T} \boldsymbol{U} \boldsymbol{I}_{F}\right\|_{\infty}
$$

Proof. We use the same reasoning as in the non-subgraph case. For a set of demands $D=\left\{\chi_{i}\right\}$, we consider $D_{\infty}^{F}$, the demands on the edges in $F$ used by $\operatorname{OPT}^{F}(D)$. Then, it is the case that $\operatorname{OPT}^{F}(D)=$ $\operatorname{OPT}^{F}\left(D_{\infty}\right)$ and we know that cost of obliviously routing $D_{P}$ is greater than the cost of obliviously routing $D$. Therefore, we have:

$$
\begin{aligned}
& \rho^{F}=\max _{x \in \mathbb{R}^{E}}: \boldsymbol{I}_{E \backslash F} \boldsymbol{x}=0 \\
&=\max _{\boldsymbol{y} \in \mathbb{R}^{E}}: \boldsymbol{I}_{E \backslash F}\left|\boldsymbol{U}^{-1} \boldsymbol{A} \boldsymbol{B}^{T} \mathbf{1}_{e} \boldsymbol{x}_{e}\right| \|_{\infty} \\
&\left\|\boldsymbol{U}^{-1} \boldsymbol{x}\right\|_{\infty} \\
&=\max _{\boldsymbol{y} \in \mathbb{R}^{E}} \frac{\left\|\sum_{e \in E}\left|\boldsymbol{U}^{-1} \boldsymbol{A} \boldsymbol{B}^{T} \boldsymbol{U} \mathbf{1}_{e} \boldsymbol{y}_{e}\right|\right\|_{\infty}}{\|\boldsymbol{y}\|_{\infty}} \\
&\|\boldsymbol{y}\|_{\infty} \boldsymbol{A} \boldsymbol{B}^{T} \boldsymbol{U} \boldsymbol{I}_{F} \mathbf{1}_{e} \boldsymbol{y}_{e} \mid \|_{\infty} \\
&\left\|\left|\boldsymbol{U}^{-1} \boldsymbol{A} \boldsymbol{B}^{T} \boldsymbol{U} \boldsymbol{I}_{F}\right|\right\|_{\infty}=\left\|\boldsymbol{U}^{-1} \boldsymbol{A} \boldsymbol{B}^{T} \boldsymbol{U} \boldsymbol{I}_{F}\right\|_{\infty}
\end{aligned}
$$

(Having $y_{e} \neq 0$ for $e \in E \backslash F$ decreases the ratio.)

## ■ 12.5.2 Electrical-Flow Routings

In this section, we define the notion of electrical-flow routing and prove the results necessary to construct flow sparsifiers. Recall that $\boldsymbol{R}$ is the diagonal matrix of resistances and the Laplacian $\mathcal{L}$ is defined as $\boldsymbol{B}^{T} \boldsymbol{R}^{-1} \boldsymbol{B}$. For the rest of this section, we assume that resistances are set as $\boldsymbol{R}=\boldsymbol{U}^{-1}$.
Definition 12.5.5. Consider a graph $G=(V, E, \mu)$ and set the edge resistances as $r_{e}=\frac{1}{\mu_{e}}$ for all $e \in E$. The oblivious electrical-flow routing strategy is the linear operator $\boldsymbol{A}_{\mathcal{E}}$ defined as

$$
\boldsymbol{A}_{\mathcal{E}} \stackrel{\text { def }}{=} \boldsymbol{R}^{-1} \boldsymbol{B} \mathcal{L}^{+}
$$

In words, the electrical-flow routing strategy is the routing scheme that, for each demand $\chi$ sends the electrical flow with boundary condition $\chi$ on the graph $G$ with resistances $\boldsymbol{R}=\boldsymbol{U}^{-1}$.

For the electrical-flow routing strategy $\boldsymbol{A}_{\mathcal{E}}$, the upper bound on the competitive ratio $\rho\left(\boldsymbol{A}_{\mathcal{E}}\right)$ in Lemma 12.4.3 can be rephrased in terms of the voltages induced on $G$ by electrically routing an edge $e \in E$. This interpretation appears in [139, 158].
Lemma 12.5.6. Let $\boldsymbol{A}_{\mathcal{E}}$ be the electrical-flow routing strategy. For an edge $e \in E$, we let the voltage vector $\boldsymbol{v}_{e} \in \mathbb{R}^{V}$ be given by $\boldsymbol{v}_{e} \stackrel{\text { def }}{=} \mathcal{L}^{+} \boldsymbol{\chi}_{e}$. We then have

$$
\rho\left(\boldsymbol{A}_{\mathcal{E}}\right)=\max _{e \in E} \sum_{(a, b) \in E} \frac{\left|v_{e}(a)-v_{e}(b)\right|}{r_{a b}}
$$

Proof. We have:

$$
\rho\left(\boldsymbol{A}_{\mathcal{E}}\right)=\left\|\boldsymbol{B} \mathcal{L}^{+} \boldsymbol{B}^{T} \boldsymbol{R}^{-1}\right\|_{\infty}=\max _{e \in E}\left\|\boldsymbol{R}^{-1} \boldsymbol{B} \mathcal{L}^{+} \boldsymbol{B}^{T} \mathbf{1}_{e}\right\|_{1}=\max _{e \in E} \sum_{(a, b) \in E} \frac{\left|v_{e}(a)-v_{e}(b)\right|}{r_{a b}}
$$

The same reasoning can be extended to the subgraph-routing case to obtain the following lemma.
Lemma 12.5.7. For $F \subseteq E$ and $\boldsymbol{R}=\boldsymbol{U}^{-1}$ we have

$$
\rho^{F}\left(\boldsymbol{A}_{\mathcal{E}}\right)=\max _{e \in E} \sum_{(a, b) \in F} \frac{\left|v_{e}(a)-v_{e}(b)\right|}{r_{a b}} .
$$

Proof. As before, we have:

$$
\begin{align*}
\rho^{F}\left(\boldsymbol{A}_{\mathcal{E}}\right) & =\left\|\boldsymbol{B} \mathcal{L}^{+} \boldsymbol{B}^{T} \boldsymbol{R}^{-1} \boldsymbol{I}_{F}\right\|_{\infty}  \tag{ByLemma12.5.4}\\
& =\max _{e \in E}\left\|\boldsymbol{I}_{F} \boldsymbol{R}^{-1} \boldsymbol{B} \mathcal{L}^{+} \boldsymbol{B}^{T} \mathbf{1}_{e}\right\|_{1}=\max _{e \in E} \sum_{(a, b) \in F} \frac{\left|v_{e}(a)-v_{e}(b)\right|}{r_{a b}}
\end{align*}
$$

### 12.5.2.1 Bounding the Congestion

In this section, we prove that we can bound the $F$-competitive ratio of the oblivious electrical-routing strategy as long as the edges $F$ that the optimum flow is allowed to route over are contained within an induced expander $G(U)=(U, E(U))$ for some $U \subseteq V$. Towards this we provide and prove the following lemma. This is a generalization of a similar lemma proved in [139].

Lemma 12.5.8. For weighted graph $G=(V, E, w)$ with integer weights and vertex subset $U \subseteq V$ the following holds:

$$
\rho^{F}\left(\boldsymbol{A}_{\mathcal{E}}\right) \leq \frac{8 \log (\operatorname{vol}(G(U)))}{\Phi(G(U))^{2}}
$$

Proof. By Lemma 12.5.7, for every edge $e \in E$,

$$
\rho^{F}\left(\boldsymbol{A}_{\mathcal{E}}\right) \leq\left\|\boldsymbol{I}_{E(U)} \boldsymbol{R}^{-1} \boldsymbol{B} \mathcal{L}^{+} \boldsymbol{\chi}_{e}\right\|_{1}
$$

Fix any edge $e \in E$ and let $v \stackrel{\text { def }}{=} \mathcal{L}^{+} \chi_{e}$. Recall that with this definition

$$
\begin{equation*}
\left\|\boldsymbol{I}_{E(U)} \boldsymbol{R}^{-1} \boldsymbol{B} \mathcal{L}^{+} \boldsymbol{\chi}_{e}\right\|_{1}=\sum_{(a, b) \in E(U)} \frac{|v(a)-v(b)|}{r_{a b}}=\sum_{(a, b) \in E(U)} w_{a b} \cdot|v(a)-v(b)| \tag{12.3}
\end{equation*}
$$

We define the following vertex subsets:

$$
\forall x \in \mathbb{R} \quad: \quad S_{\bar{x}}^{\leq} \xlongequal{\text { def }}\{a \in U \mid v(a) \leq x\} \quad \text { and } \quad S_{\bar{x}}^{\geq} \stackrel{\text { def }}{=}\{a \in U \mid v(a) \geq x\}
$$

Since adding a multiple of the all-ones vector to $v$ does not change the quantity of interest in Equation 12.3 , we can assume without loss of generality that

$$
\operatorname{vol}_{G(U)}\left(S_{0}^{\geq}\right) \geq \frac{1}{2}(\operatorname{vol}(G(U))) \quad \text { and } \quad \operatorname{vol}_{G(U)}\left(S_{0}^{\leq}\right) \geq \frac{1}{2}(\operatorname{vol}(G(U))) .
$$

For any vertex subset $S \subseteq U$, we denote the flow out of $S$ and the weight out of $S$ by

$$
f(S) \stackrel{\text { def }}{=} \sum_{e=(a, b) \in E(U) \cap \partial(S)} w_{e}|v(a)-v(b)|, \quad \text { and } \quad w(S) \stackrel{\text { def }}{=} \sum_{e \in E(U) \cap \partial(S)} w_{e} .
$$

At this point, we define a collections of subsets $\left\{C_{i} \in S_{0}^{\geq}\right\}$. For an increasing sequence of real numbers
$\left\{c_{i}\right\}$, we let $C_{i} \stackrel{\text { def }}{=} S_{c_{i}}^{\geq}$and we define the sequence $\left\{c_{i}\right\}$ inductively as follows:

$$
c_{0} \stackrel{\text { def }}{=} 0 \quad, \quad c_{i} \stackrel{\text { def }}{=} c_{i-1}+\Delta_{i-1} \quad, \text { and } \quad \Delta_{i} \stackrel{\text { def }}{=} 2 \frac{f\left(C_{i}\right)}{w\left(C_{i}\right)} .
$$

In words, the $c_{i+1}$ equals the sum of $c_{i}$ and an increase $\Delta_{i}$ which depends on how much the cut $\delta\left(C_{i}\right) \cap E(U)$ was congested by the electrical flow.

Now, $l_{i} \stackrel{\text { def }}{=} w\left(\partial_{E(U)}\left(C_{i-1}\right)-\partial_{E(U)}\left(C_{i}\right)\right)$, i.e. the weight of the edges in $E(U)$ cut by $C_{i-1}$ but not by $C_{i}$. We get

$$
\begin{array}{rlr}
\operatorname{vol}\left(C_{i+1}\right) & \leq \operatorname{vol}\left(C_{i}\right)-l_{i} & \\
& \leq \operatorname{vol}\left(C_{i}\right)-\frac{w\left(C_{i}\right)}{2} & \text { (By choice of } \left.l_{i} \text { and } \Delta_{i}\right) \\
& \leq \operatorname{vol}\left(C_{i}\right)-\frac{1}{2} \operatorname{vol}\left(C_{i}\right) \Phi(G(U)) & \text { (Definition of conductance) }
\end{array}
$$

Applying this inductively and using our assumption on $\operatorname{vol}\left(S_{0}^{\geq}\right)$we have that

$$
\operatorname{vol}\left(C_{i}\right) \leq\left(1-\frac{1}{2} \Phi(G(U))\right)^{i} \operatorname{vol}\left(C_{0}\right) \leq \frac{1}{2}\left(1-\frac{1}{2} \Phi(G(U))\right)^{i} \operatorname{vol}(G(U))
$$

Since $\phi(G(U)) \in(0,1)$, for $i+1=\frac{2 \log (\operatorname{vol}(G(U)))}{\Phi(G(U))}$ we have that $\operatorname{vol}\left(S_{i}\right) \leq \frac{1}{2}$. Since $\operatorname{vol}\left(S_{i}\right)$ decreases monotonically with $i$, if we let $r$ be the smallest value such that $C_{r+1}=\emptyset$, we must have

$$
r \leq \frac{2 \cdot \log (\operatorname{vol}(G(U)))}{\Phi(G(U))}
$$

Since $v$ corresponds to a unit flow, we know that $f\left(C_{i}\right) \leq 1$ for all $i$. Moreover, by the definition of conductance we know that $w\left(C_{i}\right) \geq \Phi(G(U)) \cdot \operatorname{vol}\left(C_{i}\right)$. Therefore,

$$
\Delta_{i} \leq \frac{2}{\Phi(G(U)) \cdot \operatorname{vol}\left(C_{i}\right)}
$$

We can now bound the contribution of $C_{0}^{\geq}$to the volume of the linear embedding $v$. In the following, for a vertex $a \in V$, we let $d(a) \stackrel{\text { def }}{=} \sum_{e=\{a, b\} \in E(U)} w_{e}$ be the degree of $a$ in $E(U)$.

$$
\begin{aligned}
\sum_{a \in C_{0}^{\geq}} d(a) v(a) & =\sum_{i=0}^{r}\left[\sum_{a \in C_{i}-C_{i+1}} d(a) v(a)\right] \\
& \leq \sum_{i=0}^{r}\left[\sum_{a \in C_{i}-C_{i+1}} d(a)\left(\sum_{j=0}^{i} \Delta_{j}\right)\right] \quad \quad \quad\left(\text { By definition of } C_{i}\right) \\
& \leq \sum_{i=0}^{r}\left[\left(\operatorname{vol}\left(C_{i}\right)-\operatorname{vol}\left(C_{i+1}\right)\right) \cdot\left(\sum_{j=0}^{i} \Delta_{j}\right)\right] \\
& \left.=\sum_{i=0}^{r} \operatorname{vol}\left(C_{i}\right) \Delta_{i} \leq \frac{2 r}{\Phi(G(U))} \quad \text { (Rearrangement and fact that } \operatorname{vol}\left(C_{r+1}\right)=0\right)
\end{aligned}
$$

By repeating the same argument on $S_{0}^{\leq}$, we get that $\sum_{a \in S_{0}^{\leq}} d(a) v(a) \leq \frac{2 r}{\Phi(G(U))}$. Putting this all
together yields

$$
\left\|\boldsymbol{I}_{E(U)} \boldsymbol{R}^{-1} \boldsymbol{B} \mathcal{L}^{+} \boldsymbol{\chi}_{e}\right\|=\sum_{(a, b) \in G(U)} w_{a b} \cdot|v(a)-v(b)| \leq \sum_{a \in G(U)} d(a) v(a) \leq \frac{4 r}{\Phi(G(U))}
$$

From this lemma and Lemma 12.5.7, the following is immediate:
Lemma 12.5.9. Let $F \subseteq E$ be contained within some vertex induced subgraph $G(U)$, then for $\boldsymbol{R}=$ $\boldsymbol{U}^{-1}$ we have

$$
\rho^{F}\left(\boldsymbol{R}^{-1} \boldsymbol{B} \mathcal{L}^{+}\right) \leq \rho^{E(U)}\left(\boldsymbol{R}^{-1} \boldsymbol{B} \mathcal{L}^{+}\right) \leq \frac{8 \log (\operatorname{vol}(G(U)))}{\Phi(G(U))^{2}}
$$

## - 12.5.3 Construction and Analysis of Flow Sparsifiers

In the remainder of this section we show how to produce an efficient $O\left(\log ^{c}\right)$-flow sparsifier for some fixed constant $c$, proving Theorem 12.5.3. In this chapter, we make no attempt to optimize the value of $c$. For the rest of this section, we again assume that we choose the resistance of an edge to be the the inverse of its capacity, i.e. $\boldsymbol{U}=\boldsymbol{W}=\boldsymbol{R}^{-1}$.

As discussed before, our approach follows closely that of Spielman and Teng [244] to the construction of spectral sparsifiers. The first step of this line of attack is to reduce the problem to the unweighted case.

Lemma 12.5.10. Given an $(h, \varepsilon, \alpha)$-flow-sparsifier algorithm for unweighted graphs, it is possible to construct an $(h \cdot \log U, \varepsilon, \alpha)$-flow-sparsifier algorithm for weighted graphs $G=(V, E, \mu)$ with capacity ratio $U$ obeying

$$
U=\frac{\max _{e \in E} \mu_{e}}{\min _{e \in E} \mu_{e}}=\operatorname{poly}(|V|)
$$

Proof. We write each edge in binary so that $G=\sum_{i=0}^{\log U} 2^{i} G_{i}$ for some unweighted graphs $\left\{G_{i}=\right.$ $\left.\left(V, E_{i}\right\}_{i \in[\log U]}\right\}$, where $\left|E_{i}\right| \leq m$ for all $i$. We now apply the unweighted flow-sparsifier to each $G_{i}$ in turn to obtain graphs $\left\{G_{i}^{\prime}\right\}$. We let $G^{\prime} \stackrel{\text { def }}{=} \sum_{i=0}^{\log U} 2^{i} G_{i}^{\prime}$ be the weighted flow-sparsified graph. By the assumption on the unweighted flow-sparsifier, each $G_{i}^{\prime}$ is $h$-sparse, so that $G^{\prime}$ must have at most $h \cdot \log U$ edges. Similarly, $G^{\prime}$ is an $\varepsilon$-cut approximation of $G$, as each $G_{i}^{\prime}$ is an $\varepsilon$-cut approximation of the corresponding $G_{i}$. Letting $\boldsymbol{M}_{i}$ be the embedding of $G_{i}^{\prime}$ into $G_{i}$, we can consider the embedding $\boldsymbol{M}=\sum_{i=0}^{\log U} 2^{i} \boldsymbol{M}_{i}$ of $G^{\prime}$ into $G$. As each $\boldsymbol{M}_{\boldsymbol{i}}$ has congestion bounded by $\alpha$, it must be the case that $\boldsymbol{M}$ also has congestion bounded by $\alpha$. The time to run the weighted flow sparsifier and to invoke $\boldsymbol{M}$ is now $\tilde{O}(m) \cdot \log U=\tilde{O}(m)$ by our assumption on $U$.

The next step is to construct a routine which flow-sparsifies a constant fraction of the edges of $E$. This routine will then be applied iteratively to produce the final flow-sparsifier.

Lemma 12.5.11. On input an unweighted graph $G=(V, E)$, there is an algorithm that runs in $\tilde{O}(m)$ and computes a partition of $E$ into $(F, \bar{F})$, an edge set $F^{\prime} \subseteq F$ with weight vector $\overrightarrow{w_{F^{\prime}}} \in$ $\mathbb{R}^{E}$, support $\left(w_{F^{\prime}}\right)=F^{\prime}$, and an embedding $\boldsymbol{H}: \mathbb{R}^{F^{\prime}} \rightarrow \mathbb{R}^{E}$ with the following properties:

1. F contains most of the volume of $G$, i.e.

$$
|F| \geq \frac{|E|}{2}
$$

2. $F^{\prime}$ contains only $\tilde{O}(n)$ edges, i.e. $\left|F^{\prime}\right| \leq \tilde{O}(n)$.
3. The weights $w_{F^{\prime}}$ are bounded

$$
\forall e \in F^{\prime} \quad, \quad \frac{1}{\operatorname{poly}(n)} \leq w_{F^{\prime}}(e) \leq n
$$

4. The graph $H^{\prime}=\left(V, F^{\prime}, w_{F^{\prime}}\right)$ is an $\varepsilon$-cut approximation to $H=(V, F)$, i.e. for all $S \subseteq V$ :

$$
(1-\varepsilon)\left|\partial_{H}(S)\right| \leq w_{F^{\prime}}\left(\partial_{H^{\prime}}(S)\right) \leq(1+\varepsilon)\left|\partial_{H}(S)\right|
$$

5. The embedding $\boldsymbol{H}$ from $H=\left(V, F^{\prime}, w_{F^{\prime}}\right)$ to $G$ has bounded congestion

$$
\operatorname{cong}(\boldsymbol{H})=\widetilde{O}(1)
$$

and can be applied in time $\tilde{O}(m)$.
Given Lemma 12.5 .10 and Lemma 12.5.11, it is straightforward to complete the proof of Theorem 12.5.3.

Proof. Using Lemma 12.5.10, we reduce the objective of Theorem 12.5 .3 to running a $(\widetilde{O}(n), \varepsilon, \widetilde{O}(1))$ flow sparsifier on $\log U$ unweighted graphs. To construct this unweighted flow sparsifier, we apply Lemma 12.5 .11 iteratively as follows. Starting with the instance unweighted graph $G_{1}=\left(V, E_{1}\right)$, we run the algorithm of Lemma 12.5 .11 on the current graph $G_{t}=\left(V, E_{t}\right)$ to produce the sets $F_{t}$ and $F_{t}^{\prime}$, the weight vector $w_{F_{t}^{\prime}}$ and the embedding $\boldsymbol{H}_{t}: \mathbb{R}^{F_{t}^{\prime}} \rightarrow R^{E}$. To proceed to the next iteration, we then define $E_{t+1} \stackrel{\text { def }}{=} E_{t} \backslash F_{t}$ and move on to $G_{t+1}$.

By Lemma 12.5.11, at every iteration $t,\left|F_{t}\right| \geq \frac{1}{2} \cdot\left|E_{t}\right|$, so that $\left|E_{t+1}\right| \leq \frac{1}{2} \cdot\left|E_{t}\right|$. This shows that there can be at most $T \leq \log \left(\left|E_{1}\right|\right)=O(\log n)$ iterations.

After the last iteration $T$, we have effectively partitioned $E_{1}$ into disjoint subsets $\left\{F_{t}\right\}_{t \in[T]}$, where each $F_{t}$ is well-approximated but the weighted edgeset $F_{t}^{\prime}$. We then output the weighted graph $G^{\prime}=$ $\left(V, E^{\prime} \stackrel{\text { def }}{=} \cap_{t=1}^{T} F_{t}^{\prime}, w^{\prime} \stackrel{\text { def }}{=} \sum_{t=1}^{T} w_{F_{t}^{\prime}}\right)$, which is the sum of the disjoint weighted edges sets $\left\{F_{t}^{\prime}\right\}_{t \in[T]}$. We also output the embedding $M: \mathbb{R}^{E^{\prime}} \rightarrow \mathbb{R}^{E}$ from $G^{\prime}$ to $G$, defined as the direct sum

$$
\boldsymbol{M}=\bigoplus_{t=1}^{T} \boldsymbol{H}_{\boldsymbol{t}}
$$

In words, $\boldsymbol{M}$ maps an edge $e^{\prime} \in E^{\prime}$ by finding $t$ for which $e^{\prime} \in F_{t}^{\prime}$ and applying the corresponding $\boldsymbol{H}_{\boldsymbol{t}}$.
We are now ready to prove that this algorithm with output $G^{\prime}$ and $\boldsymbol{M}$ is an efficient $(\widetilde{O}(n), \varepsilon, \widetilde{O}(n))$ flow sparsifier. To bound the capacity ratio $U^{\prime}$ of $G^{\prime}$, we notice that

$$
U^{\prime} \leq \max _{t} \frac{\max _{e \in F_{t}^{\prime}} w_{F_{t}^{\prime}}(e)}{\min _{e \in F_{t}^{\prime}} w_{F_{t}^{\prime}}(e)} \leq \operatorname{poly}(n)
$$

where we used the fact that the sets $F_{t}^{\prime}$ are disjoint and the guarantee on the range of $w_{F_{t}^{\prime}}$.
Next, we bound the sparsity of $G^{\prime}$. By Lemma $12.5 .11, F_{t}^{\prime}$ contains at most $\widetilde{O}(n)$ edges. As a result, we get the required bound

$$
\left|E^{\prime}\right|=\sum_{t=1}^{T}\left|F_{t}^{\prime}\right| \leq \widetilde{O}(T n)=\widetilde{O}(n)
$$

For the cut approximation, we consider any $S \subseteq V$. By the cut guarantee of Lemma 12.5.11, we have that, for all $t \in[T]$,

$$
(1-\varepsilon)\left|\partial_{G}(S) \cap F_{t}\right| \leq w_{F_{t}^{\prime}}\left(\partial_{G}(S) \cap F_{t}^{\prime}\right) \leq(1+\varepsilon)\left|\partial_{G}(S) \cap F_{t}\right|
$$

Summing over all $t$, as $E^{\prime}=\dot{\bigcup} F_{t}^{\prime}$ and $E=\dot{U} F_{t}$, we obtain the required approximation

$$
(1-\varepsilon)\left|\partial_{G}(S)\right| \leq w^{\prime}\left(\partial_{G^{\prime}}(S)\right) \leq(1+\varepsilon)\left|\partial_{G}(S)\right| .
$$

The congestion of $\boldsymbol{M}$ can be bounded as follows

$$
\operatorname{cong}(\boldsymbol{M}) \leq \sum_{t=1}^{T} \operatorname{cong}\left(\boldsymbol{H}_{\boldsymbol{t}}\right)=\widetilde{O}(T)=\widetilde{O}(1)
$$

To conclude the proof, we address the efficiency of the flow sparsifier. The algorithm applies the routine of Lemma 12.5 .11 for $T=\widetilde{O}(1)$ times and hence runs in time $\widetilde{O}(m)$, as required. Invoking the embedding $\boldsymbol{M}$ requires invoking each of the $T$ embeddings $\boldsymbol{H}_{\boldsymbol{t}}$. This takes time $\widetilde{O}(T m)=\widetilde{O}(m)$.

### 12.5.3.1 Flow Sparsification of Unweighted Graphs: Proof of Lemma 12.5.11

In this subsection, we prove Lemma 12.5.11. Our starting point is the following decomposition statement, which shows that we can form a partition of an unweighted graph where most edges do not cross the boundaries and the subgraphs induced within each set of this partition are near-expanders. The following lemma is implicit in Spielman and Teng's local clustering approach to spectral sparsification [244].
Lemma 12.5.12 (Decomposition Lemma). For an unweighted graph $G=(V, E)$, in $\tilde{O}(m)$-time we can produce a partition $V_{1}, \ldots, V_{k}$ of $V$ and a collection of sets $S_{1}, \ldots, S_{k} \subseteq V$ with the following properties:

- For all $i, S_{i}$ is contained in $V_{i}$.
- For all $i$, there exists a set $T_{i}$ with $S_{i} \subseteq T_{i} \subseteq V_{i}$, such that

$$
\Phi\left(G\left(T_{i}\right)\right) \geq \Omega\left(\frac{1}{\log ^{2} n}\right)
$$

- At least half of the edges are found within the sets $\left\{S_{i}\right\}$, i.e.

$$
\sum_{i=1}^{k}\left|E\left(S_{i}\right)\right|=\sum_{i=1}^{k}\left|\left\{e=\{a, b\}: a \in S_{i}, b \in S_{i}\right\}\right| \geq \frac{1}{2}|E| .
$$

To design an algorithm satisfying the requirements of Lemma 12.5 .11 , we start by appling the Decomposition Lemma to our unweighted input graph $G=(V, E)$ to obtain the partition $\left\{V_{i}\right\}_{i \in[k]}$ and the sets $\left\{S_{i}\right\}_{i \in[k]}$. We let $G_{i} \stackrel{\text { def }}{=}\left(S_{i}, E\left(S_{i}\right)\right)$. To reduce the number of edges, while preseving cuts, we apply a spectral sparsification algorithm to each $G_{i}$. Concretely, by applying the spectral sparsification by effective resistances of Spielman and Srivastava [242] to each $G_{i}$, we obtain weighted graphs $G_{i}^{\prime}=\left(S_{i}, E_{i}^{\prime} \subseteq E\left(S_{i}\right), w_{i}^{\prime}\right)$ in time $\sum_{i=1}^{k} \widetilde{O}\left(\left|E\left(S_{i}\right)\right|\right) \leq \widetilde{O}(|E|)$ with $\left|E_{i}^{\prime}\right| \leq \widetilde{O}\left(\left|S_{i}\right|\right)$ and the property that cuts are preserved ${ }^{6}$ for all $i$ :

$$
\forall S \subseteq S_{i} \quad, \quad(1-\varepsilon) \cdot\left|\partial_{G_{i}}(S)\right| \leq w_{i}^{\prime}\left(\partial_{G_{i}^{\prime}}(S)\right) \leq(1+\varepsilon) \cdot\left|\partial_{G_{i}}(S)\right|
$$

Moreover, the spectral sparsification of [242] constructs the weights $\left\{w_{i}^{\prime}(e)\right\}_{e \in E_{i}^{\prime}}$ such that

$$
\forall e \in E_{i}^{\prime} \quad, \quad \frac{1}{\operatorname{poly}(n)} \leq \frac{1}{\operatorname{poly}\left(\left|S_{i}\right|\right)} \leq w_{i}^{\prime}(e) \leq\left|S_{i}\right| \leq n .
$$

[^27]To complete the description of the algorithm, we output the partition $(F, \bar{F})$ of $E$, where

$$
F \stackrel{\text { def }}{=} \bigcup_{i=1}^{k} E\left(S_{i}\right)
$$

We also output the set of weighted sparsified edges $F^{\prime}$.

$$
F^{\prime} \stackrel{\text { def }}{=} \bigcup_{i=1}^{k} E_{i}^{\prime}
$$

The weight $w_{F^{\prime}}(e)$ of edge $e \in F^{\prime}$ is given by finding $i$ such that $e \in E_{i}^{\prime}$ and setting $w_{F^{\prime}}(e)=w_{i}^{\prime}(e)$.
We now depart from Spielman and Teng's construction by endowing our $F^{\prime}$ with an embedding onto $G$. The embedding $\boldsymbol{H}: \mathbb{R}^{F^{\prime}} \rightarrow \mathbb{R}^{E}$ of the graph $H=\left(V, F^{\prime}, w_{F^{\prime}}\right)$ to $G$ is constructed by using the oblivious electrical-flow routing of $E\left(S_{i}\right)$ into $G\left(V_{i}\right)$. More specifically, as the sets $\left\{V_{i}\right\}$ partition $V$, the embedding $\boldsymbol{H}$ can be expressed as the following direct sum over the orthogonal subspaces $\left\{\mathbb{R}^{\left.E\left(V_{i}\right) \times E_{i}^{\prime}\right\} \text {. }}\right.$

$$
\boldsymbol{H} \stackrel{\text { def }}{=}\left(\bigoplus_{i=1}^{k} \boldsymbol{B}_{E\left(V_{i}\right)} \mathcal{L}_{G\left(V_{i}\right)}^{+} \boldsymbol{B}_{E\left(V_{i}\right)}^{T} \boldsymbol{I}_{\left(E\left(V_{i}\right), E_{i}^{\prime}\right)}\right)
$$

where $\boldsymbol{I}_{\left(E\left(V_{i}\right), E_{i}^{\prime}\right)}$ is the identity mapping of the edges $E_{i}^{\prime} \subseteq E\left(V_{i}\right)$ of $F^{\prime}$ over $V_{i}$ to the edges $E\left(V_{i}\right)$ of $V_{i}$ in $G$. Notice that there is no dependence on the resistances over $G$ as $G$ is unweighted.

This complete the description of the algorithm. We are now ready to give the proof of Lemma 12.5.11.
[Proof of Lemma 12.5.11]. The algorithm described above performs a decomposition of the input graph $G=(V, E)$ in time $\widetilde{O}(m)$ by the Decomposition Lemma. By the result of Spielman and Srivastava [242], each $G_{i}$ is sparsified in time $\widetilde{O}\left(\left|E\left(S_{i}\right)\right|\right)$. Hence, the sparsification step requires time $\widetilde{O}(m)$ as well. This shows that the algorithm runs in $\widetilde{O}(m)$-time, as required.

By the Decomposition Lemma, we know that $|F|=\sum_{i=1}^{k}\left|E\left(S_{i}\right)\right| \geq \frac{|E|}{2}$, which satisfies the requirement of the Lemma. Moreover, by the spectral sparsification result, we know that $\left|F^{\prime}\right|=\sum_{i=1}^{k}\left|E_{i}^{\prime}\right| \leq$ $\sum_{i=1}^{k} \widetilde{O}\left(\left|S_{i}\right|\right) \leq \widetilde{O}(n)$, as required. We also saw that by construction the weights $w_{F^{\prime}}$ are bounded:

$$
\forall e \in F^{\prime} \quad, \quad \frac{1}{\operatorname{poly}(n)} \leq w_{F^{\prime}}(e) \leq n
$$

To obtain the cut-approximation guarantee, we use the fact that for every $i$, by spectral sparsification,

$$
\forall S \subseteq S_{i} \quad, \quad(1-\varepsilon) \cdot\left|\partial_{G_{i}}(S)\right| \leq w_{i}^{\prime}\left(\partial_{G_{i}^{\prime}}(S)\right) \leq(1+\varepsilon) \cdot\left|\partial_{G_{i}}(S)\right|
$$

We have $H^{\prime}=\left(V, F^{\prime}, w_{F^{\prime}}\right)$ and $H=(V, F)$. Consider now $T \subseteq V$ and apply the previous bound to $T \cap S_{i}$ for all $i$. Because $F^{\prime} \subseteq F=\cup_{i=1}^{k} E\left(S_{i}\right)$, we have that summing over the $k$ bounds yields

$$
\forall T \subseteq V \quad, \quad(1-\varepsilon)\left|\partial_{H}(T)\right| \leq w_{F^{\prime}}\left(\partial_{H^{\prime}}(T)\right) \leq(1+\varepsilon)\left|\partial_{H}(T)\right|
$$

which is the desired cut-approximaton guarantee.
Finally, we are left to prove that the embedding $\boldsymbol{H}$ from $H^{\prime}=\left(V, F^{\prime}, w_{F^{\prime}}\right)$ to $G=(V, E)$ has low congestion and can be applied efficiently. By definition of congestion,

$$
\operatorname{cong}(\boldsymbol{H})=\max _{\boldsymbol{x} \in \mathbb{R}^{F^{\prime}}} \frac{\|\boldsymbol{H} \boldsymbol{x}\|_{\infty}}{\left\|\boldsymbol{U}_{F^{\prime}}^{-1} \boldsymbol{x}\right\|_{\infty}}=\| \| \boldsymbol{H}\left|\boldsymbol{U}_{F^{\prime}} \mathbf{1}_{F^{\prime}}\left\|_{\infty}=\right\| \bigoplus_{i=1}^{k} \boldsymbol{B}_{E\left(V_{i}\right)} \mathcal{L}_{G\left(V_{i}\right)}^{+} \boldsymbol{B}_{E\left(V_{i}\right)}^{T} \boldsymbol{I}_{\left(E\left(V_{i}\right), E_{i}^{\prime}\right)}\right| \boldsymbol{U}_{F^{\prime}} \mathbf{1}_{F^{\prime}} \|_{\infty}
$$

Decomposing $\mathbb{R}^{E}$ into the subspaces $\left\{\mathbb{R}^{E\left(V_{i}\right)}\right\}$ and $\mathbb{R}^{F^{\prime}}$ into the subspaces $\left\{\mathbb{R}^{E_{i}^{\prime}}\right\}$ we have:

$$
\operatorname{cong}(\boldsymbol{H}) \leq \max _{i \in[k]}\| \| \boldsymbol{B}_{E\left(V_{i}\right)} \mathcal{L}_{G\left(V_{i}\right)}^{+} \boldsymbol{B}_{E\left(V_{i}\right)}^{T} \boldsymbol{I}_{\left(E\left(V_{i}\right), E_{i}^{\prime}\right)} \mid \boldsymbol{U}_{E_{i}^{\prime}} \mathbf{1}_{E_{i}^{\prime}} \|_{\infty}
$$

For each $i \in[k]$, consider now the set of demands $D_{i}$ over $V_{i}, D_{i} \stackrel{\text { def }}{=}\left\{\chi_{e}\right\}_{e \in E_{i}^{\prime}}$, given by the edges of $E_{i}^{\prime}$ with their capacities $w_{i}^{\prime}$. That is, $\chi_{e} \in \mathbb{R}^{V_{i}}$ is the demand corresponding to edge $e \in E_{i}^{\prime}$ with weight $w_{i}^{\prime}(e)$. Consider also the electrical routing $\boldsymbol{A}_{\mathcal{E}, i}=\boldsymbol{B}_{E\left(V_{i}\right)} \mathcal{L}_{G\left(V_{i}\right)}^{+}$over $G\left(V_{i}\right)$. Then:

$$
\operatorname{cong}(\boldsymbol{H}) \leq \max _{i \in[k]} \operatorname{cong}\left(\boldsymbol{A}_{\mathcal{E}, i} D_{i}\right)
$$

Notice that, by construction, $D_{i}$ is routable in $G_{i}^{\prime}=\left(S_{i}, E_{i}^{\prime}, w_{i}^{\prime}\right)$ and $\mathrm{OPT}_{G_{i}^{\prime}}\left(D_{i}\right)=1$. But, by our use of spectral sparsifiers in the construction, $G_{i}^{\prime}$ is an $\varepsilon$-cut approximation of $G_{i}$. Hence, by the flow-cut gap of Aumann and Rabani [21], we have:

$$
\mathrm{OPT}_{G_{i}}\left(D_{i}\right) \leq O\left(\log \left(\left|D_{i}\right|\right)\right) \cdot \mathrm{OPT}_{G_{i}^{\prime}}\left(D_{i}\right) \leq \widetilde{O}(1)
$$

When we route $D_{i}$ oblivious in $G\left(V_{i}\right)$, we can consider the $E\left(S_{i}\right)$-competitive ratio $\rho^{E\left(S_{i}\right)}\left(\boldsymbol{A}_{\mathcal{E}, i}\right)$ of the electrical routing $\boldsymbol{A}_{\mathcal{E}, i}=\boldsymbol{B}_{E\left(V_{i}\right)} \mathcal{L}_{G\left(V_{i}\right)}^{+}$, as $D_{i}$ is routable in $E\left(S_{i}\right)$, because $E_{i}^{\prime} \subseteq E\left(S_{i}\right)$. We have

$$
\operatorname{cong}(\boldsymbol{H}) \leq \max _{i \in[k]} \rho_{G\left(V_{i}\right)}^{E\left(S_{i}\right)}\left(\boldsymbol{A}_{\mathcal{E}, i}\right) \cdot \operatorname{OPT}_{G\left(V_{i}\right)}^{E\left(S_{i}\right)}\left(D_{i}\right)=\max _{i \in[k]} \rho_{G\left(V_{i}\right)}^{E\left(S_{i}\right)}\left(\boldsymbol{A}_{\mathcal{E}, i}\right) \cdot \mathrm{OPT}_{G_{i}}\left(D_{i}\right)
$$

Finally, putting these bounds together, we have:

$$
\operatorname{cong}(\boldsymbol{H}) \leq \max _{i \in[k]} \rho_{G\left(V_{i}\right)}^{E\left(S_{i}\right)}\left(\boldsymbol{A}_{\mathcal{E}, i}\right) \cdot \operatorname{OPT}_{G_{i}}\left(D_{i}\right) \leq \widetilde{O}(1) \cdot \max _{i \in[k]} \rho_{G\left(V_{i}\right)}^{E\left(S_{i}\right)}\left(\boldsymbol{A}_{\mathcal{E}, i}\right)
$$

But, by the Decomposition Lemma, there exists $T_{i}$ with $S_{i} \subseteq T_{i} \subseteq V_{i}$ such that

$$
\Phi\left(G\left(T_{i}\right)\right) \geq \Omega\left(\frac{1}{\log ^{2} n}\right)
$$

Then, by Lemma 12.5.9, we have that:

$$
\rho_{G\left(V_{i}\right)}^{E\left(S_{i}\right)}\left(\boldsymbol{A}_{\mathcal{E}, i}\right) \leq O\left(\frac{\log \operatorname{vol}\left(G\left(T_{i}\right)\right)}{\Phi\left(G\left(T_{i}\right)\right)^{2}}\right) \leq \widetilde{O}(1)
$$

This concludes the proof that $\operatorname{cong}(\boldsymbol{H}) \leq \widetilde{O}(1)$. To complete the proof of the Lemma, we just notice that $\boldsymbol{H}$ can be invoked in time $\widetilde{O}(m)$. A call of $\boldsymbol{H}$ involves solving $k$-electrical-problems, one for each $G\left(V_{i}\right)$. This can be done in time $\sum_{i=1}^{k} \widetilde{O}\left(\left|E\left(V_{i}\right)\right|\right) \leq \widetilde{O}(m)$, using any of the nearly-linear Laplacian system solvers available, such as [141].

## ■ 12.6 Removing Vertices in Oblivious Routing Construction

In this section we show how to reduce computing an efficient oblivious routing on a graph $G=(V, E)$ to computing an oblivious routing for $t$ graphs with $\widetilde{O}\left(\frac{|V|}{t}\right)$ vertices and at most $|E|$ edges. Formally we show

Theorem 12.6.1. Let $G=(V, E, \boldsymbol{\mu})$ be an undirected capacitated graph with capacity ratio $U$. For all $t>0$ in $\widetilde{O}(t \cdot|E|)$ time we can compute graphs $G_{1}, \ldots, G_{t}$ each with at most $\widetilde{O}\left(\frac{|E| \log (U)}{t}\right)$ vertices, at most $|E|$ edges, and capacity ratio at most $|V| \cdot U$, such that given oblivious routings $\boldsymbol{A}_{i}$ for each
$G_{i}$, in $\widetilde{O}(t \cdot|E|)$ time we can compute an oblivious routing $\boldsymbol{A} \in \mathbb{R}^{E \times V}$ on $G$ such that

$$
\mathcal{T}(\boldsymbol{A})=\widetilde{O}\left(t \cdot|E|+\sum_{i=1}^{t} \mathcal{T}\left(\boldsymbol{A}_{i}\right)\right) \quad \text { and } \quad \rho(\boldsymbol{A})=\widetilde{O}\left(\max _{i} \rho\left(\boldsymbol{A}_{i}\right)\right)
$$

We break this proof into several parts. First we show how to embed $G$ into a collection of $t$ graphs consisting of trees minus some edges which we call patrial tree embeddings (Section 12.6.1). Then we show how to embed a partial tree embedding in an "almost $j$-tree" [182], that is a graph consisting of a tree and a subgraph on at most $j$ vertices, for $j=2 t$ (Section 12.6.2). Finally, we show how to reduce oblivious routing on an almost $j$-tree to oblivious routing on a graph with at most $O(j)$ vertices by removing degree-1 and degree-2 vertices (Section 12.6.3). Finally, in Section 12.6.4 we put this all together to prove Theorem 12.4.10.

We remark that much of the ideas in the section were either highly influenced from [182] or are direct restatements of theorems from [182] adapted to our setting. We encourage the reader to look over that paper for further details regarding the techniques used in this section.

## ■ 12.6.1 From Graphs to Partial Tree Embeddings

To prove Theorem 12.4.10, we make heavy use of spanning trees and various properties of them. In particular, we use the facts that for every pair of vertices there is a unique tree path connecting them, that every edge in the tree induces a cut in the graph, and that we can embed a graph in a tree by simply routing ever edge over its tree path and that the congestion of this embedding will be determined by the load the edges place on tree edges. We define these quantities formally below.

Definition 12.6.2 (Tree Path). For undirected graph $G=(V, E)$, spanning tree $T$, and all $a, b \in V$ we let $P_{a, b} \subseteq E$ denote the unique path from $a$ to $b$ using only edges in $T$ and we let $\boldsymbol{p}_{a, b} \in \mathbb{R}^{E}$ denote the vector representation of this path corresponding to the unique vector sending one one unit from $a$ to $b$ that is nonzero only on $T$ (i.e. $\boldsymbol{B}^{T} \boldsymbol{p}_{a, b}=\boldsymbol{\chi}_{a, b}$ and $\forall e \in E \backslash T$ we have $\boldsymbol{p}_{a, b}(e)=0$ ).

Definition 12.6.3 (Tree Cuts). For undirected $G=(V, E)$ and spanning tree $T \subseteq E$ the edges cut by e, $\partial_{T}(F)$, and the edges cut by $F, \partial_{T}(e)$, are given by

$$
\partial_{T}(e) \stackrel{\text { def }}{=}\left\{e^{\prime} \in E \mid e^{\prime} \in P_{e}\right\} \quad \text { and } \quad \partial_{T}(F) \stackrel{\text { def }}{=} \cup_{e \in F} \partial(e)
$$

Definition 12.6.4 (Tree Load). For undirected capacitated $G=(V, E, \boldsymbol{\mu})$ and spanning tree $T \subseteq E$ the load on edge $e \in E$ by $T, \operatorname{cong}_{T}(e)$ is given by $\operatorname{load}_{T}(e)=\sum_{e^{\prime} \in E \mid e \in P_{e^{\prime}}} \boldsymbol{\mu}_{e^{\prime}}$

While these properties do highlight the fact that we could just embed our graph into a collection of trees to simplify the structure of our graph, this approach suffers from a high computational cost [225]. Instead we show that we can embed parts of the graph onto collections of trees at a lower computational cost but higher complexity. In particular we will consider what we call partial tree embeddings.

Definition 12.6.5 ([Partial Tree Embedding $\left.\left.{ }^{7}\right]\right)$. For undirected capacititated graph $G=(V, E, \boldsymbol{\mu})$, spanning tree $T$ and spanning tree subset $F \subseteq T$ we define the partial tree embedding graph $H=$ $H(G, T, F)=\left(V, E^{\prime}, \mu^{\prime}\right)$ to a be a graph on the same vertex set where $E^{\prime}=T \cup \partial_{T}(F)$ and

$$
\forall e \in E^{\prime} \quad: \quad \boldsymbol{\mu}^{\prime}(e)= \begin{cases}\operatorname{load}_{T}(e) & \text { if } e \in T \backslash F . \\ \boldsymbol{\mu}(e) & \text { otherwise }\end{cases}
$$

[^28]Furthermore, we let $\boldsymbol{M}_{H} \in \mathbb{R}^{E^{\prime} \times E}$ denote the embedding from $G$ to $H(G, T, F)$ where edges not cut by $F$ are routed over the tree and other edges are mapped to themselves.

$$
\forall e \in E \quad: \quad \boldsymbol{M}_{H}(e)= \begin{cases}\boldsymbol{p}_{e} & e \notin \partial_{T}(F) \\ \mathbf{1}_{e} & \text { otherwise }\end{cases}
$$

and we let $\boldsymbol{M}_{H}^{\prime} \in \mathbb{R}^{E \times E^{\prime}}$ denote the embeding from $H$ to $G$ that simply maps edges in $H$ to their corresponding edges in $G$, i.e. $\forall e \in E^{\prime}, \boldsymbol{M}_{H}^{\prime}(e)=\mathbf{1}_{e}$.

Note that by definition cong $\left(\boldsymbol{M}_{H}\right) \leq 1$, i.e. a graph embeds into its partial tree embedding with no congestion. However, to get embedding guarantees in the other direction more work is required. For this purpose we use a lemma from Madry [182] saying that we can construct a convex combination or a distribution of partial tree embeddings we can get such a guarantee.

Lemma 12.6.6 ([Probabilistic Partial Tree Embedding $\left.\left.{ }^{8}\right]\right)$. For any undirected capacitated graph $G=(V, E, \boldsymbol{\mu})$ and any $t>0$ in $\widetilde{O}(t \cdot m)$ time we can find a collection of partial tree embeddings $H_{1}=H\left(G, T_{1}, F_{1}\right), \ldots, H_{t}=H\left(G, T_{t}, F_{t}\right)$ and coefficients $\lambda_{i} \geq 0$ with $\sum_{i} \lambda_{i}=1$ such that $\forall i \in[t]$ we have $\left|F_{i}\right|=\widetilde{O}\left(\frac{m \log U}{t}\right)$ and such that $\sum_{i} \lambda_{i} \boldsymbol{M}_{H_{i}}^{\prime}$ embeds $G^{\prime}=\sum_{i} \lambda_{i} G_{i}$ into $G$ with congestion $\widetilde{O}(1)$.

Using this lemma, we can prove that we can reduce constructing an oblivious routing for a graph to constructing oblivious routings on several partial tree embeddings.

Lemma 12.6.7. Let the $H_{i}$ be graphs produced by Lemma 12.6.6 and for all $i$ let $\boldsymbol{A}_{i}$ be an oblivious routing algorithm for $H_{i}$. It follows that $\boldsymbol{A}=\sum_{i} \lambda_{i} \boldsymbol{M}_{H_{i}}^{\prime} \boldsymbol{A}_{i}$ is an oblivious routing on $G$ with $\rho(\boldsymbol{A}) \leq$ $\widetilde{O}\left(\max _{i} \rho\left(\boldsymbol{A}_{i}\right) \log n\right)$ and $\mathcal{T}(\boldsymbol{A})=O\left(\sum_{i} \mathcal{T}\left(\boldsymbol{A}_{i}\right)\right)$.

Proof. The proof is similar to the proof of Lemma 12.4.8. For all $i$ let $\boldsymbol{U}_{i}$ denote the capacity matrix of graph $G_{i}$. Then using Lemma 12.4.3 we get

$$
\rho(\boldsymbol{A})=\left\|\boldsymbol{U}^{-1} \boldsymbol{A} \boldsymbol{B}^{T} \boldsymbol{U}\right\|_{\infty}=\left\|\sum_{i=1}^{t} \lambda_{i} \boldsymbol{U}^{-1} \boldsymbol{M}_{H_{i}}^{\prime} \boldsymbol{A}_{i} \boldsymbol{B}^{T} \boldsymbol{U}\right\|_{\infty}
$$

Using that $\boldsymbol{M}_{H_{i}}$ is an embedding and therefore $\boldsymbol{B}_{H_{i}}^{T} \boldsymbol{M}_{H_{i}}=\boldsymbol{B}^{T}$ we get

$$
\rho(\boldsymbol{A})=\left\|\sum_{i=1}^{t} \lambda_{i} \boldsymbol{U}^{-1} \boldsymbol{M}_{H_{i}}^{\prime} \boldsymbol{A}_{i} \boldsymbol{B}_{H_{i}}^{T} \boldsymbol{M}_{H_{i}} \boldsymbol{U}\right\|_{\infty} \leq \max _{j, k}\left\|\sum_{i=1}^{t} \lambda_{i} \boldsymbol{U}^{-1} \boldsymbol{M}_{i}^{\prime} \boldsymbol{U}_{i}\right\|_{\infty} \cdot \rho\left(\boldsymbol{A}_{j}\right) \cdot \operatorname{cong}\left(\boldsymbol{M}_{H_{k}}\right)
$$

Since $\sum_{i} \lambda_{i} \boldsymbol{M}_{H_{i}}^{\prime}$ is an embedding of congestion of at most $\widetilde{O}(1)$ and $\operatorname{cong}\left(\boldsymbol{M}_{H_{k}}\right) \leq 1$ we have the desired result.

## ■ 12.6.2 From Partial Tree Embeddings To Almost-j-trees

Here we show how to reduce constructing an oblivious routing for a partial tree embedding to constructing an oblivious routing for what Madry [182] calls an "almost $j$-tree," the union of a tree plus a subgraph on at most $j$ vertices. First we define such objects and then we prove the reduction.

Definition 12.6.8 (Almost -tree). We call a graph $G=(V, E)$ an almost $j$-tree if there is a spanning tree $T \subseteq E$ such that the endpoints of $E \backslash T$ include at most $j$ vertices.

[^29]Lemma 12.6.9. For undirected capacitated $G=(V, E, \boldsymbol{\mu})$ and partial tree embedding $H=H(G, T, F)$ in $\widetilde{O}(|E|)$ time we can construct an almost $2 \cdot|F|$-tree $G^{\prime}=\left(V, E^{\prime}, \mu^{\prime}\right)$ with $\left|E^{\prime}\right| \leq|E|$ and an embedding $\boldsymbol{M}^{\prime}$ from $G^{\prime}$ to $H$ such that $H$ is embeddable into $G^{\prime}$ with congestion 2, $\operatorname{cong}\left(\boldsymbol{M}^{\prime}\right)=2$, and $\mathcal{T}\left(\boldsymbol{M}^{\prime}\right)=$ $\widetilde{O}(|E|)$.

Proof. For every $e=(a, b) \in E$, we let $v^{1}(e) \in V$ denote the first vertex on tree path $P_{(a, b)}$ incident to $F$ and we let $v^{2}(e) \in V$ denote the last vertex incident to $F$ on tree path $P_{(a, b)}$. Note that for every $e=(a, b) \in T$ we have that $\left(v^{1}(e), v^{2}(e)\right)=e$.

We define $G^{\prime}=\left(V, E^{\prime}, \boldsymbol{\mu}^{\prime}\right)$ to simply be the graph that consists of all these $\left(v^{1}(e), v^{2}(e)\right)$ pairs

$$
E^{\prime}=\left\{(a, b) \mid \exists e \in E \text { such that }(a, b)=\left(v^{1}(e), v^{2}(e)\right)\right\}
$$

and we define the weights to simply be the sums

$$
\forall e^{\prime} \in E^{\prime} \quad: \quad \boldsymbol{\mu}^{\prime}\left(e^{\prime}\right) \stackrel{\text { def }}{=} \sum_{e \in E \mid} \sum_{e=\left(v^{1}\left(e^{\prime}\right), v^{2}\left(e^{\prime}\right)\right)} \boldsymbol{\mu}(e)
$$

Now to embed $H$ in $G^{\prime}$ we define $\boldsymbol{M}$ by

$$
\forall e=(a, b) \in E \quad: \quad \boldsymbol{M} \mathbf{1}_{e}=\boldsymbol{p}_{a, v^{1}(e)}+\mathbf{1}_{\left(v^{1}(e), v^{2}(e)\right)}+\boldsymbol{p}_{v^{2}(e), b}
$$

and to embed $G^{\prime}$ in $H$ we define $\boldsymbol{M}^{\prime}$ by

$$
\forall e^{\prime} \in E \quad: \quad \boldsymbol{M}^{\prime} \mathbf{1}_{e^{\prime}}=\sum_{e=(a, b) \in E} \sum_{e^{\prime}=\left(v^{1}(e), v^{2}(e)\right)} \frac{\boldsymbol{\mu}(e)}{\boldsymbol{\mu}^{\prime}\left(e^{\prime}\right)}\left[\boldsymbol{p}_{v^{1}(e), a}+\mathbf{1}_{(a, b)}+\boldsymbol{p}_{b, v^{2}(e)}\right]
$$

In other words we route edges in $H$ along the tree until we encounter nodes in $F$ and then we route them along added edges and we simply route the other way for the reverse embedding. By construction clearly the congestion of the embedding in either direction is 2 .

To bound the running time, we note that by having every edge $e$ in $H$ maintain its $v^{1}(e)$ and $v^{2}(e)$ information, having every edge $e^{\prime}$ in $E^{\prime}$ maintain the set $\left\{e \in E \mid e^{\prime}=\left(v^{1}(e), v^{2}(e)\right)\right\}$ in a list, and using link cut trees [240] or the static tree structure in [141] to update information along tree paths we can obtain the desired value of $\mathcal{T}\left(\boldsymbol{M}^{\prime}\right)$.

## ■ 12.6.3 From Almost-J Trees to Less Vertices

Here we show that by "greedy elimination" [243] [147] [146], i.e. removing all degree 1 and degree 2 vertices in $O(m)$ time we can reduce oblivious routing in almost- $j$-trees to oblivious routing in graphs with $O(j)$ vertices while only losing $O(1)$ in the competitive ratio. Again, we remark that the lemmas in this section are derived heavily from [182] but repeated for completeness and to prove additional properties that we will need for our purposes.

We start by showing that an almost- $j$-tree with no degree 1 or degree 2 vertices has at most $O(j)$ vertices.

Lemma 12.6.10. For any almost $j$-tree $G=(V, E)$ with no degree 1 or degree 2 vertices, we have $|V| \leq 3 j-2$.

Proof. Since $G$ is an almost $j$-tree, there is some $J \subseteq V$ with $|J| \leq j$ such that the removal of all edges with both endpoints in $J$ creates a forest. Now, since $K=V-J$ is incident only to forest edges clearly the sum of the degrees of the vertices in $K$ is at most $2(|V|-1)$ (otherwise there would be a cycle). However, since the minimum degree in $G$ is 3 , clearly this sum is at least $3(|V|-j)$. Combining yields that $3|V|-3 j \leq 2|V|-2$.

Next, we show how to remove degree one vertices efficiently.
Lemma 12.6.11 (Removing Degree One Vertices). Let $G=(V, E, \boldsymbol{\mu})$ be an unweighted capacitated graph, let $a \in V$ be a degree 1 vertex, let $e=(a, b) \in E$ be the single edge incident to a and let $G^{\prime}=\left(V^{\prime}, E^{\prime}, \boldsymbol{\mu}^{\prime}\right)$ be the graph that results from simply removing e and a, i.e. $V^{\prime}=V \backslash\{a\}$ and $E^{\prime}=E \backslash\{e\}$. Given $a \in V$ and an oblivious routing algorithm $\boldsymbol{A}^{\prime}$ in $G^{\prime}$ in $O(1)$ time we can construct an oblivious routing algorithm $\boldsymbol{A}$ in $G$ such that

$$
\mathcal{T}(\boldsymbol{A})=O\left(\mathcal{T}\left(\boldsymbol{A}^{\prime}\right)+1\right) \quad, \text { and } \quad \rho(\boldsymbol{A})=\rho\left(\boldsymbol{A}^{\prime}\right)
$$

Proof. For any demand vector $\chi$, the only way to route demand at $a$ in $G$ is over $e$. Therefore, if $\boldsymbol{B} \boldsymbol{f}=\boldsymbol{\chi}$ then $\boldsymbol{f}(e)=\boldsymbol{\chi}$. Therefore, to get an oblivious routing algorithm on $G$, we can simply send demand at $a$ over edge $e$, modify the demand at $b$ accordingly, and then run the oblivious routing algorithm on $G^{\prime}$ on the remaining vertices. The routing algorithm we get is the following

$$
\boldsymbol{A} \stackrel{\text { def }}{=} \boldsymbol{I}_{E^{\prime} \rightarrow E} \boldsymbol{A}^{\prime}\left(\boldsymbol{I}+\mathbf{1}_{b} \mathbf{1}_{a}^{T}\right)+\mathbf{1}_{e} \mathbf{1}_{a}^{T}
$$

Since all routing algorithms send this flow on $e$ we get that $\rho(\boldsymbol{A})=\rho\left(\boldsymbol{A}^{\prime}\right)$ and since the above operators not counting $\boldsymbol{A}$ have only $O(1)$ entries that are not the identity we can clearly implement the operations in the desired running time.

Using the above lemma we show how to remove all degree 1 and 2 vertices in $O(m)$ time while only increasing the congestion by $O(1)$.

Lemma 12.6.12 (Greedy Elimination). Let $G=(V, E, \boldsymbol{\mu})$ be an unweighted capacitated graph and let $G^{\prime}=\left(V^{\prime}, E^{\prime}, \boldsymbol{\mu}^{\prime}\right)$ be the graph the results from iteratively removing vertices of degree 1 and replacing degree 2 vertices with an edge connecting its neighbors of the minimum capacity of its adjacent edges. We can construct $G^{\prime}$ in $O(m)$ time and given an oblivious routing algorithm $\boldsymbol{A}^{\prime}$ in $G^{\prime}$ in $O(1)$ time we can construct an oblivious routing algorithm $\boldsymbol{A}$ in $G$ such that ${ }^{9}$

$$
\mathcal{T}(\boldsymbol{A})=O\left(\mathcal{T}\left(\boldsymbol{A}^{\prime}\right)+|E|\right) \quad, \text { and } \quad \rho(\boldsymbol{A}) \leq 4 \cdot \rho\left(\boldsymbol{A}^{\prime}\right)
$$

Proof. First we repeatedly apply Lemma 12.6 .11 repeatedly to in reduce to the case that there are no degree 1 vertices. By simply array of the degrees of every vertex and a list of degree 1 vertices this can be done in $O(m)$ time. We denote the result of these operations by graph $K$.

Next, we repeatedly find degree two vertices that have not been explored and explore this vertices neighbors to get a path of vertices, $a_{1}, a_{2}, \ldots, a_{k} \in V$ for $k>3$ such that each vertex $a_{2}, \ldots, a_{k-1}$ is of degree two. We then compute $j=\arg \min _{i \in[k-1]} \boldsymbol{\mu}\left(a_{i}, a_{i+1}\right)$, remove edge $\left(a_{j}, a_{j+1}\right)$ and add an edge $\left(a_{1}, a_{k}\right)$ of capacity $\boldsymbol{\mu}\left(a_{j}, a_{j+1}\right)$. We denote the result of doing this for all degree two vertices by $K^{\prime}$ and note that again by careful implementation this can be performed in $O(\mathrm{~m})$ time.

Note that clearly $K$ is embeddable in $K^{\prime}$ with congestion 2 just by routing every edge over itself except the removed edges which we route by the path plus the added edges. Furthermore, $K^{\prime}$ is embeddable in $K$ with congestion 2 again by routing every edge on itself except for the edges which we added which we route back over the paths they came from. Furthermore, we note that clearly this embedding and the transpose of this operator is computable in $O(m)$ time.

Finally, by again repeatedly applying Lemma 12.6 .11 to $K^{\prime}$ until there are no degree 1 vertices we get a graph $G^{\prime}$ that has no degree one or degree two vertices (since nothing decreased the degree of vertices with degree more than two). Furthermore, by Lemma 12.6.11 and by Lemma 12.4.8 we see that we can compose these operators to compute $\boldsymbol{A}$ with the desired properties.

[^30]
## ■ 12.6.4 Putting It All Together

Here we put together the previous components to prove the main theorem of this section.
Node Reduction Theorem 12.4.10. Using Lemma 12.6.6, we can construct $G^{\prime}=\sum_{i=1}^{t} \lambda_{i} G_{i}$ and embeddings $\boldsymbol{M}_{1}, \ldots, \boldsymbol{M}_{t}$ from $G_{i}$ to $G$. Next we can apply Lemma 12.6.9 to each $G_{i}$ to get almost- $j$-trees $G_{1}^{\prime}, \ldots, G_{t}^{\prime}$ and embeddings $\boldsymbol{M}_{1}^{\prime}, \ldots, \boldsymbol{M}_{t}^{\prime}$ from $G_{i}^{\prime}$ to $G_{i}$. Furthermore, using Lemma 12.6 .12 we can construction graphs $G_{1}^{\prime \prime}, \ldots, G_{t}^{\prime \prime}$ with the desired properties (the congestion ratio property follows from the fact that we only add capacities during these reductions)

Now given oblivious routing algorithms $\boldsymbol{A}_{1}^{\prime \prime}, \ldots, \boldsymbol{A}_{t}^{\prime \prime}$ on the $G_{i}^{\prime \prime}$ and again by Lemma 12.6 .12 we could get oblivious routing algorithms $\boldsymbol{A}_{1}^{\prime}, \ldots, \boldsymbol{A}_{t}^{\prime}$ on the $G_{i}^{\prime}$ with constant times more congestion. Finally, by the guarantees of Lemma 12.4 .8 we have that $\boldsymbol{A} \stackrel{\text { def }}{=} \sum_{i=1}^{t} \lambda \boldsymbol{M}_{i} \boldsymbol{M}_{i}^{\prime} \boldsymbol{A}_{i}^{\prime}$ is an oblivious routing algorithm that satisfies the requirements.

## ■ 12.7 Nonlinear Projection and Maximum Concurrent Flow

## ■ 12.7.1 Gradient Descent Method for Nonlinear Projection Problem

In this section, we strengthen and generalize the MaxFlow algorithm to a more general setting. We believe this algorithm may be of independent interest as it includes maximum concurrent flow problem, the compressive sensing problem, etc. For some norms, e.g. $\|\cdot\|_{1}$ as typically of interest compressive sensing, the Nesterov algorithm [208] can be used to replace gradient descent. However, this kind of accelerated method is not known in the general norm settings as good proxy function may not exist at all. Even worse, in the non-smooth regime, the minimization problem on the $\|\cdot\|_{p}$ space with $p>2$ is difficult under some oracle assumption [201]. For these reasons we focus here on the gradient descent method which is always applicable.

Given a norm $\|\cdot\|$, we wish to solve the what we call the non-linear projection problem

$$
\min _{\boldsymbol{x} \in L}\|\boldsymbol{x}-\boldsymbol{y}\|
$$

where $\boldsymbol{y}$ is an given point and $L$ is a linear subspace. We assume the following:
Assumption 12.7.1.

1. There are a family of convex differentiable functions $f_{t}$ such that for all $\boldsymbol{x} \in L$, we have

$$
\|\boldsymbol{x}\| \leq f_{t}(\boldsymbol{x}) \leq\|\boldsymbol{x}\|+K t
$$

and the Lipschitz constant of $\nabla f_{t}$ is $\frac{1}{t}$.
2. There is a projection matrix $\boldsymbol{P}$ onto the subspace $L$.

In other words we assume that there is a family of regularized objective functions $f_{t}$ and a projection matrix $\boldsymbol{P}$, which we can think of as an approximation algorithm of this projection problem.

Now, let $\boldsymbol{x}^{*}$ be a minimizer of $\min _{\boldsymbol{x} \in L}\|\boldsymbol{x}-\boldsymbol{y}\|$. Since $\boldsymbol{x}^{*} \in L$, we have $\boldsymbol{P} \boldsymbol{x}^{*}=\boldsymbol{x}^{*}$ and hence

$$
\begin{align*}
\|\boldsymbol{P} \boldsymbol{y}-\boldsymbol{y}\| & \leq\left\|\boldsymbol{y}-\boldsymbol{x}^{*}\right\|+\left\|\boldsymbol{x}^{*}-\boldsymbol{P} \boldsymbol{y}\right\| \\
& \leq\left\|\boldsymbol{y}-\boldsymbol{x}^{*}\right\|+\left\|\boldsymbol{P} \boldsymbol{x}^{*}-\boldsymbol{P} \boldsymbol{y}\right\| \\
& \leq(1+\|\boldsymbol{P}\|) \min _{\boldsymbol{x} \in L}\|\boldsymbol{x}-\boldsymbol{y}\| . \tag{12.4}
\end{align*}
$$

Therefore, the approximation ratio of $\boldsymbol{P}$ is $1+\|\boldsymbol{P}\|$ and we see that our problem is to show that we can solve nonlinear projection using a decent linear projection matrix. Our algorithm for solving this problem is shown in NonlinearProjection.

```
Algorithm 34: NonlinearProjection
    Input: a point \(\boldsymbol{y}\) and OPT \(=\min _{\boldsymbol{x} \in L}\|\boldsymbol{x}-\boldsymbol{y}\|\).
    Let \(\boldsymbol{y}_{0}=(\boldsymbol{I}-\boldsymbol{P}) \vec{y}\) and \(\boldsymbol{x}_{0}=0\).
    for \(j=0, \cdots\), until \(2^{-j}\|\boldsymbol{P}\| \leq \frac{1}{2}\) do
        If \(2^{-j}\|\boldsymbol{P}\|>1\), then let \(t_{j}=\frac{2^{-(j+2)}\|\boldsymbol{P}\| \mathrm{OPT}}{K}\) and \(k_{j}=3200\|\boldsymbol{P}\|^{2} K\).
        If \(2^{-j}\|\boldsymbol{P}\| \leq 1\), then let \(t_{j}=\frac{\varepsilon \mathrm{OPT}}{2 K}\) and \(k_{j}=\frac{800\|\boldsymbol{P}\|^{2} K}{\varepsilon^{2}}\).
        Let \(g_{j}(\boldsymbol{x})=f_{t_{j}}\left(\boldsymbol{P} \boldsymbol{x}-\boldsymbol{y}_{j}\right)\) and \(\boldsymbol{x}_{0}=0\).
        for \(i=0, \cdots, k_{j}-1\) do
            \(\boldsymbol{x}_{i+1}=\boldsymbol{x}_{i}-\frac{t}{\|\boldsymbol{P}\|^{2}}\left(\nabla g_{j}\left(\boldsymbol{x}_{i}\right)\right)^{\#}\).
        end
        Let \(\boldsymbol{y}_{j+1}=\boldsymbol{y}_{j}-\boldsymbol{P} \boldsymbol{x}_{k_{j}}\).
    end
    Output: \(\boldsymbol{y}-\boldsymbol{y}_{\text {last }}\).
```

Note that this algorithm and its proof are quite similar to Theorem 12.3.4 but modified to scale parameters over an outer loop. By changing the parameter $t$ we can decrease the dependence of the initial error. ${ }^{10}$

Theorem 12.7.2. Assume the conditions in Assumption 12.7.1 are satisfied. Let $\mathcal{T}$ be the time needed to compute $\boldsymbol{P} x$ and $\boldsymbol{P}^{T} x$ and $x^{\#}$. Then, NonlinearProjection outputs a vector $\boldsymbol{x}$ with $\|\boldsymbol{x}\| \leq(1+\varepsilon) \min _{\boldsymbol{x} \in L}\|\boldsymbol{x}-\boldsymbol{y}\|$ and the algorithm takes time

$$
O\left(\|\boldsymbol{P}\|^{2} K(\mathcal{T}+m)\left(\frac{1}{\varepsilon^{2}}+\log \|\boldsymbol{P}\|\right)\right) .
$$

Proof. We prove by induction on $j$ that when $2^{-(j-1)}\|\boldsymbol{P}\| \geq 1$ we have $\left\|\boldsymbol{y}_{j}\right\| \leq\left(1+2^{-j}\|\boldsymbol{P}\|\right)$ OPT.
For the base case $(j=0),(12.7 .2)$ shows that $\left\|\boldsymbol{y}_{0}\right\| \leq(1+\|\boldsymbol{P}\|)$ OPT.
For the inductive case we assume that the assertion holds for some $j$. We start by bounding the corresponding $R$ in Theorem 12.3.1 for $g_{j}$, which we denote $R_{j}$. Note that

$$
g_{j}\left(\boldsymbol{x}_{0}\right)=f_{t_{j}}\left(-\boldsymbol{y}_{j}\right) \leq\left\|\boldsymbol{y}_{j}\right\|+K t_{j} \leq\left(1+2^{-j}\|\boldsymbol{P}\|\right) \mathrm{OPT}+K t_{j} .
$$

Hence, the condition that $g_{j}(\boldsymbol{x}) \leq g_{j}\left(\boldsymbol{x}_{0}\right)$ implies that

$$
\left\|\boldsymbol{P} \boldsymbol{x}-\boldsymbol{y}_{j}\right\| \leq\left(1+2^{-j}\|\boldsymbol{P}\|\right) \mathrm{OPT}+K t_{j} .
$$

Take any $\boldsymbol{y} \in X^{*}$, let $\boldsymbol{c}=\boldsymbol{x}-\boldsymbol{P} \boldsymbol{x}+\boldsymbol{y}$, and note that $\boldsymbol{P} \boldsymbol{c}=\boldsymbol{P} \boldsymbol{y}$ and therefore $\boldsymbol{c} \in X^{*}$. Using these facts, we can bound $R_{j}$ as follows

$$
\begin{aligned}
R_{j} & =\max _{\boldsymbol{x} \in \mathbb{R}^{E}: g_{j}(\boldsymbol{x}) \leq g_{j}\left(\boldsymbol{x}_{0}\right)}\|\boldsymbol{x}-\boldsymbol{c}\| \\
& \leq \max _{\boldsymbol{x} \in \mathbb{R}^{E}: g_{j}(\boldsymbol{x}) \leq g_{j}\left(\boldsymbol{x}_{0}\right)}\|\boldsymbol{P} \boldsymbol{x}\|+\|\boldsymbol{P} \boldsymbol{y}\| \\
& \leq 2\left\|\boldsymbol{y}_{0}\right\|+\left\|\boldsymbol{P} \boldsymbol{x}-\boldsymbol{y}_{j}\right\|+\left\|\boldsymbol{P} \boldsymbol{y}-\boldsymbol{y}_{j}\right\| \\
& \leq 4\left(1+2^{-j}\|\boldsymbol{P}\|\right) \mathrm{OPT}+2 K t_{j} .
\end{aligned}
$$

[^31]Similar to Lemma 12.3.3, the Lipschitz constant $L_{j}$ of $g_{j}$ is $\|\boldsymbol{P}\|^{2} / t_{j}$. Hence, Theorem 12.3.1 shows that

$$
\begin{aligned}
g_{j}\left(\boldsymbol{x}_{k_{j}}\right) & \leq \min _{\boldsymbol{x}} g_{j}(\boldsymbol{x})+\frac{2 \cdot L_{j} \cdot R_{j}^{2}}{k_{j}+4} \\
& \leq \min _{\boldsymbol{x}}\left\|\boldsymbol{P} \vec{x}-\overrightarrow{y_{j}}\right\|+\frac{2 \cdot L_{j} \cdot R_{j}^{2}}{k_{j}+4}+K t_{j}
\end{aligned}
$$

So, we have

$$
\begin{aligned}
\left\|\boldsymbol{P} \boldsymbol{x}_{k_{j}}-\boldsymbol{y}_{j}\right\| & \leq f_{t_{j}}\left(\boldsymbol{P} \boldsymbol{x}_{k_{j}}-\boldsymbol{y}_{j}\right) \\
& \leq \mathrm{OPT}+K t_{j}+\frac{2\|\boldsymbol{P}\|^{2}}{t_{j}\left(k_{j}+4\right)}\left(4\left(1+2^{-j}\|\boldsymbol{P}\|\right) \mathrm{OPT}+2 K t_{j}\right)^{2} .
\end{aligned}
$$

When $2^{-j}\|\boldsymbol{P}\|>1$, we have $t_{j}=\frac{2^{-(j+2)}\|\boldsymbol{P}\| \text { OPT }}{K}$ and $k_{j}=3200\|\boldsymbol{P}\|^{2} K$ and hence

$$
\left\|\boldsymbol{y}_{j+1}\right\|=\left\|\boldsymbol{P} \boldsymbol{x}_{k_{j}}-\boldsymbol{y}_{j}\right\| \leq\left(1+2^{-j-1}\|\boldsymbol{P}\|\right) \mathrm{OPT} .
$$

When $2^{-j}\|\boldsymbol{P}\| \leq 1$, we have $t_{j}=\frac{\varepsilon \mathrm{OPT}}{2 K}$ and $k_{j}=\frac{800\|\boldsymbol{P}\|^{2} K}{\varepsilon^{2}}$ and hence

$$
\left\|\boldsymbol{y}_{j+1}\right\|=\left\|\boldsymbol{P} \boldsymbol{x}_{k_{j}}-\boldsymbol{y}_{j}\right\| \leq(1+\varepsilon) \mathrm{OPT} .
$$

Since $\boldsymbol{y}_{\text {last }}$ is $\boldsymbol{y}$ plus some vectors in $L, \boldsymbol{y}-\boldsymbol{y}_{\text {last }} \in L$ and $\left\|\boldsymbol{t}-\boldsymbol{y}_{\text {last }}-\boldsymbol{y}\right\|=\left\|\boldsymbol{y}_{\text {last }}\right\| \leq(1+\varepsilon)$ OPT.

## ■ 12.7.2 Maximum Concurrent Flow

For an arbitrary set of demands $\chi_{i} \in \mathbb{R}^{V}$ with $\sum_{v \in V} \chi_{i}(v)=0$ for $i=1, \cdots, k$, we wish to solve the following maximum concurrent flow problem

$$
\max _{\alpha \in R, f \in \mathbb{R}^{E}} \alpha \quad \text { subject to } \quad \boldsymbol{B}^{T} \boldsymbol{f}_{i}=\alpha \boldsymbol{\chi}_{i} \quad \text { and } \quad\left\|\boldsymbol{U}^{-1} \sum_{i=1}^{k}\left|\boldsymbol{f}_{i}\right|\right\|_{\infty} \leq 1 .
$$

Similar to Section 12.3.2, it is equivalent to the problem

$$
\min _{c \in \mathbb{R}^{E X[k]}}\left\|\sum_{i=1}^{k}\left|\boldsymbol{\alpha}_{i}+(\boldsymbol{Q} \boldsymbol{x})_{i}\right|\right\|_{\infty}
$$

where $\boldsymbol{Q}$ is a projection matrix onto the subspace $\left\{\boldsymbol{B}^{T} \boldsymbol{U} \boldsymbol{x}_{i}=0\right\}$, the output maximum concurrent flow is

$$
\boldsymbol{f}_{i}(\boldsymbol{x})=\boldsymbol{U}\left(\boldsymbol{\alpha}_{i}+(\boldsymbol{Q} \boldsymbol{x})_{i}\right) /\left\|\sum_{i=1}^{k}\left|\boldsymbol{\alpha}_{i}+(\boldsymbol{Q} \boldsymbol{x})_{i}\right|\right\|_{\infty}
$$

and $\boldsymbol{U} \boldsymbol{\alpha}_{i}$ is any flow such that $\boldsymbol{B}^{T} \boldsymbol{U} \overrightarrow{\alpha_{i}}=\boldsymbol{\chi}_{i}$. In order to apply NonlinearProjection, we need to find a regularized norm and a good projection matrix. Let us define the norm

$$
\|\boldsymbol{x}\|_{1 ; \infty}=\max _{e \in E} \sum_{i=1}^{k}\left|x_{i}(e)\right|
$$

The problem is simply $\|\boldsymbol{\alpha}+\boldsymbol{Q} \boldsymbol{x}\|_{1 ; \infty}$ where $\boldsymbol{Q}$ is a projection matrix from $\mathbb{R}^{E \times[k]}$ to $\mathbb{R}^{E \times[k]}$ onto some subspace. Since each copy $\mathbb{R}^{E}$ is same, there is no reason that there is coupling in $\boldsymbol{Q}$ between different copies of $\mathbb{R}^{E}$. In the next lemma, we formalize this by the fact that any good projection
matrix $\boldsymbol{P}$ onto the subspace $\left\{\boldsymbol{B}^{T} \boldsymbol{U} \vec{x}=0\right\} \subset \mathbb{R}^{E}$ extends to a good projection $\boldsymbol{Q}$ onto the subspace $\left\{\boldsymbol{B}^{T} \boldsymbol{U} \boldsymbol{x}_{i}=0\right\} \subset \mathbb{R}^{E \times[k]}$. Therefore, we can simply extends the good circulation projection $\boldsymbol{P}$ by formula $(\boldsymbol{Q} \boldsymbol{x})_{i}=\boldsymbol{P} \boldsymbol{x}_{i}$. Thus, the only last piece needed is a regularized $\|\cdot\|_{1 ; \infty}$. However, it turns out that smoothing via conjugate does not work well in this case because the dual space of $\|\cdot\|_{1 ; \infty}$ involves with $\|\cdot\|_{\infty}$, which is unfavorable for this kind of smoothing procedure. It can be shown that there is no such good regularized $\|\cdot\|_{1 ; \infty}$. Therefore, we could not do $O\left(m^{1+o(1)} k / \varepsilon^{2}\right)$ using this approach, however, $O\left(m^{1+o(1)} k^{2} / \varepsilon^{2}\right)$ is possible by using a bad regularized $\|\cdot\|_{1 ; \infty}$. We believe the dependence of $k$ can be improved to $3 / 2$ using this approach by suitable using Nesterov algorithm because the $\|\cdot\|_{1}$ space caused by the multicommodity is a favorable geometry for accelerated methods.

Lemma 12.7.3. Let $\operatorname{smaxL} 1_{t}(\boldsymbol{x})=\operatorname{smax}_{t}\left(\sum_{i=1}^{k} \sqrt{\left(x_{i}(e)\right)^{2}+t^{2}}\right)$. It is a convex continuously differentiable function. The Lipschitz constant of $\nabla s m a x L 1_{t}$ is $\frac{2}{t}$ and

$$
\|\boldsymbol{x}\|_{1 ; \infty}-t \ln (2 m) \leq \operatorname{smax}^{2} 1_{t}(\boldsymbol{x}) \leq\|\boldsymbol{x}\|_{1 ; \infty}+k t .
$$

Proof. 1) It is clear that $\operatorname{smax} L 1_{t}$ is smooth.
2) $\operatorname{smax} L 1_{t}$ is convex.

Since $\operatorname{smax}_{t}$ is increasing for positive values and $\sqrt{x^{2}+t^{2}}$ is convex, for any $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{E \times[k]}$ and $0 \leq t \leq 1$, we have

$$
\begin{aligned}
\operatorname{smax} 1_{t}(t \boldsymbol{x}+(1-t) \boldsymbol{y}) & =\operatorname{smax}_{t}\left(\sum_{i=1}^{k} \sqrt{\left(\left(t x_{i}+(1-t) y_{i}\right)(e)\right)^{2}+t^{2}}\right) \\
& \leq \operatorname{smax}_{t}\left(\sum_{i=1}^{k}\left(t \sqrt{\left(x_{i}(e)\right)^{2}+t^{2}}+(1-t) \sqrt{\left(y_{i}(e)\right)^{2}+t^{2}}\right)\right) \\
& \leq t \operatorname{smax}_{t} L 1_{t}(\boldsymbol{x})+(1-t) \operatorname{smax}^{2} L 1_{t}(\boldsymbol{y}) .
\end{aligned}
$$

3) The Lipschitz constant of $\nabla \operatorname{smax} L 1_{t}$ is $\frac{2}{t}$.

Note that $\operatorname{smax}_{t}$ (not its gradient) has Lipschitz constant 1 because for any $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{E}$,

$$
\begin{aligned}
& \left|\operatorname{smax}_{t}(\boldsymbol{x})-\operatorname{smax}_{t}(\boldsymbol{y})\right| \\
= & \left|t \ln \left(\frac{\sum_{e \in E}\left(\exp \left(-\frac{x(e)}{t}\right)+\exp \left(\frac{x(e)}{t}\right)\right)}{2 m}\right)-t \ln \left(\frac{\sum_{e \in E}\left(\exp \left(-\frac{y(e)}{t}\right)+\exp \left(\frac{y(e)}{t}\right)\right)}{2 m}\right)\right| \\
= & t\left|\ln \left(\frac{\sum_{e \in E}\left(\exp \left(-\frac{x(e)}{t}\right)+\exp \left(\frac{x(e)}{t}\right)\right)}{\sum_{e \in E}\left(\exp \left(-\frac{y(e)}{t}\right)+\exp \left(\frac{y(e)}{t}\right)\right)}\right)\right| \\
\leq & t\left|\ln \left(\max _{e \in E} \exp \left(\frac{|x-y|(e)}{t}\right)\right)\right| \\
= & \|\boldsymbol{x}-\boldsymbol{y}\|_{\infty} .
\end{aligned}
$$

Also, by the definition of derivative, for any $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{n}$ and $t \in \mathbb{R}$, we have

$$
\operatorname{smax}_{t}(\boldsymbol{x}+t \boldsymbol{y})-\operatorname{smax}_{t}(\boldsymbol{x})=t\left\langle\nabla \operatorname{smax}_{t}(\boldsymbol{x}), \boldsymbol{y}\right\rangle+o(t) .
$$

and it implies $\left|\left\langle\nabla \operatorname{smax}_{t}(\boldsymbol{x}), \boldsymbol{y}\right\rangle\right| \leq\|\boldsymbol{y}\|_{\infty}$ for arbitrary $\boldsymbol{y}$ and hence

$$
\begin{equation*}
\left\|\nabla \operatorname{smax}_{t}(\boldsymbol{x})\right\|_{1} \leq 1 \tag{12.5}
\end{equation*}
$$

For notational simplicity, let $s_{1}=\operatorname{smax} L 1_{t}, s_{2}=\operatorname{smax}_{t}$ and $s_{3}(x)=\sqrt{x^{2}+t^{2}}$. Thus, we have

$$
s_{1}(\vec{x})=s_{2}\left(\sum_{i=1}^{k} s_{3}\left(x_{i}(e)\right)\right) .
$$

Now, we want to prove

$$
\left\|\nabla s_{1}(\boldsymbol{x})-\nabla s_{1}(\boldsymbol{y})\right\|_{\infty ; 1} \leq \frac{2}{t}\|\boldsymbol{x}-\boldsymbol{y}\|_{1 ; \infty}
$$

Note that

$$
\frac{\partial s_{1}(\boldsymbol{x})}{\partial x_{i}(e)}=\frac{\partial s_{2}}{\partial e}\left(\sum_{j} s_{3}\left(x_{j}(e)\right)\right) \frac{d s_{3}}{d x}\left(x_{i}(e)\right) .
$$

Hence, we have

$$
\begin{aligned}
& \left\|\nabla s_{1}(\boldsymbol{x})-\nabla s_{1}(\boldsymbol{y})\right\|_{\infty ; 1} \\
= & \sum_{e} \max _{i}\left|\frac{\partial s_{2}}{\partial e}\left(\sum_{j} s_{3}\left(x_{j}(e)\right)\right) \frac{d s_{3}}{d x}\left(x_{i}(e)\right)-\frac{\partial s_{2}}{\partial e}\left(\sum_{j} s_{3}\left(y_{j}(e)\right)\right) \frac{d s_{3}}{d x}\left(y_{i}(e)\right)\right| \\
\leq & \sum_{e}\left|\frac{\partial s_{2}}{\partial e}\left(\sum_{j} s_{3}\left(x_{j}(e)\right)\right)\right| \max _{i}\left|\frac{d s_{3}}{d x}\left(x_{i}(e)\right)-\frac{d s_{3}}{d x}\left(y_{i}(e)\right)\right| \\
& +\sum_{e} \max _{i} \frac{d s_{3}}{d x}\left(y_{i}(e)\right)\left|\frac{\partial s_{2}}{\partial e}\left(\sum_{j} s_{3}\left(x_{j}(e)\right)\right)-\frac{\partial s_{2}}{\partial e}\left(\sum_{j} s_{3}\left(y_{j}(e)\right)\right)\right| .
\end{aligned}
$$

Since $s_{3}$ has $\frac{1}{t}$-Lipschitz gradient, we have

$$
\left|\frac{d s_{3}}{d x}(\boldsymbol{x})-\frac{d s_{3}}{d x}(\boldsymbol{y})\right| \leq \frac{1}{t}|\boldsymbol{x}-\boldsymbol{y}| .
$$

By (12.5), we have

$$
\sum_{e}\left|\frac{\partial s_{2}}{\partial e}(\boldsymbol{x}(e))\right| \leq 1
$$

Hence, we have

$$
\begin{aligned}
& \sum_{e} \max _{i}\left|\frac{d s_{3}}{d x}\left(x_{i}(e)\right)-\frac{d s_{3}}{d x}\left(y_{i}(e)\right)\right|\left|\frac{\partial s_{2}}{\partial e}\left(\sum_{i} s_{3}\left(x_{i}(e)\right)\right)\right| \\
\leq & \frac{1}{t} \max _{i, e}\left|x_{i}(e)-y_{i}(e)\right| \sum_{e}\left|\frac{\partial s_{3}}{\partial e}\left(\sum_{i} s_{3}\left(x_{i}(e)\right)\right)\right| \\
= & \frac{1}{t}\|\vec{x}-\vec{y}\|_{1 ; \infty} .
\end{aligned}
$$

Since $s_{3}$ is 1 -Lipschitz, we have

$$
\left|\frac{d s_{3}}{d x}\right| \leq 1
$$

Since $s_{2}$ has $\frac{1}{t}$-Lipschitz gradient in $\|\cdot\|_{\infty}$, we have

$$
\sum_{e}\left|\frac{\partial s_{2}}{\partial e}(x)-\frac{\partial s_{2}}{\partial e}(y)\right| \leq \frac{1}{t}\|\boldsymbol{x}-\boldsymbol{y}\|_{\infty} .
$$

Hence, we have

$$
\begin{aligned}
& \sum_{e} \max _{i} \frac{d s_{3}}{d x}\left(y_{i}(e)\right)\left|\frac{\partial s_{2}}{\partial e}\left(\sum_{j} s_{3}\left(x_{j}(e)\right)\right)-\frac{\partial s_{2}}{\partial e}\left(\sum_{j} s_{3}\left(y_{j}(e)\right)\right)\right| \\
\leq & \sum_{e}\left|\frac{\partial s_{2}}{\partial e}\left(\sum_{j} s_{3}\left(x_{j}(e)\right)\right)-\frac{\partial s_{2}}{\partial e}\left(\sum_{j} s_{3}\left(y_{j}(e)\right)\right)\right| \\
\leq & \frac{1}{t}\left\|\sum_{i} s_{3}\left(x_{i}(e)\right)-\sum s_{3}\left(y_{i}(e)\right)\right\|_{\infty} \\
\leq & \frac{1}{t}\left\|\sum_{i}\left|x_{i}(e)-y_{i}(e)\right|\right\| \\
= & \frac{1}{t}\|\boldsymbol{x}-\boldsymbol{y}\|_{1 ; \infty}
\end{aligned}
$$

Therefore, we have

$$
\left\|\nabla s_{1}(\boldsymbol{x})-\nabla s_{1}(\boldsymbol{y})\right\|_{\infty ; 1} \leq \frac{2}{t}\|\boldsymbol{x}-\boldsymbol{y}\|_{1 ; \infty} .
$$

4) Using the fact that

$$
\|x(e)\| \leq \sum_{i=1}^{k} \sqrt{\left(x_{i}(e)\right)^{2}+t^{2}} \leq\|x(e)\|_{1}+k t
$$

and smax is 1-Lipschitz, we have

$$
\|\boldsymbol{x}\|_{1 ; \infty}-t \ln (2 m) \leq \operatorname{smax}_{2} L 1_{t}(\boldsymbol{x}) \leq\|\boldsymbol{x}\|_{1 ; \infty}+k t .
$$

The last thing needed is to check is that the \# operator is easy to compute.
Lemma 12.7.4. In $\|\cdot\|_{1 ; \infty}$, the $\#$ operator is given by an explicit formula

$$
\left(\boldsymbol{x}^{\#}\right)_{i}(e)=\left\{\begin{array}{ll}
\|\boldsymbol{x}\|_{\infty ; 1} \operatorname{sign}\left(x_{i}(e)\right) & \text { if } i \text { is the smallest index such that } \max _{j} x_{j}(e)=x_{i}(e) \\
0 & \text { otherwises }
\end{array} .\right.
$$

Proof. It can be proved by direct computation.
Now, all the conditions in the Assumption 12.7.1 are satisfied. Therefore, Theorem 12.7.2 and Theorem 12.4.12 gives us the following theorem:

Theorem 12.7.5. Given an undirected capacitated graph $G=(V, E, \mu)$ with capacity ratio $U$. Assume $U=\operatorname{poly}(|V|)$. There is an algorithm finds an $(1-\varepsilon)$ approximate Maximum Concurrent Flow in time

$$
O\left(\frac{k^{2}}{\varepsilon^{2}}|E| 2^{O(\sqrt{\log |V| \log \log |V|})}\right)
$$

Proof. Let $\boldsymbol{A}$ be the oblivious routing algorithm given by Theorem 12.4.12. And we have $\rho(\boldsymbol{A}) \leq$ $2^{O(\sqrt{\log |V| \log \log |V|})}$. Let us define the scaled circulation projection matrix $\boldsymbol{P}=\boldsymbol{I}-\boldsymbol{U} \boldsymbol{A} \boldsymbol{B}^{T} \boldsymbol{U}^{-1}$. Lemma 12.4.5 shows that $\|\boldsymbol{P}\|_{\infty} \leq 1+2^{O(\sqrt{\log |V| \log \log |V|})}$.

