

Flows, Submodularity, Sparsity, and Beyond: Continuous Optimization Insights into Discrete Problems

by

Kyriakos Axiotis

Diploma, National Technical University of Athens (2016)
S.M., Massachusetts Institute of Technology (2019)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2022

© Massachusetts Institute of Technology 2022. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
August 25, 2022

Certified by.....
Aleksander Mądry
Cadence Design Systems Professor of Computing
Thesis Supervisor

Accepted by
Leslie A. Kolodziejcki
Professor of Electrical Engineering and Computer Science
Chair, Department Committee on Graduate Students

Flows, Submodularity, Sparsity, and Beyond: Continuous Optimization Insights into Discrete Problems

by
Kyriakos Axiotis

Submitted to the Department of Electrical Engineering and Computer Science
on August 25, 2022, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

In this thesis we build on connections between discrete and continuous optimization. In the first part of the thesis we propose faster second-order convex optimization algorithms for classical graph algorithmic problems. Our main contribution is to show that the runtime of interior point methods is closely connected to spectral connectivity notions in the underlying graph, such as electrical conductance and effective resistance. We explore these connections along two orthogonal directions: Making manual interventions to the graph to improve connectivity, or keeping track of connectivity so as to make faster updates. These ideas lead to the first runtime improvement for the minimum cost flow problem in more than 10 years, as well as faster algorithms for problems like negative-weight shortest path and minimum cost perfect matching.

In the second part of the thesis, we investigate efficient optimization algorithms for problems relevant to machine learning that have some discrete element, such as sparse or low rank structure. We introduce a new technique, called adaptive regularization, which eliminates the sparsity performance degradation caused by ℓ_2 projections onto structured non-convex domains, like the set of sparse vectors or low rank matrices. This leads to improving the sparsity guarantee of one of the most well known sparse optimization algorithms, IHT.

Thesis Supervisor: Aleksander Mądry
Title: Cadence Design Systems Professor of Computing

Acknowledgments

This thesis is dedicated to the people who inspired, helped, and supported me through all these years.

To my advisor Aleksander Mađry for providing extremely valuable guidance and support for the completion of this thesis. He was one of the few people who always emphasized the importance of shooting high. This helped me concentrate on big questions and understand that meaningful research progress can only be achieved by resisting easy answers.

To my collaborator Adrian Vladu without whom this thesis would surely not have been possible. He taught me a lot of things and his inquisitive spirit led us to countless hours of learning and discovery that shaped my research philosophy.

To my collaborator Maxim Sviridenko who encouraged me to look at the more practical aspects of my work and taught me how important it is to have practical foundations for the research problems that I pick. It's hard to believe that I spent less than 3 months at Yahoo!

To my thesis committee members, Costis Daskalakis and Piotr Indyk.

To Christos Tzamos, who served as a researcher role model since my high school years and for the stimulating collaborations and discussions ever since.

To my Boston friends Christos, Vasso, Themis, Dimitris, Shibani, Giorgos, Manolis, Katerina, Ilias, Konstantina, Sophie, Chara, Dimitris.

To Dimitris, Fotis, Alex, Chenyang, Giorgos, Christos, Panagiotis, with whom we spent a unique semester at Berkeley.

To my lab and office mates Michael, Dimitris, Alex, Arturs, Nishanth, Tselil, Gautam, Andrew, Logan, Alex, Brynmor, Shibani, Quanquan.

To Dimitris Fotakis, who believed in me and fostered a unique research environment at CoReLab NTUA, and Stathis Zachos who was a source of energy and provided a sense of community.

To my friends in Athens.

To my parents Sakis and Eleni and my sister Victoria for providing an immensely supportive home environment full of love.

Finally, to my partner Eva, who has been a pillar of love and support throughout these years and someone I always know I can count on, as well as our dog Pavlito who is a constant source of joy.

Contents

1	Introduction	11
1.1	Structure of the Thesis	14
2	Background	17
2.1	Basic Notation	17
2.2	Convex Analysis	18
2.3	Graphs and Flows	20
2.4	Electrical Flows and Laplacians	21
I	Optimization on Graphs	24
3	Circulation Control for Faster Minimum Cost Flow in Unit-Capacity Graphs	25
3.1	Introduction	25
3.1.1	Our Contributions	25
3.1.2	Previous Work	26
3.1.3	Organization of the Chapter	27
3.2	Preliminaries	27
3.3	Minimum Cost Flow by Circulation Improvement	27
3.3.1	LP Formulation and Interior Point Method	27
3.3.2	Optimality and Duality Gap	30
3.3.3	Initialization	30
3.3.4	The Augmentation Procedure and Routing the Residual	30
3.3.5	Vanilla Interior Point Method	34
3.4	A Faster Algorithm for Minimum Cost Flow	35
3.4.1	Regularized Newton Step	35
3.4.2	Choice of Regularization Parameters	41
3.4.3	Weight Invariants	42
3.4.4	Executing the Interior Point Method	43
3.5	Repairing the Flow	52
3.6	Improving the Running Time to $m^{4/3+o(1)} \log W$	55
3.6.1	Bounding Congestion	59
3.6.2	Making Progress	62
3.6.3	Wrapping Up	63

4	Faster Sparse Minimum Cost Flow by Electrical Flow Localization	65
4.1	Introduction	65
4.1.1	Previous work	65
4.1.2	Our result	66
4.1.3	High level overview of our approach	66
4.2	Preliminaries	70
4.2.1	Random walks	70
4.3	Interior Point Method with Dynamic Data Structures	73
4.3.1	LP formulation and background	73
4.3.2	Making progress with approximate electrical flows	74
4.3.3	The LOCATOR data structure	75
4.3.4	The minimum cost flow algorithm	76
4.3.5	Proof of Theorem 4.1.1	79
4.4	An Efficient $(\alpha, \beta, \varepsilon)$ -LOCATOR	80
4.4.1	Moving demands to the sparsifier	81
4.4.2	ε -Important edges	81
4.4.3	Proving Lemma 4.3.8	83
4.5	The Demand Projection Data Structure	89
4.5.1	Estimating $\pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{q_s}{\sqrt{r}} \right)$	90
4.5.2	Estimating $\pi^C(\mathbf{1}_v)$	92
4.5.3	Proof of Lemma 4.4.14	93
5	Decomposable Submodular Function Minimization via Maximum Flow	95
5.1	Introduction	95
5.1.1	Our Results	96
5.1.2	Previous Work	97
5.2	Background and Preliminaries	98
5.2.1	Submodular Set Functions and Convex Analysis	98
5.2.2	Overview of Approach	100
5.3	Parametric Min s, t -Cut	101
5.4	Parametric Decomposable Submodular Minimization via Base Polytope Approximations	103
5.4.1	Algorithm Overview	103
5.4.2	Removing Assumptions	103
5.4.3	From Parametric Minimum Cut to Cut Base Polytope Optimization	103
5.4.4	Dual Progress Analysis in One Step	104
5.4.5	Main Theorem	105
II	Optimization Under Sparsity Constraints	106
6	Sparse Convex Optimization via Adaptively Regularized Hard Thresholding	107
6.1	Introduction	107
6.1.1	Sparse Optimization	107
6.1.2	Sparse Solution and Support Recovery	108
6.1.3	Our Work	109
6.1.4	Comparison to Previous Work	110
6.1.5	Runtime discussion	111
6.1.6	Naming Conventions	111

6.2	Preliminaries	111
6.2.1	Algorithms	113
6.3	Adaptively Regularized Hard Thresholding (ARHT)	115
6.3.1	Overview and Main Theorem	115
6.3.2	Bounding Type 2 Iterations	119
6.3.3	Corollaries	126
6.4	Analysis of Orthogonal Matching Pursuit with Replacement (OMPR)	127
6.4.1	Overview and Main Theorem	127
6.4.2	Progress Lemma and Theorem Proof	129
6.4.3	Corollaries of Theorem 6.4.2	134
6.5	Lower Bounds	135
6.5.1	$\Omega(s^* \kappa)$ Lower Bound due to [63]	135
6.5.2	$\Omega(s^* \kappa^2)$ Lower Bound for OMPR	136
6.6	Experiments	137
6.6.1	Overview	137
6.6.2	Setup Details	144
6.6.3	Implementation Details	144
7	Iterative Hard Thresholding with Adaptive Regularization: Sparser Solutions Without Sacrificing Runtime	147
7.1	Introduction	147
7.1.1	Iterative Hard Thresholding (IHT)	148
7.1.2	Reconciling Sparsity and Efficiency: Regularized IHT	148
7.1.3	Beyond Sparsity: Low Rank Optimization	149
7.1.4	Related Work	150
7.2	The Adaptive Regularization Method	150
7.2.1	Regularized IHT	150
7.2.2	Learning Weights	151
7.2.3	Beyond Sparsity: Learning Subspaces	151
7.3	Sparse Optimization Using Regularized IHT	152
7.4	Low Rank Optimization Using Regularized Local Search	154
7.5	Experiments	155
7.5.1	Real data	157
7.5.2	Synthetic data	157
8	Local Search Algorithms for Rank-Constrained Convex Optimization	161
8.1	Introduction	161
8.2	Algorithms & Theoretical Guarantees	162
8.2.1	Greedy	163
8.2.2	Local Search	163
8.2.3	Algorithmic adjustments	164
8.3	Optimization Applications	166
8.3.1	Low-rank approximation on observed set	166
8.3.2	Robust principal component analysis (RPCA)	166
8.4	Machine Learning Applications	167
8.4.1	Regularization techniques	167
8.4.2	Matrix Completion with Random Noise	168
8.4.3	Recommender Systems	168

8.5	Conclusions	170
9	Appendix	185
9.1	Appendix for Chapter 3	185
9.1.1	Initializing the Interior Point Method	185
9.1.2	Preconditioning Proof	186
9.1.3	Solving the Mixed Objective	191
9.1.4	Strengthening the Mixed Objective Solver	194
9.1.5	Deferred Proofs	197
9.2	Appendix for Chapter 4	207
9.2.1	Maintaining the Schur Complement	207
9.2.2	Auxiliary Lemmas	208
9.2.3	Deferred Proofs from Section 4.3	209
9.2.4	Deferred Proofs from Section 4.4	226
9.2.5	Deferred Proofs from Section 4.5	239
9.2.6	The CHECKER Data Structure	247
9.3	Appendix for Chapter 5	249
9.3.1	Approximating Submodular Set Functions with Graph Cuts	249
9.3.2	Parametric Submodular Minimization via Optimization on the Base Polytope	252
9.3.3	Parametric s - t Cuts	256
9.3.4	Removing Assumptions on F_i	264
9.3.5	Deferred Proofs	265
9.4	Appendix for Chapter 6	273
9.4.1	Deferred Proofs	273
9.5	Appendix for Chapter 7	274
9.5.1	Python implementation	274
9.5.2	Proof of Theorem 7.1.1	275
9.5.3	Proof of Lemma 7.3.1	278
9.5.4	Low Rank Minimization	283
9.5.5	Lower Bounds	296
9.6	Appendix for Chapter 8	297

Chapter 1

Introduction

Optimization is a central ingredient in many disciplines, including engineering, science, economics, and operations research, with tasks such as linear programming (LP) or semi-definite programming (SDP) being among the most commonly encountered. Owing to this broad applicability of optimization and especially due to the surge in machine learning (ML) research that is being observed in the last decade, there has been tremendous work in analyzing various aspects of optimization algorithms, as well as developing new ones to tackle emergent challenges. Even though optimization algorithms are generally built to operate in a *continuous* (high-dimensional) space, a lot of practical problems naturally induce *discrete* constraints, such as ones related to selecting a subset of variables, constraints coming from graph structure, or integrality restrictions. Incorporating discrete constraints generally introduces significant challenges, and can easily make a problem depart from the realm of polynomial solvability. Yet, it has been demonstrated time and time again that insights from the continuous world are often a game-changer for solving discrete problems. Classical examples include LP relaxations for approximation algorithms and primal-dual algorithms, and more recent ones are solving approximate maximum flow with gradient descent [152, 97], or k -server with mirror descent [27], among many others. In this thesis, we will introduce new connections between continuous and discrete optimization for a variety of problems, including graph algorithms, submodular optimization, and sparse optimization.

Part I: Optimization on Graphs

Traditional graph algorithms, which make use of combinatorial notions such as paths and cuts and have been extensively studied for decades, seem to be limited by long-standing efficiency barriers. On the other hand, continuous optimization has emerged as a versatile and powerful toolbox to help overcome these limitations. Casting discrete graph problems as continuous optimization tasks allows one to argue about the algorithm in the language of high-dimensional geometry, by applying the bedrock of optimization—gradient descent. This view reveals the geometric properties which affect the performance of the algorithm, effects which the algorithm designer can then try to mitigate by employing either graph-theoretic interventions, or optimization techniques like regularization. This conceptual approach has led to breakthroughs in many problems, including approximate maximum flow [39, 152, 97], minimum cost flow in dense graphs [106, 26, 25], maximum flow in sparse graphs [118, 117, 115, 95, 73], and minimum cost flow in sparse graphs [44, 42, 10, 11]. A catalyst to these results has been the groundbreaking work of Spielman and Teng [155] for solving Laplacian linear systems, which introduced a quiver of modern (spectral) graph-algorithmic ideas.

Interestingly, the flux of innovative ideas between the discrete and the continuous world goes both ways. In other words, not only does continuous optimization prove useful in the study of

graph algorithms, but conversely the study of graph algorithms also gives us deep insights into more general optimization tasks like linear programming, insights which might have been too abstract to conceive in the general setting. At the same time, graph structure is becoming increasingly relevant in machine learning.

Minimum Cost Flow Arguably, the problem that reigns in this area is the minimum cost flow problem, whose goal is to find the cheapest way to route a given demand in a capacitated graph. Its importance stems not only from the fact that it is ubiquitous in fields such as transportation and finance, but also because of its generality, as it contains problems like maximum flow, negative-weight shortest path, and minimum-cost bipartite matching as special cases. For many years, the state of the art runtime for sparse graphs was $\tilde{O}\left(m^{3/2} \log^{O(1)}(U + W)\right)$, where m is the number of edges in the graph and U, W are upper bounds on edge capacities and costs respectively, and it resulted from an appropriate instantiation of interior point method [44]. A natural question to ask then is:

Can we speed up interior point methods for minimum cost flow?

In order to understand what this question entails, it is important to have some insights into how interior point methods work. Path-following interior point methods involve an iterative process, each step of which involves solving a linear system of the form $\mathbf{A}^\top \mathbf{V} \mathbf{A} \mathbf{x} = \mathbf{b}$ (1), where \mathbf{A} is the constraint matrix of the LP and $\mathbf{V} = \text{diag}(\mathbf{v})$ is a diagonal matrix of positive weights, one per constraint. In particular, there are two ways in which the structure of the matrix $\mathbf{A}^\top \mathbf{V} \mathbf{A}$ influences the total runtime. One is, it affects how fast we are able to solve all the linear systems of the form (1) that arise throughout the algorithm, and the other one is how well-conditioned the solutions to these systems are, which in turn affects the largest value that the step size can be set to, and consequently determines the number of iterations.

When we are working with graphs, (1) is commonly a Laplacian system on the input graph with edge weights given by \mathbf{v} . Our first contribution is to relate the *electrical conductance* of this graph with the magnitude of the entries in \mathbf{x} . In particular, high conductance implies that \mathbf{x} has small entries, which can be used to obtain an efficient algorithm for (a modified) interior point method on high-conductance graphs.

Unfortunately, guaranteeing that the conductance is always high is not an easy task, because the edge weights can change significantly throughout the algorithm. Besides, some of the weights eventually converge to 0 (this is because some constraints become tight), so it is guaranteed that some edges have arbitrarily bad conductance. In order to alleviate this issue, we modify the graph to improve its conductance. After necessary post-processing to repair these modifications, this leads to an $\tilde{O}\left(m^{4/3} \log W\right)$ time algorithm for minimum cost flow on a graph with m edges, unit capacities, and costs bounded by W in absolute value. It also leads to an algorithm with the same runtime for *negative-weight shortest path* and *minimum cost bipartite \mathbf{b} -matching*, improving the runtimes of all these problems.

Even though the unit-capacity result makes considerable progress, it is hardly satisfying for cases when the edge capacities are large. This is because the magnitude of perturbations in the problem parameters depends polynomially on the edge capacities, and polynomial perturbations lead to a polynomial runtime overhead. To alleviate this issue, we refrain from trying to manually improve the graph's conductance, and instead focus on solving (1) faster, and in particular, in *sublinear* time. This dynamic problem can be cast more generally as maintaining the solution to a *dynamically-changing Laplacian system* in sublinear time. Interestingly, it turns out that the runtime of the updates required to dynamically solve (1) is also closely related to the magnitude of entries of the solution

vector \mathbf{x} . This leads to sublinear runtime, albeit only for unit-capacity graphs. Generalizing this idea to graphs with general capacities requires a substantially more fine-grained analysis, which reveals that the requirement for the entries of \mathbf{x} to be small is too strong, and can instead be replaced by an electrical energy argument. This leads to a runtime of $\tilde{O}(m^{3/2-\delta} \log(U+W))$ for some small $\delta > 0$ for minimum cost flow, which is the first runtime improvement in more than ten years.