Let the multi-commodity circulation projection matrix $\boldsymbol{Q}: \mathbb{R}^{E \times[k]} \rightarrow \mathbb{R}^{E \times[k]}$ defined by $(\boldsymbol{Q x})_{i}=$ $\boldsymbol{P} \boldsymbol{x}_{i}$. Note that the definition of $\|\boldsymbol{Q}\|_{1 ; \infty}$ is similar to $\rho(\boldsymbol{Q})$. By similar proof as Lemma 12.4.3, we have $\|\boldsymbol{Q}\|_{1 ; \infty}=\|\boldsymbol{P}\|_{\infty}$. Hence, we have $\|\boldsymbol{Q}\|_{1 ; \infty} \leq 1+2^{O(\sqrt{\log |V| \log \log |V|})}$. Also, since $\boldsymbol{P}$ is a projection matrix on the subspace $\left\{\boldsymbol{x} \in \mathbb{R}^{E}: \boldsymbol{B}^{T} \boldsymbol{U} \boldsymbol{x}=0\right\}, \boldsymbol{Q}$ is a projection matrix on the subspace $\left\{\boldsymbol{x} \in \mathbb{R}^{E \times[k]}: \boldsymbol{B}^{T} \boldsymbol{U} \boldsymbol{x}_{i}=0\right\}$.

By Lemma 12.7.3, the function $\operatorname{smax} L 1_{t}(\boldsymbol{x})$ is a convex continuously differentiable function such that the Lipschitz constant of $\nabla \operatorname{smax} L 1_{t}$ is $\frac{2}{t}$ and

$$
\|\boldsymbol{x}\|_{1 ; \infty}-t \ln (2 m) \leq \operatorname{smax}^{2} L 1_{t}(\boldsymbol{x}) \leq\|\boldsymbol{x}\|_{1 ; \infty}+k t .
$$

Given an arbitrary set of demands $\chi_{i} \in \mathbb{R}^{V}$, we find a vector $\vec{y}$ such that

$$
\boldsymbol{B}^{T} \boldsymbol{U} \boldsymbol{y}=-\boldsymbol{\chi}_{i}
$$

Then, we use the NonlinearProjection to solve

$$
\min _{\boldsymbol{B}^{T} \boldsymbol{U} \boldsymbol{x}=0}\|\boldsymbol{x}-\boldsymbol{y}\|_{1 ; \infty}
$$

using a family of functions $\operatorname{smax} L 1_{t}(\boldsymbol{x})+t \ln (2 n)$ and the projection matrix $\boldsymbol{Q}$. Since each iteration involves calculation of gradients and \# operator, it takes $O(m k)$ each iteration. And it takes $\tilde{O}\left(\|\boldsymbol{Q}\|_{1 ; \infty}^{2} K / \varepsilon^{2}\right)$ iterations in total where $K=k+\ln (2 m)$. In total, it NonlinearProjection outputs a $(1+\varepsilon)$ approximate minimizer $\boldsymbol{x}$ in time

$$
O\left(\frac{k^{2}}{\varepsilon^{2}}|E| 2^{O(\sqrt{\log |V| \log \log |V|})}\right)
$$

And it gives a $(1-\varepsilon)$ approximate maximum concurrent flow $\vec{f}_{i}$ by a direct formula.

## ■ 12.8 Appendix

In this section, we present some basic fact used in this chapter about norm and dual norm. Also, we presented some lemmas about convex functions with Lipschitz gradient. See [35, 206] for comprehensive discussion.

## - 12.8.1 Norms

Fact 12.8.1. $\boldsymbol{x}=0 \quad \Leftrightarrow \boldsymbol{x}^{\#}=0$.
Proof. If $\boldsymbol{x}=0$ then $\forall \boldsymbol{s} \neq 0$ we have $\langle\boldsymbol{x}, \boldsymbol{s}\rangle-\frac{1}{2}\|\boldsymbol{s}\|^{2}<0$ but $\langle\boldsymbol{x}, \boldsymbol{x}\rangle-\frac{1}{2}\|\boldsymbol{x}\|^{2}=0$. So we have $\boldsymbol{x}^{\#}=0$. If $\boldsymbol{x} \neq 0$ then let $\boldsymbol{s}=\frac{\langle\boldsymbol{x}, \boldsymbol{x}\rangle}{\|\boldsymbol{x}\|^{2}} \boldsymbol{x}$ with this choice we have $\langle\boldsymbol{x}, \boldsymbol{s}\rangle-\frac{1}{2}\|\boldsymbol{s}\|^{2}=\frac{1}{2} \frac{\langle\boldsymbol{x}, \boldsymbol{x}\rangle^{2}}{\|\boldsymbol{x}\|^{2}}>0$. However, for $\boldsymbol{s}=0$ we have that $\langle\boldsymbol{x}, \boldsymbol{s}\rangle-\frac{1}{2}\|\boldsymbol{s}\|^{2}=0$ therefore we have $\boldsymbol{x}^{\#} \neq 0$.

Fact 12.8.2. $\forall x \in \mathbb{R}^{n} \quad: \quad\left\langle x, x^{\#}\right\rangle=\left\|x^{\#}\right\|^{2}$.
Proof. If $\boldsymbol{x}=0$ then $\boldsymbol{x}^{\#}=0$ by Fact 12.8.1 and we have the result. Otherwise, again by Fact 12.8.1
we know that $\boldsymbol{x}^{\#} \neq 0$ and therefore by the definition of $\boldsymbol{x}^{\#}$ we have

$$
1=\underset{c \in \mathbb{R}}{\arg \max }\left\langle x, c \cdot x^{\#}\right\rangle-\frac{1}{2}\left\|c \cdot x^{\#}\right\|^{2}=\underset{c \in \mathbb{R}}{\arg \max } c \cdot\left\langle x, x^{\#}\right\rangle-\frac{c^{2}}{2}\left\|x^{\#}\right\|^{2}
$$

Setting the derivative of with respect to $c$ to 0 we get that $1=c=\frac{\left\langle\boldsymbol{x}, \boldsymbol{x}^{\#}\right\rangle}{\left\|\boldsymbol{x}^{\#}\right\|^{2}}$.
Fact 12.8.3. $\forall x \in \mathbb{R}^{n} \quad: \quad\|x\|^{*}=\left\|x^{\#}\right\|$.
Proof. Note that if $\boldsymbol{x}=0$ then the fact follows from Fact (12.8.1) otherwise we have

$$
\|\boldsymbol{x}\|^{*}=\max _{\|y\| \leq 1}\langle x, y\rangle=\max _{\|y\|=1}\langle x, y\rangle \leq \max _{y \in \mathbb{R}^{n}} \frac{\langle x, y\rangle}{\|y\|}
$$

From this it is clear that $\|\boldsymbol{x}\|^{*} \geq\left\|\boldsymbol{x}^{\#}\right\|$. To see the other direction consider a $\boldsymbol{y}$ that maximizes the above and let $\boldsymbol{z}=\frac{\langle\boldsymbol{x}, \boldsymbol{y}\rangle}{\|\boldsymbol{y}\|^{2}} \boldsymbol{y}$

$$
\langle\boldsymbol{x}, \boldsymbol{z}\rangle-\frac{1}{2}\|\boldsymbol{z}\|^{2} \leq\left\langle\boldsymbol{x}, \boldsymbol{x}^{\#}\right\rangle-\frac{1}{2}\left\|\boldsymbol{x}^{\#}\right\|^{2}
$$

and therefore

$$
\|\boldsymbol{x}\|^{* 2}-\frac{1}{2}\|\boldsymbol{x}\|^{* 2} \leq \frac{1}{2}\left\|\boldsymbol{x}^{\#}\right\|^{2}
$$

Fact 12.8.4 (Cauchy Schwarz). $\forall \boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{n} \quad: \quad\langle\boldsymbol{y}, \boldsymbol{x}\rangle \leq\|\boldsymbol{y}\|^{*}\|\boldsymbol{x}\|$.
Proof. By the definition of dual norm, for all $\|\boldsymbol{x}\|=1$, we have $\langle\boldsymbol{y}, \boldsymbol{x}\rangle \leq\|\boldsymbol{y}\|^{*}$. Hence, it follows by linearity of both side.

## ■ 12.8.2 Functions with Lipschitz Gradient

Lemma 12.8.5. Let $f$ be a continuously differentiable convex function. Then, the following are equivalence:

$$
\forall \boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{n} \quad: \quad\|\nabla f(\boldsymbol{x})-\nabla f(\boldsymbol{y})\|^{*} \leq L \cdot\|\boldsymbol{x}-\boldsymbol{y}\|
$$

and

$$
\forall \boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{n} \quad: \quad f(\boldsymbol{x}) \leq f(\boldsymbol{y})+\langle\nabla f(\boldsymbol{y}), \boldsymbol{x}-\boldsymbol{y}\rangle+\frac{L}{2}\|\boldsymbol{x}-\boldsymbol{y}\|^{2} .
$$

For any such $f$ and any $\boldsymbol{x} \in \mathbb{R}^{n}$, we have

$$
f\left(\boldsymbol{x}-\frac{1}{L} \nabla f(\boldsymbol{x})^{\#}\right) \leq f(\boldsymbol{x})-\frac{1}{2 L}\|\nabla f(\boldsymbol{x})\|^{* 2} .
$$

Proof. From the first condition, we have

$$
\begin{aligned}
f(\boldsymbol{y}) & =f(\boldsymbol{x})+\langle\nabla f(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\rangle+\int_{0}^{1}\langle\nabla f(\boldsymbol{x}+t(\boldsymbol{y}-\boldsymbol{x}))-\nabla f(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\rangle d t \\
& \leq f(\boldsymbol{x})+\langle\nabla f(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\rangle+\int_{0}^{1}\|\nabla f(\boldsymbol{x}+t(\boldsymbol{y}-\boldsymbol{x}))-\nabla f(\boldsymbol{x})\|^{*}\|\boldsymbol{y}-\boldsymbol{x}\| d t \\
& \leq f(x)+\langle\nabla f(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\rangle+\int_{0}^{1} L t\|\boldsymbol{y}-\boldsymbol{x}\|^{2} d t \\
& =f(x)+\langle\nabla f(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\rangle+\frac{L}{2}\|\boldsymbol{y}-\boldsymbol{x}\|^{2} .
\end{aligned}
$$

Given the second condition. For any $\boldsymbol{x} \in \mathbb{R}^{n}$. let $\phi_{\boldsymbol{x}}(\boldsymbol{y})=f(\boldsymbol{y})-\langle\nabla f(\boldsymbol{x}), \boldsymbol{y}\rangle$. From the convexity of $f$, for any $\boldsymbol{y} \in \mathbb{R}^{n}$

$$
f(\boldsymbol{y})-f(\boldsymbol{x}) \geq\langle\nabla f(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\rangle .
$$

Hence, $\boldsymbol{x}$ is a minimizer of $\phi_{\boldsymbol{x}}$. Hence, we have

$$
\begin{aligned}
\phi_{\boldsymbol{x}}(\boldsymbol{x}) & \leq \phi_{\boldsymbol{x}}\left(\boldsymbol{y}-\frac{1}{L} \nabla \phi_{\boldsymbol{x}}(\boldsymbol{y})^{\#}\right) \\
& \leq \phi_{\boldsymbol{x}}(\boldsymbol{y})-\left\langle\nabla \phi_{\boldsymbol{x}}(\boldsymbol{y}), \frac{1}{L} \nabla \phi_{\boldsymbol{x}}(\boldsymbol{y})^{\#}\right\rangle+\frac{L}{2}\left\|\frac{1}{L} \nabla \phi_{\boldsymbol{x}}(\boldsymbol{y})^{\#}\right\|^{2} \quad \text { (First part of this lemma) } \\
& =\phi_{\boldsymbol{x}}(\boldsymbol{y})-\frac{1}{2 L}\left\|\nabla \phi_{\boldsymbol{x}}(\boldsymbol{y})^{\#}\right\|^{2} \\
& =\phi_{\boldsymbol{x}}(\boldsymbol{y})-\frac{1}{2 L}\left(\left\|\nabla \phi_{\boldsymbol{x}}(\boldsymbol{y})\right\|^{*}\right)^{2} .
\end{aligned}
$$

Hence,

$$
f(\boldsymbol{y}) \geq f(\boldsymbol{x})+\langle\nabla f(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\rangle+\frac{1}{2 L}\left(\|\nabla f(\boldsymbol{y})-\nabla f(\boldsymbol{x})\|^{*}\right)^{2} .
$$

Adding up this inequality with $\boldsymbol{x}$ and $\boldsymbol{y}$ interchanged, we have

$$
\begin{aligned}
\frac{1}{L}\left(\|\nabla f(\boldsymbol{y})-\nabla f(\boldsymbol{x})\|^{*}\right)^{2} & \leq\langle\nabla f(\boldsymbol{y})-\nabla f(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\rangle \\
& \leq\|\nabla f(\boldsymbol{y})-\nabla f(\boldsymbol{x})\|^{*}\|\boldsymbol{y}-\boldsymbol{x}\|
\end{aligned}
$$

The last inequality follows from similar proof in above for $\phi_{\boldsymbol{x}}$.
The next lemma relate the Hessian of function with the Lipschitz parameter $L$ and this lemma gives us a easy way to compute $L$.

Lemma 12.8.6. Let $f$ be a twice differentiable function such that for any $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{n}$

$$
0 \leq \boldsymbol{y}^{T}\left(\nabla^{2} f(\boldsymbol{x})\right) \boldsymbol{y} \leq L\|\boldsymbol{y}\|^{2}
$$

Then, $f$ is convex and the gradient of $f$ is Lipschitz continuous with Lipschitz parameter $L$.
Proof. Similarly to Lemma 12.8.5, we have

$$
\begin{aligned}
f(\boldsymbol{y}) & =f(\boldsymbol{x})+\langle\nabla f(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\rangle+\int_{0}^{1}\langle\nabla f(\boldsymbol{x}+t(\boldsymbol{y}-\boldsymbol{x}))-\nabla f(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\rangle d t \\
& =f(\boldsymbol{x})+\langle\nabla f(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\rangle+\int_{0}^{1} t(\boldsymbol{y}-\boldsymbol{x})^{T} \nabla^{2} f\left(\boldsymbol{x}+\theta_{t}(\boldsymbol{y}-\boldsymbol{x})\right)(\boldsymbol{y}-\boldsymbol{x}) d t
\end{aligned}
$$

where the $0 \leq \theta_{t} \leq t$ comes from mean value theorem. By the assumption, we have

$$
\begin{aligned}
f(\boldsymbol{x})+\langle\nabla f(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\rangle & \leq f(\boldsymbol{y}) \\
& \leq f(\boldsymbol{x})+\langle\nabla f(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\rangle+\int_{0}^{1} t L\|\boldsymbol{y}-\boldsymbol{x}\|^{2} d t \\
& \leq f(\boldsymbol{x})+\langle\nabla f(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\rangle+\frac{L}{2}\|\boldsymbol{y}-\boldsymbol{x}\|^{2} .
\end{aligned}
$$

And the conclusion follows from Lemma 12.8.5.

# Connection Laplacian In Nearly-Linear Time 

## ■ 13.1 Introduction

In this chapter, we introduce the sparsified Cholesky and sparsified multigrid algorithms for solving systems of linear equations. Two advantages of these algorithms over other nearly-linear time algorithms for solving systems of equations in Laplacian matrices [256, 243, 147, 146, 141, 57] are:

1. They give nearly-linear time algorithms for solving systems of equations in a much broader class of matrices-the connection Laplacians and Hermitian block diagonally dominant matrices. Connection Laplacians $[238,29]$ are a generalization of graph Laplacians that arise in many applications, including celebrated work on cryo-electron microscopy [237, 235, 275], phase retrieval [7, 186], and many image processing problems (e.g. [221, 16]). Previous algorithms for solving systems of equations in graph Laplacians cannot be extended to solve equations in connection Laplacians because the previous algorithms relied on some form of low stretch spanning trees, a concept that has no analog for the more general connection Laplacians.
2. They provide linear-sized approximate inverses of connection Laplacian matrices. That is, for every $n$-dimensional connection Laplacian $M$, the sparsified Cholesky factorization algorithm produces an block-upper-triangular matrix $\boldsymbol{U}$ and block diagonal $\boldsymbol{D}$ with $O(n)$ nonzero entries so that $\boldsymbol{U}^{T} \boldsymbol{D} \boldsymbol{U}$ is a constant-factor approximation of $\boldsymbol{M}$. Such matrices $\boldsymbol{U}$ and $\boldsymbol{D}$ allow one to solve systems of equations in $M$ to accuracy $\varepsilon$ in time $O\left(m \log \varepsilon^{-1}\right)$, where $m$ is the number of nonzero entries of $\boldsymbol{M}$. Even for ordinary Laplacian matrices, the existence of such approximate inverses is entirely new. The one caveat of this result is that we do not yet know how to compute those approximate inverses in nearly linear time.

The sparsified Cholesky and sparsified multigrid algorithms work by sparsifying the matrices produced during Gaussian elimination. Recall that Cholesky factorization is the version of Gaussian elimination for symmetric matrices, and that the high running time of Gaussian elimination comes from fill-new nonzero matrix entries that are created by row operations. If Gaussian elimination never produced rows with a super-constant number of entries, then it would run in linear time. The sparsified Cholesky algorithm accelerates Gaussian elimination by sparsifying the rows that are produced by elimination, thereby guaranteeing that the elimination will be fast. Sparsified Cholesky is inspired by one of the major advances in algorithms for solving linear equations in Laplacian matrices-the Incomplete Cholesky factorization (ICC) [191]. However, ICC merely drops entries produced by elimination, whereas sparsification also carefully increases ones that remain. The difference is crucial, and is why ICC does not provide a nearly-linear time solver.

To control the error introduced by sparsification, we have to be careful not to do it too often. This means that our algorithm actually chooses a large set of rows and columns to eliminate at once, and then sparsifies the result. This is the basis of our first algorithms, which establish the existence of linear time solvers and linear sized approximate inverses, after precomputation. This precomputation
is analogous to the procedure of computing a matrix inverse: the approximate inverse takes much less time to apply than to compute.

To produce entire algorithms that run in nearly linear time (both to compute and apply the approximate inverse) requires a little more work. To avoid the work of computing the matrix obtained by eliminating the large set of rows and columns, we design a fast algorithm for approximating it quickly. This resulting algorithm produces a solver routine that does a little more work at each level, and so resembles a multigrid V-cycle [252]. We call the resulting algorithm the sparsified multigrid. We note that Krishnan, Fattal, and Szeliski [150] present experimental results from the use of a sparsification heuristic in a multigrid algorithm for solving problems in computer vision.

Our new algorithms are most closely related to the Laplacian solver recently introduced by Peng and Spielman [223]: unlike the other Laplacian solvers, they rely only on sparsification and do not directly rely on "support theory" preconditioners or any form of low stretch spanning trees. This is why our algorithms can solve a much broader family of linear systems. To sparsify without using graph theoretic algorithms, we employ the algorithm in Chapter 3 that allows us to sparsify a matrix by solving systems of equations in a subsampled matrix.

## ■ 13.1.1 Connection Laplacians and Block DD Matrices

In this section, we define block diagonally dominant (bDD) matrices-the most general family of matrices such that the associated systems of linear equations can be solved by our algorithms. We begin by defining our motivating case: the connection Laplacians

Connection Laplacians may be thought of as a generalization of graph Laplacians where every vertex is associated with a vector, instead of a real number, and every edge is associated with a unitary matrix. Like graph Laplacians, they describe a natural quadratic form. Let $\boldsymbol{M}_{[i, j]} \in \mathbb{C}^{r \times r}$ be the unitary matrix associated with edge ( $i, j$ ), and let $w_{i, j}$ be the (nonnegative) weight of edge $(i, j)$. We require that $\boldsymbol{M}_{[i, j]}=M_{[j, i]}^{*}$, where $*$ denotes the conjugate transpose. The quadratic form associated with this connection Laplacian is a function of vectors $\boldsymbol{v}^{(i)} \in \mathbb{C}^{r}$, one for each vertex $i$, that equals

$$
\sum_{(i, j) \in E} \boldsymbol{w}_{i, j}\left\|\boldsymbol{v}^{(i)}-\boldsymbol{M}_{[i, j]} \boldsymbol{v}^{(j)}\right\|^{2}
$$

The matrix corresponding to this quadratic form is a block matrix with blocks $\boldsymbol{M}_{[i, j]}$. Most applications of the connection Laplacian require one to either solve systems of linear equations in this matrix, or to compute approximations of its smallest eigenvalues and eigenvectors. By applying the inverse power method (or inverse Lanczos), we can perform these eigenvector calculations by solving a logarithmic number of linear systems in the matrix (see [243, Section 7]).

The matrices obtained from the connection Laplacian are a special case of block diagonally dominant ( bDD ) matrices, which we now define. Throughout this chapter, we consider block matrices having entries in $\mathbb{C}^{r \times r}$, where $r>0$ is a fixed integer. We say that $M \in\left(\mathbb{C}^{r \times r}\right)^{m \times n}$ has $m$ block-rows, and $n$ block-columns. For $i \in[m], j \in[n]$, we let $\boldsymbol{M}_{[i, j]} \in \mathbb{C}^{r \times r}$ denote the $i, j$-block in $\boldsymbol{M}$, and $\boldsymbol{M}_{[j]}$ denote the $j^{\text {th }}$ block-column, i.e., $\boldsymbol{M}_{[j]}=\left[\left(\boldsymbol{M}_{[1, j]}\right)^{*},\left(\boldsymbol{M}_{[2, j]}\right)^{*}, \ldots,\left(\boldsymbol{M}_{[m, j]}\right)^{*}\right]^{*}$. For sets $F \subseteq[m], C \subseteq[n]$, we let $M_{[F, C]} \in\left(\mathbb{C}^{r \times r}\right)^{|F| \times|C|}$ denote the block-submatrix with blocks $M_{[i, j]}$ for $i \in F, j \in C . M$ is block-diagonal, if $\boldsymbol{M}_{[i, j]}=0$ for $i \neq j$. We emphasize that all computations are done over $\mathbb{C}$, not over a matrix group.

Definition 13.1.1. A Hermitian block-matrix $M \in\left(\mathbb{C}^{r \times r}\right)^{n \times n}$ is block diagonally dominant (or bDD) if

$$
\text { for all } i \in[n], \quad \boldsymbol{M}_{[i, i]} \succeq \boldsymbol{I}_{r} \cdot \sum_{j: j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\|_{2},
$$

where $\boldsymbol{I}_{r} \in \mathbb{C}^{r \times r}$ denotes the identity matrix.

For the rest of this chapter, we use $\|\cdot\|$ to denote $\|\cdot\|_{2}$. Equivalently, a Hermitian $\boldsymbol{M}$ is a bDD matrix iff it can be written as $\boldsymbol{D}-\boldsymbol{A}$ where $\boldsymbol{D}$ is block-diagonal, and $\boldsymbol{D}_{[i, i]} \succeq \boldsymbol{I}_{r} \sum_{j}\left\|\boldsymbol{A}_{[i, j]}\right\|$ (see Lemma 13.6.1).
Throughout this chapter we treat $r$ as a constant. The hidden dependence on $r$, of the running times of the algorithms we present, is polynomial. The major results of this chapter are the following.

Theorem 13.1.2 (Sparsified Multigrid). There is an algorithm that, when given a bDD matrix $M$ with $n$ block rows and $m$ nonzero blocks, produces a solver for $\boldsymbol{M}$ in $O\left(m \log n+n \log ^{2+o(1)} n\right)$ work and $O\left(n^{o(1)}\right)$ depth so that the solver finds $\varepsilon$-approximate solutions to systems of equations in $\boldsymbol{M}$ in $O\left(\left(m+n \log ^{1+o(1)} n\right) \log (1 / \varepsilon)\right)$ work and $O\left(\log ^{2} n \log \log n \log (1 / \varepsilon)\right)$ depth.

Previously, the existence of nearly-linear time solvers was even unknown for the 1-dimensional case $(r=1)$ when the off-diagonal entries were allowed to be complex numbers.

We can use the above algorithm to find approximations of the smallest eigenvalues and eigenvectors of such matrices at an additional logarithmic cost.

Theorem 13.1.3 (Sparsified Cholesky). For every bDD matrix $M$ with $n$ block-rows there exists a diagonal matrix $\boldsymbol{D}$ and an upper triangular matrix $\boldsymbol{U}$ with $O(n)$ nonzero blocks so that

$$
\boldsymbol{U}^{T} \boldsymbol{D} \boldsymbol{U} \approx_{3 / 4} \boldsymbol{M}
$$

Moreover, linear equations in $\boldsymbol{U}, \boldsymbol{U}^{T}$, and $\boldsymbol{D}$ can be solved with linear work in depth $O\left(\log ^{2} n\right)$, and these matrices can be computed in polynomial time.

These matrices allow one to solve systems of linear equations in $M$ to $\varepsilon$-accuracy in parallel time $O\left(\log ^{2} n \log ^{-1} \varepsilon\right)$ and work $O\left(m \log ^{-1} \varepsilon\right)$. Results of this form were previously unknown even for graph Laplacians.

In the next two sections we explain the ideas used to prove these theorems. Proofs may be found in the appendices that follow.

## ■ 13.2 Preliminaries

We say that $\boldsymbol{A}$ is an $\varepsilon$-approximation of $\boldsymbol{B}$, written $\boldsymbol{A} \approx_{\varepsilon} \boldsymbol{B}$, if $e^{\varepsilon} \boldsymbol{B} \succcurlyeq \boldsymbol{A} \succcurlyeq e^{-\varepsilon} \boldsymbol{B}$. This relation is symmetric. We say that $\tilde{\boldsymbol{x}}$ is an $\varepsilon$-approximate solution to the system $\boldsymbol{A} \boldsymbol{x}=\boldsymbol{b}$ if $\left\|\tilde{\boldsymbol{x}}-\boldsymbol{A}^{-1} \boldsymbol{b}\right\|_{\boldsymbol{A}} \leq$ $\varepsilon\|\boldsymbol{x}\|_{\boldsymbol{A}}$, where $\|\boldsymbol{x}\|_{\boldsymbol{A}}=\left(\boldsymbol{x}^{T} \boldsymbol{A} \boldsymbol{x}\right)^{1 / 2}$. If $\varepsilon<1 / 2, \boldsymbol{A} \approx_{\varepsilon} \boldsymbol{B}$ and $\boldsymbol{B} \tilde{\boldsymbol{x}}=\boldsymbol{b}$, then $\tilde{\boldsymbol{x}}$ is a $2 \varepsilon$ approximate solution to $\boldsymbol{A} \boldsymbol{x}=\boldsymbol{b}$.

Fact 13.2.1. If $\boldsymbol{A} \approx_{\varepsilon} \boldsymbol{B}$ and $\boldsymbol{B} \approx_{\delta} \boldsymbol{C}$, then $\boldsymbol{A} \approx_{\varepsilon+\delta} \boldsymbol{C}$.
We now address one technicality of dealing with $b D D$ matrices: it is not immediate whether or not a $b D D$ matrix is singular. Moreover, if it is singular, the structure of its null space is not immediately clear either. Throughout the rest of this chapter, we will consider $b D D$ matrices to which a small multiple of the identity have been added. These matrices will be nonsingular. To reduce the problem of solving equations in a general $b D D$ matrix $\boldsymbol{M}$ to that of solving equations in a nonsingular matrix, we require an estimate of the smallest nonzero eigenvalue of $\boldsymbol{M}$.

Fact 13.2.2. Suppose that all nonzero eigenvalues of $\boldsymbol{M}$ are at least $\mu$ and $\boldsymbol{Z} \approx_{\varepsilon}(\boldsymbol{M}+\varepsilon \mu \boldsymbol{I})^{-1}$ for some $0<\varepsilon<1 / 2$. Given any $\boldsymbol{b}$ such that $\boldsymbol{M} \boldsymbol{x}=\boldsymbol{b}$ for some $\boldsymbol{x}$, we have

$$
\|\boldsymbol{x}-\boldsymbol{Z} \boldsymbol{b}\|_{\boldsymbol{M}}^{2} \leq 6 \varepsilon\|x\|_{\boldsymbol{M}}^{2}
$$

Hence, we can solve systems in $\boldsymbol{M}$ by approximately solving systems in $\boldsymbol{M}+\varepsilon \mu \boldsymbol{I}$. Any lower bound $\mu$ on the smallest nonzero eigenvalue of $M$ will suffice. It only impacts the running times of our algorithms in the numerical precision with which one must carry out the computations.

The above fact allows us to solve systems in $M$ that have a solution. If $M$ is singular and we want to apply its pseudoinverse (that is, to find the closest solution in the range of $\boldsymbol{M}$ ), we can do so by pre and post multiplying by $M$ to project onto its range. The resulting algorithm, which is implicit in the following fact, requires applying a solver for $M$ three times. It also takes $O(\log \kappa)$ times as long to run, where $\kappa$ is the finite condition number ${ }^{1}$ of $\boldsymbol{M}$.

Fact 13.2.3. Let $\kappa$ be an upper bound on the finite condition number of $\boldsymbol{M}$. Given an error parameter $0<\varepsilon<1$, let $\delta=\varepsilon / 56 \kappa^{3}$. If all nonzero eigenvalues of $M$ are at least $\mu$, and if $\boldsymbol{Z} \approx_{\delta}(\boldsymbol{M}+\varepsilon \mu \boldsymbol{I})^{-1}$, then

$$
M Z^{3} M \approx_{4 \varepsilon} M^{+}
$$

## ■ 13.3 Overview of the Algorithms

We index the block rows and columns of a block matrix by a set of vertices (or indices) $V$. When we perform an elimination, we eliminate a set $F \subset V$, and let $C=V-F$. Here, $F$ stands for "fine" and $C$ stands for "coarse". In contrast with standard multigrid methods, we will have $|F| \leq|C|$.

To describe block-Cholesky factorization, we write the matrix $M$ with the rows and columns in $F$ first:

$$
\boldsymbol{M}=\left[\begin{array}{ll}
\boldsymbol{M}_{[F, F]} & \boldsymbol{M}_{[F, C]} \\
\boldsymbol{M}_{[C, F]} & \boldsymbol{M}_{[C, C]}
\end{array}\right] .
$$

Cholesky factorization writes the inverse of this matrix as

$$
\boldsymbol{M}^{-1}=\left[\begin{array}{cc}
\boldsymbol{I} & -\boldsymbol{M}_{[F, F]}^{-1} \boldsymbol{M}_{[F, C]}  \tag{13.1}\\
0 & \boldsymbol{I}
\end{array}\right]\left[\begin{array}{cc}
\boldsymbol{M}_{[F, F]}^{-1} & 0 \\
0 & S c(\boldsymbol{M}, F)^{-1}
\end{array}\right]\left[\begin{array}{cc}
\boldsymbol{I} & 0 \\
-\boldsymbol{M}_{[C, F]} \boldsymbol{M}_{[F, F]}^{-1} & \boldsymbol{I}
\end{array}\right],
$$

where

$$
S c(\boldsymbol{M}, F) \stackrel{\text { def }}{=} \boldsymbol{M}_{[C, C]}-\boldsymbol{M}_{[C, F]} \boldsymbol{M}_{[F, F]}^{-1} \boldsymbol{M}_{[F, C]}
$$

is the Schur complement of $M$ with respect to $F$.
Our algorithms rely on two fundamental facts about $b D D$ matrices: that the Schur complement of a $b D D$ matrix is a $b D D$ matrix (Lemma 13.6.3) and that one can sparsify $b D D$ matrices. The following theorem is implicit in [46].

Theorem 13.3.1. For every $\varepsilon \leq 1$, every $b D D$ matrix $M$ with $n r$-dimensional block rows and columns can be $\varepsilon$-approximated by a bDD matrix having at most $10 \mathrm{nr} / \varepsilon^{2}$ nonzero blocks.

We use the identity (13.1) to reduce the problem of solving a system of equations in $M$ to that of solving equations in its Schur complement. The easiest part of this is multiplication by $\boldsymbol{M}_{[C, F]}$ : the time is proportional to the number of nonzero entries in the submatrix, and we can sparsify $M$ to guarantee that this is small.

The costlier part of the reduction is the application of the inverse of $\boldsymbol{M}_{[F, F]}$ three times. This would be fast if $\boldsymbol{M}_{[F, F]}$ were block-diagonal, which corresponds to $F$ being an independent set. We cannot find a sufficiently large independent set $F$, but we can find a large set that is almost independent. This results in a matrix $\boldsymbol{M}_{[F, F]}$ that is well approximated by its diagonal, and thus linear equations

[^32]in this matrix can be quickly solved to high accuracy by a few Jacobi iterations (see Theorem 13.3.1 and Section 13.9).

We can prove the existence of linear sized approximate inverses by explicitly computing $S c(\boldsymbol{M}, F)$, sparsifying it, and then recursively applying the algorithm just described. To make this algorithm efficient, we must compute a sparse approximation to $S c(\boldsymbol{M}, F)$ without constructing $S c(\boldsymbol{M}, F)$. This is the problem of spectral vertex sparsification, and we provide a fast algorithm for this task in Sections 13.3.4 and 13.11.

## ■ 13.3.1 Schur Complement Chains

We encode this recursive algorithm by a Schur complement chain (SCC). An SCC defines a linear operator that can be used to approximately solve equations in an initial matrix $\boldsymbol{M}^{(0)}$. If the $b D D$ matrix $\boldsymbol{M}$ is sparse, then it is the same as $\boldsymbol{M}^{(0)}$; if not, then $\boldsymbol{M}^{(0)}$ is a sparse approximation to $\boldsymbol{M}$. Let $F_{1}$ be the first set of vertices eliminated, $\boldsymbol{M}^{(1)}$ a sparse approximation to the Schur complement of $\boldsymbol{M}^{(0)}$ with respect to $F_{1}$, and $\boldsymbol{Z}^{(1)}$ an operators that approximates the inverse of $\boldsymbol{M}_{\left[F_{1}, F_{1}\right]}^{(0)}$.
Definition 13.3.2 (Schur Complement Chain). An $\varepsilon$-Schur complement chain ( $\varepsilon$-SCC) for a matrix $\boldsymbol{M}^{0}$ indexed by vertex set $V$ is a sequence of operators and subsets,

$$
\left(\left(\boldsymbol{M}^{(1)}, \boldsymbol{Z}^{(1)}\right), \ldots,\left(\boldsymbol{M}^{(d)}, \boldsymbol{Z}^{(d)}\right) ; F_{1}, \ldots, F_{d}\right)
$$

so that for $C_{0}=V$ and $C_{i+1}=C_{i} \backslash F_{i+1},\left|C_{d}\right| \leq 1000$ and for $1 \leq i \leq d$,

$$
\boldsymbol{M}^{(i)} \approx_{\varepsilon_{i}} S c\left(\boldsymbol{M}^{(i-1)}, F_{i}\right) \quad \text { and } \quad 0 \preceq\left(\boldsymbol{Z}^{(i)}\right)^{-1}-\boldsymbol{M}_{\left[F_{i}, F_{i}\right]}^{(i-1)} \preceq \varepsilon_{i} \cdot S c\left(\boldsymbol{M}^{(i-1)}, C_{i}\right) .
$$

The algorithm ApplyChain, described in Section 13.7, applies an SCC to solve equations in $\boldsymbol{M}^{(0)}$ in the natural way and satisfies the following guarantee.
Lemma 13.3.1. Consider an $\boldsymbol{\varepsilon}$-SCC for $\boldsymbol{M}^{(0)}$ where $\boldsymbol{M}^{(i)}$ and $\boldsymbol{Z}^{(i)}$ can be applied to a vector in work $W_{\boldsymbol{M}^{(i)}}, W_{\boldsymbol{D}^{(i)}}$ and depth $D_{\boldsymbol{M}^{(i)}}, D_{\boldsymbol{Z}^{(i)}}$ respectively.

The algorithm ApplyChain $\left(\left(\boldsymbol{M}^{(1)}, \boldsymbol{Z}^{(1)}\right), \ldots,\left(\boldsymbol{M}^{(d)}, \boldsymbol{Z}^{(d)}\right) ; F_{1}, \ldots, F_{d} ; \boldsymbol{b}\right)$ corresponds to a linear operator $\boldsymbol{W}$ acting on $\boldsymbol{b}$ such that

1. $\boldsymbol{W}^{-1} \approx_{\sum_{i=1}^{d} 2 \varepsilon_{i}} \boldsymbol{M}^{(0)}$, and
2. for any vector $\boldsymbol{b}$, ApplyChain $\left(\left(\boldsymbol{M}^{(1)}, \boldsymbol{Z}^{(1)}\right), \ldots,\left(\boldsymbol{M}^{(d)}, \boldsymbol{Z}^{(d)}\right) ; F_{1}, \ldots, F_{d} ; \boldsymbol{b}\right)$ runs in $O\left(\sum_{i=1}^{d}\left(D_{\boldsymbol{M}^{(i)}}+D_{\boldsymbol{Z}^{(i)}}\right)\right)$ depth and $O\left(\sum_{i=1}^{d}\left(W_{\boldsymbol{M}^{(i)}}+W_{\boldsymbol{Z}^{(i)}}\right)\right)$ work.

For an $\varepsilon$-SCC chain for $\boldsymbol{M}^{(0)}$, we define the depth and work of the chain to be the work and depth required by ApplyChain to apply the SCC.

## - 13.3.2 Choosing $F_{i}$ : $\alpha$-bDD Matrices

We must choose the set of vertices $F_{i}$ so that we can approximate the inverse of $\boldsymbol{M}_{\left[F_{i}, F_{i}\right]}^{(i)}$ by an operator $\boldsymbol{Z}^{(i)}$ that is efficiently computable. We do this by requiring that the matrix $\boldsymbol{M}_{\left[F_{i}, F_{i}\right]}^{(i)}$ be $\alpha$-block diagonally dominant $(\alpha-\mathrm{bDD})$, a term that we now define.

Definition 13.3.3. A Hermitian block-matrix $\boldsymbol{M}$ is $\alpha$-bDD if

$$
\begin{equation*}
\forall i, \quad \boldsymbol{M}_{[i, i]} \succcurlyeq(1+\alpha) \boldsymbol{I}_{r} \cdot \sum_{j: j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\| . \tag{13.2}
\end{equation*}
$$

We remark that a 0 -bDD matrix is simply a bDD matrix. In particular, for $r=1$, Laplacian matrices are $0-\mathrm{bDD}$.
By picking a subset of rows at random and discarding those that violate condition (13.2), the algorithm BDDSubset (described in Section 13.8) finds a linear sized subset $F$ of the block-rows of a bDD matrix $\boldsymbol{M}$ so that $\boldsymbol{M}_{[F, F]}$ is $\alpha$-bDD.
Lemma 13.3.2. Given a bDD matrix $\boldsymbol{M}$ with $n$ block-rows, and an $\alpha \geq 0$, BDDSUBSET computes a subset $F$ of size least $n /(8(1+\alpha))$ such that $\boldsymbol{M}_{[F, F]}$ is $\alpha$-bDD. It runs in runs in $O(m)$ expected work and $O(\log n)$ expected depth, where $m$ is the number of nonzero blocks in $\boldsymbol{M}$.

We can express an $\alpha$-bDD matrix as a sum of a block diagonal matrix and a bDD matrix so that it is well-approximated by the diagonal.
Lemma 13.3.3. Every $\alpha$-bDD matrix $\boldsymbol{M}$ can be written in the form $\boldsymbol{X}+\boldsymbol{L}$ where $\boldsymbol{X}$ is block-diagonal, $\boldsymbol{L}$ is bDD, and $\boldsymbol{X} \succcurlyeq \frac{\alpha}{2} \boldsymbol{L}$.

As $\boldsymbol{L}$ is positive semidefinite, $(1+2 / \alpha) \boldsymbol{X} \succcurlyeq \boldsymbol{M} \succcurlyeq \boldsymbol{X}$, which means that $\boldsymbol{X}$ is a good approximation of $\boldsymbol{M}$ when $\alpha$ is reasonably big. As block-diagonal matrices like $\boldsymbol{X}$ are easy to invert, systems in these well-conditioned matrices can be solved rapidly using preconditioned iterative methods. In Section 13.9, we show that a variant of Jacobi iteration provides an operator that satisfies the requirements of Definition 13.3.2.

Theorem 13.3.1. Let $\boldsymbol{M}$ be a bDD matrix with index set $V$, and let $F \subseteq V$ such that $\boldsymbol{M}_{[F, F]}$ is $\alpha-b D D$ for some $\alpha \geq 4$, and has $m_{F F}$ nonzero blocks. The algorithm $\operatorname{JACOBI}\left(\varepsilon, \boldsymbol{M}_{[F, F]}, \boldsymbol{b}\right)$ acts as a linear operator $\boldsymbol{Z}$ on $\boldsymbol{b}$ that satisfies

$$
0 \preceq(\boldsymbol{Z})^{-1}-\boldsymbol{M}_{[F, F]} \preceq \varepsilon \cdot S c(\boldsymbol{M}, V \backslash F) .
$$

The algorithm takes $O\left(m_{F F} \log \left(\frac{1}{\varepsilon}\right)\right)$ work and $O\left(\log n \log \left(\frac{1}{\varepsilon}\right)\right)$ depth.
We now explain how Theorems 13.3.1 and 13.3.1 allow us to construct Schur complement chains that can be applied in nearly linear time. We optimize the construction in the next section.

Theorem 13.3.1 tells us that there is a $b D D$ matrix $M^{(0)}$ with $O\left(n / \varepsilon^{2}\right)$ nonzero blocks that that $\varepsilon$-approximates $\boldsymbol{M}$, and that for every $i$ there is a $b D D$ matrix $\boldsymbol{M}^{(i+1)}$ with $O\left(\left|C_{i}\right| / \varepsilon^{2}\right)$ nonzero blocks that is an $\varepsilon$-approximation of $S c\left(\boldsymbol{M}^{(i)}, F_{i}\right)$. We will pick $\varepsilon$ later. Lemma 13.3 .2 provides a set $F_{i}$ containing a constant fraction of the block-rows of $\boldsymbol{M}^{(i)}$ so that $\boldsymbol{M}_{\left[F_{i}, F_{i}\right]}^{(i)}$ is 4-bDD. Theorem 13.3.1 then provides an operator that solves systems in $\boldsymbol{M}_{\left[F_{i}, F_{i}\right]}^{(i)}$ to $\varepsilon$ accuracy in time $O\left(\log \varepsilon^{-1}\right)$ times the number of nonzero entries in $\boldsymbol{M}_{\left[F_{i}, F_{i}\right]}^{(i)}$. This is at most the number of nonzero entries in $\boldsymbol{M}^{(i)}$, and thus at most $O\left(\left|C_{i}\right| / \varepsilon^{2}\right)$. As each $F_{i}$ contains at least a constant fraction of the rows of $\boldsymbol{M}^{(i)}$, the depth of the recursion, $d$, is $O(\log n)$. Thus, we can obtain constant accuracy by setting $\varepsilon=\Theta(1 / \log n)$. The time required to apply the resulting SCC would thus be $O\left(n \log ^{2} n \log \log n\right)$.

We can reduce the running time by setting $\varepsilon_{1}$ to a constant and allowing it to shrink as $i$ increases. For example, setting $\varepsilon_{i}=1 / 2(i+1)^{2}$ results in a linear time algorithm that produces a constant-factor approximation of the inverse of $\boldsymbol{M}$. We refine this idea in the next section.

## ■ 13.3.3 Linear Sized Approximate Inverses

In this section we sketch the proof of Theorem 13.1.3, which tells us that every $b D D$ matrix has a linear-sized approximate inverse. The rest of the details appear in Section 13.10.

The linear-sized approximate inverse of a $b D D$ matrix $M$ with $n$ block rows and columns is provided by a $3 / 4$-approximate of the form $\boldsymbol{U}^{T} \boldsymbol{D} \boldsymbol{U}$ where $\boldsymbol{D}$ is block diagonal and $\boldsymbol{U}$ is block upper-triangular
and has $O(n)$ nonzero blocks. As systems of linear equations in block-triangular matrices like $\boldsymbol{U}$ and $\boldsymbol{U}^{T}$ can be solved in time proportional to their number of nonzero blocks, this provides a linear time algorithm for computing a $3 / 4$ approximation of the inverse of $M$. By iteratively refining the solutions provided by the approximate inverse, this allows one to find $\varepsilon$-accurate solutions to systems of equations in $M$ in time $O\left(m \log \varepsilon^{-1}\right)$.

The matrix $\boldsymbol{U}$ that we construct has meta-block structure that allows solves in $\boldsymbol{U}$ and $\boldsymbol{U}^{T}$ to be performed with linear work and depth $O\left(\log ^{2} n\right)$. This results in a parallel algorithm for solving equations in $M$ in work $O\left(m \log \varepsilon^{-1}\right)$ and depth $O\left(\log ^{2} n \log \varepsilon^{-1}\right)$. We remark that with some additional work (analogous to Section 7 of [160]), one could reduce this depth to $O\left(\log n \log \log n \log \varepsilon^{-1}\right)$.

The key to constructing $\boldsymbol{U}$ is realizing that the algorithm Jacobi corresponds to multiplying by the matrix $\boldsymbol{Z}^{(k)}$ defined in equation (13.7). Moreover, this matrix is a polynomial in $\boldsymbol{X}$ and $\boldsymbol{L}$ of degree $\log (3 / \varepsilon)$, where $\boldsymbol{M}_{\left[F_{i}, F_{i}\right]}^{(i)}=\boldsymbol{X}+\boldsymbol{L}$ where $\boldsymbol{L}$ is bDD, and $\boldsymbol{X}$ block diagonal such that $\boldsymbol{X} \succcurlyeq 2 \boldsymbol{L}$. To force $\boldsymbol{Z}^{k}$ to be a sparse matrix, we require $\boldsymbol{L}$ be sparse.

If we use the algorithm BSDDSUBSET to choose $F_{i}$, then $L$ need not be sparse. However, this problem is easily remedied by forbidding algorithm BSDDSUBSET from choosing any vertex of more than twice average degree. Thus, we can ensure that all $\boldsymbol{Z}^{(i)}$ are sparse.

Lemma 13.3.4. For every $b D D$ matrix $M$ and every $\alpha \geq 0$, there is a subset $F$ of size at least $\frac{n}{16(1+\alpha)}$ such that $M_{[F, F]}$ is $\alpha-b D D$ and the number of nonzeros blocks in each block-row of $F$ at most twice the average number of nonzero blocks in a block-row of $\boldsymbol{M}$.

Proof. Discard every block-row of $\boldsymbol{M}$ that has more than twice the average number of nonzeros blocks per row-block. Then remove the corresponding row blocks. The remaining matrix has dimension at least $n / 2$. We can now use Lemma 13.3.2 to find an $\alpha$-bDD submatrix.

We obtain the $\boldsymbol{U}^{T} \boldsymbol{D} \boldsymbol{U}$ factorization by applying the inverse of the factorization (13.1):

$$
\boldsymbol{M}=\left[\begin{array}{cc}
\boldsymbol{I} & 0 \\
\boldsymbol{M}_{[F, F]}^{-1} \boldsymbol{M}_{[F, C]} & \boldsymbol{I}
\end{array}\right]\left[\begin{array}{cc}
\boldsymbol{M}_{[F, F]} & 0 \\
0 & S c(\boldsymbol{M}, F)
\end{array}\right]\left[\begin{array}{cc}
\boldsymbol{I} & \boldsymbol{M}_{[C, F]} \boldsymbol{M}_{[F, F]}^{-1} \\
0 & \boldsymbol{I}
\end{array}\right] .
$$

In the left and right triangular matrices we replace $\boldsymbol{M}_{[F, F]}^{-1}$ with the polynomial we obtain from Jacobi. In the middle matrix, it suffices to approximate $\boldsymbol{M}_{[F, F]}$ by a block diagonal matrix, and $S c(M, F)$ by a factorization of its sparse approximation given by Theorem 13.3.1. The details, along with a careful setting of $\varepsilon_{i}$, are carried out in Section 13.10.

## ■ 13.3.4 Spectral Vertex Sparsification Algorithm

In this section, we outline a procedure ApproxSchur that efficiently approximates the Schur complement of a bDD matrix $M$ w.r.t. a set of indices $F$ s.t. $\boldsymbol{M}_{[F, F]}$ is $\alpha$-bDD. The following lemma summarizes the guarantees of ApproxSchur.

Lemma 13.3.4. Let $M$ be a bDD matrix with index set $V$, and $m$ nonzero blocks. Let $F \subseteq V$ be such that $\boldsymbol{M}_{[F, F]}$ is $\alpha$-bDD for some $\alpha \geq 4$. The algorithm $\operatorname{ApproxSchur}(\boldsymbol{M}, F, \varepsilon)$, returns a matrix $\widetilde{\boldsymbol{M}}_{S C}$ s.t.

1. $\widetilde{\boldsymbol{M}}_{S C}$ has $O\left(m\left(\varepsilon^{-1} \log \log \varepsilon^{-1}\right)^{O\left(\log \log \varepsilon^{-1}\right)}\right)$ nonzero blocks, and
2. $\widetilde{\boldsymbol{M}}_{S C} \approx_{\varepsilon} S c(\boldsymbol{M}, F)$,
in $O\left(m\left(\varepsilon^{-1} \log \log \varepsilon^{-1}\right)^{O\left(\log \log \varepsilon^{-1}\right)}\right)$ work and $O\left(\log n\left(\log \log \varepsilon^{-1}\right)\right)$ depth.

We sketch a proof of the above lemma in this section. A complete proof and pseudocode for ApproxSchur are given in Section 13.11.
First consider a very simple special case: where $F$ is a singleton, $F=\{i\}$. Let $C=V \backslash F$ be the remaining indices.

$$
S c(\boldsymbol{M}, i)=\boldsymbol{M}_{[C, C]}-\boldsymbol{M}_{[C, i]} \boldsymbol{M}_{[i, i]}^{-1} \boldsymbol{M}_{[i, C]}
$$

Thus, if $\boldsymbol{M}_{[C, i]}$ has $k$ nonzero blocks, $S c(\boldsymbol{M}, i)$ could have $k^{2}$ additional nonzero blocks compared to $\boldsymbol{M}_{[C, C]}$, potentially making it dense. If $\boldsymbol{M}$ were a graph Laplacian, then $\boldsymbol{M}_{[C, i]} \boldsymbol{M}_{[i, i]}^{-1} \boldsymbol{M}_{[i, C]}$ would represent the adjacency matrix of a weighted clique. In Section 13.14, we construct weighted expanders that allow us to $\varepsilon$-approximate $S c(\boldsymbol{M}, i)$ in this case using $m+O\left(k \varepsilon^{-4}\right)$ edges. In Section 13.11.4, we show how to use such weighted expanders to sparsify $S c(\boldsymbol{M}, i)$ when $\boldsymbol{M}$ is a bDD matrix.

This reduction can also be performed in parallel. If $F$ is such that $\boldsymbol{M}_{[F, F]}$ is block diagonal, we can approximate $S c(\boldsymbol{M}, F)$ by expressing $\boldsymbol{M}_{[C, i]} \boldsymbol{M}_{[i, i]}^{-1} \boldsymbol{M}_{[i, C]}$ as $\sum_{i \in F} \boldsymbol{M}_{[C, i]} \boldsymbol{M}_{[i, i]}^{-1} \boldsymbol{M}_{[i, C]}$, and using $|F|$ weighted expanders. However, $\boldsymbol{M}_{[F, F]}$ may not be diagonal. Instead, we give a procedure SchurSquare that generates $\boldsymbol{M}^{\prime}$ that is better approximated by its diagonal.

Invoking SchurSquare a few times leads a sequence of matrices $M_{[F, F]}^{(0)}, M_{[F, F]}^{(1)}, \ldots, M_{[F, F]}^{(i)}$. We will show that $\boldsymbol{M}_{[F, F]}^{(i)}$ is $\varepsilon$-approximated by its diagonal and we call the procedure LASTSTEP to approximate $S c\left(\boldsymbol{M}^{(i)}, F\right)$. An additional caveat is that replacing $\boldsymbol{M}_{[F, F]}^{(i)}$ by its diagonal at this step gives errors that are difficult to bound. We discuss the correct approximation below

SchurSquare is based on a squaring identity for matrix inverse developed in [223]. Given a splitting of $\boldsymbol{M}_{[F, F]}$ into $\boldsymbol{D}-\boldsymbol{A}$, where $\boldsymbol{D}$ is block-diagonal, and $\boldsymbol{A}$ has its diagonal blocks as zero, it relies on the fact that the matrix

$$
\boldsymbol{M}_{2}=\frac{1}{2}\left[\begin{array}{cc}
\boldsymbol{D}-\boldsymbol{A} \boldsymbol{D}^{-1} \boldsymbol{A} & \boldsymbol{M}_{[F, C]}+\boldsymbol{A} \boldsymbol{D}^{-1} \boldsymbol{M}_{[F, C]}  \tag{13.3}\\
\boldsymbol{M}_{[C, F]}+\boldsymbol{M}_{[C, F]} \boldsymbol{D}^{-1} \boldsymbol{A} & 2 \boldsymbol{M}_{[C, C]}-\boldsymbol{M}_{[C, F]} \boldsymbol{D}^{-1} \boldsymbol{M}_{[F, C]}
\end{array}\right]
$$

satisfies $S c(\boldsymbol{M}, F)=S c\left(\boldsymbol{M}_{2}, F\right)$. Furthermore, we can show that if $\boldsymbol{M}$ is $\alpha-\mathrm{bDD}, \boldsymbol{D}-\boldsymbol{A} \boldsymbol{D}^{-1} \boldsymbol{A}$ is $\alpha^{2}$-bDD, which indicates that the block on $[F, F]$ rapidly approaches being diagonal.

As $\boldsymbol{M}_{2}$ may be dense, we construct a sparse approximation to it. Since $\boldsymbol{D}$ is diagonal, we can construct sparse approximations to the blocks $\left(\boldsymbol{M}_{2}\right)_{[F, F]}$ and $\left(\boldsymbol{M}_{2}\right)_{[C, C]}$ in a manner analogous to the case of diagonal $\boldsymbol{M}_{[F, F]}$. Similarly, we use bipartite expanders to construct sparse approximations to $\left(\boldsymbol{M}_{2}\right)_{[C, F]}$ and $\left(\boldsymbol{M}_{2}\right)_{[F, C]}$.

Our sequence of calls to SchurSquare terminates with $M_{[F, F]}^{(i)}$ being roughly $\varepsilon^{-1}$-bDD. We then return LastStep $\left(\boldsymbol{M}^{(i)}, F, \varepsilon^{-1}, \varepsilon\right)$. As mentioned above, we cannot just replace $\boldsymbol{M}_{[F, F]}^{(i)}$ by its diagonal. Instead, LastStep performs one step of squaring similar to SchurSquare with a key difference: Rather than expressing $\boldsymbol{M}_{[F, F]}^{(i)}$ as $\boldsymbol{D}-\boldsymbol{A}$, it expresses it as $\boldsymbol{X}+\boldsymbol{L}$, where $\boldsymbol{X}$ is block-diagonal, and $\boldsymbol{L}$ is just barely bDD. With this splitting, it constructs $\boldsymbol{M}^{(l a s t)}$ after performing one iteration similar to Eq. (13.3). After this step, it replaces the $\boldsymbol{M}_{[F, F]}^{(\text {last })}$ block with the block-diagonal matrix $\boldsymbol{X}$. Again, we directly produce sparse approximations to $\boldsymbol{M}^{(\text {last })}$ and its Schur complement via weighted (bipartite) expanders. A precise description and proofs are given in Section 13.11.2.

## ■ 13.3.5 Sparsifying bDD Matrices

The main technical hurdle left to address is how we sparsify bDD matrices. We to do this both to approximate $\boldsymbol{M}$ by $\boldsymbol{M}^{(0)}$, if $\boldsymbol{M}$ is not already sparse, and to ensure that all the matrices $\boldsymbol{M}^{(i)}$ remain sparse. While the spectral vertex sparsification algorithm described in the previous section allows us
to compute an approximation to a Schur complement $S c\left(M^{(i)}, F_{i}\right)$, it is sparse only when $M^{(i)}$ is already sparse. As we iteratively apply this procedure, the density of the matrices produced will grow unacceptably. We overcome this problem by occasionally applying another sparsification routine that substantially decreases the number of nonzero blocks. The cost of this sparsification routine is that it requires solving systems of equations in sparse $b D D$ matrices. We, of course, do this recursively.

Our sparsification procedure begins by generalizing the observation that graph Laplacians can be sparsified by sampling edges with probabilities determined by their effective resistances [242]. There is a block analog of leverage scores (13.36) that provides probabilities of sampling blocks so that the resulting sampled matrix approximates the original and has $O(n \log n)$ nonzero blocks with high probability. To compute this block analog of leverage scores we employ the procedure developed in Chapter 3.

Once we generalize their results to block matrices, they show that we can obtain sufficiently good estimates of the block leverage scores by computing leverage scores in a bDD matrix obtained by randomly subsampling blocks of the original. The block leverage scores in this matrix are obtained by solving a logarithmic number of linear equations in this subsampled matrix. Thus, our sparsification procedure requires constructing a solver for a subsampled matrix and then applying that solver a logarithmic number of times. We compute this solver recursively.

There is a tradeoff between the number of nonzero blocks in the subsampled system and in the resulting approximation of the original matrix. If the original matrix has $m$ nonzero blocks and we subsample to a system of $m / K$ nonzero blocks, then we obtain an $\varepsilon$-approximation of the original matrix with $O\left(K \varepsilon^{-2} n \log n\right)$ nonzero blocks.

The details of the analysis of the undersampling procedure appear in Section 13.12.

## ■ 13.3.6 Our Algorithm

We now explain how we prove Theorem 13.1.2. The details supporting this exposition appear in Section 13.13. Our main goal is to control the density of the Schur complement chain as we repeatedly invoke Lemma 13.3.4.

Starting from some $\boldsymbol{M}^{(0)}$, we compute sets $F_{i}$ (via calls to BDDSUBSET), approximate solvers $\boldsymbol{Z}^{(i)}$ (via JACOBI), and approximations of Schur complements $\boldsymbol{M}^{(i)}$ (via ApproxSchur), until we obtain a matrix $\boldsymbol{M}^{(i)}$ such that its dimension is a smaller than that of $\boldsymbol{M}^{(0)}$ by a large constant factor (like 4). While the dimension of $\boldsymbol{M}^{(i)}$ is much smaller, its number of nonzero blocks is potentially larger by an even larger factor. We use the procedure described in the previous section to sparsify it. This sparsification procedure produces a sparse approximation of the matrix at the cost of solving systems of equations in a subsampled version of that matrix. Some care is required to balance the cost of the resulting recursion.

We now sketch an analysis of a nearly linear time algorithm that results from a simple choice of parameters. We optimize the parameter choice and analysis in Section 13.13. Let $n$ be the dimension of $\boldsymbol{M}^{(0)}$ and let $m$ be its number of nonzero blocks. To begin, assume that $m \leq n \Delta \log ^{3} n$, for a $\Delta$ to be specified later. We call this the sparse case, and address the case of dense $M$ later.

We consider fixing $\varepsilon_{i}=c / \log n$ for all $i$, for some constant $c$. As the depth of the Schur complement chain is $O(\log n)$, this results in a solver with constant accuracy. A constant number of iterations of the procedure described above are required to produce an $\boldsymbol{M}^{(i)}$ whose dimension is a factor of 4 smaller than $\boldsymbol{M}^{(0)}$. Lemma 13.3 .4 tells us that the edge density of this $\boldsymbol{M}^{(i)}$ is potentially higher than that of $\boldsymbol{M}^{(0)}$ by a factor of

$$
O\left(\left(\varepsilon^{-1} \log \log \varepsilon^{-1}\right)^{O\left(\log \log \varepsilon^{-1}\right)}\right)=\exp \left(O\left(\log \log ^{2} n\right)\right)=n^{o(1)} .
$$

Set $\Delta$ to be this factor. We use the sparsification procedure from the previous section to guarantee that no matrix in the chain has density higher than that of this matrix, which is upper bounded by $(m / n) \Delta=\Delta^{2} \log ^{3} n$.

Setting $K=2 \Delta$, the subsampling produces a matrix of density half that of $M^{(0)}$, and it produces a sparse approximation of $M^{(i)}$ of density $O\left(K \varepsilon_{i}^{-2} \log n\right) \leq O\left(\Delta \log ^{3} n\right)$, which, by setting constants appropriately, we can also force to be half that of $\boldsymbol{M}^{(0)}$. In order to perform the sparsification procedure, we need to construct a Schur complement chain for the subsampled matrix, and then use it to solve $O(\log n)$ systems of linear equations. The cost of using this chain to solve equations in the subsampled system is at most $O\left(n \Delta^{2} \log ^{4} n\right)$, and the cost of using the solutions to these equations to sparsify $M^{(i)}$ is $O(m \log n)$. The cost of the calls to BDDSUbSET and ApproxSchur are proportional to the number of edges in the matrices, which is $O\left(n \Delta \log ^{4} n\right)$.

We repeat this procedure all the way down the chain, only using sparsification when the dimension of $\boldsymbol{M}^{(i)}$ shrinks by a factor of 4 . Since none of the matrices that we generate have density higher than $\Delta \log ^{3} n$, we remain in the sparse case. Let $T_{\text {sparse }}(n)$ be the time required to construct a solver chain on systems of size $n$ with $m \leq n \Delta \log ^{3} n$. We obtain the following recurrence
$T_{\text {sparse }}(n) \leq 2 T_{\text {sparse }}(n / 4)+n \Delta^{2} \log ^{4} n+m \log n+m \Delta \leq 2 T_{\text {sparse }}(n / 4)+n \Delta^{2} \log ^{4} n+n \Delta \log n+n \Delta^{2}$, which gives

$$
T(n)_{\text {sparse }} \leq O\left(n \Delta^{2}+n \Delta^{2} \log ^{4} n\right) \leq n^{1+o(1)} .
$$

To handle the case of dense $M$, we repeatedly sparsify while keeping $n$ fixed until we obtain a matrix with fewer than $n \Delta \log ^{3} n$ edges, at which point we switch to the algorithm described above. The running time of this algorithm on a graph with $m$ edges, $T_{\text {dense }}(m)$, satisfies the recurrence

$$
T_{\text {dense }}(m) \leq \begin{cases}T_{\text {sparse }}(n) & \text { if } m \leq n \Delta \log ^{3} n, \text { and } \\ 2 T_{\text {dense }}(m / 2)+n \Delta^{2} \log ^{4} n+m \log n+m \Delta & \text { otherwise. }\end{cases}
$$

Thus $T_{\text {dense }}(m)$ is upper bounded bounded by $O\left(m n^{o(1)}+n^{1+o(1)}\right)$.
We tighten this bound in Section 13.13 to prove Theorem 13.1.2 by carefully choosing the parameters to accompany a sequence $\boldsymbol{\varepsilon}$ that starts constant and decreases slowly.

## ■ 13.4 Summary

We introduce a new approach to solving systems of linear equations that gives the first nearly linear time algorithms for solving systems proof that connection Laplacians have linear-sized approximate inverses. This was unknown even for graph Laplacians.

Our algorithms build on ideas introduced in [223] and are a break from those used in the previous work on solving systems of equations in graph Laplacians [256, 243, 147, 146, 141, 57]. Those algorithms all rest on support theory [34], originally introduced by Vaidya [256], and rely on the fact that the Laplacian of one edge is approximated by the Laplacian of a path between its endpoints. No analogous fact is true for connection Laplacians, even those with complex entries for $r=1$.

Instead, our algorithms rely on many new ideas, the first being that of sparsifying the matrices that appear during elimination. Other critical ideas are finding $\alpha$-bDD subsets of vertices to eliminate in bulk, approximating Schur complements without computing them explicitly, and the use of sub-sampling to sparsify in a recursive fashion. To efficiently compute approximations of the Schur complements, we introduce a new operation that transforms a matrix into one with the same Schur complement but a much better conditioned upper block (13.3). To obtain the sharp bounds in Theorem 13.1.2, we exploit a new linear-time algorithm for constructing linear-sized sparse approximations to
implicitly represented weighted cliques whose edge weights are products of weights at vertices (Section 13.14), and extend this to the analog for bDD matrices (Section 13.11.4).

## ■ 13.5 Background

Proof. (of Fact 13.2.2) Note that

$$
\|\boldsymbol{x}-\boldsymbol{Z} b\|_{M}^{2}=\boldsymbol{x}^{*} \boldsymbol{M} \boldsymbol{x}-2 \boldsymbol{x}^{*} \boldsymbol{M} \boldsymbol{Z} \boldsymbol{M} \boldsymbol{x}+\boldsymbol{x}^{*} \boldsymbol{M} \boldsymbol{Z} \boldsymbol{M} \boldsymbol{Z} \boldsymbol{M} \boldsymbol{x}
$$

Since all nonzero eigenvalues of $\boldsymbol{M}$ are at least $\mu$, the eigenvalues of $\boldsymbol{M}^{1 / 2}(\boldsymbol{M}+\varepsilon \mu \boldsymbol{I})^{-1} \boldsymbol{M}^{1 / 2}$ lie between $1 /(1+\varepsilon)$ and 1 . Using $\boldsymbol{Z} \approx_{\varepsilon}(\boldsymbol{M}+\varepsilon \mu \boldsymbol{I})^{-1}$, we see that the eigenvalues of $\boldsymbol{M}^{1 / 2} \boldsymbol{Z} \boldsymbol{M}^{1 / 2}$ lie between $e^{-2 \varepsilon}$ and $e^{\varepsilon}$. Using $0<\varepsilon<1 / 2$, we have

$$
\|\boldsymbol{x}-\boldsymbol{Z} \boldsymbol{b}\|_{\boldsymbol{M}}^{2} \leq\left(1-2 e^{-2 \varepsilon}+e^{2 \varepsilon}\right) \boldsymbol{x}^{*} \boldsymbol{M} \boldsymbol{x} \leq 6 \varepsilon\|\boldsymbol{x}\|_{\boldsymbol{M}}^{2}
$$

Fact 13.5.1. Let $\boldsymbol{A}$ be a matrix of condition number $\kappa$ and let $\boldsymbol{A} \approx_{\varepsilon} \boldsymbol{B}$ for $\varepsilon \leq\left(56 \kappa^{3}\right)^{-1}$. Then, $\boldsymbol{A}^{3} \approx_{28 \kappa^{3} \varepsilon} \boldsymbol{B}^{3}$.

Proof. First, observe that $\boldsymbol{A} \approx_{\varepsilon} \boldsymbol{B}$ implies that $\|\boldsymbol{B}\| \leq\left(1+e^{\varepsilon}\right)\|\boldsymbol{A}\| \leq 2\|\boldsymbol{A}\|$. It also implies that $\|\boldsymbol{A}-\boldsymbol{B}\| \leq 2 \varepsilon\|\boldsymbol{A}\|$. As

$$
\boldsymbol{A}^{2}-\boldsymbol{B}^{2}=\frac{1}{2}(\boldsymbol{A}-\boldsymbol{B})(\boldsymbol{A}+\boldsymbol{B})+\frac{1}{2}(\boldsymbol{A}+\boldsymbol{B})(\boldsymbol{A}-\boldsymbol{B})
$$

$\left\|\boldsymbol{A}^{2}-\boldsymbol{B}^{2}\right\| \leq 6 \varepsilon\|\boldsymbol{A}\|^{2}$. Similarly, as

$$
\boldsymbol{A}^{3}-\boldsymbol{B}^{3}=\frac{1}{2}(\boldsymbol{A}-\boldsymbol{B})\left(\boldsymbol{A}^{2}+\boldsymbol{B}^{2}\right)+\frac{1}{2}(\boldsymbol{A}+\boldsymbol{B})\left(\boldsymbol{A}^{2}-\boldsymbol{B}^{2}\right)
$$

$\left\|\boldsymbol{A}^{3}-\boldsymbol{B}^{3}\right\| \leq 28 \varepsilon\|\boldsymbol{A}\|^{3}$.
Let $\kappa=\|\boldsymbol{A}\| / \lambda_{\min }(\boldsymbol{A})$. The above relation implies that

$$
\left\|\boldsymbol{A}^{3}-\boldsymbol{B}^{3}\right\| \leq 28 \varepsilon \kappa^{3} \lambda_{\min }(\boldsymbol{A})^{3}=28 \varepsilon \kappa^{3} \lambda_{\min }\left(\boldsymbol{A}^{3}\right)
$$

Thus, as $28 \varepsilon \kappa^{3} \leq 1 / 2$,

$$
\boldsymbol{A}^{3} \approx_{56 \varepsilon \kappa^{3}} \boldsymbol{B}^{3}
$$

Proof. (of Fact 13.2.3)
By Fact 13.5.1, using $\delta=\frac{\varepsilon}{56 \kappa^{3}}$,

$$
\boldsymbol{M} \boldsymbol{Z}^{3} \boldsymbol{M} \approx_{\varepsilon} \boldsymbol{M}(\boldsymbol{M}+\varepsilon \mu I)^{-3} \boldsymbol{M}
$$

$\boldsymbol{M}$ has an eigendecomposition in the same basis as $(\boldsymbol{M}+\varepsilon \mu I)^{-3}$, and so it follows that $\boldsymbol{M}(\boldsymbol{M}+$ $\varepsilon \mu I)^{-3} \boldsymbol{M}$ has the same eigenbasis, and the same null space as $\boldsymbol{M}$.

When $\lambda^{-1}$ is the eigenvalue of $\boldsymbol{M}^{+}$of an eigenvector $\boldsymbol{v}$, the corresponding eigenvalue of $\boldsymbol{M}(\boldsymbol{M}+$ $\varepsilon \mu I)^{-3} \boldsymbol{M}$ is $\beta \stackrel{\text { def }}{=} \frac{\lambda^{2}}{(\lambda+\varepsilon \mu)^{3}}$ and

$$
\lambda^{-1}>\beta \geq \frac{\lambda^{2}}{(1+\varepsilon)^{3} \lambda^{3}}=e^{-3 \varepsilon} \lambda^{-1}
$$

So $\boldsymbol{M}(\boldsymbol{M}+\varepsilon \mu I)^{-3} \boldsymbol{M} \approx_{3 \varepsilon} \boldsymbol{M}^{+}$, and $\boldsymbol{M} \boldsymbol{Z}^{3} \boldsymbol{M} \approx_{4 \varepsilon} \boldsymbol{M}^{+}$.

Fact 13.5.2. For every $\boldsymbol{d} \in \mathbb{C}^{r \times r}$, we can find $\boldsymbol{Q}^{(1)}, \boldsymbol{Q}^{(2)} \in \mathbb{C}^{r \times r}$ where

$$
\left(\boldsymbol{Q}^{(1)}\right)\left(\boldsymbol{Q}^{(1)}\right)^{*}=\left(\boldsymbol{Q}^{(1)}\right)^{*}\left(\boldsymbol{Q}^{(1)}\right)=\left(\boldsymbol{Q}^{(2)}\right)\left(\boldsymbol{Q}^{(2)}\right)^{*}=\left(\boldsymbol{Q}^{(2)}\right)^{*}\left(\boldsymbol{Q}^{(2)}\right)=\boldsymbol{I}_{r},
$$

such that

$$
\boldsymbol{d}=\frac{1}{2}\|\boldsymbol{d}\|\left(\boldsymbol{Q}^{(1)}+\boldsymbol{Q}^{(2)}\right) .
$$

Proof. Let $\hat{\boldsymbol{d}}=\frac{1}{\|\boldsymbol{d}\|} \boldsymbol{d}$. Thus, $\|\hat{\boldsymbol{d}}\|=1$. Write $\hat{\boldsymbol{d}}$ using its singular value decomposition as $\boldsymbol{U} \boldsymbol{D} \boldsymbol{V}^{*}$, where $\boldsymbol{D}, \boldsymbol{U}, \boldsymbol{V} \in \mathbb{C}^{r \times r}, \boldsymbol{D}$ is a real diagonal matrix with the singular values of $\boldsymbol{d}$ on the diagonal, and $\boldsymbol{U} \boldsymbol{U}^{*}=\boldsymbol{U}^{*} \boldsymbol{U}=\boldsymbol{V} \boldsymbol{V}^{*}=\boldsymbol{V}^{*} \boldsymbol{V}=\boldsymbol{I}_{r}$. Since $\|\hat{\boldsymbol{d}}\|=1$, we have $\boldsymbol{D}_{j, j} \in[0,1]$ for all $j \in[r]$. Thus, there exists a real $\theta_{j}$ such that $\cos \theta_{j}=\boldsymbol{D}_{j, j}$. If we let $\boldsymbol{D}^{(1)}, \boldsymbol{D}^{(2)}$ be diagonal matrices usch that $\boldsymbol{D}_{j, j}^{(1)}=\exp \left(i \theta_{j}\right), \boldsymbol{D}_{j, j}^{(2)}=\exp \left(-i \theta_{j}\right)$, we have $\boldsymbol{D}=\frac{1}{2}\left(\boldsymbol{D}^{(1)}+\boldsymbol{D}^{(2)}\right)$. Moreover,

$$
\left(\boldsymbol{D}^{(1)}\right)\left(\boldsymbol{D}^{(1)}\right)^{*}=\left(\boldsymbol{D}^{(1)}\right)^{*}\left(\boldsymbol{D}^{(1)}\right)=\left(\boldsymbol{D}^{(2)}\right)\left(\boldsymbol{D}^{(2)}\right)^{*}=\left(\boldsymbol{D}^{(2)}\right)^{*}\left(\boldsymbol{D}^{(2)}\right)=\boldsymbol{I}_{r} .
$$

Letting $\boldsymbol{Q}^{(k)}=\boldsymbol{U} \boldsymbol{D}^{(k)} \boldsymbol{V}^{*}$ for $k=1,2$, we get $\hat{\boldsymbol{d}}=\frac{1}{2}\left(\boldsymbol{Q}^{(1)}+\boldsymbol{Q}^{(2)}\right)$ and hence $\boldsymbol{d}=\frac{1}{2}\|\boldsymbol{d}\|\left(\boldsymbol{Q}^{(1)}+\boldsymbol{Q}^{(2)}\right)$. Moreover,

$$
\left(\boldsymbol{Q}^{(1)}\right)\left(\boldsymbol{Q}^{(1)}\right)^{*}=\left(\boldsymbol{Q}^{(1)}\right)^{*}\left(\boldsymbol{Q}^{(1)}\right)=\left(\boldsymbol{Q}^{(2)}\right)\left(\boldsymbol{Q}^{(2)}\right)^{*}=\left(\boldsymbol{Q}^{(2)}\right)^{*}\left(\boldsymbol{Q}^{(2)}\right)=\boldsymbol{I}_{r} .
$$

Fact 13.5.3 (Lemma B.1. from [194]). If $\boldsymbol{M}$ and $\widetilde{M}$ are positive semidefinite matrices satisfying $M \preceq \widetilde{M}$, then

$$
S c(\boldsymbol{M}, F) \preceq S c(\widetilde{\boldsymbol{M}}, F) .
$$

This fact can be proven via an energy minimization definition of Schur complement. More details on this formulation can be found in [194].

## ■ 13.6 Block Diagonally Dominant Matrices

In this section, we prove a few basic facts about bDD matrices. The following lemma gives an equivalent definition of bDD matrices.

Lemma 13.6.1. A Hermitian block-matrix $\boldsymbol{M} \in\left(\mathbb{C}^{r \times r}\right)^{n \times n}$ is bDD iff it can be written as $\boldsymbol{D}-\boldsymbol{A}$ where $\boldsymbol{D}$ is block-diagonal, $\boldsymbol{A}$ is Hermitian, and $\boldsymbol{D}_{[i, i]} \succeq \boldsymbol{I}_{r} \sum_{j}\left\|\boldsymbol{A}_{[i, j]}\right\|$, for all $i \in V$.

Proof. The only if direction is easy. For bDD matrix $\boldsymbol{M}$, if we let $\boldsymbol{D}$ be the block diagonal matrix such that $\boldsymbol{D}_{[i, i]}=\boldsymbol{M}_{[i, i]}$, and $\boldsymbol{A}$ be the matrix such that

$$
\boldsymbol{A}_{[i, j]}= \begin{cases}0 & i=j \\ \boldsymbol{M}_{[i, j]} & i \neq j\end{cases}
$$

It is immediate that $\boldsymbol{A}$ is Hermitian. Moreover, $\boldsymbol{M}$ is bDD implies that $\boldsymbol{D}_{[i, i]} \succeq \boldsymbol{I}_{r} \sum_{j}\left\|\boldsymbol{A}_{[i, j]}\right\|$.
For the if direction. Suppose we have $\boldsymbol{D}, \boldsymbol{A}$ such that $\boldsymbol{D}$ is block-diagonal, and for all $i \boldsymbol{D}_{[i, i]} \succeq$ $\boldsymbol{I}_{r} \sum_{j}\left\|\boldsymbol{A}_{[i, j]}\right\|$. Thus, letting $\boldsymbol{M}=\boldsymbol{D}-\boldsymbol{A}$, we have, for all $i \in V$,

$$
\boldsymbol{I}_{r} \cdot \sum_{j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\|=\boldsymbol{I}_{r} \cdot \sum_{j \neq i}\left\|\boldsymbol{A}_{[i, j]}\right\| \preccurlyeq \boldsymbol{D}_{[i, i]}-\boldsymbol{I}_{r} \cdot\left\|\boldsymbol{A}_{[i, i]}\right\| \preccurlyeq \boldsymbol{D}_{[i, i]}-\boldsymbol{A}_{[i, i]}=\boldsymbol{M}_{[i, i]},
$$

where we used that since $\boldsymbol{A}, \boldsymbol{M}$ are Hermitian, $\boldsymbol{D}$ is also Hermitian, and thus $\boldsymbol{D}_{[i, i]}, \boldsymbol{A}_{[i, i]}$ are Hermitian.

This immediately implies the corollary that flipping the sign of off-diagonal blocks preserves bDDness.

Corollary 13.6.2. Given a bDD matrix $\boldsymbol{M}$, write it as $\boldsymbol{D}-\boldsymbol{A}$, where $\boldsymbol{D}$ is a block-diagonal, and $\boldsymbol{A}$ has its diagonal blocks as zero. Then, $\boldsymbol{D}+\boldsymbol{A}$ is also PSD.

Proof. First observe that for all $i,(\boldsymbol{D}+\boldsymbol{A})_{[i, i]}=(\boldsymbol{D}-\boldsymbol{A})_{[i, i]}=\boldsymbol{M}_{[i, i]}$, i.e., their diagonal blocks are identical. Moreover, for all $i \neq j$, we have $\left\|(\boldsymbol{D}+\boldsymbol{A})_{[i, j]}\right\|=\left\|(\boldsymbol{D}-\boldsymbol{A})_{[i, j]}\right\|=\left\|\boldsymbol{M}_{[i, j]}\right\|$. Thus $\boldsymbol{D}+\boldsymbol{A}$ is also bDD, and hence PSD.

Next, we show that the class of bDD matrices is closed under Schur complement.
Lemma 13.6.3. The class of $b D D$ matrices is closed under Schur complement.
Proof. Since Schur complementation does not depend on the order of indices eliminated, it suffices to prove that for any bDD matrix $\boldsymbol{M} \in\left(\mathbb{C}^{r \times r}\right)^{n \times n}, S c(\boldsymbol{M}, 1)$ is a bDD matrix. Let $C=V \backslash\{1\}$.

We have $S c(\boldsymbol{M}, 1)=\boldsymbol{M}_{[C, C]}-\boldsymbol{M}_{[C, 1]} \boldsymbol{M}_{[1,1]}^{-1} \boldsymbol{M}_{[1, C]}$. Let $\boldsymbol{D} \in\left(\mathbb{C}^{r \times r}\right)^{C \times C}$ be the block diagonal matrix such that $\boldsymbol{D}_{[i, i]}=\boldsymbol{M}_{[i, i]}$ for $i \in C$. Expressing $S c(\boldsymbol{M}, 1)$ as $\boldsymbol{D}+\left(\boldsymbol{M}_{[C, C]}-\boldsymbol{D}+\boldsymbol{M}_{[C, 1]} \boldsymbol{M}_{[1,1]}^{-1} \boldsymbol{M}_{[1, C]}\right)$, we have for any $i \in C$,

$$
\begin{aligned}
\sum_{j \in C} \|\left(\boldsymbol{M}_{[C, C]}\right. & \left.-\boldsymbol{D}+\boldsymbol{M}_{[C, 1]} \boldsymbol{M}_{[1,1]}^{-1} \boldsymbol{M}_{[1, C]}\right)_{[i, j]}\left\|\leq\left(\sum_{j \in C: j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\|\right)+\sum_{j \in C}\right\| \boldsymbol{M}_{[i, 1]} \boldsymbol{M}_{[1,1]}^{-1} \boldsymbol{M}_{[1, j]} \| \\
& \leq\left(\sum_{j \in C: j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\|\right)+\left\|\boldsymbol{M}_{[i, 1]}\right\|\left\|\boldsymbol{M}_{[1,1]}^{-1}\right\| \sum_{j \in C}\left\|\boldsymbol{M}_{[1, j]}\right\| \\
& \leq\left(\sum_{j \in C: j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\|\right)+\left\|\boldsymbol{M}_{[i, 1]}\right\|=\sum_{j \in V: j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\|,
\end{aligned}
$$

where the last inequality uses $\left\|\boldsymbol{M}_{[1,1]}^{-1}\right\|\left(\sum_{j \neq 1}\left\|\boldsymbol{M}_{[1, j]}\right\|\right) \leq 1$, since $\boldsymbol{M}_{[1,1]} \succcurlyeq \boldsymbol{I}_{r} \cdot \sum_{j \neq 1}\left\|\boldsymbol{M}_{[1, j]}\right\|$. Thus using Lemma 13.6.1, we have $S c(\boldsymbol{M}, 1)=\boldsymbol{D}-\left(-\left(\boldsymbol{M}_{[C, C]}-\boldsymbol{D}+\boldsymbol{M}_{[C, 1]} \boldsymbol{M}_{[1,1]}^{-1} \boldsymbol{M}_{[1, C]}\right)\right)$ is bDD.

The next definition describes a special form that we can express any bDD matrix in, which will occasionally be useful.

Definition 13.6.4. A matrix $\boldsymbol{B} \in\left(\mathbb{C}^{r \times r}\right)^{n \times m}$ is called a unitary edge-vertex transfer matrix, when each block column of $\boldsymbol{B}_{[e]}$ has exactly two nonzero blocks $\boldsymbol{U}_{e}, \boldsymbol{V}_{e} \in \mathbb{C}^{r \times r}$ s.t. $\boldsymbol{U}_{e} \boldsymbol{U}_{e}^{*}=\boldsymbol{V}_{e} \boldsymbol{V}_{e}^{*}=w_{e} \boldsymbol{I}_{r}$ where $w_{e} \geq 0$.

Lemma 13.6.5. Every $b D D$ matrix $\boldsymbol{M} \in\left(\mathbb{C}^{r \times r}\right)^{n \times n}$ with $m$ nonzero off-diagonal blocks can be written as $\boldsymbol{X}+\boldsymbol{B} \boldsymbol{B}^{*}$ where $\boldsymbol{B} \in\left(\mathbb{C}^{r \times r}\right)^{n \times 2 m}$ is a unitary edge-vertex transfer matrix and $\boldsymbol{X}$ is a block diagonal PSD matrix. This implies that every bDD matrix is PSD. Furthermore, for every block diagonal matrix $\boldsymbol{Y}$ s.t. $\boldsymbol{M}-\boldsymbol{Y}$ is $b D D$, we have $\boldsymbol{X} \succcurlyeq \boldsymbol{Y}$. This decomposition can be found in $O(m)$ time and $O(\log n)$ depth.

Proof. Consider a pair $\{i, j\} \in V \times V$ such that $i \neq j$, and $\boldsymbol{M}_{[i, j]} \neq 0$. Using Fact 13.5.2, we can write such a $\boldsymbol{M}_{[i, j]}$ as $\frac{1}{2}\left\|\boldsymbol{M}_{[i, j]}\right\|\left(\boldsymbol{Q}_{\{i, j\}}^{(1)}+\boldsymbol{Q}_{\{i, j\}}^{(2)}\right)$, where $\boldsymbol{Q}_{\{i, j\}}^{(1)}\left(\boldsymbol{Q}_{\{i, j\}}^{(1)}\right)^{*}=\boldsymbol{Q}_{\{i, j\}}^{(2)}\left(\boldsymbol{Q}_{\{i, j\}}^{(2)}\right)^{*}=\boldsymbol{I}_{r}$. we construct two vectors vector $\boldsymbol{B}_{\{i, j\}}^{(1)}, \boldsymbol{B}_{\{i, j\}}^{(2)} \in\left(\mathbb{C}^{r \times r}\right)^{n}$ such that for $k=1,2,\left(\boldsymbol{B}_{\{i, j\}}^{(k)}\right)_{[i]}=\frac{1}{\sqrt{2}}\left\|\boldsymbol{M}_{[i, j]}\right\|^{1 / 2} \boldsymbol{I}_{r}$,
$\left(\boldsymbol{B}_{\{i, j\}}^{(k)}\right)_{[j]}=\frac{1}{\sqrt{2}}\left\|\boldsymbol{M}_{[i, j]}\right\|^{1 / 2}\left(\boldsymbol{Q}_{\{i, j\}}^{(k)}\right)^{*}$, and all other blocks are zero. We can verify that for all $k, \ell \in V$,

$$
\left(\boldsymbol{B}_{\{i, j\}}^{(1)}\left(\boldsymbol{B}_{\{i, j\}}^{(1)}\right)^{*}+\boldsymbol{B}_{\{i, j\}}^{(2)}\left(\boldsymbol{B}_{\{i, j\}}^{(2)}\right)^{*}\right)_{[k, \ell]}= \begin{cases}\left\|\boldsymbol{M}_{[i, j]}\right\| \boldsymbol{I}_{r} & k=\ell=i, \\ \left\|\boldsymbol{M}_{[i, j]}\right\| \boldsymbol{I}_{r} & k=\ell=j, \\ \boldsymbol{M}_{[i, j]} & k=i, \ell=j, \\ \boldsymbol{M}_{[i, j]}^{*}=\boldsymbol{M}_{[j, i]} & k=j, \ell=i, \\ 0 & \text { otherwise }\end{cases}
$$

We let $\boldsymbol{X}=\boldsymbol{M}-\sum_{\{i, j\}: i \neq j}\left(\boldsymbol{B}_{\{i, j\}}^{(1)}\left(\boldsymbol{B}_{\{i, j\}}^{(1)}\right)^{*}+\boldsymbol{B}_{\{i, j\}}^{(2)}\left(\boldsymbol{B}_{\{i, j\}}^{(2)}\right)^{*}\right)$, which must be block-diagonal. We now show that for all $i \in V$, the block $\boldsymbol{X}_{[i, i]}$ is PSD. We have for all $i \in V$,

$$
\boldsymbol{X}_{[i, i]}=\boldsymbol{M}_{[i, i]}-\boldsymbol{I}_{r} \cdot \sum_{j: j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\| \succcurlyeq 0
$$

where the last inequality holds since $M$ is bDD.
Thus, if we define $\boldsymbol{B} \in\left(\mathbb{C}^{r \times r}\right)^{n \times 2 m}$ such that its columns are all the vectors $\boldsymbol{B}_{\{i, j\}}^{(1)}, \boldsymbol{B}_{\{i, j\}}^{(2)}$ defined above, we have $\boldsymbol{M}=\boldsymbol{X}+\boldsymbol{B} \boldsymbol{B}^{*}$, and every column of $\boldsymbol{B}$ has exactly 2 nonzero blocks.

To show that for every block diagonal $\boldsymbol{Y}$ s.t. $\boldsymbol{M}-\boldsymbol{Y}$ is $\mathrm{bDD}, \boldsymbol{X} \succcurlyeq \boldsymbol{Y}$, first consider applying the decomposition described above to $\boldsymbol{M}-\boldsymbol{Y}$ instead of $\boldsymbol{M}$. Since the construction of $\boldsymbol{B}$ only depends on the off-diagonal blocks, we get $\boldsymbol{M}-\boldsymbol{Y}=\boldsymbol{\Lambda}+\boldsymbol{B} \boldsymbol{B}^{*}$, where $\boldsymbol{\Lambda}$ is block diagonal and PSD. So, $\boldsymbol{X}-\boldsymbol{Y}=\boldsymbol{M}-\boldsymbol{B} \boldsymbol{B}^{*}-\boldsymbol{Y}=\boldsymbol{\Lambda} \succcurlyeq 0$.

It is immediate that the decomposition can be found in $O(m)$ time and $O(\log n)$ depth.

## ■13.7 Schur Complement Chains

In this section, we give a proof of Lemma 13.3.1. We restate the lemma here for convenience.
Lemma 13.3.1. Consider an $\boldsymbol{\varepsilon}$-SCC for $\boldsymbol{M}^{(0)}$ where $\boldsymbol{M}^{(i)}$ and $\boldsymbol{Z}^{(i)}$ can be applied to a vector in work $W_{\boldsymbol{M}^{(i)}}, W_{\boldsymbol{D}^{(i)}}$ and depth $D_{\boldsymbol{M}^{(i)}}, D_{\boldsymbol{Z}^{(i)}}$ respectively.

The algorithm ApplyChain $\left(\left(\boldsymbol{M}^{(1)}, \boldsymbol{Z}^{(1)}\right), \ldots,\left(\boldsymbol{M}^{(d)}, \boldsymbol{Z}^{(d)}\right) ; F_{1}, \ldots, F_{d} ; \boldsymbol{b}\right)$ corresponds to a linear operator $\boldsymbol{W}$ acting on $\boldsymbol{b}$ such that

1. $\boldsymbol{W}^{-1} \approx_{\sum_{i=1}^{d} 2 \varepsilon_{i}} M^{(0)}$, and
2. for any vector $\boldsymbol{b}$, ApplyChain $\left(\left(\boldsymbol{M}^{(1)}, \boldsymbol{Z}^{(1)}\right), \ldots,\left(\boldsymbol{M}^{(d)}, \boldsymbol{Z}^{(d)}\right) ; F_{1}, \ldots, F_{d} ; \boldsymbol{b}\right)$ runs in $O\left(\sum_{i=1}^{d}\left(D_{\boldsymbol{M}^{(i)}}+D_{\boldsymbol{Z}^{(i)}}\right)\right)$ depth and $O\left(\sum_{i=1}^{d}\left(W_{\boldsymbol{M}^{(i)}}+W_{\boldsymbol{Z}^{(i)}}\right)\right)$ work.
The pseudocode for procedure ApplyChain that uses an $\varepsilon$-vertex sparsifier chain to approximately solve a system of equations in $\boldsymbol{M}^{(0)}$ is given in Figure 13-1.