Beyond flows The question about speeding up linear programs is also interesting to be asked in a broader context. When can we efficiently solve (and how efficiently can we solve) a convex program with an exponential number of constraints? A prime example of such a problem is *submodular function minimization (SFM)*. Although still under the umbrella of convex optimization, this problem is much harder than graph problems, and the fastest known algorithm for it runs in $\tilde{O}(n^3)$, where n is the size of the ground set of the function [107]. Even though this runtime is too slow for large-scale applications, the SFM instances that arise in practice often have special structure. Specifically, the function to be minimized is often *decomposable*, i.e. can be written as the sum of much simpler (e.g. $O(1)$ -sized) submodular functions. Examples include hypergraph cut with $O(1)$ -sized edges, image segmentation, and MAP inference.

Algorithms for this problem have used ideas from both traditional (discrete) algorithms (e.g. augmenting paths), and continuous optimization algorithms (e.g. coordinate descent). As none of these approaches have an ideal runtime, but they have complementary strengths and weaknesses, the question arises:

Can we bridge discrete and continuous approaches for the decomposable submodular minimization problem?

This part of the thesis will study how to achieve this by designing a hybrid discrete-continuous optimization algorithm, which works by a combination of a preconditioned iterative method and $\tilde{O}(1)$ calls to a maximum flow oracle. This is quite surprising, as the submodular minimization problem is not directly related to graphs. Using the recent breakthrough result of [37] which gives an almost-linear time maximum flow algorithm, this leads to an almost-linear time algorithm for decomposable submodular minimization, when all functions in the decomposition are $\tilde{O}(1)$ -sized.

Part II: Optimization Under Sparsity Constraints

How much can we compress a machine learning model while not sacrificing predictive performance? How fast can we produce such a model and what are the types of structured sparsity constraints that we can impose? What is the tradeoff between sparsity and runtime?

Modern machine learning tasks deal with huge amounts of data. For this reason, it is important for the model that is being produced to be as concise as possible, something that is especially crucial in applications with limited processing and storage capabilities. These questions naturally arise in the field of machine learning, and are also relevant in the fields of optimization and compressed sensing. While training a machine learning model, e.g. a linear model, is a continuous optimization task, it is often useful to include additional discrete constraints. Examples include forcing an upper bound on the number of features that are used, imposing a low rank structure to the solution, or injecting relational knowledge by a pre-specified graph. Such use cases are encountered for example in ML feature engineering, neural network model compression, ML design under budget constraints, and relationship-aware models for biology.

Sparse optimization Sparse optimization refers to the task of minimizing a function $f(\mathbf{x})$ under a *sparsity* constraint $\|\mathbf{x}\|_0 \leq s$, which restricts the vector \mathbf{x} to have at most s non-zero entries. As is often the case, the introduction of the discrete sparsity constraint results to the problem being NP-hard. The most common way to overcome this issue is to relax the non-convex sparsity constraint to its convex relaxation, i.e. the ℓ_1 constraint $\|\mathbf{x}\|_1 \leq \lambda$, making the problem amenable to convex optimization algorithms while at the same time returning a solution of relaxed sparsity $s' \geq s$. Other approaches include selecting coordinates incrementally in a greedy fashion (e.g. the orthogonal matching pursuit (OMP) algorithm [135] or Frank-Wolfe methods [69]), and performing projected gradient descent on the space of s -sparse solutions (e.g. the iterative hard thresholding (IHT) algorithm [19]). Yet, none of these are ideal: For convex f , OMP guarantees sparsity $s' \leq O\left(\kappa \log \frac{f(\mathbf{0}) - f(\mathbf{x}^*)}{\varepsilon}\right) \cdot s$, while IHT guarantees $s' \leq O(\kappa^2) \cdot s$, where κ is the (restricted) condition number of f . It is natural to ask whether we can get the best of both worlds, i.e. no dependence on the desired error tolerance ε , and a linear instead of quadratic dependence on the condition number.

Is there a sparse convex optimization algorithm that guarantees sparsity $s' \leq O(\kappa) \cdot s$?

In the second part of the thesis, we will show that this question can be answered affirmatively. In fact, by a more careful analysis, we will show that a simple modification to the IHT algorithm is enough to achieve this bound, thus essentially reducing the quadratic condition number dependence of IHT to linear, with no significant runtime overhead. This is made possible by a new technique that we call *adaptive regularization*, which avoids the sparsity performance degradation caused by projections onto structured non-convex domains, like the set of s -sparse vectors.

Low rank optimization Beyond sparsity, different measures often better capture the sparsity structure of a particular problem. For matrix problems, the most common measure of simplicity is rank, which is the sparsity of the matrix spectrum. The low rank optimization problem asks to minimize a function $f(\mathbf{A})$ acting on matrices, under the constraint that $\text{rank}(\mathbf{A}) \leq r$. Matrix completion, low rank matrix approximation, and robust principal component analysis, which are central problems in machine learning and signal processing, are all captured by the low rank optimization problem by appropriate choice of f . Algorithms for these problems are closely connected to sparse optimization algorithms, and indeed most algorithms for the latter have analogues for the former, albeit with some additional difficulties. In this thesis, we propose simple greedy and local search algorithms for low rank convex optimization, and show their theoretical and practical advantage.

1.1 Structure of the Thesis

Part I: Optimization on Graphs. In the first part of the thesis, we examine connections between flow algorithms and convex optimization, with the goal on one hand to design and analyze faster graph algorithms that are based on convex optimization, and simultaneously on the other hand to determine how purely graph algorithmic ideas tie back to the efficiency of convex optimization algorithms, even providing insights for problems not directly related to graphs.

- In **Chapter 3**, we will develop a way to speed up the number of iterations of interior point method when applied to the *unit-capacity minimum cost flow* problem, which leads to improving the runtime from $\tilde{O}(m^{10/7})$ [42] to $\tilde{O}(m^{4/3})$. By known reductions [42], this leads

to similarly faster algorithms for negative-weight shortest paths and minimum cost \mathbf{b} -matching. This chapter is based on the paper *Circulation Control for Faster Minimum Cost Flow in Unit-Capacity Graphs* [10], which is joint work with Aleksander Mądry and Adrian Vladu.

- In **Chapter 4**, we will continue with the minimum cost flow problem, this time with general (polynomially bounded) capacities. We will show that, by closely examining structural properties of electrical flows using spectral graph theory, we can speed up the amortized runtime of each interior point method iteration by discarding a large fraction of the graph. It is based on the paper *Faster Sparse Minimum Cost Flow by Electrical Flow Localization* [11], which is joint work with Aleksander Mądry and Adrian Vladu.
- In **Chapter 5**, we will study the connection between flow problems and submodular minimization. We will show that, when the submodular function is decomposable as the sum of $\tilde{O}(1)$ -sized submodular functions, it can be minimized by $\tilde{O}(1)$ calls to any maximum flow oracle. It is based on the paper *Decomposable Submodular Function Minimization via Maximum Flow* [9], which is joint work with Adam Karczmarz, Anish Mukherjee, Piotr Sankowski, and Adrian Vladu.

Part II: Optimization Under Sparsity Constraints. In the second part of the thesis, we deal with problems arising from resource-constrained machine learning and compressed sensing. We study optimization under sparsity constraints, which are discrete constraints that enforce a bound on the size of the predictor model. Our goal will be to present efficient optimization algorithms that can approach the optimal sparsity level. All the papers on which the second part of the thesis is based are joint works with Maxim Sviridenko.

- In **Chapter 6** we develop algorithms for sparse convex optimization, where the goal is to optimize a convex function under the constraint that the solution has a bounded number of non-zero entries. It will be based on a new technique which we call *adaptive regularization*, and is used to improve the sparsity of the returned solution, based on the paper *Sparse Convex Optimization via Adaptively Regularized Hard Thresholding* [14].
- Then, in **Chapter 7** we show how to alleviate the high computational overhead of the algorithm in the previous chapter, while achieving a similar sparsity bound. This is based on *Iterative Hard Thresholding with Adaptive Regularization: Sparser Solutions Without Sacrificing Runtime* [12].
- In **Chapter 8** we study convex optimization with a low rank constraint, which is related to problems like low rank approximation and robust principal components analysis. This is based on the paper *Local Search Algorithms for Rank-Constrained Convex Optimization* [13].

Chapter 2

Background

In this chapter, we will briefly present the background material necessary for following the rest of the thesis. This consists of commonly used notation, fundamental notions from convex analysis that will be extensively used throughout, as well as notation and fundamental theorems about graphs—one of the main objects this thesis is concerned about.

2.1 Basic Notation

Vectors. For any $k \in \mathbb{Z}_{\geq 0}$, we denote $[k] = \{1, 2, \dots, k\}$. We will use **bold** to refer to vectors or matrices. We denote by $\mathbf{0}$ the all-zero vector, $\mathbf{1}$ the all-one vector, \mathbf{O} the all-zero matrix, and by \mathbf{I} the identity matrix (with dimensions understood from the context). Additionally, we will denote by $\mathbf{1}_i$ the i -th basis vector, i.e. the vector that is 0 everywhere except at position i , where it is 1.

We apply scalar operations to vectors with the interpretation that they are applied element-wise. For example, given two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, \mathbf{x}/\mathbf{y} represents the vector whose i^{th} entry is x_i/y_i . Similarly, $\mathbf{x}\mathbf{y}$ is the element-wise product of \mathbf{x} and \mathbf{y} , and \mathbf{x}^2 is the element-wise square of vector \mathbf{x} . We use the inline notation $(\mathbf{x}; \mathbf{y})$ to represent the concatenation of vectors \mathbf{x} and \mathbf{y} . For any $\mathbf{x} \in \mathbb{R}^n$, we also denote the *support* of \mathbf{x} by $\text{supp}(\mathbf{x}) := \{i : x_i \neq 0\}$.

Where not ambiguous, we use the uppercase symbol corresponding to a symbol denoting a vector, to represent the diagonal matrix of that vector. In other words, given $\mathbf{r} \in \mathbb{R}^n$, $\mathbf{R} = \text{diag}(\mathbf{r})$ is a matrix whose diagonal entries are r_1, \dots, r_n and off-diagonal entries are 0.

Norms and Inner Products. Given two vectors \mathbf{x} and \mathbf{y} of the same dimension, we use $\langle \mathbf{x}, \mathbf{y} \rangle$ or $\mathbf{x}^\top \mathbf{y}$ to denote their inner product. For any $p \in \mathbb{R}_{\geq 1}$, we denote the ℓ_p norm of some vector $\mathbf{x} \in \mathbb{R}^n$ as

$$\|\mathbf{x}\|_p = \left(\sum_{i=1}^n |x_i|^p \right)^{1/p},$$

as well as the special cases $\|\mathbf{x}\|_0 = |\{i \in [n] : x_i \neq 0\}|$ and $\|\mathbf{x}\|_\infty = \max_{i \in [n]} |x_i|$. $\|\mathbf{x}\|_0$ is also called the *sparsity* of \mathbf{x} . For any ℓ_p norm and a non-negative weight vector $\mathbf{w} \in \mathbb{R}_{\geq 0}^n$, we also define the *weighted ℓ_p norm* of a vector $\mathbf{x} \in \mathbb{R}^n$ as:

$$\|\mathbf{x}\|_{\mathbf{w}, p} = \left(\sum_{i=1}^n w_i |x_i|^p \right)^{1/p}.$$

For two matrices $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{m \times n}$, we denote by $\langle \mathbf{A}, \mathbf{B} \rangle$ their *Frobenius inner product*, i.e. $\sum_{i=1}^m \sum_{j=1}^n A_{ij} B_{ij}$, which is also equal to the trace $\text{Tr}(\mathbf{A}^\top \mathbf{B})$. We also let $\|\mathbf{A}\|_2$ be the *spectral norm* of \mathbf{A} , i.e. the largest singular value of \mathbf{A} , $\|\mathbf{A}\|_F := \sqrt{\langle \mathbf{A}, \mathbf{A} \rangle}$ be the *Frobenius norm* of \mathbf{A} , and $\|\mathbf{A}\|_*$ be the *nuclear norm* of \mathbf{A} , i.e. the sum of its singular values.

Definition 2.1.1 (Dual norm). *Given a norm $\|\cdot\|$, its dual norm is defined as*

$$\|\mathbf{y}\|_* = \max_{\|\mathbf{x}\| \leq 1} \langle \mathbf{y}, \mathbf{x} \rangle.$$

Restrictions and Thresholding. For any vector $\mathbf{x} \in \mathbb{R}^n$ and set $S \subseteq [n]$, we denote by \mathbf{x}_S the vector that results from \mathbf{x} after zeroing out all the entries except those in indices given by S . For any $t \in \mathbb{R}$ and $\mathbf{x} \in \mathbb{R}^n$, we denote by $\mathbf{x}_{\geq t}$ the vector that results from setting all the entries of \mathbf{x} that are less than t to 0. For a differentiable function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, its gradient $\nabla f(\mathbf{x})$, and a set of indices $S \subseteq [n]$, we denote $\nabla_S f(\mathbf{x}) := (\nabla f(\mathbf{x}))_S$. We also define the *thresholding operator* $H_s(\mathbf{x})$ to be equal to \mathbf{x}_R , where R are the s entries of \mathbf{x} with largest absolute value (breaking ties arbitrarily). We override the thresholding operator $H_r(\mathbf{A})$ when the argument is a matrix \mathbf{A} , defining $H_r(\mathbf{A}) = \mathbf{U} \text{diag}(H_r(\boldsymbol{\lambda})) \mathbf{V}^\top$, where $\mathbf{U} \text{diag}(\boldsymbol{\lambda}) \mathbf{V}^\top$ is the singular value decomposition of \mathbf{A} . In other words, $H_r(\mathbf{A})$ only keeps the top r singular components of \mathbf{A} .

For a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, a subset of rows $C \subseteq [m]$ and subset of columns $F \subseteq [n]$, we denote by \mathbf{A}_{CF} the matrix that results by setting all the elements A_{ij} where $i \notin C$ or $j \notin F$ to 0 and keeping the rest intact.

2.2 Convex Analysis

One of the useful features of using continuous optimization is that it replaces the ad hoc potentials and analytic tools of discrete algorithms by principled, generic mathematical properties and tools. For example, one can analyze algorithms for a discrete optimization problem by turning it into a continuous function in high-dimensional space, then studying fundamental analytic properties of this function, such as convexity, smoothness, and Lipschitzness. In this section, we will go over the most widely used of these properties. As we will be mostly dealing with differentiable functions throughout the thesis, we tailor our definitions for differentiable functions. Given a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, we will denote its gradient at \mathbf{x} by $\nabla f(\mathbf{x}) \in \mathbb{R}^n$ and its *Hessian* at \mathbf{x} (when defined) by $\nabla^2 f(\mathbf{x}) \in \mathbb{R}^{n \times n}$.

Convexity. The most important notion will undoubtedly be *convexity*. This is partially because it comes hand in hand with a polynomial time algorithm for minimizing such a function. In fact, when posed with an algorithmic problem about which not much is known, the principled approach is to first attempt to reduce it to minimizing a convex function. Let us get some insight into what convexity is. A set $S \subseteq \mathbb{R}^n$ is convex if it is closed under convex combinations, i.e. for any $\mathbf{x}, \mathbf{y} \in S$ and $\tau \in [0, 1]$, the solution $\mathbf{z} = \tau \mathbf{x} + (1 - \tau) \mathbf{y}$ that lies in the straight line between \mathbf{x} and \mathbf{y} , also belongs to S . A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is then convex if all its level sets are convex. Alternatively, if any linear Taylor approximation of f lower bounds the value of f everywhere, i.e. for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ we have $f(\mathbf{y}) \geq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle$.

The power of this simple property comes from the fact that it is always possible to reach an *optimal* solution by following a straight line from any *feasible* solution, while always staying feasible. What is somewhat of a mystery is why this property, which on its surface can appear quite restrictive, is so commonly encountered in practice, and even when it fails to hold, the optimization frameworks

that are built on convex foundations are still practically applicable and with surprisingly good performance.

The most basic yet important algorithm for convex functions is the *gradient descent algorithm*, which minimizes a convex function by repeatedly moving in the negative direction of the gradient: $\mathbf{x}' = \mathbf{x} - \eta \nabla f(\mathbf{x})$, where $\eta > 0$ is a problem-dependent (and often iteration-dependent) parameter called the *step size*.

Lipschitzness and smoothness. While convexity is sufficient to obtain a polynomial time algorithm, functions encountered often have more “niceness” properties. In particular, if the value or the gradient of a function change erratically while moving in the problem domain, it can be easy to *overshoot* the optimum, thus requiring the algorithm to use very small step sizes and compromising efficiency.

It is known that, as long as the function is *Lipschitz*, the (sub)gradient descent (or more generally *mirror descent*) algorithm can be analyzed to give useful convergence bounds. The Lipschitzness property is expressed by a bound on the norm of the gradient of the function. and it gives a bound on how fast the value of the function can change in the underlying norm. Specifically, $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is called *L-Lipschitz* with respect to a norm $\|\cdot\|$ if for any $\mathbf{x} \in \mathbb{R}^n$ we have $\|\nabla f(\mathbf{x})\|_* \leq L$.

In some cases, it is beneficial to exploit the *smoothness* of the function, i.e. how fast the *gradient* changes in the underlying norm. Algorithmically, this gives us a predictable change in the gradient along the update step, and thus also a lower bound on the step size and a convergence analysis of gradient descent. Smoothness also gives an upper bound on the error of the first-order Taylor expansion of f . $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is called β -*smooth* with respect to a norm $\|\cdot\|$ if for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ we have $f(\mathbf{y}) \leq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + (\beta/2) \|\mathbf{y} - \mathbf{x}\|^2$.

Strong convexity and condition number. Sometimes, the rate of change of the gradient can also be *lower bounded*. Together with smoothness, this property called *strong convexity* implies that not only the gradient is correlated with the direction to the optimal solution, but is strongly correlated. Together, the smoothness and strong convexity properties lead to very favorable, linear convergence rates for gradient descent. Concretely, $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is called α -*strongly convex* with respect to a norm $\|\cdot\|$ if for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ we have $f(\mathbf{y}) \geq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + (\alpha/2) \|\mathbf{y} - \mathbf{x}\|^2$. If f is both smooth and strongly convex with respect to the ℓ_2 norm, then we can define the *condition number* $\kappa \geq 1$ as β/α , where $\beta > 0$ is the smallest constant (or infimum) such that f is β -smooth and $\alpha > 0$ is the largest constant (or supremum) such that f is α -strongly convex.

Sparse supports. In some applications like compressed sensing, the function is only going to be well-conditioned along *sparse* directions. This leads to a generalization of all the previous definitions. Specifically, a function is called β -*smooth* at sparsity level $s \geq 1$ if for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ with $\|\mathbf{y} - \mathbf{x}\|_0 \leq s$ we have

$$f(\mathbf{y}) \leq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + (\beta/2) \|\mathbf{y} - \mathbf{x}\|^2 .$$

We denote the smallest such constant β by β_s and call it the *restricted smoothness constant* (RSS constant). Completely analogously, f is called α -*strongly convex* at sparsity level $s \geq 1$ if for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ with $\|\mathbf{y} - \mathbf{x}\|_0 \leq s$ we have

$$f(\mathbf{y}) \geq f(\mathbf{x}) + \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle + (\alpha/2) \|\mathbf{y} - \mathbf{x}\|^2 .$$

We denote the largest such constant α by α_s and call it the *restricted strong convexity constant* (RSC constant). Then, the *restricted condition number* (RCN) of f is denoted by $\kappa_s \geq 1$ and defined as β_s/α_s . The properties defined above are also closely connected to the well-known restricted isometry property, which is commonly used in compressed sensing. We say that a function f has the *restricted isometry property (RIP)* at sparsity level s if $\beta_s, \alpha_s \in \mathbb{R}_{>0}$. Then, the *RIP constant* of f at sparsity level s is defined as $\delta_s = \frac{\kappa_s - 1}{\kappa_s + 1}$.¹

Projections. The previously outlined gradient descent algorithm works for the problem of unconstrained minimization (i.e. minimizing over \mathbb{R}^n). If there are additional constraints, the minimization is over a subset of \mathbb{R}^n . This is tackled by incorporating an additional step to gradient descent, which projects the solution iterate onto the feasible set. The resulting algorithm is often called *projected gradient descent*. It is thus important to introduce notation for projections. Given a subspace $\mathcal{V} \subseteq \mathbb{R}^n$, we will denote the orthogonal (ℓ_2) projection operator onto \mathcal{V} as $\mathbf{I}_{\mathcal{V}} \in \mathbb{R}^{n \times n}$. In particular, for any matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ we denote by $\text{im}(\mathbf{A}) = \{\mathbf{A}\mathbf{x} \mid \mathbf{x} \in \mathbb{R}^n\}$ the *image* of \mathbf{A} and by $\text{ker}(\mathbf{A}) = \{\mathbf{x} \mid \mathbf{A}^\top \mathbf{x} = \mathbf{0}\}$ the *kernel* of \mathbf{A} . Therefore, $\mathbf{I}_{\text{im}(\mathbf{A})} = \mathbf{A}(\mathbf{A}^\top \mathbf{A})^+ \mathbf{A}^\top$ is the orthogonal projection onto the image of \mathbf{A} and $\mathbf{I}_{\text{ker}(\mathbf{A}^\top)} = \mathbf{I} - \mathbf{I}_{\text{im}(\mathbf{A})}$ the orthogonal projection onto the kernel of \mathbf{A}^\top , where $(\cdot)^+$ denotes the Moore-Penrose matrix pseudoinverse.

2.3 Graphs and Flows

Graphs are significant in their own right because of their direct applications, but a perhaps even more crucial aspect of their significance comes because of their role in the context of bridging discrete and continuous optimization. They are the most simple mathematical objects that have a meaningful and non-trivial dual nature: Either as *combinatorial* objects consisting of vertices and edges connecting these, or as *algebraic* objects representing linear operators in a high-dimensional space.

Graphs as sets of vertices and edges. A (directed) graph $G = (V, E)$ is a mathematical object consisting of a set of vertices V and a set of edges E . Each edge (or sometimes arc, in the case of directed graphs) e of G is a link that connects a certain (ordered) pair of vertices (u, v) to each other. u is called the *tail* of edge e , and v is called its *head*. We denote by $E^+(u)$ the set of edges that have u as their tail and $E^-(u)$ those that have u as their head. When undirected, will write $e \sim v$ to denote the set of edges $e \in E$ that have v as their endpoint. A *walk* in G is a sequence of vertices, where each pair of successive vertices are connected by an edge, and its length is the number of vertices minus 1. A *path* is simply a walk without repeated vertices. We also use the notation $G = (V, E, \mathbf{c})$ to denote a *weighted* graph, where $\mathbf{c} \in \mathbb{R}^{|E|}$ is a cost vector and each edge e is associated by a real-valued cost c_e . When a graph $G(V, E)$ is clear from context, we will use $n = |V|$ to denote the number of vertices and $m = |E|$ to denote the number of edges.

Graphs as linear operators. An alternative way to view a graph is to encode it as a matrix. One common way to do so is by the *adjacency matrix* of a graph. This matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ has one row and one column for each vertex, and a real number in the corresponding position (u, v) that is equal to 0 if the edge (u, v) does not exist in the graph, and 1 (or some other number if the graph is weighted) if the edge exists. One way to glimpse at the usefulness of the adjacency matrix is to raise it to a natural power \mathbf{A}^k and notice that the entry at position (u, v) of this matrix is equal to the total number of distinct walks from u to v with length exactly k . The cumbersomeness of

¹We note that this is a more general definition than the one usually given, but it is equivalent up to a rescaling of f .

computing the same quantity in the combinatorial version shows the power of the dual nature of graphs. For various technical reasons, we will be instead using a related but different matrix called the *incidence matrix* $\mathbf{B} \in \mathbb{R}^{m \times n}$. The rows of this matrix correspond to edges and the columns to vertices, and the row corresponding to edge $e = (u, v)$ is all-zero except 1 at position u and -1 at position v (if the graph is undirected, then it can be made directed by arbitrarily orienting each edge). When applied to vectors, \mathbf{B} turns vertex potentials into edge potential drops, and \mathbf{B}^\top turns flows into vertex demands.

Flows. Given a graph G we view a flow in G as a vector $\mathbf{f} \in \mathbb{R}^m$ that assigns a value to each arc of G . If this value is negative we interpret it as having a flow of $|f_e|$ flowing in the direction opposite to the arc orientation. This convention is especially useful when discussing flows in undirected graphs.

We will be working with flows in G that satisfy a certain *demand* $\mathbf{d} \in \mathbb{R}^n$ such that $\sum_u d_u = 0$. We say that a flow \mathbf{f} satisfies or routes demand \mathbf{d} if it satisfies the flow conservation constraints with respect to the demands. That is:

$$\sum_{e \in E^+(u)} f_e - \sum_{e \in E^-(u)} f_e = d_u, \quad \text{for all } u \in V, \quad (2.1)$$

or simply $\mathbf{B}^\top \mathbf{f} = \mathbf{d}$ by using the incidence matrix. Intuitively, these constraints enforce that the net balance of the total in-flow into vertex u and the total out-flow leaving that vertex is equal to d_u . A flow for which the demand vector \mathbf{d} is zero everywhere is called a *circulation*.

We say that a flow \mathbf{f} is *feasible* (or that it respects capacities) in G if it obeys the capacity constraints:

$$0 \leq f_e \leq u_e, \quad \text{for all } e \in E, \quad (2.2)$$

where $\mathbf{u} \in \mathbb{R}^m$ is a vector of arc capacities.

We can now arrive at the definition of the minimum cost flow problem, which is one of the most widely studied and general flow problems:

Definition 2.3.1 (Minimum cost flow). *Given a directed graph $G = (V, E, \mathbf{c})$ with demands $\mathbf{d} \in \mathbb{R}^n$ and capacities $\mathbf{u} \in \mathbb{R}_{>0}^m$, the minimum cost flow problem asks to compute a flow \mathbf{f} that*

- routes the demand: $\mathbf{B}^\top \mathbf{f} = \mathbf{d}$,
- respects the capacities: $\mathbf{0} \leq \mathbf{f} \leq \mathbf{u}$, and
- minimizes the cost: $\langle \mathbf{c}, \mathbf{f} \rangle$.

We will denote such an instance of minimum cost flow by the tuple $(G, \mathbf{c}, \mathbf{d}, \mathbf{u})$.

If we drop the capacity constraint, we obtain the *transshipment problem*. If we also replace the demand by a unit $s - t$ demand, we obtain the *shortest path problem*. If we instead drop the cost minimization constraint and replace the demand by an (appropriately scaled) $s - t$ demand, we obtain the *maximum flow problem*. We can thus see why the above formulation is so widely studied and powerful.

2.4 Electrical Flows and Laplacians

Electrical flows. The flow problems mentioned in the previous section can be formulated as ℓ_∞ or ℓ_1 flow minimization under appropriate constraints. For various reasons, it is natural and useful

to also study ℓ_2 flow minimization. Natural, because the ℓ_2 norm is the objective minimized by flow processes in nature, e.g. in electricity, hydraulics, and random walks. It also leads to a more robust solution. For example, penalizing the ℓ_1 norm of the flow provides an optimal solution for computing the shortest path in a routing problem. However, when subject to real world situations like unexpected network changes or diversity requirements, it is desirable to have a more robust solution (i.e. consisting of more than a single path), even if it is not optimal. As for the usefulness, ℓ_2 minimization can serve as a proxy for ℓ_1 or ℓ_∞ minimization, thus allowing one to solve these with a potentially simpler ℓ_2 oracle.

Concretely, given a graph $G(V, E)$ with edge resistances $\mathbf{r} \in \mathbb{R}_{>0}^m$, a flow $\mathbf{f} \in \mathbb{R}^m$ is called an *electrical flow* if it minimizes the *energy* $\sum_{e \in E} r_e f_e^2$ among all flows routing the demand of \mathbf{f} (i.e. $\mathbf{d} = \mathbf{B}^\top \mathbf{f}$). Electrical flows are very closely connected to *graph Laplacians*. In particular, the electrical flow \mathbf{f} induced by the demand \mathbf{d} with resistances \mathbf{r} is given by $\mathbf{f} = \mathbf{R}^{-1} \mathbf{B} \mathbf{L}^+ \mathbf{d}$, where $\mathbf{L} = \mathbf{B}^\top \mathbf{R}^{-1} \mathbf{B}$ is the Laplacian of G with resistances \mathbf{r} .

Potential embeddings. The dual notion to that of a flow is the *potential embedding*. In fact, the energy optimality of a flow \mathbf{f} can be certified by a potential embedding $\phi \in \mathbb{R}^n$ that satisfies Ohm's law for all edges $e = (u, v)$, i.e. $\phi_v - \phi_u = r_e f_e$, or just $\mathbf{B}\phi = \mathbf{r}\mathbf{f}$ by using the incidence matrix. Given resistances \mathbf{r} , we will say that a potential embedding ϕ *induces* the flow $\mathbf{f} = \frac{\mathbf{B}\phi}{\mathbf{r}}$, and conversely, given an electrical flow, we will say that it induces the potentials that certify its energy optimality. The *energy of a potential embedding* is defined to be equal to the energy of the electrical flow induced by ϕ :

$$E_{\mathbf{r}}(\phi) = \sum_{e \in E} \frac{(\mathbf{B}\phi)_e^2}{r_e}.$$

It should be noted that even though every potential embedding induces an electrical flow with some demand, not every flow is electrical. For example, a flow cycle can never be electrical. This can also be witnessed by the fact that potential embeddings live in n -dimensional space but flows live in m -dimensional space, so there are “more” flows than potential embeddings.

Electrical energy statements. For any valid demand and resistances, there exists an electrical flow that routes the demand. It useful to abstract away from the electrical flow, which is the solution to an ℓ_2 minimization problem, and directly deal with the energy required to route a particular demand. This will be a very useful potential throughout, since it is much more stable than tracking the electrical flow.

For any valid demand ($\langle \mathbf{1}, \mathbf{d} \rangle = 0$) and resistances $\mathbf{r} > \mathbf{0}$, we denote by

$$\mathcal{E}_{\mathbf{r}}(\mathbf{d}) = \min_{\phi: \mathbf{L}\phi = \mathbf{d}} E_{\mathbf{r}}(\phi)$$

the total energy that is required to route the demand \mathbf{d} with resistances \mathbf{r} , where $\mathbf{L} = \mathbf{B}^\top \mathbf{R}^{-1} \mathbf{B}$ is the Laplacian of G . We extend this definition for a \mathbf{d} that is not a demand vector ($\langle \mathbf{1}, \mathbf{d} \rangle \neq 0$), as $\mathcal{E}_{\mathbf{r}}(\mathbf{d}) = \mathcal{E}_{\mathbf{r}}\left(\mathbf{d} - \frac{\langle \mathbf{1}, \mathbf{d} \rangle}{n} \cdot \mathbf{1}\right)$. We now outline various properties of energy, that show its predictability and make it a suitable quantity to be used as a potential in the analysis of algorithms.

Fact 2.4.1 (Energy statements). *Consider a graph $G(V, E)$ with resistances \mathbf{r} .*

- For any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, we have $\sqrt{\mathcal{E}_{\mathbf{r}}(\mathbf{x} + \mathbf{y})} \leq \sqrt{\mathcal{E}_{\mathbf{r}}(\mathbf{x})} + \sqrt{\mathcal{E}_{\mathbf{r}}(\mathbf{y})}$.

- Consider a vector $\mathbf{d} \in \mathbb{R}^n$. Then,

$$\max_{\phi: E_r(\phi) \leq 1} \langle \mathbf{d}, \phi \rangle = \sqrt{\mathcal{E}_r(\mathbf{d})}.$$

- For any resistances $\mathbf{r}' \leq \alpha \mathbf{r}$ for some $\alpha \geq 1$ and any $\mathbf{d} \in \mathbb{R}^n$, we have $\mathcal{E}_{\mathbf{r}'}(\mathbf{d}) \leq \alpha \cdot \mathcal{E}_r(\mathbf{d})$

Proof. We let $\mathbf{L} = \mathbf{B}^\top \mathbf{R}^{-1} \mathbf{B}$ be the Laplacian of G with resistances \mathbf{r} . For the first one, we have $\sqrt{\mathcal{E}_r(\mathbf{x} + \mathbf{y})} = \|\mathbf{x} + \mathbf{y}\|_{\mathbf{L}^+} \leq \|\mathbf{x}\|_{\mathbf{L}^+} + \|\mathbf{y}\|_{\mathbf{L}^+} = \sqrt{\mathcal{E}_r(\mathbf{x})} + \sqrt{\mathcal{E}_r(\mathbf{y})}$, where we used the triangle inequality.

The second one follows since $E_r(\phi) = \|\phi\|_{\mathbf{L}}^2$ and $\mathcal{E}_r(\mathbf{d}) = \|\mathbf{d}\|_{\mathbf{L}^+}^2$ and the norms $\|\cdot\|_{\mathbf{L}}$ and $\|\cdot\|_{\mathbf{L}^+}$ are dual.