Proof. We begin by observing that the output vector $\boldsymbol{x}^{(0)}$ is a linear transformation of the input vector $\boldsymbol{b}^{(0)}$. Let $\boldsymbol{W}^{(0)}$ be the matrix that realizes this transformation. Similarly, for $1 \leq i \leq d$, define $\boldsymbol{W}^{(i)}$ to be the matrix so that

$$
\boldsymbol{x}^{(i)}=\boldsymbol{W}^{(i)} \boldsymbol{b}^{(i)}
$$

An examination of the algorithm reveals that

$$
\begin{equation*}
\boldsymbol{W}^{(d)}=\left(\boldsymbol{M}^{(d)}\right)^{-1} \tag{13.4}
\end{equation*}
$$

```
    Initialize \(\boldsymbol{b}^{(0)} \leftarrow \boldsymbol{b}\).
    for \(i=1, \ldots, d\) do
        \(\boldsymbol{x}_{\left[F_{i}\right]}^{(i-1)} \leftarrow \boldsymbol{Z}^{(i)} \boldsymbol{b}_{\left[F_{i}\right]}^{(i-1)}\),
        \(\boldsymbol{b}^{(i)} \leftarrow \boldsymbol{b}_{\left[C_{i}\right]}^{(i-1)}-\boldsymbol{M}_{\left[C_{i}, F_{i}\right]}^{(i-1)} \boldsymbol{x}_{\left[F_{i}\right]}^{(i-1)}\).
    end
    \(\boldsymbol{x}^{(d)} \leftarrow\left(\boldsymbol{M}^{(d)}\right)^{-1} \boldsymbol{b}^{(d)}\).
    for \(i=d, \ldots, 1\) do
        \(\boldsymbol{x}_{\left[C_{i}\right]}^{(i-1)} \leftarrow \boldsymbol{x}^{(i)}\).
        \(\boldsymbol{x}_{\left[F_{i}\right]}^{(i-1)} \leftarrow \boldsymbol{x}_{\left[F_{i}\right]}^{(i-1)}-\boldsymbol{Z}^{(i)} \boldsymbol{M}_{\left[F_{i}, C_{i}\right]}^{(i-1)} \boldsymbol{x}^{(i)}\).
    end
```

Algorithm 35: $\boldsymbol{x}^{(0)}=\operatorname{ApplyChain}\left(\left(\boldsymbol{M}^{(1)}, \boldsymbol{Z}^{(1)}\right), \ldots,\left(\boldsymbol{M}^{(d)}, \boldsymbol{Z}^{(d)}\right) ; F_{1}, \ldots, F_{d}\right)$

Figure 13-1: Solver Algorithm using Vertex Sparsifier Chain
and for $1 \leq i \leq d$,

$$
\boldsymbol{W}^{(i-1)}=\left[\begin{array}{cc}
\boldsymbol{I} & -\boldsymbol{Z}^{(i)} \boldsymbol{M}_{\left[F_{i}, C_{i}\right]}^{(i-1)}  \tag{13.5}\\
0 & \boldsymbol{I}
\end{array}\right]\left[\begin{array}{cc}
\boldsymbol{Z}^{(i)} & 0 \\
0 & \boldsymbol{W}^{(i)}
\end{array}\right]\left[\begin{array}{cc}
\boldsymbol{I} & 0 \\
-\boldsymbol{M}_{\left[C_{i}, F_{i}\right]}^{(i-1)} \boldsymbol{Z}^{(i)} & \boldsymbol{I}
\end{array}\right] .
$$

We will now prove by backwards induction on $i$ that

$$
\left(\boldsymbol{W}^{(i)}\right)^{-1} \approx_{\sum_{j=i+1}^{d} 2 \varepsilon_{j}} \boldsymbol{M}^{(i)} .
$$

The base case of $i=d$ follows from (13.4). Using the definition of an $\varepsilon$-SCC, we know that $0 \preceq$ $\left(\boldsymbol{Z}^{(i)}\right)^{-1}-\boldsymbol{M}_{\left[F_{i}, F_{i}\right]}^{(i-1)} \preceq \varepsilon_{i} \cdot S c\left(\boldsymbol{M}^{(i-1)}, C_{i}\right)$. We show in Lemma 13.7.1 that this implies

$$
\begin{aligned}
& \quad\left[\begin{array}{cc}
\boldsymbol{I} & -\boldsymbol{Z}^{(i)} \boldsymbol{M}_{\left[F_{i}, C_{i}\right]}^{(i-1)} \\
0 & \boldsymbol{I}
\end{array}\right]\left[\begin{array}{cc}
\boldsymbol{Z}^{(i)} & 0 \\
0 & S c\left(\boldsymbol{M}^{(i-1)}, F_{i}\right)^{-1}
\end{array}\right]\left[\begin{array}{cc}
\boldsymbol{I} & 0 \\
-\boldsymbol{M}_{\left[C_{i}, F_{i}\right]}^{(i-1)} \boldsymbol{Z}^{(i)} & \boldsymbol{I}
\end{array}\right] \approx_{\varepsilon_{i}}\left(\boldsymbol{M}^{(i-1)}\right)^{-1} . \\
& \text { As } \boldsymbol{M}^{(i)} \\
& \approx_{\varepsilon_{i}} S c\left(\boldsymbol{M}^{(i-1)}, F_{i}\right), \\
& \\
& \\
& \quad\left[\begin{array}{cc}
\boldsymbol{I} & -\boldsymbol{Z}^{(i)} \boldsymbol{M}_{\left[F_{i}, C_{i}\right]}^{(i-1)} \\
0 & \boldsymbol{I}
\end{array}\right]\left[\begin{array}{cc}
\boldsymbol{Z}^{(i)} & 0 \\
0 & \left(\boldsymbol{M}^{(i)}\right)^{-1}
\end{array}\right]\left[\begin{array}{cc}
\boldsymbol{I} & 0 \\
-\boldsymbol{M}_{\left[C_{i}, F_{i}\right]}^{(i-1)} \boldsymbol{Z}^{(i)} & \boldsymbol{I}
\end{array}\right] \approx_{2 \varepsilon_{i}}\left(\boldsymbol{M}^{(i-1)}\right)^{-1} .
\end{aligned}
$$

By combining this identity with (13.5) and our inductive hypothesis, we obtain

$$
\boldsymbol{W}^{(i-1)} \approx_{\sum_{j=i}^{d} 2 \varepsilon_{j}}\left(\boldsymbol{M}^{(i-1)}\right)^{-1}
$$

Thus, by induction, we obtain

$$
\boldsymbol{W}^{(0)} \approx_{\sum_{j=1}^{d} 2 \varepsilon_{j}}\left(\boldsymbol{M}^{(0)}\right)^{-1}
$$

The whole algorithm involves a constant number of applications of $\boldsymbol{Z}^{(i)}, \boldsymbol{M}_{\left[F_{i}, C_{i}\right]}^{(i-1)}$, and $\boldsymbol{M}_{\left[C_{i}, F_{i}\right]}^{(i-1)}$. We observe that in order to compute $\boldsymbol{M}_{\left[F_{i}, C_{i}\right]}^{(i-1)} x_{C_{i}}$, using a multiplication procedure for $\boldsymbol{M}^{(i-1)}$, we can
pad $x_{C_{i}}$ with zeros, multiply by $\boldsymbol{M}^{(i-1)}$, and read off the answer on the indices in $C_{i}$. Similarly, we can multiply vectors with $\boldsymbol{M}_{\left[C_{i}, F_{i}\right]}^{(i-1)}$. This immediately gives the claimed bounds on work and depth.

We now prove the deferred claims from the above proof.
Lemma 13.7.1. Let $M \in\left(\mathbb{C}^{r \times r}\right)^{n \times n}$ be a bDD matrix, $F \subseteq V$ be a subset of the indices, and $\boldsymbol{Z} \in\left(\mathbb{C}^{r \times r}\right)^{|F| \times|F|}$ be an hermitian operator satisfying $0 \preceq \boldsymbol{Z}^{-1}-\boldsymbol{M}_{[F, F]} \preceq \varepsilon \cdot S c(\boldsymbol{M}, C)$. Then,

$$
\left[\begin{array}{cc}
\boldsymbol{I} & -\boldsymbol{Z} \boldsymbol{M}_{[F, C]} \\
0 & \boldsymbol{I}
\end{array}\right]\left[\begin{array}{cc}
\boldsymbol{Z} & 0 \\
0 & S c(\boldsymbol{M}, F)^{-1}
\end{array}\right]\left[\begin{array}{cc}
\boldsymbol{I} & 0 \\
-\boldsymbol{M}_{[C, F]} \boldsymbol{Z} & \boldsymbol{I}
\end{array}\right] \approx_{\varepsilon} \boldsymbol{M}^{-1}
$$

Proof. Define

$$
\widehat{\boldsymbol{M}}=\left[\begin{array}{cc}
(\boldsymbol{Z})^{-1} & \boldsymbol{M}_{[F, C]} \\
\boldsymbol{M}_{[C, F]} & \boldsymbol{M}_{[C, C]}
\end{array}\right]
$$

Using Lemma 13.7.2, we know that the assumption on $\boldsymbol{Z}$ is equivalent to

$$
\boldsymbol{M} \preccurlyeq \widehat{\boldsymbol{M}} \preccurlyeq(1+\varepsilon) \boldsymbol{M} .
$$

By Eq. (13.1), this implies

$$
\boldsymbol{M}^{-1} \succcurlyeq\left[\begin{array}{cc}
\boldsymbol{I} & -\boldsymbol{Z} \boldsymbol{M}_{[F, C]} \\
0 & \boldsymbol{I}
\end{array}\right]\left[\begin{array}{cc}
\boldsymbol{Z} & 0 \\
0 & S c(\widehat{\boldsymbol{M}}, F)^{-1}
\end{array}\right]\left[\begin{array}{cc}
\boldsymbol{I} & 0 \\
-\boldsymbol{M}_{[C, F]} & \boldsymbol{I}
\end{array}\right] \succcurlyeq(1+\varepsilon)^{-1} \boldsymbol{M}^{-1}
$$

From Facts 13.5.3, we know that

$$
S c(\boldsymbol{M}, F)^{-1} \succcurlyeq S c(\widehat{\boldsymbol{M}}, F)^{-1} \succcurlyeq(1+\varepsilon)^{-1} S c(\boldsymbol{M}, F)^{-1} .
$$

Now, we substitute this inequality into the one above and obtain

$$
(1+\varepsilon) \boldsymbol{M}^{-1} \succcurlyeq\left[\begin{array}{cc}
\boldsymbol{I} & -\boldsymbol{Z} \boldsymbol{M}_{[F, C]} \\
0 & \boldsymbol{I}
\end{array}\right]\left[\begin{array}{cc}
\boldsymbol{Z} & 0 \\
0 & S c(\boldsymbol{M}, F)^{-1}
\end{array}\right]\left[\begin{array}{cc}
\boldsymbol{I} & 0 \\
-\boldsymbol{M}_{[C, F]} \boldsymbol{Z} & \boldsymbol{I}
\end{array}\right] \succcurlyeq(1+\varepsilon)^{-1} \boldsymbol{M}^{-1}
$$

which implies the lemma.
Lemma 13.7.2. Given a bDD matrix $\boldsymbol{M}$, a partition of its indices $(F, C)$ such that $\boldsymbol{M}_{[F, F]}, \boldsymbol{M}_{[C, C]}$ are invertible, and an invertible hermitian operator $\boldsymbol{Z} \in\left(\mathbb{C}^{r \times r}\right)^{|F| \times|F|}$, the following two conditions are equivalent:

1. $0 \preceq(\boldsymbol{Z})^{-1}-\boldsymbol{M}_{[F, F]} \preceq \varepsilon \cdot S c(\boldsymbol{M}, C)$.
2. 

$$
\boldsymbol{M} \preccurlyeq\left[\begin{array}{cc}
(\boldsymbol{Z})^{-1} & \boldsymbol{M}_{[F, C]} \\
\boldsymbol{M}_{[C, F]} & \boldsymbol{M}_{[C, C]}
\end{array}\right] \preccurlyeq(1+\varepsilon) \boldsymbol{M} .
$$

Proof. Writing

$$
M=\left[\begin{array}{cc}
\boldsymbol{M}_{[F, F]} & \boldsymbol{M}_{[F, C]} \\
\boldsymbol{M}_{[C, F]} & \boldsymbol{M}_{[C, C]}
\end{array}\right]
$$

condition 2 is equivalent to

$$
0 \preccurlyeq\left[\begin{array}{cc}
(\boldsymbol{Z})^{-1}-\boldsymbol{M}_{[F, F]} & 0 \\
0 & 0
\end{array}\right] \preccurlyeq \varepsilon \boldsymbol{M} .
$$

The left inequality in this statement is equivalent to the left inequality in condition 1. Thus, it suffices to prove the right sides are equivalent.

To this end, using $\boldsymbol{M}_{[F, C]}=\boldsymbol{M}_{[C, F]}^{*}$, it suffices to prove that $\forall x \in\left(\mathbb{C}^{r}\right)^{|F|}, y \in\left(\mathbb{C}^{r}\right)^{|C|}$,

$$
x^{*}\left((\boldsymbol{Z})^{-1}-\boldsymbol{M}_{[F, F]}\right) x \leq \varepsilon\left(x^{*} \boldsymbol{M}_{[F, F]} x+2 x^{*} \boldsymbol{M}_{[F, C]} y+y^{*} \boldsymbol{M}_{[C, C]} y\right)
$$

This is equivalent to proving $\forall x \in\left(\mathbb{C}^{r}\right)^{|F|}$,

$$
x^{*}\left((\boldsymbol{Z})^{-1}-\boldsymbol{M}_{[F, F]}\right) x \leq \varepsilon \inf _{y \in\left(\mathbb{C}^{r}\right)^{|C|}}\left(x^{*} \boldsymbol{M}_{[F, F]} x+2 x^{*} \boldsymbol{M}_{[F, C]} y+y^{*} \boldsymbol{M}_{[C, C]} y\right)
$$

Since $\boldsymbol{M}_{[C, C]} \succcurlyeq 0$, the rhs is a convex function of $y$. The minimum is achieved at $y=-\boldsymbol{M}_{[C, C]}^{-1} \boldsymbol{M}_{[F, C]}^{*} x$, and we obtain the equivalent condition

$$
x^{*}\left((\boldsymbol{Z})^{-1}-\boldsymbol{M}_{[F, F]}\right) x \leq \varepsilon x^{*}\left(\boldsymbol{M}_{[F, F]}-\boldsymbol{M}_{[F, C]} \boldsymbol{M}_{[C, C]}^{-1} \boldsymbol{M}_{[F, C]}^{*}\right) x
$$

Since $S c(\boldsymbol{M}, C)=\boldsymbol{M}_{[F, F]}-\boldsymbol{M}_{[F, C]} \boldsymbol{M}_{[C, C]}^{-1} \boldsymbol{M}_{[C, F]}=\boldsymbol{M}_{[F, F]}-\boldsymbol{M}_{[F, C]} \boldsymbol{M}_{[C, C]}^{-1} \boldsymbol{M}_{[F, C]}^{*}$, we obtain our claim.

## ■ 13.8 Finding $\alpha$-bDD Subsets

In this section we check that a simple randomized sampling procedure leads to $\alpha$-bDD subsets. Specifically we will prove Lemma 13.3.2:
Lemma 13.3.2. Given a bDD matrix $M$ with $n$ block-rows, and an $\alpha \geq 0$, BDDSUBSET computes a subset $F$ of size least $n /(8(1+\alpha))$ such that $\boldsymbol{M}_{[F, F]}$ is $\alpha$-bDD. It runs in runs in $O(m)$ expected work and $O(\log n)$ expected depth, where $m$ is the number of nonzero blocks in $\boldsymbol{M}$.

Pseudocode for this routine is given in Figure 13-2.

```
Algorithm 36: \(F=\operatorname{BDDSubSEt}(\boldsymbol{M}, \alpha)\), where \(\boldsymbol{M}\) is a bDD matrix with \(n\) rows.
    Let \(F^{\prime}\) be a uniform random subset of \(\{1, \ldots, n\}\) of size \(\frac{n}{4(1+\alpha)}\).
    Set
\[
F=\left\{i \in F^{\prime} \text { such that } \sum_{j: j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\| \geq(1+\alpha) \sum_{j \in F^{\prime}: j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\|\right\} .
\]
If \(|F|<\frac{n}{8(1+\alpha)}\), goto Step 36.
```


## Output $F$.

Figure 13-2: Routine for finding an $\alpha$-strongly block diagonally dominant submatrix
We first show that the set returned is guaranteed to be $\alpha$-strongly block diagonally dominant.
Lemma 13.8.1. If BSDDSUBSET terminates, it returns $F$ such that $\boldsymbol{M}_{[F, F]}$ is $\alpha-b D D$.
Proof. Consider some $i \in F$, the criteria for including $i$ in $F$ in Step 36 gives:

$$
\sum_{j: j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\| \geq(1+\alpha) \sum_{j \in F^{\prime}, j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\| \geq(1+\alpha) \sum_{j \in F, j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\|,
$$

where the last inequality follows since $F$ is a subset of $F^{\prime}$.
Incorporating this into the definition of $\boldsymbol{M}$ being bDD gives

$$
\boldsymbol{M}_{[i, i]} \succcurlyeq \boldsymbol{I}_{r} \cdot \sum_{j: j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\| \geq(1+\alpha) \boldsymbol{I}_{r} \cdot \sum_{j \in F, j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\| .
$$

which means $\boldsymbol{M}_{[F, F]}$ is $\alpha$-bDD.

It remains to show that the algorithm finds a big $F$ quickly. This can be done by upper bounding the expected size of $F$, or the probability of a single index being in $F$.

Lemma 13.8.2. For any $i$, we have

$$
\operatorname{Pr}\left[i \in F^{\prime} \text { and } i \notin F\right] \leq \frac{1}{16(1+\alpha)}
$$

Proof. This event only happens if $i \in F^{\prime}$ and

$$
\begin{equation*}
\sum_{j \in F^{\prime}, j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\|>\frac{1}{1+\alpha} \sum_{j: j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\| . \tag{13.6}
\end{equation*}
$$

Conditioning on $i$ being selected initially, or $i \in F^{\prime}$, the probability that each other $j \neq i$ is in $F^{\prime}$ is

$$
\frac{1}{n-1}\left(\frac{n}{4(1+\alpha)}-1\right)<\frac{1}{4(\alpha+1)}
$$

which gives:

$$
\mathbb{E}\left[\sum_{j \in F^{\prime}, j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\| \mid i \in F^{\prime}\right]<\frac{1}{4(1+\alpha)} \sum_{j: j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\|
$$

Markov's inequality then gives:

$$
\operatorname{Pr}\left[\left.\sum_{j \in F^{\prime}, j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\|>\frac{1}{1+\alpha} \sum_{j: j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\| \right\rvert\, i \in F^{\prime}\right]<1 / 4
$$

and thus

$$
\operatorname{Pr}\left[i \in F^{\prime} \text { and } i \notin F\right]=\operatorname{Pr}\left[i \notin F \mid i \in F^{\prime}\right] \operatorname{Pr}\left[i \in F^{\prime}\right]<\frac{1}{4} \frac{1}{4(1+\alpha)}=\frac{1}{16(1+\alpha)}
$$

Combining these two bounds gives Lemma 13.3.2.

Proof. (of Lemma 13.3.2) Applying Linearity of Expectation to Lemma 13.8 .2 gives

$$
\mathbb{E}\left[\left|F^{\prime} \backslash F\right|\right] \leq \frac{n}{16(1+\alpha)}
$$

Markov's inequality then gives

$$
\operatorname{Pr}\left[\left|F^{\prime} \backslash F\right| \geq \frac{n}{8(1+\alpha)}\right]<1 / 2
$$

So, with probability at least $1 / 2,|F| \geq n /(8(1+\alpha))$, and the algorithm will pass the test in line 3 . Thus, the expected number of iterations made by the algorithm is at most 2 . The claimed bounds on the expected work and depth of the algorithm follow.

## ■ 13.9 Jacobi Iteration on $\alpha$-bDD Matrices

From an $\alpha$-bDD set $F$, we will construct an operator $\boldsymbol{Z}$ that approximates $\boldsymbol{M}_{[F, F]}^{-1}$ and that can be applied quickly. Specifically, we will show:

Theorem 13.3.1. Let $\boldsymbol{M}$ be a bDD matrix with index set $V$, and let $F \subseteq V$ such that $\boldsymbol{M}_{[F, F]}$ is $\alpha-b D D$ for some $\alpha \geq 4$, and has $m_{F F}$ nonzero blocks. The algorithm $\operatorname{JACOBI}\left(\varepsilon, \boldsymbol{M}_{[F, F]}, \boldsymbol{b}\right)$ acts as a linear operator $\boldsymbol{Z}$ on $\boldsymbol{b}$ that satisfies

$$
0 \preceq(\boldsymbol{Z})^{-1}-\boldsymbol{M}_{[F, F]} \preceq \varepsilon \cdot S c(\boldsymbol{M}, V \backslash F) .
$$

The algorithm takes $O\left(m_{F F} \log \left(\frac{1}{\varepsilon}\right)\right)$ work and $O\left(\log n \log \left(\frac{1}{\varepsilon}\right)\right)$ depth.
Pseudocode of this routine is given in Figure 13-3.

```
Algorithm 37: \(\boldsymbol{x}=\mathrm{JACOBI}(\varepsilon, \boldsymbol{M}, \boldsymbol{b})\)
    Create the matrix \(L\) where
\[
\boldsymbol{L}_{[i, j]}=\left\{\begin{array}{lc}
\boldsymbol{I}_{r} \cdot \sum_{j: j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\| & \text { if } i=j, \\
\boldsymbol{M}_{[i, j]} & \text { otherwise } .
\end{array}\right.
\]
```

Set $\boldsymbol{X}=\boldsymbol{M}-\boldsymbol{L}$.
Set $k$ to be an odd integer that is greater than $\log (3 / \varepsilon)$.
Set $\boldsymbol{x}^{(0)}=\boldsymbol{X}^{-1} \boldsymbol{b}$.
for $i=1 \cdots k$ do
Set $\boldsymbol{x}^{(i)}=-\boldsymbol{X}^{-1} \boldsymbol{L} \boldsymbol{x}^{(i-1)}+\boldsymbol{X}^{-1} \boldsymbol{b}$.
end
Output $\boldsymbol{x}^{(k)}$.
Figure 13-3: Jacobi Iteration for Solving Linear Systems in an $\alpha$-bDD Matrix
We first verify that any $\alpha$-bDD matrix has a good block-diagonal preconditioner.
Lemma 13.3.3. Every $\alpha$-bDD matrix $\boldsymbol{M}$ can be written in the form $\boldsymbol{X}+\boldsymbol{L}$ where $\boldsymbol{X}$ is block-diagonal, $\boldsymbol{L}$ is bDD, and $\boldsymbol{X} \succcurlyeq \frac{\alpha}{2} \boldsymbol{L}$.

Proof. Write $\boldsymbol{L}=\boldsymbol{Y}-\boldsymbol{A}$ where $\boldsymbol{Y}$ is block-diagonal and $\boldsymbol{A}$ has its diagonal blocks as zeros. Note that $\boldsymbol{L}$ is bDD by definition. Thus, by Corollary 13.6.5, $\boldsymbol{L}=\boldsymbol{Y}-\boldsymbol{A} \succcurlyeq 0$.

Using Lemma 13.6.2, we know that $\boldsymbol{Y}+\boldsymbol{A} \succcurlyeq 0$, or $\boldsymbol{Y} \succcurlyeq-\boldsymbol{A}$. This implies $2 \boldsymbol{Y} \succcurlyeq \boldsymbol{L}$.
As $\boldsymbol{M}$ is $\alpha$-strongly diagonally dominant and $\boldsymbol{M}_{[i, i]}=\boldsymbol{X}_{[i, i]}+\boldsymbol{Y}_{[i, i]}$, we have

$$
(\boldsymbol{X}+\boldsymbol{Y})_{[i, i]} \succcurlyeq(1+\alpha) \boldsymbol{I}_{r} \cdot \sum_{j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\|=(1+\alpha) \boldsymbol{Y}_{[i, i]} .
$$

Manipulating this then gives:

$$
\boldsymbol{X} \succcurlyeq \alpha \boldsymbol{Y} \succcurlyeq \frac{\alpha}{2} \boldsymbol{L} .
$$

We now move on to measuring the quality of the operator generated by Jacobi. It can be checked
that running it $k$ steps gives the operator

$$
\begin{equation*}
\boldsymbol{Z}^{(k)} \stackrel{\text { def }}{=} \sum_{i=0}^{k} \boldsymbol{X}_{[F, F]}^{-1}\left(-\boldsymbol{L}_{[F, F]} \boldsymbol{X}_{[F, F]}^{-1}\right)^{i} \tag{13.7}
\end{equation*}
$$

which is equivalent to evaluating a truncation of the Neumann series for $\boldsymbol{M}^{-1}$
Lemma 13.9.1. Let $\boldsymbol{M}$ be a matrix with splitting $\boldsymbol{M}=\boldsymbol{X}+\boldsymbol{L}$ where $0 \preccurlyeq \boldsymbol{L} \preccurlyeq \beta \boldsymbol{X}$ for some parameter $1>\beta>0$. Then, for odd $k$ and for $\boldsymbol{Z}^{(k)}$ as defined in (13.7) we have:

$$
\begin{equation*}
\boldsymbol{X}+\boldsymbol{L} \preceq\left(\boldsymbol{Z}^{(k)}\right)^{-1} \preceq \boldsymbol{X}+(1+\delta) \boldsymbol{L} \tag{13.8}
\end{equation*}
$$

where

$$
\delta=\beta^{k} \frac{1+\beta}{1-\beta^{k+1}}
$$

Proof. The left-hand inequality is equivalent to the statement that all the eigenvalues of $\boldsymbol{Z}^{(k)}(\boldsymbol{X}+\boldsymbol{L})$ are at most 1 (see [34, Lemma 2.2] or [243, Proposition 3.3]). To see that this is the case, expand

$$
\begin{aligned}
\boldsymbol{Z}^{(k)}(\boldsymbol{X}+\boldsymbol{L}) & =\left(\sum_{i=0}^{k} \boldsymbol{X}^{-1}\left(-\boldsymbol{L} \boldsymbol{X}^{-1}\right)^{i}\right)(\boldsymbol{X}+\boldsymbol{L}) \\
& =\sum_{i=0}^{k}\left(-\boldsymbol{X}^{-1} \boldsymbol{L}\right)^{i}-\sum_{i=1}^{k+1}\left(-\boldsymbol{X}^{-1} \boldsymbol{L}\right)^{i} \\
& =\boldsymbol{I}-\left(\boldsymbol{X}^{-1} \boldsymbol{L}\right)^{k+1}
\end{aligned}
$$

As all the eigenvalues of an even power of a matrix are nonnegative, all of the eigenvalues of this last matrix are at most 1 .

Similarly, the other inequality is equivalent to the assertion that all of the eigenvalues of $\boldsymbol{Z}^{(k)}(\boldsymbol{X}+$ $(1+\delta) \boldsymbol{L})$ are at least one. Expanding this product yields

$$
\left(\sum_{i=0}^{k} \boldsymbol{X}^{-1}\left(-\boldsymbol{L} \boldsymbol{X}^{-1}\right)^{i}\right)(\boldsymbol{X}+(1+\delta) \boldsymbol{L})=\boldsymbol{I}-\left(\boldsymbol{X}^{-1} \boldsymbol{L}\right)^{k+1}+\delta \sum_{i=0}^{k}(-1)^{i}\left(\boldsymbol{X}^{-1} \boldsymbol{L}\right)^{i+1}
$$

The eigenvalues of this matrix are precisely the numbers

$$
\begin{equation*}
1-\lambda^{k+1}+\delta \sum_{i=0}^{k}(-1)^{i} \lambda^{i+1} \tag{13.9}
\end{equation*}
$$

where $\lambda$ ranges over the eigenvalues of $\boldsymbol{X}^{-1} \boldsymbol{L}$. The assumption $\boldsymbol{L} \preccurlyeq \beta \boldsymbol{X}$ implies that the eigenvalues of $\boldsymbol{X}^{-1} \boldsymbol{L}$ are at most $\beta$, so $0 \leq \lambda \leq \beta$. We have chosen the value of $\delta$ precisely to guarantee that, under this condition on $\lambda$, the value of (13.9) is at least 1 .

This error crucially depends only on $L$, which for any choice of $F$ can be upper bounded by $\boldsymbol{M}$. Propagating the error this way allows us to prove the guarantees for JAcobi

Proof. (of Theorem 13.3.1) Consider the matrix $\boldsymbol{L}_{[F, F]}$ generated when calling JACOBI with $\boldsymbol{M}_{[F, F]}$. $\boldsymbol{M}$ being bDD means for each $i$ we have

$$
\boldsymbol{M}_{[i, i]} \succcurlyeq \boldsymbol{I}_{r} \cdot \sum_{j \neq i}\left\|\boldsymbol{M}_{[i, j]}\right\|
$$

which means for all $i \in F$

$$
\boldsymbol{M}_{[i, i]}-\boldsymbol{I}_{r} \cdot \sum_{j \neq i, j \in F}\left\|\boldsymbol{M}_{[i, j]}\right\| \succcurlyeq \boldsymbol{I}_{r} \cdot \sum_{j \notin F}\left\|\boldsymbol{M}_{[i, j]}\right\| .
$$

Therefore if we extend $\boldsymbol{L}_{[F, F]}$ onto the full matrix by putting zeros everywhere else, we have $\boldsymbol{L} \preccurlyeq \boldsymbol{M}$. Fact 13.5.3 then gives $\boldsymbol{L}_{[F, F]} \preccurlyeq S c(\boldsymbol{M}, V \backslash F)$.

Lemma 13.3.3 gives that $\boldsymbol{X}_{[F, F]} \succcurlyeq \frac{\alpha}{2} \boldsymbol{L}_{[F, F]}$. As $\alpha \geq 4$, we can invoke Lemma 13.9.1 with $\beta=\frac{1}{2}$. Since $\frac{1+\beta}{1-\beta^{k+1}} \leq\left(\frac{3}{2}\right) /\left(\frac{1}{2}\right) \leq 3$, our choice of $k=\log (3 / \varepsilon)$ gives the desired error. Each of these steps involve a matrix vector multiplication in $\boldsymbol{L}_{[F, F]}$ and two linear system solves in $\boldsymbol{X}_{[F, F]}$. The former takes $O(m)$ work and $O(\log n)$ depth since the blocks of $\boldsymbol{L}_{[F, F]}$ are a subset of the blocks of $\boldsymbol{M}$, while the latter takes $O(n)$ work and $O(\log n)$ depth due to $\boldsymbol{X}_{[F, F]}$ being block-diagonal.

## ■ 13.10 Existence of Linear-Sized Approximate Inverses

In this section we prove Theorem 13.1.3, which tells us that every $b D D$ matrix has a linear-sized approximate inverse. In particular, this implies that for every $b D D$ matrix there is a linear-time algorithm that approximately solves systems of equations in that matrix. To save space, we will not dwell on this algorithm, but rather will develop the linear-sized approximate inverses directly. There is some cost in doing so: there are very large constants in the linear-sized approximate inverses that are not present in a simpler linear-time solver.

To obtain a $\boldsymbol{U}^{T} \boldsymbol{D} \boldsymbol{U}$ factorization from an $\boldsymbol{\varepsilon}$-SCC in which each $\boldsymbol{M}^{(i)}{ }_{\left[F_{i} F_{i}\right]}$ is 4 -bDD, we employ the procedure in Figure 13-4.

```
Algorithm 38: \((\boldsymbol{D}, \boldsymbol{U})=\) Decompose \(\left(\boldsymbol{M}^{(1)}, \ldots, \boldsymbol{M}^{(d)}, F_{1}, \ldots, F_{d-1}\right)\), where each \(\boldsymbol{M}^{(i)}\) is a
bDD matrix and each \(M^{(i)}{ }_{\left[F_{i} F_{i}\right]}\) is 4-bDD.
    Use (13.7) to compute the matrix \(\boldsymbol{Z}^{(i)}\) such that \(\operatorname{JACOBI}\left(\varepsilon_{i}, \boldsymbol{M}, \boldsymbol{b}\right)=\boldsymbol{Z}^{(i)} \boldsymbol{b}\).
    For each \(i<d\), write \(\boldsymbol{M}^{(i)}=\boldsymbol{X}^{(i)}+\boldsymbol{L}^{(i)}\), where \(\boldsymbol{X}^{(i)}\) is block diagonal and \(\boldsymbol{L}^{(i)}\) is \(b D D\), as in
        Jacobi.
    Let \(\boldsymbol{X}^{(d)}=\boldsymbol{I}_{\left|C_{d-1}\right|}\) and let \(\widehat{\boldsymbol{U}}\) be the upper-triangular Cholesky factor of \(\boldsymbol{M}^{(d)}\).
    Let \(\boldsymbol{D}\) be the block diagonal matrix with \(\boldsymbol{D}_{\left[F_{i}, F_{i}\right]}=\boldsymbol{X}^{(i)}\), for \(1 \leq i<d\), and
        \(\boldsymbol{D}_{\left[C_{d-1}, C_{d-1}\right]}=\boldsymbol{I}_{\left|C_{d-1}\right|}\).
    Let \(\boldsymbol{U}\) be the upper-triangular matrix with \(1 s\) on the diagonal, \(\boldsymbol{U}_{\left[C_{d-1}, C_{d-1}\right]}=\widehat{\boldsymbol{U}}\), and
    \(\boldsymbol{U}_{\left[F_{i}, C_{i}\right]}=\boldsymbol{Z}^{(i)} \boldsymbol{M}^{(i)}{ }_{\left[F_{i}, C_{i}\right]}\), for \(1 \leq i<d\).
```

Figure 13-4: Converting a vertex sparsifier chain into $\boldsymbol{U}$ and $\boldsymbol{D}$.

Lemma 13.10.1. On input an $\varepsilon$-SCC of $\boldsymbol{M}^{(0)}$ in which each $\boldsymbol{M}^{(i)}{ }_{\left[F_{i} F_{i}\right]}$ is 4 -bDD, the algorithm Decompose produces matrices $\boldsymbol{D}$ and $\boldsymbol{U}$ such that

$$
\boldsymbol{U}^{T} \boldsymbol{D} \boldsymbol{U} \approx_{\gamma} \boldsymbol{M}
$$

where

$$
\gamma \leq 2 \sum_{i=0}^{d-1} \varepsilon_{i}+\max _{i} \varepsilon_{i}+1 / 2
$$

Proof. Consider the inverse of the operator $\boldsymbol{W}=\boldsymbol{W}^{(1)}$ realized by the algorithm ApplyChain, and the operators $\boldsymbol{W}^{(i)}$ that appear in the proof of Lemma 13.3.1.

We have

$$
\left(\boldsymbol{W}^{(i)}\right)^{-1}=\left[\begin{array}{cc}
\boldsymbol{I} & 0 \\
\boldsymbol{M}_{\left[C_{i}, F_{i}\right]} \boldsymbol{Z}^{(i)} & \boldsymbol{I}
\end{array}\right]\left[\begin{array}{cc}
\left(\boldsymbol{Z}^{(i)}\right)^{-1} & 0 \\
0 & \left(\boldsymbol{W}^{(i+1)}\right)^{-1}
\end{array}\right]\left[\begin{array}{cc}
\boldsymbol{I} & \boldsymbol{Z}^{(i)} \boldsymbol{M}_{\left[F_{i}, C_{i}\right]} \\
0 & \boldsymbol{I}
\end{array}\right]
$$

and

$$
\left(\boldsymbol{W}^{(d)}\right)^{-1}=\boldsymbol{M}^{(d)}=\widehat{\boldsymbol{U}}^{T} \widehat{\boldsymbol{U}}
$$

After expanding and multiplying the matrices in this recursive factorization, we obtain

$$
\left(\boldsymbol{W}^{(1)}\right)^{-1}=\boldsymbol{U}^{T}\left[\begin{array}{cccc}
\left(\boldsymbol{Z}^{(1)}\right)^{-1} & \cdots & 0 & 0 \\
0 & \ddots & 0 & 0 \\
0 & \cdots & \left(\boldsymbol{Z}^{(d-1)}\right)^{-1} & 0 \\
0 & \cdots & 0 & \boldsymbol{I}_{\left|C_{d-1}\right|}
\end{array}\right] \boldsymbol{U}
$$

Moreover, we know that this latter matrix is a $2 \sum_{i=0}^{d-1} \varepsilon_{i}$ approximation of $\boldsymbol{M}$. It remains to determine the impact of replacing the matrix in the middle of this expression with $\boldsymbol{D}$.

Lemma 13.9.1 implies that each $\boldsymbol{M}_{\left[F_{i} F_{i}\right]} \approx_{\varepsilon_{i}}\left(\boldsymbol{Z}^{(i)}\right)^{-1}$ and Lemma 13.3.3 implies that $\boldsymbol{X}^{(i)} \approx_{1 / 2}$ $\boldsymbol{M}_{\left[F_{i} F_{i}\right]}$. So, the loss in approximation quality when we substitute the diagonal matrices is $\max _{i} \varepsilon_{i}+$ $1 / 2$.

Invoking this decomposition procedure in conjunction with the the near-optimal sparsification routine from Theorem 13.3 .1 gives a nearly-linear work routine. Repeatedly picking subsets using Lemma 13.3.4 gives then gives the linear sized decomposition.
Theorem 13.1.3 (Sparsified Cholesky). For every bDD matrix $M$ with $n$ block-rows there exists a diagonal matrix $\boldsymbol{D}$ and an upper triangular matrix $\boldsymbol{U}$ with $O(n)$ nonzero blocks so that

$$
\boldsymbol{U}^{T} \boldsymbol{D} \boldsymbol{U} \approx_{3 / 4} \boldsymbol{M}
$$

Moreover, linear equations in $\boldsymbol{U}, \boldsymbol{U}^{T}$, and $\boldsymbol{D}$ can be solved with linear work in depth $O\left(\log ^{2} n\right)$, and these matrices can be computed in polynomial time.

Proof. We set $\alpha=4$ throughout and $\varepsilon_{i}=1 / 8(i+2)^{2}$. Theorem 13.3.1 then guarantees that the average number of nonzero blocks in each column of $\boldsymbol{M}^{(i)}$ is at most $10 r / \varepsilon_{i}^{2}=640 r(i+2)^{4}$. If we now apply Lemma 13.3 .4 to find 4 -diagonally dominant subsets $F_{i}$ of each $\boldsymbol{M}^{(i)}$, we find that each such subset contains at least a $1 / 80$ fraction of the block columns of its matrix and that each column and row of $\boldsymbol{M}^{(i)}$ indexed by $F$ has at most $1280 r(i+2)^{4}$ nonzero entries. This implies that each row of $\boldsymbol{Z}^{(i)} \boldsymbol{M}_{\left[F_{i}, C_{i}\right]}$ has at most $\left(1280 r(i+2)^{4}\right)^{k_{i}+1}$ nonzero entries.

Let $n_{i}$ denote the number of block columns of $\boldsymbol{M}^{(i)}$. By induction, we know that

$$
n_{i} \leq n\left(1-\frac{1}{80}\right)^{i-1}
$$

So, the total number of nonzero blocks in $\boldsymbol{U}$ is at most

$$
\sum_{i=1}^{d} n_{i}\left(1280 r(i+2)^{4}\right)^{k_{i}+1} \leq n \sum_{i=1}^{d}\left(1-\frac{1}{80}\right)^{i-1}\left(1280 r(i+2)^{4}\right)^{k_{i}+1}
$$

We will show that the term multiplying $n$ in this later expression is upper bounded by a constant. To see this, note that $k_{i} \leq 1+\log \left(2 \varepsilon_{i}^{-1}\right) \leq \nu \log (i+1)$ for some constant $\nu$. So, there is some other constant $\mu$ for which

$$
\left(1280 r(i+2)^{4}\right)^{k_{i}+1} \leq \exp \left(\mu \log ^{2}(i+1)\right)
$$

This implies that the sum is at most

$$
\sum_{i \geq 1} \exp \left(\mu \log ^{2}(i+1)-i / 80\right)
$$

which is bounded by a constant.
To bound the quality of the approximation, we compute

$$
2 \sum_{i} \varepsilon_{i}+\max _{i} \varepsilon_{i}+1 / 2=2 \sum_{i} 1 / 8(i+2)^{2}+1 / 72+1 / 2<3 / 4 .
$$

The claimed bound on the work to perform backwards and forwards substitution with $\boldsymbol{U}$ is standard: these operations require work linear in the number of nonzero entries of $\boldsymbol{U}$. The bound on the depth follows from the fact that the substitutions can be performed level-by-level, take depth $O(\log n)$ for each level, and the number of levels, $d$, is logarithmic in $n$.

## ■ 13.11 Spectral Vertex Sparsification Algorithm

In this section, we give a proof of the following lemma that immediately implies Lemma 13.3.4.
Lemma 13.11.1. Let $M$ be a bDD matrix with index set $V$, and $m$ nonzero blocks. Let $F \subseteq V$ be such that $\boldsymbol{M}_{[F, F]}$ is $\alpha-b D D$ for some $\alpha \geq 4$. The algorithm $\operatorname{ApproxSchur}(\boldsymbol{M}, F, \varepsilon)$, returns a matrix $\widetilde{\boldsymbol{M}}_{S C}$ s.t.

1. $\widetilde{\boldsymbol{M}}_{S C}$ has $O\left(m\left(\varepsilon^{-1} \log \log \varepsilon^{-1}\right)^{O\left(\log \log \varepsilon^{-1}\right)}\right)$ nonzero blocks, and
2. $\widetilde{\boldsymbol{M}}_{S C} \approx_{\varepsilon} S c(\boldsymbol{M}, F)$,
in $O\left(m\left(\varepsilon^{-1} \log \log \varepsilon^{-1}\right)^{O\left(\log \log \varepsilon^{-1}\right)}\right)$ work and $O\left(\log n\left(\log \log \varepsilon^{-1}\right)\right)$ depth.
Moreover, if $\boldsymbol{T}$ is a matrix such that only the submatrix $\boldsymbol{T}_{[C, C]}$ is nonzero, and $\boldsymbol{M}+\boldsymbol{T}$ is $b D D$, then $\operatorname{ApproxSchur}(\boldsymbol{M}+\boldsymbol{T}, F, \varepsilon)=\operatorname{ApproxSchur}(\boldsymbol{M}, F, \varepsilon)+\boldsymbol{T}_{[C, C]}$.

We show how to sparsify the Schur complement of $M$ after eliminating a set of indices $F$ such that $\boldsymbol{M}_{[F, F]}$ is an $\alpha$-bDD matrix. The procedure ApproxSchur is described in Figure 13-5 and uses two key subroutines SchurSquare and LastStep. SchurSquare allows us to approximate $S c(M, F)$ as the Schur complement of another matrix $\boldsymbol{M}_{1}$ such that $\boldsymbol{M}_{1[F, F]}$ is roughly $\alpha^{2}$-bDD. LaStStep allows us to approximate $S c(\boldsymbol{M}, F)$ with $O(\varepsilon)$ error, when $\boldsymbol{M}_{[F, F]}$ is roughly $1 / \varepsilon$-bDD.

Lemma 13.11.2. Let $M$ be an bDD matrix with index set $V$, and $m$ nonzero blocks and let $F \subseteq V$ be such that $\boldsymbol{M}_{[F, F]}$ is $\alpha-b D D$ for some $\alpha \geq 4$. Given $\varepsilon<1 / 2$, the algorithm $\operatorname{SchurSquare}(\boldsymbol{M}, F, \varepsilon)$, returns a bDD matrix $\boldsymbol{M}_{1}$ in $O\left(m \varepsilon^{-4}\right)$ work and $O(\log n)$ depth, such that

1. $S c(\boldsymbol{M}, F) \approx_{\varepsilon} S c\left(\boldsymbol{M}_{1}, F\right)$, and
2. $\left(\boldsymbol{M}_{1}\right)_{[F, F]}$ is $\alpha^{2} / 2-b D D$.
3. $\boldsymbol{M}_{1}$ has $O\left(m \varepsilon^{-4}\right)$ nonzero blocks,
```
Algorithm 39: \(\widetilde{\boldsymbol{M}}_{S C}=\operatorname{ApproxSchur}(\boldsymbol{M}, F, \alpha, \varepsilon)\)
    Initialize \(\boldsymbol{M}^{(0)} \leftarrow \boldsymbol{M}, d=2 \log _{2} \log _{\alpha / 2}\left(4 \varepsilon^{-1}\right)\)
    for \(i=1, \cdots, d\) do
        \(\boldsymbol{M}^{(i)} \leftarrow \operatorname{SchurSquare}\left(\boldsymbol{M}^{(i-1)}, F, \frac{\varepsilon}{2 d}\right)\).
    end
    \(\widetilde{\boldsymbol{M}}_{S C} \leftarrow \operatorname{LASTStEp}\left(\boldsymbol{M}^{(d)}, F, 4 \varepsilon^{-1}, \frac{\varepsilon}{4}\right)\).
    Output \(\widetilde{M}_{S C}\).
```

Figure 13-5: Pseudocode for Computing Spectral Vertex Sparsifiers

If $\boldsymbol{T}$ is a matrix such that only the submatrix $\boldsymbol{T}_{[C, C]}$ is nonzero, and $\boldsymbol{M}+\boldsymbol{T}$ is $b D D$, then

$$
\operatorname{SchurSquare}(\boldsymbol{M}+\boldsymbol{T}, F, \varepsilon)=\operatorname{SchurSquare}(\boldsymbol{M}, F, \varepsilon)+\boldsymbol{T}_{[C, C]} .
$$

We can repeatedly applying the above lemma, to approximate $S c(\boldsymbol{M}, F)$ as $S c\left(\boldsymbol{M}_{1}, F\right)$, where $M_{1}$ is $O\left(\varepsilon^{-1}\right)$-bDD. LastStep allows us to approximate the Schur complement for such a strongly block diagonally dominant matrix $\boldsymbol{M}_{1}$. The guarantees of LASTSTEP are given by the following lemma.

Lemma 13.11.3. Let $M$ be an bDD matrix with index set $V$, and $m$ nonzero blocks and let $F \subseteq$ $V$ be such that $\boldsymbol{M}_{[F, F]}$ is $\alpha$-bDD for some $\alpha \geq 4$. There exist a procedure LASTSTEP such that $\operatorname{LaStStep}(\boldsymbol{M}, F, \alpha, \varepsilon)$ returns in $O\left(m \varepsilon^{-8}\right)$ work and $O(\log n)$ depth a matrix $\widetilde{\boldsymbol{M}}_{S C}$ s.t. $\widetilde{\boldsymbol{M}}_{S C}$ has $O\left(m \varepsilon^{-8}\right)$ nonzero blocks and $\widetilde{\boldsymbol{M}}_{S C} \approx_{\varepsilon+2 / \alpha} S c(\boldsymbol{M}, F)$. If $\boldsymbol{T}$ is a matrix such that only the submatrix $\boldsymbol{T}_{[C, C]}$ is nonzero, and $\boldsymbol{M}+\boldsymbol{T}$ is $b D D$, then $\operatorname{LastStep}(\boldsymbol{M}+\boldsymbol{T}, F, \alpha, \varepsilon)=\operatorname{LastStep}(\boldsymbol{M}, F, \alpha, \varepsilon)+$ $\boldsymbol{T}_{[C, C]}$.

Combining the above two lemmas, we obtain a proof of Lemma 13.3.4.
Proof. (of Lemma 13.11.1). By induction, after $i$ steps of the main loop in ApproxSchur,

$$
S c(\boldsymbol{M}, F) \approx_{\frac{\frac{\varepsilon}{2 d}}{}} S c\left(M^{(i)}, F\right) .
$$

Lemma 13.11 .2 also implies that $\boldsymbol{M}^{(i)}$ is $2\left(\frac{\alpha}{2}\right)^{2^{i}}$-bDD. Thus, we have that $\boldsymbol{M}_{F F}^{(d)}$ is $8 \varepsilon^{-1}$-strongly diagonally dominant at the last step. Hence, Lemma 13.11.3 then gives

$$
S c\left(\boldsymbol{M}^{(d)}, F\right) \approx_{\frac{1}{2} \varepsilon} \operatorname{LASTSTEP}\left(\boldsymbol{M}^{(d)}, F, 4 \varepsilon^{-1}, \frac{\varepsilon}{4}\right) .
$$

Composing this bound with the guarantees of the iterations then gives the bound on overall error.
The property that if $\boldsymbol{T}$ is a matrix such that only the submatrix $\boldsymbol{T}_{[C, C]}$ is nonzero, and $\boldsymbol{M}+$ $\boldsymbol{T}$ is bDD , then $\operatorname{ApproxSchur}(\boldsymbol{M}+\boldsymbol{T}, F, \varepsilon)=\operatorname{ApproxSchur}(\boldsymbol{M}, F, \varepsilon)+\boldsymbol{T}_{[C, C]}$ follows from Lemma 13.11.2 and 13.11.3, which ensure that this property holds for all our calls to SchurSquare and LastStep.

The work of these steps, and the size of the output graph follow from Lemma 13.11.2 and 13.11.3.

## ■ 13.11.1 Iterative Squaring and Sparsification

In this section, we give a proof of Lemma 13.11.2. At the core of procedure is a squaring identity that preserves Schur complements, and efficient sparsification of special classes of bDD matrices that we
call product demand block-Laplacians.
Definition 13.11.4. The product demand block-Laplacian of a vector $\boldsymbol{d} \in\left(\mathbb{C}^{r \times r}\right)^{n}$, is a bDD matrix $\boldsymbol{L}_{G(\boldsymbol{d})} \in\left(\mathbb{C}^{r \times r}\right)^{n \times n}$, defined as

$$
\left(\boldsymbol{L}_{G(d)}\right)_{[i, j]}= \begin{cases}-\boldsymbol{d}_{[i]} \boldsymbol{d}_{[j]}^{*} & i \neq j, \\ \boldsymbol{I}_{r} \cdot\left\|\boldsymbol{d}_{[i]}\right\| \sum_{k: k \neq i}\left\|\boldsymbol{d}_{[k]}\right\| & \text { otherwise } .\end{cases}
$$

Definition 13.11.5. Given a vector $\boldsymbol{d} \in\left(\mathbb{C}^{r \times r}\right)^{n}$ and an index set $F \subseteq V$, let $C=V \backslash F$. The bipartite product demand block-Laplacian of $(\boldsymbol{d}, F)$ is a bDD matrix $\boldsymbol{L}_{G(\boldsymbol{d}), F} \in\left(\mathbb{C}^{r \times r}\right)^{n \times n}$, defined as

$$
\left(\boldsymbol{L}_{G(\boldsymbol{d}, F)}\right)_{[i, j]}= \begin{cases}\boldsymbol{I}_{r} \cdot\left\|\boldsymbol{d}_{[i]}\right\| \sum_{k \in C}\left\|\boldsymbol{d}_{[k]}\right\| & i=j, \text { and } i \in F \\ \boldsymbol{I}_{r} \cdot\left\|\boldsymbol{d}_{[i]}\right\| \sum_{k \in F}\left\|\boldsymbol{d}_{[k]}\right\| & i=j, \text { and } i \in C \\ 0 & i \neq j, \text { and }(i, j \in F \text { or } i, j \in C) \\ -\boldsymbol{d}_{[i]} \boldsymbol{d}_{[j]}^{*} & i \neq j, \text { otherwise } .\end{cases}
$$

In Section 13.11.4, we prove the following lemmas that allows us to efficiently construct sparse approximations to these matrices.

Lemma 13.11.6. There is a routine CliqueSparsification $(\boldsymbol{d}, \varepsilon)$ such that for any demand vector $\boldsymbol{d} \in\left(\mathbb{C}^{r \times r}\right)^{n}$, and $\varepsilon>0$, CLIqUESPARSIFICATION $(\boldsymbol{d}, \varepsilon)$ returns in $O\left(n \varepsilon^{-4}\right)$ work and $O(\log n)$ depth a bDD matrix $\boldsymbol{L}_{H}$ with $O\left(n \varepsilon^{-4}\right)$ nonzero blocks such that

$$
\boldsymbol{L}_{H} \approx_{\varepsilon} \boldsymbol{L}_{G(\boldsymbol{d})}
$$

Lemma 13.11.7. There is a routine BipartiteCliqueSparsification $(\boldsymbol{d}, F, \varepsilon)$ such that for any demand vector $\boldsymbol{d} \in\left(\mathbb{C}^{r \times r}\right)^{n}, \varepsilon>0$, and $F \subseteq V$, BipartiteCliqueSparsification $(\boldsymbol{d}, \varepsilon)$ returns in $O\left(n \varepsilon^{-4}\right)$ work and $O(\log n)$ depth a bDD matrix $\boldsymbol{L}_{H}$ with $O\left(n \varepsilon^{-4}\right)$ nonzero blocks such that

$$
\boldsymbol{L}_{H} \approx_{\varepsilon} \boldsymbol{L}_{G(\boldsymbol{d}, F)}
$$

Moreover, $\left(\boldsymbol{L}_{H}\right)_{[F, F]}$, and $\left(\boldsymbol{L}_{H}\right)_{[C, C]}$ are block-diagonal, where $C=V \backslash F$.
We now use these efficient sparse approximations to give a proof of the guarantees of the procedure SchurSquare.

```
Algorithm 40: \(\boldsymbol{M}_{1}=\operatorname{SchurSquare}(\boldsymbol{M}, F, \varepsilon)\)
    Construct sparse representations of \(\boldsymbol{d}_{F}^{(j)}, \boldsymbol{d}_{C}^{(j)}, \boldsymbol{f}^{(j)}\) defined by Eqs. (13.11), (13.12), and (13.13).
    Construct \(\boldsymbol{M}_{4}\) as given by Eq. (13.14).
    Compute \(\boldsymbol{M}_{3}\) as given by Eq. (13.15).
    For all \(j \in F\), let \(\widetilde{\boldsymbol{L}}_{G\left(d_{F}^{(j)}\right)} \leftarrow\) CLIQUESparsification \(\left(\boldsymbol{d}_{F}, \varepsilon\right)\).
    For all \(j \in F\), let \(\tilde{\boldsymbol{L}}_{G\left(\boldsymbol{d}_{C}^{(j)}\right)} \leftarrow\) CliqueSparsification \(\left(\boldsymbol{d}_{C}^{(j)}, \varepsilon\right)\).
    For all \(j \in F\), let \(\tilde{\boldsymbol{L}}_{G\left(\boldsymbol{f}^{(j)}, F\right)} \leftarrow \operatorname{BipartiteCliqueSparsification}\left(\boldsymbol{f}^{(j)}, \varepsilon\right)\).
    Construct and return \(\boldsymbol{M}_{1}\) defined by Eq. (13.16).
```

Figure 13-6: Pseudocode for procedure SchurSquare Iterative Squaring and Sparsification

Proof. (of Lemma 13.11.2) Let $C$ denote $V \backslash F$. Write $\boldsymbol{M}_{[F, F]}$ as $\boldsymbol{D}-\boldsymbol{A}$, where $\boldsymbol{D}$ is a block-diagonal, and $\boldsymbol{A}$ has its diagonal blocks as zero. The proof is based on the identity $S c(\boldsymbol{M}, F)=S c\left(\boldsymbol{M}_{2}, F\right)$, where $\boldsymbol{M}_{2}$ is the following matrix.

$$
\boldsymbol{M}_{2}=\frac{1}{2}\left[\begin{array}{cc}
\boldsymbol{D}-\boldsymbol{A} \boldsymbol{D}^{-1} \boldsymbol{A} & \boldsymbol{M}_{[F, C]}+\boldsymbol{A} \boldsymbol{D}^{-1} \boldsymbol{M}_{[F, C]}  \tag{13.10}\\
\boldsymbol{M}_{[C, F]}+\boldsymbol{M}_{[C, F]} \boldsymbol{D}^{-1} \boldsymbol{A} & 2 \boldsymbol{M}_{[C, C]}-\boldsymbol{M}_{[C, F]} \boldsymbol{D}^{-1} \boldsymbol{M}_{[F, C]}
\end{array}\right] .
$$

It is straightforward to prove that $\boldsymbol{M}_{2}$ satisfies $S c(\boldsymbol{M}, F)=S c\left(\boldsymbol{M}_{2}, F\right)$ (see Lemma 13.11.8). is also straightforward to show that $\boldsymbol{D}-\boldsymbol{A} \boldsymbol{D}^{-1} \boldsymbol{A}$ is $\left((\alpha+1)^{2}-1\right)$-bDD. Thus, $\boldsymbol{M}_{2}$ satisfies the first two requirements.

However, $\boldsymbol{M}_{2}$ is likely to be a dense matrix, and we will not construct it in full. The key observation is that $\boldsymbol{M}_{2}$ can be written as a sum of an explicit sparse bDD matrix, and several product demand block-Laplacians. Formally, for every $j \in F$, we define $\boldsymbol{d}_{F}^{(j)}, \boldsymbol{d}_{C}^{(j)} \in\left(\mathbb{C}^{r \times r}\right)^{n}$ as follows

$$
\begin{align*}
\left(\boldsymbol{d}_{F}^{(j)}\right)_{[i]} & = \begin{cases}\boldsymbol{A}_{[i, j]} \boldsymbol{D}_{[j, j]}^{-1 / 2} & i \in F, \\
0 & i \in C .\end{cases}  \tag{13.11}\\
\left(\boldsymbol{d}_{C}^{(j)}\right)_{[i]} & = \begin{cases}0 & i \in F, \\
\boldsymbol{M}_{[i, j]} \boldsymbol{D}_{[j, j]}^{-1 / 2} & i \in C .\end{cases} \tag{13.12}
\end{align*}
$$

For every $j \in F$, we need to define a bipartite product demand block-Laplacian given by $\boldsymbol{f}^{(j)} \in\left(\mathbb{C}^{r \times r}\right)^{n}$ defined as

$$
\boldsymbol{f}_{[i]}^{(j)}= \begin{cases}\boldsymbol{A}_{[i, j]} \boldsymbol{D}_{[j, j]}^{-1 / 2} & i \in F,  \tag{13.13}\\ -\boldsymbol{M}_{[i, j]} \boldsymbol{D}_{[j, j]}^{-1 / 2} & i \in C .\end{cases}
$$

Letting $\operatorname{deg}_{j}$ denote the number of nonzero blocks in $\boldsymbol{M}_{[j]}$, the number of nonzero blocks in each of $\boldsymbol{d}_{F}^{(j)}, \boldsymbol{d}_{C}^{(j)}, \boldsymbol{f}^{(j)}$ is at most $\operatorname{deg}_{j}$. We construct a sparse representation of each of them using $O\left(\operatorname{deg}_{j}\right)$ work. Thus, the total number of nonzero blocks in all these block-vectors is at most $O(m)$ and we can explicitly construct sparse representations for them using $O(m)$ work and $O(\log n)$ depth.

We can now express $\boldsymbol{M}_{2}$ as

$$
\boldsymbol{M}_{2}=\boldsymbol{M}_{3}+\frac{1}{2} \sum_{j \in F} \boldsymbol{L}_{G\left(\boldsymbol{d}_{F}^{(j)}\right)}+\frac{1}{2} \sum_{j \in F} \boldsymbol{L}_{G\left(\boldsymbol{d}_{C}^{(j)}\right)}+\frac{1}{2} \sum_{j \in F} \boldsymbol{L}_{G\left(\boldsymbol{f}^{(j)}, F\right)} .
$$

For all $i, j \in V$, we compute $\beta_{i, j}=\left\|\boldsymbol{f}_{i}^{(j)}\right\|$. Define $\boldsymbol{M}_{4}$ to be the block diagonal matrix such that

$$
\begin{equation*}
\boldsymbol{M}_{4[i, i]}=\frac{1}{2} \boldsymbol{I}_{r} \cdot \sum_{j \in F} \beta_{i, j}\left(\left(\sum_{k \in V} \beta_{k, j}\right)-\beta_{i, j}\right) . \tag{13.14}
\end{equation*}
$$

Since at most $m$ of $\beta_{i, j}$ are nonzero, and we can compute $\beta_{i, j}$ and construct $\boldsymbol{M}_{4}$ using using $O(m)$ work and $O(\log n)$ depth. It is easy to verify that we can express $\boldsymbol{M}_{3}$ explicitly as

$$
\boldsymbol{M}_{3}=\frac{1}{2}\left[\begin{array}{cc}
\boldsymbol{D} & \boldsymbol{M}_{[F, C]}  \tag{13.15}\\
\boldsymbol{M}_{[C, F]} & 2 \boldsymbol{M}_{[C, C]}
\end{array}\right]-\boldsymbol{M}_{4} .
$$

For each $j$, we use Lemmas 13.11 .6 and 13.11 .7 to construct, in $O\left(\operatorname{deg}_{j} \varepsilon^{-4}\right)$ work and $O(\log n)$ depth, bDD matrices $\tilde{\boldsymbol{L}}_{G\left(\boldsymbol{d}_{F}^{(j)}\right)}, \tilde{\boldsymbol{L}}_{G\left(\boldsymbol{d}_{C}^{(j)}\right)}, \tilde{\boldsymbol{L}}_{G\left(\boldsymbol{f}^{(j)}, F\right)}$, each with $O\left(\operatorname{deg}_{j} \varepsilon^{-4}\right)$ nonzero blocks such that

$$
\tilde{\boldsymbol{L}}_{G\left(d_{F}^{(j)}\right)} \approx_{\varepsilon} \boldsymbol{L}_{G\left(d_{F}^{(j)}\right)}, \quad \tilde{\boldsymbol{L}}_{G\left(d_{C}^{(j)}\right)} \approx_{\varepsilon} \boldsymbol{L}_{G\left(d_{C}^{(j)}\right)}, \quad \text { and } \quad \tilde{\boldsymbol{L}}_{G\left(f^{(j)}, F\right)} \approx_{\varepsilon} \boldsymbol{L}_{G\left(f^{(j)}, F\right)}
$$

We can now construct the matrix

$$
\begin{equation*}
\boldsymbol{M}_{1}=\boldsymbol{M}_{3}+\frac{1}{2} \sum_{j \in F} \tilde{\boldsymbol{L}}_{G\left(\boldsymbol{d}_{F}^{(j)}\right)}+\frac{1}{2} \sum_{j \in F} \tilde{\boldsymbol{L}}_{G\left(\boldsymbol{d}_{C}^{(j)}\right)}+\frac{1}{2} \sum_{j \in F} \tilde{\boldsymbol{L}}_{G\left(\boldsymbol{f}^{(j)}, F\right)}, \tag{13.16}
\end{equation*}
$$

in $O\left(m \varepsilon^{-4}\right)$ time and $O(\log n)$ depth. $\boldsymbol{M}_{1}$ has $O\left(m \varepsilon^{-4}\right)$ nonzero blocks and is our required matrix.
We first show that $\boldsymbol{M}_{3}$ is bDD. Using $\beta_{k, j} \leq\left\|\boldsymbol{D}_{[j, j]}^{-1 / 2} \boldsymbol{M}_{[i, k]}\right\|$ for all $j \in F, k \in V$, we have for all $i \in F$,

$$
\begin{equation*}
2\left\|\boldsymbol{M}_{4[i, i]}\right\| \leq \sum_{j \in F} \beta_{i, j} \sum_{k \in V} \beta_{k, j} \leq \sum_{j \in F}\left\|\boldsymbol{A}_{[i, j]}\right\|\left\|\boldsymbol{D}_{[j, j]}^{-1 / 2}\right\|^{2} \sum_{k \in V}\left\|\boldsymbol{M}_{[j, k]}\right\| \leq \sum_{j \in F}\left\|\boldsymbol{A}_{[i, j]}\right\|, \tag{13.17}
\end{equation*}
$$

where the last inequality uses that $M$ is bDD . Again, using $\boldsymbol{M}$ is bDD , for all $i \in F$, we have

$$
\boldsymbol{D}_{[i, i]} \succcurlyeq \boldsymbol{I}_{r} \sum_{j \in F}\left\|\boldsymbol{A}_{[i, j]}\right\|+\boldsymbol{I}_{r} \sum_{k \in C}\left\|\boldsymbol{M}_{[i, k]}\right\| .
$$

Now combining with Eqs. (13.15) and (13.17), we obtain $\boldsymbol{M}_{3[i, i]} \succcurlyeq \boldsymbol{I}_{r} \sum_{k \in C}\left\|\boldsymbol{M}_{[i, k]}\right\|$. Thus, $\boldsymbol{M}_{3}$ is bDD , and hence $\boldsymbol{M}_{1}$ is bDD.

We also have $\boldsymbol{M}_{1} \approx_{\varepsilon} \boldsymbol{M}_{2}$. Thus, using Fact 13.5.3, we obtain $\operatorname{Sc}\left(\boldsymbol{M}_{1}, F\right) \approx_{\varepsilon} S c\left(\boldsymbol{M}_{2}, F\right)$. It remains to show that $M_{1[F, F]}$ is $\alpha^{2} / 2$-bDD.

The key observation is that only the matrices $\tilde{\boldsymbol{L}}_{G\left(d_{C}^{(j)}\right)}$ contribute to off-diagonal blocks in $\boldsymbol{M}_{1[F, F]}$. By construction,

$$
\forall i, j \in F, \quad \boldsymbol{L}_{G\left(\boldsymbol{d}_{F}^{(j)}\right)_{[i, i]}}=\boldsymbol{I}_{r} \sum_{k \in F: k \neq i} \beta_{i, j} \beta_{k, j} .
$$

Thus, for all $i \in F$

$$
\begin{align*}
\sum_{j \in F} \| \boldsymbol{L}_{G\left(\boldsymbol{d}_{F}^{(j)}\right)}^{[i, i]}
\end{align*}\left\|\leq \sum_{j \in F} \sum_{k \in F} \beta_{i, j} \beta_{k, j} \leq \sum_{j \in F}\right\| \boldsymbol{A}_{[i, j]}\| \| \boldsymbol{D}_{[j, j]}^{-1 / 2}\left\|^{2} \sum_{k \in F}\right\| \boldsymbol{A}_{[j, k]} \|
$$

where the last inequality follows since $\boldsymbol{D}-\boldsymbol{A}=\boldsymbol{M}_{[F, F]}$ is $\alpha$-bDD. Moreover $\boldsymbol{D}-\boldsymbol{A}$ being $\alpha$-bDD implies that for each $i \in F, \boldsymbol{D}_{[i, i]} \succcurlyeq(\alpha+1) \boldsymbol{I}_{r} \sum_{j \in F}\left\|\boldsymbol{A}_{[i, j]}\right\|$. Thus, using Eqs. (13.15) and (13.17), we obtain, $2 \boldsymbol{M}_{3[i, i]} \succcurlyeq \alpha \boldsymbol{I}_{r} \sum_{j \in F}\left\|\boldsymbol{A}_{[i, j]}\right\|$.

Since for each j, $\tilde{\boldsymbol{L}}_{G\left(\boldsymbol{d}^{(j)}\right)} \approx_{\varepsilon} \tilde{\boldsymbol{L}}_{G\left(\boldsymbol{d}^{(j)}\right)}$, we have $\left(\tilde{\boldsymbol{L}}_{G\left(\boldsymbol{d}^{(j)}\right)}\right)_{[i, i]} \preccurlyeq e^{\varepsilon}\left(\boldsymbol{L}_{G\left(\boldsymbol{d}^{(j)}\right)}\right)_{[i, i]}$. Thus, $\sum_{j} \tilde{\boldsymbol{L}}_{G\left(\boldsymbol{d}_{F}^{(j)}\right)}$ is bDD with sums of norms of off-diagonal blocks at most $e^{\varepsilon}(1+\alpha)^{-1} \sum_{j \in F}\left\|\boldsymbol{A}_{[i, j]}\right\|$. Since these are the only matrices that contribute to off-diagonal blocks in $\boldsymbol{M}_{1[F, F]}$, we get that $\boldsymbol{M}_{1[F, F]}$ is at least


It is easy to verify that our transformations maintain that if $\boldsymbol{T}$ is a matrix that is only nonzero inside the submatrix $\boldsymbol{T}_{[C, C]}$, and $\boldsymbol{M}+\boldsymbol{T}$ is bDD , then

$$
\operatorname{SchurSquare}(\boldsymbol{M}+\boldsymbol{T}, F, \varepsilon)=\operatorname{SchurSquare}(\boldsymbol{M}, F, \varepsilon)+\boldsymbol{T}_{[C, C]} .
$$

We now prove the claims deferred from the above proof.
Lemma 13.11.8. The matrix $\boldsymbol{M}_{2}$ defined by Eq. (13.10) satisfies $S c\left(\boldsymbol{M}_{2}, F\right)=S c(\boldsymbol{M}, F)$.
Proof. We need the following identity from [223]:

$$
\begin{equation*}
(\boldsymbol{D}-\boldsymbol{A})^{-1}=1 / 2 \cdot\left(\boldsymbol{D}^{-1}+\left(\boldsymbol{I}+\boldsymbol{D}^{-1} \boldsymbol{A}\right)\left(\boldsymbol{D}-\boldsymbol{A} \boldsymbol{D}^{-1} \boldsymbol{A}\right)^{-1}\left(\boldsymbol{I}+\boldsymbol{A} \boldsymbol{D}^{-1}\right)\right) \tag{13.19}
\end{equation*}
$$

We have,

$$
\begin{aligned}
S c\left(\boldsymbol{M}_{2}, F\right)= & \boldsymbol{M}_{[C, C]}-1 / 2 \cdot \boldsymbol{M}_{[C, F]} \boldsymbol{D}^{-1} \boldsymbol{M}_{[F, C]} \\
& \quad-1 / 2 \cdot \boldsymbol{M}_{[C, F]}\left(\boldsymbol{I}+\boldsymbol{D}^{-1} \boldsymbol{A}\right)\left(\boldsymbol{D}-\boldsymbol{A} \boldsymbol{D}^{-1} \boldsymbol{A}\right)^{-1}\left(\boldsymbol{I}+\boldsymbol{A} \boldsymbol{D}^{-1}\right) \boldsymbol{M}_{[F, C]} \\
= & \boldsymbol{M}_{[C, C]} \\
& \quad-1 / 2 \cdot \boldsymbol{M}_{[C, F]}\left(\boldsymbol{D}^{-1}+\left(\boldsymbol{I}+\boldsymbol{D}^{-1} \boldsymbol{A}\right)\left(\boldsymbol{D}-\boldsymbol{A} \boldsymbol{D}^{-1} \boldsymbol{A}\right)^{-1}\left(\boldsymbol{I}+\boldsymbol{A} \boldsymbol{D}^{-1}\right)\right) \boldsymbol{M}_{[F, C]} \\
= & \boldsymbol{M}_{[C, C]}-\boldsymbol{M}_{[C, F]}(\boldsymbol{D}-\boldsymbol{A})^{-1} \boldsymbol{M}_{[F, C]} \\
= & S c(\boldsymbol{M}, F) .
\end{aligned}
$$

## ■ 13.11.2 Schur Complement w.r.t. Highly $\alpha$-bDD Submatrices

```
Algorithm 41: \(\widetilde{\boldsymbol{M}}_{S C}=\operatorname{LASTSTEP}(\boldsymbol{M}, F, \alpha, \varepsilon)\)
    Compute \(\boldsymbol{X}, \boldsymbol{D}\), and \(\boldsymbol{A}\) as given by equations (13.20), (13.21), and (13.22).
    Construct \(\boldsymbol{Y}\) as defined by equations (13.33), (13.34), and (13.35).
    Construct sparse vectors \(\boldsymbol{d}^{(j)}, \boldsymbol{g}^{(j)}\) for all \(j \in F\) as defined by equations (13.26) and (13.27).
    For all \(j \in F\), let \(\tilde{\boldsymbol{L}}_{G\left(\boldsymbol{d}^{(j)}, F\right)} \leftarrow \operatorname{BipartiteCliqueSparsification~}\left(\boldsymbol{d}^{(j)}, F, \varepsilon / 2\right)\)
    For all \(j \in F\), let \(\tilde{\boldsymbol{L}}_{G\left(\boldsymbol{g}^{(j)}\right)} \leftarrow\) CliqueSparsification \(\left(\boldsymbol{g}^{(j)}, \varepsilon / 2\right)\).
    Compute \(\boldsymbol{R}\) as given by equation (13.29).
    Construct sparse vectors \(\boldsymbol{r}^{(j)}\) for all \(j \in F\) as defined by equations (13.30).
    For all \(j \in F\), let \(\tilde{\boldsymbol{L}}_{G\left(\boldsymbol{r}^{(j)}\right)} \leftarrow \operatorname{CliqueSparsification~}\left(\boldsymbol{r}^{(j)}, \varepsilon / 2\right)\).
    Compute \(\boldsymbol{S}\) as given by equation (13.31).
    Output
    \(\widetilde{\boldsymbol{M}}_{S C}=\boldsymbol{S}+\sum_{i} \tilde{\boldsymbol{L}}_{G\left(\boldsymbol{r}^{(i)}\right)}\).
```

Figure 13-7: Pseudocode for procedure LASTSTEP: Computing an approximate Schur complement w.r.t. a highly $\alpha$-bDD submatrices.

In this section, we describe the LastStep procedure for computing an approximate Schur complement of a bDD matrix $M$ with a highly $\alpha$-bDD submatrix $\boldsymbol{M}_{[F, F]}$. LastStep is the final step of the ApproxSchur algorithm. The key element of the procedure is a formula for approximating the inverse of $\boldsymbol{M}_{[F, F]}$ that is leveraged to approximate the Schur complement of $\boldsymbol{M}$ as the Schur complement of matrix with the $F F$ submatrix being block diagonal.

We prove guarantees for the LASTSTEP algorithm as stated in Lemma 13.11.3.
One could attempt to deal with the highly $\alpha$-bDD matrix at the last step by directly replacing it with its diagonal, but this is problematic. Consider the case where $F$ contains $u$ and $v$ with a weight $\varepsilon$ edge between them, and $u$ and $v$ are connected to $u^{\prime}$ and $v^{\prime}$ in $C$ by weight 1 edges respectively. Keeping only the diagonal results in a Schur complement that disconnects $u^{\prime}$ and $v^{\prime}$. This however can be fixed by taking a step of random walk within $F$.

Given a bDD matrix $\boldsymbol{M}$, s.t. $\boldsymbol{M}_{[F, F]}$ is $\alpha$-bDD we define a block diagonal matrix $\boldsymbol{X} \in\left(\mathbb{C}^{r \times r}\right)^{|F| \times|F|}$ s.t. for each $i \in F$

$$
\begin{equation*}
\boldsymbol{X}_{[i, i]}=\frac{\alpha}{\alpha+1} \boldsymbol{M}_{[i, i]} \tag{13.20}
\end{equation*}
$$

and another block diagonal matrix $\boldsymbol{D} \in\left(\mathbb{C}^{r \times r}\right)^{|F| \times|F|}$ s.t. for each $i \in F$

$$
\begin{equation*}
\boldsymbol{D}_{[i, i]}=\frac{1}{\alpha+1} \boldsymbol{M}_{[i, i]} \tag{13.21}
\end{equation*}
$$

and we define a matrix $\boldsymbol{A} \in\left(\mathbb{C}^{r \times r}\right)^{|F| \times|F|}$

$$
\boldsymbol{A}_{[i, j]}= \begin{cases}0 & i=j  \tag{13.22}\\ -\boldsymbol{M}_{[i, j]} & \text { otherwise } .\end{cases}
$$

Thus $\boldsymbol{M}_{[F, F]}=\boldsymbol{X}+\boldsymbol{D}-\boldsymbol{A}$. One can check that because $\boldsymbol{M}$ is bDD and $\boldsymbol{M}_{[F, F]}$ is $\alpha$-bDD, it follows that $\boldsymbol{D}-\boldsymbol{A}$ is bDD and the matrix

$$
\left[\begin{array}{cc}
\boldsymbol{X} & \boldsymbol{M}_{[F, C]} \\
\boldsymbol{M}_{[C, F]} & \boldsymbol{M}_{[C, C]}
\end{array}\right]
$$

is also bDD .
We will consider the linear operator

$$
\begin{equation*}
\boldsymbol{Z}^{(l a s t)} \stackrel{\text { def }}{=} \frac{1}{2} \boldsymbol{X}^{-1}+\frac{1}{2} \boldsymbol{X}^{-1}(\boldsymbol{X}-\boldsymbol{D}+\boldsymbol{A}) \boldsymbol{X}^{-1}(\boldsymbol{X}-\boldsymbol{D}+\boldsymbol{A}) \boldsymbol{X}^{-1} \tag{13.23}
\end{equation*}
$$

We define

$$
\boldsymbol{M}^{(l a s t)}=\left(\begin{array}{cc}
\left(\boldsymbol{Z}^{(l a s t)}\right)^{-1} & \boldsymbol{M}_{[F, C]} \\
\boldsymbol{M}_{[F, C]} & \boldsymbol{M}_{[C, C]}
\end{array}\right)
$$

## Lemma 13.11.9.

$$
\boldsymbol{M} \preccurlyeq \boldsymbol{M}^{(l a s t)} \preccurlyeq\left(1+\frac{2}{\alpha}\right) \boldsymbol{M} .
$$

We defer the proof of Lemma 13.11.9 to Section 13.11.3.
To utilize $\boldsymbol{Z}^{\text {(last) }}$, define

$$
\boldsymbol{M}_{2}^{(l a s t)} \stackrel{\text { def }}{=}\left[\begin{array}{cc}
\frac{1}{2} \boldsymbol{X} & \frac{1}{2}(\boldsymbol{X}-\boldsymbol{D}+\boldsymbol{A}) \boldsymbol{X}^{-1} \boldsymbol{M}_{[F, C]}  \tag{13.24}\\
\frac{1}{2} \boldsymbol{M}_{[C, F]} \boldsymbol{X}^{-1}(\boldsymbol{X}-\boldsymbol{D}+\boldsymbol{A}) & \boldsymbol{M}_{[C, C]}-\frac{1}{2} \boldsymbol{M}_{[C, F]} \boldsymbol{X}^{-1} \boldsymbol{M}_{[F, C]} .
\end{array}\right]
$$

and note

$$
\begin{equation*}
S c\left(\boldsymbol{M}^{(l a s t)}, F\right)=S c\left(\boldsymbol{M}_{2}^{(l a s t)}, F\right) \tag{13.25}
\end{equation*}
$$

Lemma 13.11.9 tells us that for large enough $\alpha$, we can approximate the Schur complement of $\boldsymbol{M}$ by approximating the the Schur complement of $\boldsymbol{M}_{2}^{(l a s t)}$.

The next lemma tells us that $M_{2}$ is bDD and that we can write the matrix as a sum of an explicit bDD matrix and sparse implicitly represented product demand block-Laplacians and bipartite product demand block-Laplacians.

Lemma 13.11.10. Consider a $b D D$ matrix $M$, where $\boldsymbol{M}_{[F, F]}$ is $\alpha-b D D$ for some $\alpha \geq 4$, and let $\boldsymbol{M}_{2}^{(l a s t)}$ be the associated matrix defined by equation (13.24). Let $m$ be the number of nonzero blocks of $\boldsymbol{M}$.

For $j \in F$, we define $\boldsymbol{d}^{(i)} \in\left(\mathbb{C}^{r \times r}\right)^{n}$

$$
\boldsymbol{d}_{[i]}^{(j)}= \begin{cases}\boldsymbol{A}_{[i, j]} \boldsymbol{X}_{[j, j]}^{-1 / 2} & \text { for } i \in F  \tag{13.26}\\ \boldsymbol{M}_{[i, j]} \boldsymbol{X}_{[j, j]}^{-1 / 2} & \text { for } i \in C\end{cases}
$$

For $j \in F$, we define $\boldsymbol{g}^{(j)} \in\left(\mathbb{C}^{r \times r}\right)^{n}$

$$
\boldsymbol{g}_{[i]}^{(j)}= \begin{cases}0 & \text { for } i \in F  \tag{13.27}\\ \boldsymbol{M}_{[i, j]} \boldsymbol{X}_{[j, j]}^{-1 / 2} & \text { for } i \in C\end{cases}
$$

Then

$$
\boldsymbol{M}_{2}^{(l a s t)}=\boldsymbol{Y}+\frac{1}{2} \sum_{j \in F} \boldsymbol{L}_{G\left(\boldsymbol{g}^{(i)}\right)}+\frac{1}{2} \sum_{j \in F} \boldsymbol{L}_{G\left(\boldsymbol{d}^{(i)}, F\right)}
$$

where $\boldsymbol{Y}$ is bDD and has $O(m)$ nonzero blocks, and the total number of nonzero blocks in $\boldsymbol{d}^{(i)}$ and $\boldsymbol{g}^{(i)}$ for all $i$ combined is also $O(m)$.
$\boldsymbol{Y}$ as well as $\boldsymbol{d}^{(i)}$ and $\boldsymbol{g}^{(i)}$ for all $i$ can be computed in $O(m)$ time and $O(\log n)$ depth.
If $\boldsymbol{T}$ is a matrix that is only nonzero inside the submatrix $\boldsymbol{T}_{[C, C]}$, then if we apply the transformation of Eq. 13.24, to $\boldsymbol{M}+\boldsymbol{T}$ instead of $\boldsymbol{M}$, we find $(\boldsymbol{M}+\boldsymbol{T})_{2}^{(l a s t)}=\boldsymbol{M}_{2}^{(\text {last })}+\boldsymbol{T}_{[C, C]}$, and in particular $\boldsymbol{Y}(\boldsymbol{M}+\boldsymbol{T})=\boldsymbol{Y}(\boldsymbol{M})+\boldsymbol{T}_{[C, C]}$.

We defer the proof of Lemma 13.11.10 to Section 13.11.3.
Proof. (of Lemma 13.11.3) The procedure LastStep (M,F, $\alpha, \varepsilon)$ first computes $\boldsymbol{Y}$ and $\boldsymbol{d}^{(i)}$ and $\boldsymbol{g}^{(i)}$ for all $i$ s.t.

$$
\begin{equation*}
\boldsymbol{M}_{2}^{(\text {last })}=\boldsymbol{Y}+\frac{1}{2} \sum_{j \in F} \boldsymbol{L}_{G\left(\boldsymbol{g}^{(i)}\right)}+\frac{1}{2} \sum_{j \in F} \boldsymbol{L}_{G\left(\boldsymbol{d}^{(i)}, F\right)}, \tag{13.28}
\end{equation*}
$$

where $\boldsymbol{M}_{2}^{(l a s t)} \approx_{2 / \alpha} \boldsymbol{M}$. Let $n_{\boldsymbol{d}^{(i)}}$ and $n_{\boldsymbol{g}^{(i)}}$ denote the number of nonzero blocks in each $\boldsymbol{d}^{(i)}$ and $\boldsymbol{g}^{(i)}$. By Lemma 13.11.6, we may call CliqueSparsification $\left(\boldsymbol{g}^{(i)}, \varepsilon / 2\right)$ in $O\left(n_{\boldsymbol{g}^{(i)}} \varepsilon^{-4}\right)$ time to sparsify $\boldsymbol{L}_{G\left(\boldsymbol{g}^{(i)}\right)}$, producing a bDD matrix $\tilde{\boldsymbol{L}}_{G\left(\boldsymbol{g}^{(i)}\right)} \approx_{\varepsilon / 2} \boldsymbol{L}_{G\left(\boldsymbol{g}^{(i)}\right)}$ with $O\left(n_{\boldsymbol{g}^{(i)}} \varepsilon^{-4}\right)$ nonzero blocks. By Lemma 13.11.7, we may call BipartiteCliqueSparsification $\left(\boldsymbol{d}^{(i)}, F, \varepsilon / 2\right)$ in $O\left(n_{\boldsymbol{d}^{(i)}} \varepsilon^{-4}\right)$ time to sparsify $\boldsymbol{L}_{G\left(\boldsymbol{d}^{(i)}, F\right)}$, producing a bDD matrix $\tilde{\boldsymbol{L}}_{G\left(\boldsymbol{d}^{(i)}, F\right)} \approx_{\varepsilon / 2} \boldsymbol{L}_{G\left(\boldsymbol{d}^{(i)}, F\right)}$ with $O\left(n_{\boldsymbol{d}^{(i)}} \varepsilon^{-4}\right)$ nonzero blocks. The total running time of this is $O\left(\varepsilon^{-4} \sum_{i}\left(n_{\boldsymbol{d}^{(i)}}+n_{\boldsymbol{g}^{(i)}}\right)\right)=O\left(\varepsilon^{-4} m\right)$, and the total number of nonzero blocks in

$$
\frac{1}{2} \sum_{j \in F} \tilde{\boldsymbol{L}}_{G\left(\boldsymbol{g}^{(i)}\right)}+\frac{1}{2} \sum_{j \in F} \tilde{\boldsymbol{L}}_{G\left(\boldsymbol{d}^{(i)}, F\right)},
$$

is $O\left(\varepsilon^{-4}\left(\sum_{i} n_{\boldsymbol{d}^{(i)}}+n_{\boldsymbol{g}^{(i)}}\right)\right)=O\left(\varepsilon^{-4} m\right)$. We define

$$
\begin{equation*}
\boldsymbol{R}=\boldsymbol{Y}+\frac{1}{2} \sum_{j \in F} \tilde{\boldsymbol{L}}_{G\left(\boldsymbol{g}^{(i)}\right)}+\frac{1}{2} \sum_{j \in F} \tilde{\boldsymbol{L}}_{G\left(\boldsymbol{d}^{(i)}, F\right)} . \tag{13.29}
\end{equation*}
$$

which we can compute in $O\left(\varepsilon^{-4} m\right)$ time and $O(\log n)$ depth. We have $\boldsymbol{M} \approx_{2 / \alpha} \boldsymbol{M}_{2}^{(l a s t)}$ and $\boldsymbol{M}_{2}^{(l a s t)} \approx_{\varepsilon / 2}$ $\boldsymbol{R}$, so that $\boldsymbol{M} \approx_{2 / \alpha+\varepsilon / 2} \boldsymbol{R}$. It follows from Fact 13.5.3 that $S c(\boldsymbol{M}, F) \approx_{2 / \alpha+\varepsilon / 2} S c(\boldsymbol{R}, F)$.

Because the sparsifiers computed by BipartitecliqueSparsificationpreserve the graph bipartition, $\boldsymbol{R}_{[F, F]}$ is block diagonal.

We can use the block diagonal structure of $\boldsymbol{R}_{[F, F]}$ quickly compute a sparse approximation to $S c(\boldsymbol{R}, F)$.

For $j \in F$, we define $\boldsymbol{g}^{(j)} \in\left(\mathbb{C}^{r \times r}\right)^{|C|}$

$$
\begin{equation*}
\boldsymbol{r}_{i}^{(j)}=\boldsymbol{R}_{[i, j]} \boldsymbol{R}_{[j, j]}^{-1 / 2} \text { for } i \in C \tag{13.30}
\end{equation*}
$$

Then

$$
S c(\boldsymbol{R}, F)=\boldsymbol{R}_{[C, C]}-\boldsymbol{R}_{[C, F]} \boldsymbol{R}_{[F, F]}^{-1} \boldsymbol{R}_{[F, C]}=\boldsymbol{S}+\sum_{j \in F} \boldsymbol{L}_{G\left(\boldsymbol{r}^{(i)}\right)} .
$$

where

$$
\begin{align*}
\boldsymbol{S}= & \boldsymbol{R}_{[C, C]}-\operatorname{Diag}_{i \in C}\left(\boldsymbol{I}_{r} \sum_{j \in F} \sum_{k \in C \backslash\{i\}}\left\|\boldsymbol{R}_{[i, j]} \boldsymbol{R}_{[j, j]}^{-1 / 2}\right\|\left\|\boldsymbol{R}_{[k, j]} \boldsymbol{R}_{[j, j]}^{-1 / 2}\right\|\right)  \tag{13.31}\\
& -\operatorname{Diag}_{i \in C} \sum_{j \in F}\left(\boldsymbol{R}_{[i, j]} \boldsymbol{R}_{[j, j]}^{-1} \boldsymbol{R}_{[j, i]]}\right) .
\end{align*}
$$

Let us define, for each block row $i \in F, \rho_{i}=\sum_{j \in C}\left\|\boldsymbol{R}_{[i, j]}\right\|$, and for each block row $i \in C$, $\rho_{i}=\sum_{j \in F}\left\|\boldsymbol{R}_{[i, j]}\right\|$. From $\boldsymbol{R}$ being bDD, we then conclude for each $i \in F$

$$
\boldsymbol{R}_{[i, i]} \succcurlyeq \boldsymbol{I}_{r} \rho_{i}
$$

and for each $i \in C$

$$
\boldsymbol{R}_{[i, i]} \succcurlyeq \boldsymbol{I}_{r}\left(\rho_{i}+\sum_{j \in C}\left\|\boldsymbol{R}_{[i, j]}\right\|\right) .
$$

With this in mind, we check that each for $i \in C$ of $\boldsymbol{S}$ is bDD.
The diagonal block satisfies

$$
\begin{aligned}
\boldsymbol{S}_{[i, i]} & \succcurlyeq \boldsymbol{I}_{r}\left(\sum_{j \in C}\left\|\boldsymbol{R}_{[i, j]}\right\|+\rho_{i}-\sum_{j \in F} \sum_{k \in C}\left\|\boldsymbol{R}_{[i, j]} \boldsymbol{R}^{-1 / 2}{ }_{[j, j]}\right\|\left\|\boldsymbol{R}_{[k, j]} \boldsymbol{R}^{-1 / 2}{ }_{[j, j]}\right\|\right) \\
& \succcurlyeq \boldsymbol{I}_{r}\left(\sum_{j \in C}\left\|\boldsymbol{R}_{[i, j]}\right\|+\rho_{i}-\sum_{j \in F} \sum_{k \in C} \frac{1}{\rho_{j}}\left\|\boldsymbol{R}_{[i, j]}\right\|\left\|\boldsymbol{R}_{[k, j]}\right\|\right) \\
& \succcurlyeq \boldsymbol{I}_{r}\left(\sum_{j \in C}\left\|\boldsymbol{R}_{[i, j]}\right\|+\rho_{i}-\sum_{j \in F} \frac{1}{\rho_{j}}\left\|\boldsymbol{R}_{[i, j]}\right\|\right) \\
& \succcurlyeq \boldsymbol{I}_{r} \sum_{j \in C}\left\|\boldsymbol{R}_{[i, j]}\right\|
\end{aligned}
$$

We can compute $\boldsymbol{S}$ and all $\boldsymbol{r}_{i}^{(j)}$ in $O\left(m \varepsilon^{-4}\right)$ time and $O(\log n)$ depth, since this is an upper bound to the number of nonzero blocks in $\boldsymbol{R}$. Let $n_{\boldsymbol{r}^{(i)}}$ denote the number of nonzero blocks in $\boldsymbol{r}^{(i)}$. By Lemma 13.11.6, we may call CliqueSparsification $\left(\boldsymbol{r}^{(i)}, \varepsilon / 2\right)$ in $O\left(n_{\boldsymbol{r}^{(i)}} \varepsilon^{-4}\right)$ time to sparsify $\boldsymbol{L}_{G\left(\boldsymbol{r}^{(i)}\right)}$, producing a bDD matrix $\tilde{\boldsymbol{L}}_{G\left(\boldsymbol{r}^{(i)}\right)} \approx_{\varepsilon / 2} \boldsymbol{L}_{G\left(\boldsymbol{r}^{(i)}\right)}$ with $O\left(n_{\boldsymbol{r}^{(i)}} \varepsilon^{-4}\right)$ nonzero blocks.