For the third one, we note that $(\mathbf{B}^\top \mathbf{R}'^{-1} \mathbf{B})^+ \preceq \alpha \cdot (\mathbf{B}^\top \mathbf{R}^{-1} \mathbf{B})^+$, and so $\mathcal{E}_{\mathbf{r}'}(\mathbf{d}) = \|\mathbf{d}\|_{(\mathbf{B}^\top \mathbf{R}'^{-1} \mathbf{B})^+}^2 \leq \alpha \|\mathbf{d}\|_{(\mathbf{B}^\top \mathbf{R}^{-1} \mathbf{B})^+}^2 = \alpha \cdot \mathcal{E}_r(\mathbf{d})$. \square

Other than its use as a potential, we will also use energy as a *metric*. Just like in the context of shortest paths the right distance metric to use is the shortest path distance, in the context of electrical flows the right distance metric is the *effective resistance*. Roughly, effective resistance measures how close two vertices (or sets of vertices) of the graph are by how much energy is needed to transport flow between them. In contrast to shortest path distance, which only depends on the length of the shortest path between the two vertices, effective resistance also takes the number of alternative (edge-disjoint) paths into account.

Definition 2.4.2 (Effective resistances). Consider a graph $G(V, E)$ with resistances \mathbf{r} and any pair of vertices $u, v \in V$. We denote by $R_{eff}(u, v)$ the energy required to route 1 unit of flow from u to v , i.e. $R_{eff}(u, v) = \mathcal{E}_r(\mathbf{1}_u - \mathbf{1}_v)$. This is called the effective resistance between u and v . We extend this definition to work with vertex subsets $X, Y \subseteq V$, such that $R_{eff}(X, Y)$ is the effective resistance between the vertices x, y that result from contracting X and Y . When used as an argument of R_{eff} , an edge $e = (u, v) \in E$ is treated as the vertex subset $\{u, v\}$.

Asymptotic notation. We use $\tilde{O}(t)$ to denote a quantity that is $O(t \log^k t)$ for some constant $k \geq 0$. When the main problem parameter is clear from context (e.g. the number of edges in a graph), we will use $\tilde{O}(1)$ to denote a quantity that is $O(\log^k m)$ for some constant $k \geq 0$.

Part I

Optimization on Graphs

Chapter 3

Circulation Control for Faster Minimum Cost Flow in Unit-Capacity Graphs

3.1 Introduction

Finding the least costly way to route a demand through a network is a fundamental algorithmic primitive. Within the context of algorithmic graph theory it is captured as the *minimum cost flow problem*, in which given a graph with costs on its arcs and a set of demands on its vertices, one needs to find a flow that routes the demand while minimizing its cost. This problem has received significant attention [4] and inspired the development of new algorithmic techniques. For example, Orlin’s network simplex algorithm [133] offered an explanation of the excellent behavior that the simplex method exhibits in practice when applied to flow problems. More broadly, the recent progress on algorithms for the flow problems [44, 39, 118, 105, 152, 97, 117, 136, 42, 150, 153, 115, 114] has been an instance of the general approach to graph algorithms that leverages the tools of continuous optimization, rather than classical combinatorial techniques. Also, there exist efficient reductions that enable us to leverage algorithms for the minimum cost flow problem to solve a host of other fundamental problems, including the maximum flow problem, the minimum cost bipartite matching problem, and the shortest path problem with negative weights.

3.1.1 Our Contributions

In this chapter, we present an $m^{4/3+o(1)} \log W$ -time algorithm for the minimum cost flow problem in graphs with unit capacities, where W denotes the bound on the magnitude of the arc costs. This improves upon the previously known $\tilde{O}(m^{10/7} \log W)$ running time bound of Cohen et al. [42] and matches the running times of the recent algorithms due to Liu and Sidford [115, 114] for the unit capacity maximum flow problem.

Similarly to most of the relevant prior work, our algorithm at its core relies on an interior point method, but the variant of the interior point method we design and employ is directly attuned to the combinatorial properties of the graph. In particular, in contrast to [42], we do not rely on a reduction to the bipartite perfect \mathbf{b} -matching problem (which requires a sophisticated analysis). Instead, our algorithm operates directly in the space of circulations of the original graph.

One can also draw an analogy between the network simplex method [133] and ours. The former navigated the corners of a feasible polytope and improved an existing suboptimal solution through pushing flow around cycles. In contrast, we iteratively improve our existing suboptimal solution by augmenting it with circulations, but navigate through the strict interior of the polytope, seeking to

keep a specific condition called *centrality* satisfied. Also, while in the network simplex case, the key difficulty is in finding the right pivoting rule, our approach shifts the attention towards finding the right circulation to augment the flow with so as to maintain the centrality invariant.

A key ingredient of our approach is a custom preconditioning method, which enables us to control the flows we use to update the solution in each iteration. We derive a new way to tie the conductance of the graph to a certain guarantee on the flows computed in the preconditioned graph. This allows us to perform a better, tighter analysis of the quality of the preconditioner we use.

On a more technical level, our work provides a number of insights into the underlying interior point method. In particular, in our $m^{11/8+o(1)} \log W$ -time algorithm (that we develop first), the progress steps we perform in order to reduce the duality gap of our current solution are cast as a refinement procedure, which simply attempts to correct a residual. This procedure is very similar to iterative refinement—widely used in the more restricted case of minimizing convex quadratic functions [167, 81]. Also, in contrast to the classic approach for maintaining constraint feasibility during the interior point method update step—which relies on controlling the ℓ_2 norm of the relative updates to the slack variables—we want to perform steps for which it is only guaranteed that these relative updates are small in ℓ_∞ norm. To this end, we employ a custom residual correction procedure that works by re-weighting the capacity constraints. (It is worth noting that a similar procedure has been used in [115].¹)

This paves the way for the final algorithm that has the further improved running time of $m^{4/3+o(1)} \log W$. As a matter of fact, the key bottleneck to obtaining a faster algorithm using the above approach is the need to ensure that the residual error in the solution obtained after performing a step bounded in ℓ_∞ norm can be reduced to zero. This requires increasing the weights on the constraint barriers, and these weight increases are exactly what limits the exponent in the running time to 11/8. The step problem we need to solve, however, is well conditioned within a local ℓ_∞ ball around the current iterate. Therefore, being able to certify that the point returned by solving the step problem optimally lies within this local ℓ_∞ ball, implies that we can efficiently find it using a direct optimization subroutine. This latter observation is the key insight in the very recent preprint by Liu-Sidford [114] that enables them to improve the running time for maximum flow in graphs with unit capacity. We employ this insight in our setting in order to obtain an improved running time for unit-capacity minimum cost flow as well.

Finally, in order to guarantee that the ℓ_∞ norm of each progress step is indeed as small as needed, we employ a convex optimization subroutine with mixed ℓ_2 and ℓ_p terms [103], instead of solving a linear system of equations in each update step as is typically done. (Such subroutine was similarly used by Liu and Sidford [115, 114], in a slightly different form.)

3.1.2 Previous Work

Due to the size of the existing literature on the studied problems, we focus our discussion only on the works that are the most relevant to our results and refer the reader to [74] and Section 1.2 in [42] for a more detailed discussion.

In 2013, Mađry [118] developed an algorithm that produces an optimal solution to the unit capacity maximum flow problem in $\tilde{O}(m^{10/7})$ time and thus improves over a long standing running time barrier of $\tilde{O}(n^{3/2})$ in the case of sparse graphs. An important characteristic of this approach was the careful tracking of an electrical energy quantity which allowed to control the step size. The underlying approach was then simplified by providing a more direct correspondence between the update steps of the interior point method and computing augmenting paths via electrical flow

¹While our analysis aims to enforce a small ℓ_2 norm of the residual error, [115] seek to control the ℓ_4 norm of the congestion vector. These techniques turn out to be largely equivalent.

computations [117]. This framework has been also extended to a more general setting of unit capacity minimum cost flow [42], achieving a running time of $\tilde{O}(m^{10/7} \log W)$, where W upper bounds the largest cost of a graph edge in absolute value.

In a different context, motivated by new developments involving regression problems [52, 28, 1], Kyng et al. [103] studied the ℓ_p regression problems on graphs, obtaining an algorithm which runs in $m^{1+o(1)}$ time for a range of large values of p . This algorithm's running time was subsequently further improved by Adil and Sachdeva [2].

Liu and Sidford [115] have recently obtained an improved algorithm for the unit capacity maximum flow problem with a running time of $m^{11/8+o(1)}$. One of their key insights was that the work on ℓ_p -regression problems enables treating energy control as a self-contained problem in each iteration of the interior point method, rather than maintaining energy as a global potential over the whole course of the algorithm, which was the case in previous work. Then, in their recent follow-up work, Liu and Sidford [114] strengthen the step problem primitive by directly optimizing a regularized log barrier function as opposed to performing a sequence of regularized Newton step. This led to a running time of $m^{4/3+o(1)}$ for the unit capacity maximum flow problem.

3.1.3 Organization of the Chapter

We begin with technical preliminaries in Section 3.2. In Section 3.3, we present our interior point framework specialized to minimum cost flow, and provide a basic analysis which yields an algorithm running in $\tilde{O}(m^{3/2} \log W)$ time. We further refine our framework in Section 3.4, where we develop the key tools needed for our results, giving a faster, $m^{11/8+o(1)} \log W$ -time algorithm for obtaining the solution to a slightly perturbed instance of the original minimum cost flow problem. In Section 3.5 we then show how to use existing combinatorial techniques to repair this perturbed instance. Finally, in Section 3.6, we demonstrate how to combine the framework developed in the previous sections with an insight from the recent work of Liu and Sidford [114] to achieve the final running time of $m^{4/3+o(1)} \log W$.

3.2 Preliminaries

The *unit capacity minimum cost flow problem* is to find a flow $\mathbf{f} \in \mathbb{R}^m$ that meets the *unit capacity constraints* $0 \leq f_e \leq 1$ for all $e \in E$ and routes the demand \mathbf{d} , while minimizing the *cost* $\sum_{e \in E} c_e f_e$.

Cycle Basis. A set of circulations \mathcal{C} in G is called a *cycle basis* if any circulation in G can be expressed as a linear combination of circulations in \mathcal{C} . If G is connected, the dimension of a cycle basis of G is $m - n + 1$.

3.3 Minimum Cost Flow by Circulation Improvement

In this section we present our (customized) interior point method-based framework for solving the minimum cost flow problem, setting the foundations for the faster algorithm of Section 3.4.

3.3.1 LP Formulation and Interior Point Method

We first cast the minimum cost flow problem as a linear program that we then proceed to solve using an interior point method.

LP formulation. It will be useful to consider the parametrization of a flow in terms of the circulation space of the graph. The goal of this re-parametrization is to initialize the interior point method with an initial flow \mathbf{f}_0 which routes the prescribed demand \mathbf{d} , then iteratively improve it by adding circulations to get a flow which routes the same demand \mathbf{d} but has lower duality gap. It is noteworthy that the specific parametrization of the circulation space is irrelevant to the interior point method, due to its affine invariance. We will elaborate on this point later. For us it will be a useful tool for understanding the centrality condition arising from the interior point method and applying more aggressive progress steps.

Given the (connected) underlying graph $G = (V, E)$, let $\mathbf{C} \in \mathbb{R}^{m \times (m-n+1)}$ be a matrix whose columns encode the characteristic vectors of a basis for G 's circulation space.

In order to construct such a matrix, we let $C_1, C_2, \dots, C_{m-n+1}$ be an arbitrary cycle basis for G , where we ignore the arc orientations. An easy way to produce one is to consider a spanning tree $T \subseteq G$. For each arc $(u, v) \in E$ which is not in T , consider the unique path in T connecting v and u . The arcs on this path along with the arc (u, v) determine a cycle in the basis. More specifically, consider the set of arcs of G present in C_i , sorted according to the order in which they are visited along the cycle, starting with the off-tree arc (u, v) , then continuing with those witnessed along the tree path from v to u . If an arc $e \in E$ has the opposite orientation to the one corresponding to the traversal of the cycle, we represent it as \bar{e} , otherwise we write it just as e .

Now, letting C_i consist of a subset of arcs in E , each of which appears either with its original orientation e , or the opposite orientation \bar{e} , we write the i^{th} column of matrix \mathbf{C} as follows.

$$\mathbf{C}_{e,i} = \begin{cases} 1, & \text{if } e \in C_i, \\ -1, & \text{if } \bar{e} \in C_i, \\ 0, & \text{otherwise.} \end{cases}$$

We can now use \mathbf{C} to represent any circulation in G . Given any $\mathbf{x} \in \mathbb{R}^{m-n+1}$ we have that $\mathbf{f} = \mathbf{C}\mathbf{x}$ is a circulation. Furthermore the sign of each coordinate f_e , $e = (u, v) \in E$, shows whether \mathbf{f}_e is a flow that runs in the same direction as e or vice-versa, i.e. $f_e > 0$ if \mathbf{f} carries flow from u to v , and similarly $f_e < 0$ if \mathbf{f} carries flow from v to u . On the other hand, for any circulation $\mathbf{f} \in \mathbb{R}^m$ there exists an $\mathbf{x} \in \mathbb{R}^{m-n+1}$ such that $\mathbf{f} = \mathbf{C}\mathbf{x}$ (or in other words the image of \mathbf{C} is the space of circulations).

Now let \mathbf{f}_0 be a flow in G such that for each arc $e \in E$, $0 < (\mathbf{f}_0)_e < 1$, and \mathbf{f}_0 routes the demand \mathbf{d} . The minimum cost flow problem can be cast as the following linear program:

$$\begin{aligned} \min \langle \mathbf{c}, \mathbf{C}\mathbf{x} \rangle & \\ \mathbf{0} \leq \mathbf{f}_0 + \mathbf{C}\mathbf{x} \leq \mathbf{1}. & \end{aligned} \tag{3.1}$$

We see that the objective value of this linear program differs by a term of $\langle \mathbf{c}, \mathbf{f}_0 \rangle$ from the original objective. We did not include it here, since it is a constant. It is useful to also consider its dual:

$$\begin{aligned} \max -\langle \mathbf{1} - \mathbf{f}_0, \mathbf{y}^+ \rangle - \langle \mathbf{f}_0, \mathbf{y}^- \rangle & \\ \mathbf{C}^\top (\mathbf{y}^+ - \mathbf{y}^-) = -\mathbf{C}^\top \mathbf{c} & \\ \mathbf{y}^+, \mathbf{y}^- \geq \mathbf{0}. & \end{aligned} \tag{3.2}$$

The objective we are left to solve simply suggests that in order to find the minimum cost flow in the graph with unit capacities, we equivalently have to find the minimum cost circulation in the residual graph under shifted capacity constraints. This carries a significant similarity with the *network simplex* algorithm [133], which has been used in the past as a specialization of the

simplex method to the minimum cost flow problem. It essentially consisted of maintaining a solution routing the prescribed demand \mathbf{d} , and iteratively improving it by pushing flow around a cycle, while satisfying capacity constraints. Rather than performing such updates, which always maintain a flow on the boundary of polytope corresponding to the set of feasible solutions, the interior point method maintains a more sophisticated condition on these intermediate solutions. Another similar approach can be found in [98], where updates are iteratively pushed around cycles in order to solve Laplacian linear systems.

Like these methods, our approach will be to repeatedly improve the cost of the solution by pushing augmenting circulations. Crucially, maintaining a solution centrality condition, stemming from interior point methods, will allow us to make significant progress during each augmentation step. Figure 3-1 is the high-level procedure for this algorithm. It consists of an initialization procedure INITIALIZE which is described in Section 3.3.3, repeated circulation augmentations using procedure AUGMENT as described in Section 3.3.4, and finally a standard procedure REPAIR to round the returned solution with low duality gap to an optimal integral one (described in Section 3.5). The faster algorithm of Section 3.4 will also follow the same format, but the AUGMENT and REPAIR routines will be more sophisticated.

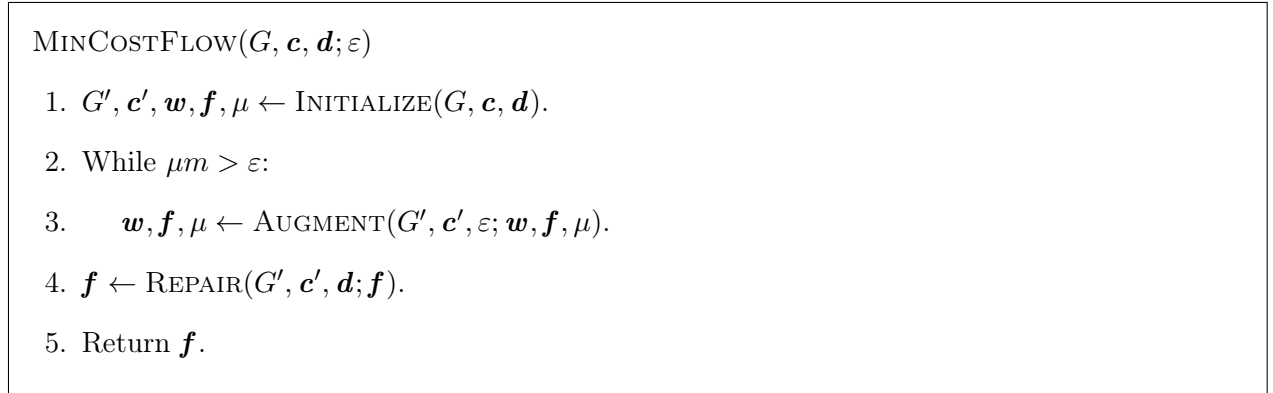


Figure 3-1: Minimum Cost Flow by Circulation Improvement

Barrier Formulation. In order to apply an interior point method on (3.1), we need to replace the feasibility constraints by a convex barrier function. We seek a nearly optimal solution, i.e. one that has small duality gap. The vanilla interior point method consists of iteratively finding the optimizer \mathbf{x}_μ for a family of functions parametrized by $\mu > 0$

$$\min_{\mathbf{x} \in \mathbb{R}^{m-n+1}} F_\mu^{\mathbf{w}}(\mathbf{x}) = \frac{1}{\mu} \cdot \langle \mathbf{c}, \mathbf{C}\mathbf{x} \rangle - \sum_{e \in E} (w_e^+ \cdot \log(\mathbf{1} - \mathbf{f}_0 - \mathbf{C}\mathbf{x})_e + w_e^- \cdot \log(\mathbf{f}_0 + \mathbf{C}\mathbf{x})_e) . \quad (3.3)$$

where $w_e^+, w_e^- > 0$ are weights on the flow capacity constraints. In order to find the optimizer \mathbf{x}_μ , one performs Newton method on $F_\mu^{\mathbf{w}}$, after warm starting with $\mathbf{x}_{\mu(1+\delta)}$ for some $\delta > 0$.

While classical methods maintain $\mathbf{w} = \mathbf{1}$ at all times, this extra parameter has been introduced in previous work in order to allow the method to make progress more aggressively. To simplify notation we define the slack vector $\mathbf{s} = (\mathbf{s}^+; \mathbf{s}^-)$ as

$$\mathbf{s}^+ = \mathbf{1} - \mathbf{f}_0 - \mathbf{C}\mathbf{x} , \quad (3.4)$$

$$\mathbf{s}^- = \mathbf{f}_0 + \mathbf{C}\mathbf{x} , \quad (3.5)$$

representing the upper slack (i.e. the distance of the current flow $\mathbf{f} = \mathbf{f}_0 + \mathbf{C}\mathbf{x}$ to the upper capacity constraint of $\mathbf{f} \leq \mathbf{1}$) and the lower slack (i.e. the distance from the current flow to the lower capacity constraint $\mathbf{0} \leq \mathbf{f}$). We will use the vector $\mathbf{w} = (\mathbf{w}^+; \mathbf{w}^-)$ to represent the weights for the two sets of barriers that we are using.

3.3.2 Optimality and Duality Gap

In order to describe the method and analyze it, it is important to understand the optimality conditions for F_μ^w . We say that a vector \mathbf{x} which minimizes F_μ^w is *central* (or satisfies *centrality*). This condition is described in the following lemma, whose proof can be found in Appendix 9.1.5.

Lemma 3.3.1. *The vector \mathbf{x} is a minimizer for F_μ^w if and only if*

$$\mathbf{C}^\top \left(\frac{\mathbf{w}^+}{\mathbf{s}^+} - \frac{\mathbf{w}^-}{\mathbf{s}^-} \right) = -\frac{\mathbf{C}^\top \mathbf{c}}{\mu}. \quad (3.6)$$

Furthermore the vector $\mathbf{y} = (\mathbf{y}^+; \mathbf{y}^-)$ with $\mathbf{y}^+ = \mu \cdot \frac{\mathbf{w}^+}{\mathbf{s}^+}$, $\mathbf{y}^- = \mu \cdot \frac{\mathbf{w}^-}{\mathbf{s}^-}$ is a feasible dual vector, and the duality gap of the primal-dual solution (\mathbf{x}, \mathbf{y}) is exactly $\mu \|\mathbf{w}\|_1$.

Maintaining the centrality condition (3.6) will be the key challenge in obtaining a faster interior point method for this linear program. This emphasizes the fact that the aim of this method is to construct a feasible set of slacks $\mathbf{s}^+ = \mathbf{1} - \mathbf{f}_0 - \mathbf{C}\mathbf{x} > \mathbf{0}$ and $\mathbf{s}^- = \mathbf{f}_0 + \mathbf{C}\mathbf{x} > \mathbf{0}$ such that $\mathbf{C}^\top \left(\frac{\mathbf{w}^+}{\mathbf{s}^+} - \frac{\mathbf{w}^-}{\mathbf{s}^-} \right) = -\frac{\mathbf{C}^\top \mathbf{c}}{\mu}$ for a very small $\mu > 0$. It is important to note that even though the existence of such an \mathbf{x} needs to be guaranteed, it is not necessary to explicitly maintain it. This will be apparent in the definition below.

Definition 3.3.2 (μ -central flow). *Given weights $\mathbf{w} = (\mathbf{w}^+; \mathbf{w}^-) \in \mathbb{R}_{>0}^{2m}$, a flow $\mathbf{0} < \mathbf{f} < \mathbf{1}$ is called μ -central with respect to \mathbf{w} if for some cycle basis matrix $\mathbf{C} \in \mathbb{R}^{m \times (m-n+1)}$,*

$$\mathbf{C}^\top \left(\frac{\mathbf{w}^+}{\mathbf{1} - \mathbf{f}} - \frac{\mathbf{w}^-}{\mathbf{f}} \right) = -\frac{\mathbf{C}^\top \mathbf{c}}{\mu} \quad (3.7)$$

for some $\mu > 0$. We will call the parameter μ the centrality of \mathbf{f} with respect to \mathbf{w} .

It should be noted that the precise choice of cycle basis \mathbf{C} in the above definition is irrelevant, as the property is invariant under the choice of cycle basis.

3.3.3 Initialization

The initialization procedure description and analysis is standard and thus deferred to Section 9.1.1. From now on we assume that G is the graph produced by the procedure in Section 9.1.1, together with a μ -central flow with $\mu \leq 2\|\mathbf{c}\|_2$.

3.3.4 The Augmentation Procedure and Routing the Residual

The AUGMENT procedure can be seen in Figure 3-2. It consists of computing an augmenting flow $\tilde{\mathbf{f}}$ by solving a linear system, augmenting the current solution by that flow, and finally calling a correction procedure in order to enforce the centrality of the solution.

We think of the interior point method as iteratively augmenting a feasible flow $\mathbf{0} < \mathbf{f} < \mathbf{1}$ satisfying $\mathbf{C}^\top \left(\frac{\mathbf{w}^+}{\mathbf{1} - \mathbf{f}} - \frac{\mathbf{w}^-}{\mathbf{f}} \right) = -\frac{\mathbf{C}^\top \mathbf{c}}{\mu}$ into another flow \mathbf{f}' such that $\mathbf{C}^\top \left(\frac{\mathbf{w}^+}{\mathbf{1} - \mathbf{f}'} - \frac{\mathbf{w}^-}{\mathbf{f}'} \right) = -\frac{\mathbf{C}^\top \mathbf{c}}{\mu'}$

AUGMENT($G, \mathbf{c}, \varepsilon; \mathbf{w}, \mathbf{f}, \mu$)

- Given \mathbf{f} : μ -central flow with respect to weights \mathbf{w} .
- Returns \mathbf{f}'' : μ' -central flow with respect to weights \mathbf{w}' .

1. Let $\tilde{\mathbf{f}}$ be a solution to (3.10), where

$$\mathbf{h} = \delta \left(\frac{\mathbf{w}^+}{\mathbf{1} - \mathbf{f}} - \frac{\mathbf{w}^-}{\mathbf{f}} \right).$$

2. Compute the congestion vector $\boldsymbol{\rho} = (\boldsymbol{\rho}^+; \boldsymbol{\rho}^-)$ as $\rho_e^+ = \frac{\tilde{f}_e}{1 - \tilde{f}_e}$, $\rho_e^- = \frac{-\tilde{f}_e}{\tilde{f}_e}$, cf. (3.11).

3. Augment flow $\mathbf{f}' = \mathbf{f} + \tilde{\mathbf{f}}$.

4. Correct residual given by $\mathbf{h}' = - \left(\frac{\mathbf{w}^+}{\mathbf{1} - \mathbf{f}'} - \frac{\mathbf{w}^-}{\mathbf{f}'} + \frac{\mathbf{c}}{\mu/(1+\delta)} \right)$ using Lemma 3.3.9 and get new weights \mathbf{w}' , flow \mathbf{f}'' , and centrality parameter μ' .

5. Return $\mathbf{w}', \mathbf{f}'', \mu'$.

Figure 3-2: Circulation improvement step

where $\mu' = \mu/(1 + \delta)$ for some $\delta > 0$. The magnitude of δ for which we are able to do so dictates the rate of the convergence of the method, since in $O(\delta^{-1})$ such iterative steps, the parameter μ is reduced to half, and hence the duality gap also reduces by a factor of $1/2$, per Lemma 3.3.1.

We can interpret such an iterative step as a residual-fixing procedure. Given a flow \mathbf{f} with slacks $\mathbf{s}^+ = \mathbf{1} - \mathbf{f} > \mathbf{0}$, $\mathbf{s}^- = \mathbf{f} > \mathbf{0}$, we consider the residual

$$\mathbf{C}^\top \left(\frac{\mathbf{w}^+}{\mathbf{s}^+} - \frac{\mathbf{w}^-}{\mathbf{s}^-} + \frac{\mathbf{c}}{\mu'} \right) = \nabla F_{\mu'}^{\mathbf{w}}(\mathbf{x}),$$

which is exactly the amount by which the target condition (3.6) fails to be satisfied. We denote the current residual as $\nabla F_{\mu'}^{\mathbf{w}}(\mathbf{x}) = -\mathbf{C}^\top \mathbf{h}$ for some $\mathbf{h} \in \mathbb{R}^m$. The goal of the iterative step is therefore to provide a feasible update rule to the flow such that the residual shrinks in some metric. To do so, we must first define a few useful notions.

Definition 3.3.3 (Energy of the residual). *Given a residual $-\mathbf{C}^\top \mathbf{h}$ for some $\mathbf{h} \in \mathbb{R}^m$ and positive vectors $\mathbf{w} = (\mathbf{w}^+; \mathbf{w}^-) \in \mathbb{R}^{2m}$, $\mathbf{s} = (\mathbf{s}^+; \mathbf{s}^-) \in \mathbb{R}^{2m}$ we define the energy to route \mathbf{h} with resistances determined by (\mathbf{w}, \mathbf{s}) as*

$$\mathcal{E}_{\mathbf{w}, \mathbf{s}}(\mathbf{h}) = \min_{\tilde{\mathbf{y}}: \mathbf{C}^\top(\tilde{\mathbf{y}} + \mathbf{h}) = \mathbf{0}} \frac{1}{2} \sum_{e \in E} \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right)^{-1} (\tilde{y}_e)^2. \quad (3.8)$$

The following lemma gives equivalent formulations for the energy which will be useful for our analysis. Its proof can be found in Appendix 9.1.5.

Lemma 3.3.4 (Equivalent energy formulations). *Given $\mathbf{h} \in \mathbb{R}^m$ and $\mathbf{w}, \mathbf{s} \in \mathbb{R}^{2m}$, one can write:*

$$\mathcal{E}_{\mathbf{w}, \mathbf{s}}(\mathbf{h}) = \max_{\tilde{\mathbf{f}} = \mathbf{C}\tilde{\mathbf{x}}} \langle \mathbf{h}, \tilde{\mathbf{f}} \rangle - \frac{1}{2} \sum_{e \in E} \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) (\tilde{f}_e)^2. \quad (3.9)$$

Equivalently, the following conditions are satisfied:

$$\begin{aligned} \tilde{\mathbf{f}} &= \mathbf{C}\tilde{\mathbf{x}}, \\ \mathbf{C}^\top \left(\frac{\mathbf{w}^+}{(\mathbf{s}^+)^2} + \frac{\mathbf{w}^-}{(\mathbf{s}^-)^2} \right) \tilde{\mathbf{f}} &= \mathbf{C}^\top \mathbf{h}. \end{aligned} \quad (3.10)$$

The latter equality can be re-stated in terms of the congestion vector

$$\boldsymbol{\rho} := \left(\rho^+ = \frac{\tilde{\mathbf{f}}}{\mathbf{s}^+}; \rho^- = \frac{-\tilde{\mathbf{f}}}{\mathbf{s}^-} \right) \quad (3.11)$$

as

$$\mathbf{C}^\top \left(\frac{\mathbf{w}^+ \boldsymbol{\rho}^+}{\mathbf{s}^+} - \frac{\mathbf{w}^- \boldsymbol{\rho}^-}{\mathbf{s}^-} \right) = \mathbf{C}^\top \mathbf{h}. \quad (3.12)$$

The energy can then be written as

$$\mathcal{E}_{\mathbf{w}, \mathbf{s}}(\mathbf{h}) = \frac{1}{2} \sum_{e \in E} (w_e^+ (\rho_e^+)^2 + w_e^- (\rho_e^-)^2) = \frac{1}{2} \sum_{e \in E} \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) (\tilde{f}_e)^2.$$

Using Lemma 3.3.4 we can prove that under certain conditions we can update the flow while simultaneously maintaining its feasibility and reducing the residual. Let us first define a residual correction step.

Definition 3.3.5 (Residual correction). *Let $\mathbf{w}, \mathbf{s} \in \mathbb{R}^{2m}$ where $\mathbf{w} > \mathbf{0}$ and $\mathbf{f} = \mathbf{f}_0 + \mathbf{C}\mathbf{x}$ is a feasible flow with slacks $\mathbf{s}^+ = \mathbf{1} - \mathbf{f} > \mathbf{0}$, $\mathbf{s}^- = \mathbf{f} > \mathbf{0}$, and let $\nabla F_\mu^{\mathbf{w}}(\mathbf{x}) = -\mathbf{C}^\top \mathbf{h}$ be the corresponding residual. A residual correction step is defined as an update to the \mathbf{x} vector, and implicitly to the flow vector \mathbf{f} via:*

$$\mathbf{x}' = \mathbf{x} + \tilde{\mathbf{x}}, \quad (3.13)$$

$$\mathbf{f}' = \mathbf{f} + \mathbf{C}\tilde{\mathbf{x}} = \mathbf{f} + \tilde{\mathbf{f}}, \quad (3.14)$$

where $\tilde{\mathbf{x}}$ is the solution to the linear system

$$\tilde{\mathbf{f}} = \mathbf{C}\tilde{\mathbf{x}}, \quad (3.15)$$

$$\mathbf{C}^\top \left(\frac{\mathbf{w}^+}{(\mathbf{s}^+)^2} + \frac{\mathbf{w}^-}{(\mathbf{s}^-)^2} \right) \tilde{\mathbf{f}} = \mathbf{C}^\top \mathbf{h}. \quad (3.16)$$

Since $\tilde{\mathbf{f}}$ is a circulation, this shows that the residual correction steps of the vanilla interior point method preserve the demand \mathbf{d} by adding an augmenting circulation to the current flow \mathbf{f} . It is also important to ensure that such updates do not break the LP feasibility constraints, i.e. $\mathbf{0} \leq \mathbf{f} \leq \mathbf{1}$ at all times. This will be made true by appropriately scaling the residual, thus enforcing $\|\boldsymbol{\rho}\|_\infty < 1/4$, or equivalently $|\tilde{f}_e| < \frac{1}{4} \min\{1 - f_e, f_e\}$ for all $e \in E$.

It is worth noting that the flow $\tilde{\mathbf{f}}$ which corresponds to solving the linear system from (3.13-3.14) can be computed in $\tilde{O}(m)$ time using a fast Laplacian solver [155, 98, 41, 137]. This may not be immediately obvious given the cycle basis formulation. But, in fact, this follows very easily from writing the linear system solve as a convex quadratic minimization problem. We do not go into further detail here, since we will elaborate more on this topic in Section 3.4.

In order to analyze the algorithm, we use energy as a potential function.

Lemma 3.3.6 (Energy after residual correction). *Let $\mathbf{w} \in \mathbb{R}^{2m}$ where $\mathbf{w} > \mathbf{0}$, and $\mathbf{f} = \mathbf{f}_0 + \mathbf{C}\mathbf{x}$ be a flow vector with slacks $\mathbf{s} > \mathbf{0}$. Let $\mathbf{c} \in \mathbb{R}^m$ and the residual $-\mathbf{C}^\top \mathbf{h} = \nabla F_\mu^{\mathbf{w}}(\mathbf{x}) = \mathbf{C}^\top \left(\frac{\mathbf{w}^+}{\mathbf{s}^+} - \frac{\mathbf{w}^-}{\mathbf{s}^-} + \frac{\mathbf{c}}{\mu} \right)$ for some $\mu > 0$ and $\boldsymbol{\rho} \in \mathbb{R}^{2m}$ be the corresponding congestion vector. Then a residual correction step produces a new flow $\mathbf{f}' = \mathbf{f}_0 + \mathbf{C}\mathbf{x}'$ with slacks \mathbf{s}' and residual $-\mathbf{C}^\top \mathbf{h}' = \nabla F_\mu^{\mathbf{w}}(\mathbf{x}') = \mathbf{C}^\top \left(\frac{\mathbf{w}^+}{(\mathbf{s}^+)' } - \frac{\mathbf{w}^-}{(\mathbf{s}^-)' } + \frac{\mathbf{c}}{\mu} \right)$ such that*

$$\mathcal{E}_{\mathbf{w}, \mathbf{s}'}(\mathbf{h}') \leq \frac{1}{2} \sum_{e \in E} (w_e^+ (\rho_e^+)^4 + w_e^- (\rho_e^-)^4).$$

The proof can be found in Appendix 9.1.5. As a corollary, we show that if the energy required to route the residual is small to begin with, after performing a step from Lemma 3.3.6, it quickly contracts.