The total number of nonzero blocks in

$$
\boldsymbol{S}+\sum_{i} \tilde{\boldsymbol{L}}_{G\left(\boldsymbol{r}^{(i)}\right)}
$$

is also $O\left(\varepsilon^{-8} m\right)$, and $\boldsymbol{S}+\sum_{i} \tilde{\boldsymbol{L}}_{G\left(\boldsymbol{r}^{(i)}\right)} \approx_{\varepsilon / 2} S c(\boldsymbol{R}, F)$. So

$$
\widetilde{\boldsymbol{M}}_{S C}=\boldsymbol{S}+\sum_{i} \tilde{\boldsymbol{L}}_{G\left(\boldsymbol{r}^{(i)}\right)} \approx_{\varepsilon+2 / \alpha} S c(\boldsymbol{M}, F)
$$

The total amount of work to compute these $\tilde{\boldsymbol{L}}_{G\left(\boldsymbol{r}^{(i)}\right)}$ is $O\left(\varepsilon^{-4}\left(\sum_{i} n_{\boldsymbol{r}^{(i)}}\right)\right)=O\left(\varepsilon^{-8} m\right)$. The depth for this computation is $O(\log n)$.

Suppose $\boldsymbol{T}$ is a matrix that is only nonzero inside the submatrix $\boldsymbol{T}_{[C, C]}$, and $\boldsymbol{M}+\boldsymbol{T}$ is bDD . We can show that LastStep $(\boldsymbol{M}+\boldsymbol{T}, F, \alpha, \varepsilon)=\operatorname{LastStep}(\boldsymbol{M}, F, \alpha, \varepsilon)+\boldsymbol{T}_{[C, C]}$, by first noting that
this type of property holds for $\boldsymbol{Y}$ by Lemma 13.11.10, and from this concluding that similarly if we consider $\boldsymbol{R}$ as a function of $\boldsymbol{M}$ then $\boldsymbol{R}_{[C, C]}(\boldsymbol{M}+\boldsymbol{T})=\boldsymbol{R}_{[C, C]}(\boldsymbol{M})+\boldsymbol{T}_{[C, C]}$, and finally considering $\widetilde{\boldsymbol{M}}_{S C}$ as a function of $\boldsymbol{M}$, we can then easily show that $(\widetilde{\boldsymbol{M}+\boldsymbol{T}})_{S C}=\widetilde{\boldsymbol{M}}_{S C}+\boldsymbol{T}_{[C, C]}$.

## ■ 13.11.3 Deferred Proofs from Section 13.11.2

To help us prove Lemma 13.11.9, we first prove the next lemma.
Lemma 13.11.11. If $\boldsymbol{M}_{[F, F]}=\boldsymbol{X}+\boldsymbol{D}-\boldsymbol{A}$ be a $\alpha$-bDD matrix for some $\alpha \geq 4$, then the operator $\boldsymbol{Z}^{(l a s t)}$ as defined in Equation 13.23 satisfies:

$$
\boldsymbol{M}_{[F, F]} \preccurlyeq\left(\boldsymbol{Z}^{(l a s t)}\right)^{-1} \preccurlyeq \boldsymbol{M}_{[F, F]}+\frac{2}{\alpha}(\boldsymbol{D}-\boldsymbol{A}) .
$$

Proof. Composing both sides by $\boldsymbol{X}^{-1 / 2}$ and substituting in $\mathcal{L}=\boldsymbol{X}^{-1 / 2}(\boldsymbol{D}-\boldsymbol{A}) \boldsymbol{X}^{-1 / 2}$ means it suffices to show

$$
\boldsymbol{I}+\mathcal{L} \preccurlyeq\left(\frac{1}{2} \boldsymbol{I}+\frac{1}{2}(\boldsymbol{I}-\mathcal{L})^{2}\right)^{-1} \preccurlyeq \boldsymbol{I}+\mathcal{L}+\frac{2}{\alpha} \mathcal{L} .
$$

We can use the fact that $\boldsymbol{M}_{[F, F]}$ is $\alpha$-strongly diagonally dominant to show $0 \preccurlyeq \boldsymbol{L} \preccurlyeq \frac{2}{\alpha} \boldsymbol{X}$, and equivalently $0 \preccurlyeq \mathcal{L} \preccurlyeq \frac{2}{\alpha} \boldsymbol{I}$, as follows:

Firstly, $\boldsymbol{D}-\boldsymbol{A} \succcurlyeq 0$ as $\boldsymbol{D}-\boldsymbol{A}$ is bDD, similarly, $\boldsymbol{D}+\boldsymbol{A} \succcurlyeq 0$ as $\boldsymbol{D}+\boldsymbol{A}$ is also bDD. From the latter $\boldsymbol{D} \succcurlyeq-\boldsymbol{A}$, so $2 \boldsymbol{D} \succcurlyeq \boldsymbol{D}-\boldsymbol{A}$. Finally, $\boldsymbol{X} \succcurlyeq \alpha \boldsymbol{D} \succcurlyeq \frac{\alpha}{2}(\boldsymbol{D}-\boldsymbol{A})$.

As $\mathcal{L}$ and $\boldsymbol{I}$ commute, the spectral theorem means it suffices to show this for any scalar $0 \leq t \leq \frac{2}{\alpha}$. Note that

$$
\frac{1}{2}+\frac{1}{2}(1-t)^{2}=1-t+\frac{1}{2} t^{2}
$$

Taking the difference between the inverse of this and the 'true' value of $1+t$ gives:

$$
\left(1-t+\frac{1}{2} t^{2}\right)^{-1}-(1+t)=\frac{1-(1+t)\left(1-t+\frac{1}{2} t^{2}\right)}{1-t+\frac{1}{2} t^{2}}=\frac{\frac{1}{2} t^{2}(1-t)}{1-t+\frac{1}{2} t^{2}}
$$

Incorporating the assumption that $0 \leq t \leq \frac{2}{\alpha}$ and $\alpha \geq 4$ gives that the denominator is at least $1-\frac{2}{\alpha} \geq \frac{1}{2}$, and the numerator term can be bounded by $0 \leq \frac{t^{2}}{2}(1-t) \leq \frac{t}{\alpha}$. Combining these two bounds then gives the result.

Lemma 13.11 .9 allows us to extend the approximation of $\boldsymbol{M}_{[F, F]}$ by the inverse of $\boldsymbol{Z}^{(l a s t)}$ to the entire matrix $M$.

Proof. (of Lemma 13.11.9) Recall that when a matrix $\boldsymbol{T}$ is PSD,

$$
\left(\begin{array}{ll}
\boldsymbol{T} & 0  \tag{13.32}\\
0 & 0
\end{array}\right) \succcurlyeq 0 .
$$

The left-hand inequality of our lemma follows immediately from Eq. 13.32 and the left-hand side of the guarantee of Lemma 13.11.11. To prove the right-hand inequality we apply Eq. 13.32 and the right-hand side of the guarantee of Lemma 13.11.11. to conclude

$$
\left(\begin{array}{cc}
\left(\boldsymbol{Z}^{(l a s t)}\right)^{-1} & \boldsymbol{M}_{[F, C]} \\
\boldsymbol{M}_{[F, C]} & \boldsymbol{M}_{[C, C]}
\end{array}\right) \preccurlyeq\left(\begin{array}{cc}
\boldsymbol{M}_{[F, F]}+\frac{2}{\alpha} \boldsymbol{L} & \boldsymbol{M}_{[F, C]} \\
\boldsymbol{M}_{[F, C]} & \boldsymbol{M}_{[C, C]}
\end{array}\right)=\boldsymbol{M}+\frac{2}{\alpha}\left(\begin{array}{ll}
\boldsymbol{L} & 0 \\
0 & 0
\end{array}\right) .
$$

The matrix $\left[\begin{array}{cc}\boldsymbol{X} & \boldsymbol{M}_{[F, C]} \\ \boldsymbol{M}_{[C, F]} & \boldsymbol{M}_{[C, C]}\end{array}\right]$ is bDD and hence PSD. It follows that $\left(\begin{array}{cc}\boldsymbol{D}-\boldsymbol{A} & 0 \\ 0 & 0\end{array}\right) \preccurlyeq \boldsymbol{M}$, by which we may conclude that $\boldsymbol{M}+\frac{2}{\alpha}\left(\begin{array}{cc}\boldsymbol{D}-\boldsymbol{A} & 0 \\ 0 & 0\end{array}\right) \preccurlyeq \boldsymbol{M}+\frac{2}{\alpha} \boldsymbol{M}$.

Proof. (of Lemma 13.11.10)
Each product demand clique $\boldsymbol{L}_{G\left(\boldsymbol{g}^{(i)}\right)}$ and bipartite product demand clique $\boldsymbol{L}_{G\left(\boldsymbol{d}^{(i)}, F\right)}$ is bDD.
We now have to find an expression for $\boldsymbol{Y}$ and show that $\boldsymbol{Y}$ is bDD. Let us write $\boldsymbol{Y}$ in terms of its blocks From

$$
\boldsymbol{Y}=\boldsymbol{M}-\left(\frac{1}{2} \sum_{j \in F} \boldsymbol{L}_{G\left(\boldsymbol{g}^{(i)}\right)}+\frac{1}{2} \sum_{j \in F} \boldsymbol{L}_{G\left(\boldsymbol{d}^{(i)}, F\right)}\right)
$$

it then follows by a simple check that

$$
\begin{equation*}
\boldsymbol{Y}_{[C, F]}^{*}=\boldsymbol{Y}_{[F, C]}=\frac{1}{2}(\boldsymbol{X}-\boldsymbol{D}) \boldsymbol{X}^{-1} \boldsymbol{M}_{[F, C]}=\frac{1}{2}\left(1-\frac{1}{\alpha}\right) \boldsymbol{M}_{[F, C]} \tag{13.33}
\end{equation*}
$$

and that

$$
\begin{equation*}
\boldsymbol{Y}_{[F, F]}=\frac{1}{2} \boldsymbol{X}-\frac{1}{2} \operatorname{Diag}_{i \in F}\left(\boldsymbol{I}_{r} \sum_{j \in F} \sum_{k \in C}\left\|\boldsymbol{A}_{[i, j]} \boldsymbol{X}_{[j, j]}^{-1 / 2}\right\|\left\|\boldsymbol{M}_{[k, j]} \boldsymbol{X}^{-1 / 2}{ }_{[j, j]}\right\|\right) \tag{13.34}
\end{equation*}
$$

and

$$
\begin{align*}
\boldsymbol{Y}_{[C, C]}= & \boldsymbol{M}_{[C, C]}-\frac{1}{2} \operatorname{Diag}_{i \in C}\left(\boldsymbol{M}_{[i, F]} \boldsymbol{X}^{-1} \boldsymbol{M}_{[F, i]}\right) \\
& -\frac{1}{2} \operatorname{Diag}_{i \in C}\left(\boldsymbol{\boldsymbol { I } _ { r }} \sum_{j \in F} \sum_{k \in C \backslash\{i\}}\left\|\boldsymbol{M}_{[i, j]} \boldsymbol{X}_{[j, j]}^{-1 / 2}\right\|\left\|\boldsymbol{M}_{[k, j]} \boldsymbol{X}_{[j, j]}^{-1 / 2}\right\|\right) \\
& -\frac{1}{2} \operatorname{Diag}_{i \in C}\left(\boldsymbol{I}_{r} \sum_{j \in F} \sum_{k \in F}\left\|\boldsymbol{M}_{[i, j]} \boldsymbol{X}_{[j, j]}^{-1 / 2}\right\| \| \boldsymbol{A}_{[k, j]} \boldsymbol{X}^{-1 / 2}\left[{ }_{[j, j]} \|\right) .\right. \tag{13.35}
\end{align*}
$$

Next, we check that $\boldsymbol{Y}$ is bDD. First, let us define, for each block row $i \in F, \rho_{i}=\sum_{j \in F}\left\|\boldsymbol{A}_{[i, j]}\right\|$. From $\boldsymbol{M}_{[F, F]}$ being $\alpha$-bDD, we then conclude $\boldsymbol{M}_{[i, i]} \succcurlyeq \boldsymbol{I}_{r}(1+\alpha) \rho_{i}$. And from $\boldsymbol{M}$ being bDD, we find that $\sum_{j \in C}\left\|\boldsymbol{M}_{[i, j]}\right\| \leq \alpha \rho_{i}$. We now check that each block row $i \in F$ of $\boldsymbol{Y}$ is bDD. The off-diagonal block norm sum is at most $\frac{1}{2}\left(1-\frac{1}{\alpha}\right) \alpha \rho_{i}$.

The diagonal satisfies

$$
\begin{aligned}
\boldsymbol{Y}_{[i, i]} & \succcurlyeq \boldsymbol{I}_{r}\left(\frac{1}{2} \alpha \rho_{i}-\frac{1}{2}\left(\sum_{j \in F} \sum_{k \in C}\left\|\boldsymbol{A}_{[i, j]} \boldsymbol{X}_{[j, j]}^{-1 / 2}\right\|\left\|\boldsymbol{M}_{[k, j]} \boldsymbol{X}_{[j, j]}^{-1 / 2}\right\|\right)\right) \\
& \succcurlyeq \boldsymbol{I}_{r}\left(\frac{1}{2} \alpha \rho_{i}-\frac{1}{2}\left(\sum_{j \in F} \sum_{k \in C} \frac{1}{\alpha \rho_{j}}\left\|\boldsymbol{A}_{[i, j]}\right\|\left\|\boldsymbol{M}_{[k, j]}\right\|\right)\right) \\
& \succcurlyeq \boldsymbol{I}_{r}\left(\frac{1}{2} \alpha \rho_{i}-\frac{1}{2}\left(\sum_{j \in F}\left\|\boldsymbol{A}_{[i, j]}\right\|\right)\right)=\boldsymbol{I}_{r}\left(\frac{1}{2}\left(1-\frac{1}{\alpha}\right) \alpha \rho_{i} .\right)
\end{aligned}
$$

So the block rows with $i \in F$ are bDD.

Next we check the block rows for $i \in C$. Let us define for each $i \in C, \rho_{i}=\sum_{j \in F}\left\|\boldsymbol{M}_{[i, j]}\right\|$. Thus for $\boldsymbol{Y}$, the sum of block norms of row $i$ over columns $j \in F$

$$
\sum_{j \in F}\left\|\boldsymbol{Y}_{[i, j]}\right\|=\frac{1}{2}\left(1-\frac{1}{\alpha}\right) \sum_{j \in F}\left\|\boldsymbol{M}_{[i, j]}\right\|=\frac{1}{2}\left(1-\frac{1}{\alpha}\right) \rho_{i} .
$$

For $\boldsymbol{Y}$, the sum of block norms of row $i$ over columns $j \in C$ is $\sum_{j \in C \backslash\{i\}}\left\|\boldsymbol{M}_{[i, j]}\right\|$. So the total sum of the blocks norms of off-diagonals is

$$
\frac{1}{2} \rho_{i}\left(1-\frac{1}{\alpha}\right)+\sum_{j \in C \backslash\{i\}}\left\|\boldsymbol{M}_{[i, j]}\right\|
$$

The diagonal block satisfies

$$
\begin{aligned}
\boldsymbol{Y}_{[i, i]} & \succcurlyeq \boldsymbol{I}_{r}\left(\rho_{i}-\frac{1}{2}\left\|\boldsymbol{M}_{[i, F]} \boldsymbol{X}^{-1} \boldsymbol{M}_{[F, i]}\right\|\right. \\
& -\frac{1}{2}\left(\sum_{j \in F} \sum_{k \in C \backslash\{i\}}\left\|\boldsymbol{M}_{[i, j]} \boldsymbol{X}_{[j, j]}^{-1 / 2}\right\|\left\|\boldsymbol{M}_{[k, j]} \boldsymbol{X}_{[j, j]}^{-1 / 2}\right\|\right) \\
& \left.-\frac{1}{2}\left(\sum_{j \in F} \sum_{k \in F}\left\|\boldsymbol{M}_{[i, j]} \boldsymbol{X}_{[j, j]}^{-1 / 2}\right\|\left\|\boldsymbol{A}_{[k, j]} \boldsymbol{X}_{[j, j]}^{-1 / 2}\right\|\right)\right) \\
& \succcurlyeq \boldsymbol{I}_{r}\left(\rho_{i}+\sum_{j \in C \backslash\{i\}}\left\|\boldsymbol{M}_{[i, j]}\right\|\right. \\
& -\frac{1}{2}\left(\sum_{j \in F} \sum_{k \in C} \frac{1}{\alpha \rho_{j}}\left\|\boldsymbol{M}_{[i, j]}\right\|\left\|\boldsymbol{M}_{[k, j]}\right\|\right) \\
& \left.-\frac{1}{2}\left(\sum_{j \in F} \sum_{k \in F} \frac{1}{\alpha \rho_{j}}\left\|\boldsymbol{M}_{[i, j]}\right\|\left\|\boldsymbol{A}_{[k, j]}\right\|\right)\right) \\
& \succcurlyeq \boldsymbol{I}_{r}\left(\rho_{i}+\sum_{j \in C \backslash\{i\}}\left\|\boldsymbol{M}_{[i, j]}\right\|-\frac{1}{2} \rho_{i}-\frac{1}{2} \frac{1}{\alpha} \rho_{i}\right) \\
& =\boldsymbol{I}_{r}\left(\frac{1}{2} \rho_{i}\left(1-\frac{1}{\alpha}\right)+\sum_{j \in C \backslash\{i\}}\left\|\boldsymbol{M}_{[i, j]}\right\|\right) .
\end{aligned}
$$

For these block rows are also bDD , and hence $\boldsymbol{Y}$ is bDD . It is clear from the definitions that $\boldsymbol{Y}$ as well as $\boldsymbol{d}^{(i)}$ and $\boldsymbol{g}^{(i)}$ for all $i$ can be computed in $O(m)$ time and $O(\log n)$ depth.

It is easy to verify that if $\boldsymbol{T}$ is a matrix that is only nonzero inside the submatrix $\boldsymbol{T}_{[C, C]}$, then if we apply the transformation of Eq. 13.24 , to $\boldsymbol{M}+\boldsymbol{T}$ instead of $\boldsymbol{M}$, we find $(\boldsymbol{M}+\boldsymbol{T})_{2}^{(\text {last })}=\boldsymbol{M}_{2}^{(\text {last })}+\boldsymbol{T}_{[C, C]}$, and in particular $\boldsymbol{Y}(\boldsymbol{M}+\boldsymbol{T})=\boldsymbol{Y}(\boldsymbol{M})+\boldsymbol{T}_{[C, C]}$.

## ■ 13.11.4 Sparsifying Product Demand Block-Laplacians

In this section, we show how to efficiently sparsify product demand block-Laplacians and their bipartite analogs. We prove the following key lemma later in this section that allows us to transfer results on
graph sparsification to sparsifying these product block-Laplacians.
Lemma 13.11.12. Suppose we have two graphs on $n$ vertices, $H^{(1)}$ and $H^{(2)}$ such that $H^{(1)} \approx_{\varepsilon} H^{(2)}$. Given $\boldsymbol{d} \in\left(\mathbb{C}^{r \times r}\right)^{n}$, define the bDD matrix $\boldsymbol{L}^{(\ell)} \in\left(\mathbb{C}^{r \times r}\right)^{n \times n}$ for each $\ell=1,2$, as

$$
\boldsymbol{L}^{(\ell)}{ }_{[i, j]}= \begin{cases}-\frac{h_{i, j}^{(\ell)}}{\left\|d_{[i]}\right\|\left\|d_{[j]}\right\|} \boldsymbol{d}_{[i]} \boldsymbol{d}_{[j]}^{*} & i \neq j, \\ \boldsymbol{I}_{r} \cdot \sum_{k: k \neq i} h_{i, k}^{(\ell)} & \text { otherwise, }\end{cases}
$$

where $h_{i, j}^{(\ell)}$ denotes the weight of the edge $i, j$ in $H^{(\ell)}$. Then, $\boldsymbol{L}^{(1)} \approx_{\varepsilon} \boldsymbol{L}^{(2)}$.
We now introduce scalar versions of product block-Laplacian matrices that will be useful.
Definition 13.11.13. The product demand graph of a vector $\boldsymbol{d} \in\left(\mathbb{R}_{>0}\right)^{n}, G(\boldsymbol{d})$, is a complete weighted graph on $n$ vertices whose weight between vertices $i$ and $j$ is given by $\boldsymbol{w}_{i j}=\boldsymbol{d}_{i} \boldsymbol{d}_{j}$.

The Laplacian of $G(\boldsymbol{d})$, denoted $\boldsymbol{L}_{G(\boldsymbol{d})}$ is called the product demand Laplacian of $\boldsymbol{d}$.
Definition 13.11.14. The bipartite product demand graph of two vectors $\boldsymbol{d}^{A} \in\left(\mathbb{R}_{>0}\right)^{|A|}, \boldsymbol{d}^{B} \in$ $\left(\mathbb{R}_{>0}\right)^{|B|}, G\left(\boldsymbol{d}^{A}, \boldsymbol{d}^{B}\right)$, is a weighted bipartite graph on vertices $A \cup B$, whose weight between vertices $i \in A$ and $j \in B$ is given by $\boldsymbol{w}_{i j}=\boldsymbol{d}_{i}^{A} \boldsymbol{d}_{j}^{B}$.

The Laplacian of $G\left(\boldsymbol{d}^{A}, \boldsymbol{d}^{B}\right)$, denoted $\boldsymbol{L}_{G\left(\boldsymbol{d}^{A}, \boldsymbol{d}^{B}\right)}$ is called the bipartite product demand Laplacian of $\left(\boldsymbol{d}^{A}, \boldsymbol{d}^{B}\right)$.

In Section 13.14, we give results on efficiently constructing approximations to product demand Laplacians that can be summarized as follows:
Lemma 13.11.1. There is a routine WeightedExpander $(\boldsymbol{d}, \varepsilon)$ such that for any $\varepsilon>0$, and a demand vector $\boldsymbol{d} \in\left(\Re_{>0}\right)^{n}$, WeightedExpander $(\boldsymbol{d}, \varepsilon)$ returns in $O\left(n \varepsilon^{-4}\right)$ work and $O(\log n)$ depth a graph $H$ with $O\left(n \varepsilon^{-4}\right)$ edges such that

$$
\boldsymbol{L}_{H} \approx_{\varepsilon} \boldsymbol{L}_{G(\boldsymbol{d})}
$$

Lemma 13.11.2. There is a routine WeightedBipartiteExpander $\left(\boldsymbol{d}^{A}, \boldsymbol{d}^{B}, \varepsilon\right)$ such that for any demand vectors $\boldsymbol{d}^{A}$ and $\boldsymbol{d}^{B}$ of total length $n$, and a parameter $\varepsilon$, it returns in $O\left(n \varepsilon^{-4}\right)$ work and $O(\log n)$ depth a bipartite graph $H$ between $A$ and $B$ with $O\left(n \varepsilon^{-4}\right)$ edges such that

$$
\boldsymbol{L}_{H} \approx_{\varepsilon} \boldsymbol{L}_{G\left(\boldsymbol{d}^{A}, \boldsymbol{d}^{B}\right)} .
$$

Finally, we need to define an operation on graphs: Given a graph $G$, define $G^{(K 2)}$ to be the graph obtained by duplicating each vertex in $G$, and for each edge $(i, j)$ in $G$, add a $2 \times 2$ bipartite clique between the two copies of $i$ and $j$.

We now combine the above construction of sparsifiers fo product demand graphs with Lemma 13.11.12 to efficiently construct sparse approximations to product demand block-Laplacians.

Proof. (of Lemma 13.11.6) The procedure CliqueSparsification $(\boldsymbol{d}, \varepsilon)$ returns the matrix $\boldsymbol{L}$ where

$$
\boldsymbol{L}_{[i, j]}= \begin{cases}-\frac{h_{i, j}}{\left\|d_{[i]}\right\|\left\|d_{[j]}\right\|} \boldsymbol{d}_{[i]} \boldsymbol{d}_{[j]}^{*} & i \neq j, \\ \boldsymbol{I}_{r} \cdot \sum_{k: k \neq i} h_{i, k} & \text { otherwise. }\end{cases}
$$

By construction $\boldsymbol{L}_{H} \approx_{\boldsymbol{\varepsilon}} \boldsymbol{L}_{G(\boldsymbol{w})}$. Thus, applying Lemma 13.11.12, with $H^{(1)}=H$ and $H^{(2)}=G(\boldsymbol{w})$, we know that $\boldsymbol{L} \approx_{\varepsilon} \boldsymbol{L}^{(2)}$, where $\boldsymbol{L}^{(2)}$ is given by

$$
\boldsymbol{L}_{[i, j]}^{(2)}= \begin{cases}-\frac{\boldsymbol{w}_{i} \boldsymbol{w}_{j}}{\left\|\boldsymbol{d}_{[i]}\right\|\left\|d_{[j]}\right\|} \boldsymbol{d}_{[i]} \boldsymbol{d}_{[j]}^{*} & i \neq j, \\ \boldsymbol{I}_{r} \cdot \boldsymbol{w}_{i} \sum_{k: k \neq i} \boldsymbol{w}_{k} & \text { otherwise. }\end{cases}
$$

```
Algorithm 42: \(\boldsymbol{L}=\) CliqueSparsification \((\boldsymbol{d}, \varepsilon)\)
    Initialize \(\boldsymbol{L} \in\left(\mathbb{C}^{r \times r}\right)^{n \times n}\) to 0 .
    Construct \(\boldsymbol{w}\) where \(\boldsymbol{w}_{i}=\left\|\boldsymbol{d}_{[i]}\right\|\).
    \(H \leftarrow\) WeightedExpander \((\boldsymbol{w}, \varepsilon)\).
    for edge \((i, j) \in H\), with weight \(h_{i, j}\), do
        \(\boldsymbol{L}_{[i, i]} \leftarrow \boldsymbol{L}_{[i, i]}+h_{i, j} \boldsymbol{I}_{r}\).
        \(\boldsymbol{L}_{[j, j]} \leftarrow \boldsymbol{L}_{[j, j]}+h_{i, j} \boldsymbol{I}_{r}\).
        \(\boldsymbol{L}_{[i, j]} \leftarrow \boldsymbol{L}_{[i, j]}-\frac{h_{i, j}}{\boldsymbol{w}_{i} \boldsymbol{w}_{j}} \boldsymbol{d}_{[i]} \boldsymbol{d}_{[j]}^{*}\).
        \(\boldsymbol{L}_{[j, i]} \leftarrow \boldsymbol{L}_{[j, i]}-\frac{h_{i, j}}{\boldsymbol{w}_{i} \boldsymbol{w}_{j}} \boldsymbol{d}_{[j]} \boldsymbol{d}_{[i]}^{*}\).
    end
    Output \(L\).
```

Figure 13-8: Pseudocode for Sparsifying Product Demand Block-Laplacians

```
Algorithm 43: \(\boldsymbol{L}=\) BipartiteCliqueSparsification \((b, F, \varepsilon)\)
    Initialize \(\boldsymbol{L} \in\left(\mathbb{C}^{r \times r}\right)^{n \times n}\) to 0 .
    Construct \(\boldsymbol{w}\) where \(\boldsymbol{w}_{i}=\left\|\boldsymbol{d}_{[i]}\right\|\). Let \(C=V \backslash F\).
    \(H \leftarrow\) WeightedBipartiteExpander \(\left(\left.\boldsymbol{w}\right|_{F},\left.\boldsymbol{w}\right|_{C}, \varepsilon\right)\).
    for edge \((i, j) \in H\), with weight \(h_{i, j}\), do
        \(\boldsymbol{L}_{[i, i]} \leftarrow \boldsymbol{L}_{[i, i]}+h_{i, j} \boldsymbol{I}_{r}\).
        \(\boldsymbol{L}_{[j, j]} \leftarrow \boldsymbol{L}_{[j, j]}+h_{i, j} \boldsymbol{I}_{r}\).
        \(\boldsymbol{L}_{[i, j]} \leftarrow \boldsymbol{L}_{[i, j]}-\frac{h_{i, j}}{\boldsymbol{w}_{i} \boldsymbol{w}_{j}} \boldsymbol{d}_{[i]} \boldsymbol{d}_{[j]}^{*}\).
        \(\boldsymbol{L}_{[j, i]} \leftarrow \boldsymbol{L}_{[j, i]}-\frac{h_{i, j}}{\boldsymbol{w}_{i} \boldsymbol{w}_{j}} \boldsymbol{d}_{[j]} \boldsymbol{d}_{[i]}^{*}\).
    end
    Output \(L\).
```

Figure 13-9: Pseudocode for Sparsifying Bipartite Product Demand Block-Laplacians

Since $\boldsymbol{w}_{i}=\left\|\boldsymbol{d}_{[i]}\right\|$, we have $\boldsymbol{L}^{(2)}=\boldsymbol{L}_{G(\boldsymbol{d})}$, proving our claim.

Proof. (of Lemma 13.11.7) The procedure BipartiteCliqueSparsification $(\boldsymbol{d}, F, \varepsilon)$ returns the matrix $L$ where

$$
\boldsymbol{L}_{[i, j]}= \begin{cases}-\frac{h_{i, j}}{\left\|d_{[i]}\right\|\left\|d_{[j]}\right\|} \boldsymbol{d}_{[i]} \boldsymbol{d}_{[j]}^{*} & i \neq j, \\ \boldsymbol{I}_{r} \cdot \sum_{k: k \neq i} h_{i, k} & \text { otherwise. }\end{cases}
$$

Since $H$ returned by WeightedBipartiteExpander is guaranteed to be bipartite using Lemma 13.11.2, we obtain that $\left(\boldsymbol{L}_{H}\right)_{[F, F]}$, and $\left(\boldsymbol{L}_{H}\right)_{[C, C]}$ are block-diagonal.

By construction $\boldsymbol{L}_{H} \approx_{\varepsilon} \boldsymbol{L}_{G\left(\left.\boldsymbol{w}\right|_{F},\left.\boldsymbol{w}\right|_{C}\right)}$. Thus, applying Lemma 13.11.12, with $H^{(1)}=H$ and $H^{(2)}=$ $G\left(\left.\boldsymbol{w}\right|_{F},\left.\boldsymbol{w}\right|_{C}\right)$, we know that $\boldsymbol{L} \approx_{\varepsilon} \boldsymbol{L}^{(2)}$, where $\boldsymbol{L}^{(2)}$ is given by

$$
\boldsymbol{L}^{(2)}{ }_{[i, j]}= \begin{cases}\boldsymbol{I}_{r} \cdot \boldsymbol{w}_{i} \sum_{k \in C} \boldsymbol{w}_{k} & i=j, \text { and } i \in F, \\ \boldsymbol{I}_{r} \cdot \boldsymbol{w}_{i} \sum_{k \in F} \boldsymbol{w}_{k} & i=j, \text { and } i \in C, \\ 0 & i \neq j, \text { and }(i, j \in F \text { or } i, j \in C), \\ -\frac{w_{i} \boldsymbol{w}_{j}}{\left\|d_{[i]]}\right\|\left\|d_{[j]]}\right\|} \boldsymbol{d}_{[i]} \boldsymbol{d}_{[j]}^{*} & i \neq j, \text { otherwise. }\end{cases}
$$

Since $\boldsymbol{w}_{i}=\left\|\boldsymbol{d}_{[i]}\right\|$, we have $\boldsymbol{L}^{(2)}=\boldsymbol{L}_{G\left(\left.\boldsymbol{d}\right|_{F},\left.\boldsymbol{d}\right|_{C}\right)}$, proving our claim.
Finally, we give a proof of Lemma 13.11.12.

Proof. (of Lemma 13.11.12) Using Fact 13.5.2, we can write each $\boldsymbol{d}_{[i]}$ as $\frac{1}{2}\left\|\boldsymbol{d}_{[i]}\right\|\left(\boldsymbol{Q}_{i}^{(1)}+\boldsymbol{Q}_{i}^{(2)}\right)$, where $\boldsymbol{Q}_{i}^{(1)}\left(\boldsymbol{Q}_{i}^{(1)}\right)^{*}=\boldsymbol{Q}_{i}^{(2)}\left(\boldsymbol{Q}_{i}^{(2)}\right)^{*}=\boldsymbol{I}_{r}$. Construct the matrix $\boldsymbol{F} \in\left(\mathbb{C}^{r \times r}\right)^{2 n \times 2 n}$ as follows:

$$
\boldsymbol{F}=\left[\begin{array}{ccccc}
\boldsymbol{Q}_{1}^{(1)} & \boldsymbol{Q}_{1}^{(2)} & 0 & 0 & \cdots \\
0 & 0 & \boldsymbol{Q}_{2}^{(1)} & \boldsymbol{Q}_{2}^{(2)} & \cdots \\
\vdots & \vdots & & & \ddots
\end{array}\right]
$$

Now, observe that for each $\ell=1,2$,

$$
\boldsymbol{L}^{(\ell)}{ }_{[i, i]}=\frac{1}{2}\left(\sum_{k: k \neq i} h_{i, k}^{(\ell)}\right) \cdot\left(\boldsymbol{Q}_{i}^{(1)}\left(\boldsymbol{Q}_{i}^{(1)}\right)^{*}+\boldsymbol{Q}_{i}^{(2)}\left(\boldsymbol{Q}_{i}^{(2)}\right)^{*}\right)=\frac{1}{4}\left(\boldsymbol{F}\left(\boldsymbol{L}_{\left(H^{(\ell)}\right)^{(K 2)}} \otimes \boldsymbol{I}_{r}\right) \boldsymbol{F}^{*}\right)_{[i, i]},
$$

where $\boldsymbol{L}_{\left(H^{(\ell)}\right)^{(K 2)}}$ is the Laplacian of the graph $\left(H^{(\ell)}\right)^{(K 2)}$ defined above, and $\otimes$ tensor product. Also, for $i \neq j$,

$$
\boldsymbol{L}^{(\ell)}{ }_{[i, j]}=-\frac{h_{i, j}^{(\ell)}}{\boldsymbol{w}_{i} \boldsymbol{w}_{j}} \boldsymbol{d}_{[i]} \boldsymbol{d}_{[j]}^{*}=-\frac{h_{i, j}^{(\ell)}}{4} \cdot\left(\boldsymbol{Q}_{i}^{(1)}+\boldsymbol{Q}_{i}^{(2)}\right)\left(\boldsymbol{Q}_{j}^{(1)}+\boldsymbol{Q}_{j}^{(2)}\right)^{*}=\frac{1}{4}\left(\boldsymbol{F}\left(\boldsymbol{L}_{\left(H^{(\ell)}\right)^{(K 2)}} \otimes \boldsymbol{I}_{r}\right) \boldsymbol{F}^{*}\right)_{[i, j]},
$$

Thus,

$$
\boldsymbol{L}^{(\ell)}=\frac{1}{4} \boldsymbol{F}\left(\boldsymbol{L}_{\left(H^{(\ell)}\right)^{(K 2)}} \otimes \boldsymbol{I}_{r}\right) \boldsymbol{F}^{*} .
$$

By assumption, we have $\boldsymbol{L}_{H^{(1)}} \approx_{\varepsilon} \boldsymbol{L}_{H^{(2)}}$. Using Lemma 13.11.15, we get that $\boldsymbol{L}_{H^{(1)}\left(K^{2}\right)} \approx_{\varepsilon} L_{H^{(2)}\left(K^{2}\right)}$. This implies

$$
\boldsymbol{L}_{\left(H^{(1)}\right)^{(K 2)}} \otimes \boldsymbol{I}_{r} \approx_{\varepsilon} \boldsymbol{L}_{\left(H^{(2)}\right)^{(K 2)}} \otimes \boldsymbol{I}_{r},
$$

and thus, $\boldsymbol{L}^{(1)} \approx_{\varepsilon} \boldsymbol{L}^{(2)}$.

### 13.11.5 Constructing sparsifiers for lifts of graphs

Given a graph $G$, define $G^{(K 2)}$ to be the graph obtained by duplicating each vertex in $G$, and for each edge $(i, j)$ in $G$, add a $2 \times 2$ bipartite clique between the two copies of $i$ and $j$.

Lemma 13.11.15. If $H$ is a sparsifier for $G$, i.e., $L_{G} \approx_{\varepsilon} L_{H}$, then $H^{(K 2)}$ is a sparsifier for $G^{(K 2)}$, i.e. $L_{G^{(K 2)}} \approx_{\varepsilon} L_{H^{(K 2)}}$.

Proof. Since $L_{G} \approx_{\varepsilon} L_{H}$, we have $e_{i}^{\top} L_{G} e_{i} \approx_{\varepsilon} e_{i}^{\top} L_{H} e_{i}$. Thus, if $D_{G}$ denotes the diagonal matrix of degrees of $G$, we have $D_{G} \approx_{\varepsilon} D_{H}$. The Laplacian for $G^{(K 2)}$ is

$$
L_{G^{(K 2)}}=L_{G} \otimes\left[\begin{array}{ll}
1 & 1 \\
1 & 1
\end{array}\right]+D_{G} \otimes\left[\begin{array}{cc}
1 & -1 \\
-1 & 1
\end{array}\right] .
$$

Similarly,

$$
L_{H^{(K 2)}}=L_{H} \otimes\left[\begin{array}{ll}
1 & 1 \\
1 & 1
\end{array}\right]+D_{H} \otimes\left[\begin{array}{cc}
1 & -1 \\
-1 & 1
\end{array}\right] .
$$

Observe that we can write

$$
L_{G} \otimes\left[\begin{array}{cc}
1 & 1 \\
1 & 1
\end{array}\right]=\left[\begin{array}{ccccc}
1 & 1 & 0 & 0 & \ldots \\
0 & 0 & 1 & 1 & \ldots \\
\vdots & \vdots & & \ddots &
\end{array}\right] L_{G}\left[\begin{array}{ccccc}
1 & 1 & 0 & 0 & \ldots \\
0 & 0 & 1 & 1 & \ldots \\
\vdots & \vdots & & \ddots &
\end{array}\right]^{\top}
$$

Since $L_{G} \approx_{\varepsilon} L_{H}$, this implies

$$
L_{G} \otimes\left[\begin{array}{ll}
1 & 1 \\
1 & 1
\end{array}\right] \approx_{\varepsilon} L_{H} \otimes\left[\begin{array}{ll}
1 & 1 \\
1 & 1
\end{array}\right]
$$

Similarly, we get,

$$
D_{G} \otimes\left[\begin{array}{cc}
1 & -1 \\
-1 & 1
\end{array}\right] \approx_{\varepsilon} D_{H} \otimes\left[\begin{array}{cc}
1 & -1 \\
-1 & 1
\end{array}\right]
$$

Adding the above two, we get $L_{G^{(K 2)}} \approx_{\varepsilon} L_{H^{(K 2)}}$.

## ■ 13.12 Estimating Leverage Scores by Undersampling

We will control the densities of all intermediate bDD matrices using the uniform sampling technique introduced in Chapter 3. It relies on sampling columns ${ }^{2}$ of matrices by upper bounds of their true leverage scores,

$$
\boldsymbol{\tau}_{i}(\boldsymbol{A})=\boldsymbol{a}_{i}^{T}\left(\boldsymbol{A} \boldsymbol{A}^{*}\right)^{-1} \boldsymbol{a}_{i}
$$

These upper bounds are measured w.r.t. a different matrix, giving generalized leverage scores of the form:

$$
\tau_{i}^{\boldsymbol{B}}(\boldsymbol{A})= \begin{cases}\boldsymbol{a}_{i}^{*}\left(\boldsymbol{B} \boldsymbol{B}^{*}\right)^{-1} \boldsymbol{a}_{i} & \text { if } \boldsymbol{a}_{i} \perp \operatorname{ker}(\boldsymbol{B}), \\ 1 & \text { otherwise }\end{cases}
$$

We introduced unitary edge-vertex transfer matrices in Definition 13.6.4. Sparsifying bDD matrices can be transformed into the more general setting described above via a unitary edge-vertex transfer matrix, which is analogous to the edge-vertex incidence matrix.

Lemma 13.6.5 proves that every bDD matrix $\boldsymbol{M} \in\left(\mathbb{C}^{r \times r}\right)^{n \times n}$ with $m$ nonzero off-diagonal blocks can be written as $\boldsymbol{X}+\boldsymbol{B} \boldsymbol{B}^{*}$ where $\boldsymbol{B} \in\left(\mathbb{C}^{r \times r}\right)^{n \times 2 m}$ is a unitary edge-vertex transfer matrix and $\boldsymbol{X}$ is a block diagonal PSD matrix. Additionally, for every block diagonal matrix $\boldsymbol{Y}$ s.t. $\boldsymbol{M}-\boldsymbol{Y}$ is bDD, we have $\boldsymbol{X} \succcurlyeq \boldsymbol{Y}$. This decomposition can be found in $O(m)$ time and $O(\log n)$ depth. We rely on this $\boldsymbol{X}$ to detect some cases of high leverage scores as samples of $\boldsymbol{B}$ may have lower rank.

We will reduce the number of nonzero blocks in $\boldsymbol{M}$ by sampling columns blocks from this matrix. This is more restrictive than sampling individual columns. Nonetheless, it can be checked via matrix concentration bounds [3, 251] that it suffices to sample the block by analogs of leverage scores:

$$
\begin{equation*}
\boldsymbol{\tau}_{[i]}=\operatorname{tr}\left(\boldsymbol{B}_{[i]}^{*}\left(\boldsymbol{X}+\boldsymbol{B} \boldsymbol{B}^{*}\right)^{-1} \boldsymbol{B}_{[i]}\right) \tag{13.36}
\end{equation*}
$$

As in Chapter 3, we recursively estimate upper bounds for these scores, leading to the pseudocode given in Figure 13-10.
Lemma 13.12.1. Given a positive definite bDD matrix $M \in\left(\mathbb{C}^{r \times r}\right)^{n \times n}$ with $m$ nonzero blocks.
Assume that for any positive definite bDD matrix $\boldsymbol{M}^{\prime} \in\left(\mathbb{C}^{r \times r}\right)^{n \times n}$ with $\widehat{m}$ nonzero blocks, we can find an implicit representation of a matrix $\boldsymbol{W}$ such that $\boldsymbol{W} \approx_{1}\left(\boldsymbol{M}^{\prime}\right)^{-1}$ in depth $d_{\text {con }}(\widehat{m}, n)$ and work $w_{\text {con }}(\widehat{m}, n)$ and for any vector $b$, we can evaluate $\boldsymbol{W} \boldsymbol{b}$ in depth $d_{\text {eval }}(\widehat{m}, n)$ and work $w_{\text {eval }}(\widehat{m}, n)$.

[^33]Algorithm 44: $\widetilde{\boldsymbol{M}}=\operatorname{SpARSIFY}(\boldsymbol{M}, \varepsilon, K)$
Write $\boldsymbol{M}$ into the form $\boldsymbol{X}+\boldsymbol{B} \boldsymbol{B}^{*}$ where $\boldsymbol{X}$ is a block diagonal positive definite matrix $\boldsymbol{X}$ and $\boldsymbol{B}$ is an unitary edge-vertex transfer matrix.
Uniformly sample $\frac{m}{K}$ column blocks of $\boldsymbol{B}$ to form $\boldsymbol{C}$.
Find an implicit representation of $\left(\boldsymbol{X}+\boldsymbol{C} \boldsymbol{C}^{*}\right)^{-1}, \boldsymbol{W}$.
Approximate $\operatorname{tr}\left(\boldsymbol{B}_{[i]}^{*}\left(\boldsymbol{X}+\boldsymbol{C} \boldsymbol{C}^{*}\right)^{-1} \boldsymbol{B}_{[i]}\right)$ using Johnson-Lindenstrauss projections.
This involves applying $\boldsymbol{W}$ to $O(\log n)$ random vectors.
Use these estimates to sample blocks of $\boldsymbol{B}$ to form $\tilde{\boldsymbol{B}}$.
Return $\boldsymbol{X}+\tilde{\boldsymbol{B}} \tilde{\boldsymbol{B}}^{*}$.
Figure 13-10: Sparsification via. Uniform Sampling

For any $K \geq 1,1 \geq \varepsilon>0$, the algorithm $\operatorname{Sparsify}(M, \varepsilon, K)$ outputs an explicit positive definite bDD matrix $\widetilde{\boldsymbol{M}}$ with $O\left(K n \log n / \varepsilon^{2}\right)$ nonzero blocks and $\widetilde{\boldsymbol{M}} \approx_{\varepsilon} \boldsymbol{M}$.

Also, this algorithm runs in

$$
d_{\text {con }}\left(\frac{m}{K}, n\right)+O\left(d_{\text {eval }}\left(\frac{m}{K}, n\right)+\log n\right)
$$

depth and

$$
w_{\text {con }}\left(\frac{m}{K}, n\right)+O\left(w_{\text {eval }}\left(\frac{m}{K}, n\right) \log n+m \log n\right)
$$

work.
The guarantees of this process can be obtained from the main result from Chapter 3:
Theorem 13.12.2. Let $\boldsymbol{A}$ be a $n$ by matrix, and $\alpha \in(0,1]$ a density parameter. Consider the matrix $\boldsymbol{A}^{\prime}$ consisting of a random $\alpha$ fraction of the columns of $\boldsymbol{A}$. Then, with high probability we have $\boldsymbol{u}_{i}=\tau_{i}^{\boldsymbol{A}^{\prime}}(\boldsymbol{A})$ is a vector of leverage score overestimates, i.e. $\tau_{i}(\boldsymbol{A}) \leq \boldsymbol{u}_{i}$, and

$$
\sum_{i=1}^{m} \boldsymbol{u}_{i} \leq \frac{3 n}{\alpha}
$$

Proof. (Sketch of Lemma 13.12.1) Instead of sampling $m / K$ column blocks, consider sampling $m / K$ individual columns to form $\widehat{\boldsymbol{C}}$. Theorem 13.12.2 gives that the total leverage scores of all the columns of $\boldsymbol{B}$ w.r.t. $\widehat{\boldsymbol{C}}$ is at most $O\left(n K r^{2}\right)$.

As $\boldsymbol{C}$ is formed by taking blocks instead of columns, we have

$$
X+C C^{*} \succcurlyeq X+\widehat{C} \widehat{C}^{*}
$$

so the individual leverage scores computed w.r.t. $\boldsymbol{C}$ (and in turn $\boldsymbol{W}$ ) sums up to less than the ones computed w.r.t. $\widehat{\boldsymbol{C}}$.

Let us use $\boldsymbol{b}_{[i], j} \in \mathbb{C}^{n}$ to denote the $j^{\text {th }}$ column of the $\boldsymbol{B}_{[i]}$ block column. Note that

$$
\boldsymbol{B}_{[i]} \boldsymbol{B}_{[i]}^{*}=\sum_{j} \boldsymbol{b}_{[i], j} \boldsymbol{b}_{[i], j}^{*}
$$

Define $\tau_{[i], j}^{\boldsymbol{X}+\boldsymbol{C} \boldsymbol{C}^{*}}=\operatorname{tr}\left(\boldsymbol{b}_{[i], j}^{*}\left(\boldsymbol{X}+\boldsymbol{C} \boldsymbol{C}^{*}\right)^{-1} \boldsymbol{b}_{[i], j}\right)$.
We then have

$$
\operatorname{tr}\left(\boldsymbol{B}_{[i]}^{*}\left(\boldsymbol{X}+\boldsymbol{C} \boldsymbol{C}^{*}\right)^{-1} \boldsymbol{B}_{[i]}\right)=\sum_{j} \tau_{[i], j}^{\boldsymbol{X}+\boldsymbol{C} \boldsymbol{C}^{*}}
$$

so the total sum of the estimates we obtain is at most $O\left(n r^{2} K\right)$, which gives $\tilde{\boldsymbol{B}}$ with $O\left(K n \log n \varepsilon^{-2}\right)$ blocks.

To approximate $\sum_{j} \tau_{[i], j}^{\boldsymbol{X}+\boldsymbol{C C}} \boldsymbol{C}^{*}$, we apply a standard technique given by [242, 23, 170]. The rough idea is to write

$$
\tau_{[i], j}^{\boldsymbol{X}+\boldsymbol{C} C^{*}}=\left\|\boldsymbol{C}^{*}\left(\boldsymbol{X}+\boldsymbol{C} \boldsymbol{C}^{*}\right)^{-1} \boldsymbol{b}_{[i], j}\right\|^{2}+\left\|\sqrt{\boldsymbol{X}}\left(\boldsymbol{X}+\boldsymbol{C} \boldsymbol{C}^{*}\right)^{-1} \boldsymbol{b}_{[i], j}\right\|^{2}
$$

By Johnson-Lindenstrauss Lemma, for a random Gaussian matrix $\boldsymbol{G}$ with $\Theta(\log (n))$ rows, we know that

$$
\tau_{[i], j}^{\boldsymbol{X}+\boldsymbol{C} C^{*}} \approx_{2}\left\|\boldsymbol{G} \boldsymbol{C}^{*}\left(\boldsymbol{X}+\boldsymbol{C} \boldsymbol{C}^{*}\right)^{-1} \boldsymbol{b}_{[i], j}\right\|^{2}+\left\|\boldsymbol{G} \sqrt{\boldsymbol{X}}\left(\boldsymbol{X}+\boldsymbol{C} \boldsymbol{C}^{*}\right)^{-1} \boldsymbol{b}_{[i], j}\right\|^{2}
$$

for high probability. Since $\boldsymbol{G}$ has $O(\log (n))$ rows, this can be approximated by applying the approximate inverse $\boldsymbol{W}$ to $O(\log (n))$ vectors. Hence, step 3 runs in $O\left(d_{\text {eval }}\left(\frac{m}{K}, n\right)+\log n\right)$ depth and $O\left(w_{\text {eval }}\left(\frac{m}{K}, n\right) \log n+m \log n\right)$ work.

We remark that the extra factor of $r^{2}$ can likely be improved by modifying the proof of Theorem 13.12 .2 to work with blocks. We omit this improvement for simplicity, especially since we're already treating $r$ as a constant.

## ■ 13.13 Recursive Construction of Schur Complement Chains

We now give the full details for the recursive algorithm that proves the running times as stated in Theorem 13.1.2:

Theorem 13.1.2 (Sparsified Multigrid). There is an algorithm that, when given a bDD matrix $\boldsymbol{M}$ with $n$ block rows and $m$ nonzero blocks, produces a solver for $\boldsymbol{M}$ in $O\left(m \log n+n \log ^{2+o(1)} n\right)$ work and $O\left(n^{o(1)}\right)$ depth so that the solver finds $\varepsilon$-approximate solutions to systems of equations in $M$ in $O\left(\left(m+n \log ^{1+o(1)} n\right) \log (1 / \varepsilon)\right)$ work and $O\left(\log ^{2} n \log \log n \log (1 / \varepsilon)\right)$ depth.

This argument can be viewed as a more sophisticated version of the one in Section 13.3.6. The main idea is to invoke routines for reducing edges and vertices in a slightly unbalanced recursion, where several steps of vertex reductions take place before a single edge reduction. The routines that we will call are:

1. BDDSUbSET given by Lemma 13.3.2 proven in Section 13.8 for finding a large set of $\alpha$-bDD subset.
2. ApproxSchur given by Lemma 13.11 .1 proven in Section 13.11 that gives sparse approximations to Schur complements.
3. Sparsify given by Lemma 13.12 .1 proven in Section 13.12 that allows us to sparsify a bDD matrix by recursing on a uniform subsample of its non-zero blocks.

An additional level of complication comes from the fact that the approximation guarantees of our constructions rely on gradually decreasing errors down the Schur complement chain. This means that the density increases faster and larger reduction factors $K$ are required as the iteration goes on. Pseudocode of our algorithm is given in Figure 13-11.

Note that since $\boldsymbol{M}^{(0)}$ may be initially dense, we first make a recursive call to SPARSIFY before computing the approximate Schur complements.

In our analysis, we will use $n^{(i)}$ to denote the number of non-zero column/row blocks in $M^{(i)}$, and $m^{(i)}$ to denote the number of non-zero blocks in $\boldsymbol{M}^{(i)}$. These are analogous to dimension and

```
Algorithm 45: \(\left(\boldsymbol{M}^{(1)}, \boldsymbol{M}^{(2)}, \cdots ; F_{1}, F_{2}, \cdots\right)=\) RecursiveConstruct \(\left(\boldsymbol{M}^{(0)}\right)\)
    Set \(i \leftarrow 0\), and \(k\) and \(c\) are parameters to be decided.
    while \(M^{(i)}\) has more than \(\Theta(1)\) blocks do
        If \(i \bmod k=0\), Then \(\boldsymbol{M}^{(i)} \leftarrow \operatorname{Sparsify}\left(\boldsymbol{M}^{(i)},(i+8)^{-2}, 2^{2 c k \log ^{2}(i+1)}\right)\).
        \(F_{i+1} \leftarrow \operatorname{BDDSUBSET}\left(M^{(i)}, 4\right)\).
        \(\boldsymbol{M}^{(i+1)} \leftarrow \operatorname{ApproxSchur}\left(\boldsymbol{M}^{(i)}, F_{i+1}, 4,(i+8)^{-2}\right)\).
        \(i \leftarrow i+1\).
    end
```

Figure 13-11: Pseudocode for Recursively Constructing Schur Complement Chains, the recursive calls happen through Sparsify, which constructs its own Schur Complement Chains to perform the sparsification
number of non-zeros in the matrix. It is also useful to refer to the steps between calls to Sparsify as phases. Specifically, phase $j$ consists of iterations $(j-1) k$ to $j k-1$. We will also use $K_{j}$ to denote the reduction factor used to perform the sparsification at the start of phase $j$, aka.

$$
K_{j}=2^{2 c k \log ^{2}((j-1) k+1)}
$$

A further technicality is that Lemma 13.12.1 require a strictly positive definite block-diagonal part to facilitate the detection of vectors in the null space of the sample. We do so by padding $M$ with a small copy of the identity, and will check below that this copy stays throughout the course of this algorithm.

Lemma 13.13.1. For any bDD matrix $\boldsymbol{M}^{(0)}$ expressible as $\xi \boldsymbol{I}+\boldsymbol{B}^{(0)}\left(\boldsymbol{B}^{(0)}\right)^{*}$ for some $\xi>0$ and unitary edge-vertex transfer matrix $\boldsymbol{B}$, all intermediate matrices $\boldsymbol{M}^{\prime}$ generated during the algorithm RecursiveConstruct $\left(\boldsymbol{M}^{(0)}\right)$ can be expressed as $\xi \boldsymbol{I}+\boldsymbol{B}^{\prime}\left(\boldsymbol{B}^{\prime}\right)^{*}$.

Proof. There are two mechanisms by which this recursive algorithm generates new matrices: through uniform sampling within Sparsify and via ApproxSchur. We will show inductively down the algorithmic calls that all matrices satisfy this property.

Lemma 13.12 .1 gives that this $\boldsymbol{X}$ is preserved in the sample.
For calls to ApproxSchur $(\boldsymbol{M}, C)$, the inductive hypothesis gives that $\boldsymbol{M}_{[C, C]}$ is expressible as $\widehat{\boldsymbol{M}}+\xi \boldsymbol{I}_{C}$. The last condition in Lemma 13.3.4 then gives that the result also has a $\xi \boldsymbol{I}_{C}$ part, which gives the inductive hypothesis for it as well.

As the interaction between this padding and our recursion is minimal, we will simply state the input conditions as $\boldsymbol{M}=\xi \boldsymbol{I}+\boldsymbol{B} \boldsymbol{B}^{*}$ for some $\xi>0$. We will first bound the sizes of the matrices in the Schur complement chain produced by RecursiveConstruct.

Lemma 13.13.2. For any $b D D$ matrix $\boldsymbol{M}^{(0)}=\xi \boldsymbol{I}+\boldsymbol{B}^{(0)}\left(\boldsymbol{B}^{(0)}\right)^{*}$ for some $\xi>0$, the algorithm RECURSIVECONSTRUCT $\left(\boldsymbol{M}^{(0)}\right)$ returns a Schur complement chain $\left(\boldsymbol{M}^{(1)}, \boldsymbol{M}^{(2)}, \cdots ; F_{1}, F_{2}, \cdots\right)$ such that:

1. $n^{(i)} \leq \beta^{i} n^{(0)}$ for some absolute constant $\beta<1$.
2. The number of non-zero blocks in any iteration $i$ of phase $j$ is at most $2^{3 c k \log ^{2}(j k)} n^{((j-1) k)} \log n^{((j-1) k)}$.

Proof. Lemma 13.3.2 and the choice of $\alpha=4$ ensures $\left|F_{k}\right|=\Omega\left(n^{(i)}\right)$, which means there is constant $\beta<1$ such that $n^{(i)} \leq \beta^{k} n$. Furthermore, we do not increase vertex count at any point in this recursion, and all recursive calls are made to smaller graphs. Therefore, the recursion terminates.

Lemma 13.11 .1 shows that after computing each approximate Schur complement,

$$
\begin{aligned}
m^{(i+1)} & =O\left(m^{(i)}\left(i^{2} \log (i+8)\right)^{O(\log (i+8))}\right) \\
& \leq 2^{O\left(\log ^{2}(i+1)\right)} m^{(i)}
\end{aligned}
$$

Hence, by picking $c$ appropriately we can guarantee that the density increases by a factor of most $2^{c \log ^{2}(i+1)}$ during each iteration. This size increase is controlled by calls to SPARSIFY. Specifically, Lemma 13.12 .1 gives that at the start of phase $j$ we have:

$$
\frac{m^{((j-1) k)}}{n^{((j-1) k)} \log n^{((j-1) k)}} \leq 2^{2 c k \log ^{2}((j-1) k)}
$$

Then, since we go at most $k$ steps without calling Sparsify, this increase in density can be bounded by:

$$
2^{2 c k \log ^{2}((j-1) k)} 2^{c \log ^{2}(j k-1)} \cdots 2^{c \log ^{2}((j-1) k+1)} \leq 2^{3 c k \log ^{2}(j k)}
$$

Lemma 13.13.3. For any bDD matrix $\boldsymbol{M}^{(0)}=\xi \boldsymbol{I}+\boldsymbol{B}^{(0)}\left(\boldsymbol{B}^{(0)}\right)^{*}$ for some $\xi>0$ that has $n$ blocks, the algorithm RECURSIVEConstruct $\left(\boldsymbol{M}^{(0)}\right)$ returns a Schur Complement Chain $\left(\boldsymbol{M}^{(1)}, \ldots ; F_{1}, \ldots\right.$ ) whose corresponding the linear operator $\boldsymbol{W}$ satisfies

$$
\boldsymbol{W} \approx_{O(1)}\left(\boldsymbol{M}^{(0)}\right)^{-1}
$$

Also, we can evaluate $\boldsymbol{W} \boldsymbol{b}$ in $O\left(\log ^{2} n \log \log n\right)$ depth and $2^{O\left(k \log ^{2} k\right)} n \log n$ work for any vector $\boldsymbol{b}$.
Proof. We first bound the quality of approximation between $\boldsymbol{W}$ and $\left(\boldsymbol{M}^{(0)}\right)^{-1}$. The approximate Schur complement $\boldsymbol{M}^{(i+1)}$ was constructed so that $\boldsymbol{M}^{(i+1)} \approx_{(i+8)^{-2}} S c\left(\boldsymbol{M}^{(i)}, F_{i}\right)$. The other source of error, SpARSIFY, is called only for some $i$. In those iterations, Lemma 13.12 .1 guarantee that $\boldsymbol{M}^{(i)}$ changes only by $(i+8)^{-2}$ factor. This means overall we have $\boldsymbol{M}^{(i+1)} \approx_{2(i+8)^{-2}} S c\left(\boldsymbol{M}^{(i)}, F_{i}\right)$. By Lemma 13.3.1, we have:

$$
\boldsymbol{W} \approx_{1 / 2+4 \sum_{i}(i+8)^{-2}}\left(\boldsymbol{M}^{(0)}\right)^{-1}
$$

and it can be checked that $\sum_{i}(i+8)^{-2}$ is a constant.
The cost of ApplyChain is dominated by the sequence of calls to Jacobi. As each $F_{i}$ is chosen to be 4-bDD, the number of iterations required is $O\left(\log \left(\varepsilon_{i}\right)\right)=O(\log i)$. As matrix-vector multiplications take $O(\log n)$ depth, the total depth can be bounded by.

$$
O\left(\sum_{i} \log i \log n^{(i)}\right)=O\left(\log ^{2} n \log \log n\right)
$$

The total work of these steps depend on the number of non-zero blocks, $m^{(i)}$. Substituting the bounds from Lemma 13.13.2 into Jacobi gives a total of:

$$
O\left(\sum_{i} 2^{3 c k \log ^{2}(i+k)} n^{(i)} \log n^{(i)} \log i\right) \leq 2^{O\left(k \log ^{2} k\right)} n \log n
$$

where the inequality follows from the fact that the $n^{(i)}$ s are geometrically decreasing.

This allows us to view the additional overhead of SpARSIFY as a black box, and analyze the total cost incurred by the non-recursive parts of RECURSIVECONSTRUCT during each phase.

Lemma 13.13.4. The total cost of RECURSIVECONSTRUCT during phase $j$, including the linear system solves made by Sparsify at iteration $(j-1) k$ (but not its recursive invocation to RECURSIVEConstruct) is

$$
2^{O\left(k \log ^{2}(j k)\right)} n^{((j-1) k)} \log ^{2} n^{((j-1) k)}
$$

in work and

$$
O\left(k \log (j k) \log ^{2} n^{((j-1) k)}\right)+O\left(\log ^{2} n^{((j-1) k)} \log \log n^{((j-1) k)}\right)
$$

in depth.
Proof. Let $m^{(i)}$ and $n^{(i)}$ be the number of non zeros and variables in $\boldsymbol{M}^{(i)}$ before the Sparsify call if there is. Lemmas 13.3 .2 and 13.11 .1 show that the $i^{\text {th }}$ iteration takes $O\left(m^{(i)}+m^{(i+1)}\right)$ work and $O\left(\log i \log n^{(i)}\right)$ depth. By Lemma 13.13 .2 the cost during these iterations excluding the call to SPARSIFY is:

$$
\begin{aligned}
\sum_{i=(j-1) k}^{j k-1} O\left(m^{(i)}+m^{(i+1)}\right) & \leq \sum_{i=(j-1) k}^{j k-1} 2^{O\left(k \log ^{2} i\right)} n^{(i)} \log n^{(i)} \\
& \leq 2^{O\left(k \log ^{2}(j k)\right)} n^{((j-1) k)} \log n^{((j-1) k)}
\end{aligned}
$$

work and

$$
\sum_{i=(j-1) k}^{j k-1} O\left(\log i \log n^{(i)}\right) \leq O\left(k \log ((j-1) k) \log n^{((j-1) k)}\right)
$$

depth.
We now consider the call to Sparsify made during iteration $(j-1) k$. Access to a fast solver for the sampled bDD matrix is obtained via recursive calls to RECURSIVEConstruct. The guarantees of the chain given by Lemma 13.13.3 above means each solve takes depth

$$
d_{\text {eval }}=\log ^{2} n^{((j-1) k)} \log \log n^{((j-1) k)}
$$

and work

$$
w_{\text {eval }}=n^{((j-1) k)} \log n^{((j-1) k)}
$$

Incorporating these parameters into Lemma 13.12 .1 allows us to bound the overhead from these solves by

$$
2^{\left.O\left(k \log ^{2}(j k)\right)\right)} n^{((j-1) k)} \log ^{2} n^{((j-1) k)}
$$

work and

$$
O\left(k \log (j k) \log n^{(j k)}\right)+O\left(\log ^{2} n^{((j-1) k)} \log \log \left(n^{((j-1) k)}\right)\right)
$$

depth.
Note that at the end of the $j^{t h}$ phase, the time required to construct an extra Schur complement chain for the Sparsify call is less than the remaining cost after the $j^{t h}$ phase. This is the reason why we use $2^{2 c k \log ^{2}(i+1)}$ as the reduction factor for the SPARSIFY call. The following theorem takes account for the recursive call and show the total running time for the algorithm.

Lemma 13.13.5. With high probability, RecursiveConstruct $\left(\boldsymbol{M}^{(0)}\right)$ takes $2^{O(\log n / k)}$ depth and $m \log n+2^{O\left(k \log ^{2} k\right)} n \log ^{2} n$ work.

Proof. Lemma 13.13 .2 shows that the call to SPARSIFY at the start of each phase $j$ requires the construction of an extra Schur complement chain on a matrix with $n^{((j-1) k)}$ row/column blocks and at most $2^{c k \log ^{2}((j-1) k)} n^{((j-2) k)} \log n^{((j-2) k)}$ non-zeros blocks. The guarantees of Lemma 13.12 .1 gives that the number of non-zero block in the sparsified matrix is at most

$$
C 2^{2 c k \log ^{2}((j-1) k)} n^{((j-1) k)} \log n^{((j-1) k)}
$$

for some absolute constant $C$. Therefore the cost of this additional call is less than the cost of constructing the rest of the phases during the construction process. Therefore, every recursive call except the first one can be viewed as an extra branching factor of 2 at each subsequent phase.

Depth can be bounded by the total number of recursive invocations to RECURSIVEConstruct. The fact that $n^{(i)}$ is geometrically decreasing means there are $O(\log n / k)$ phases. Choosing $k$ so that $k \log ^{2} k=o(\log n)$ gives a depth of:

$$
\begin{aligned}
& \sum_{j=1}^{O(\log n / k)} 2^{j}\left(k \log (j k) \log n^{(j k)}+\log ^{2} n^{(j k)} \log \log n^{(j k)}\right) \\
= & 2^{O(\log n / k)} O\left(k \log \log n \log n^{(\text {last })}+\log ^{2} n^{(l a s t)} \log \log n^{(l a s t)}\right) \\
= & 2^{O(\log n / k)} O\left(k^{2} \log \log n+k^{2} \log k\right) \\
= & 2^{O(\log n / k)} k^{2} \log \log (n) \\
= & 2^{O(\log n / k)} .
\end{aligned}
$$

For bounding work, we start with the sparse case since all intermediate matrices generated during the construction process have density at most $2^{O\left(k \log ^{2} k\right)}$. In this case, the extra branching factor of 2 at each phase can be accounted for by weighting the cost of iteration $j$ by $2^{j}$, giving:

$$
\begin{aligned}
& \sum_{j=1}^{O(\log n / k)} 2^{j}\left(2^{O\left(k \log ^{2}(j k)\right)} n^{(j k)} \log ^{2} n^{(j k)}\right) \\
= & 2^{O\left(k \log ^{2} k\right)} n \log ^{2} n .
\end{aligned}
$$

For the dense case, the first recursive call to $\operatorname{SpARSIFy}\left(\boldsymbol{M}^{(0)}, O(1), 2^{2 c k}\right)$ is made to a graph whose edge count is much less. This leads to a geometric series, and an overhead of $O(m \log n)$ work at each step. As this can happen at most $O(\log n)$ times, it gives an additional factor of $O(\log n)$ in depth, which is still $2^{O(\log n / k)}$. The work obeys the recurrence:

$$
W(m)_{\text {dense }} \leq \begin{cases}2^{O\left(k \log ^{2} k\right)} n \log ^{2} n & \text { if } m \leq 2^{O\left(k \log ^{2} k\right)} n \log n, \text { and } \\ W_{\text {dense }}(m / 2)+m \log n+2^{O\left(k \log ^{2} k\right)} n \log ^{2} n & \text { otherwise }\end{cases}
$$

which solves to:

$$
W_{\text {dense }}(m) \leq O(m \log n)+2^{O\left(k \log ^{2} k\right)} n \log ^{2} n
$$

To obtain Theorem 13.1.2, we simply choose an appropriate initial padding and set the parameter $k$.

Proof. (of Theorem 13.1.2) Lemma 13.2 .2 allows us to solve the linear system $\boldsymbol{M}+\varepsilon \mu \boldsymbol{I}$ instead. Invoking Lemma 13.13 .5 with $k=\log \log \log n$ gives a depth of $2^{O(\log n / \log \log \log n)}$, and work of $m \log n+$ $2^{O\left(\log \log \log n \log \log \log \log ^{4} n\right)} n \log ^{2} n$. The depth term can be simplified to $n^{o(1)}$ while the work term can be simplified to $2^{O\left(\log \log \log ^{2} n\right)}=\log ^{o(1)} n$.

## ■ 13.14 Weighted Expander Constructions

In this section, we give a linear time algorithm for computing linear sized spectral sparsifiers of complete and bipartite product demand graphs. These routines give Lemmas 13.11 .2 and 13.11.1, which are crucial for controlling the densities of intermediate matrices in the spectral vertex sparsification algorithm from Section 13.11. Recall that the product demand graph with vertex set $V$ and demands $\boldsymbol{d}: V \rightarrow \mathbb{R}_{>0}$ is the complete graph in which the weight of edge $(u, v)$ is the product $d_{u} d_{v}$. Similarly, the bipartite demand graph with vertex set $U \cup V$ and demands $d: U \cup V \rightarrow \mathbb{R}_{>0}$ is the complete bipartite graph on which the weight of the edge $(u, v)$ is the product $d_{u} d_{v}$. Our routines are based on reductions to the unweighted, uniform case. In particular, we

1. Split all of the high demand vertices into many vertices that all have the same demand. This demand will still be the highest.
2. Given a graph in which almost all of the vertices have the same highest demand, we
(a) drop all of the edges between vertices of lower demand,
(b) replace the complete graph between the vertices of highest demand with an expander, and
(c) replace the bipartite graph between the high and low demand vertices with a union of stars.
3. To finish, we merge back together the vertices that split off from each original vertex.

We start by showing how to construct the expanders that we need for step (2b). We state formally and analyze the rest of the algorithm for the complete case in the following two sections. We explain how to handle the bipartite case in Section 13.14.3.

Expanders give good approximations to unweighted complete graphs, and our constructions will use the spectrally best expanders-Ramanujan graphs. These are defined in terms of the eigenvalues of their adjacency matrices. We recall that the adjacency matrix of every $d$-regular graph has eigenvalue $d$ with multiplicity 1 corresponding to the constant eigenvector. If the graph is bipartite, then it also has an eigenvalue of $-d$ corresponding to an eigenvector that takes value 1 on one side of the bipartition and -1 on the other side. These are called the trivial eigenvalues. A $d$-regular graph is called a Ramanujan graph if all of its non-trivial eigenvalues have absolute value at most $2 \sqrt{d-1}$. Ramanujan graphs were constructed independently by Margulis [188] and Lubotzky, Phillips, and Sarnak [181]. The following theorem and proposition summarizes part of their results.

Theorem 13.14.1. Let $p$ and $q$ be unequal primes congruent to 1 modulo 4. If $p$ is a quadratic residue modulo $q$, then there is a non-bipartite Ramanujan graph of degree $p+1$ with $q^{2}(q-1) / 2$ vertices. If $p$ is not a quadratic residue modulo $q$, then there is a bipartite Ramanujan graph of degree $p+1$ with $q^{2}(q-1)$ vertices.

The construction is explicit.
Proposition 13.14.2. If $p<q$, then the graph guaranteed to exist by Theorem 13.14.1 can be constructed in parallel depth $O(\log n)$ and work $O(n)$, where $n$ is its number of vertices.

Proof. When $p$ is a quadratic residue modulo $q$, the graph is a Cayley graph of $\operatorname{PSL}(2, Z / q Z)$. In the other case, it is a Cayley graph of $\operatorname{PGL}(2, Z / q Z)$. In both cases, the generators are determined by the $p+1$ solutions to the equation $p=a_{0}^{2}+a_{1}^{2}+a_{2}^{2}+a_{3}^{2}$ where $a_{0}>0$ is odd and $a_{1}, a_{2}$, and $a_{3}$ are even. Clearly, all of the numbers $a_{0}, a_{1}, a_{2}$ and $a_{3}$ must be at most $\sqrt{p}$. So, we can compute a list of all sums $a_{0}^{2}+a_{1}^{2}$ and all of the sums $a_{2}^{2}+a_{3}^{2}$ with work $O(p)$, and thus a list of all $p+1$ solutions with work $O\left(p^{2}\right)<O(n)$.

As the construction requires arithmetic modulo $q$, it is convenient to compute the entire multiplication table modulo $q$. This takes time $O\left(q^{2}\right)<O(n)$. The construction also requires the computation of a square root of -1 modulo $q$, which may be computed from the multiplication table. Given this data, the list of edges attached to each vertex of the graph may be produced using linear work and logarithmic depth.

For our purposes, there are three obstacles to using these graphs:

1. They do not come in every degree.
2. They do not come in every number of vertices.
3. Some are bipartite and some are not.

We handle the first two issues by observing that the primes congruent to 1 modulo 4 are sufficiently dense. To address the third issue, we give a procedure to convert a non-bipartite expander into a bipartite expander, and vice versa.

An upper bound on the gaps between consecutive primes congruent to 1 modulo 4 can be obtained from the following theorem of Tchudakoff.

Theorem 13.14.3. For two integers $a$ and $b$, let $p_{i}$ be the ith prime congruent to a modulo $b$. For every $\varepsilon>0$,

$$
p_{i+1}-p_{i} \leq O\left(p_{i}^{3 / 4+\varepsilon}\right) .
$$

Corollary 13.14.4. There exists an $n_{0}$ so that for all $n \geq n_{0}$ there is a prime congruent to 1 modulo 4 between $n$ and $2 n$.

We now explain how we convert between bipartite and non-bipartite expander graphs. To convert a non-bipartite expander into a bipartite expander, we take its double-cover. We recall that if $G=(V, E)$ is a graph with adjacency matrix $\boldsymbol{A}$, then its double-cover is the graph with adjacency matrix

$$
\left(\begin{array}{cc}
0 & \boldsymbol{A} \\
\boldsymbol{A}^{T} & 0
\end{array}\right)
$$

It is immediate from this construction that the eigenvalues of the adjacency matrix of the double-cover are the union of the eigenvalues of $\boldsymbol{A}$ with the eigenvalues of $-\boldsymbol{A}$.

Proposition 13.14.5. Let $G$ be a connected, d-regular graph in which all matrix eigenvalues other than $d$ are bounded in absolute value by $\lambda$. Then, all non-trivial adjacency matrix eigenvalues of the double-cover of $G$ are also bounded in absolute value by $\lambda$.

To convert a bipartite expander into a non-bipartite expander, we will simply collapse the two vertex sets onto one another. If $G=(U \cup V, E)$ is a bipartite graph, we specify how the vertices of $V$ are mapped onto $U$ by a permutation $\pi: V \rightarrow U$. We then define the collapse of $G$ induced by $\pi$ to be the graph with vertex set $U$ and edge set

$$
\{(u, \pi(v)):(u, v) \in E\} .
$$

Note that the collapse will have self-loops at vertices $u$ for which $(u, v) \in E$ and $u=\pi(v)$. We assign a weight of 2 to every self loop. When a double-edge would be created, that is when $\left(\pi(v), \pi^{-1}(u)\right)$ is also an edge in the graph, we give the edge a weight of 2 . Thus, the collapse can be a weighted graph.

Proposition 13.14.6. Let $G$ be a d-regular bipartite graph with all non-trivial adjacency matrix eigenvalues bounded by $\lambda$, and let $H$ be a collapse of $G$. Then, every vertex in $H$ has weighted degree $2 d$ and all adjacency matrix eigenvalues of $H$ other than $d$ are bounded in absolute value by $2 \lambda$.

Proof. To prove the bound on the eigenvalues, let $G$ have adjacency matrix

$$
\left(\begin{array}{cc}
0 & \boldsymbol{A} \\
\boldsymbol{A}^{T} & 0
\end{array}\right) .
$$

After possibly rearranging rows and columns, we may assume that the adjacency matrix of the collapse is given by

$$
\boldsymbol{A}+\boldsymbol{A}^{T}
$$

Note that the self-loops, if they exist, correspond to diagonal entries of value 2 . Now, let $\boldsymbol{x}$ be a unit vector orthogonal to the all-1s vector. We have

$$
\boldsymbol{x}^{T}\left(\boldsymbol{A}+\boldsymbol{A}^{T}\right) \boldsymbol{x}=\binom{\boldsymbol{x}}{\boldsymbol{x}}^{T}\left(\begin{array}{cc}
0 & \boldsymbol{A} \\
\boldsymbol{A}^{T} & 0
\end{array}\right)\binom{\boldsymbol{x}}{\boldsymbol{x}} \leq \lambda\left\|\binom{\boldsymbol{x}}{\boldsymbol{x}}\right\|^{2} \leq 2 \lambda,
$$

as the vector $[\boldsymbol{x} ; \boldsymbol{x}]$ is orthogonal to the eigenvectors of the trivial eigenvalues of the adjacency matrix of $G$.

We now state how bounds on the eigenvalues of the adjacency matrices of graphs lead to approximations of complete graphs and complete bipartite graphs.

Proposition 13.14.7. Let $G$ be a graph with $n$ vertices, possibly with self-loops and weighted edges, such that every vertex of $G$ has weighted degree $d$ and such that all non-trivial eigenvalues of the adjacency matrix of $G$ have absolute value at most $\lambda \leq d / 2$. If $G$ is not bipartite, then $(n / d) \boldsymbol{L}_{G}$ is an $\varepsilon$-approximation of $K_{n}$ for $\varepsilon=(2 \ln 2)(\lambda) / d$. If $G$ is bipartite, then $(n / d) \boldsymbol{L}_{G}$ is an $\varepsilon$-approximation of $K_{n, n}$ for $\varepsilon=(2 \ln 2)(\lambda) / d$.

Proof. Let $\boldsymbol{A}$ be the adjacency matrix of $G$. Then,

$$
\boldsymbol{L}_{G}=d \boldsymbol{I}-\boldsymbol{A} .
$$

In the non-bipartite case, we observe that all of the non-zero eigenvalues of $\boldsymbol{L}_{K_{n}}$ are $n$, so for all vectors $x$ orthogonal to the constant vector,

$$
x^{T} \boldsymbol{L}_{K_{n}} x=n x^{T} x .
$$

As all of the non-zero eigenvalues of $\boldsymbol{L}_{G}$ are between $d-\lambda$ and $d+\lambda$, for all vectors $x$ orthogonal to the constant vector

$$
n\left(1-\frac{\lambda}{d}\right) x^{T} x \leq x^{T}(n / d) \boldsymbol{L}_{G} x \leq n\left(1+\frac{\lambda}{d}\right) x^{T} x .
$$

Thus,

$$
\left(1-\frac{\lambda}{d}\right) \boldsymbol{L}_{K_{n}} \preccurlyeq \boldsymbol{L}_{G} \preccurlyeq\left(1+\frac{\lambda}{d}\right) \boldsymbol{L}_{K_{n}} .
$$

In the bipartite case, we naturally assume that the bipartition is the same in both $G$ and $K_{n, n}$. Now, let $\boldsymbol{x}$ be any vector on the vertex set of $G$. Both the graphs $K_{n, n}$ and $(n / d) G$ have Laplacian
matrix eigenvalue 0 with the constant eigenvector, and eigenvalue $2 n$ with eigenvector $[\mathbf{1} ;-\mathbf{1}]$. The other eigenvalues of the Laplacian of $K_{n, n}$ are $n$, while the other eigenvalues of the Laplacian of $(n / d) G$ are between

$$
n\left(1-\frac{\lambda}{d}\right) \quad \text { and } \quad n\left(1+\frac{\lambda}{d}\right)
$$

Thus,

$$
\left(1-\frac{\lambda}{d}\right) \boldsymbol{L}_{K_{n, n}} \preccurlyeq \boldsymbol{L}_{G} \preccurlyeq\left(1+\frac{\lambda}{d}\right) \boldsymbol{L}_{K_{n, n}}
$$

The proposition now follows from our choice of $\varepsilon$, which guarantees that

$$
e^{-\varepsilon} \leq 1-\lambda / d \quad \text { and } \quad 1+\lambda / d \leq e^{\varepsilon}
$$

provided that $\lambda / d \leq 1 / 2$.

Lemma 13.14.8. There are algorithms that on input $n$ and $\varepsilon>n^{-1 / 6}$ produce a graph having $O\left(n / \varepsilon^{2}\right)$ edges that is an $O(\varepsilon)$ approximation of $K_{n^{\prime}}$ or $K_{n^{\prime}, n^{\prime}}$ for some $n \leq n^{\prime} \leq 8 n$. These algorithms run in $O(\log n)$ depth and $O\left(n / \varepsilon^{2}\right)$ work.

Proof. We first consider the problem of constructing an approximation of $K_{n^{\prime}, n^{\prime}}$. By Corollary 13.14.4 there is a constant $n_{0}$ so that if $n>n_{0}$, then there is a prime $q$ that is equivalent to 1 modulo 4 so that $q^{2}(q-1)$ is between and $n$ and $8 n$. Let $q$ be such a prime and let $n^{\prime}=q^{2}(q-1)$. Similarly, for $\varepsilon$ sufficiently small, there is a prime $p$ equivalent to 1 modulo 4 that is between $\varepsilon^{-2} / 2$ and $\varepsilon^{-2}$. Our algorithm should construct the corresponding Ramanujan graph, as described in Theorem 13.14.1 and Proposition 13.14.2. If the graph is bipartite, then Proposition 13.14 .7 tells us that it provides the desired approximation of $K_{n^{\prime}, n^{\prime}}$. If the graph is not bipartite, then we form its double cover to obtain a bipartite graph and use Proposition 13.14.5 and Proposition 13.14 .7 to see that it provides the desired approximation of $K_{n^{\prime}, n^{\prime}}$.

The non-bipartite case is similar, except that we require a prime $q$ so that $q^{2}(q-1) / 2$ is between $n$ and $8 n$, and we use a collapse to convert a bipartite expander to a non-bipartite one, as analyzed in Proposition 13.14.6.

For the existence results in Section 13.10, we just need to know that there exist graphs of low degree that are good approximations of complete graphs. We may obtain them from the recent theorem of Marcus, Spielman and Srivastava that there exist bipartite Ramanujan graphs of every degree and number of vertices [187].

Lemma 13.14.9. For every integer $n$ and even integer $d$, there is a weighted graph on $n$ vertices of degree at most $d$ that is a $4 / \sqrt{d}$ approximation of $K_{n}$.

Proof. The main theorem of [187] tells us that there is a bipartite Ramanujan graph on $2 n$ vertices of degree $k$ for every $k \leq n$. By Propositions 13.14 .6 and 13.14.7, a collapse of this graph is a weighted graph of degree at most $2 k$ that is a $(4 \ln 2) / \sqrt{k}$ approximation of $K_{n, n}$. The result now follows by setting $d=2 k$.

## ■ 13.14.1 Sparsifying Complete Product Demand Graphs

In the rest of the section, we will adapt these expander constructions to weighted settings. We start with product demand graphs.

## Algorithm 46: $G^{\prime}=$ WeightedExpander $(d, \varepsilon)$

1. Let $\hat{n}$ be the least integer greater than $2 n / \varepsilon^{2}$ such that the algorithm described in Lemma 13.14.8 produces an $\varepsilon$-approximation of $K_{\hat{n}}$.
2. Let $t=\frac{\sum_{k} d_{k}}{\hat{n}}$.
3. Create a new product demand graph $\widehat{G}$ with demand vector $\hat{\boldsymbol{d}}$ by splitting each vertex $i$ into a set of $\left\lceil d_{i} / t\right\rceil$ vertices, $S_{i}$ :
(a) $\left\lfloor d_{i} / t\right\rfloor$ vertices with demand $t$.
(b) one vertex with demand $d_{i}-t\left\lfloor d_{i} / t\right\rfloor$.
4. Let $H$ be a set of $\hat{n}$ vertices in $\widehat{G}$ with demand $t$, and let $L$ contain the other vertices. Set $k=|L|$.
5. Partition $H$ arbitrarily into sets $V_{1}, \ldots, V_{k}$, so that $\left|V_{i}\right| \geq\lfloor\hat{n} / k\rfloor$ for all $1 \leq i \leq k$.
6. Use the algorithm described in Lemma 13.14 .8 to produce $\tilde{K}_{H H}$, an $\varepsilon$-approximation of the complete graph on $H$. Set

$$
\widetilde{G}=t^{2} \tilde{K}_{H H}+\sum_{l \in L} \frac{|H|}{\left|V_{l}\right|} \sum_{h \in V_{l}} \hat{d}_{l} \hat{d}_{h}(l, h)
$$

7. Let $G^{\prime}$ be the graph obtained by collapsing together all vertices in each set $S_{i}$.

Lemma 13.11.1. There is a routine $\operatorname{WeightedExpander}(\boldsymbol{d}, \varepsilon) \operatorname{such}$ that for any $\varepsilon>0$, and a demand vector $\boldsymbol{d} \in\left(\Re_{>0}\right)^{n}$, WEIGHTEDEXPANDER $(\boldsymbol{d}, \varepsilon)$ returns in $O\left(n \varepsilon^{-4}\right)$ work and $O(\log n)$ depth a graph $H$ with $O\left(n \varepsilon^{-4}\right)$ edges such that

$$
\boldsymbol{L}_{H} \approx_{\varepsilon} \boldsymbol{L}_{G(d)}
$$

Our algorithm for sparsifying complete product demand graphs begins by splitting the vertices of highest demands into many vertices. By splitting a vertex, we mean replacing it by many vertices whose demands sum to its original demand. In this way, we obtain a larger product demand graph. We observe that we can obtain a sparsifier of the original graph by sparsifying the larger graph, and then collapsing back together the vertices that were split.

Proposition 13.14.10. Let $G$ be a product demand graph with vertex set $\{1, \ldots, n\}$ and demands $\boldsymbol{d}$, and let $\widehat{G}=(\widehat{V}, \widehat{E})$ be a product demand graph with demands $\hat{\boldsymbol{d}}$. If there is a partition of $\widehat{V}$ into sets $S_{1}, \ldots, S_{n}$ so that for all $i \in V, \sum_{j \in S_{i}} \hat{d}_{j}=d_{i}$, then $\widehat{G}$ is a splitting of $G$ and there is a matrix $\boldsymbol{M}$ so that

$$
\boldsymbol{L}_{G}=\boldsymbol{M} \boldsymbol{L}_{\widehat{G}} \boldsymbol{M}^{T}
$$

Proof. The $(i, j)$ entry of matrix $\boldsymbol{M}$ is 1 if and only if $j \in S_{i}$. Otherwise, it is zero.
We now show that we can sparsify $G$ by sparsifying $\widehat{G}$.
Proposition 13.14.11. Let $\widehat{G}_{1}$ and $\widehat{G}_{2}$ be graphs on the same vertex set $\widehat{V}$ such that $\widehat{G}_{1} \approx_{\varepsilon} \widehat{G}_{2}$ for some $\varepsilon$. Let $S_{1}, \ldots, S_{n}$ be a partition of $\widehat{V}$, and let $G_{1}$ and $G_{2}$ be the graphs obtained by collapsing together all the vertices in each set $S_{i}$ and eliminating any self loops that are created. Then

$$
G_{1} \approx_{\varepsilon} G_{2}
$$

Proof. Let $\boldsymbol{M}$ be the matrix introduced in Proposition 13.14.10. The proof follows from the fact that

$$
\boldsymbol{L}_{G_{1}}=\boldsymbol{M} \boldsymbol{L}_{\widehat{G}_{1}} \boldsymbol{M}^{T} \quad \text { and } \quad \boldsymbol{L}_{G_{2}}=\boldsymbol{M} \boldsymbol{L}_{\widehat{G}_{2}} \boldsymbol{M}^{T}
$$

For distinct vertices $i$ and $j$, we let $(i, j)$ denote the graph with an edge of weight 1 between vertex $i$ and vertex $j$. If $i=j$, we let $(i, j)$ be the empty graph. With this notation, we can express the product demand graph as

$$
\sum_{i<j} d_{i} d_{j}(i, j)=\frac{1}{2} \sum_{i, j \in V} d_{i} d_{j}(i, j)
$$

This notation also allows us to precisely express our algorithm for sparsifying product demand graphs. This section and the next are devoted to the analysis of this algorithm. Given Proposition 13.14.11, we just need to show that $\widetilde{G}$ is a good approximation to $\widehat{G}$.

Proposition 13.14.12. The number of vertices in $\widehat{G}$ is at most $n+\hat{n}$.
Proof. The number of vertices in $\widehat{G}$ is

$$
\sum_{i \in V}\left\lceil d_{i} / t\right\rceil \leq n+\sum_{i \in V} d_{i} / t=n+\hat{n}
$$

So, $k \leq n$ and $\hat{n} \geq 2 k / \varepsilon^{2}$. That is, $|H| \geq 2|L| / \varepsilon^{2}$. In the next section, we prove the lemmas that show that for these special product demand graphs $\widehat{G}$ in which almost all weights are the maximum, our algorithm produces a graph $\widetilde{G}$ that is a good approximation of $\widehat{G}$.

Proof. (of Lemma 13.11.1) The number of vertices in the graph $\widehat{G}$ will be between $n+2 n / \varepsilon^{2}$ and $n+16 n / \varepsilon^{2}$. So, the algorithm described in Lemma 13.14 .8 will take $O(\log n)$ depth and $O\left(n / \varepsilon^{4}\right)$ work to produce an $\varepsilon$ approximation of the complete graph on $\hat{n}$ vertices. This dominates the computational cost of the algorithm.

Proposition 13.14 .11 tells us that $G^{\prime}$ approximates $G$ at least as well as $\widetilde{G}$ approximates $\widehat{G}$. To bound how well $\widetilde{G}$ approximates $\widehat{G}$, we use two lemmas that are stated in the next section. Lemma 13.14.14 shows that

$$
\widehat{G}_{H H}+\widehat{G}_{L H} \approx_{O\left(\varepsilon^{2}\right)} \widehat{G} .
$$

Lemma 13.14 .16 shows that

$$
\widehat{G}_{H H}+\widehat{G}_{L H} \approx_{4 \varepsilon} \widehat{G}_{H H}+\sum_{l \in L} \frac{|H|}{\left|V_{l}\right|} \sum_{h \in V_{l}} \hat{d}_{l} \hat{d}_{h}(l, h) .
$$

And, we already know that $t^{2} \tilde{K}$ is an $\varepsilon$-approximation of $\widehat{G}_{H H}$. Fact 13.2 .1 says that we can combine these three approximations to conclude that $\widetilde{G}$ is an $O(\varepsilon)$-approximation of $\widehat{G}$.

## ■ 13.14.2 Product demand graphs with most weights maximal

In this section, we consider product demand graphs in which almost all weights are the maximum. For simplicity, we make a slight change of notation from the previous section. We drop the hats, we let $n$ be the number of vertices in the product demand graph, and we order the demands so that

$$
d_{1} \leq d_{2} \leq \cdots \leq d_{k} \leq d_{k+1}=\cdots=d_{n}=1 .
$$

We let $L=\{1, \ldots, k\}$ and $H=\{k+1, \ldots, n\}$ be the set of low and high demand vertices, respectively. Let $G$ be the product demand graph corresponding to $d$, and let $G_{L L}, G_{H H}$ and $G_{L H}$ be the subgraphs containing the low-low, high-high and low-high edges respectively. We now show that little is lost by dropping the edges in $G_{L L}$ when $k$ is small.

Our analysis will make frequent use of the following Poincare inequality:
Lemma 13.14.13. Let $c(u, v)$ be an edge of weight $c$ and let $P$ be a path from from $u$ to $v$ consisting of edges of weights $c_{1}, c_{2}, \cdots, c_{k}$. Then

$$
c(u, v) \preceq c\left(\sum c_{i}^{-1}\right) P .
$$

As the weights of the edges we consider in this section are determined by the demands of their vertices, we introduce the notation

$$
(i, j)_{2}=d_{i} d_{j}(i, j)
$$

With this notation, we can express the product demand graph as

$$
\sum_{i<j}(i, j)_{2}=\frac{1}{2} \sum_{i, j \in V}(i, j)_{2} .
$$

Lemma 13.14.14. If $|L| \leq|H|$, then

$$
G_{H H}+G_{L H} \approx_{\frac{|L|}{|H|}} G
$$

Proof. The lower bound $G_{H H}+G_{L H} \preceq G_{H H}+G_{L H}+G_{L L}$ follows from $G_{L L} \succeq 0$.

Using lemma 13.14.13 and the assumptions $d_{l} \leq 1$ for $l \in L$ and and $d_{h}=1$ for $h \in H$, we derive for every $l_{1}, l_{2} \in L$,

$$
\left(l_{1}, l_{2}\right)_{2}=\frac{1}{|H|^{2}} \sum_{h_{1}, h_{2} \in H}\left(l_{1}, l_{2}\right)_{2}
$$

(by Lemma 13.14.13)

$$
\begin{aligned}
& \preceq \frac{1}{|H|^{2}} \sum_{h_{1}, h_{2} \in H} d_{l_{1}} d_{l_{2}}\left(\frac{1}{d_{l_{1}} d_{h_{1}}}+\frac{1}{d_{h_{1}} d_{h_{2}}}+\frac{1}{d_{h_{2}} d_{l_{2}}}\right)\left(\left(l_{1}, h_{1}\right)_{2}+\left(h_{1}, h_{2}\right)_{2}+\left(h_{2}, l_{2}\right)_{2}\right) \\
& \preceq \frac{3}{|H|^{2}} \sum_{h_{1}, h_{2} \in H}\left(\left(l_{1}, h_{1}\right)_{2}+\left(h_{1}, h_{2}\right)_{2}+\left(h_{2}, l_{2}\right)_{2}\right) \\
& =\frac{3}{|H|} \sum_{h \in H}\left(\left(l_{1}, h\right)_{2}+\left(l_{2}, h\right)_{2}\right)+\frac{6}{|H|^{2}} G_{H H} .
\end{aligned}
$$

So,

$$
\begin{aligned}
G_{L L} & =\frac{1}{2} \sum_{l_{1}, l_{2} \in L}\left(l_{1}, l_{2}\right)_{2} \\
& \preceq \frac{1}{2} \sum_{l_{1}, l_{2}}\left(\frac{3}{|H|} \sum_{h \in H}\left(\left(l_{1}, h\right)_{2}+\left(l_{2}, h\right)_{2}\right)+\frac{6}{|H|^{2}} G_{H H}\right) \\
& =\frac{3|L|}{|H|} G_{L H}+\frac{3|L|^{2}}{|H|^{2}} G_{H H} .
\end{aligned}
$$

The assumption $|L| \leq|H|$ then allows us to conclude

$$
G_{H H}+G_{L H}+G_{L L} \preceq\left(1+3 \frac{|L|}{|H|}\right)\left(G_{H H}+G_{L H}\right) .
$$

Using a similar technique, we will show that the edges between $L$ and $H$ can be replaced by the union of a small number of stars. In particular, we will partition the vertices of $H$ into $k$ sets, and for each of these sets we will create one star connecting the vertices in that set to a corresponding vertex in $L$.

We employ the following consequence of the Poincare inequality in Lemma 13.14.13.
Lemma 13.14.15. For any $\varepsilon \leq 1, l \in L$ and $h_{1}, h_{2} \in H$,

$$
\varepsilon\left(h_{1}, l\right)_{2}+(1 / 2)\left(h_{1}, h_{2}\right)_{2} \approx_{4 \sqrt{\varepsilon}} \varepsilon\left(h_{2}, l\right)_{2}+(1 / 2)\left(h_{1}, h_{2}\right)_{2}
$$

Proof. By applying Lemma 13.14.13 and recalling that $d_{h_{1}}=d_{h_{2}}=1$ and $d_{l} \leq 1$, we compute

$$
\begin{aligned}
\left(h_{1}, l\right)_{2} & \preceq d_{h_{1}} d_{l}\left(\frac{\sqrt{\varepsilon}}{d_{h_{1}} d_{h_{2}}}+\frac{1}{d_{h_{2}} d_{l}}\right)\left(\frac{1}{\sqrt{\varepsilon}}\left(h_{1}, h_{2}\right)_{2}+\left(h_{2}, l\right)_{2}\right) \\
& \preceq \frac{1+\sqrt{\varepsilon}}{\sqrt{\varepsilon}}\left(h_{1}, h_{2}\right)_{2}+(1+\sqrt{\varepsilon})\left(h_{2}, l\right)_{2} \\
& \preceq(1+\sqrt{\varepsilon})\left(h_{2}, l\right)_{2}+\frac{2}{\sqrt{\varepsilon}}\left(h_{1}, h_{2}\right)_{2} .
\end{aligned}
$$

Multiplying both sides by $\varepsilon$ and adding $(1 / 2)\left(h_{1}, h_{2}\right)_{2}$ then gives

$$
\begin{aligned}
\varepsilon\left(h_{1}, l\right)_{2}+(1 / 2)\left(h_{1}, h_{2}\right)_{2} & \preccurlyeq(1+\sqrt{\varepsilon}) \varepsilon\left(h_{2}, l\right)_{2}+(2 \sqrt{\varepsilon}+1 / 2)\left(h_{1}, h_{2}\right)_{2} \\
& \preccurlyeq(1+4 \sqrt{\varepsilon})\left(\varepsilon\left(h_{2}, l\right)_{2}+(1 / 2)\left(h_{1}, h_{2}\right)_{2}\right) \\
& \preccurlyeq e^{4 \sqrt{\varepsilon}}\left(\varepsilon\left(h_{2}, l\right)_{2}+(1 / 2)\left(h_{1}, h_{2}\right)_{2}\right) .
\end{aligned}
$$

By symmetry, we also have

$$
\varepsilon\left(h_{2}, l\right)_{2}+(1 / 2)\left(h_{1}, h_{2}\right)_{2} \preccurlyeq e^{4 \sqrt{\varepsilon}}\left(\varepsilon\left(h_{1}, l\right)_{2}+(1 / 2)\left(h_{1}, h_{2}\right)_{2}\right)
$$

Lemma 13.14.16. Recall that $L=\{1, \ldots, k\}$ and let $V_{1}, \ldots, V_{k}$ be a partition of $H=\{k+1, \ldots, n\}$ so that $\left|V_{l}\right| \geq s$ for all $l$. Then,

$$
G_{H H}+G_{L H} \approx_{4 / \sqrt{s}} G_{H H}+\sum_{l \in L} \frac{|H|}{\left|V_{l}\right|} \sum_{h \in V_{l}}(l, h)_{2}
$$

Proof. Observe that

$$
G_{L H}=\sum_{l \in L} \sum_{h \in H}(l, h)_{2}
$$

For each $l \in L, h_{1} \in H$ and $h_{2} \in V_{l}$ we apply Lemma 13.14 .15 to show that

$$
\frac{1}{\left|V_{l}\right|}\left(l, h_{1}\right)_{2}+\frac{1}{2}\left(h_{1}, h_{2}\right)_{2} \approx_{4 / \sqrt{s}} \frac{1}{\left|V_{l}\right|}\left(l, h_{2}\right)_{2}+\frac{1}{2}\left(h_{1}, h_{2}\right)_{2}
$$

Summing this approximation over all $h_{2} \in V_{l}$ gives

$$
\left(l, h_{1}\right)_{2}+\sum_{h_{2} \in V_{l}} \frac{1}{2}\left(h_{1}, h_{2}\right)_{2} \approx_{4 / \sqrt{s}} \sum_{h_{2} \in V_{l}}\left(\frac{1}{\left|V_{l}\right|}\left(l, h_{2}\right)_{2}+\frac{1}{2}\left(h_{1}, h_{2}\right)_{2}\right)
$$

Summing the left-hand side of this this approximation over all $l \in L$ and $h_{1} \in H$ gives

$$
\sum_{l \in L, h_{1} \in H}\left(l, h_{1}\right)_{2}+\sum_{h_{2} \in V_{l}} \frac{1}{2}\left(h_{1}, h_{2}\right)_{2}=\sum_{l \in L, h_{1} \in H}\left(l, h_{1}\right)_{2}+\frac{1}{2} \sum_{h_{1} \in H, l \in L} \sum_{h_{2} \in V_{l}}\left(h_{1}, h_{2}\right)_{2}=G_{L H}+G_{H H} .
$$

On the other hand, the sum of the right-hand terms gives

$$
G_{H H}+\sum_{l \in L, h_{1} \in H} \sum_{h_{2} \in V_{l}} \frac{1}{\left|V_{l}\right|}\left(l, h_{2}\right)_{2}=G_{H H}+\sum_{l \in L} \sum_{h_{2} \in V_{l}} \frac{|H|}{\left|V_{l}\right|}\left(l, h_{2}\right)_{2}
$$

## ■ 13.14.3 Weighted Bipartite Expanders

This construction extends analogously to bipartite product graphs. The bipartite product demand graph of vectors $\left(\boldsymbol{d}^{A}, \boldsymbol{d}^{B}\right)$ is a complete bipartite graph whose weight between vertices $i \in A$ and $j \in B$ is given by $w_{i j}=d_{i}^{A} d_{j}^{B}$. The main result that we will prove is:
Lemma 13.11.2. There is a routine WEIGHTEDBipartiteExpander $\left(\boldsymbol{d}^{A}, \boldsymbol{d}^{B}, \varepsilon\right)$ such that for any demand vectors $\boldsymbol{d}^{A}$ and $\boldsymbol{d}^{B}$ of total length $n$, and a parameter $\varepsilon$, it returns in $O\left(n \varepsilon^{-4}\right)$ work and $O(\log n)$ depth a bipartite graph $H$ between $A$ and $B$ with $O\left(n \varepsilon^{-4}\right)$ edges such that

$$
\boldsymbol{L}_{H} \approx_{\varepsilon} \boldsymbol{L}_{G\left(d^{A}, \boldsymbol{d}^{B}\right)}
$$

Without loss of generality, we will assume $d_{1}^{A} \geq d_{2}^{A} \geq \cdots \geq d_{n^{A}}^{A}$ and $d_{1}^{B} \geq d_{2}^{B} \geq \cdots \geq d_{n^{B}}^{B}$. As the weights of the edges we consider in this section are determined by the demands of their vertices, we introduce the notation

$$
(i, j)_{2}=d_{i}^{A} d_{j}^{B}(i, j)
$$

Our construction is based on a similar observation that if most vertices on $A$ side have $d_{i}^{A}$ equaling to $d_{1}^{A}$ and most vertices on $B$ side have $d_{i}^{B}$ equaling to $d_{1}^{B}$, then the uniform demand graph on these vertices dominates the graph.

## Algorithm 47: $G^{\prime}=\mathrm{WeightedBipartiteExpander}\left(\boldsymbol{d}^{A}, \boldsymbol{d}^{B}, \varepsilon\right)$

1. Let $n^{\prime}=\max \left(n^{A}, n^{B}\right)$ and $\hat{n}$ be the least integer greater than $2 n^{\prime} / \varepsilon^{2}$ such that the algorithm described in Lemma 13.14 .8 produces an $\varepsilon$-approximation of $K_{\hat{n}, \hat{n}}$.
2. Let $t^{A}=\frac{\sum_{k} d_{k}^{A}}{\hat{n}}$ and $t^{B}=\frac{\sum_{k} d_{k}^{B}}{\hat{n}}$.
3. Create a new bipartite demand graph $\widehat{G}$ with demands $\hat{\boldsymbol{d}}^{A}$ and $\hat{\boldsymbol{d}}^{B}$ follows:
(a) On the side $A$ of the graph, for each vertex $i$, create a subset $S_{i}$ consisting of $\left\lceil d_{i}^{A} / t^{A}\right\rceil$ vertices:
i. $\left\lfloor d_{i}^{A} / t^{A}\right\rfloor$ with demand $t^{A}$.
ii. one vertex with demand $d_{i}^{A}-t^{A}\left\lfloor d_{i}^{A} / t^{A}\right\rfloor$.
(b) Let $H^{A}$ contain $\hat{n}$ vertices of $A$ of with demand $t^{A}$, and let $L^{A}$ contain the rest. Set $k^{A}=\left|L^{A}\right|$.
(c) Create the side $B$ of the graph with partition $H^{B}, L^{B}$ and demand vector $\hat{\boldsymbol{d}}^{B}$ similarly.
4. Partition $H^{A}$ into sets of size $\left|V_{i}^{A}\right| \geq\left\lfloor\hat{n} / k^{A}\right\rfloor$, one corresponding to each vertex $l \in L^{A}$. Partition $V_{B}$ similarly.
5. Let $\tilde{K}_{H^{A} H^{B}}$ be a bipartite expander produced by Lemma 13.14 .8 that $\varepsilon$-approximates $K_{\hat{n} \hat{n}}$, identified with the vertices $H^{A}$ and $H^{B}$.
6. Set

$$
\widetilde{G}=t^{A} t^{B} \tilde{K}+\sum_{l \in L^{A}} \frac{\left|H^{B}\right|}{\left|V_{l}^{B}\right|} \sum_{h \in V_{l}^{B}} \hat{d}_{l}^{A} \hat{d}_{h}^{B}(l, h)+\sum_{l \in L^{B}} \frac{\left|H^{A}\right|}{\left|V_{l}^{A}\right|} \sum_{h \in V_{l}^{A}} \hat{d}_{l}^{B} \hat{d}_{h}^{A}(l, h)
$$

7. Let $G^{\prime}$ be the graph obtained by collapsing together all vertices in each set $S_{i}^{A}$ and $S_{i}^{B}$.

Similarly to the non-bipartite case, the Poincare inequality show that the edges between low demand vertices can be completely omitted if there are many high demand vertices which allows the demand routes through high demand vertices.
Lemma 13.14.17. Let $G$ be the bipartite product demand graph of the demand $\left(\boldsymbol{d}_{i}^{A}, \boldsymbol{d}_{j}^{B}\right)$. Let $H^{A}$ a subset of vertices on $A$ side with demand higher than the set of remaining vertices $L^{A}$ on $A$ side. Define $H^{B}, L^{B}$ similarly. Assume that $\left|L^{A}\right| \leq\left|H^{A}\right|$ and $\left|L^{B}\right| \leq\left|H^{B}\right|$, then

$$
G_{H^{A} H^{B}}+G_{H^{A} L^{B}}+G_{L^{A} H^{B}} \approx_{3 \max }\left(\frac{\left|L^{A}\right|}{\left|H^{A}\right|} \cdot \frac{\left|L^{B}\right|}{\left|H^{B}\right|}\right) \text {. }
$$

Proof. The proof is analogous to Lemma 13.14.14, but with the upper bound modified for bipartite
graphs.
For every edge $l_{A}, l_{B}$, we embed it evenly into paths of the form $l_{A}, h_{B}, h_{A}, l_{B}$ over all choices of $h_{A}$ and $h_{B}$. The support of this embedding can be calculated using Lemma 13.14.13, and the overall accounting follows in the same manner as Lemma 13.14.14.

It remains to show that the edges between low demand and high demand vertices can be compressed into a few edges. The proof here is also analogous to Lemma 13.14.15: we use the Poincare inequality to show that all demands can routes through high demand vertices. The structure of the bipartite graph makes it helpful to further abstract these inequalities via the following Lemma for four edges.

Lemma 13.14.18. Let $G$ be the bipartite product demand graph of the demand $\left(\boldsymbol{d}_{i}^{A}, \boldsymbol{d}_{j}^{B}\right)$. Given $h_{A}, l_{A} \in A$ and $h_{B, 1}, h_{B, 2} \in B$. Assume that $d_{h_{A}}^{A}=d_{h_{B, 1}}^{B}=d_{h_{B, 2}}^{B} \geq d_{l_{A}}^{A}$. For any $\varepsilon<1$, we have

$$
\varepsilon\left(l_{A}, h_{B, 1}\right)_{2}+\left(h_{A}, h_{B, 2}\right)_{2}+\left(h_{A}, h_{B, 1}\right)_{2} \approx_{3 \sqrt{\varepsilon}} \varepsilon\left(l_{A}, h_{B, 2}\right)_{2}+\left(h_{A}, h_{B, 2}\right)_{2}+\left(h_{A}, h_{B, 1}\right)_{2} .
$$

Proof. Using Lemma 13.14.13 and $d_{h_{A}}^{A}=d_{h_{B, 1}}^{B}=d_{h_{B, 2}}^{B} \geq d_{l_{A}}^{A}$, we have

$$
\begin{aligned}
& \left(l_{A}, h_{B, 1}\right)_{2} \\
& \preceq d_{l_{A}}^{A} d_{h_{B, 1}}^{B}\left(\frac{1}{d_{l_{A}}^{A} d_{h_{B, 2}}^{B}}+\frac{\sqrt{\varepsilon}}{d_{h_{A}}^{A} d_{h_{B, 2}}^{B}}+\frac{\sqrt{\varepsilon}}{d_{h_{A}}^{A} d_{h_{B, 1}}^{B}}\right)\left(\left(l_{A}, h_{B, 2}\right)_{2}+\frac{1}{\sqrt{\varepsilon}}\left(h_{A}, h_{B, 2}\right)_{2}+\frac{1}{\sqrt{\varepsilon}}\left(h_{A}, h_{B, 1}\right)_{2}\right) \\
& \preceq(1+2 \sqrt{\varepsilon})\left(l_{A}, h_{B, 2}\right)_{2}+\frac{1+2 \sqrt{\varepsilon}}{\sqrt{\varepsilon}}\left(h_{A}, h_{B, 2}\right)_{2}+\frac{1+2 \sqrt{\varepsilon}}{\sqrt{\varepsilon}}\left(h_{A}, h_{B, 1}\right)_{2} .
\end{aligned}
$$

Therefore,

$$
\left.\left.\left.\left.\varepsilon\left(l_{A}, h_{B, 1}\right)_{2}+\mathbf{(} h_{A}, h_{B, 2}\right)_{2}+\mathbf{(} h_{A}, h_{B, 1}\right)_{2} \preceq(1+3 \sqrt{\varepsilon})\left(\varepsilon\left(l_{A}, h_{B, 2}\right)_{2}+\mathbf{(} h_{A}, h_{B, 2}\right)_{2}+\mathbf{(} h_{A}, h_{B, 1}\right)_{2}\right) .
$$

The other side is similar due to the symmetry.
(of Lemma 13.11.2) The proof is analogous to Lemma 13.11.1. After the splitting, the demands in $H^{A}$ are higher than the demands in $L^{A}$ and so is $H^{B}$ to $L^{B}$. Therefore, Lemma 13.14 .17 shows that that

Proof.

$$
\widehat{G}_{H^{A} H^{B}}+\widehat{G}_{H^{A} L^{B}}+\widehat{G}_{L^{A} H^{B}} \approx_{3 \varepsilon^{2} / 2} \widehat{G} .
$$

By a proof analogous to Lemma 13.14.16, one can use Lemma 13.14.18 to show that

$$
\widehat{G}_{H^{A} H^{B}}+\widehat{G}_{H^{A} L^{B}}+\widehat{G}_{L^{A} H^{B}} \approx_{O(\varepsilon)} \widehat{G}_{H^{A} H^{B}}+\frac{\left|H^{B}\right|}{\left|V_{l}^{B}\right|} \sum_{h \in V_{l}^{B}} \hat{d}_{l}^{A} \hat{d}_{h}^{B}(l, h)+\sum_{l \in L^{B}} \frac{\left|H^{A}\right|}{\left|V_{l}^{A}\right|} \sum_{h \in V_{l}^{A}} \hat{d}_{l}^{B} \hat{d}_{h}^{A}(l, h) .
$$

And, we already know that $t^{A} t^{B} \tilde{K}$ is an $\varepsilon$-approximation of $\widehat{G}_{H^{A} H^{B}}$. Fact 13.2 .1 says that we can combine these three approximations to conclude that $\widetilde{G}$ is an $O(\varepsilon)$-approximation of $\widehat{G}$.

Part IV
Geometry


# Sampling In Sub-Quadratic Steps 

## ■ 14.1 Introduction

In this chapter, we study the problem of sampling a high-dimensional polytope, a fundamental algorithmic problem with many applications. The problem can be solved in randomized polynomial time. Progress on the more general problem of sampling a convex body given only by a membership oracle $[78,79,175,176,177,131,179,178,261]$ has lead to a set of general-purpose techniques, both for algorithms and for analysis in high dimension. All known algorithms are based on sampling by discrete-time Markov chains. These include the ball walk [174], hit-and-run [241, 179] and the Dikin walk [132], the last requiring stronger access than a membership oracle. In each case, the main challenge is analyzing the mixing time of the Markov chain. For a polytope defined by $m$ inequalities in $\mathbb{R}^{n}$, the current best complexity of sampling is roughly the minimum of $n^{3} \cdot m n$ and $m n \cdot m n^{\omega-1}$ where the first factor in each term is the mixing time and the second factor is the time to implement one step. In fact, the bound of $n^{3}$ on the mixing time (achieved by the ball walk and hit-and-run) holds for arbitrary convex bodies, and $O(m n)$ is just the time to implement a membership oracle. The second term is for the Dikin walk, for which Kannan and Narayanan showed a mixing time of $O(m n)$ for the Dikin walk [132], with each step implementable in roughly matrix multiplication time. For general convex bodies given by membership oracles, $\Omega\left(n^{2}\right)$ is a lower bound on the number of oracle calls for all known walks. A quadratic upper bound would essentially follow from a positive resolution of the KLS hyperplane conjecture (we mention that [59] show a mixing bound of $\tilde{O}\left(n^{2}\right)$ for the ball walk for sampling from a Gaussian distribution restricted to a convex body). The quadratic barrier seems inherent for sampling convex bodies given by membership oracles, holding even for cubes and cylinders for the known walks based on membership oracles. It has not been surpassed thus far for explicitly described polytopes by any means.

For a polytope in $\mathbb{R}^{n}$, the Euclidean perspective is natural and predominant. The game has been to define a process on the points of the polytope so that the distribution of the current point quickly approaches the uniform (or a desired stationary distribution). The difficulty is that for points near the boundary of a body, steps are necessarily small due to the nature of volume distribution in high dimension. The Dikin walk departs from the standard perspective by making the distribution of the next step depend on the distance(s) of the current point to the boundary of the polytope. At each step, the process picks a random point from a suitable ellipsoid that is guaranteed to almost lie inside. This process adapts to the boundary, but runs into the similar difficulties - the ellipsoid has to shrink as the point approaches the boundary in order to ensure that (a) the stationary distribution is close to uniform and (b) the 1-step distribution has to be smooth, both necessary properties for rapidly converging to the uniform distribution. While the walk has the appealing property of being affineinvariant, and thus avoiding having to explicitly round the polytope, the current best upper bound is still quadratic, even for cylinders.

An alternative approach for sampling is the simulation of Brownian motion with boundary reflection [114, 67, 64, 43]. While there has been much study of this process, several difficulties arise in
turning it into an efficient algorithm. In particular, if the current point is close to the boundary of the polytope, extra care is needed in simulation and the process effectively slows down. However, if it is deep inside the polytope, we should expect that Brownian motion locally looks like a Gaussian distribution and hence it is easier to simulate. This suggests that the standard Euclidean metric, which does not take into account distance to the boundary, is perhaps not the right notion for getting a fast sampling algorithm.

In this chapter, we combine the use of stochastic differential equations (SDE) with non-Euclidean geometries (Riemannian manifolds) to break the quadratic barrier for mixing in polytopes. As a result we obtain significantly faster sampling algorithms.

Roughly speaking, our work is based on three key conceptual ideas. The first is the use of a Riemannian metric rather than the Euclidean metric. This allows us to scale space as we get closer to the boundary and incorporate boundary information much more smoothly. This idea was already used by Narayanan [198] to extend the Dikin walk to more general classes of convex bodies. The relevant metrics are induced by Hessians of convex barrier functions, objects that have been put to remarkable use for efficient linear and convex optimization [211]. The second idea is to simulate a Stochastic Differential Equation (SDE) on the manifold corresponding to the metric, via its geodesics (shortest-path curves, to be defined presently). Unlike straight lines, geodesics bend away from the boundary and this allows us to take larger steps while staying inside the polytope. The third idea is to use a modification of standard Brownian motion via a drift term, i.e., rather than centering the next step at the current point, we first shift the current point deterministically, then take a random step. This drift term compensates the changes of the step size and this makes the process closer to symmetric. Taken together, these ideas allow us to simulate an SDE by a discrete series of ordinary differential equations (ODE), which we are able to solve efficiently to the required accuracy. In order to state our contributions and results more precisely, we introduce some background, under three headings.

Riemannian Geometry. A manifold can be viewed as a surface embedded in a Euclidean space. Each point in the manifold (on the surface), has a tangent space (the linear approximate of the surface at that point) and a local metric. For a point $x$ in a manifold $M$, the metric at $x$ is defined by a positive definite matrix $g(x)$ and the length of a vector $u$ in the tangent space $T_{x} M$ is defined as $\|u\|_{x} \stackrel{\text { def }}{=} u^{T} g(x) u$. By integration, the length of any curve on the manifold is defined as $\int\left\|\frac{d c}{d t}\right\|_{c(t)}$. A basic fact about Riemannian manifolds is that for any point in the manifold, in any direction (from the tangent space), there is locally a shortest path (geodesic) starting in that direction. In Euclidean space, this is just a straight line in the starting direction. Previous random walks involve generating a random direction and going along a straight line in that direction. However such straight lines do not take into account the local geometry, while geodesics do. We give formal definitions in Section 14.2.1.

Hessian Manifolds. In this chapter, we are concerned primarily with Riemannian manifolds induced by the Hessians of smooth (infinitely differentiable) convex functions. More precisely, for any such function $\phi$, the metric (of the manifold induced by $\phi$ ) at a point $x \in M$ is given by the Hessian of $\phi$ at $x$, i.e., $\nabla^{2} \phi(x)$. Since $\phi$ is convex, the Hessian is positive definite and hence the Riemannian manifold induced by $\phi$ is well-defined and is called a Hessian manifold. In the context of convex optimization, we are interested in a class of convex functions called self-concordant barriers. Such convex function is smooth in a certain sense and blows up on the boundary of a certain convex set. The class of Hessian manifolds corresponding to self-concordant barriers has been studied and used to study interior-point methods (IPM) [137, 214, 212].

The barrier we are particularly interested is the logarithmic barrier. For a polytope $A x>b$, with
$A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^{m}$, for any $x$ in the polytope, the logarithmic barrier is given by

$$
\phi(x)=-\sum_{i} \ln (A x-b)_{i} .
$$

In this chapter, we study the logarithmic barrier and show how to use this to develop a faster algorithm for sampling polytopes.

Stochastic Differential Equations. Given a self-concordant barrier $\phi$ on a convex set $K$, there is an "unique" Brownian motion with drift on the Hessian manifold $M$ induced by $\phi$ that has uniform stationary distribution. In the Euclidean coordinate system, the SDE is given by

$$
\begin{equation*}
d x_{t}=\mu\left(x_{t}\right) d t+\left(\nabla^{2} \phi\left(x_{t}\right)\right)^{-1 / 2} d W_{t} \tag{14.1}
\end{equation*}
$$

where the first term, called drift, is given by:

$$
\begin{equation*}
\mu_{i}\left(x_{t}\right)=\frac{1}{2} \sum_{j=1}^{n} \frac{\partial}{\partial x_{j}}\left(\left(\nabla^{2} \phi\left(x_{t}\right)\right)^{-1}\right)_{i j} \tag{14.2}
\end{equation*}
$$

This suggests an approach for generating a random point in a polytope, namely to simulate the SDE. The running time of such an algorithm depends on the convergence rate of the SDE and the cost of simulating the SDE in discrete time steps.

Since the SDE is defined on the Riemannian manifold $M$, it is natural to consider the following geodesic walk:

$$
\begin{equation*}
x^{(j+1)}=\exp _{X^{(j)}}\left(\sqrt{h} w+\frac{h}{2} \mu\left(x^{(j)}\right)\right) \tag{14.3}
\end{equation*}
$$

where $\exp _{x^{(j)}}$ is a map from $T_{x^{(j)}} M$ back to the manifold, $w$ is a random Gaussian vector on $T_{x^{(j)}} M$, $\mu\left(x^{(j)}\right) \in T_{x^{(j)}} M$ is the drift term and $h$ is the step size. The coefficient of the drift term depends on the coordinate system we use; as we show in Lemma 14.3.1, for the coordinate system induced by a geodesic, the drift term is $\mu / 2$ instead of $\mu$ as in (14.1). The Gaussian vector $w$ has mean 0 and variance 1 in the metric at $x$, i.e. $\mathbf{E}_{w}\|w\|_{x}^{2}=1$. We write it as $w \sim N_{x}(0, I)$.

It can be shown that this discrete walk converges to (14.1) as $h \rightarrow 0$ and it converges in a rate faster than the walk suggested by Euclidean coordinates (Theorem 14.3.3), namely, $x^{(j+1)}=$ $x^{(j)}+\sqrt{h} w+h \mu\left(x^{(j)}\right)$. (Note the drift here is proportional to $h$ and not $h / 2$.) This is the reason we study the geodesic walk.

## ■ 14.1.1 Algorithm

The algorithm is a discrete-time simulation of the geodesic process (14.3). For step-size $h$ chosen in advance, let $p(x \xrightarrow{w} y)$ be the probability density (in Euclidean coordinates) of going from $x$ to $y$ using the local step $w$. In general, the stationary distribution of the geodesic process is not uniform and it is difficult to analyze the stationary distribution unless $h$ is very small, which would lead to a high number of steps. To get around this issue, we use the standard method of rejection sampling to get a uniform stationary distribution. We call this geodesic walk.

To implement the geodesic walk, we need to compute the exponential map $\exp _{x}$ and $p(x \xrightarrow{w} y)$ efficiently. We show how to implement this in Section 14.4.6.Our goal is to make sure that each iteration of geodesic walks only uses matrix multiplication and matrix inverse for $O\left(\log ^{O(1)} m\right)$ many $O(m) \times O(m)$-size matrices, and the rejection probability is small.

It turns out that the problems of computing the exponential map and $p(x \xrightarrow{w} y)$ are similar; both involve solving some ordinary differential equations (ODEs) to accuracy $1 / n^{\Theta(1)}$. Hence, one can view

```
Algorithm 48: Geodesic Walk (See Algo 49 for details)
    Pick a Gaussian random vector \(w \sim N_{x}(0,1)\), i.e. \(\mathbf{E}_{w}\|w\|_{x}^{2}=1\).
    Compute \(y=\exp _{x}\left(\sqrt{h} w+\frac{h}{2} \mu(x)\right)\) where \(\mu(x)\) is given by (14.2).
    Let \(p(x \xrightarrow{w} y)\) is the probability density of going from \(x\) to \(y\) using the above step \(w\).
    Compute a corresponding \(w^{\prime}\) s.t. \(x=\exp _{y}\left(\sqrt{h} w^{\prime}+\frac{h}{2} \mu(y)\right)\).
    With probability \(\min \left(1, \frac{p\left(y \xrightarrow{w^{\prime}} x\right)}{p(x \rightarrow y)}\right)\), go to \(y\);
    Otherwise, stay at \(x\).
```

the geodesic walk as reducing the problem of simulating an SDE 14.1 to solving a sequence of ODEs. Although solving ODEs is well-studied, existing literature seems quite implicit about the dependence on the dimension and hence it is difficult to apply it directly. In Section 14.6, we rederive some existing result about solving ODEs efficiently, but with quantitative estimates of the dependence on the dimension and desired error.

## ■ 14.1.2 Main results

In this chapter, we analyze the geodesic walk for the logarithmic barrier. The convergence analysis will need tools from Riemannian geometry and stochastic calculus, while the implementation uses efficient (approximate) solution of ODEs. Both aspects appear to be of independent interest. For the reader unfamiliar with these topics, we include an exposition of the relevant background.

We analyze geodesic walk in general and give a bound on the mixing time in terms of a set of manifold parameters (Theorem 14.4.2). Applying this on the logarithmic barrier, we obtain a faster sampling algorithm, going below the $m n$ mixing time of the Dikin walk, while maintaining the same per-step complexity.

Theorem 14.1.1 (Sampling with logarithmic barrier). Given a polytope $\{A x \geq b\}$ with an uniformly random initial point in the polytope. Using the logarithmic barrier, we can sample another uniformly random point in the polytope via geodesic walk using $\tilde{O}\left(m n^{\frac{3}{4}}\right)$ steps, with each step taking $\tilde{O}\left(m n^{\omega-1}\right)$ time to implement.

The implementation of each step of sampling is based on the efficient solution of high-dimensional ODEs (Theorem 14.6.4); this might have other applications.

## ■ 14.1.3 Discussion and open problems

At a high level, our algorithm is based on the following sequence of natural choices. First, we consider an Brownian motion that gives uniform stationary distribution and such Brownian motion is unique for a given metric (Fokker-Planck equation). Since the set we sample is a polytope, we use the metric given by the Hessian of a self-concordant barrier, a well-studied class of metrics in convex optimization. This allows us to reduce the sampling problem to the problem of simulating an $\mathrm{SDE}^{1}$. To simulate an SDE, we apply the Milstein method, well-known in that field. To implement the Milstein method, we perform a change of coordinates to make the metric locally constant and this makes most of the terms in the Milstein approximation method vanish. These coordinates are called normal coordinates

[^34]and can be calculated by geodesics (Lemma 14.3.1). After all of this, we obtain the step of our walk (14.3).

There are two choices which are not natural. First, it is unclear that Hessians of self-concordant barriers are the best metrics for the sampling problems; after all, these barriers are designed for solving linear programs. Second, it is unclear that Milstein method is the best choice. There are other numerical method for SDE with better convergence rates, such as higher-order Runge-Kutta schemes. However, these methods are very complex and it is not clear how to implement the steps of these methods in $\tilde{O}\left(m n^{\omega-1}\right)$ time.

Clearly, we are not simultaneously experts on Riemannian geometry, numerical SDE, numerical ODE and convex geometry; we view this chapter as a sampling algorithm that connects different areas while improving the state of the art. Although the sampling problem is a harder problem than solving linear programs, the step size of geodesic walk we use is larger than that of the short-step interior point method. We hope that the relations revealed in this chapter might be useful for further development of both samplers and linear programming solvers.

There are several avenues for improving the sampling complexity further. One is to take longer steps and use a higher-order simulation of the SDE that is accurate up to a larger distance. Another part that is not tight in the analysis is the isoperimetry. At the moment, we incur a linear factor in the mixing time due to the isoperimetry ( $m$ for the $\log$ barrier, $\tilde{O}(n)$ for the LS barrier). As far as we know, this factor might in fact be as small as $O(1)$. We make this precise via the following generalization of the KLS hyperplane conjecture. If true, it would directly improve our mixing time bound to $\tilde{O}\left(n^{\frac{3}{4}}\right)$.

Conjecture 14.1.2. For any Hessian manifold $M$ with metric d, let the isoperimetric ratio of a subset $S$ w.r.t. d be defined as

$$
\psi_{d}(S)=\inf _{\varepsilon>0} \frac{\operatorname{vol}(\{x \in K \backslash S, d(x, y) \leq \varepsilon\})}{\varepsilon \cdot \min \{\operatorname{vol}(S), \operatorname{vol}(K \backslash S)\}}
$$

and the isoperimetric ratio of $d$ as $\psi=\inf _{S \subset K} \psi_{d}(S)$. Then there is a subset $S$ defined as a halfspace intersected with $K$ with $\psi_{d}(S)=O(\psi)$.

## - 14.1.4 Outline

In the next section, we recall the basic definitions and properties of Riemannian geometry and Hessian manifolds. Following that, in Section 14.3, we derive the discrete-time geodesic walk, showing how the magic formula naturally arises. In Section 14.4 , we prove the main mixing bound for the geodesic walk in terms of a set of manifold parameters. In subsequent sections, we bound these parameters for the logarithmic barrier, thus proving Theorem 14.1.1. The algorithm for solving ODEs (collocation method) is presented and analyzed in Section 14.6.

## ■ 14.2 Preliminaries

Throughout the chapter, we use lowercase letter for vectors and vector fields and uppercase letter for matrices and tensors. We use $e_{k}$ to denote coordinate vectors. We use $\frac{d}{d t}$ for the usual derivative, e.g. $\frac{d f(c(t))}{d t}$ is the derivative of some function along a curve parametrized by $t, \nabla$ for the connection (manifold derivative, defined below which takes into account the local metric), $D_{v}$ for the directional derivative wrt to the vector (or vector field) $v$, and $D_{t}$ if the parametrization is clear from the context. We use $g$ for the local metric. Given a point $x \in M, g$ is a matrix with entries $g_{i j}$. Its inverse has


Figure 14-1: Riemannian Manifold and Tangent Space
entries $g^{i j}$. Also $n$ is the dimension, $m$ the number of inequalities, $\gamma$ is a geodesic, and $\phi$ is a smooth convex function.

## ■ 14.2.1 Basic Definitions of Riemannian geometry

Here we first introduce basic notions of Riemannian geometry. In this chapter, we use lowercase letter for vectors and vector fields and uppercase letter for matrices (this is not the convention used in Riemannian geometry). One can think of a manifold $M$ as a $n$-dimensional "surface" in $\mathbb{R}^{m}$ for some $m \geq n$.

1. Tangent space $T_{p} M$ : For any point $p$, the tangent space $T_{p} M$ of $M$ at point $p$ is a linear subspace of $\mathbb{R}^{m}$. Intuitively, $T_{p} M$ is the vector space of possible directions that are tangential to the manifold at $x$. Equivalently, it can be thought as the first-order linear approximation of the manifold $M$ at $p$. For any curve $c$ on $M$, the direction $\frac{d}{d t} c(t)$ is tangent to $M$ and hence lies in $T_{c(t)} M$. When it is clear from context, we define $c^{\prime}(t)=\frac{d c}{d t}(t)$. For any open subset $M$ of $\mathbb{R}^{n}$, we can identify $T_{p} M$ with $\mathbb{R}^{n}$ because all directions can be realized by derivatives of some curves in $\mathbb{R}^{n}$.
2. Riemannian metric: For any $v, u \in T_{p} M$, the inner product (Riemannian metric) at $p$ is given by $\langle v, u\rangle_{p}$ and this allows us to define the norm of a vector $\|v\|_{p}=\sqrt{\langle v, v\rangle_{p}}$. We call a manifold a Riemannian manifold if it is equipped with a Riemannian metric. When it is clear from context, we define $\langle v, u\rangle=\langle v, u\rangle_{p}$. In $\mathbb{R}^{n},\langle v, u\rangle_{p}$ is the usual $\ell_{2}$ inner product.
3. Pushforward $d$ : Given a function $f$ from a manifold $M$ to a manifold $N$, we define $d f(x)$ as the linear map from $T_{x} M$ to $T_{f(x)} N$ such that

$$
d f(x)\left(c^{\prime}(0)\right)=(f \circ c)^{\prime}(0)
$$

for any curve $c$ on $M$ starting at $x=c(0)$. When $M$ and $N$ are Euclidean spaces, $d f(x)$ is the Jacobian of $f$ at $x$.
4. Hessian manifold: We call $M$ a Hessian manifold if $M$ is an open subset of $\mathbb{R}^{n}$ with the Rieman-
nian metric at a point $p \in M$ defined by

$$
\langle v, u\rangle_{p}=v^{T} \nabla^{2} \phi(p) u
$$

where $v, u \in T_{p} M$ and $\phi$ is a smooth convex function on $M$.
5. Length: For any curve $c:[0,1] \rightarrow M$, we define its length by

$$
L(c)=\int_{0}^{1}\left\|\frac{d}{d t} c(t)\right\|_{c(t)} d t
$$

6. Distance: For any $x, y \in M$, we define $d(x, y)$ be the infimum of the lengths of all paths connecting $x$ and $y$. In $\mathbb{R}^{n}, d(x, y)=\|x-y\|_{2}$.
7. Geodesic: We call a curve $\gamma(t):[a, b] \rightarrow M$ a geodesic if
(a) The curve $\gamma(t)$ is parameterized with constant velocity. Namely, there is some $c$ such that $\left\|\frac{d}{d t} \gamma(t)\right\|_{\gamma(t)}=c$ for all $a \leq t \leq b$
(b) The curve is the locally shortest length curve between $\gamma(a)$ and $\gamma(b)$. Namely, for any family of curve $c(t, s)$ with $c(t, 0)=\gamma(t)$ and $c(0, a)=\gamma(a)$ and $c(0, b)=\gamma(b)$, we have that $\left.\frac{d}{d s}\right|_{s=0} \int_{a}^{b}\left\|\frac{d}{d t} c(t, s)\right\|_{c(t, s)} d t=0$.

Note that, if $\gamma(t)$ is a geodesic, then $\gamma(\alpha t)$ is a geodesic for any $\alpha$. Intuitively, geodesics are local shortest-path curves. In $\mathbb{R}^{n}$, geodesics are straight lines.
8. Exponential map: The map $\exp _{p}: T_{p} M \rightarrow M$ is defined as

$$
\exp _{p}(v)=\gamma_{v}(1)
$$

where $\gamma_{v}$ is the unique geodesic starting at $p$ with initial velocity $\gamma_{v}^{\prime}(0)$ equals to $v$. The exponential map takes a straight line $t v \in T_{p} M$ to a geodesic $\gamma_{t v}(1)=\gamma_{v}(t) \in M$. Intuitively, the exponential map can be thought as vector addition in a manifold. In $\mathbb{R}^{n}$, we have $\exp _{p}(v)=p+v$.
9. Parallel transport: Given any geodesic $c(t)$ and a vector $v$ such that $\left\langle v, c^{\prime}(0)\right\rangle_{c(0)}=0$, we define the parallel transport of $v$ along $c(t)$ by the following process: Take $h$ to be infinitesimally small and $v_{0}=v$. For $i=1,2, \cdots, 1 / h$, we let $v_{i h}$ be the vector orthogonal to $c^{\prime}(i h)$ that minimizes the distance between $\exp _{c(i h)}\left(h v_{i h}\right)$ and $\exp _{c((i-1) h)}\left(h v_{(i-1) h}\right)$. Intuitively, the parallel transport finds the vectors on the curve such that their end points are closest to the end points of $v$. For general vector $v \in T_{c^{\prime}(0)}$, we write $v=\alpha c^{\prime}(0)+w$ and we define the parallel transport of $v$ along $c(t)$ is the sum of $\alpha c^{\prime}(t)$ and the parallel transport of $w$ along $c(t)$. For non-geodesic curve, see the definition in Fact 14.2.1.
10. Orthonormal frame: Given a vector fields $v_{1}, v_{2}, \cdots, v_{n}$ on a subset of $M$, we call $\left\{v_{i}\right\}_{i=1}^{n}$ is an orthonormal frame if $\left\langle v_{i}, v_{j}\right\rangle_{x}=\delta_{i j}$ for all $x$. Given a curve $c(t)$ and an orthonormal frame at $c(0)$, we can extend it on the whole curve by parallel transport and it remains orthonormal on the whole curve.
11. Directional derivatives and Levi-Civita connection: Given any vector $v \in T_{p} M$ and a vector field $u$ in a neighborhood of $p$. Let $\gamma_{v}$ is the unique geodesic starting at $p$ with initial velocity $\gamma_{v}^{\prime}(0)=v$, we define

$$
\nabla_{v} u=\lim _{h \rightarrow 0} \frac{u(h)-u(0)}{h}
$$

where $u(h)$ is the parallel transport of $u(\gamma(h))$ from $\gamma(h)$ to $\gamma(0)$. Intuitively, Levi-Civita connection is the directional derivative of $u$ along direction $v$, taking the metric into account. In particular, for $\mathbb{R}^{n}$, we have $\nabla_{v} u(x)=\frac{d}{d t} u(x+t v)$. When $u$ is defined on a curve $c$, we define $D_{t} u=\nabla_{c^{\prime}(0)} u$. In $\mathbb{R}^{n}$, we have $D_{t} u(\gamma(t))=\frac{d}{d t} u(\gamma(t))$. We reserve $\frac{d}{d t}$ for the usual derivative with Euclidean coordinates.

We list some basic facts about the definitions introduced above that are useful for computation and intuition.

Fact 14.2.1. Given a manifold $M$, a curve $c(t) \in M$, a vector $v$ and vector fields $u, w$ on $M$, we have the following:

1. (alternative definition of parallel transport) $v(t)$ is the parallel transport of $v$ along $c(t)$ if and only if $\nabla_{c^{\prime}(t)} v(t)=0$.
2. (alternative definition of geodesic) $c$ is a geodesic if and only if $\nabla_{c^{\prime}(t)} c^{\prime}(t)=0$.
3. (linearity) $\nabla_{v}(u+w)=\nabla_{v} u+\nabla_{v} w$.
4. (product rule) For any scalar-valued function $f, \nabla_{v}(f \cdot u)=\frac{\partial f}{\partial v} u+f \cdot \nabla_{v} u$.
5. (metric preserving) $\frac{d}{d t}\langle u, w\rangle_{c(t)}=\left\langle D_{t} u, w\right\rangle+\left\langle u, D_{t} w\right\rangle$.
6. (torsion free-ness) For any map $c(t, s)$ from a subset of $\mathbb{R}^{2}$ to $M$, we have that $D_{s} \frac{\partial c}{\partial t}=D_{t} \frac{\partial c}{\partial s}$ where $D_{s}=\nabla_{\frac{\partial c}{\partial s}}$ and $D_{t}=\nabla_{\frac{\partial c}{\partial t}}$.
7. (alternative definition of Levi-Civita connection) $\nabla_{v} u$ is the unique linear mapping from the product of vector and vector field to vector field that satisfies (3), (4), (5) and (6).

### 14.2.1.1 Curvature

Here, we define various notions of curvature. Roughly speaking, they measure the amount by which a manifold deviates from Euclidean space.

Given vector $u, v \in T_{p} M$, in this section, we define $u v$ be the point obtained from moving from $x$ along direction $u$ with distance $\|u\|_{p}$ (using geodesic), then moving along direction " $v$ " with distance $\|v\|_{p}$ where " $v$ " is the parallel transport of $v$ along the path $u$. In $\mathbb{R}^{n}, u v$ is exactly $p+u+v$ and hence $u v=v u$, namely, parallelograms close up. For a manifold, parallelograms almost close up, namely, $d(u v, v u)=o(\|u\|\|v\|)$. This property is called being torsion-free.

1. Riemann curvature tensor: Three-dimensional parallelepipeds might not close up, and the curvature tensor measures how far they are from closing up. Given vector $u, v, w \in T_{p} M$, we define $u v w$ as the point obtained from moving from $u v$ along direction " $w$ " for distance $\|w\|_{p}$ where " $w$ " is the parallel transport of $w$ along the path $u v$. In a manifold, parallelepipeds do not close up and the Riemann curvature tensor how much uvw deviates from vuw. Formally, for vector fields $v, w$, we define $\tau_{v} w$ be the parallel transport of $w$ along the vector field $v$ for one unit of time. Given vector field $v, w, u$, we define the Riemann curvature tensor by

$$
\begin{equation*}
R(u, v) w=\left.\frac{d}{d s} \frac{d}{d t} \tau_{s u}^{-1} \tau_{t v}^{-1} \tau_{s u} \tau_{t v} w\right|_{t, s=0} . \tag{14.4}
\end{equation*}
$$

Riemann curvature tensor is a tensor, namely, $R(u, v) w$ at point $p$ depends only on $u(p), v(p)$ and $w(p)$.


Figure 14-2: Riemann curvature tensor measures the deviation of parallelepipeds (14.4) and Ricci curvature measures the change of volumes (14.5).
2. Ricci curvature: Given a vector $v \in T_{p} M$, the Ricci curvature Ric $(v)$ measures if the geodesics starting around $p$ with direction $v$ converge together. Positive Ricci curvature indicates the geodesics converge while negative curvature indicates they diverge. For example, let $S(0)$ be a small shape around $p$ and $S(t)$ be the set of point obtained from moving $S(0)$ along direction $v$ with $t$ unit of time. Then,

$$
\begin{equation*}
\operatorname{vol} S(t)=\operatorname{vol} S(0)\left(1-\frac{t^{2}}{2} \operatorname{Ric}(v)+\text { smaller terms }\right) \tag{14.5}
\end{equation*}
$$

Formally, we define

$$
\operatorname{Ric}(v)=\sum_{u_{i}}\left\langle R\left(v, u_{i}\right) u_{i}, v\right\rangle
$$

where $u_{i}$ is an orthonormal basis of $T_{p} M$. Equivalently, we have $\operatorname{Ric}(v)=\mathbf{E}_{u \sim N(0, I)}\langle R(v, u) u, v\rangle$. For $\mathbb{R}^{n}, \operatorname{Ric}(v)=0$. For a sphere in $n+1$ dimension with radius $r, \operatorname{Ric}(v)=\frac{n-1}{r^{2}}\|v\|^{2}$.

Fact 14.2.2 (Alternative definition of Riemann curvature tensor). Given vector fields $x, y, z$ on $M$. Then, we have

$$
R(x, y) z=\nabla_{x} \nabla_{y} z-\nabla_{y} \nabla_{x} z-\nabla_{[x, y]} z
$$

where $[x, y]$ is the vector field such that $\nabla_{[x, y]} f=\nabla_{x} \nabla_{y} f-\nabla_{y} \nabla_{x} f$ for any smooth scalar function $f$.
Given any $M$-valued function $c(t, s)$, we have vector fields $\frac{\partial c}{\partial t}$ and $\frac{\partial c}{\partial s}$ on $M$. Then, we have that

$$
R\left(\frac{\partial c}{\partial t}, \frac{\partial c}{\partial s}\right) z=\nabla_{\frac{\partial c}{\partial t}} \nabla_{\frac{\partial c}{\partial s}} z-\nabla_{\frac{\partial c}{\partial s}} \nabla_{\frac{\partial c}{\partial t}} z
$$

Sometimes, we simply write it as $R\left(\partial_{t} c, \partial_{s} c\right) z=D_{t} D_{s} z-D_{s} D_{t} z$.

Fact 14.2.3. Given vector fields $v, u, w, z$ on $M$. Then, we have

$$
\langle R(v, u) w, z\rangle=\langle R(w, z) v, u\rangle=-\langle R(u, v) w, z\rangle=-\langle R(v, u) z, w\rangle
$$

### 14.2.1.2 Jacobi field

Let $c:[0, \ell] \rightarrow M$ be a geodesic and $c(t, s)$ be a variation of $c(t)$ (i.e. $c(t, 0)=c(t))$ such that $c_{s}(t) \stackrel{\text { def }}{=} c(t, s)$ is a geodesic for all $s$. Then, $\left.x(t) \stackrel{\text { def }}{=} \frac{\partial}{\partial s} c(t, s)\right|_{s=0}$ satisfies the following equation

$$
D_{t} D_{t} x+R\left(x, \frac{d c}{d t}\right) \frac{d c}{d t}=0
$$

where $R(\cdot, \cdot)$. is Riemann curvature tensor defined before. We call any vector field satisfying the equation above a Jacobi field. In $\mathbb{R}^{n}$, geodesics are straight lines and Jacobi field is linear, namely, $x(t)=x(0)+x^{\prime}(0) t$.

Fact 14.2.4. Every Jacobi field $x$ can be split into the tangential part $x_{1}$ and the normal part $x_{2}$ such that

1. $x=x_{1}+x_{2}$,
2. $x_{1}$ and $x_{2}$ are Jacobi fields, namely, $D_{t} D_{t} x_{1}+R\left(x_{1}, \frac{d c}{d t}\right) \frac{d c}{d t}=0$ and $D_{t} D_{t} x_{2}+R\left(x_{2}, \frac{d c}{d t}\right) \frac{d c}{d t}=0$,
3. $x_{1}$ is parallel to $\frac{d c}{d t}$ and is linear, namely, $x_{1}(t)=\left(\left\langle x(0), \frac{d c}{d t}(0)\right\rangle_{c(0)}+\left\langle D_{t} x(0), \frac{d c}{d t}(0)\right\rangle_{c(0)} t\right) \frac{d c}{d t}(t)$,
4. $x_{2}$ is orthogonal to $\frac{d c}{d t}$, namely, $\left\langle x_{2}(t), \frac{d c}{d t}(t)\right\rangle_{c(t)}=0$.

### 14.2.1.3 Hessian manifolds

Recall that a manifold is called Hessian if it is a subset of $\mathbb{R}^{n}$ and its metric is given by $g_{i j}=\frac{\partial^{2}}{\partial x^{i} \partial x^{j}} \phi$ for some smooth convex function $\phi$. We let $g^{i j}$ be entries of the inverse matrix of $g_{i j}$. For example, we have $\sum_{j} g^{i j} g_{j k}=\delta_{i k}$. We use $\phi_{i j}$ to denote $\frac{\partial^{2}}{\partial x^{i} \partial x^{j}} \phi$ and $\phi_{i j k}$ to denote $\frac{\partial^{2}}{\partial x^{i} \partial x^{j} \partial x^{k}} \phi$.

Since Hessian manifold is a subset of Euclidean space, we identify tangent spaces $T_{p} M$ by Euclidean coordinates. The following lemma give formulas for Levi-Civita connection and curvature under the Euclidean coordinates.

Lemma 14.2.5 ([250]). Given a Hessian manifold $M$, vector fields $v, u, w, z$ on $M$, we have the following:

1. (Levi-Civita connection) $\nabla_{v} u=\sum_{i k} v_{i} \frac{\partial u_{k}}{\partial x_{i}} e_{k}+\sum_{i j k} v_{i} u_{j} \Gamma_{i j}^{k} e_{k}$ where $e_{k}$ are coordinate vectors and the Christoffel symbol

$$
\Gamma_{i j}^{k}=\frac{1}{2} \sum_{l} g^{k l} \phi_{i j l}
$$

2. (Riemann curvature tensor) $\langle R(u, v) w, z\rangle=\sum_{i j l k} R_{k l i j} u_{i} v_{i} w_{l} z_{k}$ where

$$
R_{k l i j}=\frac{1}{4} \sum_{p q} g^{p q}\left(\phi_{j k p} \phi_{i l q}-\phi_{i k p} \phi_{j l q}\right) .
$$

3. (Ricci curvature) Ric $(v)=\frac{1}{4} \sum_{i j l k p q} g^{p q} g^{j l}\left(\phi_{j k p} \phi_{i l q}-\phi_{i k p} \phi_{j l q}\right) v^{i} v^{k}$.

In this chapter, geodesic is used everywhere and we will implicitly using the following lemma in all of our calculations.

Lemma 14.2.6 ([214, Cor 3.1]). If $\phi$ is a self-concordant function, namely, $\left|D^{3} f(x)[h, h, h]\right| \leq$ $2\left|D^{2} f(x)[h, h]\right|^{3 / 2}$, then the corresponding Hessian manifold $M$ is geodesically complete, namely, for any $p \in M$, the exponential map is defined on the entire tangent space $T_{p} M$ and for any two points $p, q \in M$, there is a length minimizing geodesic connecting them.

In particular, for a polytope $M=\{x: A x>b\}$, the Hessian manifold induced by the function $\phi(x)=-\sum_{i} \log \left(a_{i}^{T} x-b_{i}\right)$ is geodesically complete.

### 14.2.1.4 Normal coordinates

For any manifold $M$, and any $p \in M$, the exponential map $\exp _{x}$ maps from $T_{x} M$ to $M$. Since $T_{x} M$ is isomorphic to $\mathbb{R}^{n}, \exp _{x}^{-1}$ gives a local coordinate system of $M$ around $x$. We call this system the normal coordinates at $x$. In a normal coordinate system, the metric is locally constant.

Lemma 14.2.7. In normal coordinates, we have

$$
g_{i j}=\delta_{i j}-\frac{1}{3} \sum_{k l} R_{i k j l} x^{k} x^{l}+O\left(|x|^{3}\right) .
$$

For a Hessian manifold, one can do a linear transformation on the normal coordinate to make this coincide with Euclidean coordinate up to the first order.

Lemma 14.2.8. Given a Hessian manifold $M$ and any point $x \in M$. We pick a basis of $T_{x} M$ such that

$$
\exp _{x}(t v)=x+t v+O\left(t^{2}\right)
$$

Let a local coordinate system $F: M \rightarrow \mathbb{R}^{n}$ defined by $F(y)=\exp _{x}^{-1}$. Then, we have

$$
D F[h]=h \text { and } D^{2} F_{k}(x)[h, h]=h^{T} \Gamma^{k} h
$$

where $F_{k}$ is the $k^{\text {th }}$ coordinate of $F$ and $\Gamma^{k}$ is the matrix with entries $\Gamma_{i j}^{k}$ defined in Lemma 14.2.5.

## ■ 14.2.2 Stochastic calculus

A stochastic differential equation (SDE) describes a stochastic process over a domain $\Omega$. It has the form $d x_{t}=\mu\left(x_{t}, t\right) d t+\sigma\left(x_{t}, t\right) d W_{t}$ where $x_{t}$ is the current point at time $t, W_{t}$ is a standard Brownian motion, and $\mu\left(x_{t}, t\right), \sigma\left(x_{t}, t\right)$ are the mean and covariance of the next infinitesimal step at time $t$.

Lemma 14.2.9 (Itō's lemma). Given a $S D E d x_{t}=\mu\left(x_{t}\right) d t+\sigma\left(x_{t}\right) d W_{t}$ and any smooth function $f$, we have

$$
d f\left(t, x_{t}\right)=\left\{\frac{\partial f}{\partial t}+\langle\nabla f, \mu\rangle+\frac{1}{2} \operatorname{Tr}\left[\sigma^{T}\left(\nabla^{2} f\right) \sigma\right]\right\} d t+(\nabla f)^{T} \sigma d W_{t} .
$$

SDEs are closely related to diffusion equations:

$$
\frac{\partial}{\partial t} p(x, t)=\frac{1}{2} \nabla \cdot(A(x, t) \nabla p(x, t))
$$

where $p(x, t)$ is the density at point $x$ and time $t, \nabla \cdot$ is the usual divergence operator, $\nabla p$ is the gradient of $p$ and the matrix $A(x, t)$ represents the diffusion coefficient at point $x$ and time $t$. When $A(x, t)=I$, we get the familiar heat equation:

$$
\frac{\partial}{\partial t} p(x, t)=\frac{1}{2} \Delta p(x, t)
$$

In this chapter, the diffusion coefficient will be a symmetric positive definite matrix given by the Hessian $\left(\nabla^{2} \phi(x)\right)^{-1}$ of a convex function $\phi(x)$.

The Fokker-Planck equation connects an SDE to a diffusion equation.
Theorem 14.2.10 (Fokker-Planck equation). For any stochastic differential equation (SDE) of the form

$$
d x_{t}=\mu\left(x_{t}, t\right) d t+\sqrt{A\left(x_{t}, t\right)} d W_{t}
$$

the probability density of the SDE is given by the diffusion equation

$$
\frac{\partial}{\partial t} p(x, t)=-\sum_{i=1}^{n} \frac{\partial}{\partial x_{i}}\left[\mu_{i}(x, t) p(x, t)\right]+\frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} \frac{\partial^{2}}{\partial x_{i} x_{j}}\left[A_{i j}(x, t) p(x, t)\right] .
$$

## ■ 14.2.3 Complex analysis

A complex function is said to be (complex) analytic (equivalently, holomorphic) if it is locally defined by a convergent power series. Hartog's theorem shows that a complex function in several variables $f: \mathbb{C}^{n} \rightarrow \mathbb{C}$ is holomorphic iff it is analytic in each variable (while fixing all the other variables). For any power series expansion, we define the radius of convergence at $x$ as the largest number $r$ such that the series converges on the sphere with radius $r$ centered at $x$. Roughly speaking, complex analytic functions behave very nicely up to the radius of convergence, and one can avoid complex and tedious computations via general convergence theorems.

Theorem 14.2.11 (Cauchy's Estimates). Suppose $f$ is holomorphic on a neighborhood of the ball $B \stackrel{\text { def }}{=}\left\{z \in \mathbb{C}:\left|z-z_{0}\right| \leq r\right\}$, then we have that

$$
\left|f^{(k)}\left(z_{0}\right)\right| \leq \frac{k!}{r^{k}} \sup _{z \in B}|f(z)|
$$

In particular, for any rational function $f(z)=\frac{\prod_{i=1}^{\alpha}\left(z-a_{i}\right)}{\prod_{j=1}^{\beta}\left(z-b_{j}\right)}$ and any ball $B \stackrel{\text { def }}{=}\left\{z \in \mathbb{C}:\left|z-z_{0}\right| \leq r\right\}$ such that $b_{j} \notin B$, we have that

$$
\left|f^{(k)}\left(z_{0}\right)\right| \leq \frac{k!}{r^{k}} \sup _{z \in B}|f(z)| .
$$

A similar estimate holds for analytic functions in several variables.

## ■ 14.3 Intuition

In this section, we derive the formula of the geodesic walk from first principles.

## ■ 14.3.1 Derivation of the Geodesic walk

Given a smooth convex function $\phi$ on the convex domain $M$, namely that it is convex and is infinitely differentiable at every interior point of $M$, we consider the corresponding diffusion equation

$$
\frac{\partial}{\partial t} p(x, t)=\frac{1}{2} \nabla \cdot\left(\nabla^{2} \phi\right)^{-1} \nabla p
$$

We can expand it by

$$
\begin{aligned}
\frac{\partial}{\partial t} p(x, t) & =\frac{1}{2} \sum_{i=1}^{n} \frac{\partial}{\partial x_{i}}\left(\sum_{j=1}^{n}\left(\left(\nabla^{2} \phi\right)^{-1}\right)_{i j} \frac{\partial}{\partial x_{j}} p(x, t)\right) \\
& =\frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} \frac{\partial^{2}}{\partial x_{i} x_{j}}\left(\left(\left(\nabla^{2} \phi\right)^{-1}\right)_{i j} p(x, t)\right)-\frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} \frac{\partial}{\partial x_{i}}\left(\frac{\partial}{\partial x_{j}}\left(\left(\nabla^{2} \phi\right)^{-1}\right)_{i j} p(x, t)\right) .
\end{aligned}
$$

The uniform distribution is the stationary distribution of this diffusion equation. Now applying the Fokker-Planck equation (Theorem 14.2.10) with $A=\left(\nabla^{2} \phi\right)^{-1}$, the SDE for the above diffusion is given by:

$$
d x_{t}=\mu\left(x_{t}\right) d t+\left(\nabla^{2} \phi\left(x_{t}\right)\right)^{-1 / 2} d W_{t} .
$$

This explains the definition of (14.1). To simplify the notation, we write the SDE as

$$
\begin{equation*}
d x_{t}=\mu\left(x_{t}\right) d t+\sigma\left(x_{t}\right) d W_{t} \tag{14.6}
\end{equation*}
$$

where $\sigma\left(x_{t}\right)=\left(\nabla^{2} \phi\left(x_{t}\right)\right)^{-1 / 2}$. One way to simulate this is via the Euler-Maruyama method, namely

$$
x_{(t+1) h}=x_{t h}+\mu\left(x_{t h}\right) h+\sigma\left(x_{t h}\right) w_{t h} \sqrt{h}
$$

where $w_{t h} \sim N_{x_{t h}}(0, I)$. We find the direction we are heading and take a small step along that direction. However, if we view $M$ as a manifold, then directly adding the direction $\mu\left(x_{t h}\right) h+\sigma\left(x_{t h}\right) w_{t h} \sqrt{h}$ to $x_{t h}$ is not natural; the Euclidean coordinate is just an arbitrary coordinate system and we could pick any other coordinate systems and add the direction into $x_{t h}$, giving a different step. Instead, we take the step in normal coordinates (Section 14.2.1.4).

In particular, given an initial point $x_{0}$, we define $F=\exp _{x_{0}}^{-1}$ and we note that $F\left(x_{t}\right)$ is another SDE. To see the defining equation of this transformed SDE, we use Ito's lemma to show that the transformed SDE looks the same but with half the drift term. This explains the formulation of geodesic walk: $x^{(j+1)}=\exp _{x^{(j)}}\left(\sqrt{h} w+\frac{h}{2} \mu\left(x^{(j)}\right)\right)$.

Lemma 14.3.1. Let $F=\exp _{x_{0}}^{-1}$ and $x_{t}$ satisfies the $S D E$ (14.6) Then we have

$$
\begin{equation*}
d F\left(x_{0}\right)=\frac{1}{2} \mu\left(x_{0}\right) d t+\sigma\left(x_{0}\right) d W_{0} . \tag{14.7}
\end{equation*}
$$

Proof. Itō's lemma (Lemma 14.2.9) shows that

$$
d F_{k}\left(x_{t}\right)=\left\{\left\langle\nabla F_{k}, \mu\right\rangle+\frac{1}{2} \operatorname{Tr}\left[\sigma^{T}\left(\nabla^{2} F_{k}\right) \sigma\right]\right\} d t+\left(\nabla F_{k}\right)^{T} \sigma d W_{t}
$$

where $F_{k}$ indicates the $k^{\text {th }}$ coordinate of $F$. From Lemma 14.2.8, we have that $\left\langle\nabla F_{k}\left(x_{0}\right), \mu\right\rangle=\mu_{k}$, $\left(\nabla F_{k}\left(x_{0}\right)\right)^{T} \sigma=e_{k}^{T} \sigma$ and

$$
\begin{aligned}
\operatorname{Tr}\left[\sigma\left(x_{0}\right)^{T}\left(\nabla^{2} F_{k}\left(x_{0}\right)\right) \sigma\left(x_{0}\right)\right] & =\sum_{i} D^{2} F_{k}\left[\sigma e_{i}, \sigma e_{i}\right] \\
& =\sum_{i} e_{i}^{T} \sigma^{T} \Gamma^{k} \sigma e_{i} \\
& =\operatorname{Tr}\left(\sigma^{T} \Gamma^{k} \sigma\right) \\
& =\sum e_{i}^{T} \Gamma^{k}\left(\nabla^{2} \phi\right)^{-1} e_{i} .
\end{aligned}
$$

Now, using Lemma 14.2.5, we have that

$$
\operatorname{Tr}\left[\sigma\left(x_{0}\right)^{T}\left(\nabla^{2} F_{k}\left(x_{0}\right)\right) \sigma\left(x_{0}\right)\right]=\frac{1}{2} \sum_{i j l} g^{k l} \phi_{i j l} g^{j i}
$$

Hence, we have that

$$
d F_{k}\left(x_{0}\right)=\left\{\mu_{k}+\frac{1}{4} \sum_{i j l} g^{k l} \phi_{i j l} g^{j i}\right\} d t+e_{k}^{T} \sigma d W_{t}
$$

Recall that the drift term (14.2) is given by

$$
\begin{align*}
\mu_{k} & =\frac{1}{2} \sum_{i=1}^{n} \frac{\partial}{\partial x_{i}}\left(\left(\nabla^{2} \phi\right)^{-1}\right)_{k i} \\
& =-\frac{1}{2} \sum_{i} e_{k}^{T}\left(\nabla^{2} \phi\right)^{-1} \frac{\partial}{\partial x_{i}} \nabla^{2} \phi\left(\nabla^{2} \phi\right)^{-1} e_{i} \\
& =-\frac{1}{2} \sum_{l, i, j} g^{k l} \frac{\partial}{\partial x_{i}} \phi_{l j} g^{j i} \\
& =-\frac{1}{2} \sum_{i j l} g^{k l} \phi_{i j l} g^{j i} . \tag{14.8}
\end{align*}
$$

Therefore, we have the result.

## ■ 14.3.2 Discussion of Geodesic Walk

In this section, we relate our geodesic walk with the following commonly-used method for solving SDE.
Theorem 14.3.2 (Milstein Method). Given a SDE $d x_{t}=\mu\left(x_{t}\right) d t+\sigma\left(x_{t}\right) d W_{t}$, consider the algorithm

$$
\bar{x}_{h}=x_{0}+\mu\left(x_{0}\right) h+\sigma\left(x_{0}\right) \Delta W+\sum_{j_{1}, j_{2}=1}^{m} L^{j_{1}} \sigma^{k, j_{2}} I_{j_{1} j_{2}}
$$

where $L^{j}=\sum_{k=1}^{d} \sigma^{k, j} \frac{\partial}{\partial x^{k}}, \Delta W_{j} \sim \int_{0}^{h} d W_{t}^{j}$ and $I_{j_{1} j_{2}} \sim \int_{0}^{h} \int_{0}^{t_{1}} d W_{t_{2}}^{j_{2}} d W_{t_{1}}^{j_{1}}$. Then, we have

$$
\mathbf{E}\left\|\bar{x}_{h}-x_{h}\right\|_{2}=O\left(h^{2}\right)
$$

where the $O($.$) notation hides constants related to anything except h$.
Under normal coordinates, the metric is locally constant (Lemma 14.2.7). Due to this, the term $\sum_{j_{1}, j_{2}=1}^{m} L^{j_{1}} \sigma^{k, j_{2}} I_{j_{1} j_{2}}$ in the Milstein method vanishes. Hence, we have the following result.

Theorem 14.3.3 (Geodesic Walk). Let $X_{t}$ satisfies the SDE (14.6). Consider the discrete step

$$
\bar{x}_{h}=\exp _{x_{0}}\left(\frac{1}{2} \mu\left(x_{0}\right) h+\sigma\left(x_{0}\right) \Delta W\right)
$$

where $\Delta W \sim N(0, h I)$. Then,

$$
\mathbf{E}\left\|\bar{x}_{h}-x_{h}\right\|_{2}=O\left(h^{2}\right)
$$

where the $O($.$) notation hides constants related to anything except h$.

Proof. In normal coordinates, we consider the step

$$
\overline{F\left(x_{h}\right)}=\frac{1}{2} \mu\left(x_{0}\right) h+\sigma\left(x_{0}\right) \Delta W .
$$

Since $F\left(x_{h}\right)$ satisfies 14.7 and $\sigma$ is locally constant at $X_{0}$ (Lemma 14.2.7), Theorem 14.3.2 (Milstein Method) shows that

$$
\left\|\overline{F\left(x_{h}\right)}-F\left(x_{h}\right)\right\|_{2}=O\left(h^{2}\right) .
$$

The result follows from $\bar{x}_{h}=\exp _{X_{0}}\left(\overline{F\left(x_{h}\right)}\right), x_{h}=\exp _{X_{0}}\left(F\left(x_{h}\right)\right)$ and the smoothness of $\exp _{x_{0}}$.

For the Euler-Maruyama method, $\bar{x}_{h}=x_{0}+\mu\left(x_{0}\right) h+\sigma\left(x_{0}\right) \Delta W$, we only get $O\left(h^{1.5}\right)$ as the bound on the error, rather than $O\left(h^{2}\right)$. This is one of the reasons we use geodesic instead of just adding up the vector in Euclidean coordinates.

## ■ 14.4 Convergence of the Geodesic Walk

The geodesic walk is a Metropolis-filtered Markov chain, whose stationary distribution is the uniform distribution over the polytope to be sampled. We will prove that the conductance of this chain is large with an appropriate choice of the step-size parameter. Therefore, its mixing time to converge to the stationary distribution will be small. The proof of high conductance involves showing (a) the acceptance probability of the Metropolis filter is at least a constant (b) the induced metric satisfies a strong isoperimetric inequality (c) two points that are close in metric distance are also close in probabilistic distance, namely, the one-step distributions from them have large overlap. Besides bounding the number of steps, we also have to show that each step of the Markov chain can be implemented efficiently. We describe the implementation via reduction to solving ODEs in Section 14.4.6. We do this in later sections via an efficient algorithm for approximately solving ODEs.

In this section, we present the general conductance proof for Hessian manifolds. The proof will need bounds on seven parameters determined by the specific barrier function. In Section 14.5, we will bound these parameters for the logarithmic barrier.

For a Markov chain with state space $M$, stationary distribution $Q$ and next step distribution $P_{u}(\cdot)$ for any $u \in M$, the conductance of the Markov chain is

$$
\phi=\inf _{S \subset M} \frac{\int P_{u}(M \backslash S) d Q(u)}{\min \{Q(S), Q(M \backslash S)\}} .
$$

The conductance of an ergodic Markov chain allows us to bound its mixing time, i.e., the rate of convergence to its stationary distribution, e.g., via the following theorem of Lovász and Simonovits.

Theorem 14.4.1 ([177]). Let $Q_{t}$ be the distribution of the current point after $t$ steps of a Markov chain with stationary distribution $Q$ and conductance at least $\phi$, starting from initial distribution $Q_{0}$. Then,

$$
d_{T V}\left(Q_{t}, Q\right) \leq \sqrt{d_{0}}\left(1-\frac{\phi^{2}}{2}\right)^{t}
$$

where $d_{0}=\mathbf{E}_{Q_{0}}\left(d Q_{0}(u) / d Q(u)\right)$ is a measure of the distance of the starting distribution from the stationary and $d_{T V}$ is the total variation distance.

## ■ 14.4.1 Hessian parameters

The mixing of the walk depends on the maximum values of several smoothness parameters of the manifold. For some of them, we only need to bound the parameters with high probability over the choice of steps of the geodesic walk. For $x \in M$, let $\widetilde{\Pi}_{x}$ be the set of geodesics used in one step of the geodesic walk starting at $x$ with the parameterization $\gamma:[0, \ell] \rightarrow M$ and $\ell=\sqrt{n h}$. Parameters $D_{1}, G_{1}, R_{1}, R_{2}$ below will be bounded with high probability over the choice of the next step (the others hold for all points in the domain). Let $\Pi_{x}$ denote a $1-\exp \left(n^{-\Theta(1)}\right)$ probability subset of $\widetilde{\Pi}_{x}$ and $\Pi=\bigcup_{x \in M} \Pi_{x}$. These will be fixed when bounding the parameters for specific manifolds.

1. The maximum norm of the drift, $D_{0}=\sup _{x \in M}\|\mu(x)\|_{x}$.
2. The smoothness of drift norm, $D_{1}=\sup _{\gamma \in \Pi, 0 \leq t \leq \ell} \frac{d}{d t}\|\mu(\gamma(t))\|^{2}$.
3. The smoothness of the drift, $D_{2}=\sup _{x \in M,\|s\|_{x} \leq 1}\left\|\nabla_{s} \mu(x)\right\|_{x}$.
4. The smoothness of the volume, $\left.G_{1}=\sup _{\gamma \in \Pi, 0 \leq t \leq \ell} \mid \log \operatorname{det}(g(\gamma(t)))\right)^{\prime \prime \prime} \mid$ where $g(x)$ is the metric at $x$.
5. The smoothness of the metric, $G_{2}=\sup \frac{d(x, y)}{d_{H}(x, y)}$ where $d_{H}$ is the Hilbert distance defined in Section 14.4.3 and $d$ is the shortest path distance in $M$.
6. The stability of the Jacobian field, $R_{1}=\sup _{\gamma \in \Pi, 0 \leq t \leq \ell}\|R(t)\|_{F}$ where $R(t)_{i j}=\left\langle R\left(X_{i}, \gamma^{\prime}(t)\right) \gamma^{\prime}(t), X_{j}\right\rangle$ where $\left\{X_{i}\right\}$ is any orthonormal frame at $\gamma(t)$.
7. The smoothness of the Ricci curvature, $R_{2}=\sup _{\gamma \in \Pi}\left|\frac{d}{d s} \operatorname{Ric}\left(\gamma_{s}^{\prime}(t)\right)\right|$ (see Definition 14.4.16).

We refer to these as the smoothness parameters of a Hessian manifold. Our main theorem for convergence can be stated as follows.

Theorem 14.4.2. On a Hessian manifold with smoothness parameters $D_{0}, D_{1}, D_{2}, G_{1}, G_{2}, R_{1}, R_{2}$ and assumption 14.4.23, the geodesic walk with step size

$$
0<h \leq \Theta(1) \min \left\{\frac{1}{\left(n D_{0} R_{1}\right)^{2 / 3}}, \frac{1}{D_{2}}, \frac{1}{n R_{1}}, \frac{1}{n^{1 / 3} D_{1}^{2 / 3}}, \frac{1}{n G_{1}^{2 / 3}}, \frac{1}{\left(n R_{2}\right)^{2 / 3}}\right\}
$$

has conductance $\Omega\left(\sqrt{h} / G_{2}\right)$ and its mixing time is $O\left(G_{2}^{2} / h\right)$.

## ■ 14.4.2 1-step distribution

We derive a formula for the drift term - it is in fact a Newton step of the volumetric barrier function $\log \operatorname{det} \nabla^{2} \phi(x)$.

Lemma 14.4.3. We have

$$
\mu(x)=-\left(\nabla^{2} \phi(x)\right)^{-1} \nabla \psi(x)
$$

where $\psi(x)=\frac{1}{2} \log \operatorname{det} \nabla^{2} \phi(x)$.

Proof. We note that

$$
\begin{aligned}
\frac{\partial}{\partial x_{j}} \log \operatorname{det}\left(\nabla^{2} \phi\right)^{-1} & =\operatorname{Tr}\left(\left(\nabla^{2} \phi\right) \frac{\partial}{\partial x_{j}}\left(\nabla^{2} \phi\right)^{-1}\right) \\
& =-\operatorname{Tr}\left(\left(\nabla^{2} \phi\right)\left(\nabla^{2} \phi\right)^{-1}\left(\frac{\partial}{\partial x_{j}} \nabla^{2} \phi\right)\left(\nabla^{2} \phi\right)^{-1}\right) \\
& =-\sum_{k} e_{k}^{T}\left(\frac{\partial}{\partial x_{j}} \nabla^{2} \phi\right)\left(\nabla^{2} \phi\right)^{-1} e_{k} .
\end{aligned}
$$

Hence, we have

$$
\begin{aligned}
\frac{1}{2} e_{i}^{T}\left(\nabla^{2} \phi\right)^{-1} \nabla \log \operatorname{det}\left(\nabla^{2} \phi\right)^{-1} & =-\frac{1}{2} \sum_{j k}\left(\nabla^{2} \phi\right)_{i j}^{-1} e_{k}^{T}\left(\frac{\partial}{\partial x_{j}} \nabla^{2} \phi\right)\left(\nabla^{2} \phi\right)^{-1} e_{k} \\
& =-\frac{1}{2} \sum_{j k} e_{i}^{T}\left(\nabla^{2} \phi\right)^{-1}\left(\frac{\partial}{\partial x_{k}} \nabla^{2} \phi\right)\left(\nabla^{2} \phi\right)^{-1} e_{k}
\end{aligned}
$$

On the other hand, we have

$$
\begin{aligned}
\mu_{i} & =\frac{1}{2} \sum_{k} \frac{\partial}{\partial x_{k}}\left(\left(\nabla^{2} \phi\right)^{-1}\right)_{i k} \\
& =-\frac{1}{2} \sum_{k} e_{i}^{T}\left(\nabla^{2} \phi\right)^{-1}\left(\frac{\partial}{\partial x_{k}} \nabla^{2} \phi\right)\left(\nabla^{2} \phi\right)^{-1} e_{k} .
\end{aligned}
$$

To have a uniform stationary distribution, the geodesic walk uses a Metropolis filter. The transition probability before applying the filter is given as follows.

Lemma 14.4.4. For any $x \in M$ and $h>0$, the probability density of the 1 -step distribution from $x$ (before applying the Metropolis filter) is given by

$$
\begin{equation*}
p_{x}(y)=\sum_{v_{x}: \exp _{x}\left(v_{x}\right)=y} \operatorname{det}\left(d \exp _{x}\left(v_{x}\right)\right)^{-1} \sqrt{\frac{\operatorname{det}(g(y))}{(2 \pi h)^{n}}} \exp \left(-\frac{1}{2}\left\|\frac{v_{x}-\frac{h}{2} \mu(x)}{\sqrt{h}}\right\|_{x}^{2}\right) \tag{14.9}
\end{equation*}
$$

where $y=\exp _{x}\left(v_{x}\right)$ and $d \exp _{x}$ is the differential of the exponential map at $x$.
Proof. We prove the formula by separately considering each $v_{x} \in T_{x} M$ s.t. $\exp _{x}\left(v_{x}\right)=y$, then summing up. In the tangent space $T_{x} M$, the point $v_{x}$ is a Gaussian step. Therefore, the probability density of $v_{x}$ in $T_{x} M$ as follows.

$$
p_{x}^{T_{x} M}\left(v_{x}\right)=\frac{1}{(2 \pi h)^{n / 2}} \exp \left(-\frac{1}{2}\left\|\frac{v_{x}-\frac{h}{2} \mu(x)}{\sqrt{h}}\right\|_{x}^{2}\right) .
$$

Note that $v_{x}, \mu(x) \in T_{x} M$. Let $y=\exp _{x}\left(v_{x}\right)$. In the tangent space $T_{y} M$, we have that $y$ maps to 0 . Let $\phi: T_{x} M \rightarrow K$ defined by $F(v)=\operatorname{id}_{M \rightarrow K} \circ \exp _{x}(v)$. Here $K$ is the same set as $M$ but endowed with the Euclidean metric (we use $F$ instead of $\exp _{y}^{-1}\left(\exp _{x}(v)\right)$ since the exponential map might not be 1-1). Hence, we have

$$
d F\left(v_{x}\right)=d \operatorname{id}_{M \rightarrow K}(y) d \exp _{x}\left(v_{x}\right)
$$

The result follows from $p_{x}(y)=\operatorname{det}\left(d F\left(v_{x}\right)\right)^{-1} p_{x}^{T_{x} M}\left(v_{x}\right)$ and

$$
\begin{aligned}
\operatorname{det} d F\left(v_{x}\right) & =\operatorname{det}\left(d \operatorname{id}_{M \rightarrow K}(y)\right) \operatorname{det}\left(d \exp _{x}\left(v_{x}\right)\right) \\
& =\operatorname{det}(g(y))^{-1 / 2} \operatorname{det}\left(d \exp _{x}\left(v_{x}\right)\right)
\end{aligned}
$$

In Section 14.4.5, we bound the acceptance probability of the Metropolis filter.
Lemma 14.4.5. Under the condition on $h$ of Theorem 14.4.2,

$$
\left|\log \left(\frac{p_{x}(y)}{p_{y}(x)}\right)\right|=O\left(\sqrt{n} h^{3 / 2} D_{1}+(n h)^{3 / 2} G_{1}+\left(n h R_{1}\right)^{2}\right)<\frac{1}{4} .
$$

In Section 14.4.7, we bound the overlap of one-step distributions from nearby points.
Lemma 14.4.6. Let $x, y \in M$ be such that $d(x, y) \leq 0.1 \sqrt{h}$. Then, under the condition on $h$ of Theorem 14.4.2, the one-step distributions $P_{x}, P_{y}$ from $x, y$ satisfy

$$
d_{T V}\left(P_{x}, P_{y}\right)=O\left(n h^{3 / 2} R_{2}+h D_{2}+n h^{3 / 2} R_{1} D_{0}\right)<0.3
$$

## ■ 14.4.3 Isoperimetry

For a convex body $K$, the cross-ratio distance of $x$ and $y$ is

$$
d_{K}(x, y)=\frac{|x-y||p-q|}{|p-x||y-q|}
$$

where $p$ and $q$ are on the boundary of $K$ such that $p, x, y, q$ are on the straight line $\overline{x y}$ and are in order. In this section, we show that if the distance $d(x, y)$ induced by the Riemannian metric is upper bounded by the cross-ratio distance, then the body has good isoperimetric constant in terms of $d$. We note that although the cross-ratio distance is not a metric, the closely-related Hilbert distance is a metric:

$$
d_{H}(x, y)=\log \left(1+\frac{|x-y||p-q|}{|p-x||y-q|}\right) .
$$

Theorem 14.4.7. For a Hessian manifold $M$ with smoothness parameters $G_{2}$, for any partition of $M$ into three measurable subsets $S_{1}, S_{2}, S_{3}$, we have that

$$
\operatorname{vol}\left(S_{3}\right) \geq \frac{d\left(S_{1}, S_{2}\right)}{G_{2}} \min \left\{\operatorname{vol}\left(S_{1}\right), \operatorname{vol}\left(S_{2}\right)\right\}
$$

The theorem follows from the following isoperimetric inequality from [173], the definition of $G_{2}$ and the fact $d_{H} \leq d_{K}$.

Theorem 14.4.8 ([173]). For any convex body $K$ and any partition of $K$ into disjoint measurable subsets $S_{1}, S_{2}, S_{3}$

$$
\operatorname{vol}\left(S_{3}\right) \geq d_{K}\left(S_{1}, S_{2}\right) \operatorname{vol}\left(S_{1}\right) \operatorname{vol}\left(S_{2}\right)
$$

## - 14.4.4 Conductance and Mixing

Proof of Theorem 14.4.2. The proof follows the standard outline for geometric random walks (see e.g., [261]). Let $Q$ be the uniform distribution over $M$ and $S$ be any measurable subset of $M$. Then our goal is to show that

$$
\frac{\int_{S} P_{x}(M \backslash S) d Q(x)}{\min \{Q(S), Q(M \backslash S)\}}=\Omega\left(\frac{\sqrt{h}}{G_{2}}\right) .
$$

Since the Markov chain is time-reversible (For any two subsets $A, B, \int_{A} P_{x}(B) d x=\int_{B} P_{x}(A) d x$ ), we can write the numerator of the LHS above as

$$
\frac{1}{2}\left(\int_{S} P_{x}(M \backslash S) d Q(x)+\int_{M \backslash S} P_{x}(S) d Q(x)\right)
$$

Define

$$
\begin{aligned}
& S_{1}=\left\{x \in S: P_{x}(M \backslash S)<0.05\right\} \\
& S_{2}=\left\{x \in M \backslash S: P_{x}(S)<0.05\right\} \\
& S_{3}=M \backslash S_{1} \backslash S_{2} .
\end{aligned}
$$

We can assume wlog that $Q\left(S_{1}\right) \geq(1 / 2) Q(S)$ and $Q\left(S_{2}\right) \geq(1 / 2) Q(M \backslash S)$ (if not, the conductance is $\Omega(1)$ ).

Next, we note that for any two points $x \in S_{1}$ and $y \in S_{2}, d_{T V}\left(P_{x}, P_{y}\right)>0.9$. Therefore, by Lemma 14.4.6, we have that $d(x, y) \geq 0.1 \sqrt{h}$ and hence $d\left(S_{1}, S_{2}\right) \geq 0.1 \sqrt{h}$. Therefore, using Theorem 14.4.7,

$$
\operatorname{vol}\left(S_{3}\right) \geq \frac{0.1 \sqrt{h}}{G_{2}} \min \left\{\operatorname{vol}\left(S_{1}\right), \operatorname{vol}\left(S_{2}\right)\right\} .
$$

Going back to the conductance,

$$
\begin{aligned}
\frac{1}{2}\left(\int_{S} P_{x}(M \backslash S) d Q(x)+\int_{M \backslash S} P_{x}(S) d Q(x)\right) & \geq \frac{1}{2} \int_{S_{3}}(0.05) d Q(x) \\
& \geq \frac{\sqrt{h}}{800 G_{2}} \min \left\{\operatorname{vol}\left(S_{1}\right), \operatorname{vol}\left(S_{2}\right)\right\} \frac{1}{\operatorname{vol}(K)} \\
& =\frac{\sqrt{h}}{1600 G_{2}} \min \left\{\frac{\operatorname{vol}(S)}{\operatorname{vol}(M)}, \frac{\operatorname{vol}(M \backslash S)}{\operatorname{vol}(M)}\right\} \\
& =\frac{\sqrt{h}}{1600 G_{2}} \min \{Q(S), Q(M \backslash S)\}
\end{aligned}
$$

Therefore, $\phi(S) \geq(1 / 1600) \sqrt{h} / G_{2}$.

Corollary 14.4.9. Let $K$ be a polytope with $m$ facets. Let $Q$ be the uniform distribution over $K$ and $Q_{t}$ be the distribution obtained after $t$ steps of the geodesic walk started from a distribution $Q_{0}$ with $d_{0}=\sup _{K} \frac{d Q_{0}}{d Q}$. Then after $t>C\left(G_{2}^{2} / h\right) \log \left(\frac{d_{0}}{\varepsilon}\right)$ steps, with probability at least $1-\delta$, we have $d_{T V}\left(Q_{t}, Q\right) \leq \varepsilon$.