Corollary 3.3.7. *Let $\mathbf{w} \in \mathbb{R}^{2m}$ where $\mathbf{w} \geq \mathbf{1}$, and $\mathbf{f} = \mathbf{f}_0 + \mathbf{C}\mathbf{x}$ be a flow vector with slacks $\mathbf{s} > \mathbf{0}$ and residual $\nabla F_\mu^{\mathbf{w}}(\mathbf{x}) = -\mathbf{C}^\top \mathbf{h}$. Updating \mathbf{x} and \mathbf{f} via a residual correction step as in Lemma 3.3.6 yields a new flow $\mathbf{f}' = \mathbf{f}_0 + \mathbf{C}\mathbf{x}'$ with slacks $\mathbf{s}' > \mathbf{0}$ and residual $\nabla F_\mu^{\mathbf{w}}(\mathbf{x}') = -\mathbf{C}^\top \mathbf{h}'$ such that*

$$\mathcal{E}_{\mathbf{w}, \mathbf{s}'}(\mathbf{h}') \leq 2 \cdot \mathcal{E}_{\mathbf{w}, \mathbf{s}}(\mathbf{h})^2.$$

We defer the proof to Appendix 9.1.5.

Corollary 3.3.7 shows that with a good initialization, residual correction steps decrease the energy of the residual very fast. In other words, starting from a μ -central flow \mathbf{f} one can quickly obtain a μ' -central flow \mathbf{f}' with a smaller duality gap, i.e. with $\mu' < \mu$. This is the workhorse of the vanilla interior point method that we proceed to analyze in Section 3.3.5.

Before that, we introduce an additional residual correction step that ensures that our solution is always *exactly* central. Since residual correction reduces the residual very fast, i.e. energy gets reduced to ε in $O(\log \log \varepsilon^{-1})$ steps, we can intuitively think of it as a method which effectively removes the residual in $\tilde{O}(1)$ steps. To make this intuition rigorous, we first reduce the residual energy to $m^{-O(1)}$, then we force optimality conditions by changing the weights \mathbf{w} . While this perfect correction step is generally unnecessary, it will make the description and analysis of our algorithm somewhat cleaner since we will always be able to assume exact centrality.

Lemma 3.3.8 (Perfect correction). *Let $\mathbf{w} \in \mathbb{R}^{2m}$ be a set of weights such that $\mathbf{1} \leq \mathbf{w}$, and let $\mathbf{f} = \mathbf{f}_0 + \mathbf{C}\mathbf{x}$ be a flow vector with slacks $\mathbf{s} > \mathbf{0}$ and residual $\nabla F_\mu^{\mathbf{w}}(\mathbf{x}) = -\mathbf{C}^\top \mathbf{h}$ such that $\mathcal{E}_{\mathbf{w}, \mathbf{s}}(\mathbf{h}) \leq \varepsilon \leq 1/100$. Then one can compute weights $\mathbf{w}' \in \mathbb{R}^{2m}$ such that $\mathbf{w} \leq \mathbf{w}'$ and $\|\mathbf{w}' - \mathbf{w}\|_1 \leq \|\mathbf{w}\|_1 \cdot 4\sqrt{\varepsilon}$ for which the residual $\nabla F_{\mu'}^{\mathbf{w}'}(\mathbf{x}) = \mathbf{0}$, where $\mu' \leq \mu(1 + 2\sqrt{\varepsilon})$.*

The complete proof can be found in Appendix 9.1.5. This lemma shows that after performing residual correction until the energy becomes smaller than $m^{-20}/4$, we can slightly increase the weights from \mathbf{w} to \mathbf{w}' such that for the new objective, \mathbf{x} exactly satisfies the optimality condition, as in Equation (3.6). The effect of this perfect correction is an extremely small increase in the sum of weights and the current duality gap. While this step is not essential, it enables us to ensure that

for all the essential points in our analysis we are able to assume *exact* centrality, which comes at a negligible expense, but makes all of our proofs much cleaner. This is summarized in the following lemma, whose proof can be found in Appendix 9.1.5:

Lemma 3.3.9. *Suppose that $\nabla F_\mu^{\mathbf{w}}(\mathbf{x}) = -\mathbf{C}^\top \mathbf{g}$ for some vector \mathbf{g} , and*

$$\mathcal{E}_{\mathbf{w},s}(\mathbf{g}) \leq 1/4.$$

Then using $O(\log \log \|\mathbf{w}\|_1)$ iterations of a vanilla residual correction step, we can obtain a new instance with weights $\mathbf{w}' \geq \mathbf{w}$ and $\mu' \leq \mu(1 + \frac{1}{2}\|\mathbf{w}\|_1^{-11})$ such that

$$\nabla F_{\mu'}^{\mathbf{w}'}(\mathbf{x}') = \mathbf{0}.$$

and $\|\mathbf{w}' - \mathbf{w}\|_1 \leq \|\mathbf{w}\|_1^{-10}$.

3.3.5 Vanilla Interior Point Method

At this point we are ready to describe a basic interior point method, which requires $\tilde{O}(m^{1/2})$ iterations. The following lemma shows that once we have a μ -central flow, we can scale down μ by a significant factor such that the new residual can be routed with low energy. The proof appears in Appendix 9.1.5.

Lemma 3.3.10. *Let $\mathbf{f} = \mathbf{f}_0 + \mathbf{C}\mathbf{x}$ be a μ -central flow with respect to weights $\mathbf{w} \in \mathbb{R}^{2m}$, and with slacks \mathbf{s} . Let $\delta = \frac{1}{(2\|\mathbf{w}\|_1)^{1/2}}$, $\mu' = \mu/(1 + \delta)$, and the corresponding residual $\nabla F_{\mu'}^{\mathbf{w}}(\mathbf{x}) = -\mathbf{C}^\top \mathbf{h}'$. Then one has that*

$$\mathcal{E}_{\mathbf{w},s}(\mathbf{h}') \leq 1/4.$$

Together with Corollary 3.3.7 this enables us to recover a simple analysis of the classical $\tilde{O}(m^{1/2})$ iteration bound. This is shown in the following lemma, whose proof appears in Appendix 9.1.5.

Lemma 3.3.11. *Given a μ^0 -central flow with respect to weights $\mathbf{1}$ and $\mu^0 = m^{O(1)}$, we can obtain a minimum cost flow solution with duality gap at most $\varepsilon = 1/m^{O(1)}$ using $\tilde{O}(m^{1/2})$ calls to the residual correction procedure (Definition 3.3.5).*

As previously discussed, each iteration of the interior point method can be implemented in $\tilde{O}(m)$ time using fast Laplacian solvers. This carries over to an algorithm with a total running time of $\tilde{O}(m^{3/2} \log W)$, matching that of previous classical algorithms.

As we saw in the proof above the major shortcoming of this method consists of only being able to scale down the parameter μ by $1 + 1/\Omega(m^{1/2})$ in every step, the reason being that while advancing from μ to a smaller μ' we only rely on Lemma 3.3.10 to certify the fact that the new residual can be corrected.

Instead, in what follows we choose to employ a stronger method to produce a new point which satisfies centrality and has a smaller parameter μ . To do so we employ a more sophisticated procedure for producing the new iterate, which also forces some more drastic changes in the weights \mathbf{w} .

Finally, we make an important observation concerning the vectors that need to be maintained throughout the algorithm.

Observation 3.3.12. *A linear system solving oracle which returns the vector $\mathbf{C}\tilde{\mathbf{x}}$ rather than $\tilde{\mathbf{x}}$ when solving the linear system described in Equations (3.11), (3.12) suffices to execute the algorithm. This is because we never need to maintain the explicit solution \mathbf{x} but rather only the flow $\mathbf{f} = \mathbf{C}\mathbf{x}$ and corresponding slacks, as the interior point method works in the space of slacks.*

3.4 A Faster Algorithm for Minimum Cost Flow

Our improved algorithm will be based on the interior point method framework that was developed in Section 3.3. The main bottleneck for the running time of that algorithm stems from the fact that the augmenting circulation we compute might not allow us to decrease the duality gap by more than a factor of $1 + 1/\Omega(\sqrt{m})$, as otherwise it is generally impossible to guarantee that the circulation will never congest some edges by more than the available capacity. Hence the iteration bound of $\tilde{O}(m^{1/2})$, common to standard interior point methods.

We alleviate this difficulty by adding an ℓ_p regularization term for the augmenting flow in the objective (3.9), similarly to [115]. In [115], the authors follow the idea of [117] by computing augmenting s - t flows. A crucial ingredient is the fact that the congestion of these resulting augmenting flows is then immediately bounded by using a result from [117] which states that as long as there is enough s - t residual capacity, these flows come together with an electrical potential embedding, where no edge is too stretched.

However, this property is specific to the s - t maximum flow problem. To apply a similar argument for the minimum cost flow problem, one would need to guarantee that *all* cuts of the graph have sufficient residual capacity, which is not automatically enforced as in the case of s - t max flow. In order to enforce this cut property, we further regularize our objective in a different way. We do this by temporarily superimposing a star on top of our graph, thus obtaining an augmented graph. This transformation improves the conductance properties of the graph, ensuring that there is enough residual capacity in all cuts of the graph.

In Section 3.4.1, we describe the regularized step problem and outline the guarantees of the solution. In particular, the bias introduced by the regularizers implies that the augmenting flow is not a circulation anymore, and that we have introduced an additional residual for our solution in the barrier objective. We bound the magnitude of both of these perturbations and “undo” them at a later stage. Finally, we present our electrical stretch guarantee, which serves as the crucial ingredient in both preserving feasibility and maintaining centrality.

In Section 3.4.2 we state our choice of regularization parameters and their consequences.

Even though the electrical stretch guarantee suffices for all purposes if the interior point method barrier terms are unweighted, as soon as weights come in the guarantee is affected. In particular, for any edge whose forward and backward weights are too imbalanced, the electrical stretch and congestion bounds that we obtain loosen. In Section 3.4.3 we deal with this issue by ensuring that the forward and backward weights for each edge are always relatively balanced, while introducing an additional demand perturbation.

In Section 3.4.4 we provide the full view of the algorithm, which consists of combining all the ingredients of the previous sections, together with a residual routing scheme that includes both vanilla centering steps and constraint re-weighting to obtain an ℓ_∞ -based interior point method rather than an ℓ_4 -based one, as achieved by the vanilla algorithm.

As we mentioned, the solution obtained by the interior point method is for a minimum cost flow problem with a slightly perturbed demand. In Section 3.5, we outline an approach given in [42] that given this solution, one can turn it into an optimal solution for the original demand, as long as the total demand perturbation is small.

3.4.1 Regularized Newton Step

The initialization procedure from Section 3.3.3 produces a solution with large duality gap, i.e. $O(\mu m)$ where $\mu \leq 2\|\mathbf{c}\|_2$. Our goal will be to reduce this by gradually lowering the parameter μ , while maintaining centrality. While in general to achieve this we require solving a sequence of linear

MODIFIEDAUGMENT($G, \mathbf{c}, \varepsilon; \mathbf{w}, \mathbf{f}, \mu$)

- Given \mathbf{f} : μ -central flow with respect to weights \mathbf{w} .
- Returns \mathbf{f}'' : μ' -central flow with respect to weights \mathbf{w}''' and with perturbed demand.

1. $\mathbf{w}, \mathbf{f} \leftarrow \text{BALANCEWEIGHTS}(G; \mathbf{w}, \mathbf{f})$.

2. Augment graph G to G_* .

3. Let $\tilde{\mathbf{f}}$ be a minimizer to (3.18) where $h_e = \delta \left(\frac{w_e^+}{1-f_e} - \frac{w_e^-}{f_e} \right)$ for each $e \in E$ and $\Delta \mathbf{h}$ be the residual perturbation.

4. Augment flow $\mathbf{f}' \leftarrow \mathbf{f} + \tilde{\mathbf{f}}$.

5. Compute congestion vector $\boldsymbol{\rho} = (\boldsymbol{\rho}^+; \boldsymbol{\rho}^-)$ as $\rho_e^+ = \frac{\tilde{f}_e}{1-f_e}$, $\rho_e^- = \frac{-\tilde{f}_e}{f_e}$.

6. Compute new slacks $(\mathbf{s}^+)' = \mathbf{1} - \mathbf{f}'$ and $(\mathbf{s}^-)' = \mathbf{f}'$.

7. Correct residual for congested edges:

$$(w_e^+)' = \begin{cases} w_e^+ + \frac{(s_e^+)'}{(s_e^-)'} \cdot w_e^- (\rho_e^-)^2, & \text{if } |\rho_e^-| \geq C_\infty, \\ w_e^+, & \text{otherwise,} \end{cases}$$

$$(w_e^-)' = \begin{cases} w_e^- + \frac{(s_e^-)'}{(s_e^+)'} \cdot w_e^+ (\rho_e^+)^2, & \text{if } |\rho_e^+| \geq C_\infty, \\ w_e^-, & \text{otherwise.} \end{cases}$$

8. Correct perturbed residual given by $\mathbf{h}' = - \left(\frac{(w^+)' }{(s^+)'} - \frac{(w^-)' }{(s^-)'} + \frac{\mathbf{c}}{\mu/(1+\delta)} + \Delta \mathbf{h} \right)$ using Lemma 3.4.19 and get new weights \mathbf{w}'' , flow \mathbf{f}'' , and centrality parameter μ' .

9. Compute new slacks $(\mathbf{s}^+)'' = \mathbf{1} - \mathbf{f}''$ and $(\mathbf{s}^-)'' = \mathbf{f}''$.

10. Adjust weights to restore exact centrality:

$$(w_e^+)''' = (w_e^+)'' + \max \{0, -(s_e^+)'' \cdot \Delta h_e\},$$

$$(w_e^-)''' = (w_e^-)'' + \max \{0, (s_e^-)'' \cdot \Delta h_e\}.$$

11. Return $\mathbf{w}''', \mathbf{f}'', \mu'$.

Figure 3-3: Modified circulation improvement step

systems of equations (as we saw in Section 3.3.5), here we choose to solve a slightly perturbed linear system.

In order to do so, we modify the optimization problem from (3.9) by adding two regularization terms, which will force the produced solution to be well-behaved. In addition, we allow the newly produced flow $\tilde{\mathbf{f}}$, which we will use to update the current solution, to not be a circulation, as long as the demand it routes is small in ℓ_1 norm. While this breaks the structure of the problem we are solving, it only does so mildly – therefore once the interior point method has finished running we can repair the broken demand using combinatorial techniques.

Mixed Objective. To specify the regularized objective, we first augment the graph G with $O(m)$ extra edges, which are responsible for routing a subset of the flows that would otherwise force the output of the objective to be too degenerate.

Definition 3.4.1 (Weighted degree). *Given a graph $G(V, E)$ and a weight vector $\mathbf{w} = (\mathbf{w}^+; \mathbf{w}^-) \in \mathbb{R}^{2m}$, the weighted degree of $v \in V$ in G with respect to \mathbf{w} is defined as $d_v^{\mathbf{w}} = \sum_{e \sim v} (w_e^+ + w_e^-)$.*

Definition 3.4.2. *Given a graph $G = (V, E)$ we define the augmented graph $G_\star = (V \cup \{v_\star\}, E_\star)$, where $E_\star = E \cup E'$ and E' is obtained by constructing $\lceil d_v^{\mathbf{w}} \rceil$ parallel edges (v, v_\star) for each $v \in V$.*

Furthermore, if \mathbf{C} is a cycle basis for G , we let \mathbf{C}_\star be a cycle basis for G_\star obtained by appending columns to \mathbf{C} , i.e.

$$\mathbf{C}_\star = \begin{bmatrix} \mathbf{C} & \mathbf{P}_1 \\ \mathbf{0} & \mathbf{P}_2 \end{bmatrix}. \quad (3.17)$$

We observe that $|E'| = \sum_{v \in V} \lceil d_v^{\mathbf{w}} \rceil \leq \sum_{v \in V} (d_v^{\mathbf{w}} + 1) \leq 3\|\mathbf{w}\|_1$. We can now write the regularized objective.