## - 14.4.5 Rejection Probability

The goal of this section is to prove Lemma 14.4.5, i.e., the rejection probability of the Metropolis filter is small. For a transition from $x$ to $y$ on the manifold, the filter is applied with respect to a randomly chosen $v_{x}$, i.e., for one geodesic from $x$ to $y$. We will bound the ratio of the transition probabilities (without the filter) as follows:

$$
\begin{aligned}
\log \left(\frac{p\left(x \xrightarrow{v_{x}} y\right)}{p\left(y \rightarrow v_{y} x\right)}\right)= & \log \left(\frac{\operatorname{det}\left(d \exp _{x}\left(v_{x}\right)\right)^{-1}}{\operatorname{det}\left(d \exp _{y}\left(v_{y}\right)\right)^{-1}}\right)+\frac{1}{2} \log \operatorname{det}(g(y))-\frac{1}{2} \log \operatorname{det}(g(x)) \\
& -\frac{1}{2}\left\|\frac{v_{x}-\frac{h}{2} \mu(x)}{\sqrt{h}}\right\|_{2}^{2}+\left\|\frac{v_{y}-\frac{h}{2} \mu(y)}{\sqrt{h}}\right\|_{2}^{2}
\end{aligned}
$$

Since a geodesic has constant speed, we have $\left\|v_{x}\right\|=\left\|v_{y}\right\|$. Therefore, we have that

$$
\begin{align*}
\log \left(\frac{p\left(x \xrightarrow{v_{x}} y\right)}{p\left(y \xrightarrow{v_{y}} x\right)}\right)= & \log \left(\frac{\operatorname{det}\left(d \exp _{y}\left(v_{y}\right)\right)}{\operatorname{det}\left(d \exp _{x}\left(v_{x}\right)\right)}\right)+\frac{1}{2} \log \operatorname{det}(g(y))-\frac{1}{2} \log \operatorname{det}(g(x))  \tag{14.10}\\
& +\frac{1}{2}\left\langle v_{x}, \mu(x)\right\rangle-\frac{h}{8}\|\mu(x)\|^{2}-\frac{1}{2}\left\langle v_{y}, \mu(y)\right\rangle+\frac{h}{8}\|\mu(y)\|^{2}
\end{align*}
$$

We separate the proof into three parts:

- $\left|\|\mu(x)\|^{2}-\|\mu(y)\|^{2}\right| \leq O\left(\sqrt{n h} D_{1}\right)$, immediate from the definition of $D_{1}$ and the fact that the $\|x-y\|=O(\sqrt{n h})$ whp.
- Sec 14.4.5.1: $\left|\log \operatorname{det}(g(y))-\log \operatorname{det}(g(x))+\left\langle v_{x}, \mu(x)\right\rangle-\left\langle v_{y} \mu(y)\right\rangle\right|=O\left((n h)^{3 / 2} G_{1}\right)$,
- Sec 14.4.5.2: Assuming that $h<\frac{1}{2 n R_{1}}$, we have $\left|\log \left(\frac{\operatorname{det}\left(d \exp _{y}\left(v_{y}\right)\right)}{\operatorname{det}\left(d \exp _{x}\left(v_{x}\right)\right)}\right)\right|=O\left(\left(n h R_{1}\right)^{2}\right)$.

Together, these facts imply the lemma.

### 14.4.5.1 Trapezoidal Rule for Volumetric Barrier

Recall that the Trapezoidal rule is to approximate $\int_{0}^{h} f(t) d t$ by $\frac{h}{2}(f(0)+f(h))$. The nice thing about this rule is that the error is $O\left(h^{3}\right)$ instead of $O\left(h^{2}\right)$ because the second order term cancels by symmetry. Our main observation here is that the geodesic walk is implicitly following a trapezoidal rule on the metric and hence it has a small error.

Lemma 14.4.10. We have that

$$
\left|\int_{0}^{\ell} f^{\prime}(t) d t-\frac{\ell}{2}\left(f^{\prime}(0)+f^{\prime}(\ell)\right)\right| \leq \frac{\ell^{3}}{12} \max _{0 \leq t \leq \ell}\left|f^{\prime \prime \prime}(t)\right|
$$

Proof. Note that

$$
\begin{aligned}
\int_{0}^{\ell} f^{\prime}(t) d t-\frac{\ell}{2}\left(f^{\prime}(0)+f^{\prime}(\ell)\right) & =\int_{0}^{\ell}\left(f^{\prime}(0)+\int_{0}^{t} f^{\prime \prime}(s) d s\right) d t-\ell f^{\prime}(0)-\frac{\ell}{2} \int_{0}^{\ell} f^{\prime \prime}(s) d s \\
& =\int_{0}^{\ell} \int_{0}^{t} f^{\prime \prime}(s) d s d t-\frac{\ell}{2} \int_{0}^{\ell} f^{\prime \prime}(s) d s \\
& =\int_{0}^{\ell}\left(\frac{\ell}{2}-s\right) f^{\prime \prime}(s) d s \\
& =\int_{0}^{\ell}\left(\frac{\ell}{2}-s\right)\left(f^{\prime \prime}(0)+\int_{0}^{s} f^{\prime \prime \prime}(t) d t\right) d s \\
& =\int_{0}^{\ell}\left(\frac{\ell}{2}-s\right) \int_{0}^{s} f^{\prime \prime \prime}(t) d t d s \\
& \leq \frac{\ell^{3}}{12} \max _{0 \leq t \leq \ell}\left|f^{\prime \prime \prime}(t)\right|
\end{aligned}
$$

We apply this to the logdet function.
Lemma 14.4.11. Let $f(t)=\log \operatorname{det}\left(g(\gamma(t))\right.$ where $\gamma(t)=\exp _{x}\left(\frac{t}{\ell} v_{x}\right) \in \Pi$ (See Definition of $\Pi$ in Section 14.4.1). Then,

$$
\left|\log \operatorname{det}(g(y))-\log \operatorname{det}(g(x))+\left\langle v_{x}, \mu(x)\right\rangle_{x}-\left\langle v_{y}, \mu(y)\right\rangle_{y}\right|=O\left((n h)^{3 / 2} G_{1}\right)
$$

Proof. Let $f(t)=\log \operatorname{det} g(\gamma(t))$ and $z=\gamma(t)$. By Lemma 14.4.3, $\mu(z)=-\frac{1}{2} g(z)^{-1} \nabla f(z)$. Using this,

$$
\begin{aligned}
f^{\prime}(t) & =\left\langle\nabla_{z} \log \operatorname{det} g(z), \gamma^{\prime}(t)\right\rangle_{2} \\
& =\left\langle g(z)^{-1} \nabla_{z} \log \operatorname{det} g(z), \gamma^{\prime}(t)\right\rangle_{z} \\
& =-\left\langle 2 \mu(z), \gamma^{\prime}(t)\right\rangle_{z}
\end{aligned}
$$

Noting that $v_{x}=\ell \gamma^{\prime}(0)$ and $v_{y}=-\ell \gamma^{\prime}(\ell)$, and using Lemma 14.4.10, we have

$$
\begin{aligned}
& \left|\log \operatorname{det}(g(y))-\log \operatorname{det}(g(x))+\left\langle v_{x}, \mu(x)\right\rangle_{x}-\left\langle v_{y}, \mu(y)\right\rangle_{y}\right| \\
= & \left|\log \operatorname{det}(g(y))-\log \operatorname{det}(g(x))+\ell\left(\left\langle\gamma^{\prime}(0), \mu(x)\right\rangle_{x}+\left\langle\gamma^{\prime}(\ell), \mu(y)\right\rangle_{y}\right)\right| \\
= & \left|\int_{0}^{\ell} f^{\prime}(t) d t-\frac{\ell}{2}\left(f^{\prime}(0)+f^{\prime}(\ell)\right)\right| \\
\leq & \frac{\ell^{3}}{12} \max _{0 \leq t \leq \ell}\left|f^{\prime \prime \prime}(t)\right|=O\left((n h)^{3 / 2} G_{1}\right)
\end{aligned}
$$

### 14.4.5.2 Smoothness of exponential map

First, we show the relation between the differential of exponential map $d \exp _{x}\left(v_{x}\right)$ and the Jacobi field along the geodesic $\exp _{x}\left(t v_{x}\right)$.

Lemma 14.4.12. Given a geodesic $\gamma(t)=\exp _{x}\left(\frac{t}{\ell} v_{x}\right) \in \Pi$, let $\left\{X_{i}(t)\right\}_{i=1}^{n}$ be the parallel transport of some orthonormal frame along $\gamma(t)$. Then, for any $w \in \mathbb{R}^{n}$, we have that

$$
d \exp _{x}\left(v_{x}\right)\left(\sum w_{i} X_{i}(0)\right)=\sum_{i} \psi_{w}(\ell)_{i} X_{i}(\ell)
$$

where $\psi_{w}$ satisfies the Jacobi field along $\gamma(t)$, i.e.,

$$
\begin{align*}
\frac{d^{2}}{d t^{2}} \psi_{w}(t)+R(t) \psi_{w}(t) & =0 \\
\frac{d}{d t} \psi_{w}(0) & =w / \ell  \tag{14.11}\\
\psi_{w}(t) & =0
\end{align*}
$$

and $R(t)$ is a matrix given by $R_{i j}(t)=\left\langle R\left(X_{i}(t), \gamma^{\prime}(t)\right) \gamma^{\prime}(t), X_{j}(t)\right\rangle$.
Proof. We want to compute $d \exp _{x}\left(v_{x}\right)(w)$ for some $w \in T_{y} M$. By definition, we have that

$$
d \exp _{x}\left(v_{x}\right)(w)=\left.\frac{d}{d s} \gamma(t, s)\right|_{t=\ell, s=0}
$$

where $\gamma(t, s)=\exp _{x}\left(t v_{x} / \ell+s w\right)$. Let $\eta_{w}(t)=\left.\frac{d}{d s} \gamma(t, s)\right|_{s=0}$ be a Jacobi field given by the formula

$$
\begin{aligned}
D_{t} D_{t} \eta_{w}+R\left(\eta_{w}, \gamma^{\prime}(t)\right) \gamma^{\prime}(t) & =0, \text { for } 0 \leq t \leq \ell \\
D_{t} \eta_{w}(0) & =w / \ell, \\
\dot{\eta}_{w}(0) & =0 .
\end{aligned}
$$

Recall that the parallel transport of an orthonormal frame remains orthonormal because $\frac{d}{d t}\left\langle X_{i}, X_{j}\right\rangle=$ $\left\langle D_{t} X_{i}, X_{j}\right\rangle+\left\langle X_{i}, D_{t} X_{j}\right\rangle=0$. Since $X_{i}(t)$ is an orthonormal basis at $T_{\gamma(t)} M$, we can write

$$
\eta_{w}(t)=\sum_{i} \psi_{i}(t) X_{i}(t)
$$

Since $\eta_{w}$ satisfies the ODE above, we have

$$
D_{t} D_{t} \sum_{i} \psi_{i}(t) X_{i}(t)+R\left(\sum_{i} \psi_{i}(t) X_{i}(t), \gamma^{\prime}(t)\right) \gamma^{\prime}(t)=0 .
$$

Since $X_{i}(t)$ is a parallel transport, we have that $D_{t} D_{t}\left(\psi_{i}(t) X_{i}(t)\right)=\frac{d^{2} \psi_{i}(t)}{d t^{2}} X_{i}(t)$ and hence

$$
\sum_{i} \frac{d^{2}}{d t^{2}} \psi_{i}(t) X_{i}(t)+\sum_{i} \psi_{i}(t) R\left(X_{i}(t), \gamma^{\prime}(t)\right) \gamma^{\prime}(t)=0
$$

Let $R_{i j}(t)=\left\langle R\left(X_{i}(t), \gamma^{\prime}(t)\right) \gamma^{\prime}(t), X_{j}(t)\right\rangle$. Since $R(t)$ is a symmetric matrix (Fact 14.2.3), we have

$$
\frac{d^{2}}{d t^{2}} \psi(t)+R(t) \psi(t)=0
$$

Hence, we have

$$
d \exp _{x}\left(v_{x}\right)(w)=\psi(\ell)
$$

Next, we have a lemma about determinant.

Lemma 14.4.13. Suppose that $E$ is a matrix (not necessarily symmetric) with $\|E\|_{F} \leq \frac{1}{4}$, we have

$$
|\log \operatorname{det}(I+E)-\operatorname{Tr} E| \leq\|E\|_{F}^{2} .
$$

Proof. Let $f(t)=\log \operatorname{det}(I+t E)$. Then, by Jacobi's formula, we have

$$
\begin{aligned}
f^{\prime}(t) & =\operatorname{Tr}\left((I+t E)^{-1} E\right) \\
f^{\prime \prime}(t) & =-\operatorname{Tr}\left((I+t E)^{-1} E(I+t E)^{-1} E\right)
\end{aligned}
$$

Therefore, we have

$$
\begin{aligned}
f(1) & =f(0)+f^{\prime}(0)+\int_{0}^{1}(1-s) f^{\prime \prime}(s) d s \\
& =\operatorname{Tr}(E)+\int_{0}^{1}(1-s) f^{\prime \prime}(s) d s
\end{aligned}
$$

We note that

$$
\begin{aligned}
\left|f^{\prime \prime}(s)\right| & =\left|\operatorname{Tr}\left((I+t E)^{-1} E(I+t E)^{-1} E\right)\right| \\
& \leq\left|\operatorname{Tr}\left(E^{T}\left((I+t E)^{-1}\right)^{T}(I+t E)^{-1} E\right)\right|
\end{aligned}
$$

Since $\|E\|_{F} \leq \frac{1}{4}$, we have $\|E\|_{2} \leq \frac{1}{4}$ and hence $\left\|(I+t E)^{-1}\right\|_{2} \leq \frac{4}{3}$ for all $0 \leq t \leq 1$. Therefore, we have

$$
\left|f^{\prime \prime}(s)\right| \leq 2\left|\operatorname{Tr}\left(E^{T} E\right)\right|=2\|E\|_{F}^{2}
$$

Therefore, we have

$$
|f(1)-\operatorname{Tr} E| \leq\|E\|_{F}^{2} .
$$

Using lemma 14.4.13 and lemma 14.4.21, we have the following:
Lemma 14.4.14. Given a geodesic walk $\gamma(t)=\exp _{x}\left(\frac{t}{\ell} v_{x}\right) \in \Pi$ with step size $h$ satisfying $0<h \leq \frac{1}{n R_{1}}$. We have that $d \exp _{x}\left(v_{x}\right)$ is invertible and

$$
\begin{equation*}
\left|\log \operatorname{det}\left(d \exp _{x}\left(v_{x}\right)\right)-\int_{0}^{\ell} \frac{s(\ell-s)}{\ell} \operatorname{Ric}\left(\gamma^{\prime}(s)\right) d s\right| \leq \frac{\left(n h R_{1}\right)^{2}}{6} . \tag{14.12}
\end{equation*}
$$

Therefore, we have that

$$
\left|\log \operatorname{det}\left(d \exp _{x}\left(v_{x}\right)\right)-\log \operatorname{det}\left(d \exp _{y}\left(v_{y}\right)\right)\right| \leq \frac{\left(n h R_{1}\right)^{2}}{3}
$$

Proof. Let $\Psi$ be the solution of the $\operatorname{ODE} \Psi^{\prime \prime}(t)+R(t) \Psi(t)=0, \Psi^{\prime}(0)=I / \ell$ and $\Psi(0)=0$. We know that $\|R(t)\|_{2} \leq\|R(t)\|_{F} \leq R_{1}$ for all $0 \leq t \leq \ell$. Hence, Lemma 14.4.21 shows that

$$
\begin{align*}
\left\|\Psi(t)-\frac{t}{\ell} I\right\|_{F} & \leq \max _{0 \leq s \leq \ell}\|R(s)\|_{F}\left(\frac{t^{3}}{5}\|I / \ell\|_{2}\right) \\
& \leq \frac{1}{5} R_{1} n h \leq \frac{1}{4} \tag{14.13}
\end{align*}
$$

By the Lemma 14.4.13, we have that

$$
|\log \operatorname{det}(\Psi(\ell))-\operatorname{Tr}(\Psi(\ell)-I)| \leq\left(\frac{1}{5} R_{1} n h\right)^{2}
$$

Now, we need to estimate $\operatorname{Tr}(\Psi(\ell)-I)$. Note that

$$
\begin{aligned}
\Psi(\ell) & =\Psi(0)+\Psi^{\prime}(0) \ell+\int_{0}^{\ell}(\ell-s) R(s) \Psi(s) d s \\
& =I+\int_{0}^{\ell}(\ell-s) R(s) \Psi(s) d s
\end{aligned}
$$

Hence, we have

$$
\Psi(\ell)-I-\int_{0}^{\ell} \frac{s(\ell-s)}{\ell} R(s) d s=\int_{0}^{\ell}(\ell-s) R(s)\left(\Psi(s)-\frac{s}{\ell} I\right) d s
$$

Using Lemma (14.13), we have

$$
\begin{aligned}
\left|\operatorname{Tr}\left(\Psi(\ell)-I-\int_{0}^{\ell} \frac{s(\ell-s)}{\ell} R(s) d s\right)\right| & \leq \int_{0}^{\ell}(\ell-s)\left|\operatorname{Tr} R(s)\left(\Psi(s)-\frac{s}{\ell} I\right)\right| d s \\
& \leq \int_{0}^{\ell}(\ell-s)\|R(s)\|_{F}\left\|\Psi(s)-\frac{s}{\ell} I\right\| d s \\
& \leq \frac{\ell^{2}}{2} \cdot R_{1} \cdot \frac{1}{5} R_{1} n h \leq \frac{\left(n h R_{1}\right)^{2}}{10}
\end{aligned}
$$

Lemma 14.4.12 shows that $d \exp _{x}\left(v_{x}\right)\left(\sum w_{i} X_{i}(0)\right)=\sum_{i} \psi_{w}(\ell)_{i} X(\ell)=\sum_{i}(\Psi(\ell) w)_{i} X_{i}(\ell)$. Since $\left\{X_{i}(t)\right\}_{i=1}^{n}$ are orthonormal, this shows that

$$
d \exp _{x}\left(v_{x}\right)=X(\ell) \Psi(\ell) X(0)^{T}
$$

where $X$ is the matrix $\left[X_{1}, X_{2}, \cdots, X_{n}\right]$. Since $\|\Psi(\ell)-I\|_{F} \leq \frac{1}{5}$ (14.13), we have that $\Psi(\ell)$ is invertible and so is $d \exp _{x}\left(v_{x}\right)$.

Since $X(\ell)$ and $X(0)$ are orthonormal, we have that

$$
\log \operatorname{det}\left(d \exp _{x}\left(v_{x}\right)\right)=\log \operatorname{det} \Psi(\ell) .
$$

Therefore, this gives

$$
\left|\log \operatorname{det}\left(d \exp _{x}\left(v_{x}\right)\right)-\int_{0}^{\ell} \frac{s(\ell-s)}{\ell} \operatorname{Tr} R(s) d s\right| \leq \frac{\left(n h R_{1}\right)^{2}}{6} .
$$

By the definition of Ricci curvature and Fact 14.2.3, we have that

$$
\begin{aligned}
\operatorname{Tr} R(s) & =\sum_{i}\left\langle R\left(X_{i}(s), \gamma^{\prime}(s)\right) \gamma^{\prime}(s), X_{j}(s)\right\rangle \\
& =\operatorname{Ric}\left(\gamma^{\prime}(s)\right) .
\end{aligned}
$$

This gives the result (14.12).
Since the geodesic $\exp _{x}\left(t v_{x}\right)$ is the same as $\exp _{y}\left(t v_{y}\right)$ except for swapping the parameterization, and since $\int_{0}^{\ell} \frac{s(\ell-s)}{\ell} \operatorname{Tr} R(s) d s$ is invariant under this swap, (14.12) implies that $\log \operatorname{det}\left(d \exp _{x}\left(v_{x}\right)\right)$ is close to $\log \operatorname{det}\left(d \exp _{y}\left(v_{y}\right)\right)$.

## ■ 14.4.6 Implementation

Here, we explain in high level how to implement the geodesic walk in general via an efficient algorithm for approximately solving ODEs. Note that to implement the step, we need to compute the geodesic and compute the probability $p(x \xrightarrow{w} y)$ and $p\left(y \xrightarrow{w^{\prime}} x\right)$. From the formula (14.10), we see that they
involves the term $d \exp _{x}(w)$. Lemma 14.4.12 shows that $d \exp _{x}(w)$ can be computed by Jacobi field. Also, Lemma 14.4.12 shows that Jacobi field can be computed by a ODE if it is written in the orthonormal frame systems. Therefore, to implement the geodesic walk, we need to compute geodesic, parallel transport and Jacobi field (See Algo 49). All of these are ODEs and can be solved using the collocation method. In the later sections, we will see how to collocation method can indeed solves these ODEs in matrix multiplication time.

```
Algorithm 49: Geodesic Walk (Detailed)
    Pick a Gaussian random vector \(w \sim N_{x}(0,1)\), i.e. \(\mathbf{E}_{w}\|w\|_{x}^{2}=1\).
    /* Compute \(y=\exp _{x}\left(\sqrt{h} w+\frac{h}{2} \mu(x)\right)\) where \(\mu(x)\) is given by (14.2). */
    /* Compute a corresponding \(w^{\prime}\) s.t. \(\quad x=\exp _{y}\left(\sqrt{h} w^{\prime}+\frac{h}{2} \mu(y)\right)\). */
```

    Generate a random direction \(d=\sqrt{h} w+\frac{h}{2} \mu(x)\) where \(\mu(x)\) is given by (14.2).
    Solve the geodesic equation \(D_{\gamma^{\prime}} \gamma^{\prime}=0\) with \(\gamma(0)=x\) and \(\gamma^{\prime}(0)=d\) using collocation method
        (Sec 14.6).
    Set \(y=\gamma(1)\) and \(w^{\prime}=-\gamma^{\prime}(1)\).
    /* Compute the probability \(p(x \xrightarrow{w} y)\) of going from \(x\) to \(y\) using the step \(w\). */
    Pick an orthonormal frame \(\left\{X_{i}\right\}_{i=1}^{n}\) at \(x\).
    Compute the parallel transport of \(\left\{X_{i}\right\}_{i=1}^{n}\) along \(\gamma(t)\) using collocation method.
    Compute \(d \exp _{x}(w)\) via solve the Jacobi field equation (14.11) (in the coordinate systems
        \(\left.\left\{X_{i}(t)\right\}\right)\).
    Set \(p(x \xrightarrow{w} y)=\operatorname{det}\left(d \exp _{x}(w)\right)^{-1} \sqrt{\frac{\operatorname{det}(g(y))}{(2 \pi h)^{n}}} \exp \left(-\frac{1}{2}\left\|\frac{w-\frac{h}{2} \mu(x)}{\sqrt{h}}\right\|_{x}^{2}\right)\).
    Compute \(p\left(y \xrightarrow{w^{\prime}} x\right)\) similarly.
    With probability $\min \left(1, \frac{p\left(y \xrightarrow{w^{\prime}} x\right)}{p(x \xrightarrow{w} y)}\right)$, go to $y$; otherwise, stay at $x$.

## ■ 14.4.7 Smoothness of $p_{x}$

Here we prove Lemma 14.4.6. Recall that the probability density of going from $x$ to $y$ is given by the following formula:

$$
p_{x}(y)=\sum_{v_{x}: \exp _{x}\left(v_{x}\right)=y} \operatorname{det}\left(d \exp _{x}\left(v_{x}\right)\right)^{-1} \sqrt{\frac{\operatorname{det}(g(y))}{(2 \pi h)^{n}}} \exp \left(-\frac{1}{2}\left\|\frac{v_{x}-\frac{h}{2} \mu(x)}{\sqrt{h}}\right\|_{x}^{2}\right)
$$

To simplify the calculation, we apply Lemma 14.4 .14 and consider the following estimate of $p_{x}(y)$ instead

$$
\tilde{p}_{x}(y)=\sum_{v_{x}:{e x p p_{x}}\left(v_{x}\right)=y} \sqrt{\frac{\operatorname{det}(g(y))}{(2 \pi h)^{n}}} \exp \left(-\int_{0}^{\ell} \frac{t(\ell-t)}{\ell} \operatorname{Ric}\left(\gamma^{\prime}(t)\right) d t-\frac{1}{2}\left\|\frac{v_{x}-\frac{h}{2} \mu(x)}{\sqrt{h}}\right\|_{x}^{2}\right)
$$

where $\gamma(t)$ be the geodesic from $\gamma(0)=x$ to $\gamma(\ell)=y$. Lemma 14.4.14 shows that $\left|\log \left(\tilde{p}_{x}(y) / p_{x}(y)\right)\right| \leq$ $\frac{\left(n h R_{1}\right)^{2}}{6}$ and hence it suffices to prove the smoothness of $\tilde{p}_{x}(y)$.

Let $c(s)$ be an unit speed geodesic going from $x$ to some point $z$ infinitesimally close to $x$. For (almost) any $v_{x}$ such that $\exp _{x}\left(v_{x}\right)=y$, Lemma 14.4.24 shows that there is an unique vector field $v(s)$
on $c(s)$ such that $\exp _{c(s)}(v(s))=y$ and $v(0)=v_{x}$. Now, we define

$$
\zeta(s)=\sqrt{\frac{\operatorname{det}(g(y))}{(2 \pi h)^{n}}} \exp \left(-\int_{0}^{\ell} \frac{t(\ell-t)}{\ell} \operatorname{Ric}\left(\gamma_{s}^{\prime}(t)\right) d t-\frac{1}{2}\left\|\frac{v(s)-\frac{h}{2} \mu(c(s))}{\sqrt{h}}\right\|_{c(s)}^{2}\right)
$$

where $\exp _{c(s)}(v(s))=y$ and $\gamma_{s}(t)=\exp _{c(s)}\left(\frac{t}{\ell} v(s)\right)$. Then, we have that

$$
\tilde{p}_{c(s)}(y)=\sum_{v(s): \exp _{c(s)}(v(s))=y} \zeta(s)
$$

and hence

$$
\frac{d}{d s} \tilde{p}_{c(s)}(y)=\sum_{v(s): \exp _{c(s)}(v(s))=y}\left(-\int_{0}^{\ell} \frac{t(\ell-t)}{\ell} \frac{d}{d s} \operatorname{Ric}\left(\gamma_{s}^{\prime}(t)\right) d t-\frac{1}{2} \frac{d}{d s}\left\|\frac{v(s)-\frac{h}{2} \mu(c(s))}{\sqrt{h}}\right\|_{c(s)}^{2}\right) \zeta(s) .
$$

Hence, it suffices to bound the terms in the bracket.
In Section 14.4.7.1, we analyze $\frac{d}{d s} \operatorname{Ric}\left(\gamma_{s}^{\prime}(t)\right)$ and prove that

$$
\left|\int_{0}^{\ell} \frac{t(\ell-t)}{\ell} \frac{d}{d s} \operatorname{Ric}\left(\gamma_{s}^{\prime}(t)\right) d t\right| \leq 2 n h R_{2} .
$$

In Section 14.4.7.2, we analyze $\frac{d}{d s}\left\|\frac{v(s)-\frac{h}{2} \mu(c(s))}{\sqrt{h}}\right\|_{c(s)}^{2}$ and prove that

$$
\left.\mathbf{E}_{\gamma}\left|\frac{d}{d s}\left\|\frac{v(s)-\frac{h}{2} \mu(c(s))}{\sqrt{h}}\right\|_{c(s)}^{2}\right|_{s=0} \right\rvert\,=O\left(D_{2} \sqrt{h}+\frac{1}{\sqrt{h}}+D_{0}\right) .
$$

Therefore, we have that

$$
\left|\frac{d}{d s} \tilde{p}_{c}(y)\right|=O\left(n h R_{2}+\sqrt{h} D_{2}+\frac{1}{\sqrt{h}}+n h R_{1} D_{0}\right) \tilde{p}_{c}(y) .
$$

Using $\left|\log \left(\tilde{p}_{x}(y) / p_{x}(y)\right)\right| \leq \frac{\left(n h R_{1}\right)^{2}}{6}$, we have our result.
Theorem 14.4.15. We have that

$$
d_{T V}\left(p_{x}, p_{y}\right)=O\left(n h R_{2}+\sqrt{h} D_{2}+\frac{1}{\sqrt{h}}+n h R_{1} D_{0}\right) d(x, y)+O\left(n h R_{1}\right)^{2} .
$$

### 14.4.7.1 Smoothness of exponential map

Definition 14.4.16. Given a manifold $M$ and a geodesic walk $\gamma(t)$ on $M_{L}$ with step size $h$. Let $R_{2}$ be a constant depending on the manifold $M$ and the step size $h$ such that for any $t$ such that $0 \leq t \leq \ell$, any curve $\alpha(s)$ starting from $\gamma(t)$ and any vector field $u(s)$ on $\alpha(s)$, we have that

$$
\left.\left|\frac{d}{d s} \operatorname{Ric}(u(s))\right|_{s=0} \right\rvert\, \leq\left(\left\|\alpha^{\prime}(0)\right\|+\left\|D_{s} u(0)\right\|\right) R_{2} .
$$

Lemma 14.4.17. For $\frac{1}{n}<h \leq \frac{1}{2 n R_{1}}$, we have

$$
\left|\int_{0}^{\ell} \frac{t(\ell-t)}{\ell} \frac{d}{d s} \operatorname{Ric}\left(\gamma_{s}^{\prime}(t)\right) d t\right| \leq 2 n h R_{2}
$$

where $\gamma_{s}$ is as defined in the beginning of Section 14.4.7.
Proof. Fix any $t$ such that $0 \leq t \leq \ell$ and let $\alpha(s)=\gamma_{s}(t)$ and $u(s)=\gamma_{s}^{\prime}(t)$ be a vector field on $\alpha$. By Definition 14.4.16, we have that

$$
\left.\left|\frac{d}{d s} \operatorname{Ric}(u(s))\right|_{s=0} \right\rvert\, \leq\left(\left\|\alpha^{\prime}(0)\right\|+\left\|D_{s} u(0)\right\|\right) R_{2}
$$

Recall that $\gamma_{s}=\exp _{c(s)}\left(\frac{t}{\ell} v(s)\right)$ and hence

$$
\begin{aligned}
\frac{d \alpha}{d s}(0) & =\left.\frac{\partial}{\partial s} \exp _{c(s)}\left(\frac{t}{\ell} v(s)\right)\right|_{s=0} \\
D_{s} u(0) & =\left.D_{s} \frac{\partial}{\partial t} \exp _{c(s)}\left(\frac{t}{\ell} v(s)\right)\right|_{s=0} \\
& =\left.D_{t} \frac{\partial}{\partial s} \exp _{c(s)}\left(\frac{t}{\ell} v(s)\right)\right|_{s=0}
\end{aligned}
$$

Let $\psi(t)=\left.\frac{\partial}{\partial s} \exp _{c(s)}\left(\frac{t}{\ell} v(s)\right)\right|_{s=0}$. Since $\exp _{c(s)}\left(\frac{t}{\ell} v(s)\right)$ is a geodesic for each $s, \psi(t)$ satisfies the Jacobi field equation. It is easier to calculate using orthogonal frame. So, we pick an arbitrary orthogonal frame along the curve $\gamma(t)$ and let $\phi(t)$ be $\psi(t)$ represented in that orthogonal frame. Then, we know that

$$
\begin{aligned}
\frac{d^{2}}{d t^{2}} \phi(t)+R\left(\phi(t), \gamma^{\prime}(t)\right) \gamma^{\prime}(t) & =0 \\
\phi(0) & =\alpha \\
\phi^{\prime}(0) & =\frac{\beta}{\ell}
\end{aligned}
$$

for some $\alpha$ and $\beta$. Since $c$ is a unit speed geodesic, we have that $\alpha=\phi(0)=\frac{\partial c}{\partial s}(0)$ has $\ell_{2}$ norm 1 . Next, we note that $\beta=\ell \phi^{\prime}(0)=D_{s} v(s)$. By lemma 14.4.22, we have that $\|\beta\|_{2} \leq \frac{3}{2}$.

Therefore, Lemma 14.4.20 shows that

$$
\left\|\phi(t)-\alpha-\frac{\beta}{\ell} t\right\|_{2} \leq 2 R_{1} \ell^{2}
$$

for all $0 \leq t \leq \frac{1}{\sqrt{R_{1}}}$. Hence, we have $\|\phi(t)\|_{2} \leq 4$ for all $0 \leq t \leq \frac{1}{\sqrt{R_{1}}}$. By Lemma 14.4.20 again, we have

$$
\left\|\phi^{\prime}(t)-\beta\right\|_{2} \leq 3 R_{1} \ell
$$

for all $0 \leq t \leq \frac{1}{\sqrt{R_{1}}}$. Hence, we have $\left\|\phi^{\prime}(t)\right\|_{2} \leq 4$ for all $0 \leq t \leq \frac{1}{\sqrt{R_{1}}}$.
Therefore, we have

$$
\begin{aligned}
\left.\left|\frac{d}{d s} \operatorname{Ric}(u(s))\right|_{s=0} \right\rvert\, & \leq\left(\left\|\alpha^{\prime}(0)\right\|+\left\|D_{s} u(0)\right\|\right) R_{2} \\
& =\left(\|\phi(t)\|_{2}+\left\|\phi^{\prime}(t)\right\|_{2}\right) R_{2} \\
& \leq 8 R_{2}
\end{aligned}
$$

This gives the result.

### 14.4.7.2 One-step overlap

We parametrize the geodesic from $x$ to $y$ as $\gamma(s)$. For a fixed $z$, we have $z=\exp _{x}\left(d_{x}\right)=\exp _{\gamma(s)}\left(d_{\gamma(s)}\right)=$ $\exp _{\gamma(s)}(d(s))$.

Lemma 14.4.18. We have that

$$
\left.\mathbf{E}_{\gamma}\left|\frac{d}{d s}\left\|\frac{v(s)-\frac{h}{2} \mu(c(s))}{\sqrt{h}}\right\|_{c(s)}^{2}\right|_{s=0} \right\rvert\,=O\left(D_{2} \sqrt{h}+\frac{1}{\sqrt{h}}+D_{0}\right) .
$$

where $c(s)$ is as defined in the beginning of Section 14.4.7.
Proof. Note that

$$
\begin{equation*}
\left.\frac{d}{d s}\left\|\frac{v(s)-\frac{h}{2} \mu(c(s))}{\sqrt{h}}\right\|_{c(s)}^{2}\right|_{s=0}=\frac{2}{h}\left\langle\left. D_{s} v\right|_{s=0}-\left.\frac{h}{2} D_{s} \mu\right|_{s=0}, v-\frac{h}{2} \mu\right\rangle . \tag{14.14}
\end{equation*}
$$

Since $v-\frac{h}{2} \mu$ is a random Gaussian vector $N(0, n h I)$ in $T_{c} M$ independent of $D_{s} \mu$, we have that

$$
\begin{equation*}
\mathbf{E}\left|\left\langle\left.\frac{h}{2} D_{s} \mu\right|_{s=0}, v-\frac{h}{2} \mu\right\rangle\right| \leq O\left(\frac{\left\|\left.\frac{h}{2} D_{s} \mu\right|_{s=0}\right\| \sqrt{n h}}{\sqrt{n}}\right) \leq O\left(\frac{h D_{2} \sqrt{n h}}{2 \sqrt{n}}\right) . \tag{14.15}
\end{equation*}
$$

By Lemma 14.4.22, we know that $\left.D_{s} v\right|_{s=0}=-c^{\prime}+\zeta$ where $\zeta \perp v(0)$ and $\|\zeta\| \leq \frac{3}{2} n h R_{1}\left\|c^{\prime}\right\|$. Since $v-\frac{h}{2} \mu$ is a random Gaussian vector independent of $c^{\prime}$, we have that

$$
\begin{align*}
\mathbf{E}\left|\left\langle D_{s} v, v-\frac{h}{2} \mu\right\rangle\right| & \leq \mathbf{E}\left|\left\langle c^{\prime}, v-\frac{h}{2} \mu\right\rangle\right|+\mathbf{E} \frac{h}{2}|\langle\zeta, \mu\rangle| \\
& \leq O\left(\frac{\left\|c^{\prime}\right\| \sqrt{n h}}{\sqrt{n}}\right)+n h^{2} R_{1}\left\|c^{\prime}\right\|\|\mu\| \\
& \leq O\left(\sqrt{h}+n h^{2} R_{1} D_{0}\right) . \tag{14.16}
\end{align*}
$$

Combining the bounds (14.14), (14.15) and (14.16), we have that

$$
\left.\mathbf{E}\left|\frac{d}{d s}\left\|\frac{v(s)-\frac{h}{2} \mu(c(s))}{\sqrt{h}}\right\|_{c(s)}^{2}\right|_{s=0} \right\rvert\, \leq O\left(D_{2} \sqrt{h}+\frac{1}{\sqrt{h}}+n h R_{1} D_{0}\right) .
$$

## ■ 14.4.8 Approximate Solution of Jacobi Field

Let $\gamma(t)$ be a geodesic and $\left\{X_{i}(t)\right\}_{i=1}^{n}$ be the parallel transport of some orthonormal frame along $\gamma(t)$. As we demonstrated in the proof of Lemma 14.4.12, Jacobi fields are solutions the following matrix ODE:

$$
\begin{align*}
\frac{d^{2}}{d t^{2}} \psi(t)+R(t) \psi(t) & =0, \\
\frac{d}{d t} \psi(0) & =\beta,  \tag{14.17}\\
\psi(0) & =\alpha
\end{align*}
$$

where $\psi(t), \alpha, \beta \in \mathbb{R}^{n}$ and $R(t)$ is a matrix given by $R_{j i}(t)=\left\langle R\left(X_{i}(t), \gamma^{\prime}(t)\right) \gamma^{\prime}(t), X_{j}(t)\right\rangle$. Recall that $R(t)$ is a symmetric matrix (Fact 14.2.3).

In this section, we show give an approximate solution of the the Jacobi equation (14.17). First, we give a basic bound on the solution in terms of hyperbolic sine and cosine functions.

Lemma 14.4.19. Let $\psi$ be the solution of (14.17). Suppose that $\|R(t)\|_{2} \leq \lambda$. Then, we have that

$$
\|\psi(t)\|_{2} \leq\|\alpha\|_{2} \cosh (\sqrt{\lambda} t)+\frac{\|\beta\|_{2}}{\sqrt{\lambda}} \sinh (\sqrt{\lambda} t)
$$

for all $t \geq 0$.
Proof. Let $a(t)=\|\psi(t)\|_{2}$. Using $\|R(t)\|_{2} \leq \lambda$ and $\frac{d^{2}}{d t^{2}} \psi(t)+R(t) \psi(t)=0$, we have that

$$
a^{\prime \prime}(t) \leq \lambda a(t) .
$$

Since $a(0)=\|\alpha\|_{2}$ and $a^{\prime}(0) \leq\|\beta\|_{2}$, for all $t \geq 0$, we have that $a(t) \leq b(t)$ where $b(t)$ is the solution of the $b^{\prime \prime}(t)=\lambda b(t), b(0)=\|\alpha\|_{2}, b^{\prime}(0)=\|\beta\|_{2}$. Solving these equations, we have

$$
a(t) \leq b(t)=\|\alpha\|_{2} \cosh (\sqrt{\lambda} t)+\frac{\|\beta\|_{2}}{\sqrt{\lambda}} \sinh (\sqrt{\lambda} t)
$$

for all $t \geq 0$.
Next, we give an approximate solution of (14.17).
Lemma 14.4.20. Let $\psi$ be the solution of (14.17). Suppose that $\|R(t)\|_{2} \leq \lambda$. For any $0 \leq t \leq \frac{1}{\sqrt{\lambda}}$, we have that

$$
\|\psi(t)-\alpha-\beta t\|_{2} \leq \lambda t^{2}\|\alpha\|_{2}+\frac{\lambda t^{3}}{5}\|\beta\|_{2} \text { and }\left\|\psi^{\prime}(t)-\beta\right\|_{2} \leq 2 \lambda t\|\alpha\|_{2}+\frac{3 \lambda t^{2}}{5}\|\beta\|_{2}
$$

Proof. Note that

$$
\begin{aligned}
\psi(t) & =\psi(0)+\psi^{\prime}(0) t+\int_{0}^{t}(t-s) \psi^{\prime \prime}(s) d s \\
& =\alpha+\beta t-\int_{0}^{t}(t-s) R(s) \psi(s) d s
\end{aligned}
$$

Using Lemma 14.4.19 and $\|R(t)\|_{2} \leq \lambda$, we have that

$$
\begin{aligned}
\|\psi(t)-\alpha-\beta t\|_{2} & \leq \lambda \int_{0}^{t}(t-s)\|\psi(s)\|_{2} d s \\
& \leq \lambda \int_{0}^{t}(t-s)\left(\|\alpha\|_{2} \cosh (\sqrt{\lambda} s)+\frac{\|\beta\|_{2}}{\sqrt{\lambda}} \sinh (\sqrt{\lambda} s)\right) d s \\
& =\|\alpha\|_{2}(\cosh (\sqrt{\lambda} t)-1)+\frac{\|\beta\|_{2}}{\sqrt{\lambda}}(\sinh (\sqrt{\lambda} t)-\sqrt{\lambda} t)
\end{aligned}
$$

Since $0 \leq t \leq \frac{1}{\sqrt{\lambda}}$, we have that $|\cosh (\sqrt{\lambda} t)-1| \leq \lambda t^{2}$ and $|\sinh (\sqrt{\lambda} t)-\sqrt{\lambda} t| \leq \frac{\lambda^{3 / 2} t^{3}}{5}$. This gives the result.

Similarly, we have that $\psi^{\prime}(t)=\psi^{\prime}(0)+\int_{0}^{t} \psi^{\prime \prime}(s) d s=\psi^{\prime}(0)+\int_{0}^{t} R(s) \psi(s) d s$ and hence

$$
\begin{aligned}
\left\|\psi^{\prime}(t)-\beta\right\|_{2} & \leq \lambda \int_{0}^{t}\left(\|\alpha\|_{2} \cosh (\sqrt{\lambda} s)+\frac{\|\beta\|_{2}}{\sqrt{\lambda}} \sinh (\sqrt{\lambda} s)\right) d s \\
& \leq \sqrt{\lambda}\|\alpha\|_{2} \sinh (\sqrt{\lambda} t)+\|\beta\|_{2}(\cos (\sqrt{\lambda} t)-1) \\
& \leq 2 \lambda t\|\alpha\|_{2}+\frac{3}{5} \lambda t^{2}\|\beta\|_{2}
\end{aligned}
$$

The following is a matrix version of the above result.
Lemma 14.4.21. Let $\psi$ be the solution of

$$
\begin{aligned}
\frac{d^{2}}{d t^{2}} \Psi(t)+R(t) \Psi(t) & =0 \\
\frac{d}{d t} \Psi(0) & =B \\
\Psi(0) & =A
\end{aligned}
$$

where $\Psi, A$ and $B$ are matrices with a compatible size. Suppose that $\|R(t)\|_{2} \leq \lambda$. For any $0 \leq t \leq \frac{1}{\sqrt{\lambda}}$, we have that

$$
\|\Psi(t)-A-B t\|_{F} \leq \max _{0 \leq s \leq t}\|R(s)\|_{F}\left(t^{2}\|A\|_{2}+\frac{t^{3}}{5}\|B\|_{2}\right)
$$

Proof. Note that $\Psi(t) x$ is the solution of (14.17) with $\beta=B x$ and $\alpha=A x$. Therefore, Lemma 14.4.19 shows that

$$
\|\Psi(t) x\|_{2} \leq\left(\|A\|_{2} \cosh (\sqrt{\lambda} t)+\frac{\|B\|_{2}}{\sqrt{\lambda}} \sinh (\sqrt{\lambda} t)\right)\|x\|_{2}
$$

for all $x$. Therefore, we have that

$$
\Psi(t)^{T} \Psi(t) \preceq\left(\|A\|_{2} \cosh (\sqrt{\lambda} t)+\frac{\|B\|_{2}}{\sqrt{\lambda}} \sinh (\sqrt{\lambda} t)\right)^{2} I .
$$

Hence, we have that

$$
\Psi(t) \Psi(t)^{T} \preceq\left(\|A\|_{2} \cosh (\sqrt{\lambda} t)+\frac{\|B\|_{2}}{\sqrt{\lambda}} \sinh (\sqrt{\lambda} t)\right)^{2} I
$$

Using this, we have that

$$
\begin{aligned}
\|\Psi(t)-A-B t\|_{F} & =\left\|\int_{0}^{t}(t-s) R(s) \Psi(s) d s\right\|_{F} \\
& \leq \int_{0}^{t}(t-s) \sqrt{\operatorname{Tr} \Psi^{T}(s) R^{T}(s) R(s) \Psi(s)} d s \\
& =\int_{0}^{t}(t-s) \sqrt{\operatorname{Tr} R(s) \Psi(s) \Psi^{T}(s) R^{T}(s)} d s \\
& \leq \int_{0}^{t}(t-s)\left(\|A\|_{2} \cosh (\sqrt{\lambda} s)+\frac{\|B\|_{2}}{\sqrt{\lambda}} \sinh (\sqrt{\lambda} s)\right)\|R(s)\|_{F} d s \\
& \leq \max _{0 \leq s \leq t}\|R(s)\|_{F} \int_{0}^{t}(t-s)\left(\|A\|_{2} \cosh (\sqrt{\lambda} s)+\frac{\|B\|_{2}}{\sqrt{\lambda}} \sinh (\sqrt{\lambda} s)\right) d s \\
& =\max _{0 \leq s \leq t}\|R(s)\|_{F}\left(\|A\|_{2}(\cosh (\sqrt{\lambda} t)-1)+\frac{\|B\|_{2}}{\sqrt{\lambda}}(\sinh (\sqrt{\lambda} t)-\sqrt{\lambda} t)\right)
\end{aligned}
$$

Since $0 \leq t \leq \frac{1}{\sqrt{\lambda}}$, we have that $|\cosh (\sqrt{\lambda} t)-1| \leq \lambda t^{2}$ and $|\sinh (\sqrt{\lambda} t)-\sqrt{\lambda} t| \leq \frac{\lambda^{3 / 2} t^{3}}{5}$. This gives the result.

## ■ 14.4.9 Almost One-to-One Correspondence of Geodesics

We do not know if every pair $x, y \in M$ has a unique geodesic connecting $x$ and $y$. Due to this, the probability density $p_{x}$ at $y$ on $\mathcal{M}$ can be the sum over all possible geodesics connect $x$ and $y$. The goal
of this section is to show there is a 1-1 map between geodesics paths connecting $x$ to $y$ as we move $x$.

Lemma 14.4.22. Given a geodesic walk $\gamma(t)=\exp _{x}\left(\frac{t}{\ell} v_{x}\right) \in \Pi$ with step size $h$ satisfying $\frac{1}{n} \leq h \leq$ $\frac{1}{n R_{1}}$. Let the end points $x=\gamma(0)$ and $y=\gamma(\ell)$. There is an unique smooth invertible function $v: U \subset M \rightarrow V \subset T_{x} M$ such that

$$
y=\exp _{z}(v(z))
$$

for any $z \in U$ where $U$ is a neighborhood of $x$ and $V$ is a neighborhood of $v_{x}=v(x)$. Furthermore, for any $\eta=\eta_{1}+\alpha v_{x}$ with $\eta_{1} \perp v_{x}$ and scale $\alpha$, we have that

$$
\nabla_{\eta} v(x)=-\eta+\zeta
$$

where $\|\zeta\| \leq \frac{3}{2} n h R_{1}\left\|\eta_{1}\right\| \leq \frac{3}{2}\left\|\eta_{1}\right\|$ and $\zeta \perp v(x)$. In particular, we have that $\left\|\nabla_{\eta} v(x)\right\| \leq \frac{5}{2}\|\eta\|$.

Proof. Consider the function $f(z, w)=\exp _{z}(w)$. It is continuously differentiable. From Lemma 14.4.14, the differential of $w$ at $\left(x, v_{x}\right)$ on the $w$ variables, i.e. $d \exp _{x}\left(v_{x}\right)$, is invertible. Hence, the implicit function theorem shows that there is a open neighborhood $U$ of $x$ and a unique function $v$ on $U$ such that $f(z, v(z))=f\left(z, v_{z}\right)$, i.e. $y=\exp _{x}\left(v_{x}\right)=\exp _{z}(v(z))$.

To compute $\nabla_{\eta} v$, let $c(t, s)=\exp _{\gamma(s)}\left(\frac{t}{\ell} v(\gamma(s))\right)$ be a family of geodesics with the end points $c(\ell, s)=\exp _{\gamma(s)}(v(\gamma(s)))=y$. Note that $\left.\psi(t) \stackrel{\text { def }}{=} \frac{\partial c(t, s)}{\partial s}\right|_{s=0}$ satisfies the Jacobi field equation

$$
D_{t}^{2} \psi(t)+R\left(\psi, \frac{\partial c}{\partial t}\right) \frac{\partial c}{\partial t}=0
$$

Moreover, we know that

$$
\begin{aligned}
\psi(0) & =\left.\frac{\partial c(0, s)}{\partial s}\right|_{s=0}=\gamma^{\prime}(0)=\eta \\
D_{t} \psi(0) & =\left.D_{t} \frac{\partial c(t, s)}{\partial s}\right|_{t, s=0}=\left.D_{s} \frac{\partial c(t, s)}{\partial t}\right|_{t, s=0}=\frac{1}{\ell} \nabla_{\eta} v(x) \stackrel{\text { def }}{=} \frac{\chi}{\ell} \\
\psi(\ell) & =\left.\frac{\partial c(\ell, s)}{\partial s}\right|_{s=0}=0
\end{aligned}
$$

where we used $D_{t} \frac{\partial}{\partial s}=D_{s} \frac{\partial}{\partial t}$ (torsion free-ness, Fact 14.2.1).
From Fact 14.2.4, we know that $\psi$ can be split into the tangential part $\psi_{1}$ and the normal part $\psi_{2}$. For the tangential part, we know that

$$
\begin{aligned}
\psi_{1}(t) & =\left(\left\langle\psi(0), \frac{d c}{d t}(0,0)\right\rangle_{c(0)}+\left\langle D_{t} \psi(0), \frac{d c}{d t}(0,0)\right\rangle_{c(0)} t\right) \frac{d c}{d t}(t) \\
& =\left(\left\langle\eta, v_{x}\right\rangle_{x}+\left\langle\frac{\chi}{\ell}, v_{x}\right\rangle_{x} t\right) \frac{d c}{d t}(t) \\
& =\left(\alpha\left\|v_{x}\right\|^{2}+\left\langle\frac{\chi}{\ell}, v_{x}\right\rangle_{x} t\right) \frac{d c}{d t}(t)
\end{aligned}
$$

Since $\psi_{1}(\ell)=0$, we have that

$$
\begin{equation*}
\left\langle\chi, v_{x}\right\rangle=-\alpha\left\|v_{x}\right\|^{2} \tag{14.18}
\end{equation*}
$$

For the normal part $\psi_{2}$, it is easier to calculate using orthogonal frames. So, we pick an arbitrary orthogonal frame parallel transported along the curve $c(t, 0)$ and let $\bar{\psi}(t)$ be $\psi_{2}(t)$ represented in that
orthogonal frame. Hence, we have that

$$
\begin{aligned}
\frac{d^{2}}{d t^{2}} \bar{\psi}(t)+R\left(\bar{\psi}(t), \frac{\partial c}{\partial t}\right) \frac{\partial c}{\partial t} & =0 \\
\bar{\psi}(0) & =\eta_{1} \\
\bar{\psi}^{\prime}(0) & =\frac{\chi_{1}}{\ell} \\
\bar{\psi}(\ell) & =0
\end{aligned}
$$

where $\chi_{1}=\chi-\left\langle\chi, \frac{v_{x}}{\left\|v_{x}\right\|}\right\rangle \frac{v_{x}}{\left\|v_{x}\right\|}=\chi+\alpha v_{x}$ (14.18). Using $h \leq \frac{1}{n R_{1}}$, Lemma 14.4.20 shows that

$$
\begin{aligned}
\left\|\eta_{1}+\chi_{1}\right\|_{2}=\left\|\bar{\psi}(\ell)-\eta_{1}-\chi_{1}\right\|_{2} & \leq R_{1} \ell^{2}\left\|\eta_{1}\right\|_{2}+\frac{R_{1} \ell^{2}}{5}\left\|\chi_{1}\right\|_{2} \\
& \leq \frac{6}{5} R_{1} \ell^{2}\left\|\eta_{1}\right\|_{2}+\frac{1}{5}\left\|\chi_{1}\right\|_{2}
\end{aligned}
$$

Hence, we have that $\left\|\eta_{1}+\chi_{1}\right\|_{2} \leq \frac{3}{2} R_{1} \ell^{2}\left\|\eta_{1}\right\|_{2}$. Therefore, we have that

$$
\begin{aligned}
\nabla_{\eta} v(x) & =\chi \\
& =-\alpha v_{x}+\chi_{1} \\
& =-\eta+\left(\eta_{1}+\chi_{1}\right)
\end{aligned}
$$

where $\left\|\chi_{1}+\eta_{1}\right\|_{2} \leq \frac{3}{2}\left\|\eta_{1}\right\|_{2}$. Furthermore, we have that both $\eta_{1}$ and $\chi_{1}$ are orthogonal to $\eta$.
Remark. If the above lemma holds without the assumption $\gamma \in \Pi$, this would imply uniqueness of geodesics.

The following lemma shows there is a 1-1 map between geodesics paths connecting $x$ to $y$ as we move $x$. When we move $x$, the geodesic $\gamma$ from $x$ to $y$ no longer insides $\Xi$. To avoid this issue, we need to assume $\left\langle R\left(X_{i}, \gamma^{\prime}\right) \gamma^{\prime}, X_{j}\right\rangle$ is mildly smooth in the following sense.

Assumption 14.4.23. Given a manifold $M$. There is some constant $C$ such that for any a curve $\alpha(s) \in M$ and a vector field $u(s)$ on $\alpha(s)$, let $A(s)_{i j}=\left\langle R\left(X_{i}, u(s)\right) u(s), X_{j}\right\rangle$ where $\left\{X_{i}\right\}$ is any orthonormal frame at $\alpha(s)$, we have

$$
\left|\frac{d}{d s}\|A(s)\|_{F}\right| \leq\left(\left\|\alpha^{\prime}(0)\right\|+\left\|D_{s} u(0)\right\|\right) n^{C} .
$$

Note that this assumption implies $R_{1}=n^{\Theta(1)}$ by picking the curve is a singleton and the vector field $u(s)=s \cdot u$.

Lemma 14.4.24. Given a geodesic $\gamma(t)=\exp _{x}\left(\frac{t}{\ell} v_{x}\right)$ on $M$ with step size $h$ satisfying $\frac{1}{n} \leq h \leq \frac{1}{2 n R_{1}}$, let $c(s)$ be any geodesic starting at $\gamma(0)$. Let $x=c(0)=\gamma(0)$ and $z=c(1)$. Suppose that the length of the geodesic is less than $n^{-\Theta(1)}$. Then, there is a unique vector field $v$ on $c$ such that

$$
y=\exp _{c(s)}(v(s))
$$

This vector field is uniquely determined by the geodesic and any $v(s)$ on this vector field.
Proof. Let $s_{\text {max }}$ be the supremum of $s$ such that $v(s)$ can be defined continuously and $y=\exp _{c(s)}(v(s))$. Note that $s_{\max }>0$ as Lemma 14.4.22 shows that there is a neighborhood $N$ at $x$ and a vector field $u$ on $N$ such that for any $z \in N$, we have that

$$
y=\exp _{z}(u(z)) .
$$

Note that the the vector field $u$ is unique and hence the extension we found for $v$ is unique. Also, we know that $\left\|D_{t} v(t)\right\|=\left\|\nabla_{c^{\prime}} u(x)\right\| \leq \frac{5}{2} L$ where $L$ is the length of $c$. Therefore, by continuity, $v\left(s_{\max }\right)$ is well-defined.

If $s_{\max }>1$, then we are done. Otherwise, we want to apply Lemma 14.4.22 at $s_{\max }$. However, the curve $\exp _{c\left(s_{\max }\right)}\left(\frac{t}{\ell} v\left(s_{\text {max }}\right)\right)$ may not lie in $\Pi$. Note that the only place Lemma 14.4.22 uses the assumption $\gamma=\exp _{x}\left(\frac{t}{\ell} v_{x}\right) \in \Pi$ is to apply the bound $R_{1}=\sup _{\gamma \in \Pi, 0 \leq t \leq \ell}\|R(t)\|_{F}$.

To bound $\|R(t)\|_{F}$, we note that $\left\|D_{t} v(t)\right\|=\left\|\nabla_{c^{\prime}} u(x)\right\| \leq \Theta(L)$ where $L$ is the length of $c$. Hence, both the starting point and the direction of the geodesic move by $\Theta(L)$. Now, we can apply Lemma 14.4.19 to prove that along the geodesic, all the points move by $\Theta(L)$ for $0 \leq t \leq \ell$. Then, we can apply Assumption 14.4.23 to conclude that $\|R(t)\|_{F}$ only changes by at most $\Theta\left(n^{C} L\right)$. Therefore, if $L \leq \frac{R_{1}}{n^{C}}=n^{-C}$, then $\|R(t)\|_{F}$ is still bounded by $\Theta\left(R_{1}\right)$ and Lemma 14.4.22 can be applied at $s_{\max }<1$. This is a contradiction. Hence, $s_{\max } \geq 1$. Its uniqueness follows from Lemma 14.4.22.

## ■ 14.5 Logarithmic Barrier

For any polytope $M=\{A x>b\}$, the logarithmic barrier function $\phi(x)$ is defined as

$$
\phi(x)=-\sum_{i=1}^{m} \log \left(a_{i}^{T} x-b_{i}\right) .
$$

We denote the Hessian manifold induced by the logarithmic barrier on $M$ by $M_{L}$. The goal of this section is to analyze the geodisic walk on $M_{L}$.

In section 14.5.1, we give explicit formulas of various Riemannian geometry concepts on $M_{L}$. In Section 14.5.2, we explain the geodesic walk on $M_{L}$. In Sections 14.5.3 to 14.5.7, we bound the necessary constant required by Theorem 14.4.2. This shows that

Theorem 14.5.1. The geodesic walk on $M_{L}$ with step size $h \leq \frac{c}{n^{3 / 4}}$ has mixing time is $O(m / h)$ for some universal constant $c$.

In the sections afterward, we show how to implement geodesic walk and calculate the rejection probability. To implement these, we apply the techniques developed in Section 14.6 to solve the corresponding ODEs. In this section, we first show that geodesic, parallel transport and Jacobi Field are complex analytic (Section 14.5.9), then we bound their radius of convergence (Section 14.5.9.1). This will help us bound the approximation error of the collocation method developed in Section 14.6.

Theorem 14.5.2. We can implement one step of geodesic walk with step size $h \leq \frac{c}{\sqrt{n}}$ in time $O\left(m n^{\omega-1} \log ^{2}(n)\right)$ for a universal constant $c$.

## $\square$ 14.5.1 Riemannian geometry on $M_{L}$

We use the following definitions throughout this section:

- $s_{x}=A x-b, S_{x}=\operatorname{Diag}\left(s_{x}\right)$.
- $A_{x}=S_{x}^{-1} A$.
- $s_{x, v}=A_{x} v, S_{x, v}=\operatorname{Diag}\left(A_{x} v\right)$.
- $P_{x}=A_{x}\left(A_{x}^{T} A_{x}\right)^{-1} A_{x}^{T}, \sigma_{x}=\operatorname{Diag}\left(P_{x}\right)$.
- $\left(P_{x}^{(2)}\right)_{i j}=\left(P_{x}\right)_{i j}^{2}$.
- For brevity (overloading notation), we define $s_{\gamma^{\prime}}=s_{\gamma, \gamma^{\prime}}, s_{\gamma^{\prime \prime}}=s_{\gamma, \gamma^{\prime \prime}}, S_{\gamma^{\prime}}=S_{\gamma, \gamma^{\prime}}$ and $S_{\gamma^{\prime \prime}}=S_{\gamma, \gamma^{\prime \prime}}$ for a curve $\gamma(t)$.

Since the Hessian manifold $M_{L}$ is naturally embedded in $\mathbb{R}^{n}$, we identify $T_{x} M_{L}$ by Euclidean coordinates unless otherwise stated. Therefore, we have that

$$
\begin{aligned}
\langle u, v\rangle_{x} & =u^{T} \nabla^{2} \phi(x) v \\
& =u^{T} A_{x}^{T} A_{x} v .
\end{aligned}
$$

Lemma 14.5.3. Let $u(t)$ be a vector field defined on a curve $\gamma(t)$ in $M_{L}$. Then, we have that

$$
\nabla_{\gamma^{\prime}} u=\frac{d u}{d t}-\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime}} s_{\gamma, u}
$$

In particular, the geodesic equation on $M_{L}$ is given by

$$
\begin{equation*}
\gamma^{\prime \prime}=\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} s_{\gamma^{\prime}}^{2} \tag{14.19}
\end{equation*}
$$

and the equation for parallel transport on a curve $\gamma(t)$ is given by

$$
\begin{equation*}
\frac{d}{d t} v(t)=\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime}} A_{\gamma} v \tag{14.20}
\end{equation*}
$$

Proof. By Lemma 14.2.5, the Christoffel symbols respect to the Euclidean coordinates is given by

$$
\begin{aligned}
\Gamma_{i j}^{k} & =\frac{1}{2} \sum_{l} g^{k l} \phi_{i j l} \\
& =-\sum_{z} \sum_{l} e_{k}^{T}\left(A_{x}^{T} A_{x}\right)^{-1} e_{l}\left(A_{x} e_{i}\right)_{z}\left(A_{x} e_{j}\right)_{z}\left(A_{x} e_{l}\right)_{z} \\
& =-e_{k}^{T}\left(A_{x}^{T} A_{x}\right)^{-1} A_{x}^{T}\left(\left(A_{x} e_{i}\right)\left(A_{x} e_{j}\right)\right) .
\end{aligned}
$$

Recall that the Levi-Civita connection is given by

$$
\nabla_{v} u=\sum_{i k} v_{i} \frac{\partial u_{k}}{\partial x^{i}} e_{k}+\sum_{i j k} v_{i} u_{j} \Gamma_{i j}^{k} e_{k}
$$

Therefore, we have

$$
\nabla_{v} u=\sum_{i k} v^{i} \frac{\partial u^{k}}{\partial x^{i}} e_{k}-\left(A_{x}^{T} A_{x}\right)^{-1} A_{x}^{T} S_{x, v} s_{x, u} .
$$

Since $U$ is a vector field defined on a curve $\gamma(t)$, we have that $\sum_{i k} \gamma_{i}^{\prime} \frac{\partial u_{k}}{\partial x^{i}} e_{k}=\frac{d u}{d t}$.
The geodesic equation follows from $\nabla_{\gamma^{\prime}} \gamma^{\prime}=0$ and the parallel transport equation follows from $\nabla_{\gamma^{\prime}} v=0$.

Lemma 14.5.4. Given $u, v, w, z \in T_{x} M_{L}$, the Riemann Curvature Tensor is given by

$$
\langle R(u, v) w, z\rangle=\left(s_{x, u} s_{x, w}\right)^{T} A_{x}\left(A_{x}^{T} A_{x}\right)^{-1} A_{x}^{T}\left(s_{x, v} s_{x, z}\right)-\left(s_{x, u} s_{x, z}\right)^{T} A_{x}\left(A_{x}^{T} A_{x}\right)^{-1} A_{x}^{T}\left(s_{x, v} s_{x, w}\right)
$$

and the Ricci curvature is given by

$$
\operatorname{Ric}(u)=s_{x, u}^{T} P_{x}^{(2)} s_{x, u}-\sigma_{x}^{T} P_{x} s_{x, u}^{2} .
$$

Proof. By Lemma 14.2.5, we have that

$$
\begin{aligned}
\langle R(u, v) w, z\rangle & =\frac{1}{4} \sum_{p q i j l k} g^{p q}\left(\phi_{j k p} \phi_{i l q}-\phi_{i k p} \phi_{j l q}\right) u^{i} v^{j} w^{l} z^{k} \\
& =\sum_{p q} g^{p q}\left(e_{p}^{T} A_{x}^{T}\left(s_{x, v} s_{x, z}\right) e_{q}^{T} A_{x}^{T}\left(s_{x, u} s_{x, w}\right)-e_{p}^{T} A_{x}^{T}\left(s_{x, u} s_{x, z}\right) e_{q}^{T} A_{x}^{T}\left(s_{x, v} s_{x, w}\right)\right) \\
& =\left(s_{x, u} s_{x, w}\right)^{T} A_{x}\left(A_{x}^{T} A_{x}\right)^{-1} A_{x}^{T}\left(s_{x, v} s_{x, z}\right)-\left(s_{x, u} s_{x, z}\right)^{T} A_{x}\left(A_{x}^{T} A_{x}\right)^{-1} A_{x}^{T}\left(s_{x, v} s_{x, w}\right)
\end{aligned}
$$

and

$$
\begin{aligned}
\operatorname{Ric}(u) & =\sum_{j l} g^{j l}\left\langle R\left(u, \frac{\partial}{\partial x^{j}}\right), \frac{\partial}{\partial x^{l}}, u\right\rangle \\
& =\sum_{j l} g^{j l}\left(s_{x, u} s_{x, e_{l}}\right)^{T} A_{x}\left(A_{x}^{T} A_{x}\right)^{-1} A_{x}^{T}\left(s_{x, e_{j}} s_{x, u}\right)-\left(s_{x, u}^{2}\right)^{T} A_{x}\left(A_{x}^{T} A_{x}\right)^{-1} A_{x}^{T}\left(s_{x, e_{j}} s_{x, e_{l}}\right) \\
& =\operatorname{Tr} S_{x, u} A_{x}\left(A_{x}^{T} A_{x}\right)^{-1} A_{x}^{T} S_{x, u} A_{x}\left(A_{x}^{T} A_{x}\right)^{-1} A_{x}-\operatorname{Tr}\left(A_{x}^{T} \operatorname{Diag}\left(P_{x} s_{x, u}^{2}\right) A_{x}\left(A_{x}^{T} A_{x}\right)^{-1}\right) \\
& =s_{x, u}^{T} P_{x}^{(2)} s_{x, u}-\sigma_{x}^{T} P_{x} s_{x, u}^{2} .
\end{aligned}
$$

Lemma 14.5.5. Given a geodesic $\gamma(t)$ on $M_{L}$ and an orthogonal frame $\left\{x_{i}\right\}_{i=1}^{n}$ on $\gamma(t)$. The Jacobi field equation (in the orthogonal frame coordinate) is given by

$$
\frac{d^{2} u}{d t^{2}}+X^{-1}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1}\left(A_{\gamma}^{T} S_{\gamma^{\prime}} P_{\gamma} S_{\gamma^{\prime}} A_{\gamma}-A_{\gamma}^{T} \operatorname{Diag}\left(P_{\gamma} s_{\gamma^{\prime}}^{2}\right) A_{\gamma}\right) X u=0
$$

where $X(t)=\left[x_{1}(t), x_{2}(t), \cdots, x_{n}(t)\right]$.
Proof. The equation for Jacobi field along is

$$
D_{t}^{2} v+R\left(v, \gamma^{\prime}\right) \gamma^{\prime}=0
$$

By Lemma 14.5.4, under Euclidean coordinates, we have

$$
D_{t}^{2} v+\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1}\left(A_{\gamma}^{T} S_{\gamma^{\prime}} P_{\gamma} S_{\gamma^{\prime}} A_{\gamma}-A_{\gamma}^{T} \operatorname{diag}\left(P_{\gamma} s_{\gamma^{\prime}}^{2}\right) A_{\gamma}\right) v=0
$$

We write $v$ in terms of the orthogonal frame, namely, $v(t)=X(t) u(t)$ where $u(t) \in \mathbb{R}^{n}$. Then, we have that $D_{t}^{2} v=X(t) \frac{d^{2} u(t)}{d t^{2}}$. Hence, under the orthogonal frame coordinate, we have that

$$
\frac{d^{2} u}{d t^{2}}+X^{-1}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1}\left(A_{\gamma}^{T} S_{\gamma^{\prime}} P_{\gamma} S_{\gamma^{\prime}} A_{\gamma}-A_{\gamma}^{T} \operatorname{Diag}\left(P_{\gamma} s_{\gamma^{\prime}}^{2}\right) A_{\gamma}\right) X u=0 .
$$

Now, we show the relation between the distance induced by the metric and the Hilbert metric.
Lemma 14.5.6. For any $x, y \in M_{L}$, we have that

$$
\frac{d(x, y)}{d_{H}(x, y)} \leq \sqrt{m}
$$

Hence, we have that $G_{2} \leq \sqrt{m}$.
Proof. First, we note that it suffices to prove that $\frac{d(x, y)}{d_{H}(x, y)} \leq(1+O(\varepsilon)) \sqrt{m}$ for any $x, y \in M_{L}$ with $d(x, y) \leq \varepsilon$. Then, one can run a limiting argument as follows. Let $x_{t}=t \cdot x+(1-t) \cdot y$, then we have
that

$$
d_{H}(x, y)=\lim _{n \rightarrow \infty} \sum_{i=0}^{n-1} d_{H}\left(x_{k / n}, x_{(k+1) / n}\right) .
$$

Since $d_{H}\left(x_{k / n}, x_{(k+1) / n}\right)=O_{x, y}\left(\frac{1}{n}\right)$, we have that

$$
\begin{aligned}
d_{H}(x, y) & =\lim _{n \rightarrow \infty} \sum_{i=0}^{n-1} d_{H}\left(x_{k / n}, x_{(k+1) / n}\right) \\
& \geq \lim _{n \rightarrow \infty} \frac{\left(1-O_{x, y}\left(\frac{1}{n}\right)\right)}{\sqrt{m}} \sum_{i=0}^{n-1} d\left(x_{k / n}, x_{(k+1) / n}\right) \\
& =\frac{1}{\sqrt{m}} \lim _{n \rightarrow \infty} \sum_{i=0}^{n-1} d\left(x_{k / n}, x_{(k+1) / n}\right) \\
& \geq \frac{1}{\sqrt{m}} \lim _{n \rightarrow \infty} d(x, y)=\frac{d(x, y)}{\sqrt{m}} .
\end{aligned}
$$

Now, we can assume $d(x, y) \leq \varepsilon$.Let $p$ and $q$ are on the boundary of $M_{L}$ such that $p, x, y, q$ are on the straight line $\overline{x y}$ and are in order. Without loss of generality, we assume $p$ is closer to $x$. Then, we have that $p \in M_{L} \cap\left(x-M_{L}\right)$, equivalently, we have that $\left|a_{i}^{T} p-a_{i}^{T} x\right| \leq a_{i}^{T} x-b_{i}$ for all $i$ and hence $\left\|A_{x}(p-x)\right\|_{\infty} \leq 1$. Therefore, we have that

$$
\left\|A_{x}(p-x)\right\|_{2} \leq \sqrt{m}\left\|A_{x}(p-x)\right\|_{\infty} \leq \sqrt{m} .
$$

Since $p, x, y$ are on the same line, we have that

$$
\frac{\|x-y\|_{2}\|p-q\|_{2}}{\|p-x\|_{2}\|y-q\|_{2}} \geq \frac{\|x-y\|_{2}}{\|p-x\|_{2}}=\frac{\left\|A_{x}(x-y)\right\|_{2}}{\left\|A_{x}(p-x)\right\|_{2}} \geq \frac{\left\|A_{x}(x-y)\right\|_{2}}{\sqrt{m}} .
$$

Since $d(x, y)<\varepsilon$, Lemma 3.1 in [214] shows that

$$
d(x, y) \leq-\log \left(1-\left\|A_{x}(x-y)\right\|_{2}\right)
$$

Hence, we have that $\left\|A_{x}(x-y)\right\|_{2} \geq(1-O(\varepsilon)) d(x, y)$ and hence

$$
d_{H}(x, y)=\frac{\|x-y\|_{2}\|p-q\|_{2}}{\|p-x\|_{2}\|y-q\|_{2}} \geq(1-O(\varepsilon)) \frac{d(x, y)}{\sqrt{m}} .
$$

## ■ 14.5.2 Geodesic Walk on $M_{L}$

Recall that the geodesic walk is given by

$$
x^{(j+1)}=\exp _{x^{(j)}}\left(\sqrt{h} w+\frac{h}{2} \mu\left(x^{(j)}\right)\right)
$$

where $w \sim N(0, I)$ (in the local coordinates at $x^{(j)}$ ). In many proofs in this section, we consider the geodesic $\gamma$ from $x^{(j)}$ to $x^{(j+1)}$ with the initial velocity

$$
\gamma^{\prime}(0)=\frac{w}{\sqrt{n}}+\frac{1}{2} \sqrt{\frac{h}{n}} \mu(x) .
$$

The scaling is to make the speed of geodesic $\gamma$ close to one. Since $w$ is a Gaussian vector, we have that $0.9 \leq\left\|\gamma^{\prime}(0)\right\|_{\gamma(0)} \leq 1.1$ with high probability. Due to this rescaling, the geodesic is defined from 0 to $\ell=\sqrt{n h}$.

We often work on the Euclidean coordinates. In this case, the geodesic walk is given by the following formula:

Lemma 14.5.7. Given $x \in M_{L}$ and step size $h>0$, one step of the geodesic walk starting at $x$ is given by the solution $\gamma(\ell)$ of the following geodesic equation

$$
\begin{aligned}
\gamma^{\prime \prime}(t) & =\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} s_{\gamma^{\prime}}^{2} \text { for } 0 \leq t \leq \ell \\
\gamma^{\prime}(0) & =\frac{w}{\sqrt{n}}+\frac{1}{2} \sqrt{\frac{h}{n}} \mu(x) \\
\gamma(0) & =x
\end{aligned}
$$

where $\ell=\sqrt{n h}, w \sim N\left(0,\left(A_{x}^{T} A_{x}\right)^{-1}\right)$ and $\mu(x)=\left(A_{x}^{T} A_{x}\right)^{-1} A_{x}^{T} \sigma_{x}$.
Proof. The geodesic equation is given by Lemma 14.5.3. From (14.2), the drift term is given by

$$
\begin{aligned}
\mu_{i}(x) & =\frac{1}{2} \sum_{j=1}^{n} \frac{\partial}{\partial x_{j}}\left(\left(\nabla^{2} \phi\left(X_{t}\right)\right)^{-1}\right)_{i j}=\frac{1}{2} \sum_{j=1}^{n} \frac{\partial}{\partial x_{j}}\left(\left(A_{x}^{T} A_{x}\right)^{-1}\right)_{i j} \\
& =\sum_{j} e_{i}^{T}\left(A_{x}^{T} A_{x}\right)^{-1} A_{x}^{T} S_{x, e_{j}} A_{x}\left(A_{x}^{T} A_{x}\right)^{-1} e_{j} \\
& =\sum_{j} \sum_{k} V_{i k} e_{k}^{T} A_{x} e_{j} V_{j k} \quad \text { where } V=\left(A_{x}^{T} A_{x}\right)^{-1} A_{x}^{T} \\
& =\sum_{k} V_{i k} e_{k}^{T} A_{x} \sum_{j} e_{j}\left(e_{k}^{T} A_{x}\left(A_{x}^{T} A_{x}\right)^{-1} e_{j}\right) \\
& =\sum_{k} V_{i k} e_{k}^{T} A_{x} \sum_{j} e_{j} e_{j}^{T}\left(A_{x}^{T} A_{x}\right)^{-1} A_{x}^{T} e_{k} \\
& =\sum_{k} V_{i k} e_{k}^{T} A_{x}\left(A_{x}^{T} A_{x}\right)^{-1} A_{x}^{T} e_{k} \\
& =V_{i} \sigma_{x}=e_{i}^{T}\left(A_{x}^{T} A_{x}\right)^{-1} A_{x}^{T} \sigma_{x} .
\end{aligned}
$$

## ■ 14.5.3 Randomness and Stability of the Geodesic

In this and subsequent sections, we will frequently use the following elementary calculus facts (using only the chain/product rules and formula for derivative of inverse of a matrix):

$$
\begin{aligned}
\frac{d A_{x}}{d t} & =-S_{\gamma^{\prime}} A_{x} \\
\frac{d P_{x}}{d t} & =-S_{\gamma^{\prime}} P_{x}-P_{x} S_{\gamma^{\prime}}+2 P_{x} S_{\gamma^{\prime}} P_{x} \\
\frac{d S_{\gamma^{\prime}}}{d t} & =\operatorname{Diag}\left(-S_{\gamma^{\prime}} A_{x} \gamma^{\prime}+A_{x} \gamma^{\prime \prime}\right)=-S_{\gamma^{\prime}}^{2}+\operatorname{Diag}\left(P_{x}\right) S_{\gamma^{\prime}}^{2} \\
\frac{d \sigma_{x}}{d t} & =\operatorname{Diag}\left(\frac{d P_{x}}{d t}\right)
\end{aligned}
$$

We also use these matrix inequalities: $\operatorname{Tr}(A B)=\operatorname{Tr}(B A), \operatorname{Tr}(P A P) \leq \operatorname{Tr}(A)$ for any psd matrix $A$; $\operatorname{Tr}\left(A B A^{T}\right) \leq \operatorname{Tr}\left(A Z A^{T}\right)$ for any $B \preceq Z$; the Cauchy-Schwartz, namely, $\operatorname{Tr}(A B) \leq \operatorname{Tr}\left(A A^{T}\right)^{\frac{1}{2}} \operatorname{Tr}\left(B^{T} B\right)^{\frac{1}{2}}$. We also use $P^{2}=P$ since $P$ is a projection matrix.

We note that $\left\|\gamma^{\prime}(0)\right\|_{4}$ is small because it is the sum of a random vector plus a small drift term. Here, we would like to prove that the geodesic is stable in $\ell^{4}$ norm. This fact is very important in the proof and hence we first establish this here.

Lemma 14.5.8. Let $\gamma$ be a geodesic in $M_{L}$ starting at $x$. Let $v_{4}=\left\|s_{\gamma^{\prime}(0)}\right\|_{4}$. Then, for $0 \leq t \leq \frac{1}{10 v_{4}}$, we have that

1. $\left\|s_{\gamma^{\prime}(t)}\right\|_{4} \leq 1.25 v_{4}$.
2. $\left\|\gamma^{\prime \prime}\right\|_{\gamma}^{2} \leq 3 v_{4}^{4}$.

Proof. Let $u(t)=\left\|s_{\gamma^{\prime}(t)}\right\|_{4}$. Then, we have

$$
\begin{align*}
\frac{d u}{d t} & \leq\left\|\frac{d}{d t}\left(A_{\gamma} \gamma^{\prime}\right)\right\|_{4} \\
& =\left\|A_{\gamma} \gamma^{\prime \prime}-\left(A_{\gamma} \gamma^{\prime}\right)^{2}\right\|_{4} \\
& \leq\left\|A_{\gamma} \gamma^{\prime \prime}\right\|_{4}+u^{2}(t) \tag{14.21}
\end{align*}
$$

Under the Euclidean coordinates, the geodesic equation is given by

$$
\gamma^{\prime \prime}=\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} s_{\gamma^{\prime}}^{2}
$$

Hence, we have that

$$
\begin{align*}
\left\|\gamma^{\prime \prime}\right\|_{\gamma}^{2} & =\left(s_{\gamma^{\prime}}^{2}\right)^{T} A_{\gamma}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1}\left(A_{\gamma}^{T} A_{\gamma}\right)\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} s_{\gamma^{\prime}}^{2} \\
& \leq \sum_{i}\left(s_{\gamma^{\prime}}^{4}\right)_{i}=u^{4}(t) . \tag{14.22}
\end{align*}
$$

Therefore, we have

$$
\left\|A_{\gamma} \gamma^{\prime \prime}\right\|_{4} \leq\left\|A_{\gamma} \gamma^{\prime \prime}\right\|_{2} \leq u^{2}(t) .
$$

Plugging it into (14.21), we have that

$$
\frac{d u}{d t} \leq u^{2}(t)+u^{2}(t) \leq 2 u^{2}(t) .
$$

Since $u(0)=v_{4}$, for all $t \geq 0$, we have

$$
\begin{equation*}
u(t) \leq \frac{v_{4}}{1-2 v_{4} t} . \tag{14.23}
\end{equation*}
$$

For $0 \leq t \leq \frac{1}{10 v_{4}}$, we have $u(t) \leq 1.25 v_{4}$ and this gives the first inequality. Using (14.22), we get the second inequality.

Lemma 14.5.9. For $p \geq 1$, we have

$$
P_{x \sim N(0, I)}\left(\|A x\|_{p}^{p} \leq\left(\left(\frac{2^{p / 2} \Gamma\left(\frac{p+1}{2}\right)}{\sqrt{\pi}} \sum_{i}\left\|a_{i}\right\|_{2}^{p}\right)^{1 / p}+\|A\|_{2 \rightarrow p} t\right)^{p}\right) \leq 1-\exp \left(-\frac{t^{2}}{2}\right) .
$$

In particular, we have

$$
P_{x \sim N(0, I)}\left(\|A x\|_{2}^{2} \leq\left(\|A\|_{F}+\|A\|_{2} t\right)^{2}\right) \leq 1-\exp \left(-\frac{t^{2}}{2}\right)
$$

and

$$
P_{x \sim N(0, I)}\left(\|A x\|_{4}^{4} \leq\left(\left(3 \sum_{i}\left\|a_{i}\right\|_{2}^{4}\right)^{1 / 4}+\|A\|_{2 \rightarrow 4} t\right)^{4}\right) \leq 1-\exp \left(-\frac{t^{2}}{2}\right) .
$$

Proof. Let $F(x)=\|A x\|_{p}$. Since $|F(x)-F(y)| \leq\|A\|_{2 \rightarrow p}\|x-y\|_{2}$, Gaussian concentration shows that

$$
P\left(F(x) \leq \mathbf{E} F(x)+\|A\|_{2 \rightarrow p} t\right) \leq 1-\exp \left(-\frac{t^{2}}{2}\right) .
$$

Since $x^{p}$ is convex, we have that

$$
\begin{aligned}
\mathbf{E}\|A x\|_{p} & \leq\left(\mathbf{E}\|A x\|_{p}^{p}\right)^{1 / p}=\left(\sum_{i} \mathbf{E}\left|a_{i}^{T} x\right|^{p}\right)^{1 / p} \\
& =\left(\mathbf{E}_{t \sim N(0,1)}|t|^{p} \sum_{i}\left\|a_{i}\right\|_{2}^{p}\right)^{1 / p} \\
& =\left(\frac{2^{p / 2} \Gamma\left(\frac{p+1}{2}\right)}{\sqrt{\pi}} \sum_{i}\left\|a_{i}\right\|_{2}^{p}\right)^{1 / p} .
\end{aligned}
$$

Hence, we have that

$$
P\left(\|A x\|_{p}^{p} \leq\left(\left(\frac{2^{p / 2} \Gamma\left(\frac{p+1}{2}\right)}{\sqrt{\pi}} \sum_{i}\left\|a_{i}\right\|_{2}^{p}\right)^{1 / p}+\|A\|_{2 \rightarrow p} t\right)^{p}\right) \leq 1-\exp \left(-\frac{t^{2}}{2}\right) .
$$

Using the lemma above, we prove that within small time, the direction of geodesic is a random vector plus a small vector.

Lemma 14.5.10. Let $\gamma$ given by a geodesic walk on $M_{L}$ with step size $h$ satisfying $h \leq \frac{1}{256 \sqrt{n}}$. Let $u$ be the Gaussian part of the initial velocity $\gamma^{\prime}(0)$. If $\left\|A_{\gamma(0)} u\right\|_{2} \leq \frac{3}{2}$ and $\left\|A_{\gamma} u\right\|_{4} \leq \frac{1.55}{n^{1 / 4}}$, then for any $0 \leq t \leq \ell$,

1. $\left\|A_{\gamma} \gamma^{\prime}\right\|_{2} \leq 2$.
2. $\left\|A_{\gamma} \gamma^{\prime}\right\|_{4} \leq 2 n^{-1 / 4}$.
3. $\left\|\gamma^{\prime \prime}\right\|_{\gamma}^{2} \leq 15 n^{-1}$.
4. $\left\|A_{\gamma} \gamma^{\prime}\right\|_{\infty} \leq 6 \sqrt{h}$.

Furthermore, $\left\|A_{\gamma(0)} u\right\|_{2} \leq \frac{3}{2}$ and $\left\|A_{\gamma} u\right\|_{4} \leq \frac{1.55}{n^{1 / 4}}$ holds with probability at least $1-\exp (-\sqrt{n} / 100)$.

Proof. From the definition of the geodesic walk (Lemma 14.5.7), we have that

$$
\gamma^{\prime}(0)=u+v
$$

where $u \sim N\left(0, \frac{1}{n}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1}\right)$ and $v=\frac{1}{2} \sqrt{\frac{h}{n}}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} \sigma_{\gamma}$. Note that

$$
\begin{aligned}
\left\|A_{\gamma} \gamma^{\prime}(0)\right\|_{2} & \leq\left\|A_{\gamma} u\right\|_{2}+\left\|\frac{1}{2} \sqrt{\frac{h}{n}} A_{\gamma}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} \sigma_{\gamma}\right\|_{2} \\
& =\left\|A_{\gamma} u\right\|_{2}+\frac{1}{2} \sqrt{\frac{h}{n}} \sqrt{\sigma_{\gamma}^{T} A_{\gamma}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} \sigma_{\gamma}} .
\end{aligned}
$$

Note that $A_{\gamma} u=B x$ where $B=\frac{1}{\sqrt{n}} A_{\gamma}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1 / 2}$ and $x \sim N(0, I)$. Note that $\|B x\|_{2}^{2}=\frac{1}{n}\|x\|_{2}^{2}$ for all $x$ and hence $\|B\|_{F}^{2}=1$ and $\|B\|_{2 \rightarrow 2}=\frac{1}{\sqrt{n}}$. Hence, Lemma 14.5.9 shows that

$$
P\left(\left\|A_{\gamma} u\right\|_{2}^{2} \leq\left(1+\frac{t}{\sqrt{n}}\right)^{2}\right) \leq 1-\exp \left(-\frac{t^{2}}{2}\right)
$$

Hence, we have

$$
\left\|A_{\gamma} \gamma^{\prime}\right\|_{2} \leq \frac{3}{2}+\frac{1}{2} \sqrt{h} \leq 2
$$

with probability $1-\exp (-n / 8)$. Since geodesic preserves the speed, this gives the first inequality.
Next, we note that $\sum_{i}\left\|e_{i}^{T} B\right\|_{2}^{4}=\frac{1}{n^{2}} \sum_{i}\left(\sigma_{\gamma}\right)_{i}^{2} \leq \frac{1}{n}$ and $\|B\|_{2 \rightarrow 4} \leq\|B\|_{2 \rightarrow 2}=\frac{1}{\sqrt{n}}$. Hence, Lemma 14.5.9 shows that

$$
P\left(\left\|A_{\gamma} u\right\|_{4}^{4} \leq\left(\left(\frac{3}{n}\right)^{1 / 4}+\frac{t}{\sqrt{n}}\right)^{4}\right) \leq 1-\exp \left(-\frac{t^{2}}{2}\right) .
$$

Hence, we have that

$$
v_{4} \stackrel{\text { def }}{=}\left\|A_{\gamma} \gamma^{\prime}(0)\right\|_{4} \leq \frac{1.55}{n^{1 / 4}}+\frac{1}{2} \sqrt{h} \leq \frac{1.6}{n^{1 / 4}}
$$

with probability at least $1-\exp (-\sqrt{n} / 100)$. Since $\ell=\sqrt{n h} \leq \frac{n^{1 / 4}}{16}=\frac{1}{10 v_{4}}$, Lemma 14.5 .8 shows the middle two inequalities.