Definition 3.4.3. *Given a vector \mathbf{h} , we define the regularized objective as*

$$\begin{aligned} \max_{\tilde{\mathbf{f}} = \mathbf{C}_\star \tilde{\mathbf{x}}} \sum_{e \in E} h_e \cdot \tilde{f}_e - \frac{1}{2} \sum_{e \in E} (\tilde{f}_e)^2 \cdot \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) \\ - \frac{R_\star}{2} \sum_{e \in E'} (\tilde{f}_e)^2 - \frac{R_p}{p} \sum_{e \in E \cup E'} (\tilde{f}_e)^p, \end{aligned} \quad (3.18)$$

where $p > 2$ is an even positive integer, and R_\star, R_p are some appropriately chosen non-negative scalars.

While this objective might seem difficult to handle, the fact that we are solving a problem on graphs makes it feasible for our purposes. In particular, the works of [103, 2] show that this objective can be solved to high precision in time $O(m^{1+o(1)})$, whenever p is sufficiently large. We will make this statement more rigorous, but for simplicity let us for now assume that we can solve (3.18) exactly. In Appendix 9.1.3 we will show how to handle the solver error.

Let us now understand the effect of the augmenting edges E' . Since they allow routing some of the flow through v_\star , if we look at the restriction of $\tilde{\mathbf{f}}$ to the edges of G we see that it stops being a circulation. Let $\tilde{\mathbf{d}}$ be the demand routed by the restriction of $\tilde{\mathbf{f}}$ to G . We will see that $\tilde{\mathbf{f}}$ satisfies optimality conditions for an objective similar to (3.18) among all flows that route the demand $\tilde{\mathbf{d}}$ in G .

Before that, we give a useful lemma that, given a residual $-\mathbf{C}^\top \mathbf{h}$, can be used to certify an upper bound on the energy required to route it. We capture this via the following definition.

Definition 3.4.4. Given a vector \mathbf{h} , weights \mathbf{w} , and slacks \mathbf{s} , we define

$$\mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}) = \frac{1}{2} \sum_{e \in E} h_e^2 \cdot \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right)^{-1}. \quad (3.19)$$

Lemma 3.4.5. Given weights \mathbf{w} , slacks \mathbf{s} , and a residual $-\mathbf{C}^\top \mathbf{h}$, we have that

$$\mathcal{E}_{\mathbf{w}, \mathbf{s}}(\mathbf{h}) \leq \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}).$$

Furthermore, if $\mathbf{h} = \delta \left(\frac{\mathbf{w}^+}{s^+} - \frac{\mathbf{w}^-}{s^-} \right)$, we have that

$$\mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}) \leq \frac{1}{2} \delta^2 \|\mathbf{w}\|_1.$$

The proof of this lemma is given in Appendix 9.1.5. We are now ready to state the lemma that gives guarantees for the restriction of $\tilde{\mathbf{f}}$ to G .

Lemma 3.4.6 (Optimality in the non-augmented graph). Let $\tilde{\mathbf{f}}_\star = \mathbf{C}_\star \tilde{\mathbf{x}}_\star$ be the optimizer of the regularized objective from (3.18), and let $\tilde{\mathbf{f}}$ be its restriction to the edges of G . Let $\tilde{\mathbf{d}}$ be the demand routed by $\tilde{\mathbf{f}}$ in G . Then $\tilde{\mathbf{f}}$ optimizes the objective

$$\max_{\substack{\tilde{\mathbf{f}}: \\ \tilde{\mathbf{f}} \text{ routes } \tilde{\mathbf{d}} \text{ in } G}} \langle \mathbf{h}, \tilde{\mathbf{f}} \rangle - \frac{1}{2} \sum_{e \in E} (\tilde{f}_e)^2 \cdot \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) - \frac{R_p}{p} \sum_{e \in E} (\tilde{f}_e)^p. \quad (3.20)$$

Furthermore

$$\mathbf{C}^\top \left(\frac{\mathbf{w}^+}{(s^+)^2} + \frac{\mathbf{w}^-}{(s^-)^2} \right) \cdot \tilde{\mathbf{f}} = \mathbf{C}^\top (\mathbf{h} + \Delta \mathbf{h}), \quad (3.21)$$

where $\Delta \mathbf{h} = -R_p (\tilde{\mathbf{f}})^{p-1}$, and

$$\|\tilde{\mathbf{d}}\|_1 \leq \left(\frac{6 \|\mathbf{w}\|_1 \cdot \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s})}{R_\star} \right)^{1/2}, \quad (3.22)$$

$$\|\tilde{\mathbf{f}}_\star\|_p \leq \left(\frac{p \cdot \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s})}{R_p} \right)^{1/p}. \quad (3.23)$$

Finally, the energy required to route the perturbed residual can be bounded by the energy required to route the original residual:

$$\frac{1}{2} \sum_{e \in E} (\tilde{f}_e)^2 \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) \leq 4 \cdot \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}). \quad (3.24)$$

Proof. Let $\tilde{\mathbf{f}}_\star = \tilde{\mathbf{f}} + \tilde{\mathbf{f}}'$ where $\tilde{\mathbf{f}}'$ is the restriction of $\tilde{\mathbf{f}}_\star$ to the edges incident to v_\star . By computing the first order derivative in $\tilde{\mathbf{x}}_\star$, optimality conditions for (3.18) imply that for any circulation \mathbf{g} in G_\star

one has that

$$\left\langle \mathbf{g}, \begin{bmatrix} \mathbf{h} - \tilde{\mathbf{f}} \cdot \left(\frac{\mathbf{w}^+}{(s^+)^2} + \frac{\mathbf{w}^-}{(s^-)^2} \right) - R_p \cdot (\tilde{\mathbf{f}})^{p-1} \\ -R_\star \cdot \tilde{\mathbf{f}}' - R_p \cdot (\tilde{\mathbf{f}}')^{p-1} \end{bmatrix} \right\rangle = 0. \quad (3.25)$$

Therefore, restricting ourselves to circulations supported only in the non-preconditioned graph G , one has that for any circulation in $\mathbf{g}' = \mathbf{C}\mathbf{z}$ in G :

$$\left\langle \mathbf{g}', \mathbf{h} - \tilde{\mathbf{f}} \cdot \left(\frac{\mathbf{w}^+}{(s^+)^2} + \frac{\mathbf{w}^-}{(s^-)^2} \right) - R_p \cdot (\tilde{\mathbf{f}})^{p-1} \right\rangle = 0, \quad (3.26)$$

and equivalently

$$\left\langle \mathbf{z}, \mathbf{C}^\top \left(\mathbf{h} - \tilde{\mathbf{f}} \cdot \left(\frac{\mathbf{w}^+}{(s^+)^2} + \frac{\mathbf{w}^-}{(s^-)^2} \right) - R_p \cdot (\tilde{\mathbf{f}})^{p-1} \right) \right\rangle = 0. \quad (3.27)$$

Since this holds for any test vector \mathbf{z} , it must be that the second term in the inner product is 0. Rearranging, it yields the identity from (3.21).

Now we verify that this is the first order optimality condition for the objective in (3.20). Parametrizing the flows that route \mathbf{d} via $\mathbf{f} = \mathbf{C}\mathbf{x} + \tilde{\mathbf{f}}_d$ where $\tilde{\mathbf{f}}_d$ is an arbitrary flow which routes \mathbf{d} in G , and setting the derivative with respect to \mathbf{x} equal to $\mathbf{0}$, we obtain exactly (3.21).

Let us proceed to bound the norm of the demand routed by $\tilde{\mathbf{f}}$. Consider the value of the objective in (3.18) after truncating it to only the first two terms, which we can write as:

$$\sum_{e \in E} h_e \cdot \tilde{f}_e - \frac{1}{2} \sum_{e \in E} (\tilde{f}_e)^2 \cdot \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) \quad (3.28)$$

$$\leq \frac{1}{2} \sum_{e \in E} h_e^2 \cdot \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right)^{-1} \quad (3.29)$$

$$= \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}), \quad (3.30)$$

where we used the fact that $\langle \mathbf{a}, \mathbf{b} \rangle \leq \frac{1}{2} \|\mathbf{a}\|^2 + \frac{1}{2} \|\mathbf{b}\|^2$.

Note that the value of the regularized objective (3.18) is at least 0 since we can always substitute $\tilde{\mathbf{x}} = \mathbf{0}$ and obtain exactly 0. By re-arranging,

$$\frac{R_\star}{2} \sum_{e \in E'} (\tilde{f}'_e)^2 \quad (3.31)$$

$$\leq \sum_{e \in E} h_e \cdot \tilde{f}_e - \frac{1}{2} \sum_{e \in E} (\tilde{f}_e)^2 \cdot \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) - \frac{R_p}{p} \sum_{e \in E \cup E'} (\tilde{f}'_e)^p \quad (3.32)$$

$$\leq \sum_{e \in E} h_e \cdot \tilde{f}_e - \frac{1}{2} \sum_{e \in E} (\tilde{f}_e)^2 \cdot \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) \quad (3.33)$$

$$\leq \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}), \quad (3.34)$$

where we also used the fact that the last term of (3.32) is non-positive and (3.30). Therefore (3.34)

enables us to upper bound

$$\sum_{e \in E'} (\tilde{f}'_e)^2 \leq \frac{2}{R_\star} \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}), \quad (3.35)$$

which implies that

$$\sum_{e \in E'} |\tilde{f}'_e| \leq |E'|^{1/2} \cdot \left(\sum_{e \in E'} (\tilde{f}'_e)^2 \right)^{1/2} \quad (3.36)$$

$$\leq \left(3 \|\mathbf{w}\|_1 \cdot \frac{2}{R_\star} \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}) \right)^{1/2} \quad (3.37)$$

$$= \left(\frac{6 \|\mathbf{w}\|_1 \cdot \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s})}{R_\star} \right)^{1/2}, \quad (3.38)$$

a quantity that upper bounds the demand perturbation. Using a similar argument we can upper bound $\|\tilde{\mathbf{f}}_\star\|_p$. We have

$$\frac{R_p}{p} \sum_{e \in E \cup E'} (\tilde{f}_\star)_e^p \quad (3.39)$$

$$\leq \sum_{e \in E} h_e \cdot \tilde{f}_e - \frac{1}{2} \sum_{e \in E} (\tilde{f}_e)^2 \cdot \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) - \frac{R_\star}{2} \sum_{e \in E'} (\tilde{f}'_e)^2 \quad (3.40)$$

$$\leq \sum_{e \in E} h_e \cdot \tilde{f}_e - \frac{1}{2} \sum_{e \in E} (\tilde{f}_e)^2 \cdot \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) \quad (3.41)$$

$$\leq \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}), \quad (3.42)$$

thus concluding that

$$\|\tilde{\mathbf{f}}_\star\|_p = \left(\sum_{e \in E \cup E'} (\tilde{f}_\star)_e^p \right)^{1/p} \leq \left(\frac{p \cdot \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s})}{R_p} \right)^{1/p}. \quad (3.43)$$

Finally, once more using the same argument, we have that

$$\begin{aligned} & \frac{1}{4} \sum_{e \in E} (\tilde{f}_e)^2 \cdot \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) \\ & \leq \sum_{e \in E} h_e \cdot \tilde{f}_e - \frac{1}{4} \sum_{e \in E} (\tilde{f}_e)^2 \cdot \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) - \frac{R_\star}{2} \sum_{e \in E'} (\tilde{f}'_e)^2 - \frac{R_p}{p} \sum_{e \in E \cup E'} (\tilde{f}_\star)_e^p \\ & \leq \sum_{e \in E} h_e \cdot \tilde{f}_e - \frac{1}{4} \sum_{e \in E} (\tilde{f}_e)^2 \cdot \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) \\ & \leq 2 \cdot \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}), \end{aligned}$$

therefore

$$\frac{1}{2} \sum_{e \in E} (\tilde{f}_e)^2 \cdot \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) \leq 4 \cdot \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}).$$

□

Finally, we present an important property of the solution of the regularized Newton step, which will be crucial for obtaining the final result.

Lemma 3.4.7. *Let $\tilde{\mathbf{f}}_\star$ be the solution of the regularized objective (3.18) and $\tilde{\mathbf{f}}$ its restriction on G , and suppose that $\|\mathbf{w}\|_1 \geq 3$. Then one has that over the edges $e \in E$:*

$$\left| \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} + R_p \cdot \tilde{f}_e^{p-2} \right) \tilde{f}_e - h_e \right| \leq \hat{\gamma}, \quad (3.44)$$

where

$$\hat{\gamma} = \left(R_\star + R_p \cdot \|\tilde{\mathbf{f}}_\star\|_\infty^{p-2} \right)^{1/2} \cdot \left\| \frac{\mathbf{h}}{\sqrt{(\mathbf{w}^+ + \mathbf{w}^-) \left(\frac{\mathbf{w}^+}{(s^+)^2} + \frac{\mathbf{w}^-}{(s^-)^2} \right)}} \right\|_\infty \cdot 32 \log \|\mathbf{w}\|_1.$$

Furthermore, this implies that

$$\left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) \cdot |\tilde{f}_e| \leq |h_e| + \hat{\gamma}. \quad (3.45)$$

Since the proof is technical, we defer it to Section 9.1.2.

3.4.2 Choice of Regularization Parameters

Now we state our choice of regularization parameters R_p , R_\star and their consequences. In particular, they affect the ℓ_∞ norm of the flow $\tilde{\mathbf{f}}$ we obtain by solving the regularized objective, and the preconditioning guarantee (Lemma 3.4.7) which will be essential to obtain the correct trade-off between iteration complexity and demand perturbation.

Definition 3.4.8. *Given a weight vector \mathbf{w} , we will use the following values for the parameters p , R_p and R_\star :*

$$\begin{aligned} p &= \min \left\{ k \in 2\mathbb{Z} : k \geq (\log m)^{1/3} \right\}, \\ R_p &= p \cdot (10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \cdot \log \|\mathbf{w}\|_1)^{p+1}, \\ R_\star &= 3 \cdot \delta^2 \|\mathbf{w}\|_1^2. \end{aligned}$$

This immediately yields the following useful corollaries.

Corollary 3.4.9. *For our specific choice of p , the regularized objective from (3.18) can be solved to high precision in $m^{1+o(1)}$ time. We can, furthermore, assume an exact solution.*

Proof. We use Theorem 9.1.5 which was proved in [103, 2]. The objective in (3.18) matches exactly the type handled there. For our choice of p , the running time is

$$2^{O(p^{3/2})} m^{1+O(1/\sqrt{p})} = 2^{O\sqrt{\log n}} m^{1+1/(\log^{1/6} n)} = m^{1+o(1)}.$$

Furthermore, the resulting solution has a quasi-polynomially small error $2^{-(\log m)^{O(1)}}$, which can be neglected in our analysis, per the discussion in Section 9.1.3. \square

The proofs of the following two corollaries appear in Appendix 9.1.5 and 9.1.5, respectively.

Corollary 3.4.10. *Let $\tilde{\mathbf{f}}_\star$ be the solution obtained by solving (3.18) and $\tilde{\mathbf{f}}$ be its restriction to G . For our specific choice of regularization parameters we have*

$$\|\tilde{\mathbf{f}}\|_p \leq \frac{1}{10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \cdot \log \|\mathbf{w}\|_1}.$$

Corollary 3.4.11. *For the choice of $\mathbf{h} = \delta \left(\frac{\mathbf{w}^+}{s^+} - \frac{\mathbf{w}^-}{s^-} \right)$, and as long as $\delta \leq \|\mathbf{w}\|_1^{-(1/4+o(1))}$ the $\hat{\gamma}$ in Lemma 3.4.7 can be upper bounded by*

$$\gamma = \delta^2 \|\mathbf{w}\|_1 \cdot 32\sqrt{6} \cdot \log \|\mathbf{w}\|_1.$$

3.4.3 Weight Invariants

Total weight invariant. Our interior point method will inherently increase edge weights. However, the total increase has to remain bounded by $O(m)$ in order to be able to guarantee an upper bound on the energy $\mathcal{E}_{\mathbf{w},s}(\mathbf{h})$. We state the following invariant that we intend to always enforce, in order to ensure that this is the case:

Invariant 3.4.12. *The sum of weights is bounded: $\|\mathbf{w}\|_1 \leq 3m$.*

Weight balancing. Before proceeding with the description of the method, we define a notion that will be used by the algorithm to deal with severe edge weight imbalances. Such imbalances can limit the usability of Lemma 3.4.7 and lead to the augmenting flow $\tilde{\mathbf{f}}$ being infeasible or expensive to correct.