Lemma 14.5 .8 shows that $\left\|A_{\gamma} \gamma^{\prime \prime}\right\|_{\infty} \leq\left\|\gamma^{\prime \prime}\right\|_{\gamma(t)} \leq \sqrt{3} v_{4}^{2} \leq 5 n^{-1 / 2}$ and hence

$$
\begin{aligned}
\left\|A_{\gamma(t)} \gamma^{\prime}(t)\right\|_{\infty} & \leq\left\|A_{\gamma} \gamma^{\prime}(0)\right\|_{\infty}+\int_{0}^{\ell}\left\|A_{\gamma(t)} \gamma^{\prime \prime}(t)\right\|_{\infty} d t \\
& \leq \sqrt{h}+5 \ell n^{-1 / 2} \\
& =6 \sqrt{h} .
\end{aligned}
$$

## - 14.5.4 Stability of Drift

We begin with $D_{0}$ and $D_{1}$.
Lemma 14.5.11. For any $x \in M_{L}$, we have that $\|\mu(x)\|_{x}^{2} \leq n$. Hence, $D_{0} \leq \sqrt{n}$.
Proof. We have

$$
\|\mu(x)\|_{x}^{2}=\sigma_{x}^{T} A_{x}\left(A_{x}^{T} A_{x}\right)^{-1} A_{x} \sigma_{x} \leq \sum_{i}\left(\sigma_{x}\right)_{i}^{2} \leq n .
$$

Lemma 14.5.12. Let $\gamma(t)$ be a geodesic walk on $M_{L}$ with step size $h$ satisfying $\frac{1}{n} \leq h \leq \frac{1}{256 \sqrt{n}}$. Then, we have that

$$
D_{1}=\sup _{0 \leq t \leq \ell}\left|\frac{d}{d t}\|\mu(\gamma(t))\|^{2}\right| \leq 32 n \sqrt{h} .
$$

with probability at least $1-\exp (-\sqrt{n} / 100)$ over the choice of the next step $\gamma \in \widetilde{\Xi}_{x}$ for any $x \in M_{L}$.
Proof. We want to bound $\left|\|\mu(x)\|_{x}^{2}-\|\mu(y)\|_{y}^{2}\right|$. Note that

$$
\begin{aligned}
\|\mu(\gamma(t))\|_{\gamma}^{2} & =\sigma_{\gamma}^{T} A_{\gamma}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma} \sigma_{\gamma} \\
& =1^{T} \operatorname{Diag}\left(P_{\gamma}\right) P_{\gamma} \operatorname{Diag}\left(P_{\gamma}\right) 1 .
\end{aligned}
$$

Using $\frac{d}{d t} P_{\gamma}=-S_{\gamma^{\prime}} P_{\gamma}-P_{\gamma} S_{\gamma^{\prime}}+2 P_{\gamma} S_{\gamma^{\prime}} P_{\gamma}$, we have

$$
\begin{aligned}
\frac{d}{d t}\|\mu(\gamma(t))\|_{\gamma}^{2}= & -2 \cdot 1^{T} \operatorname{Diag}\left(S_{\gamma^{\prime}} P_{\gamma}\right) P_{\gamma} \operatorname{Diag}\left(P_{\gamma}\right) 1 \\
& -2 \cdot 1^{T} \operatorname{Diag}\left(P_{\gamma} S_{\gamma^{\prime}}\right) P_{\gamma} \operatorname{Diag}\left(P_{\gamma}\right) 1 \\
& +4 \cdot 1^{T} \operatorname{Diag}\left(P_{\gamma} S_{\gamma^{\prime}} P_{\gamma}\right) P_{\gamma} \operatorname{Diag}\left(P_{\gamma}\right) 1 \\
& -1^{T} \operatorname{Diag}\left(P_{\gamma}\right) S_{\gamma^{\prime}} P_{\gamma} \operatorname{Diag}\left(P_{\gamma}\right) 1 \\
& -1^{T} \operatorname{Diag}\left(P_{\gamma}\right) P_{\gamma} S_{\gamma^{\prime}} \operatorname{Diag}\left(P_{\gamma}\right) 1 \\
& +2 \cdot 1^{T} \operatorname{Diag}\left(P_{\gamma}\right) P_{\gamma} S_{\gamma^{\prime}} P_{\gamma} \operatorname{Diag}\left(P_{\gamma}\right) 1 \\
= & -2 \cdot 1^{T} \operatorname{Diag}\left(S_{\gamma^{\prime}} P_{\gamma}\right) P_{\gamma} \operatorname{Diag}\left(P_{\gamma}\right) 1 \\
& -4 \cdot 1^{T} \operatorname{Diag}\left(P_{\gamma}\right) S_{\gamma^{\prime}} P_{\gamma} \operatorname{Diag}\left(P_{\gamma}\right) 1 \\
& +4 \cdot 1^{T} \operatorname{Diag}\left(P_{\gamma} S_{\gamma^{\prime}} P_{\gamma}\right) P_{\gamma} \operatorname{Diag}\left(P_{\gamma}\right) 1 \\
& +2 \cdot 1^{T} \operatorname{Diag}\left(P_{\gamma}\right) P_{\gamma} S_{\gamma^{\prime}} P_{\gamma} \operatorname{Diag}\left(P_{\gamma}\right) 1 .
\end{aligned}
$$

Now, we note that

$$
\begin{aligned}
\left|1^{T} \operatorname{Diag}\left(S_{\gamma^{\prime}} P_{\gamma}\right) P_{\gamma} \operatorname{Diag}\left(P_{\gamma}\right) 1\right| & \leq\left\|\gamma^{\prime}\right\|_{\gamma} \sqrt{n}, \\
\left|1^{T} \operatorname{Diag}\left(P_{\gamma}\right) S_{\gamma^{\prime}} P_{\gamma} \operatorname{Diag}\left(P_{\gamma}\right) 1\right| & \leq\left\|\gamma^{\prime}\right\|_{\gamma} \sqrt{n}, \\
\left|1^{T} \operatorname{Diag}\left(P_{\gamma} S_{\gamma^{\prime}} P_{\gamma}\right) P_{\gamma} \operatorname{Diag}\left(P_{\gamma}\right) 1\right| & \leq\left\|\gamma^{\prime}\right\|_{\gamma} \sqrt{n}, \\
\left|1^{T} \operatorname{Diag}\left(P_{\gamma}\right) P_{\gamma} S_{\gamma^{\prime}} P_{\gamma} \operatorname{Diag}\left(P_{\gamma}\right) 1\right| & \leq\left\|S_{\gamma^{\prime}}\right\|_{\infty} n .
\end{aligned}
$$

Since $\left\|\gamma^{\prime}\right\|_{\gamma} \leq 2$ and $\left\|A_{\gamma} \gamma^{\prime}\right\|_{\infty} \leq 6 \sqrt{h}$ (Lemma 14.5.10), we have

$$
D_{1}=\sup \left|\frac{d}{d t}\|\mu(\gamma(t))\|^{2}\right| \leq 20 \sqrt{n}+12 n \sqrt{h} \leq 32 n \sqrt{h} .
$$

## ■ 14.5.5 Smoothness of the metric

Lemma 14.5.13. Let $\gamma$ given by a geodesic walk on $M_{L}$ with step size $h$ satisfying $\frac{1}{n} \leq h \leq \frac{1}{256 \sqrt{n}}$. Let $f(t)=\log \operatorname{det}\left(A_{\gamma(t)}^{T} A_{\gamma(t)}\right)$. Then, we have

$$
G_{1}=\sup _{0 \leq t \leq \ell}\left|f^{\prime \prime \prime}(t)\right| \leq 1000 \sqrt{h}
$$

with probability at least $1-\exp (-\sqrt{n} / 100)$.

Proof. Note that

$$
f^{\prime}(t)=-2 \operatorname{Tr}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime}} A_{\gamma}
$$

Hence, we have

$$
\begin{aligned}
f^{\prime \prime}(t)= & -4 \operatorname{Tr}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime}} A_{\gamma}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime}} A_{\gamma} \\
& +6 \operatorname{Tr}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime}}^{2} A_{\gamma} \\
& -2 \operatorname{Tr}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime \prime}} A_{\gamma}
\end{aligned}
$$

So, we have

$$
\begin{aligned}
f^{\prime \prime \prime}(t)= & -16 \operatorname{Tr}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime}} A_{\gamma}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime}} A_{\gamma}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime}} A_{\gamma} \\
& +24 \operatorname{Tr}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime}}^{2} A_{\gamma}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime}} A_{\gamma} \\
& -8 \operatorname{Tr}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime \prime}} A_{\gamma}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime}} A_{\gamma} \\
& +12 \operatorname{Tr}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime}} A_{\gamma}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime}}^{2} A_{\gamma} \\
& -24 \operatorname{Tr}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime}}^{3} A_{\gamma} \\
+ & 12 \operatorname{Tr}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime}} S_{\gamma^{\prime \prime}} A_{\gamma} \\
& -4 \operatorname{Tr}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime}} A_{\gamma}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime \prime}} A_{\gamma} \\
& +4 \operatorname{Tr}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime \prime}} S_{\gamma^{\prime}} A_{\gamma} \\
& -2 \operatorname{Tr}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} \frac{d}{d t} S_{\gamma^{\prime \prime}} A_{\gamma} \\
= & -16 \operatorname{Tr} P_{\gamma} S_{\gamma^{\prime}} P_{\gamma} S_{\gamma^{\prime}} P_{\gamma} S_{\gamma^{\prime}}+36 \operatorname{Tr} P_{\gamma} S_{\gamma^{\prime}}^{2} P_{\gamma} S_{\gamma^{\prime}} \\
& -12 \operatorname{Tr} P_{\gamma} S_{\gamma^{\prime \prime}} P_{\gamma} S_{\gamma^{\prime}}-24 \operatorname{Tr} P_{\gamma} S_{\gamma^{\prime}}^{3} \\
& +16 \operatorname{Tr} P_{\gamma} S_{\gamma^{\prime}} S_{\gamma^{\prime \prime}}-2 \operatorname{Tr} P_{\gamma} \frac{d}{d t} S_{\gamma^{\prime \prime}}
\end{aligned}
$$

Hence, we have

$$
\left|f^{\prime \prime \prime}(t)\right| \leq(16+36+24) \sqrt{\operatorname{Tr} S_{\gamma^{\prime}}^{4} \operatorname{Tr} S_{\gamma^{\prime}}^{2}}+(12+16) \sqrt{\operatorname{Tr} S_{\gamma^{\prime \prime}}^{2} \operatorname{Tr} S_{\gamma^{\prime}}^{2}}+2\left|\operatorname{Tr} P_{\gamma} \frac{d}{d t} S_{\gamma^{\prime \prime}}\right|
$$

Since $\operatorname{Tr} S_{\gamma^{\prime}}^{2} \leq 2, \operatorname{Tr} S_{\gamma^{\prime}}^{4} \leq 16 n^{-1}, \operatorname{Tr} S_{\gamma^{\prime \prime}}^{2} \leq 15 n^{-1}$ (Lemma 14.5.10), we have

$$
\begin{aligned}
\left|f^{\prime \prime \prime}(t)\right| & \leq 76 \sqrt{2 \cdot 16 n^{-1}}+28 \sqrt{2 \cdot 15 n^{-1}}+2\left|\operatorname{Tr} P_{\gamma} \frac{d}{d t} S_{\gamma^{\prime \prime}}\right| \\
& \leq 600 n^{-1 / 2}+2\left|\operatorname{Tr} P_{\gamma} \frac{d}{d t} S_{\gamma^{\prime \prime}}\right|
\end{aligned}
$$

To bound the last term, we start with the geodesic equation: $s_{\gamma^{\prime \prime}}=P_{\gamma} s_{\gamma^{\prime}}^{2}$. Since $\frac{d}{d t} P_{\gamma}=-S_{\gamma^{\prime}} P_{\gamma}-$ $P_{\gamma} S_{\gamma^{\prime}}+2 P_{\gamma} S_{\gamma^{\prime}} P_{\gamma}$, we have that

$$
\begin{aligned}
\frac{d}{d t} s_{\gamma^{\prime \prime}}= & -S_{\gamma^{\prime}} P_{\gamma} s_{\gamma^{\prime}}^{2}-P_{\gamma} S_{\gamma^{\prime}} s_{\gamma^{\prime}}^{2}+2 P_{\gamma} S_{\gamma^{\prime}} P_{\gamma} s_{\gamma^{\prime}}^{2} \\
& +2 P_{\gamma} s_{\gamma^{\prime}} s_{\gamma^{\prime \prime}}-2 P_{\gamma} s_{\gamma^{\prime}}^{3} \\
= & -S_{\gamma^{\prime}} P_{\gamma} s_{\gamma^{\prime}}^{2}-3 P_{\gamma} s_{\gamma^{\prime}}^{3}+2 P_{\gamma} S_{\gamma^{\prime}} P_{\gamma} s_{\gamma^{\prime}}^{2}+2 P_{\gamma} s_{\gamma^{\prime}} s_{\gamma^{\prime \prime}} \\
= & -S_{\gamma^{\prime}} P_{\gamma} s_{\gamma^{\prime}}^{2}-3 P_{\gamma} s_{\gamma^{\prime}}^{3}+4 P_{\gamma} S_{\gamma^{\prime}} P_{\gamma} s_{\gamma^{\prime}}^{2}
\end{aligned}
$$

Hence, we have that

$$
\begin{aligned}
\left|\operatorname{Tr} P_{\gamma} \frac{d}{d t} S_{\gamma^{\prime \prime}}\right| & \leq\left|\sigma_{\gamma}^{T} S_{\gamma^{\prime}} P_{\gamma} s_{\gamma^{\prime}}^{2}\right|+3\left|\sigma_{\gamma}^{T} P_{\gamma} s_{\gamma^{\prime}}^{3}\right|+4\left|\sigma_{\gamma}^{T} P_{\gamma} S_{\gamma^{\prime}} P_{\gamma} s_{\gamma^{\prime}}^{2}\right| \\
& \leq\left|\left(S_{\gamma^{\prime}} \sigma_{\gamma}\right)^{T} P_{\gamma} s_{\gamma^{\prime}}^{2}\right|+3\left|\left(S_{\gamma^{\prime}} P_{\gamma} \sigma_{\gamma}\right)^{T} s_{\gamma^{\prime}}^{2}\right|+4\left(S_{\gamma^{\prime}} P_{\gamma} \sigma_{\gamma}\right)^{T}\left(P_{\gamma} s_{\gamma^{\prime}}^{2}\right) \\
& \leq \sqrt{\sum s_{\gamma^{\prime}}^{2} \sum s_{\gamma^{\prime}}^{4}}+28 n^{-1 / 2} \sqrt{\sigma_{\gamma}^{T} P_{\gamma} S_{\gamma^{\prime}}^{2} P_{\gamma} \sigma_{\gamma}} \\
& \leq 8 n^{-1 / 2}+168 \sqrt{h} \leq 180 \sqrt{h} .
\end{aligned}
$$

Hence, we have that $\left|f^{\prime \prime \prime}(t)\right| \leq 600 n^{-1 / 2}+360 \sqrt{h} \leq 1000 \sqrt{h}$.

## - 14.5.6 Stability of Jacobi Field

Now, we give a bound on $R_{1}$.
Lemma 14.5.14. Let $\gamma$ given by a geodesic walk on $M_{L}$ with step size $h$ satisfying $\frac{1}{n} \leq h \leq \frac{1}{256 \sqrt{n}}$. With probability at least $1-\exp (-\sqrt{n} / 100)$ over the choice of the next step $\gamma \in \widetilde{\Xi}_{x}$, we have that

$$
R_{1}=\sup _{0 \leq t \leq \ell}\|R(t)\|_{F} \leq \frac{8}{\sqrt{n}}
$$

where $R(t)_{i j}=\left\langle R\left(X_{i}(t), \gamma^{\prime}(t)\right) \gamma^{\prime}(t), X_{j}(t)\right\rangle$ and $\left\{X_{i}\right\}_{i=1}^{n}$ is any orthogonal basis at $\gamma(t)$.
Proof. Let $\overline{R(t)}$ such that $\left\langle R\left(u, \gamma^{\prime}(t)\right) \gamma^{\prime}(t), v\right\rangle=u^{T} \overline{R(t)} v$. Lemma 14.5.4 shows that

$$
u^{T} \overline{R(t)} v=s_{u}^{T} S_{\gamma^{\prime}} A_{\gamma}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime}} s_{v}-\left(s_{v} s_{u}\right)^{T} A_{\gamma}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} s_{\gamma^{\prime}}^{2} .
$$

Hence, we have

$$
\overline{R(t)}=A_{\gamma}^{T} S_{\gamma^{\prime}} P_{\gamma} S_{\gamma^{\prime}} A_{\gamma}-A_{\gamma}^{T} \operatorname{Diag}\left(P_{\gamma} s_{\gamma^{\prime}}^{2}\right) A_{\gamma} .
$$

Pick $X_{i}=\left(A^{T} A_{\gamma}\right)^{-1 / 2} e_{i}$. Then, we have

$$
R(t)=\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1 / 2}\left(A_{\gamma}^{T} S_{\gamma^{\prime}} P_{\gamma} S_{\gamma^{\prime}} A_{\gamma}-A_{\gamma}^{T} \operatorname{Diag}\left(P_{\gamma} s_{\gamma^{\prime}}^{2}\right) A_{\gamma}\right)\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1 / 2} .
$$

The claim follows from Lemma 14.5.10 and the calculation

$$
\begin{aligned}
& \|R(t)\|_{F}^{2} \\
\leq & 2\left\|\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1 / 2} A_{\gamma}^{T} S_{\gamma^{\prime}} P_{\gamma} S_{\gamma^{\prime}} A_{\gamma}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1 / 2}\right\|_{F}^{2}+2\left\|\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1 / 2} A_{\gamma}^{T} \operatorname{Diag}\left(P_{\gamma} s_{\gamma^{\prime}}^{2}\right) A_{\gamma}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1 / 2}\right\|_{F}^{2} \\
= & 2 \operatorname{Tr} P_{\gamma} S_{\gamma^{\prime}} P_{\gamma} S_{\gamma^{\prime}} P_{\gamma} S_{\gamma^{\prime}} P_{\gamma} S_{\gamma^{\prime}}+2 \operatorname{Tr} P_{\gamma} \operatorname{Diag}\left(P_{\gamma} s_{\gamma^{\prime}}^{2}\right) P_{\gamma} \operatorname{Diag}\left(P_{\gamma} s_{\gamma^{\prime}}^{2}\right) \leq 4\left\|s_{\gamma^{\prime}}\right\|_{4}^{4} .
\end{aligned}
$$

## - 14.5.7 Smoothness of One-Step Distribution

We begin with $R_{2}$.
Lemma 14.5.15. Given a geodesic $\gamma(t)$ on $M_{L}$ with step size $h$ satisfying $\frac{1}{n} \leq h \leq \frac{1}{256 \sqrt{n}}$, with probability at least $1-\exp (-\sqrt{n} / 100)$ over the choice of the next step $\gamma \in \widetilde{\Xi}_{x}$, for any $t$ such that $0 \leq t \leq \ell$, any curve $c(s)$ starting from $\gamma(t)$ and any vector field $v(s)$ on $c(s)$, we have that

$$
\left.\left|\frac{d}{d s} \operatorname{Ric}(v(s))\right|_{s=0} \right\rvert\, \leq 64 \sqrt{n h}\left\|c^{\prime}(0)\right\|+6 \sqrt{n h}\left\|D_{s} v(0)\right\| .
$$

Therefore, $R_{2}=O(\sqrt{n h})$.
Proof. By Lemma 14.5.4, we know that

$$
\begin{aligned}
\operatorname{Ric}(v(s)) & =s_{c(s), v(s)}^{T} P_{c(s)}^{(2)} s_{c(s), v(s)}-\sigma_{c(s)}^{T} P_{c(s)} s_{c(s), v(s)}^{2} \\
& =\operatorname{Tr}\left(S_{c(s), v(s)} P_{c(s)} S_{c(s), v(s)} P_{c(s)}\right)-\operatorname{Tr}\left(\operatorname{Diag}\left(P_{c(s)} s_{c(s), v(s)}^{2}\right) P_{c(s)}\right) .
\end{aligned}
$$

For simplicity, we suppress the parameter $s$ and hence, we have

$$
\operatorname{Ric}(v)=\operatorname{Tr}\left(S_{c, v} P_{c} S_{c, v} P_{c}\right)-\operatorname{Tr}\left(\operatorname{Diag}\left(P_{c} S_{c, v}^{2}\right) P_{c}\right)
$$

Now, we use that

$$
\begin{aligned}
\frac{d}{d s} P_{c} & =-S_{c^{\prime}} P_{c}-P_{c} S_{c^{\prime}}+2 P_{c} S_{c^{\prime}} P_{c} \\
\frac{d}{d s} S_{c, v} & =-S_{c^{\prime}} S_{c, v}+S_{c, v^{\prime}}
\end{aligned}
$$

and get

$$
\begin{aligned}
& \frac{d}{d s} \operatorname{Ric}(v) \\
= & -2 \operatorname{Tr}\left(S_{c, v} S_{c^{\prime}} P_{c} S_{c, v} P_{c}\right)-2 \operatorname{Tr}\left(S_{c, v} P_{c} S_{c^{\prime}} S_{c, v} P_{c}\right)+2 \operatorname{Tr}\left(S_{c, v} P_{c} S_{c^{\prime}} P_{c} S_{c, v} P_{c}\right) \\
& -2 \operatorname{Tr}\left(S_{c^{\prime}} S_{c, v} P_{c} S_{c, v} P_{c}\right)+2 \operatorname{Tr}\left(S_{c, v^{\prime}} P_{c} S_{c, v} P_{c}\right) \\
& +\operatorname{Tr}\left(\operatorname{Diag}\left(P_{c} s_{c, v}^{2} S_{c^{\prime}} P_{c}\right)+\operatorname{Tr}\left(\operatorname{Diag}\left(P_{c} s_{c, v}^{2}\right) P_{c} S_{c^{\prime}}\right)-2 \operatorname{Tr}\left(\operatorname{Diag}\left(P_{c} s_{c, v}^{2}\right) P_{c} S_{c^{\prime}} P_{c}\right)\right. \\
& +\operatorname{Tr}\left(\operatorname{Diag}\left(P_{c} S_{c^{\prime}} s_{c, v}^{2}\right) P_{c}\right)+\operatorname{Tr}\left(\operatorname{Diag}\left(S_{c^{\prime}} P_{c} s_{c, v}^{2}\right) P_{c}\right)-2 \operatorname{Tr}\left(\operatorname{Diag}\left(P_{c} S_{c^{\prime}} P_{c} s_{c, v}^{2}\right) P_{c}\right) \\
& +2 \operatorname{Tr}\left(\operatorname{Diag}\left(P_{c} S_{c, v} S_{c^{\prime}} s_{c, v}\right) P_{c}\right)-2 \operatorname{Tr}\left(\operatorname{Diag}\left(P_{c} S_{c, v} s_{c, v^{\prime}}\right) P_{c}\right) \\
= & -6 \operatorname{Tr}\left(S_{c, v} S_{c^{\prime}} P_{c} S_{c, v} P_{c}\right)+2 \operatorname{Tr}\left(S_{c, v} P_{c} S_{c^{\prime}} P_{c} S_{c, v} P_{c}\right)+2 \operatorname{Tr}\left(S_{c, v^{\prime}} P_{c} S_{c, v} P_{c}\right) \\
& +2 \operatorname{Tr}\left(\operatorname{Diag}\left(P_{c} s_{c, v}^{2}\right) S_{c^{\prime}} P_{c}\right)-2 \operatorname{Tr}\left(\operatorname{Diag}\left(P_{c} s_{c, v}^{2}\right) P_{c} S_{c^{\prime}} P_{c}\right) \\
& +3 \operatorname{Tr}\left(\operatorname{Diag}\left(P_{c} S_{c^{\prime}}^{2} s_{c, v}\right) P_{c}\right)+\operatorname{Tr}\left(\operatorname{Diag}\left(S_{c^{\prime}} P_{c} s_{c, v}^{2}\right) P_{c}\right)-2 \operatorname{Tr}\left(\operatorname{Diag}\left(P_{c} S_{c^{\prime}} P_{c} s_{c, v}^{2}\right) P_{c}\right) \\
& -2 \operatorname{Tr}\left(\operatorname{Diag}\left(P_{c} S_{c, v} s_{c, v^{\prime}}\right) P_{c}\right) .
\end{aligned}
$$

Let $\frac{d}{d s} \operatorname{Ric}(v)=(1)+(2)$ where (1) is the sum of all terms not involving $v^{\prime}$ and (2) is the sum of other terms.

For the first term (1), let $L=\left\|s_{c^{\prime}(s)}\right\|_{2}$. Then, we have that

$$
\begin{aligned}
|(1)| \leq & 6\left|\operatorname{Tr}\left(S_{c, v} S_{c^{\prime}} P_{c} S_{c, v} P_{c}\right)\right|+2\left|\operatorname{Tr}\left(S_{c, v} P_{c} S_{C^{\prime}} P_{c} S_{c, v} P_{c}\right)\right| \\
& +2\left|\operatorname{Tr}\left(\operatorname{Diag}\left(P_{c} s_{c, v}^{2}\right) S_{c^{\prime}} P_{c}\right)\right|+2\left|\operatorname{Tr}\left(\operatorname{Diag}\left(P_{c} s_{c, v}^{2}\right) P_{c} S_{c^{\prime}} P_{c}\right)\right| \\
& +3 \mid \operatorname{Tr}\left(\operatorname{Diag}\left(P_{c} S_{c^{\prime}}^{2} s_{c, v}^{2} P_{c}\right)\left|+\left|\operatorname{Tr}\left(\operatorname{Diag}\left(S_{c^{\prime}} P_{c} s_{c, v}^{2}\right) P_{c}\right)\right|+2\right| \operatorname{Tr}\left(\operatorname{Diag}\left(P_{c} S_{c^{\prime}} P_{c} s_{c, v}^{2}\right) P_{c}\right) \mid\right. \\
\leq & 6 L \sqrt{\sum_{i}\left(s_{c, v}\right)_{i}^{2}} \sqrt{\sum_{i}\left(s_{c, v}\right)_{i}^{2}}+2 L\left|\operatorname{Tr}\left(S_{c, v} P_{c} S_{c, v} P_{c}\right)\right| \\
& +2 L \sqrt{\sum_{i}\left(s_{c, v}\right)_{i}^{4}} \sqrt{\sum_{i}\left(P_{c}\right)_{i i}^{2}}+2 L \sqrt{\sum_{i}\left(s_{c, v}\right)_{i}^{4}} \sqrt{\sum_{i}\left(P_{c} S_{c^{\prime}} P_{c}\right)_{i i}^{2}} \\
& +3 L \sqrt{\sum_{i}\left(s_{c, v}\right)_{i}^{4}} \sqrt{\sum_{i}\left(P_{c}\right)_{i i}^{2}}+L \sqrt{\sum_{i}\left(s_{c, v}\right)_{i}^{4}} \sqrt{\sum_{i}\left(P_{c}\right)_{i i}^{2}}+2 L \sqrt{\sum_{i}\left(s_{c, v}\right)_{i}^{4}} \sqrt{\sum_{i}\left(P_{c}\right)_{i i}^{2}} \\
\leq & 8 L \sqrt{\sum_{i}\left(s_{c, v}\right)_{i}^{2}} \sqrt{\sum_{i}\left(s_{c, v}\right)_{i}^{2}}+10 \sqrt{n} L \sqrt{\sum_{i}\left(s_{c, v}\right)_{i}^{4}} .
\end{aligned}
$$

Since $v(0)=\gamma^{\prime}(t)$, Lemma 14.5.10 shows that $\left\|s_{c, v}\right\|_{2} \leq 2$ and $\left\|s_{c, v}\right\|_{4} \leq 2 n^{-1 / 4}$. Hence, we have

$$
|(1)| \leq 52 L \text {. }
$$

For the second term (2), we note that

$$
D_{s} v=\frac{d v}{d s}-\left(A_{c}^{T} A_{c}\right)^{-1} A_{c}^{T} S_{c^{\prime}} s_{c, v} .
$$

Therefore, we have

$$
s_{c, v^{\prime}}=A_{c}\left(D_{s} v\right)-A_{c}\left(A_{c}^{T} A_{c}\right)^{-1} A_{c}^{T} S_{c^{\prime}} s_{c, v}
$$

and hence

$$
\begin{aligned}
\left\|s_{c, v^{\prime}}\right\|_{2} & \leq\left\|D_{s} v\right\|+\left\|A_{c}\left(A_{c}^{T} A_{c}\right)^{-1} A_{c}^{T} S_{c^{\prime}} s_{c, v}\right\|_{2} \\
& \leq\left\|D_{s} v\right\|+L\left\|s_{c, v}\right\|_{2} \\
& \leq 2 L+\left\|D_{s} v\right\| .
\end{aligned}
$$

Also, note that Lemma 14.5 .10 shows that $\left\|s_{c, v}\right\|_{\infty} \leq 6 \sqrt{h}$. Hence, we have that

$$
\begin{aligned}
|(2)| & \leq 2\left|\operatorname{Tr}\left(S_{c, v^{\prime}} P_{c} S_{c, v} P_{c}\right)\right|+2\left|\operatorname{Tr}\left(\operatorname{Diag}\left(P_{c} S_{c, v} s_{c, v^{\prime}}\right) P_{c}\right)\right| \\
& \leq 2\left\|s_{c, v^{\prime}}\right\| 2\left\|s_{c, v}\right\|_{2}+2 \sqrt{n} \sqrt{\sum_{i}\left(s_{c, v^{\prime}} s_{c, v}\right)_{i}^{2}} \\
& \leq 4\left(2+\left\|D_{s} v\right\|\right)+2 \sqrt{n h}\left(2+\left\|D_{s} v\right\|\right) \\
& \leq 8+4\left\|D_{s} v\right\|+4 \sqrt{n h}+2 \sqrt{n h}\left\|D_{s} v\right\| \\
& \leq 12 \sqrt{n h}+6 \sqrt{n h}\left\|D_{s} v\right\| .
\end{aligned}
$$

Therefore, we have

$$
\begin{aligned}
\left|\frac{d}{d s} \operatorname{Ric}(v)\right| & \leq 52 L+12 \sqrt{n h} L+6 \sqrt{n h}\left\|D_{s} v\right\| \\
& \leq 64 \sqrt{n h} L+6 \sqrt{n h}\left\|D_{s} v\right\|
\end{aligned}
$$

Finally, we bound $D_{2}$.

Lemma 14.5.16. Given $x \in M_{L}$ and any curve $c(t)$ starting at $x$ with unit speed. We have that

$$
\left\|\frac{D}{d t} \mu(c(t))\right\| \leq 9 \sqrt{n}
$$

Hence, $D_{2} \leq 9 \sqrt{n}$.

Proof. Recall that $\frac{D}{d t} \mu(c(t))=\frac{d \mu(c)}{d t}-\left(A_{c}^{T} A_{c}\right)^{-1} A_{c}^{T} S_{c, \mu} s_{c^{\prime}}$. We bound the terms separately. For the second term, since $c$ is an unit speed curve, we have that $\left\|s_{c^{\prime}}\right\|_{2}=1$ and

$$
\begin{aligned}
\left\|\left(A_{c}^{T} A_{c}\right)^{-1} A_{c}^{T} S_{c^{\prime}} A_{c} \mu(c)\right\|_{c} & \leq\left\|S_{c, \mu} s_{c^{\prime}}\right\|_{2} \leq\left\|S_{c, \mu}\right\|_{\infty} \\
& =\left\|A_{c}\left(A_{c}^{T} A_{c}\right)^{-1} A_{c}^{T} \sigma_{c}\right\|_{\infty} \leq \sqrt{n} .
\end{aligned}
$$

For the first term,

$$
\begin{aligned}
\frac{d}{d t}\left(A_{c} \mu(c)\right)= & \frac{d}{d t}\left(P_{c} \boldsymbol{\operatorname { d i a g }}\left(P_{c}\right)\right) \\
= & -S_{c^{\prime}} P_{c} \boldsymbol{\operatorname { d i a g }}\left(P_{c}\right)-P_{c} S_{S^{\prime}} \boldsymbol{\operatorname { d i a g }}\left(P_{c}\right)+2 P_{c} S_{c^{\prime}} P_{c} \operatorname{diag}\left(P_{c}\right) \\
& -P_{c} \operatorname{diag}\left(S_{c^{\prime}} P_{c}\right)-P_{c} \boldsymbol{\operatorname { d i a g }}\left(P_{c} S_{c^{\prime}}\right)+2 P_{c} \boldsymbol{\operatorname { d i a g }}\left(P_{c} S_{c^{\prime}} P_{c}\right) \\
= & -S_{c^{\prime}} P_{c} \boldsymbol{\operatorname { d i a g }}\left(P_{c}\right)-2 P_{c} S_{C^{\prime}} \boldsymbol{\operatorname { d i a g }}\left(P_{c}\right)+2 P_{c^{\prime}} S_{c^{\prime}} P_{c} \operatorname{diag}\left(P_{c}\right) \\
& -P_{c} \boldsymbol{\operatorname { d i a g }}\left(P_{c} S_{c^{\prime}}\right)+2 P_{c} \operatorname{diag}\left(P_{c} S_{c^{\prime}} P_{c}\right) .
\end{aligned}
$$

Using $\left\|s_{c^{\prime}}\right\|_{2}=1$, we have

$$
\begin{aligned}
\left\|\frac{d}{d t} \mu(c)\right\|_{c} \leq & \left\|S_{c^{\prime}} A_{c} \mu(c)\right\|_{2} \\
& +\left\|S_{c^{\prime}} P_{c} \boldsymbol{\operatorname { d i a g }}\left(P_{c}\right)\right\|_{2}+2\left\|P_{c} S_{c^{\prime}} \boldsymbol{\operatorname { d i a g }}\left(P_{c}\right)\right\|_{2}+2\left\|P_{c} S_{c^{\prime}} P_{c} \boldsymbol{\operatorname { d i a g }}\left(P_{c}\right)\right\|_{2} \\
& +\left\|P_{c} \boldsymbol{\operatorname { d i a g }}\left(P_{c} S_{c^{\prime}}\right)\right\|_{2}+\left\|P_{c} \boldsymbol{\operatorname { d i a g }}\left(P_{c} S_{c^{\prime}} P_{c}\right)\right\|_{2} \\
\leq & 8 \sqrt{n} .
\end{aligned}
$$

Combining both terms, we have the result.

## ■ 14.5.8 Mixing time

Proof of Theorem 14.5.1. In the last previous sections, we proved that if $\frac{1}{n} \leq h \leq \frac{1}{256 \sqrt{n}}$

1. $D_{0} \leq O(\sqrt{n})$ (Lemma 14.5.11)
2. $D_{1}=O(n \sqrt{h})$ (Lemma 14.5.12)
3. $D_{2}=O(\sqrt{n})$ (Lemma 14.5.16)
4. $G_{1}=O(\sqrt{h})$ (Lemma 14.5.13)
5. $G_{2}=O(\sqrt{m})$ (Lemma 14.5.6)
6. $R_{1}=O(1 / \sqrt{n})$ (Lemma 14.5.14)
7. $R_{2}=O(\sqrt{n h})($ Lemma 14.5.15)

It is easy to see that all these parameters decrease if $h$ decreases. So, the condition $\frac{1}{n} \leq h$ can be ignored. We only use randomness to prove that $\left\|A_{\gamma(0)} u\right\|_{2} \leq \frac{3}{2}$ and $\left\|A_{\gamma} u\right\|_{4} \leq \frac{1.55}{n^{1 / 4}}$ (Lemma 14.5.10) and they can be checked. These bounds hold with exponentially small failure probability. Hence, the measure of $\Xi$ is large as claimed and the result follows. Theorem 14.4.2 implies that the walk has mixing time $O\left(G_{2}^{2} / h\right)$ as long as

$$
h \leq \min \left\{\frac{1}{\left(n D_{0} R_{1}\right)^{2 / 3}}, \frac{1}{D_{2}}, \frac{1}{n R_{1}}, \frac{1}{n^{1 / 3} D_{1}^{2 / 3}}, \frac{1}{n G_{1}^{2 / 3}}, \frac{1}{\left(n R_{2}\right)^{2 / 3}}\right\} \leq \frac{1}{\Theta\left(n^{3 / 4}\right)}
$$

## - 14.5.9 Complex Analyticity of geodesics, parallel transport and Jacobi fields

To show geodesics, parallel transport and Jacobi fields on $M_{L}$ are complex analytic, we note that the Cauchy-Kowalevski theorem shows that if a differential equation is complex analytic, then the equation has a unique complex analytic solution.

Theorem 14.5.17 (Simplified Version of Cauchy-Kowalevski theorem). If $f$ is a complex analytic function defined on a neighborhood of $\left(z_{0}, \alpha\right) \in \mathbb{C}^{n+1}$, then the problem

$$
\frac{d w}{d z}=f(z, w), \quad w\left(z_{0}\right)=\alpha
$$

has a unique complex analytic solution $w$ defined on a neighborhood around $z_{0}$.
Similarly, for a complex analytic function $f$ defined in a neighborhood of $\left(z_{0}, \alpha, \beta\right) \in \mathbb{C}^{2 n+1}$, the ODE

$$
\frac{d^{2} w}{d z^{2}}=f\left(z, w, \frac{d w}{d z}\right), \quad w\left(z_{0}\right)=\alpha, \quad \frac{d w}{d z}\left(z_{0}\right)=\beta
$$

has a unique complex analytic solution $w$ defined in a neighborhood around $z_{0}$.
As we see in section 14.5.1, the equations for geodesic, parallel transport and Jacobi field involve only rational functions. Since rational functions are complex analytic, the Cauchy-Kowalevski theorem shows that geodesic, orthogonal frame and Jacobi field are complex analytic.

Lemma 14.5.18. Geodesics, parallel transport and Jacobi fields are complex analytic for the Hessian manifold induced by the logarithmic barrier.

### 14.5.9.1 Radius of Convergence

Next we bound the higher-order derivatives of geodesic, parallel transport and Jacobi field, using techniques developed in Section 14.7 to. Since these are complex analytic, we get a bound for their radius of convergence.

The purpose of these derivative bounds is to show that the solutions of the corresponding ODEs are well-approximated by low-degree polynomials, where the degree of the polynomial grows as $\log \left(\frac{1}{\varepsilon}\right)$ for desired accuracy $\varepsilon$. The bound on the degree also implies that the Collocation method for solving ODEs is efficient (roughly matrix multiplication time).

### 14.5.9.2 Geodesic

Motivated from the geodesic equation under Euclidean coordinate (14.19), we define the following auxiliary function

$$
\begin{equation*}
\bar{F}_{\eta}(y, x, t)=\left(A_{x+t \eta}^{T} A_{x+t \eta}\right)^{-1} A_{x+t \eta}^{T} s_{x+t \eta, y+\eta}^{2} . \tag{14.24}
\end{equation*}
$$

The derivative bounds on geodesic rely on the smoothness of this auxiliary function.
Lemma 14.5.19. Under the normalization $A^{T} A=I, S_{x}=I$, we have that

$$
\bar{F}_{\eta}(y, x, t) \leq_{(0, x, 0)} \frac{4\left(\|A \eta\|_{4}+1\right)^{2}}{1-\max \left(8+8\|A \eta\|_{\infty}, 1\right) t}
$$

where $\bar{F}_{\eta}$ is defined in (14.24).
Proof. By the assumption that $S_{x}=I$, we have that $\left\|S_{x}\right\|_{2}=1$. Using $A^{T} A=I$, we have that

$$
\begin{aligned}
\left\|D S_{x+t \eta}(x, t)\left[\left(d_{x}, d_{t}\right)\right]\right\|_{2} & =\left\|\operatorname{Diag}\left(A d_{x}+d_{t} A \eta\right)\right\|_{2} \\
& \leq\left\|A d_{x}\right\|_{\infty}+\left\|d_{t} A \eta\right\|_{\infty} \\
& \leq\left\|d_{x}\right\|_{2}+\left|d_{t}\right|\|A \eta\|_{\infty} \\
& \leq\left(1+\|A \eta\|_{\infty}\right)\left\|\left(d_{x}, d_{t}\right)\right\|_{2} .
\end{aligned}
$$

Let $\beta=1+\|A \eta\|_{\infty}$. Then, we have that $S_{x+t \eta} \leq(x, t=0) 1+\beta t$.

By using the inverse formula (Lemma 14.7.3),

$$
S_{x+t \eta}^{-1} \leq_{(x, 0)} \frac{1}{1-\beta t}
$$

Using $A^{T} A=I$, we have that $A \leq_{(x, 0)} 1$ and hence product formula (Lemma 14.7.3) shows that

$$
A_{x+t \eta}=S_{x+t \eta}^{-1} A \leq_{(x, 0)} \frac{1}{1-\beta t}
$$

Since $A^{T} A=I$, we have that $A A^{T} \preceq I$ and hence $A_{x+t \eta}^{T} \leq_{(x, 0)} \frac{1}{1-\beta t}$. Therefore, we have that

$$
A_{x+t \eta}^{T} A_{x+t \eta} \leq_{(x, 0)} \frac{1}{(1-\beta t)^{2}}
$$

By using the inverse formula again,

$$
\left(A_{x+t \eta}^{T} A_{x+t \eta}\right)^{-1} \leq_{(x, 0)} \frac{1}{2-\frac{1}{(1-\beta t)^{2}}}=\frac{(1-\beta t)^{2}}{2(1-\beta t)^{2}-1}
$$

Hence, we have that

$$
\left(A_{x+t \eta}^{T} A_{x+t \eta}\right)^{-1} A_{x+t \eta}^{T} \leq_{(x, 0)} \frac{(1-\beta t)^{2}}{2(1-\beta t)^{2}-1} \frac{1}{1-\beta t}=\frac{1-\beta t}{2(1-\beta t)^{2}-1}
$$

Now, we consider the function $H(y)=(A(y+\eta))^{2}$. Note that

$$
\begin{aligned}
\|H(y)\|_{2} & \leq\|A(y+\eta)\|_{4}^{2} \\
\|D H(y)[d]\|_{2} & =2\|(A(y+\eta)) A d\|_{2} \leq 2\|A(y+\eta)\|_{4}\|d\|_{2} \\
\left\|D^{2} H(y)[d, d]\right\|_{2} & \leq 2\left\|(A d)^{2}\right\|_{2} \leq 2\|d\|_{2}^{2}
\end{aligned}
$$

Therefore, we have that

$$
H \leq_{y=0}\|A \eta\|_{4}^{2}+2\|A \eta\|_{4} t+t^{2}=\left(\|A \eta\|_{4}+t\right)^{2}
$$

Hence, we have that

$$
\bar{F}_{\eta}(y, x, t)=\left(A_{x+t \eta}^{T} A_{x+t \eta}\right)^{-1} A_{x+t \eta}^{T} s_{x+t \eta, y+\eta}^{2} \leq_{(0, x, 0)} \frac{(1-\beta t)\left(\|A \eta\|_{4}+t\right)^{2}}{2(1-\beta t)^{2}-1} \frac{1}{(1-\beta t)^{2}}
$$

Let $\phi(t)=\frac{\left(\|A \eta\|_{4}+t\right)^{2}}{\left(2(1-\beta t)^{2}-1\right)(1-\beta t)}$ and we write $\phi(t)=\sum_{k=0}^{\infty} a_{k} t^{k}$. For any $|z|=\frac{1}{8} \min \left(\frac{1}{\beta},\|A \eta\|_{4}+8\right)$, we have that

$$
\begin{aligned}
|\phi(z)| & \leq 3\left(\|A \eta\|_{4}+\frac{\|A \eta\|_{4}}{8}+1\right)^{2} \\
& \leq 4\left(\|A \eta\|_{4}+1\right)^{2}
\end{aligned}
$$

Theorem 14.2.11 shows that

$$
\begin{aligned}
\left|a_{k}\right| & \leq 4\left(\|A \eta\|_{4}+1\right)^{2}\left(8 \max \left(\beta,\left(\|A \eta\|_{4}+8\right)^{-1}\right)\right)^{k} \\
& \leq 4\left(\|A \eta\|_{4}+1\right)^{2}(\max (8 \beta, 1))^{k}
\end{aligned}
$$

Hence, we can instead bound $\bar{F}_{\eta}$ by

$$
\bar{F}_{\eta} \leq(0, x, 0) \frac{4\left(\|A \eta\|_{4}+1\right)^{2}}{1-\max \left(8+8\|A \eta\|_{\infty}, 1\right) t}
$$

Now, we prove the geodesic has large radius of convergence.
Lemma 14.5.20. Under the normalization $A^{T} A=I, S_{x}=I$ and Euclidean coordinate, any geodesic starting at $x$ satisfies the bound

$$
\left\|\gamma^{(k)}(0)\right\|_{2} \leq k!c^{k}
$$

for all $k \geq 2$ where $c=512\left\|A \gamma^{\prime}(0)\right\|_{4}$.
Proof. Recall that under Euclidean coordinate, the geodesic equation is given by

$$
\gamma^{\prime \prime}=\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} s_{\gamma^{\prime}}^{2} \stackrel{\text { def }}{=} F\left(\gamma^{\prime}, \gamma\right)
$$

So, we have that $\bar{F}_{\eta}(y, x, t)=\alpha^{2} F\left(\alpha^{-1} y(t)+\gamma^{\prime}(0), x(t)+\alpha t \gamma^{\prime}(0)\right)$ with $\eta=\alpha \gamma^{\prime}(0)$.
Now, we estimate $\bar{F}_{\eta}$. Lemma 14.5 .19 shows that

$$
\bar{F}_{\eta} \leq_{(0, x, 0)} \frac{4\left(\|A \eta\|_{4}+1\right)^{2}}{1-\max \left(8+8\|A \eta\|_{\infty}, 1\right) t}=\frac{4\left(\alpha\left\|A \gamma^{\prime}(0)\right\|_{4}+1\right)^{2}}{1-\max \left(8+8 \alpha\left\|A \gamma^{\prime}(0)\right\|_{\infty}, 1\right) t}
$$

Setting $\alpha=\left\|A \gamma^{\prime}(0)\right\|_{4}^{-1}$ and using $\left\|A \gamma^{\prime}(0)\right\|_{\infty} \leq\left\|A \gamma^{\prime}(0)\right\|_{4}$, we have that

$$
\bar{F}_{\eta} \leq(0, x, 0) \frac{16}{1-16 t}
$$

Lemma 14.7 .5 shows that

$$
\left\|\gamma^{(k)}(0)\right\|_{2} \leq \frac{\psi^{(k)}(0)}{\alpha^{k}}
$$

for all $k \geq 2$ where $\psi(t)$ is the solution of

$$
\psi^{\prime}(t)=\frac{16}{1-16 \psi(t)} \text { with } \psi(0)=0
$$

Solving it, we get that

$$
\psi(t)=\frac{1}{16}(1-\sqrt{1-512 t})
$$

By Theorem 14.2.11, we have that

$$
\left|\psi^{(k)}(0)\right| \leq k!(512)^{k}
$$

Hence, we have that

$$
\left\|\gamma^{(k)}(0)\right\|_{2} \leq \frac{k!(512)^{k}}{\alpha^{k}}=k!(512)^{k}\left\|A \gamma^{\prime}(0)\right\|_{4}^{k}
$$

for all $k \geq 2$.

### 14.5.9.3 Parallel Transport

Motivated from the equation for parallel transport under Euclidean coordinate 14.20 , we define the following auxiliary function

$$
\begin{equation*}
F(t)=\left(A_{\gamma(t)}^{T} A_{\gamma(t)}\right)^{-1} A_{\gamma(t)}^{T} S_{\gamma^{\prime}(t)} A_{\gamma(t)} \tag{14.25}
\end{equation*}
$$

The derivative bounds on parallel transport rely on the smoothness of this auxiliary function.

Lemma 14.5.21. Given a geodesic $\gamma(t)$. Under the normalization that $A^{T} A=I, S_{\gamma(0)}=I$, we have that

$$
F(t) \leq_{0} \frac{c}{2(1-2 c t)^{2}-(1-c t)^{2}} \frac{1-c t}{1-2 c t}
$$

where $F$ is defined in (14.25) and $c=512\left\|A \gamma^{\prime}(0)\right\|_{4}$.
Proof. Lemma 14.5.20 shows that $\left\|\gamma^{(k)}(0)\right\|_{2} \leq k!c^{k}$ for all $k \geq 2$. Therefore, we have that

$$
S_{\gamma(t)} \leq_{0} 1+t\left\|A \gamma^{\prime}(0)\right\|_{\infty}+\sum_{k \geq 2}(c t)^{k} \leq_{0} \frac{1}{1-c t} .
$$

Using Lemma 14.7.3, we have that

$$
A_{\gamma(t)} \leq_{0} \frac{1}{2-\frac{1}{1-c t}}=\frac{1-c t}{1-2 c t}
$$

and hence

$$
\left(A_{\gamma(t)}^{T} A_{\gamma(t)}\right)^{-1} A_{\gamma(t)}^{T} \leq_{0} \frac{(1-2 c t)(1-c t)}{2(1-2 c t)^{2}-(1-c t)^{2}} .
$$

Now, we note that

$$
\begin{equation*}
\operatorname{Diag}\left(A \gamma^{\prime}(t)\right) \leq_{0}\left\|A \gamma^{\prime}(0)\right\|_{\infty}+\sum_{k \geq 1} \frac{(k+1)!c^{k+1} t^{k}}{k!} \leq_{0} \frac{c}{(1-c t)^{2}} \tag{14.26}
\end{equation*}
$$

and hence

$$
S_{\gamma^{\prime}}=S_{\gamma}^{-1} \operatorname{Diag}\left(A \gamma^{\prime}(t)\right) \leq_{0} \frac{1-c t}{1-2 c t} \frac{c}{(1-c t)^{2}}=\frac{c}{(1-2 c t)(1-c t)} .
$$

This gives the result.

Now, we prove parallel transport has large radius of convergence.
Lemma 14.5.22. Given a geodesic with $\gamma(0)=x$. Let $v(t)$ be the parallel transport of a unit vector along $\gamma(t)$. Under the normalization that $A^{T} A=I, S_{\gamma(0)}=I$, we have that

$$
\left\|v^{(k)}(0)\right\|_{2} \leq k!(16 c)^{k}
$$

for all $k \geq 1$ where $c=512\left\|A \gamma^{\prime}(0)\right\|_{4}$.
Proof. From 14.20, we have that

$$
\frac{d}{d t} v(t)=\left(A_{\gamma(t)}^{T} A_{\gamma(t)}\right)^{-1} A_{\gamma(t)}^{T} S_{\gamma^{\prime}(t)} A_{\gamma(t)} v(t)
$$

Let $u(t)=v(\alpha t)$, then we have that

$$
\begin{aligned}
u^{\prime}(t) & =\alpha v^{\prime}(\alpha t) \\
& =\alpha\left(A_{\gamma(\alpha t)}^{T} A_{\gamma(\alpha t)}\right)^{-1} A_{\gamma(\alpha t)}^{T} S_{\gamma^{\prime}(\alpha t)} A_{\gamma(\alpha t)} u(t)
\end{aligned}
$$

Let $F(x, t)=\alpha\left(A_{\gamma(\alpha t)}^{T} A_{\gamma(\alpha t)}\right)^{-1} A_{\gamma(\alpha t)}^{T} S_{\gamma^{\prime}(\alpha t)} A_{\gamma(\alpha t)} x$. Then, Lemma 14.5.21 shows that

$$
F(x, t) \leq_{(X(0), 0)} \frac{\alpha c}{2(1-2 \alpha c t)^{2}-(1-\alpha c t)^{2}} \frac{1-\alpha c t}{1-2 \alpha c t}(1+t) .
$$

Setting $\alpha=\frac{1}{8 c}$, we have that

$$
F(x, t) \leq_{(X(0), 0)} \frac{\frac{1}{8}}{2\left(1-\frac{t}{4}\right)^{2}-\left(1-\frac{t}{8}\right)^{2}} \frac{1-\frac{t}{8}}{1-\frac{t}{4}}(1+t) \leq_{0} \frac{1}{1-t}
$$

Lemma 14.7.4 shows that

$$
\left\|u^{(k)}(0)\right\|_{2} \leq \psi^{(k)}(0)
$$

for all $k \geq 1$ where $\psi(t)$ is the solution of

$$
\psi^{\prime}(t)=\frac{1}{1-\psi(t)} \text { with } \psi(0)=0
$$

Solving it, we get that

$$
\psi(t)=1-\sqrt{1-2 t}
$$

By Theorem 14.2.11, we have that for any $0 \leq t \leq \frac{1}{2}$, we have that

$$
\left\|u^{(k)}(0)\right\|_{2} \leq\left|\psi^{(k)}(0)\right| \leq k!2^{k}
$$

for all $k \geq 1$. For $k \geq 1$, we have that

$$
\left\|v^{(k)}(0)\right\|_{2} \leq k!(16 c)^{k}
$$

### 14.5.9.4 Jacobi field

Motivated from the equation for Jacobi field under orthogonal frame basis (Lemma 14.5.5), we define the following auxiliary function

$$
\begin{equation*}
F(t)=X^{-1}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1}\left(A_{\gamma}^{T} S_{\gamma^{\prime}} P_{\gamma} S_{\gamma^{\prime}} A_{\gamma}-A_{\gamma}^{T} \operatorname{diag}\left(P_{\gamma} s_{\gamma^{\prime}}^{2}\right) A_{\gamma}\right) X . \tag{14.27}
\end{equation*}
$$

The derivative bounds on Jacobi field rely on the smoothness of this auxiliary function.
Lemma 14.5.23. Given a geodesic $\gamma(t)$ and an orthogonal frame $\left\{X_{i}\right\}_{i=1}^{n}$. Under the normalization that $A^{T} A=I, S_{\gamma(0)}=I$, we have that

$$
F(t) \leq_{0} \frac{12 c^{2}}{1-64 c t}
$$

where $F$ is defined in (14.27) and $c=512\left\|A \gamma^{\prime}(0)\right\|_{4}$.
Proof. We first bound the derivatives of $s_{\gamma^{\prime}}^{2}$. Using $\left\|\gamma^{(k)}(0)\right\|_{2} \leq k!c^{k}$ for all $k \geq 2$ (Lemma 14.5.20), we have that

$$
\gamma^{\prime}(t)-\gamma^{\prime}(0) \leq_{0} \frac{d}{d t} \frac{1}{1-c t}=\frac{c}{(1-c t)^{2}} .
$$

Using $A^{T} A=I$ and $\operatorname{Diag}\left(A \gamma^{\prime}(t)\right) \leq_{0} \frac{c}{(1-c t)^{2}}(14.26)$, we have that

$$
\begin{equation*}
\operatorname{Diag}\left(A \gamma^{\prime}(t)\right) A\left(\gamma^{\prime}(t)-\gamma^{\prime}(0)\right) \leq_{0}\left(\frac{c}{(1-c t)^{2}}\right)^{2} . \tag{14.28}
\end{equation*}
$$

Next, we note that

$$
\operatorname{Diag}\left(A \gamma^{\prime}(t)\right) A \gamma^{\prime}(0)=\operatorname{Diag}\left(A \gamma^{\prime}(0)\right) A\left(\gamma^{\prime}(t)-\gamma^{\prime}(0)\right)+\operatorname{Diag}\left(A \gamma^{\prime}(0)\right)^{2} .
$$

and hence

$$
\begin{equation*}
\operatorname{Diag}\left(A \gamma^{\prime}(t)\right) A \gamma^{\prime}(0) \leq_{0} \frac{c^{2}}{(1-c t)^{2}}+c^{2} \tag{14.29}
\end{equation*}
$$

Combining (14.28) and (14.28), we get

$$
\operatorname{Diag}\left(A \gamma^{\prime}(t)\right) A\left(\gamma^{\prime}(t)\right) \leq_{0}\left(\frac{c}{(1-c t)^{2}}\right)^{2}+\frac{c^{2}}{(1-c t)^{2}}+c^{2} \leq_{0} \frac{3 c^{2}}{(1-c t)^{4}}
$$

In the proof of Lemma 14.5.21, we showed that $S_{\gamma} \leq_{0} \frac{1}{1-c t}$ and hence

$$
\begin{aligned}
s_{\gamma^{\prime}}^{2} & =S_{\gamma}^{-2} \operatorname{Diag}\left(A \gamma^{\prime}(t)\right) A\left(\gamma^{\prime}(t)\right) . \\
& \leq_{0}\left(\frac{1}{2-\frac{1}{1-c t}}\right)^{2} \frac{3 c^{2}}{(1-c t)^{4}} \\
& =\frac{3 c^{2}}{(1-c t)^{2}(2(1-c t)-1)^{2}} .
\end{aligned}
$$

In the proof of Lemma 14.5.21, we showed that $A_{\gamma} \leq_{0} \frac{1-c t}{1-2 c t}$ and $\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} \leq_{0} \frac{(1-2 c t)^{2}}{2(1-2 c t)^{2}-(1-c t)^{2}}$. Therefore, we have that

$$
P_{\gamma}=A_{\gamma}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} \leq_{0} \frac{(1-c t)^{2}}{2(1-2 c t)^{2}-(1-c t)^{2}}
$$

Hence, we have

$$
\begin{aligned}
P_{\gamma} s_{\gamma^{\prime}}^{2} & \leq_{0} \frac{(1-c t)^{2}}{2(1-2 c t)^{2}-(1-c t)^{2}} \frac{3 c^{2}}{(1-c t)^{2}(1-2 c t)^{2}} \\
& =\frac{3 c^{2}}{(1-2 c t)^{2}\left(2(1-2 c t)^{2}-(1-c t)^{2}\right)}
\end{aligned}
$$

Let $Y=\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1}\left(A_{\gamma}^{T} S_{\gamma^{\prime}} P_{\gamma} S_{\gamma^{\prime}} A_{\gamma}-A_{\gamma}^{T} \operatorname{diag}\left(P_{\gamma} s_{\gamma^{\prime}}^{2}\right) A_{\gamma}\right)$. By a similar proof, we have that

$$
\begin{aligned}
Y_{0} & \leq_{0} \frac{(1-2 c t)^{2}}{2(1-2 c t)^{2}-(1-c t)^{2}}\left(\frac{1-c t}{1-2 c t}\right)^{2} \\
& \left(\left(\frac{c}{(1-2 c t)(1-c t)}\right)^{2} \frac{(1-c t)^{2}}{2(1-2 c t)^{2}-(1-c t)^{2}}+\frac{3 c^{2}}{(1-2 c t)^{2}\left(2(1-2 c t)^{2}-(1-c t)^{2}\right)}\right) \\
& =\frac{4 c^{2}}{\left(2(1-2 c t)^{2}-(1-c t)^{2}\right)^{2}}\left(\frac{1-c t}{1-2 c t}\right)^{2}
\end{aligned}
$$

Next, we let $Z(t)=X(t) v$ for some unit vector $v$. Since $X(t)$ is a parallel transport of $X(0), Z(t)$ is a parallel transport of $X(0) v$ and hence Lemma 14.5.22 shows that $\left\|Z^{(k)}(t)\right\|_{2} \leq(16 c)^{k}$ for all $i$. For any $k$, there is unit vector $v_{k}$ such that $\left\|\frac{d}{d t^{k}} X(0)\right\|_{2}=\left\|\frac{d}{d t^{k}} X(0) v_{k}\right\|_{2}$. Hence, we have that

$$
\left\|\frac{d}{d t^{k}} X(0)\right\|_{2} \leq(16 c)^{k}
$$

Therefore, we have that $X \leq_{0} \frac{1}{1-16 c t}$. Since $X(0)=I$, we have that $X^{-1} \leq_{0} \frac{1}{2-\frac{1}{1-16 c t}}=\frac{1-16 c t}{1-32 c t}$. Thus,
we have that

$$
\begin{aligned}
F(t) & \leq_{0} \frac{4 c^{2}}{\left(2(1-2 c t)^{2}-(1-c t)^{2}\right)^{2}}\left(\frac{1-c t}{1-2 c t}\right)^{2} \frac{1}{1-32 c t} \\
& \leq_{0} \frac{12 c^{2}}{1-64 c t} .
\end{aligned}
$$

Now, we prove Jacobi field has large radius of convergence.

Lemma 14.5.24. Given a geodesic $\gamma(t)$ and an orthogonal frame $\left\{X_{i}\right\}_{i=1}^{n}$ along $\gamma$. Let $V(t)$ be a Jacobi field along $\gamma(t)$ with $V(0)=0$ and let $U(t)$ be the Jacobi field under the $X_{i}$ coordinates, namely, $V(t)=X(t) U(t)$. Under the normalization that $A^{T} A=I, S_{\gamma(0)}=I$, we have that

$$
\left\|U^{(k)}(0)\right\|_{2} \leq k!(256 c)^{k}
$$

for all $k \geq 2$ where $c=512\left\|A \gamma^{\prime}(0)\right\|_{4}$.

Proof. Note that $\frac{d^{2} U(t)}{d t^{2}}+F(t) U(t)=0$ where $F(t)$ defined in (14.27). Since the equation is linear, we can rescale $U$ and assume that $\left\|U^{\prime}(0)\right\|_{2}=\frac{1}{\alpha}$.

Let $\bar{F}(U(t), t)=\alpha^{2} F(\alpha t)\left(U(t)+\alpha t U^{\prime}(0)\right)$ with $\alpha=\frac{1}{64 c}$. Using $\alpha^{2} F(\alpha t) \leq_{0} \frac{12 c^{2} \alpha^{2}}{1-64 c \alpha t}$ (Lemma 14.5.23), $U \leq_{U=0} t$ and $\alpha t U^{\prime}(0) \leq_{t=0}\left\|\alpha U^{\prime}(0)\right\|_{2} t=t$, we have that

$$
\begin{gathered}
\bar{F} \quad \leq_{(0,0)} \frac{24 c^{2} \alpha^{2} t}{1-64 c \alpha t} \\
\\
\leq_{0} \quad \frac{t}{1-t} .
\end{gathered}
$$

Lemma 14.7 .5 shows that that

$$
\left\|U^{(k)}(t)\right\|_{2} \leq \frac{\psi^{(k)}\left(\alpha^{-1} t\right)}{\alpha^{k}}
$$

for all $k \geq 2$ where $\psi(t)$ is the solution of

$$
\psi^{\prime}(t)=\frac{2}{1-\psi(t)} \text { with } \psi(0)=0
$$

Solving it, we get that

$$
\psi(t)=1-\sqrt{1-4 t}
$$

By Theorem 14.2.11, we have that for any $0 \leq t \leq \frac{1}{4}$, we have that

$$
\left|\psi^{(k)}(0)\right| \leq k!(4)^{k} .
$$

Hence, we have that

$$
\left\|U^{(k)}(0)\right\|_{2} \leq \frac{k!4^{k}}{\alpha^{k}} \leq k!(256 c)^{k}
$$

for all $k \geq 2$.

## ■ 14.5.10 Implementation

### 14.5.10.1 Computing Geodesic Equation

To apply Theorem 14.6.4, we define

$$
\begin{equation*}
F(u, s)=A\left(A^{T} S^{-2} A\right)^{-1} A^{T} S^{-1} u^{2} \tag{14.30}
\end{equation*}
$$

The following lemma bounds the Lipschitz constant of $F$.
Lemma 14.5.25. Assuming $\frac{1}{2} \leq s_{i} \leq 2$ for all $i$. Then, we have

$$
\left\|D F(u, s)\left[d_{u}, d_{s}\right]\right\|_{2}^{2} \leq 48\|u\|_{4}^{2}\left\|d_{u}\right\|_{4}^{2}+240\|u\|_{4}^{4}\left\|d_{s}\right\|_{\infty}^{2}
$$

Proof. Let $A_{s}=S^{-1} A$ and $S_{d}=S^{-1} \operatorname{Diag}\left(d_{s}\right)$. Then, we have $F(u, s)=A\left(A_{s}^{T} A_{s}\right)^{-1} A_{s}^{T} u^{2}$. Hence, we have that

$$
\begin{aligned}
D F(u, s)\left[d_{u}, d_{s}\right]= & 2 A\left(A_{s}^{T} A_{s}\right)^{-1} A_{s}^{T} S_{d} A_{s}\left(A_{s}^{T} A_{s}\right)^{-1} A_{s}^{T} u^{2} \\
& -A\left(A_{s}^{T} A_{s}\right)^{-1} A_{s}^{T} S_{d} u^{2} \\
& +2 A\left(A_{s}^{T} A_{s}\right)^{-1} A_{s}^{T} U d_{u}
\end{aligned}
$$

Let $P=A_{s}\left(A_{s}^{T} A_{s}\right)^{-1} A_{s}^{T}$, then, we have that

$$
\begin{aligned}
\left\|D F(u, s)\left[d_{u}, d_{s}\right]\right\|_{2}^{2} \leq & 12\left(u^{2}\right)^{T} P S_{d} P S^{2} P S_{d} P u^{2} \\
& +3\left(u^{2}\right)^{T} S_{d} P S^{2} P S_{d} u^{2} \\
& +12 d_{u}^{T} U P S^{2} P U d_{u} .
\end{aligned}
$$

Using that $P \preceq I$, we have that

$$
\begin{aligned}
& \left\|D F(u, s)\left[d_{u}, d_{s}\right]\right\|_{2}^{2} \\
\leq & 12\|S\|_{\infty}^{2}\left(u^{2}\right)^{T} P S_{d}^{2} P u^{2}+3\|S\|_{\infty}^{2}\left(u^{2}\right)^{T} S_{d}^{2} u^{2}+12\|S\|_{\infty}^{2} \sum_{i}\left(d_{u}\right)_{i}^{2} u_{i}^{2} \\
\leq & 15\|S\|_{\infty}^{2}\left\|S_{d}\right\|_{\infty}^{2}\|u\|_{4}^{4}+12\|S\|_{\infty}^{2}\|u\|_{4}^{2}\left\|d_{u}\right\|_{4}^{2} \\
\leq & 15\|S\|_{\infty}^{2}\left\|S^{-1}\right\|_{\infty}^{2}\left\|d_{s}\right\|_{\infty}^{2}\|u\|_{4}^{4}+12\|S\|_{\infty}^{2}\|u\|_{4}^{2}\left\|d_{u}\right\|_{4}^{2} .
\end{aligned}
$$

Now, we use that $\frac{1}{2} \leq s_{i} \leq 2$ and get

$$
\left\|D F(u, s)\left[d_{u}, d_{s}\right]\right\|_{2}^{2} \leq 48\|u\|_{4}^{2}\left\|d_{u}\right\|_{4}^{2}+240\left\|d_{s}\right\|_{\infty}^{2}\|u\|_{4}^{4} .
$$

Now, we can apply the collocation method to obtain a good approximation of the solution $u^{*}$.
Lemma 14.5.26. Let $\gamma$ be a random geodesic generated by the geodesic walk with step size $h \leq$ $\frac{1}{10^{7} \sqrt{n}}$. In time $O\left(m n^{\omega-1} \log ^{2}(1 / \varepsilon)\right)$, we can find $\bar{\gamma}$ such that $\max _{0 \leq t \leq \ell}\|\gamma(t)-\bar{\gamma}(t)\|_{\infty} \leq \varepsilon$ and $\max _{0 \leq t \leq \ell}\left\|\gamma^{\prime}(t)-\bar{\gamma}^{\prime}\right\|_{\infty} \leq \varepsilon$ with probability at least $1-O(\exp (-\sqrt{n} / 100))$. Furthermore, $\bar{\gamma}$ is a $O(\log (1 / \varepsilon))$ degree polynomial.

Proof. Let $s(t)=A \gamma(t)-b$. By rotating the space and rescaling the rows of $A$, we assume that
$s(0)_{i}=1$ for all $i$ and $A^{T} A=I$. We define $F$ as (14.30). Then, we have that

$$
\begin{aligned}
s^{\prime \prime}(t) & =F\left(s^{\prime}, s\right) \\
s^{\prime}(0) & =A \gamma^{\prime}(0) \\
s(0) & =1
\end{aligned}
$$

We let $\alpha=40 \ell$ and

$$
\begin{aligned}
K & \stackrel{\text { def }}{=} \alpha\left\|F\left(A \gamma^{\prime}(0), 1\right)\right\|_{4}+\left\|A \gamma^{\prime}(0)\right\|_{4} \\
& \leq \alpha\left\|A\left(A^{T} A\right)^{-1} A^{T}\left(A \gamma^{\prime}(0)\right)^{2}\right\|_{2}+\left\|A \gamma^{\prime}(0)\right\|_{4} \\
& \leq \alpha\left\|A \gamma^{\prime}(0)\right\|_{4}^{2}+\left\|A \gamma^{\prime}(0)\right\|_{4} .
\end{aligned}
$$

Using $\left\|A_{x} \gamma^{\prime}\right\|_{4} \leq 2 n^{-1 / 4}$ (Lemma 14.5.10), we have that

$$
\begin{aligned}
K & \leq 2\left\|A \gamma^{\prime}(0)\right\|_{4} \leq 4 n^{-1 / 4}, \\
\alpha K & \leq 160 \ell n^{-1 / 4} \leq \frac{1}{2}
\end{aligned}
$$

For any $\left\|u-s^{\prime}(0)\right\|_{4} \leq K \leq 4 n^{-1 / 4},\|s-1\|_{4} \leq \alpha K \leq \frac{1}{2}$, Lemma 14.5.25 shows that

$$
\begin{aligned}
\left\|D F(u, s)\left[d_{u}, d_{s}\right]\right\|_{4} & \leq\left\|D F(u, s)\left[d_{u}, d_{s}\right]\right\|_{2} \\
& \leq 8\|u\|_{4}\left\|d_{u}\right\|_{4}+16\|u\|_{\|_{2}}^{2}\left\|d_{s}\right\|_{\infty} \\
& \leq 48 n^{-1 / 4}\left\|d_{u}\right\|_{4}+576 n^{-1 / 2}\left\|d_{s}\right\|_{4} .
\end{aligned}
$$

Therefore, for any $\left\|u_{1}-s^{\prime}(0)\right\|_{4} \leq K,\left\|s_{1}-1\right\|_{4} \leq \alpha K,\left\|u_{2}-s^{\prime}(0)\right\|_{4} \leq K,\left\|s_{2}-1\right\|_{4} \leq \alpha K$, we have that

$$
\begin{aligned}
\left\|F\left(u_{1}, s_{1}\right)-F\left(u_{2}, s_{2}\right)\right\|_{4} & \leq 48 n^{-1 / 4}\left\|u_{1}-u_{2}\right\|_{4}+576 n^{-1 / 2}\left\|s_{1}-s_{2}\right\|_{4} \\
& \leq \frac{1}{\alpha}\left\|u_{1}-u_{2}\right\|_{4}+\frac{1}{\alpha^{2}}\left\|s_{1}-s_{2}\right\|_{4} .
\end{aligned}
$$

Since $\gamma$ is analytic and $\left\|\gamma^{(k)}(0)\right\|_{2}=O\left(k!n^{-k / 4}\right)$ (Lemma 14.5.22), $\gamma(t)$ is $\varepsilon$ close to a polynomial with degree $O(\log (1 / \varepsilon))$ for $0 \leq t \leq o\left(n^{1 / 4}\right)$. Hence, we can apply Theorem 14.6.6 and find $\bar{\gamma}$ such that $\|\gamma-\bar{\gamma}\|_{4} \leq \varepsilon$ and $\left\|\gamma^{\prime}-\bar{\gamma}^{\prime}\right\|_{4} \leq \varepsilon$ in $O\left(n \log ^{3}(K / \varepsilon)\right)$ time plus $O\left(\log ^{2}(K / \varepsilon)\right)$ evaluations of $F$. Note that each evaluation of $F$ involves solving a linear system and hence it takes $O\left(m n^{\omega-1}\right)$. Therefore, the total running time is $O\left(m n^{\omega-1} \log ^{2}(1 / \varepsilon)\right)$.

### 14.5.10.2 Computing Parallel Transport

Lemma 14.5.27. Given $\gamma$ be a random geodesic generated by the geodesic walk with step size $h \leq$ $\frac{1}{10^{7} \sqrt{n}}$ and an unit vector $v$. Let $v(t)$ be the parallel transport of a unit vector along $\gamma(t)$. In time $O\left(m n^{\omega-1} \log ^{2}(1 / \varepsilon)\right)$, we can find $\bar{v}$ such that $\max _{0 \leq t \leq \ell}\|v(t)-\bar{v}(t)\|_{\infty} \leq \varepsilon$ with probability at least $1-O(\exp (-\sqrt{n} / 100))$. Furthermore, $\bar{v}$ is a $O(\log (1 / \varepsilon))$ degree polynomial.

Similarly, given a basis $\left\{v_{i}\right\}_{i=1}^{n}$, in time $O\left(m n^{\omega-1} \log ^{2}(1 / \varepsilon)\right)$, we can find an approximate parallel transport $\overline{v_{i}}(t)$ of $\left\{v_{i}\right\}_{i=1}^{n}$ along $\gamma(t)$ with probability at least $1-O(\exp (-\sqrt{n} / 100))$.

Proof. Recall that the equation for parallel transport (14.20) is given by

$$
\frac{d}{d t} v(t)=\left(A_{\gamma(t)}^{T} A_{\gamma(t)}\right)^{-1} A_{\gamma(t)}^{T} S_{\gamma^{\prime}(t)} A_{\gamma(t)} v .
$$

By rotating the space and rescaling the rows of $A$, we assume that $s(\gamma(0))_{i}=1$ for all $i$ and $A^{T} A=I$.

In the proof of Lemma 14.5.26, we know that $\frac{1}{2} \leq s(\gamma(t))_{i} \leq 2$ for all $0 \leq t \leq \ell$. For any unit vector $u$, we have

$$
\begin{aligned}
\left\|\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime}} A_{\gamma} u\right\|_{2} & \leq 2\left\|\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1 / 2} A_{\gamma}^{T} S_{\gamma^{\prime}} A_{\gamma} u\right\|_{2} \\
& \leq 2\left\|S_{\gamma^{\prime}}\right\|_{\infty}\left\|A_{\gamma} u\right\|_{2} \\
& \leq 4\left\|S_{\gamma^{\prime}}\right\|_{\infty} \leq 24 \sqrt{h}
\end{aligned}
$$

where we used Lemma 14.5 .10 in the last line. Using $h \leq \frac{1}{10^{7} \sqrt{n}}$, we have that

$$
\left\|\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1} A_{\gamma}^{T} S_{\gamma^{\prime}} A_{\gamma}\right\|_{2} \leq \frac{1}{20 \ell} .
$$

Since $v$ is analytic and $\left\|v^{(k)}(0)\right\|_{2}=O\left(k!n^{-k / 4}\right)$ (Lemma 14.5.22), $v(t)$ is $\varepsilon$ close to a polynomial with degree $O(\log (1 / \varepsilon))$ for $0 \leq t \leq o\left(n^{1 / 4}\right)$. Hence, we can apply Theorem 14.6.4 and find $\bar{v}$ such that $\|v-\bar{v}\|_{2} \leq \varepsilon$ in $O\left(n \log ^{3}(1 / \varepsilon)\right)$ time plus $O\left(\log ^{2}(1 / \varepsilon)\right)$ evaluations of $F$. Note that each evaluation of $F$ involves solving a linear system and hence it takes $O\left(m n^{\omega-1}\right)$. Therefore, the total running time is $O\left(m n^{\omega-1} \log ^{2}(1 / \varepsilon)\right)$.

For the last result, we note that each evaluation of $F$ becomes computing matrix inverse and performing matrix multiplication and they can be done in again $O\left(m n^{\omega-1}\right)$ time.

### 14.5.10.3 Computing Jacobi field

Lemma 14.5.28. Given $\gamma$ be a random geodesic generated by the geodesic walk with step size $h \leq$ $\frac{1}{10^{7} \sqrt{n}}$. Let $\left\{X_{i}\right\}_{i=1}^{n}$ be an orthogonal frame along $\gamma$. Let $v(t)$ be a Jacobi field along $\gamma(t)$ with $v(0)=0$ and let $u(t)$ be the Jacobi field under the $X_{i}$ coordinates, namely, $v(t)=X(t) u(t)$. In time $O\left(m n^{\omega-1} \log ^{2}(1 / \varepsilon)\right)$, we can find $\bar{u}$ such that $\max _{0 \leq t \leq \ell}\|u(t)-\bar{u}(t)\|_{\infty} \leq \varepsilon$ with probability at least $1-O(\exp (-\sqrt{n} / 100))$. Furthermore, $\bar{u}$ is a $O(\log (1 / \varepsilon))$ degree polynomial.

Proof. Recall that the equation for Jacobi field 14.5 .5 is given by

$$
\frac{d^{2} u}{d t^{2}}+X^{-1}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1}\left(A_{\gamma}^{T} S_{\gamma^{\prime}} P_{\gamma} S_{\gamma^{\prime}} A_{\gamma}-A_{\gamma}^{T} \operatorname{Diag}\left(P_{\gamma} s_{\gamma^{\prime}}^{2}\right) A_{\gamma}\right) X u=0 .
$$

By rotating the space and rescaling the rows of $A$, we assume that $s(\gamma(0))_{i}=1$ for all $i$ and $A^{T} A=I$. In the proof of Lemma 14.5.26, we know that $\frac{1}{2} \leq s(\gamma(t))_{i} \leq 2$ for all $0 \leq t \leq \ell$. Since $X$ is an orthogonal frame, we have $X X^{T}=X X^{T}=I$. Hence, for any unit vector $v$, we have

$$
\begin{aligned}
& \left\|X^{-1}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1}\left(A_{\gamma}^{T} S_{\gamma^{\prime}} P_{\gamma} S_{\gamma^{\prime}} A_{\gamma}-A_{\gamma}^{T} \operatorname{Diag}\left(P_{\gamma} s_{\gamma^{\prime}}^{2}\right) A_{\gamma}\right) X v\right\|_{2} \\
\leq & 2\left\|\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1 / 2}\left(A_{\gamma}^{T} S_{\gamma^{\prime}} P_{\gamma} S_{\gamma^{\prime}} A_{\gamma}-A_{\gamma}^{T} \operatorname{Diag}\left(P_{\gamma} s_{\gamma^{\prime}}^{\prime}\right) A_{\gamma}\right) X v\right\|_{2} \\
\leq & 2\left\|S_{\gamma^{\prime}} P_{\gamma} S_{\gamma^{\prime}} A_{\gamma} v\right\|_{2}+2\left\|\operatorname{Diag}\left(P_{\gamma} s_{\gamma^{\prime}}^{2}\right) A_{\gamma} X v\right\|_{2} .
\end{aligned}
$$

Using $\left\|S_{\gamma^{\prime}}\right\|_{2} \leq 6 \sqrt{h}$ and $\left\|\operatorname{Diag}\left(P_{\gamma} s_{\gamma^{\prime}}^{2}\right)\right\|_{2} \leq\left\|P_{\gamma} s_{\gamma^{\prime}}^{2}\right\|_{2} \leq\left\|s_{\gamma^{\prime}}\right\|_{4}^{2} \leq 4 n^{-1 / 2}$ (Lemma 14.5.10), we have that

$$
\begin{aligned}
& \left\|X^{-1}\left(A_{\gamma}^{T} A_{\gamma}\right)^{-1}\left(A_{\gamma}^{T} S_{\gamma^{\prime}} P_{\gamma} S_{\gamma^{\prime}} A_{\gamma}-A_{\gamma}^{T} \operatorname{Diag}\left(P_{\gamma} s_{\gamma^{\prime}}^{2}\right) A_{\gamma}\right) X v\right\|_{2} \\
\leq & 12 \sqrt{h}\left\|P_{\gamma} S_{\gamma^{\prime}} A_{\gamma} v\right\|_{2}+8 n^{-1 / 2}\left\|A_{\gamma} X v\right\|_{2} \\
\leq & 72 h\left\|A_{\gamma} v\right\|_{2}+16 n^{-1 / 2}\|X v\|_{2} \\
\leq & 144 h+16 n^{-1 / 2} \leq \frac{1}{(40 \ell)^{2}}
\end{aligned}
$$

where we used $h \leq \frac{1}{10^{7} \sqrt{n}}$ in the last line.
Since $u$ is analytic and $\left\|u^{(k)}(0)\right\|_{2}=O\left(k!n^{-k / 4}\right)$ (Lemma 14.5.24), $u(t)$ is $\varepsilon$ close to a polynomial with degree $O(\log (1 / \varepsilon))$ for $0 \leq t \leq o\left(n^{1 / 4}\right)$. Hence, we can apply Theorem 14.6.4 and find $\bar{u}$ such that $\|u-\bar{u}\|_{2} \leq \varepsilon$ in $O\left(n \log ^{3}(1 / \varepsilon)\right)$ time plus $O\left(\log ^{2}(1 / \varepsilon)\right)$ evaluations of $F$. Note that each evaluation of $F$ involves computing matrix inversions and matrix multiplications and hence it takes $O\left(\mathrm{mn}^{\omega-1}\right)$. Therefore, the total running time is $O\left(m n^{\omega-1} \log ^{2}(1 / \varepsilon)\right)$.

### 14.5.10.4 Computing Geodesic Walk

Proof of Theorem 14.5.2. To implement the geodesic walk, we use Lemma 14.5.26 to compute the geodesic, Lemma 14.5.27 to compute an orthogonal frame along the geodesic and Lemma 14.5.28 to compute the Jacobi field along. Using the Jacobi field, we can use (14.9) and Lemma 14.4.12 to compute the probability from $x$ to $y$ and the probability from $y$ to $x$. Using these probabilities, we can implement the rejection sampling. It suffices to compute the geodesic and the probability up to $1 / n^{O(1)}$ accuracy and hence these operations can be done in time $O\left(m n^{\omega-1} \log ^{2}(n)\right)$.

Note that we only use randomness to prove that $\left\|A_{\gamma(0)} u\right\|_{2} \leq \frac{3}{2}$ and $\left\|A_{\gamma} u\right\|_{4} \leq \frac{1.55}{n^{1 / 4}}$ (Lemma 14.5.10) and they can be checked. When we condition our walk to that, we only change the distribution by exponential small amount. Hence, this result is stated without mentioning the success probability.

### 14.6 Collocation Method for ODE

In this section, we study a collocation method for solving ordinary differential equation (ODE) and show how to solve a nice enough ODE in nearly constant number of iterations without explicitly computing higher derivatives.

## - 14.6.1 First Order ODE

We first consider the following first order ODE

$$
\begin{align*}
\frac{d}{d t} u(t) & =F(u(t), t), \text { for } 0 \leq t \leq \ell  \tag{14.31}\\
u(0) & =v
\end{align*}
$$

where $F: \mathbb{R}^{n+1} \rightarrow \mathbb{R}^{n}$ and $u(t) \in \mathbb{R}^{n}$. The idea of collocation methods is to find a degree $d$ polynomial $p$ such that

$$
\begin{align*}
\frac{d}{d t} p(t) & =F(p(t), t), \text { for } t=c_{1}, c_{2}, \cdots, c_{d}  \tag{14.32}\\
p(0) & =v
\end{align*}
$$

where $c_{1}, c_{2}, \cdots, c_{d}$ are carefully chosen distinct points on $[0, \ell]$. Here, we call $p: \mathbb{R} \rightarrow \mathbb{R}^{n}$ is a degree $d$ polynomial if $p(t)=\left[p_{1}(t) ; p_{2}(t) ; \cdots ; p_{n}(t)\right]$ and each $p_{i}(t)$ is an univariate polynomial with degree at most $d$. The first part of the proof shows the existence of a solution for the systems (14.32). To describe the algorithm, it is easier to consider an equivalent integral equation.

Lemma 14.6.1. Given distinct points $c_{1}, c_{2}, \cdots, c_{d} \in \mathbb{R}$ and $F: \mathbb{R}^{n+1} \rightarrow \mathbb{R}^{n}$, consider the nonlinear
$\operatorname{map} T: \mathbb{R}^{d \times n} \rightarrow \mathbb{R}^{d \times n}$ defined by

$$
T(\zeta)_{(i, k)}=\int_{0}^{c_{i}} \sum_{j=1}^{d} F\left(\zeta_{j}, c_{j}\right)_{k} \phi_{j}(s) d s \text { for } i \in[d], k \in[n]
$$

where $\phi_{i}(s)=\prod_{j \neq i} \frac{s-c_{j}}{c_{i}-c_{j}}$ are the Lagrange basis polynomials. Given any $\zeta \in \mathbb{R}^{d \times n}$ such that

$$
\begin{equation*}
\zeta_{i}=v+T(\zeta)_{i}, \text { for } i \in[d] \tag{14.33}
\end{equation*}
$$

the polynomial

$$
p(t)=v+\int_{0}^{t} \sum_{j=1}^{d} F\left(\zeta_{j}, c_{j}\right) \phi_{j}(s) d s
$$

is a solution of the system (14.32).
Proof. Define the polynomials $\phi_{i}(s)=\prod_{j \neq i} \frac{s-c_{j}}{c_{i}-c_{j}}$. Note that $\phi_{i}\left(c_{j}\right)=\delta_{i j}$. Therefore, we have

$$
\sum_{j=1}^{d} \alpha_{j} \phi_{j}\left(c_{i}\right)=\alpha_{i}
$$

Therefore, $p(0)=v$ and

$$
\frac{d}{d t} p\left(c_{i}\right)=\sum_{j=1}^{d} F\left(\zeta_{j}, c_{j}\right) \phi_{j}\left(c_{i}\right)=F\left(\zeta_{j}, c_{j}\right)
$$

Since $\zeta_{i}=v+T(\zeta)_{i}$, we have that

$$
\zeta_{i}=v+\int_{0}^{c_{i}} \sum_{j=1}^{d} F\left(\zeta_{j}, c_{j}\right) \phi_{j}(s) d s=p\left(c_{i}\right) .
$$

Hence, we have $\frac{d}{d t} p\left(c_{i}\right)=F\left(p\left(c_{i}\right), c_{i}\right)$. Therefore, $p$ is a solution to the system (14.32).
From Lemma 14.6.1, we see that it suffices to solve the system (14.33). We solve it by a simple fix point iteration shown in Algorithm 50.

```
Algorithm 50: CollocationMethod
    Input: An ordinary differential equation \(\frac{d}{d t} u(t)=F(u(t), t)\) for \(0 \leq t \leq \ell\) with initial condition
        \(u(0)=v\).
    Define \(T(\zeta)_{(i, k)}=\int_{0}^{c_{i}} \sum_{j=1}^{d} F\left(\zeta_{j}, c_{j}\right)_{k} \phi_{j}(s) d s\) for \(i \in[d], k \in[n]\) as defined in Lemma 14.6.1.
    \(c_{i}=\frac{\ell}{2}+\frac{\ell}{2} \cos \left(\frac{2 i-1}{2 d} \pi\right)\) for all \(i \in[d], K=40 \ell \max _{t \in[0, \ell]}\|F(v, t)\|_{p}, Z=\log _{2}(K / \varepsilon)\)
    \(\zeta_{i}^{(0)}=v+T(\bar{v})_{i}\) for all \(i \in[d]\) where \(\bar{v}=(v, v, \cdots, v) \in \mathbb{R}^{d \times n}\).
    for \(z=0, \cdots, Z\) do
        \(\zeta_{i}^{(z+1)}=v+T\left(\zeta^{(z)}\right)_{i}\) for \(i \in[d]\).
    end
    Output: \(p(t)=v+\int_{0}^{\ell} \sum_{j=1}^{d} F\left(\zeta_{j}^{(Z+1)}, c_{j}\right) \phi_{j}(s) d s\).
```

In the following lemma, we show that $T$ is a contraction mapping if $F$ is smooth enough and hence this algorithm converges linearly. We will use the following norm: $\|x\|_{\infty ; p}=\max _{i \in[d]}\left(\sum_{k \in[n]}\left|x_{(i, k)}\right|^{p}\right)^{1 / p}$.

Lemma 14.6.2. Given $x, y \in \mathbb{R}^{d \times n}$ with $d \geq 2$, suppose that $\left\|F\left(x_{i}, t\right)-F\left(y_{i}, t\right)\right\|_{p} \leq L\left\|x_{i}-y_{i}\right\|_{p}$ for all $0 \leq t \leq \ell$ and all $i \in[d]$. For $c_{k}=\frac{\ell}{2}+\frac{\ell}{2} \cos \left(\frac{2 k-1}{2 d} \pi\right)$ and $T$ defined Lemma 14.6.1, we have that

$$
\|T(x)-T(y)\|_{\infty ; p} \leq 10 \ell L\|x-y\|_{\infty ; p}
$$

Also, we have that

$$
\|T(x)\|_{\infty ; p} \leq 10 \ell \max _{i}\left\|F\left(x_{i}, c_{i}\right)\right\|_{\infty ; p}
$$

Proof. Using the definition of $T$ and $\|\cdot\|_{\infty ; p}$, we have

$$
\begin{aligned}
\|T(x)-T(y)\|_{\infty ; p} & =\max _{i}\left\|\int_{0}^{c_{i}} \sum_{j=1}^{d} F\left(x_{j}, c_{j}\right) \phi_{j}(s) d s-\int_{0}^{c_{i}} \sum_{j=1}^{d} F\left(y_{j}, c_{j}\right) \phi_{j}(s) d s\right\|_{p} \\
& \leq\left(\max _{i} \sum_{j=1}^{d}\left|\int_{0}^{c_{i}} \phi_{j}(s) d s\right|\right)\left(\max _{i}\left\|F\left(x_{i}, c_{i}\right)-F\left(y_{i}, c_{i}\right)\right\|_{p}\right) \\
& \leq L\|x-y\|_{\infty ; p}\left(\max _{i} \sum_{j=1}^{d}\left|\int_{0}^{c_{i}} \phi_{j}(s) d s\right|\right)
\end{aligned}
$$

To bound $\max _{i} \sum_{j=1}^{d}\left|\int_{0}^{c_{i}} \phi_{j}(s) d s\right|$, it is easier to work with the function $\psi_{j}(x) \stackrel{\text { def }}{=} \phi_{j}\left(\frac{2}{\ell} x-1\right)$. Note that $\psi_{k}$ is the Lagrange basis polynomials on the nodes $\left\{\cos \left(\frac{2 k-1}{2 d} \pi\right)\right\}_{k=1}^{d}$ and hence we have

$$
\psi_{k}(x)=\frac{(-1)^{k-1} \sqrt{1-x_{k}^{2}} \cos \left(d \cos ^{-1} x\right)}{d\left(x-x_{k}\right)}
$$

where $x_{k}=\cos \left(\frac{2 k-1}{2 d} \pi\right)$. Lemma 14.6.5 at the end of this section shows that $\left|\int_{-1}^{t} \psi_{k}(x) d x\right| \leq \frac{20}{d}$ for all $t$. Therefore, we have that $\max _{i} \sum_{j=1}^{d}\left|\int_{0}^{c_{i}} \phi_{j}(s) d s\right| \leq 10 \ell$. This gives the first inequality. The second inequality is similar.

In each iteration of the collation method, we need to compute $\int_{0}^{c_{i}} \sum \alpha_{j} \phi_{j}(s) d s$ for some $\alpha_{j}$. The following theorem shows that this can be done in $O(d \log (d / \varepsilon))$ time using the multipole method.
Theorem 14.6.3 ([77, Sec 5]). Let $\phi_{i}(s)$ be the Lagrange basis polynomials on the Chebyshev nodes on $[0, \ell]$, namely, $\phi_{i}(s)=\prod_{j \neq i} \frac{s-c_{j}}{c_{i}-c_{j}}$ with $c_{i}=\frac{\ell}{2}+\frac{\ell}{2} \cos \left(\frac{2 i-1}{2 d} \pi\right)$. Given a polynomial $p(s)=\sum_{j} \alpha_{j} \phi_{j}(s)$ and a point set $\left\{x_{1}, x_{2}, \cdots, x_{d}\right\}$, one can compute $t_{i}$ such that

$$
\left|t_{i}-\int_{0}^{x_{i}} p(s) d s\right| \leq \varepsilon \ell \sqrt{\sum_{i \neq j}\left(\alpha_{i}-\alpha_{j}\right)^{2}} \text { for } i \in[d]
$$

in time $O(d \log (d / \varepsilon))$.
Now we have everything to state our main result in this subsection.
Theorem 14.6.4. Let $u(t) \in \mathbb{R}^{n}$ be the solution of the $O D E$ (14.31). Suppose we are given $\varepsilon, \ell>0$, $d \geq 2$ and $1 \leq p \leq \infty$ such that

1. There is a degree d polynomial $q$ from $\mathbb{R}$ to $\mathbb{R}^{n}$ such that $q(0)=v$ and $\left\|\frac{d}{d t} u(t)-\frac{d}{d t} q(t)\right\|_{p} \leq \frac{\varepsilon}{\ell}$ for all $0 \leq t \leq \ell$.
2. We have that $\|F(x, t)-F(y, t)\|_{p} \leq \frac{1}{20 \ell}\|x-y\|_{p}$ for all $\|x-v\|_{p} \leq K$ and $\|y-v\|_{p} \leq K$ with $K=40 \ell \max _{t \in[0, \ell]}\|F(v, t)\|_{p}$.

Then, for any $t$ such that $0 \leq t \leq \ell$, Algorithm CollocationMethod outputs a degree d polynomial $p(t)$ such that $\|u(t)-p(t)\|_{p} \leq 43 \varepsilon$ in $O\left(n d \log ^{2}(d K / \varepsilon)\right)$ time and $O(d \log (K / \varepsilon))$ evaluations of $F$.

Proof. First, we estimate the initial error. Let $\zeta^{(\infty)}$ be the solution to (14.33), $\zeta^{(0)}=v+T(\bar{v})_{i}$ be the initial vector and $\bar{v}=(v, v, \cdots, v) \in \mathbb{R}^{d \times n}$. Then, we have that

$$
\begin{aligned}
\left\|\zeta^{(\infty)}-\zeta^{(0)}\right\|_{\infty ; p} & =\max _{i}\left\|\left(v+T\left(\zeta^{(\infty)}\right)_{i}\right)-\left(v+T(\bar{v})_{i}\right)\right\|_{p} \\
& =\left\|T\left(\zeta^{(\infty)}\right)-T(\bar{v})\right\|_{\infty ; p} \\
& \leq \frac{1}{2}\left\|\zeta^{(\infty)}-\bar{v}\right\|_{\infty ; p} \\
& \leq \frac{1}{2}\left(\left\|\zeta^{(\infty)}-\zeta^{(0)}\right\|_{\infty ; p}+\left\|\zeta^{(0)}-\bar{v}\right\|_{\infty ; p}\right)
\end{aligned}
$$

Therefore, we have that

$$
\begin{aligned}
\left\|\zeta^{(\infty)}-\zeta^{(0)}\right\|_{\infty ; p} & \leq\left\|\zeta^{(0)}-\bar{v}\right\|_{\infty ; p} \\
& =\|T((v, v, \cdots, v))\|_{\infty ; p} \\
& \leq 10 \ell \max _{i}\left\|F\left(v, c_{i}\right)\right\|_{p} \\
& \leq \frac{K}{4}
\end{aligned}
$$

Hence, we have that $\left\|\zeta^{(\infty)}-\bar{v}\right\|_{\infty ; p} \leq \frac{K}{2}$.
Using the assumption on $F$, Lemma 14.6.2 shows that $\|T(x)-T(y)\|_{\infty ; p} \leq \frac{1}{2}\|x-y\|_{\infty ; p}$ and hence

$$
\left\|\zeta^{(\infty)}-\zeta^{(k)}\right\|_{\infty ; p} \leq \frac{K}{2^{k+1}}
$$

Thus, it takes $\log _{2}(K / \varepsilon)$ iteration to get a point $\zeta^{(k)}$ with

$$
\begin{equation*}
\left\|\zeta^{(\infty)}-\zeta^{(k)}\right\|_{\infty ; p} \leq \varepsilon \tag{14.34}
\end{equation*}
$$

Also, this shows that $\left\|\zeta^{(\infty)}-\bar{v}\right\|_{\infty ; p} \leq K$. Hence, we only requires the assumption (2) to be satisfied in this region.

Now, we show that $\zeta^{(k)}$ is close to the solution by using the existence of $q$. By the assumption on $q$, we have that $\left\|\frac{d}{d t} u(t)-\frac{d}{d t} q(t)\right\|_{p} \leq \frac{\varepsilon}{\ell}$ and hence $\|u(t)-q(t)\|_{p} \leq \varepsilon$. Using the smoothness of $F$, we have that $\|F(u(t), t)-F(q(t), t)\|_{p} \leq \frac{\varepsilon}{\ell}$ for all $0 \leq t \leq \ell$. Therefore, we have that

$$
\begin{aligned}
\left\|\frac{d}{d t} q(t)-F(q(t), t)\right\|_{p} & \leq\left\|\frac{d}{d t} q(t)-\frac{d}{d t} u(t)\right\|_{p}+\|F(q(t), t)-F(u(t), t)\|_{p} \\
& \leq 2 \frac{\varepsilon}{\ell}
\end{aligned}
$$

for all $0 \leq t \leq 1$. Therefore, we have

$$
\begin{aligned}
q(t) & =v+\int_{0}^{t} \sum_{j=1}^{d} \frac{d}{d t} q\left(c_{j}\right) \phi_{j}(s) d s \\
& =v+\int_{0}^{t} \sum_{j=1}^{d}\left(F\left(q\left(c_{j}\right), c_{j}\right)+\delta_{j}\right) \phi_{j}(s) d s
\end{aligned}
$$

where $\left\|\delta_{i}\right\|_{p} \leq 2 \frac{\varepsilon}{\ell}$ for all $i \in[d]$. By Lemma 14.6 .2 , we have that

$$
\left\|q(t)-v-\int_{0}^{t} \sum_{j=1}^{d} F\left(q\left(c_{j}\right), c_{j}\right) \phi_{j}(s) d s\right\|_{p} \leq 20 \varepsilon
$$

Now, we compare $\bar{q}_{j} \stackrel{\text { def }}{=} q\left(c_{j}\right)$ with the approximate solution constructed by the fixpoint algorithm

$$
q^{(k)}(t)=v+\int_{0}^{t} \sum_{j=1}^{d} F\left(\zeta_{j}^{(k)}, c_{j}\right) \phi_{j}(s) d s
$$

For $0 \leq t \leq \ell$, we have

$$
\begin{align*}
\left\|q(t)-q^{(k)}(t)\right\|_{p} & \leq\left\|T(\bar{q})-T\left(\zeta^{(k)}\right)\right\|_{\infty ; p}+20 \varepsilon \\
& \leq \frac{1}{2}\left\|\bar{q}-\zeta^{(k)}\right\|_{\infty ; p}+20 \varepsilon \tag{14.35}
\end{align*}
$$

Setting $t=c_{i}$, we have

$$
\begin{aligned}
\left\|\bar{q}-\zeta^{(k+1)}\right\|_{\infty ; p} & =\max _{i}\left\|\bar{q}\left(c_{i}\right)-q^{(k)}\left(c_{i}\right)\right\|_{p} \\
& \leq \frac{1}{2}\left\|q-\zeta^{(k)}\right\|_{\infty ; p}+20 \varepsilon
\end{aligned}
$$

Since $\left\|\zeta^{(k+1)}-\zeta^{(k)}\right\|_{\infty ; p} \leq\left\|\zeta^{(k+1)}-\zeta^{(\infty)}\right\|_{\infty ; p}+\left\|\zeta^{(k)}-\zeta^{(\infty)}\right\|_{\infty ; p} \leq 2 \varepsilon$, we have

$$
\left\|\bar{q}-\zeta^{(k)}\right\|_{\infty ; p} \leq \frac{1}{2}\left\|\bar{q}-\zeta^{(k)}\right\|_{\infty ; p}+2 \varepsilon+20 \varepsilon
$$

Therefore, we have

$$
\left\|\bar{q}-\zeta^{(k)}\right\|_{\infty ; p} \leq 44 \varepsilon
$$

Putting it into (14.35), we have

$$
\left\|q(t)-q^{(k)}(t)\right\|_{p} \leq 42 \varepsilon
$$

for all $0 \leq t \leq \ell$. Using that $\|u(t)-q(t)\|_{p} \leq \varepsilon$, we have

$$
\left\|u(t)-q^{(k)}(t)\right\|_{p} \leq 43 \varepsilon
$$

Hence, it proves the guarantee.
Each iteration involves computing $v_{k}+\int_{0}^{c_{i}} \sum_{j=1}^{d} F\left(\zeta_{j}^{(z)}, c_{j}\right)_{k} \phi_{j}(s) d s$ for all $i \in[d], k \in[n]$. Note that $\sum_{j=1}^{d} F\left(\zeta_{j}^{(z)}, c_{j}\right)_{k} \phi_{j}(s)$ is a polynomial expressed by Lagrange polynomials. Theorem 14.6 .3 shows they can be computed in $O(d \log (d K / \varepsilon)))$ with $\frac{\varepsilon}{K d^{O(1)}}$ accuracy. Since there are $n$ coordinates, it takes $O(n d \log (d K / \varepsilon)))$ time plus $d$ evaluation per iteration.

The theorem above essentially says that if the solution is well approximated by a polynomial and if the $F$ has small enough Lipschitz constant, then we can reconstruct the solution efficiently. Note that this method is not useful for stochastic differential equation because Taylor expansion of the solution involves the high moments of probability distributions which is very expensive to store.

We conclude this subsection with a lemma used in the proof of Lemma 14.6.2.
Lemma 14.6.5. Given $d \geq 2$. Let $\psi(x)=\frac{\sqrt{1-y^{2}} \cos \left(d \cos ^{-1} x\right)}{d(x-y)}$ where $y=\cos \left(\frac{2 k-1}{2 d} \pi\right)$ for some integer $k \in[d]$. Then, we have that

$$
\left|\int_{-1}^{t} \psi(x) d x\right| \leq \frac{20}{d}
$$

for any $-1 \leq t \leq 1$.
Proof. Let $z_{j}=\cos (j \pi / d)$. For $j<k$, we have that

$$
\begin{aligned}
\left|\int_{z_{j-1}}^{z_{j}} \psi(x) d x\right| & \leq \frac{\sqrt{1-y^{2}}}{d}\left|\int_{z_{j-1}}^{z_{j}} \frac{\cos \left(d \cos ^{-1} x\right)}{z_{j}-y} d x\right|+\frac{\sqrt{1-y^{2}}}{d} \int_{z_{j-1}}^{z_{j}}\left|\frac{\cos \left(d \cos ^{-1} x\right)\left(z_{j}-x\right)}{\left(z_{j}-y\right)(x-y)}\right| d x \\
& \leq \frac{\sqrt{1-y^{2}}}{d}\left|\int_{z_{j-1}}^{z_{j}} \frac{\cos \left(d \cos ^{-1} x\right)}{z_{j}-y} d x\right|+\frac{\sqrt{1-y^{2}}}{d} \frac{\left(z_{j}-z_{j-1}\right)^{2}}{\left(z_{j}-y\right)^{2}} \\
& =\frac{\sqrt{1-y^{2}}}{d} \frac{z_{j}-z_{j-1}}{d^{2}-1}+\frac{\sqrt{1-y^{2}}}{d} \frac{\left(z_{j}-z_{j-1}\right)^{2}}{\left(z_{j}-y\right)^{2}} \\
& \leq \frac{\pi}{d^{2}\left(d^{2}-1\right)}+\frac{4}{d} \frac{1}{\left(k-\frac{1}{2}-j\right)^{2}} .
\end{aligned}
$$

Hence, we have that

$$
\begin{aligned}
\left|\int_{-1}^{t} \psi(x) d x\right| & \leq \sum_{k \neq j}\left|\int_{z_{j-1}}^{z_{j}} \psi(x) d x\right|+\int_{z_{k-1}}^{z_{k}}|\psi(x)| d x \\
& \leq \frac{6}{d}+\int_{z_{k-1}}^{z_{k}}|\psi(x)| d x
\end{aligned}
$$

By simple calculation, one can check that $|\psi(x)| \leq 2$ and hence

$$
\left|\int_{-1}^{t} \psi(x) d x\right| \leq \frac{20}{d}
$$

## ■ 14.6.2 Second Order ODE

Now, we consider the following second order ODE

$$
\begin{align*}
\frac{d^{2}}{d t^{2}} u(t) & =F\left(\frac{d}{d t} u(t), u(t), t\right), \text { for } 0 \leq t \leq \ell  \tag{14.36}\\
\frac{d}{d t} u(0) & =w \\
u(0) & =v
\end{align*}
$$

where $F: \mathbb{R}^{2 n+1} \rightarrow \mathbb{R}^{n}$ and $u(t) \in \mathbb{R}^{n}$. Using a standard reduction from second order ODE to first order ODE, we show how to apply our first order ODE method to second order ODE.

Theorem 14.6.6. Let $x(t) \in \mathbb{R}^{n}$ be the solution of the $O D E$ (14.36). Given some $\varepsilon, \ell>0, d \geq 2$ and $1 \leq p \leq \infty$, let $\alpha=40 \ell$ and suppose that

1. There is a degree d polynomial $q$ from $\mathbb{R}$ to $\mathbb{R}^{n}$ such that $q(0)=v, q^{\prime}(0)=w, \| \frac{d}{d t^{2}} x(t)-$ $\frac{d}{d t^{2}} q(t) \|_{p} \leq \frac{\varepsilon}{\ell^{2}}$ for all $0 \leq t \leq \ell$.
2. We have that $\|F(x, \gamma, t)-F(y, \eta, t)\|_{p} \leq \frac{1}{\alpha}\|x-y\|_{p}+\frac{1}{\alpha^{2}}\|\gamma-\eta\|_{p}$ for all $\|x-w\|_{p} \leq K,\|y-w\|_{p} \leq$ $K,\|\gamma-v\|_{p} \leq \alpha K,\|\eta-v\|_{p} \leq \alpha K$ where $K=\alpha \max _{t \in[0, \ell]}\|F(w, v, t)\|_{p}+\|w\|_{p}$.
Then, for any $t$ such that $0 \leq t \leq \ell$, in $O\left(n d \log ^{2}(d K / \varepsilon)\right)$ time plus $O(d \log (K / \varepsilon))$ evaluations of $F$, we can find $p(t), p^{\prime}(t)$ such that $\|u(t)-p(t)\|_{p}=O(\varepsilon)$ and $\left\|u^{\prime}(t)-p^{\prime}(t)\right\|_{p} \leq O(\varepsilon / \ell)$.

Proof. Let $\alpha=40 \ell$. Let $x(t)=\left(\alpha u^{\prime}(\alpha t), u(\alpha t)\right) \in \mathbb{R}^{2 n}$. Note that $x(t)$ satisfies the following ODE

$$
\begin{align*}
\frac{d}{d t} x(t) & =\bar{F}(x(t), t) \text { for } 0 \leq t \leq \ell  \tag{14.37}\\
x(0) & =(\alpha w, v) .
\end{align*}
$$

where $\bar{F}(x(t), t)=\left(\alpha^{2} F\left(\alpha^{-1} x_{(1)}(t), x_{(2)}(t), \alpha t\right), x_{(1)}(t)\right), x_{(1)}(t)$ is the first $n$ variables of $x(t)$ and $x_{(2)}(t)$ is the last $n$ variables of $x(t)$. Next, we verify the conditions of Theorem 14.6.4 for this ODE. Let $\bar{\ell}=\frac{1}{60}$ and $\bar{K}=40 \bar{\ell} \max _{t \in[0, \ell]}\|\bar{F}(x(0), t)\|_{p}$. Note that

$$
\bar{K} \leq \alpha^{2} \max _{t \in[0, \ell]}\|F(w, v, t)\|_{p}+\alpha\|w\|_{p}
$$

For any $y, z$ such that $\|y-x(0)\|_{p} \leq \bar{K}$ and $\|z-x(0)\|_{p} \leq \bar{K}$, we apply the assumption on $F$ and get that

$$
\begin{aligned}
& \|\bar{F}(y, t)-\bar{F}(z, t)\|_{p} \\
\leq & \alpha^{2}\left\|F\left(\alpha^{-1} y_{(1)}, y_{(2)}, t\right)-\bar{F}\left(\alpha^{-1} z_{(1)}, z_{(2)}, t\right)\right\|_{p}+\left\|y_{(1)}-z_{(1)}\right\|_{p} \\
\leq & \alpha^{2}\left(\frac{1}{\alpha}\left\|\alpha^{-1} y_{(1)}-\alpha^{-1} z_{(1)}\right\|_{p}+\frac{1}{\alpha^{2}}\left\|y_{(2)}-z_{(2)}\right\|_{p}\right)+\left\|y_{(1)}-z_{(1)}\right\|_{p} \\
\leq & 3\|y-z\|_{p} .
\end{aligned}
$$

Also, by our assumption on $x$, we have a polynomial $\bar{q}=\left(\alpha q^{\prime}(\alpha t), q(\alpha t)\right)$ such that

$$
\left\|\frac{d}{d t} x(t)-\frac{d}{d t} \bar{q}(t)\right\|_{p} \leq \frac{28 \varepsilon}{\bar{\ell}}
$$

where $x$ is the solution of the ODE (14.37). Therefore, Theorem 14.6.4 shows that we can compute $x(\bar{\ell})$ with $O(\varepsilon)$ error in $O\left(n d \log ^{2}(d K / \varepsilon)\right)$ time plus $O(d \log (K / \varepsilon))$ evaluations of $F$. Now, we can read off $u(\ell)$ from $x_{(2)}$ because

$$
x_{(2)}(\bar{\ell})=u(\alpha \bar{\ell})=u(\ell) .
$$

The result for $u^{\prime}(\ell)$ is similar.

## ■ 14.7 Derivative Estimations

For any smooth one dimension function $f$, we know by Tayloer's theorem that

$$
f(x)=\sum_{k=0}^{N} \frac{f^{(k)}(a)}{k!}(x-a)^{k}+\frac{1}{N!} \int_{a}^{x}(x-t)^{N} f^{(N+1)}(t) d t .
$$

This formula provides a polynomial estimate of $f$ around $a$. To analyze the accuracy of this estimate, we need to bound $\left|f^{(N+1)}(t)\right|$. In one dimension, we could simply give explicit formulas for the derivatives of $f$ and use it to estimate the remainder term. However, for functions in high dimension, it is usually too tedious. In this section, we describe some techniques for bounding the derivatives of higher dimensional functions.

The derivatives of one-variable functions can be bounded via Cauchy's estimate (Theorem (14.2.11)). In Subsection 14.7.1, we give calculus rules that reduces the problem of estimating derivatives of highdimensional functions to derivatives of one-dimensional functions. In Subsection 14.7.2, we show how to reduce bounding the derivative for an arbitrary ODE to an ODE in one dimension.

## ■ 14.7.1 Explicit Function

In this subsection, we show how to bound the derivatives of a complicated explicit function using the following object, generating upper bound. We reduce estimates of the derivatives of functions in high dimension to one variable rational functions. Since rational functions are holomorphic, one can apply Cauchy's estimates (Theorem (14.2.11)) to bound their derivatives.

Definition 14.7.1. Given a function $F$. We call that $F \leq_{x} f$ for some one variable function $f: \mathbb{R} \rightarrow \mathbb{R}$ if

$$
\begin{equation*}
\left\|D^{(k)} F(x)\left[\Delta_{1}, \Delta_{2}, \cdots, \Delta_{k}\right]\right\| \leq f^{(k)}(0) \prod_{i=1}^{k}\left\|\Delta_{i}\right\|^{k} \tag{14.38}
\end{equation*}
$$

for any $k \geq 0$ and any $\Delta_{i}$.
Remark. In general, $F \leq_{x} f$ and $f(t) \leq g(t)$ point-wise does NOT imply $F \leq_{x} g$. However, $F \leq_{x} f$ and $f \leq_{0} g$ does imply $F \leq_{x} g$.

This concept is useful for us to reduce bounding derivatives of a high dimension function to bounding derivatives of 1 dimension function. First of all, we note that upper bounds are composable.

Lemma 14.7.2. Given $F \leq_{x} f$ and $G \leq_{F(x)} g$, we have that

$$
G \circ F \leq_{x} g \circ \bar{f}
$$

where $\bar{f}(s)=f(s)-f(0)$.
Proof. Fix any $\Delta_{1}, \Delta_{2}, \cdots$ be unit vectors in the domain of $F$. Let $H(x)=G \circ F(x)$. By chain rule, we have that

$$
\begin{aligned}
D H(x)\left[\Delta_{1}\right]= & D G(F(x))\left[D F(x)\left[\Delta_{1}\right]\right], \\
D H(x)\left[\Delta_{1}, \Delta_{2}\right]= & D G(F(x))\left[D^{2} F(x)\left[\Delta_{1}, \Delta_{2}\right]\right] \\
& +D^{2} G(F(x))\left[D F(x)\left[\Delta_{1}\right], D F(x)\left[\Delta_{2}\right]\right], \\
D H(x)\left[\Delta_{1}, \Delta_{2}, \Delta_{3}\right]= & D G(F(x))\left[D^{2} F(x)\left[\Delta_{1}, \Delta_{2}, \Delta_{3}\right]\right] \\
& +D^{2} G(F(x))\left[D^{2} F(x)\left[\Delta_{1}, \Delta_{2}\right], D F(x)\left[\Delta_{3}\right]\right] \\
& +D^{2} G(F(x))\left[D^{2} F(x)\left[\Delta_{1}, \Delta_{3}\right], D F(x)\left[\Delta_{2}\right]\right] \\
& +D^{2} G(F(x))\left[D^{2} F(x)\left[\Delta_{2}, \Delta_{3}\right], D F(x)\left[\Delta_{1}\right]\right] \\
& +D^{3} G(F(x))\left[D F(x)\left[\Delta_{1}\right], D F(x)\left[\Delta_{2}\right], D F(x)\left[\Delta_{3}\right]\right],
\end{aligned}
$$

Since $G \leq_{F(x)} g$, equation (14.38) shows that

$$
\begin{aligned}
\left\|D H(x)\left[\Delta_{1}\right]\right\| \leq & g^{(1)}(0)\left\|D F(x)\left[\Delta_{1}\right]\right\|, \\
\left\|D^{2} H(x)\left[\Delta_{1}, \Delta_{2}\right]\right\| \leq & g^{(1)}(0)\left\|D^{2} F(x)\left[\Delta_{1}, \Delta_{2}\right]\right\| \\
& +g^{(2)}(0)\left\|D F(x)\left[\Delta_{1}\right]\right\|\left\|_{2}\right\| D F(x)\left[\Delta_{2}\right] \|, \\
\left\|D^{3} H(x)\left[\Delta_{1}, \Delta_{2}, \Delta_{3}\right]\right\| \leq & g^{(1)}(0)\left\|D^{2} F(x)\left[\Delta_{1}, \Delta_{2}, \Delta_{3}\right]\right\| \\
& +g^{(2)}(0)\left\|D^{2} F(x)\left[\Delta_{1}, \Delta_{2}\right]\right\|\left\|D F(x)\left[\Delta_{3}\right]\right\| \\
& +g^{(2)}(0)\left\|D^{2} F(x)\left[\Delta_{1}, \Delta_{3}\right]\right\|\left\|D F(x)\left[\Delta_{2}\right]\right\| \\
& +g^{(2)}(0)\left\|D^{2} F(x)\left[\Delta_{2}, \Delta_{3}\right]\right\|\left\|D F(x)\left[\Delta_{1}\right]\right\| \\
\vdots & +g^{(3)}(0)\left\|D F(x)\left[\Delta_{1}\right]\right\|\left\|D F(x)\left[\Delta_{2}\right]\right\|\left\|D F(x)\left[\Delta_{3}\right]\right\| .
\end{aligned}
$$

Now, we use $F \leq_{x} f$ to get

$$
\begin{aligned}
\left\|D H(x)\left[\Delta_{1}\right]\right\| & \leq g^{(1)}(0) f^{(1)}(0)=(g \circ \bar{f})^{(1)}(0) \\
\left\|D^{2} H(x)\left[\Delta_{1}, \Delta_{2}\right]\right\| & \leq g^{(1)}(0) f^{(2)}(0)+g^{(2)}(0) f^{(1)}(0)^{2}=(g \circ \bar{f})^{(2)}(0), \\
\left\|D^{3} H(x)\left[\Delta_{1}, \Delta_{2}, \Delta_{3}\right]\right\| & \leq g^{(1)}(0) f^{(3)}(0)+2 g^{(2)}(0) f^{(2)}(0) f^{(1)}(0) \\
& \vdots+g^{(3)}(0) f^{(1)}(0)^{3}=(g \circ \bar{f})^{(3)}(0) .
\end{aligned}
$$

Therefore, we have that $\left\|D^{k} H(x)\left[\Delta_{i}\right]\right\| \leq(g \circ \bar{f})^{(k)}(0)$ for all $k \geq 1$. For $k=0$, we have that $\|H(x)\|=\|G(F(x))\| \leq g(0)=g(\bar{f}(0))$.

Next, we give some extra calculus rule for generating upper bounds.

Lemma 14.7.3. Given that $H_{i} \leq_{x} h_{i}$ for all $i=1, \cdots, k$. Then, we have that

$$
\sum_{i=1}^{k} H_{i} \leq_{x} \sum_{i=1}^{k} h_{i} \text { and } \prod_{i=1}^{k} H_{i} \leq_{x} \prod_{i=1}^{k} h_{i} .
$$

Given that $H \leq_{x} h$ and $\left\|H^{-1}(x)\right\| \leq C$, we have that

$$
H^{-1} \leq_{x} \frac{1}{C^{-1}-(h(s)-h(0))} .
$$

Proof. Fix $\Delta_{1}, \Delta_{2}, \cdots$ be unit vectors. For the first claim, let $H=\sum H_{i}$, we note that

$$
D^{j} H\left[\Delta_{1}, \cdots, \Delta_{j}\right]=\sum_{i=1}^{k} D^{j} H_{i}\left[\Delta_{1}, \cdots, \Delta_{j}\right] .
$$

Therefore, we have that $\left\|D^{j} H\left[\Delta_{1}, \cdots, \Delta_{j}\right]\right\| \leq \sum_{i=1}^{k}\left\|D^{j} H_{i}\left[\Delta_{1}, \cdots, \Delta_{j}\right]\right\|$. Since $H \leq_{x} h_{i}$, we have that $\left\|D^{j} H\left[\Delta_{1}, \cdots, \Delta_{j}\right]\right\| \leq \sum_{i=1}^{k} h^{(j)}(0)$. Hence, we have $H \leq_{x} \sum_{i=1}^{k} h_{i}$.

For the second claim, we note that

$$
D^{j} G=\sum_{i_{1}+i_{2}+\cdots+i_{k}=j} \prod_{l=1}^{k} D^{i_{l}} H_{l} .
$$

Let $g=\prod_{i=1}^{k} h_{i}$. Then, we have that

$$
\left\|D^{j} G\right\| \leq \sum_{i_{1}+i_{2}+\cdots+i_{k}=j} \prod_{l=1}^{k}\left\|D^{i_{l}} H_{l}\right\| \leq \sum_{i_{1}+i_{2}+\cdots+i_{k}=j} \prod_{l=1}^{k} h_{i}^{\left(i_{l}\right)}(0)=D^{j} g(0) .
$$

Hence, $G \leq_{x} g$.
For the last claim, we first consider the function $\Phi(M)=M^{-1}$. Note that $D \Phi\left[\Delta_{1}\right]=-M^{-1} \Delta_{1} M^{-1}$, $D^{2} \Phi[\Delta]=M^{-1} \Delta_{1} M^{-1} \Delta_{2} M^{-1}+M^{-1} \Delta_{2} M^{-1} \Delta_{1} M^{-1}$ and hence

$$
\left\|D^{j} \Phi\right\| \leq j!\left\|M^{-1}\right\|^{j+1}=j!C^{j+1}
$$

Hence, we have $\Phi \leq_{M} \sum_{j=0}^{\infty} \frac{j!C^{j+1}}{j!} s^{j}=\frac{1}{C^{-1}-s}$. By Lemma 14.7.2, we see that $H^{-1} \leq_{x} \frac{1}{C^{-1}-(h(s)-h(0))}$.

## ■ 14.7.2 Explicit ODE

In this section, we study the Taylor expansion of the solution of ODE (14.31).
Lemma 14.7.4. Let $u(t)$ be the solution of the $O D E u^{\prime}(t)=F(u(t))$. Suppose that $F \leq_{u(0)} f$ and let $\psi(t)$ be the solution of the $O D E \psi^{\prime}(t)=f(\psi(t))$ and $\psi(0)=0$. Then, we have

$$
\left\|u^{(k)}(0)\right\|_{2} \leq \psi^{(k)}(0)
$$

for all $k \geq 1$.
Proof. Since $u^{\prime}(t)=F(u(t))$, we have that

$$
\begin{aligned}
u^{(2)}(t)= & D F(u(t))\left[u^{(1)}(t)\right], \\
u^{(3)}(t)= & D F(u(t))\left[u^{(2)}(t)\right]+D^{2} F(u(t))\left[u^{(1)}(t), u^{(1)}(t)\right], \\
u^{(4)}(t)= & D F(u(t))\left[u^{(3)}(t)\right]+2 D^{2} F(u(t))\left[u^{(2)}(t), u^{(1)}(t)\right] \\
& +D^{3} F(u(t))\left[u^{(1)}(t), u^{(1)}(t), u^{(1)}(t)\right]
\end{aligned}
$$

Therefore, we have

$$
\begin{aligned}
\left\|u^{(2)}(0)\right\|_{2} & \leq f^{(1)}(0)\left\|u^{(1)}(0)\right\|_{2}, \\
\left\|u^{(3)}(0)\right\|_{2} & \leq f^{(1)}(0)\left\|u^{(2)}(0)\right\|_{2}+f^{(2)}(0)\left\|u^{(1)}(0)\right\|_{2}^{2} \\
\left\|u^{(4)}(0)\right\|_{2} & \leq f^{(1)}(0)\left\|u^{(3)}(0)\right\|_{2}+2 f^{(2)}(0)\left\|u^{(2)}(0)\right\|_{2}\left\|u^{(1)}(0)\right\|_{2}+f^{(3)}(0)\left\|u^{(1)}(0)\right\|_{2}^{3}, \\
& \vdots
\end{aligned}
$$

By expanding $\psi^{\prime}(t)=f(\psi(t))$ at $t=0$, we see that

$$
\begin{aligned}
\psi^{(2)}(0) & =f^{(1)}(0) \psi^{(1)}(0) \\
\psi^{(3)}(0) & =f^{(1)}(0) \psi^{(2)}(0)+f^{(2)}(0)\left(\psi^{(1)}(0)\right)^{2} \\
\psi^{(4)}(0) & =f^{(1)}(0) \psi^{(3)}(0)+2 f^{(2)}(0) \psi^{(2)}(0) \psi^{(1)}(0)+f^{(2)}(0)\left(\psi^{(1)}(0)\right)^{3}, \\
& \vdots
\end{aligned}
$$

Since $\left\|u^{(1)}(0)\right\|_{2}=\|F(u(0))\|_{2} \leq f(0)=\psi^{(1)}(0)$, we have that $\left\|u^{(k)}(0)\right\|_{2} \leq \psi^{(k)}(0)$ for all $k \geq 1$.

Now, we apply Lemma 14.7 .4 to second order ODE.
Lemma 14.7.5. Let $u(t)$ be the solution of the $O D E u^{\prime \prime}(t)=F\left(u^{\prime}(t), u(t), t\right)$. Given some $\alpha>0$ and define $\bar{F}(y, x, t)=\alpha^{2} F\left(\alpha^{-1} y(t)+u^{\prime}(0), x(t)+\alpha t u^{\prime}(0), \alpha t\right)$. Suppose that $\bar{F} \leq_{(0, u(0), 0)} f$ and let $\psi(t)$ be the solution of the $O D E \psi^{\prime}(t)=1+\psi(t)+f(\psi(t))$ and $\psi(0)=0$. Then, we have

$$
\left\|u^{(k)}(0)\right\|_{2} \leq \frac{\psi^{(k)}(0)}{\alpha^{k}}
$$

for all $k \geq 2$.
Proof. Let $x(t)=u(\alpha t)-\alpha t u^{\prime}(0)$ and $y(t)=\alpha u^{\prime}(\alpha t)-\alpha u^{\prime}(0)$. Then, we can write the problem into
first order ODE

$$
\begin{aligned}
y^{\prime}(t) & =\alpha^{2} F\left(\alpha^{-1} y(t)+u^{\prime}(0), x(t)+\alpha t u^{\prime}(0), \alpha t\right)=\bar{F}(y(t), x(t), t), \\
x^{\prime}(t) & =y(t), \\
t^{\prime} & =1 .
\end{aligned}
$$

$\underline{\text { Let }} \mathbf{F}(y, x, t)=(\bar{F}(y, x, t), y, 1)$. Then, we have that $\left\|D^{k} \mathbf{F}\right\| \leq\left\|D^{k} 1\right\|+\left\|D^{k} y\right\|+\left\|D^{k} \bar{F}\right\|$. Using $\bar{F} \leq_{(0, u(0), 0)} f$, we have that

$$
\mathbf{F} \leq_{(0, x(0), 0)} 1+t+f
$$

By Lemma 14.7.4, we know that

$$
\left\|u^{(k)}(0)\right\|_{2}=\frac{\left\|x^{(k)}(0)\right\|_{2}}{\alpha^{k}} \leq \frac{\psi^{(k)}(0)}{\alpha^{k}} .
$$

## Bibliography

[1] Jacob Abernethy and Elad Hazan. "Faster Convex Optimization: Simulated Annealing with an Efficient Universal Barrier". In: arXiv preprint arXiv:1507.02528 (2015).
[2] Dimitris Achlioptas. "Database-friendly random projections: Johnson-Lindenstrauss with binary coins". In: Journal of computer and System Sciences 66.4 (2003), pp. 671-687.
[3] Rudolf Ahlswede and Andreas Winter. "Strong converse for identification via quantum channels". In: Information Theory, IEEE Transactions on 48.3 (2002), pp. 569-579.
[4] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. Network flows. New York, NY, USA: Elsevier North-Holland, Inc., 1989.
[5] Ravindra K. Ahuja et al. "Applications of Network Optimization". In: Network Models. Vol. 7. Handbooks in Operations Research and Management Science. North-Holland, 1995, pp. 1-75.
[6] Martin Aigner and Thomas A Dowling. "Matching theory for combinatorial geometries". In: Transactions of the American Mathematical Society 158.1 (1971), pp. 231-245.
[7] Boris Alexeev et al. "Phase retrieval with polarization". In: SIAM Journal on Imaging Sciences 7.1 (2014), pp. 35-66.
[8] Zeyuan Allen-Zhu, Yin Tat Lee, and Lorenzo Orecchia. "Using Optimization to Obtain a Width-Independent, Parallel, Simpler, and Faster Positive SDP Solver". In: arXiv preprint arXiv:1507.02259 (2015).
[9] Zeyuan Allen-Zhu, Zhenyu Liao, and Lorenzo Orecchia. "Spectral Sparsification and Regret Minimization Beyond Multiplicative Updates". In: Proceedings of the 47 th Annual ACM Symposium on Theory of Computing (STOC). 2015, pp. 237-245.
[10] Kurt M. Anstreicher. "Large step volumetric potential reduction algorithms for linear programming". In: Annals of Operations Research 62.1 (1996), pp. 521-538.
[11] Kurt M Anstreicher. "On Vaidya's volumetric cutting plane method for convex programming". In: Mathematics of Operations Research 22.1 (1997), pp. 63-89.
[12] Kurt M. Anstreicher. "The volumetric barrier for convex quadratic constraints". In: Mathematical Programming 100.3 (2004), pp. 613-662.
[13] Kurt M. Anstreicher. "The volumetric barrier for semidefinite programming". In: Mathematics of Operations Research 25.3 (2000), pp. 365-380.
[14] Kurt M. Anstreicher. "Towards a practical volumetric cutting plane method for convex programming". In: SIAM Journal on Optimization 9.1 (1998), pp. 190-206.
[15] Kurt M. Anstreicher. "Volumetric path following algorithms for linear programming". In: Math. Program. 76 (1996), pp. 245-263.
[16] Mica Arie-Nachimson et al. "Global motion estimation from point matches". In: 3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT), 2012 Second International Conference on. IEEE. 2012, pp. 81-88.
[17] Sanjeev Arora, Elad Hazan, and Satyen Kale. "Fast algorithms for approximate semidefinite programming using the multiplicative weights update method". In: Foundations of Computer Science, 2005. FOCS 2005. 46th Annual IEEE Symposium on. IEEE. 2005, pp. 339-348.
[18] Sanjeev Arora, Elad Hazan, and Satyen Kale. "The Multiplicative Weights Update Method: A Meta-Algorithm and Applications".
[19] Sanjeev Arora and Satyen Kale. "A combinatorial, primal-dual approach to semidefinite programs". In: Proceedings of the thirty-ninth annual ACM symposium on Theory of computing. ACM. 2007, pp. 227-236.
[20] David S Atkinson and Pravin M Vaidya. "A cutting plane algorithm for convex programming that uses analytic centers". In: Mathematical Programming 69.1-3 (1995), pp. 1-43.
[21] Y. Aumann and Y. Rabani. "An O(log k) Approximate Min-Cut Max-Flow Theorem and Approximation Algorithm". In: SIAM Journal on Computing 27.1 (1998), pp. 291-301.
[22] Haim Avron, Petar Maymounkov, and Sivan Toledo. "Blendenpik: Supercharging LAPACK's least-squares solver". In: SIAM J. Scientific Computing 32.3 (2010), pp. 1217-1236.
[23] Haim Avron and Sivan Toledo. "Effective stiffness: Generalizing effective resistance sampling to finite element matrices". In: arXiv preprint arXiv:1110.4437 (2011).
[24] Francis Bach. "Learning with submodular functions: A convex optimization perspective". In: Foundations and Trends in Machine Learning (2013).
[25] Mihai Badoiu, Sariel Har-Peled, and Piotr Indyk. "Approximate clustering via core-sets". In: Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada. 2002, pp. 250-257.
[26] Olivier Bahn et al. "A cutting plane method from analytic centers for stochastic programming". In: Mathematical Programming 69.1-3 (1995), pp. 45-73.
[27] Chanderjit Bajaj. "The algebraic degree of geometric optimization problems". English. In: Discrete $\xi^{3}$ Computational Geometry 3.2 (1988), pp. 177-191.
[28] Egon Balas and Chang-Sung Yu. "A Note on the Weiszfeld-Kuhn Algorithm for the General Fermat Problem". In: Managme Sci Res Report 484 (1982), pp. 1-6.
[29] Afonso S Bandeira, Amit Singer, and Daniel A Spielman. "A Cheeger inequality for the graph connection Laplacian". In: SIAM Journal on Matrix Analysis and Applications 34.4 (2013), pp. 1611-1630.
[30] Francisco Barahona and William H Cunningham. "A submodular network simplex method". In: Mathematical Programming at Oberwolfach II. Springer, 1984, pp. 9-31.
[31] Joshua Batson, Daniel Spielman, and Nikhil Srivastava. "Twice-Ramanujan sparsifiers". In: SIAM Journal on Computing 41.6 (2012), pp. 1704-1721.
[32] Stephen R Becker, Emmanuel J Candès, and Michael C Grant. "Templates for convex cone problems with applications to sparse signal recovery". In: Mathematical programming computation 3.3 (2011), pp. 165-218.
[33] András A Benczúr and David R Karger. "Approximating st minimum cuts in Õ (n 2) time". In: Proceedings of the twenty-eighth annual ACM symposium on Theory of computing. ACM. 1996, pp. 47-55.
[34] M. Bern et al. "Support-graph preconditioners". In: SIAM J. Matrix Anal. \&3 Appl 27.4 (2006), pp. 930-951.
[35] Dimitri Bertsekas. "Nonlinear programming". In: (1999).
[36] Dimitris Bertsimas and Santosh Vempala. "Solving convex programs by random walks". In: Journal of the ACM (JACM) 51.4 (2004), pp. 540-556.
[37] Prosenjit Bose, Anil Maheshwari, and Pat Morin. "Fast approximations for sums of distances, clustering and the Fermat-Weber problem". In: Computational Geometry 24.3 (2003), pp. 135146.
[38] Jean Bourgain, Joram Lindenstrauss, and V Milman. "Approximation of zonoids by zonotopes". In: Acta mathematica 162.1 (1989), pp. 73-141.
[39] Carl Brezovec, Gerard Cornuéjols, and Fred Glover. "Two algorithms for weighted matroid intersection". In: Mathematical Programming 36.1 (1986), pp. 39-53.
[40] Sébastien Bubeck. "Convex Optimization: Algorithms and Complexity". In: Foundations and Trends in Machine Learning 8.3-4 (2015), pp. 231-357.
[41] Sébastien Bubeck. "Theory of Convex Optimization for Machine Learning". In: Arxiv preprint arXiv:1405.4980 (2014).
[42] Sébastien Bubeck and Ronen Eldan. "The entropic barrier: a simple and optimal universal self-concordant barrier". In: arXiv preprint arXiv:1412.1587 (2014).
[43] Sébastien Bubeck, Ronen Eldan, and Joseph' Lehec. "Sampling from a log-concave distribution with Projected Langevin Monte Carlo". In: arXiv preprint arXiv:1507.02564 (2015).
[44] Yang Cai, Constantinos Daskalakis, and S Matthew Weinberg. "Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization". In: Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on. IEEE. 2012, pp. 130-139.
[45] Yang Cai, Constantinos Daskalakis, and S Matthew Weinberg. "Reducing revenue to welfare maximization: Approximation algorithms and other generalizations". In: Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms. SIAM. 2013, pp. 578595.
[46] Marcel K. de Carli Silva, Nicholas J. A. Harvey, and Cristiane M. Sato. "Sparse Sums of Positive Semidefinite Matrices". In: CoRR abs/1107.0088 (2011).
[47] R. Chandrasekaran and A. Tamir. "Open questions concerning Weiszfeld's algorithm for the Fermat-Weber location problem". English. In: Mathematical Programming 44.1-3 (1989), pp. 293-295.
[48] Chih-Chung Chang and Chih-Jen Lin. "LIBSVM: A library for support vector machines". In: ACM Transactions on Intelligent Systems and Technology (TIST). 2.3 (2011). Software available at http://www.csie.ntu.edu.tw/ cjlin/libsvm., p. 27.
[49] L. Paul Chew. "There Are Planar Graphs Almost as Good as the Complete Graph". In: Journal of Computer and System Sciences 39.2 (1989), pp. 205-219.
[50] Hui Han Chin et al. "Runtime guarantees for regression problems". In: ITCS. 2013, pp. 269282.
[51] Paul Christiano et al. "Electrical flows, laplacian systems, and faster approximation of maximum flow in undirected graphs". In: Proceedings of the 43 rd annual ACM symposium on Theory of computing. ACM. 2011, pp. 273-282.
[52] Nam-Kee Chung and Dong-Wan Tcha. "A dual algorithm for submodular flow problems". In: Operations research letters 10.8 (1991), pp. 489-495.
[53] Kenneth Clarkson and David Woodruff. "Low rank approximation and regression in input sparsity time". In: Proceedings of the 45th annual ACM symposium on Symposium on theory of computing. ACM. 2013, pp. 81-90.
[54] Kenneth Clarkson and David Woodruff. "Numerical linear algebra in the streaming model". In: Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC). 2009, pp. 205-214.
[55] E. S. Coakley, Vladimir Rokhlin, and Mark Tygert. "A Fast Randomized Algorithm for Orthogonal Projection". In: SIAM J. Scientific Computing 33.2 (2011), pp. 849-868.
[56] Michael B Cohen and Richard Peng. "Lp Row Sampling by Lewis Weights". In: arXiv preprint arXiv:1412.0588 (2014).
[57] Michael B. Cohen et al. "Solving SDD Linear Systems in Nearly Mlog1/2N Time". In: Proceedings of the 46 th Annual ACM Symposium on Theory of Computing. STOC '14. New York, New York: ACM, 2014, pp. 343-352.
[58] Leon Cooper and I.Norman Katz. "The Weber problem revisited". In: Computers and Mathematics with Applications 7.3 (1981), pp. 225-234.
[59] B. Cousins and S. Vempala. "A Cubic Algorithm for Computing Gaussian Volume". In: SODA. 2014, pp. 1215-1228.
[60] William H Cunningham. "Improved bounds for matroid partition and intersection algorithms". In: SIAM Journal on Computing 15.4 (1986), pp. 948-957.
[61] William H Cunningham. "On submodular function minimization". In: Combinatorica 5.3 (1985), pp. 185-192.
[62] William H Cunningham and András Frank. "A primal-dual algorithm for submodular flows". In: Mathematics of Operations Research 10.2 (1985), pp. 251-262.
[63] Samuel I Daitch and Daniel A Spielman. "Faster approximate lossy generalized flow via interior point algorithms". In: Proceedings of the 40 th annual ACM symposium on Theory of computing. ACM. 2008, pp. 451-460.
[64] Arnak S Dalalyan. "Theoretical guarantees for approximate sampling from smooth and logconcave densities". In: arXiv preprint arXiv:1412.7392 (2014).
[65] Constantinos Daskalakis and S Matthew Weinberg. "Bayesian truthful mechanisms for job scheduling from bi-criterion approximation algorithms". In: Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms. SIAM. 2015, pp. 1934-1952.
[66] James Demmel et al. "Fast matrix multiplication is stable". In: Numerische Mathematik 106.2 (2007), pp. 199-224.
[67] A. B. Dieker. "Reflected Brownian Motion". In: Wiley Encyclopedia of Operations Research and Management Science. John Wiley \& Sons, Inc., 2010.
[68] EA Dinic. "An algorithm for the solution of the max-flow problem with the polynomial estimation". In: Doklady Akademii Nauk SSSR 194.4 (1970), pp. 1277-1280.
[69] Zvi Drezner et al. "Facility location". In: Springer, 2002. Chap. The Weber problem, pp. 1-36.
[70] Petros Drineas, Ravi Kannan, and Michael W. Mahoney. "Fast Monte Carlo algorithms for matrices I: Approximating matrix multiplication". In: SIAM J. Comput. 36.1 (2006). Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 132-157.
[71] Petros Drineas, Ravi Kannan, and Michael W. Mahoney. "Fast Monte Carlo Algorithms for Matrices II: Computing a Low-Rank Approximation to a Matrix". In: SIAM J. Comput. 36.1 (2006), pp. 158-183.
[72] Petros Drineas, Ravi Kannan, and Michael W. Mahoney. "Fast Monte Carlo Algorithms for Matrices III: Computing a Compressed Approximate Matrix Decomposition". In: SIAM J. Comput. 36.1 (2006). Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 184-206.
[73] Petros Drineas and Michael W Mahoney. "Effective resistances, statistical leverage, and applications to linear equation solving". In: arXiv preprint arXiv:1005.3097 (2010).
[74] Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. "Relative-Error CUR Matrix Decompositions". In: SIAM J. Matrix Anal. Appl. 30.2 (2008), pp. 844-881.
[75] Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. "Sampling Algorithms for $\ell_{2}$ Regression and Applications". In: Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA). 2006, pp. 1127-1136.
[76] Petros Drineas, MichaelW. Mahoney, and S. Muthukrishnan. "Subspace Sampling and RelativeError Matrix Approximation: Column-Based Methods". In: Proceedings of the 10th International Workshop on Randomization and Computation (RANDOM). 2006, pp. 316-326.
[77] A Dutt, M Gu, and V Rokhlin. "Fast algorithms for polynomial interpolation, integration, and differentiation". In: SIAM Journal on Numerical Analysis 33.5 (1996), pp. 1689-1711.
[78] M. E. Dyer and A. M. Frieze. "Computing the volume of a convex body: a case where randomness provably helps". In: Proc. of AMS Symposium on Probabilistic Combinatorics and Its Applications. 1991, pp. 123-170.
[79] M. E. Dyer, A. M. Frieze, and R. Kannan. "A random polynomial-time algorithm for approximating the volume of convex bodies". In: J. ACM 38.1 (1991), pp. 1-17.
[80] Jack Edmonds. "Matroid intersection". In: Annals of discrete Mathematics 4 (1979), pp. 3949.
[81] Jack Edmonds. "Matroid partition". In: Mathematics of the Decision Sciences 11 (1968), pp. 335-345.
[82] Jack Edmonds. "Submodular functions, matroids, and certain polyhedra". In: Edited by G. Goos, J. Hartmanis, and J. van Leeuwen (1970), p. 11.
[83] Jack Edmonds and Rick Giles. "A min-max relation for submodular functions on graphs". In: Studies in Integer Programming (PL Hammer, EL Johnson and BH Korte, eds.), Ann. Discrete Math 1 (1977), pp. 185-204.
[84] Jack Edmonds and Richard M Karp. "Theoretical improvements in algorithmic efficiency for network flow problems". In: Journal of the ACM (JACM) 19.2 (1972), pp. 248-264.
[85] Shimon Even and R. Endre Tarjan. "Network Flow and Testing Graph Connectivity". In: SIAM Journal on Computing 4.4 (Dec. 1975), pp. 507-518.
[86] Dan Feldman and Michael Langberg. "A unified framework for approximating and clustering data". In: Proceedings of the forty-third annual ACM symposium on Theory of computing. ACM. 2011, pp. 569-578.
[87] Lisa Fleischer and Satoru Iwata. "A push-relabel framework for submodular function minimization and applications to parametric optimization". In: Discrete Applied Mathematics 131.2 (2003), pp. 311-322.
[88] Lisa Fleischer and Satoru Iwata. "Improved algorithms for submodular function minimization and submodular flow". In: Proceedings of the thirty-second annual ACM symposium on Theory of computing. ACM. 2000, pp. 107-116.
[89] Lisa Fleischer, Satoru Iwata, and S Thomas McCormick. "A faster capacity scaling algorithm for minimum cost submodular flow". In: Mathematical Programming 92.1 (2002), pp. 119-139.
[90] András Frank. "A weighted matroid intersection algorithm". In: Journal of Algorithms 2.4 (1981), pp. 328-336.
[91] András Frank and Éva Tardos. "An application of simultaneous Diophantine approximation in combinatorial optimization". In: Combinatorica 7.1 (1987), pp. 49-65.
[92] S. Fujishige. "Algorithms for solving the independent-flow problems". In: Journal of the Operations Research Society of Japan (1978).
[93] Satoru Fujishige. "An out-of-kilter method for submodular flows". In: Discrete applied mathematics 17.1 (1987), pp. 3-16.
[94] Satoru Fujishige and Satoru Iwata. "Algorithms for submodular flows". In: IEICE TRANSACTIONS on Information and Systems 83.3 (2000), pp. 322-329.
[95] Satoru Fujishige, Hans Röck, and Uwe Zimmermann. "A strongly polynomial algorithm for minimum cost submodular flow problems". In: Mathematics of Operations Research 14.1 (1989), pp. 60-69.
[96] Satoru Fujishige and Zhang Xiaodong. "An efficient cost scaling algorithm for the independent assignment problem". In: Journal of the Operations Research Society of Japan 38.1 (1995), pp. 124-136.
[97] Mituhiro Fukuda et al. "Exploiting sparsity in semidefinite programming via matrix completion I: General framework". In: SIAM Journal on Optimization 11.3 (2001), pp. 647-674.
[98] Zvi Galil and Éva Tardos. "An O (n $2(\mathrm{~m}+\mathrm{n} \log \mathrm{n}) \log \mathrm{n})$ min-cost flow algorithm". In: Journal of the ACM (JACM) 35.2 (1988), pp. 374-386.
[99] François Le Gall. "Powers of tensors and fast matrix multiplication". In: arXiv preprint arXiv:1401.7714 (2014).
[100] François Le Gall. "Powers of tensors and fast matrix multiplication". In: arXiv:1401.7714 (2014).
[101] Jean-Louis Goffin, Zhi-Quan Luo, and Yinyu Ye. "Complexity analysis of an interior cutting plane method for convex feasibility problems". In: SIAM Journal on Optimization 6.3 (1996), pp. 638-652.
[102] Jean-Louis Goffin and Jean-Philippe Vial. "Convex nondifferentiable optimization: A survey focused on the analytic center cutting plane method". In: Optimization Methods and Software 17.5 (2002), pp. 805-867.
[103] Jean-Louis Goffin and Jean-Philippe Vial. "Shallow, deep and very deep cuts in the analytic center cutting plane method". In: Mathematical Programming 84.1 (1999), pp. 89-103.
[104] Andrew V. Goldberg and Satish Rao. "Beyond the flow decomposition barrier". In: J. ACM 45.5 (1998), pp. 783-797.
[105] Andrew V Goldberg and Robert E Tarjan. "A new approach to the maximum-flow problem". In: Journal of the ACM (JACM) 35.4 (1988), pp. 921-940.
[106] Andrew V Goldberg and Robert E Tarjan. "Finding minimum-cost circulations by successive approximation". In: Mathematics of Operations Research 15.3 (1990), pp. 430-466.
[107] Clovis C Gonzaga. "Path-following methods for linear programming". In: SIAM review 34.2 (1992), pp. 167-224.
[108] Clóvis C Gonzaga and Elizabeth W Karas. "Fine tuning Nesterov's steepest descent algorithm for differentiable convex programming". In: Mathematical Programming 138.1-2 (2013), pp. 141-166.
[109] Martin Grötschel, László Lovász, and Alexander Schrijver. "Geometric Algorithms and Combinatorial Optimization". In: Springer, 1988.
[110] Martin Grötschel, László Lovász, and Alexander Schrijver. "The ellipsoid method and its consequences in combinatorial optimization". In: Combinatorica 1.2 (1981), pp. 169-197.
[111] B. Grünbaum. "Partitions of mass-distributions and of convex bodies by hyperplanes". In: Pacific J. Math 10.4 (1960), pp. 1257-1261.
[112] N. Halko, P. G. Martinsson, and J. A. Tropp. "Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions". In: SIAM Review 53.2 (2011), pp. 217-288.
[113] Sariel Har-Peled and Akash Kushal. "Smaller coresets for k-median and k-means clustering". In: Proceedings of the twenty-first annual symposium on Computational geometry. ACM. 2005, pp. 126-134.
[114] J. Michael Harrison. Brownian motion and stochastic flow systems. Wiley series in probability and mathematical statistics. New York: Wiley, 1985.
[115] Christoph Helmberg and Franz Rendl. "A spectral bundle method for semidefinite programming". In: SIAM Journal on Optimization 10.3 (2000), pp. 673-696.
[116] P. Indyk and Stanford University. Computer Science Dept. High-dimensional computational geometry. Stanford University, 2000.
[117] Satoru Iwata. "A capacity scaling algorithm for convex cost submodular flows". In: Mathematical programming 76.2 (1997), pp. 299-308.
[118] Satoru Iwata. "A faster scaling algorithm for minimizing submodular functions". In: SIAM Journal on Computing 32.4 (2003), pp. 833-840.
[119] Satoru Iwata. "A fully combinatorial algorithm for submodular function minimization". In: Journal of Combinatorial Theory, Series B 84.2 (2002), pp. 203-212.
[120] Satoru Iwata. "Submodular function minimization". In: Mathematical Programming 112.1 (2008), pp. 45-64.
[121] Satoru Iwata, Lisa Fleischer, and Satoru Fujishige. "A combinatorial strongly polynomial algorithm for minimizing submodular functions". In: Journal of the ACM (JACM) 48.4 (2001), pp. 761-777.
[122] Satoru Iwata, S Thomas McCormick, and Maiko Shigeno. "A fast cost scaling algorithm for submodular flow". In: Information Processing Letters 74.3 (2000), pp. 123-128.
[123] Satoru Iwata, S Thomas McCormick, and Maiko Shigeno. "A Faster Algorithm for Minimum Cost Submodular Flows." In: SODA. 1998, pp. 167-174.
[124] Satoru Iwata, S Thomas McCormick, and Maiko Shigeno. "A strongly polynomial cut canceling algorithm for the submodular flow problem". In: Integer Programming and Combinatorial Optimization. Springer, 1999, pp. 259-272.
[125] Satoru Iwata and James B Orlin. "A simple combinatorial algorithm for submodular function minimization". In: Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms. Society for Industrial and Applied Mathematics. 2009, pp. 1230-1237.
[126] Rahul Jain and Penghui Yao. "A parallel approximation algorithm for positive semidefinite programming". In: Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on. IEEE. 2011, pp. 463-471.
[127] Klaus Jansen. "Approximate strong separation with application in fractional graph coloring and preemptive scheduling". In: Theoretical Computer Science 302.1 (2003), pp. 239-256.
[128] Fritz John. "Extremum problems with inequalities as side conditions". In: Studies and Essays presented to R. Courant on his 60th Birthday (1948).
[129] R. Johnson and T. Zhang. "Accelerating Stochastic Gradient Descent using Predictive Variance Reduction". In: Advances in Neural Information Processing Systems (NIPS). 2013.
[130] William Johnson and Joram Lindenstrauss. "Extensions of Lipschitz mappings into a Hilbert space". In: Conference on modern analysis and probability. Vol. 26. Contemporary Mathematics. 1984, pp. 189-206.
[131] R. Kannan, L. Lovász, and M. Simonovits. "Random walks and an $O^{*}\left(n^{5}\right)$ volume algorithm for convex bodies". In: Random Structures and Algorithms 11 (1997), pp. 1-50.
[132] R. Kannan and H. Narayanan. "Random walks on polytopes and an affine interior point method for linear programming". In: STOC. 2009, pp. 561-570.
[133] Ravindran Kannan and Hariharan Narayanan. "Random walks on polytopes and an affine interior point method for linear programming". In: Mathematics of Operations Research 37.1 (2012), pp. 1-20.
[134] David R Karger. "Better random sampling algorithms for flows in undirected graphs". In: Proceedings of the ninth annual ACM-SIAM symposium on Discrete algorithms. Society for Industrial and Applied Mathematics. 1998, pp. 490-499.
[135] David Karger and Matthew Levine. "Random sampling in residual graphs". In: Proceedings of the thiry-fourth annual ACM symposium on Theory of computing. ACM. 2002, pp. 63-66.
[136] Narendra Karmarkar. "A new polynomial-time algorithm for linear programming". In: Proceedings of the sixteenth annual ACM symposium on Theory of computing. ACM. 1984, pp. 302311.
[137] Narendra Karmarkar. "Riemannian geometry underlying interior-point methods for linear programming". In: Contemporary Mathematics 114 (1990), pp. 51-75.
[138] Richard M Karp and Christos H Papadimitriou. "On linear characterizations of combinatorial optimization problems". In: SIAM Journal on Computing 11.4 (1982), pp. 620-632.
[139] Jonathan A. Kelner and Petar Maymounkov. "Electric routing and concurrent flow cutting". In: CoRR abs/0909.2859 (2009).
[140] Jonathan A. Kelner, Gary L. Miller, and Richard Peng. "Faster approximate multicommodity flow using quadratically coupled flows". In: Proceedings of the 44th symposium on Theory of Computing. STOC '12. New York, New York, USA: ACM, 2012, pp. 1-18.
[141] Jonathan A. Kelner et al. "A simple, combinatorial algorithm for solving SDD systems in nearly-linear time". In: Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC). 2013, pp. 911-920.
[142] Leonid G Khachiyan. "Polynomial algorithms in linear programming". In: USSR Computational Mathematics and Mathematical Physics 20.1 (1980), pp. 53-72.
[143] Leonid G Khachiyan. "Rounding of polytopes in the real number model of computation". In: Mathematics of Operations Research 21.2 (1996), pp. 307-320.
[144] LG Khachiyan, SP Tarasov, and II Erlikh. "The method of inscribed ellipsoids". In: Soviet Math. Dokl. Vol. 37. 1. 1988, pp. 226-230.
[145] Adam R Klivans and Daniel Spielman. "Randomness efficient identity testing of multivariate polynomials". In: Proceedings of the thirty-third annual ACM symposium on Theory of computing. ACM. 2001, pp. 216-223.
[146] Ioannis Koutis, Gary L. Miller, and Richard Peng. "A Nearly-m log n Time Solver for SDD Linear Systems". In: Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on. Oct. 2011, pp. 590-598.
[147] Ioannis Koutis, Gary L. Miller, and Richard Peng. "Approaching Optimality for Solving SDD Linear Systems". In: Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS). 2010, pp. 235-244.
[148] Jakob Krarup and Steven Vajda. "On Torricelli's geometrical solution to a problem of Fermat". In: IMA Journal of Management Mathematics 8.3 (1997), pp. 215-224.
[149] Andreas Krause. http://submodularity.org/.
[150] Dilip Krishnan, Raanan Fattal, and Richard Szeliski. "Efficient preconditioning of laplacian matrices for computer graphics". In: ACM Transactions on Graphics (TOG) 32.4 (2013), p. 142.
[151] Kartik Krishnan and John E Mitchell. "A unifying framework for several cutting plane methods for semidefinite programming". In: Optimization methods and software 21.1 (2006), pp. 5774.
[152] Kartik Krishnan and John E Mitchell. "Properties of a cutting plane method for semidefinite programming". In: Pacific Journal of Optimization 8.4 (2012), pp. 779-802.
[153] Kartik Krishnan and Tamás Terlaky. "Interior point and semidefinite approaches in combinatorial optimization". In: Graph theory and combinatorial optimization. Springer, 2005, pp. 101-157.
[154] Richard A. Kronmal and Arthur V. Peterson. "The Alias and Alias-rejection-mixture Methods for Generating Random Variables from Probability Distributions". In: Proceedings of the 11th Conference on Winter Simulation - Volume 1. WSC '79. San Diego, California, USA: IEEE Press, 1979, pp. 269-280.
[155] HaroldW. Kuhn. "A note on Fermat's problem". English. In: Mathematical Programming 4.1 (1973), pp. 98-107.
[156] Sanjiv Kumar, Mehryar Mohri, and Ameet Talwalkar. "Sampling Methods for the Nystrom Method". In: Journal of Machine Learning Research 13.1 (2012), pp. 981-1006.
[157] Eugene L Lawler. "Matroid intersection algorithms". In: Mathematical programming 9.1 (1975), pp. 31-56.
[158] Gregory Lawler and Hariharan Narayanan. "Mixing times and lp bounds for Oblivious Routing". In: Workshop on Analytic Algorithms and Combinatorics, (ANALCO 09). 2009.
[159] N. Le Roux, M. Schmidt, and F. Bach. "A Stochastic Gradient Method with an Exponential Convergence Rate for Strongly-Convex Optimization with Finite Training Sets". In: Advances in Neural Information Processing Systems (NIPS). 2012.
[160] Yin Tat Lee, Richard Peng, and Daniel A Spielman. "Sparsified Cholesky Solvers for SDD linear systems". In: arXiv preprint arXiv:1506.08204 (2015).
[161] Yin Tat Lee, Satish Rao, and Nikhil Srivastava. "A New Approach to Computing Maximum Flows using Electrical Flows". In: Proceedings of the 45 th symposium on Theory of Computing - STOC '13 (2013).
[162] Yin Tat Lee and Aaron Sidford. "Efficient Accelerated Coordinate Descent Methods and Faster Algorithms for Solving Linear Systems." In: The 54th Annual Symposium on Foundations of Computer Science (FOCS). 2013.
[163] Yin Tat Lee and Aaron Sidford. "Efficient Accelerated Coordinate Descent Methods and Faster Algorithms for Solving Linear Systems". In: Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science. 2013.
[164] Yin Tat Lee and Aaron Sidford. "Efficient Inverse Maintenance and Faster Algorithms for Linear Programming". In: arXiv:1503.01752 (2015).
[165] Yin Tat Lee and Aaron Sidford. "Path Finding I: Solving Linear Programs with ${ }^{\sim}$ O (sqrt (rank)) Linear System Solves". In: arXiv preprint arXiv:1312.6677 (2013).
[166] Yin Tat Lee and Aaron Sidford. "Path Finding II: An<br>~ O (m sqrt (n)) Algorithm for the Minimum Cost Flow Problem". In: arXiv preprint arXiv:1312.6713 (2013).
[167] F. Thomson Leighton and Ankur Moitra. "Extensions and limits to vertex sparsification". In: Proceedings of the $42 n d$ ACM symposium on Theory of computing. STOC '10. New York, NY, USA: ACM, 2010, pp. 47-56.
[168] A. Yu Levin. "On an algorithm for the minimization of convex functions". In: Soviet Math. Doklady (1965).
[169] D Lewis. "Finite dimensional subspaces of $L_{-}\{p\}$ ". In: Studia Mathematica 63.2 (1978), pp. 207-212.
[170] Mu Li, Gary L Miller, and Richard Peng. "Iterative row sampling". In: Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS). 2013, pp. 127-136.
[171] E Lieb and W Thirring. "Inequalities for the moments of the eigenvalues of the Schrödinger equation and their relation to Sobolev inequalities". In: Studies in Mathematical Physics: Essays in honor of Valentine Bargman, Lieb, E., Simon, B., Wightman, AS (eds.) (1976), pp. 269-303.
[172] Hendrik P. Lopuhaa and Peter J. Rousseeuw. "Breakdown Points of Affine Equivariant Estimators of Multivariate Location and Covariance Matrices". In: Ann. Statist. 19.1 (Mar. 1991), pp. 229-248.
[173] L. Lovász. "Hit-and-Run mixes fast". In: Math. Prog 86 (1998), pp. 443-461.
[174] L. Lovász. "How to compute the volume?" In: Jber. d. Dt. Math.-Verein, Jubiläumstagung 1990 (1990), pp. 138-151.
[175] L. Lovász and M. Simonovits. "Mixing rate of Markov chains, an isoperimetric inequality, and computing the volume". In: ROCS. 1990, pp. 482-491.
[176] L. Lovász and M. Simonovits. "On the randomized complexity of volume and diameter". In: Proc. 33rd IEEE Annual Symp. on Found. of Comp. Sci. 1992, pp. 482-491.
[177] L. Lovász and M. Simonovits. "Random walks in a convex body and an improved volume algorithm". In: Random Structures and Alg. Vol. 4. 1993, pp. 359-412.
[178] L. Lovász and S. Vempala. "Fast Algorithms for Logconcave Functions: sampling, Rounding, Integration and Optimization". In: FOCS. 2006, pp. 57-68.
[179] L. Lovász and S. Vempala. "Hit-and-run from a corner". In: SIAM J. Computing 35 (4 2006), pp. 985-1005.
[180] László Lovász and Santosh Vempala. "Simulated annealing in convex bodies and an $O^{*}\left(n^{4}\right)$ volume algorithm". In: J. Comput. Syst. Sci. 72.2 (2006), pp. 392-417.
[181] A. Lubotzky, R. Phillips, and P. Sarnak. "Ramanujan Graphs". In: Combinatorica 8.3 (1988), pp. 261-277.
[182] Aleksander Madry. "Fast approximation algorithms for cut-based problems in undirected graphs". In: Proceedings of the 51st Annual Symposium on Foundations of Computer Science. 2010.
[183] Aleksander Madry. "Navigating Central Path with Electrical Flows: from Flows to Matchings, and Back". In: Proceedings of the 54th Annual Symposium on Foundations of Computer Science. 2013.
[184] Michael W. Mahoney. "Randomized Algorithms for Matrices and Data". In: Foundations and Trends in Machine Learning 3.2 (2011), pp. 123-224.
[185] Michael W Mahoney and Xiangrui Meng. "Low-distortion subspace embeddings in inputsparsity time and applications to robust linear regression". In: Proceedings of the 45 th Annual ACM Symposium on Theory of Computing (STOC). 2013, pp. 91-100.
[186] Stefano Marchesini, Yu-Chao Tu, and Hau-tieng Wu. "Alternating projection, ptychographic imaging and phase synchronization". In: arXiv preprint arXiv:1402.0550 (2014).
[187] Adam W Marcus, Nikhil Srivastava, and Daniel A Spielman. "Interlacing Families IV: Bipartite Ramanujan Graphs of All Sizes". In: arXiv preprint arXiv:1505.08010 (2015). to appear in FOCS 2015.
[188] G. A. Margulis. "Explicit group theoretical constructions of combinatorial schemes and their application to the design of expanders and concentrators". In: Problems of Information Transmission 24.1 (July 1988), pp. 39-46.
[189] S McCormick. Submodular Function Minimization. 2013.
[190] S Thomas and McCormick. "Canceling most helpful total submodular cuts for submodular flow." In: IPCO. 1993, pp. 343-353.
[191] J. A. Meijerink and H. A. van der Vorst. "An Iterative Solution Method for Linear Systems of Which the Coefficient Matrix is a Symmetric M-Matrix". In: Mathematics of Computation 31.137 (1977), pp. 148-162.
[192] Xiangrui Meng, Michael A. Saunders, and Michael W. Mahoney. "LSRN: A Parallel Iterative Solver for Strongly Over- or Under-Determined Systems". In: SIAM J. Scientific Computing 36.2 (2014).
[193] Carl D Meyer Jr. "Generalized inversion of modified matrices". In: SIAM J. Appl. Math 24.3 (1973), pp. 315-323.
[194] Gary L. Miller and Richard Peng. "Approximate Maximum Flow on Separable Undirected Graphs". In: Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms. New Orleans, Louisiana, USA: SIAM, 2013, pp. 1151-1170.
[195] Mehryar Mohri and Ameet Talwalkar. "Can matrix coherence be efficiently and accurately estimated?" In: Proceedings of the 14th International Conference on Artificial Intelligence and Statistics (AISTATS). 2011, pp. 534-542.
[196] Renato DC Monteiro. "First-and second-order methods for semidefinite programming". In: Mathematical Programming 97.1-2 (2003), pp. 209-244.
[197] Kazuhide Nakata et al. "Exploiting sparsity in semidefinite programming via matrix completion II: Implementation and numerical results". In: Mathematical Programming 95.2 (2003), pp. 303-327.
[198] Hariharan Narayanan. "Randomized interior point methods for sampling and optimization". In: Annals of Applied Probability 26.1 (Feb. 2016), pp. 597-641.
[199] Hariharan Narayanan and Alexander Rakhlin. "Random walk approach to regret minimization". In: Advances in Neural Information Processing Systems. 2010, pp. 1777-1785.
[200] Jelani Nelson and Huy L. Nguyen. "OSNAP: Faster Numerical Linear Algebra Algorithms via Sparser Subspace Embeddings". In: Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS). 2013, pp. 117-126.
[201] A. Nemirovski and D. Yudin. Problem Complexity and Method Efficiency in Optimization. Wiley Interscience, 1983.
[202] Arkadi Nemirovski. "Efficient Methods in Convex Programming". In: (1994).
[203] Yu Nesterov. "A method of solving a convex programming problem with convergence rate O (1/k ${ }^{2}$ )". In: Soviet Mathematics Doklady 27.2 (1983), pp. 372-376.
[204] Yu Nesterov. "Complexity estimates of some cutting plane methods based on the analytic barrier". In: Mathematical Programming 69.1-3 (1995), pp. 149-176.
[205] Yu Nesterov. "Efficiency of coordinate descent methods on huge-scale optimization problems". In: Core discussion papers 2 (2010), p. 2010.
[206] Yu Nesterov. Introductory Lectures on Convex Optimization: A Basic Course. Vol. I. 2003.
[207] Yu Nesterov. "Rounding of convex sets and efficient gradient methods for linear programming problems". In: Available at SSRN 965658 (2004).
[208] Yu Nesterov. "Smooth minimization of non-smooth functions". In: Mathematical Programming 103.1 (2005), pp. 127-152.
[209] Yu E Nesterov and Michael J Todd. "Self-scaled barriers and interior-point methods for convex programming". In: Mathematics of Operations research 22.1 (1997), pp. 1-42.
[210] Yu Nesterov and Arkadi Nemirovski. "Conic formulation of a convex programming problem and duality". In: Optimization Methods and Software 1.2 (1992), pp. 95-115.
[211] Yu Nesterov and Arkadi Nemirovski. Interior-Point Polynomial Algorithms in Convex Programming. Studies in Applied Mathematics. SIAM, 1994.
[212] Yu Nesterov and Arkadi Nemirovski. "Primal Central Paths and Riemannian Distances for Convex Sets". In: Foundations of Computational Mathematics 8.5 (2008), pp. 533-560.
[213] Yu Nesterov and Arkadi Nemirovski. Self-concordant functions and polynomial-time methods in convex programming. USSR Academy of Sciences, Central Economic \& Mathematic Institute, 1989.
[214] Yurii E Nesterov, Michael J Todd, et al. "On the Riemannian geometry defined by selfconcordant barriers and interior-point methods". In: Foundations of Computational Mathematics 2.4 (2002), pp. 333-361.
[215] James B Orlin. "A faster strongly polynomial minimum cost flow algorithm". In: Operations research 41.2 (1993), pp. 338-350.
[216] James B Orlin. "A faster strongly polynomial time algorithm for submodular function minimization". In: Mathematical Programming 118.2 (2009), pp. 237-251.
[217] James B Orlin. "Genuinely Polynominal Simplex and Non-Simplex Algorithms for the Minimum Cost Flow Problem". In: (1984).
[218] James B Orlin. "Max flows in O (nm) time, or better". In: Proceedings of the 45th annual ACM symposium on Symposium on theory of computing. ACM. 2013, pp. 765-774.
[219] James B Orlin, John VandeVate, et al. "On a" primal" matroid intersection algorithm". In: (1983).
[220] Lawrence M. Ostresh. "On the Convergence of a Class of Iterative Methods for Solving the Weber Location Problem". In: Operations Research 26.4 (1978), pp. 597-609.
[221] Onur Ozyesil, Amit Singer, and Ronen Basri. "Stable camera motion estimation using convex programming". In: SIAM Journal on Imaging Sciences 8.2 (2015), pp. 1220-1262.
[222] Pablo A. Parrilo and Bernd Sturmfels. "Minimizing Polynomial Functions". In: DIMACS Workshop on Algorithmic and Quantitative Aspects of Real Algebraic Geometry in Mathematics and Computer Science, March 12-16, 2001, DIMACS Center, Rutgers University, Piscataway, NJ, USA. 2001, pp. 83-100.
[223] Richard Peng and Daniel A Spielman. "An Efficient Parallel Solver for SDD Linear Systems". In: arXiv preprint arXiv:1311.3286 (2013).
[224] Frank Plastria and Mohamed Elosmani. "On the convergence of the Weiszfeld algorithm for continuous single facility location allocation problems". English. In: TOP 16.2 (2008), pp. 388406.
[225] Harald Räcke. "Optimal hierarchical decompositions for congestion minimization in networks". In: Proceedings of the 40th annual ACM symposium on Theory of computing. STOC '08. Victoria, British Columbia, Canada: ACM, 2008, pp. 255-264.
[226] Srinivasan Ramaswamy and John E Mitchell. "A long step cutting plane algorithm that uses the volumetric barrier". In: Department of Mathematical Science, RPI, Troy, NY (1995).
[227] James Renegar. "A polynomial-time algorithm, based on Newton's method, for linear programming". In: Mathematical Programming 40.1-3 (1988), pp. 59-93.
[228] Vladimir Rokhlin and Mark Tygert. "A fast randomized algorithm for overdetermined linear least-squares regression". In: Proc. Natl. Acad. Sci. USA 1005.36 (2008), pp. 13212-7.
[229] Tamas Sarlos. "Improved approximation algorithms for large matrices via random projections". In: Proceedings of the 47 th Annual IEEE Symposium on Foundations of Computer Science (FOCS). 2006, pp. 143-152.
[230] M Schmidt. "minFunc: unconstrained differentiable multivariate optimization in Matlab". In: URL http://www. di. ens. fr/mschmidt/Software/minFunc. html (2012).
[231] Alexander Schrijver. "A combinatorial algorithm minimizing submodular functions in strongly polynomial time". In: Journal of Combinatorial Theory, Series B 80.2 (2000), pp. 346-355.
[232] Alexander Schrijver. Combinatorial optimization: polyhedra and efficiency. Vol. 24. Springer, 2003.
[233] Jonah Sherman. "Nearly Maximum Flows in Nearly Linear Time". In: Proceedings of the 54th Annual Symposium on Foundations of Computer Science. 2013.
[234] Maiko Shigeno and Satoru Iwata. "A dual approximation approach to weighted matroid intersection". In: Operations research letters 18.3 (1995), pp. 153-156.
[235] Yoel Shkolnisky and Amit Singer. "Viewing direction estimation in cryo-EM using synchronization". In: SIAM journal on imaging sciences 5.3 (2012), pp. 1088-1110.
[236] Naum Z Shor. "Cut-off method with space extension in convex programming problems". In: Cybernetics and systems analysis 13.1 (1977), pp. 94-96.
[237] Amit Singer and Yoel Shkolnisky. "Three-dimensional structure determination from common lines in cryo-EM by eigenvectors and semidefinite programming". In: SIAM journal on imaging sciences 4.2 (2011), pp. 543-572.
[238] Amit Singer and H-T Wu. "Vector diffusion maps and the connection Laplacian". In: Communications on pure and applied mathematics 65.8 (2012).
[239] Maurice Sion. "On general minimax theorems". In: Pacific J. Math 8.1 (1958), pp. 171-176.
[240] Daniel D. Sleator and Robert Endre Tarjan. "A data structure for dynamic trees". In: Proceedings of the thirteenth annual ACM symposium on Theory of computing. STOC '81. Milwaukee, Wisconsin, USA: ACM, 1981, pp. 114-122.
[241] R.L. Smith. "Efficient Monte-Carlo procedures for generating points uniformly distributed over bounded regions". In: Operations Res. 32 (1984), pp. 1296-1308.
[242] Daniel A. Spielman and Nikhil Srivastava. "Graph Sparsification by Effective Resistances". In: SIAM Journal on Computing 40.6 (2011), pp. 1913-1926.
[243] Daniel A. Spielman and Shang-Hua Teng. "Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems". In: Proceedings of the thirty-sixth annual ACM symposium on Theory of computing. ACM. 2004, pp. 81-90.
[244] Daniel A. Spielman and Shang-Hua Teng. "Spectral Sparsification of Graphs". In: CoRR abs/0808.4134 (2008).
[245] Michel Talagrand. "Embedding subspaces of L1 into L1". In: Proceedings of the American Mathematical Society 108.2 (1990), pp. 363-369.
[246] Michel Talagrand. "Embedding Subspaces of Lp in Lp". In: Geometric Aspects of Functional Analysis. Springer, 1995, pp. 311-326.
[247] Éva Tardos. "A strongly polynomial minimum cost circulation algorithm". In: Combinatorica 5.3 (1985), pp. 247-255.
[248] Michael J Todd. "Semidefinite optimization". In: Acta Numerica 200110 (2001), pp. 515-560.
[249] Nobuaki Tomizawa and Masao Iri. "ALGORITHM FOR DETERMINING RANK OF A TRIPLE MATRIX PRODUCT AXB WITH APPLICATION TO PROBLEM OF DISCERNING EXISTENCE OF UNIQUE SOLUTION IN A NETWORK". In: ELECTRONICS \& COMMUNICATIONS IN JAPAN 57.11 (1974), pp. 50-57.
[250] Burt Totaro. "The curvature of a Hessian metric". In: International Journal of Mathematics 15.04 (2004), pp. 369-391.
[251] Joel A. Tropp. "User-Friendly Tail Bounds for Sums of Random Matrices". In: Foundations of Computational Mathematics 12.4 (2012), pp. 389-434.
[252] Ulrich Trottenberg, Cornelius W Oosterlee, and Anton Schuller. Multigrid. Academic press, 2000.
[253] Pravin M. Vaidya. "A New Algorithm for Minimizing Convex Functions over Convex Sets". In: FOCS. 1989, pp. 338-343.
[254] Pravin M Vaidya. "An algorithm for linear programming which requires $\mathrm{O}(((\mathrm{m}+\mathrm{n}) \mathrm{n} 2+(\mathrm{m}+$ n) 1.5 n) L) arithmetic operations". In: Mathematical Programming 47.1-3 (1990), pp. 175201.
[255] Pravin M. Vaidya. "Reducing the Parallel Complexity of Certain Linear Programming Problems". In: FOCS. 1990, pp. 583-589.
[256] Pravin M. Vaidya. "Solving linear equations with symmetric diagonally dominant matrices by constructing good preconditioners." Unpublished manuscript UIUC 1990. A talk based on the manuscript was presented at the IMA Workshop on Graph Theory and Sparse Matrix Computation, October 1991, Minneapolis. 1990.
[257] Pravin M Vaidya. "Speeding-up linear programming using fast matrix multiplication". In: Foundations of Computer Science, 1989., 30th Annual Symposium on. IEEE. 1989, pp. 332337.
[258] Pravin M Vaidya and David S Atkinson. "A technique for bounding the number of iterations in path following algorithms". In: Complexity in Numerical Optimization (1993), pp. 462-489.
[259] Lieven Vandenberghe and Stephen Boyd. "Semidefinite programming". In: SIAM review 38.1 (1996), pp. 49-95.
[260] Yehuda Vardi and Cun-Hui Zhang. "The multivariate L1-median and associated data depth". In: Proceedings of the National Academy of Sciences 97.4 (2000), pp. 1423-1426.
[261] S. Vempala. "Geometric Random Walks: A Survey". In: MSRI Combinatorial and Computational Geometry 52 (2005), pp. 573-612.
[262] Santosh S Vempala. "Recent progress and open problems in algorithmic convex geometry". In: LIPIcs-Leibniz International Proceedings in Informatics. Vol. 8. Schloss Dagstuhl-LeibnizZentrum fuer Informatik. 2010.
[263] Vincenzo Viviani. "De Maximis et Minimis Geometrica Divinatio Liber 2". In: De Maximis et Minimis Geometrica Divinatio (1659).
[264] Jens Vygen. "A note on Schrijver's submodular function minimization algorithm". In: Journal of Combinatorial Theory, Series B 88.2 (2003), pp. 399-402.
[265] S. Fujishige W. Cui. "A primal algorithm for the submodular flow problem with minimum mean cycle selection". In: Journal of the Operations Research Society of Japan (1988).
[266] C Wallacher and Uwe T Zimmermann. "A polynomial cycle canceling algorithm for submodular flows". In: Mathematical programming 86.1 (1999), pp. 1-15.
[267] Alfred Weber. The Theory of the Location of Industries. Aber den I der Industrien. Chicago University Press, 1909.
[268] E. Weiszfeld. "Sur le point pour lequel la somme des distances de n points donnes est minimum". In: Tohoku Mathematical Journal (1937), pp. 355-386.
[269] Virginia Vassilevska Williams. "Multiplying matrices faster than Coppersmith-Winograd". In: Proceedings of the forty-fourth annual ACM symposium on Theory of computing. ACM. 2012, pp. 887-898.
[270] Przemyslaw Wojtaszczyk. Banach spaces for analysts. Vol. 25. Cambridge University Press, 1996.
[271] Guoliang Xue and Yinyu Ye. "An efficient algorithm for minimizing a sum of Euclidean norms with applications". In: SIAM Journal on Optimization 7 (1997), pp. 1017-1036.
[272] Yinyu Ye. "Complexity analysis of the analytic center cutting plane method that uses multiple cuts". In: Mathematical Programming 78.1 (1996), pp. 85-104.
[273] Yinyu Ye. Interior point algorithms: theory and analysis. Vol. 44. John Wiley \& Sons, 2011.
[274] David B Yudin and Arkadii S Nemirovski. "Evaluation of the information complexity of mathematical programming problems". In: Ekonomika i Matematicheskie Metody 12 (1976), pp. 128142.
[275] Zhizhen Zhao and Amit Singer. "Rotationally invariant image representation for viewing direction classification in cryo-EM". In: Journal of structural biology 186.1 (2014), pp. 153-166.
[276] U Zimmermann. "Minimization on submodular flows". In: Discrete Applied Mathematics 4.4 (1982), pp. 303-323.
[277] Uwe Zimmermann. "Negative circuits for flows and submodular flows". In: Discrete applied mathematics 36.2 (1992), pp. 179-189.
[278] Anastasios Zouzias. "A Matrix Hyperbolic Cosine Algorithm and Applications". In: Proceedings of the 39th International Colloquium on Automata, Languages and Programming (ICALP). 2012, pp. 846-858.


[^0]:    ${ }^{1}$ Throughout this thesis, our focus is on algorithms for polynomial time solvable convex optimization problems given access to a linear separation oracle. Our usage of the term cutting plane methods, should not be confused with work on integer programming.
    ${ }^{2}$ We also note that independently, Jonah Sherman produced an almost-linear time algorithm for maximum flow.

[^1]:    ${ }^{3}$ In this thesis, nearly all algorithms are randomized and nearly all time are expected running time.

[^2]:    ${ }^{1}$ We also show how to compute a $(1+\varepsilon)$-approximate geometric median in sub-linear time $O\left(d \varepsilon^{-2}\right)$.

[^3]:    ${ }^{2}$ We also note that independently, Jonah Sherman produced an almost-linear time algorithm for maximum flow.

[^4]:    ${ }^{1}$ We decribe exactly how each $\tilde{\sigma}_{i}$ is computed when we prove Theorem 3.1.1 in Section 3.4.

[^5]:    ${ }^{2}$ A simple method for handling even larger values of $n$ is outlined in [170].

[^6]:    ${ }^{1}$ Throughout this thesis $L$ denotes the standard "bit complexity" of the linear program, a quantity less than the number of bits needed to represent (6.2). For integral $\boldsymbol{A}, \boldsymbol{b}$, and $\boldsymbol{c}$ this quantity is often defined as $L \stackrel{\text { def }}{=} \log \left(1+d_{m a x}\right)+$ $\log \left(1+\max \left\{\|\boldsymbol{c}\|_{\infty},\|\boldsymbol{b}\|_{\infty}\right\}\right)$ where $d_{\max }$ is the largest absolute value of the determinant of a square sub-matrix of $\boldsymbol{A}[136]$.
    ${ }^{2}$ In this thesis we are primarily concerned with "weakly" polynomial time algorithms. The current fastest "strongly" polynomial running time for solving this problem is $O(n m)$ [218].

[^7]:    ${ }^{3}$ This expression is the dual of a linear program written in standard form. It is well known that all linear programs can be written as (6.2). Note that this notation of $m$ and $n$ differs from that in some papers. Here $m$ denotes the number of constraints and $n$ denotes the number of variables. To avoid confusion we state many of our results in terms of $\sqrt{\operatorname{rank}(\boldsymbol{A})}$ instead of $\sqrt{n}$.

[^8]:    ${ }^{4}$ Typically (6.9) is written as $\boldsymbol{A} \boldsymbol{x}=\boldsymbol{b}$ rather than $\boldsymbol{A}^{T} \boldsymbol{x}=\boldsymbol{b}$. We chose this formulation to be consistent with the previous section and to be consistent with the standard use of $n$ to denote the number of vertices and $m$ to denote the number of edges in a graph in the linear program formulation of flow problems.

[^9]:    ${ }^{5}$ The "trigonometric barrier" we use arises as the (unique) solution of the ODE $\phi^{\prime \prime \prime}=2\left(\phi^{\prime \prime}\right)^{3 / 2}$ such that the function value goes to infinity up at $u_{i}$ and $l_{i}$.

[^10]:    ${ }^{6}$ Recall that $\boldsymbol{w} \phi^{\prime}(\boldsymbol{x})$ denotes the entry-wise multiplication of the vectors $\boldsymbol{w}$ and $\boldsymbol{\phi}^{\prime}(\boldsymbol{x})$.

[^11]:    ${ }^{7}$ Note that early we assumed that $\beta<1$ and here we use much larger values of $\beta$. However, this bound on $\beta$ was primarily to assist in bounding $c_{1}$ and does not affect this proof.

[^12]:    ${ }^{8}$ There are some boundary cases we ignored for simplicity.

[^13]:    ${ }^{1}$ If $z$ is the total number of nonzero entries in the coordinates of the $a^{(i)}$ then a careful analysis of our algorithm improves our running time to $O\left(z \log ^{3} \frac{n}{\varepsilon}\right)$.

[^14]:    ${ }^{2}$ The result of [50] was stated in more general terms than given here. However, it easy to formulate the geometric median problem in their model.

[^15]:    ${ }^{3}$ It is unclear how to extend our proof for the simpler function: $\sum_{i \in[n]} \sqrt{1+t^{2}\left\|x-a^{(i)}\right\|_{2}^{2}}$.

[^16]:    ${ }^{4}$ By $\boldsymbol{A} \preceq \boldsymbol{B}$ on $S$ we mean that for all $x \in S$ we have $x^{\top} \boldsymbol{A} x \leq x^{\top} \boldsymbol{B} x$. The meaning of $\boldsymbol{A} \succeq \boldsymbol{B}$ on $S$ is analagous.

[^17]:    ${ }^{1}$ We have made no effort on improving this constant and we believe it can be improved to less than 300 using techniques in [11, 14].

[^18]:    ${ }^{1}$ Without this, the running time of interior point method depends on the bit complexity of the linear programs.

[^19]:    ${ }^{1}$ Minus signs is needed because we express our inequalities as e.g. $\boldsymbol{h}^{T} \boldsymbol{x} \leq \boldsymbol{h}^{T} \bar{x}_{h}$ whereas in Theorem 8.3.26, $\boldsymbol{a}_{i}^{T} \boldsymbol{x} \geq b_{i}$ is used. We apologize for the inconvenience.

[^20]:    ${ }^{2}$ More precisely, $B_{k}=\emptyset$ for $k>\lceil\log n\rceil$.

[^21]:    ${ }^{3}$ In the terminology of Chapter 8 , these weighted influences are the leverage scores.

[^22]:    ${ }^{1}$ This idea and analysis seems to be implicit in other work, e.g., [205] However, we could not find a clean statement like the one we need in the literature, and we have not seen it previously applied in similar settings. We believe that it will find further applications, so we state it in fairly general terms before specializing to what we need for flow problems.

[^23]:    ${ }^{2}$ Note that here and in the rest of the chapter we will focus our analysis with congestion with respect to the norm $\|\cdot\|_{\infty}$ and we will look at oblivious routing strategies that are competitive with respect to this norm. However, many of the results present are easily generalizable to other norms. These generalizations are outside the scope of this thesis.

[^24]:    ${ }^{3}$ The structure of this specific proof was modeled after a proof in [205] for a slightly different problem.

[^25]:    ${ }^{4}$ Note that the oblivous routing strategies considered in [139] [158] [225] are all linear oblivious routing strategies.

[^26]:    ${ }^{5}$ Note that our flow sparsifiers aim to reduce the number of edges, and are different from the flow sparsifiers of Leighton and Moitra [167], which work in a different setting and reduce the number of vertices.

[^27]:    ${ }^{6}$ The spectral sparsification result actually yields the stronger spectral approximation guarantee, but for our purposes the cut guarantee suffices.

[^28]:    ${ }^{7}$ This is a restatement of the $H(T, F)$ graphs in [182].

[^29]:    ${ }^{8}$ This in an adaptation of Corollary 5.6 in [182]

[^30]:    ${ }^{9}$ Note that the constant of 4 below is improved to 3 in [182].

[^31]:    ${ }^{10}$ This is an idea that has been applied previously to solve linear programming problems [207].

[^32]:    ${ }^{1}$ The finite condition number is the ratio of the largest singular value to the smallest nonzero singular value.

[^33]:    ${ }^{2}$ In Chapter 3, we sample rows of matrices. Here, we sample columns instead in order to use a more natural set of notations.

[^34]:    ${ }^{1}$ Coincidentally, when we use the best known self-concordant barrier, the canonical barrier, our SDE becomes a Brownian motion with drift on an Einstein manifold. This is similar to how physical particles mix under general relativity, a sampling algorithm that has been executed for over 10 billion years!