Definition 3.4.13 (Balanced edges). *An edge $e \in E$ is called balanced if $\max\{w_e^+, w_e^-\} \leq \delta \|\mathbf{w}\|_1$ or $\min\{w_e^+, w_e^-\} \geq 96 \cdot \delta^4 \|\mathbf{w}\|_1^2$. Otherwise it is called imbalanced.*

Even though imbalanced edges can generally emerge, we are able to balance them by manually reducing the disparity between w_e^+ and w_e^- , while breaking the flow demand by a controllable amount. We will maintain the following invariant right before solving the regularized objective:

Invariant 3.4.14. *All edges are balanced.*

In Section 3.4.4 we will see that there is a way to enforce Invariant 3.4.14 while only increasing the weight and breaking the demand by a small amount. The main motivation behind keeping edges balanced is that it implies that each edge either has a favorable stretch property, or is not too congested:

Lemma 3.4.15. *Let $\tilde{\mathbf{f}}$ be the restriction of the regularized problem (3.18) solution to G , with $\mathbf{h} = \delta \left(\frac{\mathbf{w}^+}{s^+} - \frac{\mathbf{w}^-}{s^-} \right)$ and congestion $\boldsymbol{\rho}$, where $\delta \leq \|\mathbf{w}\|_1^{-(1/4+o(1))}$. If $e \in E$ is balanced, then $\max\{|\rho_e^+|, |\rho_e^-|\} < C_\infty$ or $\left| \frac{w_e^+ \rho_e^+}{s_e^+} \right| + \left| \frac{w_e^- \rho_e^-}{s_e^-} \right| \leq 6\gamma$, where $C_\infty = \frac{1}{2\delta\sqrt{2}\|\mathbf{w}\|_1}$.*

The proof of this lemma appears in Appendix 9.1.5.

3.4.4 Executing the Interior Point Method

Having defined the regularized objective, we now show how to execute the interior point method using the solution returned by a high precision solver. Since the solution to this objective does not exactly match the guarantee required in (3.11-3.12), we will have to do some slight manual adjustments.

In the vanilla interior point method analysis that we saw earlier, we witnessed a very stringent requirement on the condition that we are able to correct a residual. Namely, we required that the energy required to route it decreases in every iteration of the correction step, which was guaranteed by the fact that after performing the first correction step the upper bound on energy $\sum w_i \rho_i^4$ is at most a small constant (i.e. 1/4).

This requirement is too strong since, as a matter of fact, the most important obstacle handled by interior point methods is preserving slack feasibility. In our specific context this means that we want to perform updates to the current flow without violating capacity constraints, which is guaranteed by a weaker ℓ_∞ bound, i.e. $\|\boldsymbol{\rho}\|_\infty \leq 1/2$. While this condition is sufficient to preserve feasibility, it is not clear that after performing the corresponding update to the flow, the energy required to route the residual will be small, so the resulting residual can be reduced to $\mathbf{0}$. Instead we can enforce this property by canceling the components of the gradient which cause this energy to be large.

Definition 3.4.16 (Perturbed residual correction). *Consider a flow \mathbf{f} with the corresponding slack vector $\mathbf{s} > \mathbf{0}$, weights \mathbf{w} and parameter $\mu > 0$, with a corresponding residual $\nabla F_\mu^{\mathbf{w}}(\mathbf{x}) = -\mathbf{C}^\top \mathbf{h}$ where $\mathbf{h} = \delta \left(\frac{\mathbf{w}^+}{\mathbf{s}^+} - \frac{\mathbf{w}^-}{\mathbf{s}^-} \right)$ and $\delta \leq \|\mathbf{w}\|_1^{-1/4}/2$. The perturbed residual correction step is defined as an update to \mathbf{f} via:*

$$\mathbf{f}' = \mathbf{f} + \tilde{\mathbf{f}}, \quad (3.46)$$

$$(\mathbf{s}^-)' = \mathbf{s}^- + \tilde{\mathbf{f}}, \quad (3.47)$$

$$(\mathbf{s}^+)' = \mathbf{s}^+ - \tilde{\mathbf{f}}, \quad (3.48)$$

where $\tilde{\mathbf{f}}$ is the solution to the linear system

$$\boldsymbol{\rho}^+ = \frac{\tilde{\mathbf{f}}}{\mathbf{s}^+}, \quad (3.49)$$

$$\boldsymbol{\rho}^- = \frac{-\tilde{\mathbf{f}}}{\mathbf{s}^-}, \quad (3.50)$$

$$\mathbf{C}^\top \left(\frac{\mathbf{w}^+ \boldsymbol{\rho}^+}{\mathbf{s}^+} - \frac{\mathbf{w}^- \boldsymbol{\rho}^-}{\mathbf{s}^-} \right) = \mathbf{C}^\top (\mathbf{h} + \Delta \mathbf{h}), \quad (3.51)$$

such that

$$\|\boldsymbol{\rho}\|_\infty \leq \frac{1}{2}, \quad (3.52)$$

for some perturbation $\Delta \mathbf{h}$, followed by the updates to the \mathbf{w} vector via:

$$(w_e^+)' = \begin{cases} w_e^+ + \frac{(s_e^+)'}{(s_e^-)'} \cdot w_e^- (\rho_e^-)^2 & \text{if } |\rho_e^-| \geq C_\infty, \\ w_e^+ & \text{otherwise,} \end{cases} \quad (3.53)$$

and

$$(w_e^-)' = \begin{cases} w_e^- + \frac{(s_e^-)'}{(s_e^+)'} \cdot w_e^+ (\rho_e^+)^2 & \text{if } |\rho_e^+| \geq C_\infty, \\ w_e^- & \text{otherwise,} \end{cases} \quad (3.54)$$

where

$$C_\infty = \frac{1}{2\delta\sqrt{2}\|\mathbf{w}\|_1}.$$

We also give a short lemma, which upper bounds the amount by which $\|\mathbf{w}\|_1$ gets increased when applying the perturbed residual correction.

Lemma 3.4.17. *Let \mathbf{w} and \mathbf{w}' be the old and new weights, respectively, as described in Definition 3.4.16. Then one has that*

$$\|\mathbf{w}' - \mathbf{w}\|_1 \leq 192\sqrt{2} \cdot \gamma \cdot (\delta^2\|\mathbf{w}\|_1)^{3/2}.$$

Proof. Let $S^+ = \{e \in E : |\rho_e^+| \geq C_\infty\}$ and $S^- = \{e \in E : |\rho_e^-| \geq C_\infty\}$. The total weight increase is equal to

$$\sum_{e \in S^+} ((w_e^-)' - w_e^-) + \sum_{e \in S^-} ((w_e^+)' - w_e^+).$$

Let us upper bound the weight increase contributed by a single edge $e \in S^+$. The respective bound will follow for an edge in S^- by symmetry. We have

$$(w_e^-)' - w_e^- = (s_e^-)' \cdot \frac{w_e^+ (\rho_e^+)^2}{(s_e^+)'} \leq 4 \cdot s_e^- \cdot \frac{w_e^+ (\rho_e^+)^2}{s_e^+} = 4 \cdot s_e^- \cdot |\rho_e^+| \cdot \left| \frac{w_e^+ \rho_e^+}{s_e^+} \right|.$$

By Lemma 9.20, all the edges in $S^+ \cup S^-$ have the low stretch property:

$$\left| \frac{w_e^+ \rho_e^+}{s_e^+} \right| + \left| \frac{w_e^- \rho_e^-}{s_e^-} \right| \leq 6\gamma,$$

and so

$$(w_e^-)' - w_e^- \leq 4 \cdot s_e^- \cdot |\rho_e^+| \cdot \left| \frac{w_e^+ \rho_e^+}{s_e^+} \right| \leq 24\gamma |\rho_e^+|.$$

The total weight increase due to S^+ is thus:

$$\sum_{e \in S^+} ((w_e^-)' - w_e^-) \leq 24\gamma \sum_{e \in S^+} |\rho_e^+|.$$

Symmetrically for S^- we get

$$\sum_{e \in S^-} ((w_e^+)' - w_e^+) \leq 24\gamma \sum_{e \in S^-} |\rho_e^-|,$$

and so

$$\|\mathbf{w}' - \mathbf{w}\|_1 \leq 24\gamma \cdot \|\boldsymbol{\rho}_{S^+ \cup S^-}\|_1.$$

Now, since by Lemma 3.4.6 the energy to route the perturbed residual is bounded by the energy to

route the original residual,

$$\|\rho\|_2^2 \leq \sum_{e \in E} (w_e^+ (\rho_e^+)^2 + w_e^- (\rho_e^-)^2) = \sum_{e \in E} (\tilde{f}_e)^2 \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) \leq 8 \cdot \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}),$$

we have that

$$\begin{aligned} \|\rho_{S^+ \cup S^-}\|_1 &\leq \|\rho\|_2^2 \cdot \max \left\{ \|1/\rho_{S^+}\|_\infty, \|1/\rho_{S^-}\|_\infty \right\} \\ &\leq 8 \cdot \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}) \cdot C_\infty^{-1} \\ &\leq 4 \cdot \delta^2 \|\mathbf{w}\|_1 \cdot 2 \cdot \delta \sqrt{\|\mathbf{w}\|_1} \\ &= 8\sqrt{2} \cdot (\delta^2 \|\mathbf{w}\|_1)^{3/2}. \end{aligned}$$

Thus we conclude that

$$\|\mathbf{w}' - \mathbf{w}\|_1 \leq 192\sqrt{2} \cdot \gamma \cdot (\delta^2 \|\mathbf{w}\|_1)^{3/2}.$$

□

The effect of the weight updates is to zero out the coordinates in the new residual that contribute a lot to energy. We formalize this intuition in the lemma below, whose proof appears in Appendix 9.1.5.

Lemma 3.4.18. *Suppose that we perform a perturbed residual correction step as described in Definition 3.4.16 and obtain a solution \mathbf{x}' with weights \mathbf{w}' , residual $\nabla F_\mu^{\mathbf{w}'}(\mathbf{x}') = -\mathbf{C}^\top \mathbf{g}$, and let $\Delta \mathbf{h}$ be the perturbation of the residual. Then we have that*

$$\mathcal{E}_{\mathbf{w}', s'}(\mathbf{g} + \Delta \mathbf{h}) \leq 1/4. \quad (3.55)$$

The next short lemma is a straightforward application of the vanilla correction and perfect correction routines, which shows that while only very slightly perturbing weights, we can obtain an instance where $\nabla F_\mu^{\mathbf{w}''}(\mathbf{x}'') - \mathbf{C}^\top \Delta \mathbf{h} = 0$.

Lemma 3.4.19. *Suppose that $\nabla F_\mu^{\mathbf{w}'}(\mathbf{x}') = -\mathbf{C}^\top \mathbf{g}$ for some vector \mathbf{g} , and*

$$\mathcal{E}_{\mathbf{w}', s'}(\mathbf{g} + \Delta \mathbf{h}) \leq 1/4.$$

Then using $O(\log \log m)$ iterations of a vanilla residual correction step, we can obtain a new instance with weights \mathbf{w}'' and $\mu' \leq \mu(1 + \frac{1}{2}\|\mathbf{w}\|_1^{-11})$ such that

$$\nabla F_{\mu'}^{\mathbf{w}''}(\mathbf{x}'') = \mathbf{C}^\top \Delta \mathbf{h}.$$

and $\|\mathbf{w}'' - \mathbf{w}'\|_1 \leq \|\mathbf{w}'\|_1^{-10}$.

Proof. We apply Lemma 3.3.9 for the perturbed function $\hat{F}_\mu^{\mathbf{w}'}(\mathbf{x}) = F_\mu^{\mathbf{w}'}(\mathbf{x}) - \langle \Delta \mathbf{h}, \mathbf{C} \mathbf{x} \rangle$ to obtain a new solution $\mathbf{f}'' = \mathbf{f}_0 + \mathbf{C} \mathbf{x}''$ and a new set of weights $\mathbf{w}'' \geq \mathbf{w}'$ such that $\mathbf{w}'' \leq \mathbf{w}'(1 - m^{-10})$ and $\nabla F_{\mu'}^{\mathbf{w}''}(\mathbf{x}'') - \mathbf{C}^\top \Delta \mathbf{h} = \nabla \hat{F}_{\mu'}^{\mathbf{w}''}(\mathbf{x}'') = \mathbf{0}$. □

Lemma 3.4.20. *Suppose we have an instance where $\nabla F_\mu^{\mathbf{w}}(\mathbf{x}) = \mathbf{C}^\top \Delta \mathbf{h}$. Then there exists a set of weights $\mathbf{w}' \geq \mathbf{w}$ such that $\nabla F_\mu^{\mathbf{w}'}(\mathbf{x}) = \mathbf{0}$ and $\|\mathbf{w}' - \mathbf{w}\|_1 \leq \|\Delta \mathbf{h}\|_1$.*

Proof. By definition we have

$$\frac{\mathbf{C}^\top \mathbf{c}}{\mu} + \mathbf{C}^\top \left(\frac{\mathbf{w}^+}{\mathbf{s}^+} - \frac{\mathbf{w}^-}{\mathbf{s}^-} \right) = \mathbf{C}^\top \Delta \mathbf{h}. \quad (3.56)$$

We can restore exact centrality by perturbing the weights via the simple update:

$$(\mathbf{w}^+)' = \mathbf{w}^+ - \mathbf{s}^+ \cdot (\Delta \mathbf{h})_{\leq 0}, \quad (3.57)$$

$$(\mathbf{w}^-)' = \mathbf{w}^- + \mathbf{s}^- \cdot (\Delta \mathbf{h})_{\geq 0}. \quad (3.58)$$

Plugging into the above identity immediately yields the desired condition:

$$\begin{aligned} \nabla F_\mu^{\mathbf{w}'}(\mathbf{x}) &= \frac{\mathbf{C}^\top \mathbf{c}}{\mu} + \mathbf{C}^\top \left(\frac{(\mathbf{w}^+)'}{\mathbf{s}^+} - \frac{(\mathbf{w}^-)'}{\mathbf{s}^-} \right) \\ &= \frac{\mathbf{C}^\top \mathbf{c}}{\mu} + \mathbf{C}^\top \left(\frac{\mathbf{w}^+}{\mathbf{s}^+} - \frac{\mathbf{w}^-}{\mathbf{s}^-} \right) + \mathbf{C}^\top \left(\frac{(\mathbf{w}^+)'}{\mathbf{s}^+} - \frac{\mathbf{w}^+}{\mathbf{s}^+} - \frac{(\mathbf{w}^-)'}{\mathbf{s}^-} + \frac{\mathbf{w}^-}{\mathbf{s}^-} \right) \\ &= \mathbf{C}^\top \Delta \mathbf{h} - \mathbf{C}^\top \Delta \mathbf{h} \\ &= \mathbf{0}. \end{aligned}$$

Finally, we bound

$$\|\mathbf{w}' - \mathbf{w}\|_1 = -\langle \mathbf{s}^+, (\Delta \mathbf{h})_{\leq 0} \rangle + \langle \mathbf{s}^-, (\Delta \mathbf{h})_{\geq 0} \rangle \leq \|\mathbf{s}\|_\infty \cdot \|\Delta \mathbf{h}\|_1 \leq \|\Delta \mathbf{h}\|_1.$$

For the final inequality we crucially used the fact that all slacks are at most 1. \square

Combining Lemmas 3.4.18 and 3.4.20 together with the vanilla correction step (Corollary 3.3.7 and Lemma 3.3.8) we can derive an improved correction step based on the solution to the regularized objective. First we show that indeed this is possible, i.e. for a particular choice of regularization parameters we obtain a step with $\|\boldsymbol{\rho}\|_\infty \leq 1/2$.

Lemma 3.4.21 (Feasibility lemma). *Suppose we have an instance with weights \mathbf{w} and slacks \mathbf{s} , residual $-\mathbf{C}^\top \mathbf{h}$ where $\mathbf{h} = \delta \left(\frac{\mathbf{w}^+}{\mathbf{s}^+} - \frac{\mathbf{w}^-}{\mathbf{s}^-} \right)$, and $\|\mathbf{w}\|_1^{-1/2} \leq \delta \leq \|\mathbf{w}\|_1^{-(1/4+o(1))}$. Then by solving the regularized objective we obtain a flow $\tilde{\mathbf{f}}$ satisfying*

$$\begin{aligned} \boldsymbol{\rho}^+ &= \frac{\tilde{\mathbf{f}}}{\mathbf{s}^+}, \\ \boldsymbol{\rho}^- &= \frac{-\tilde{\mathbf{f}}}{\mathbf{s}^-}, \\ \mathbf{C}^\top \left(\frac{\mathbf{w}^+ \boldsymbol{\rho}^+}{\mathbf{s}^+} - \frac{\mathbf{w}^- \boldsymbol{\rho}^-}{\mathbf{s}^-} \right) &= \mathbf{C}^\top (\mathbf{h} + \Delta \mathbf{h}). \end{aligned}$$

such that:

1. the congestion satisfies

$$\|\boldsymbol{\rho}\|_\infty \leq 1/2,$$

2. the perturbation $\Delta \mathbf{h}$ is bounded

$$\|\Delta \mathbf{h}\|_1 \leq p \cdot \|\mathbf{w}\|_1^{1/p} \cdot (10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \log \|\mathbf{w}\|_1)^2,$$

3. the flow $\tilde{\mathbf{f}}$ routes a demand $\tilde{\mathbf{d}}$ such that

$$\|\tilde{\mathbf{d}}\|_1 \leq 1.$$

Proof.

1. Let us first verify the crucial feature of $\tilde{\mathbf{f}}$, namely that it ensures a low congestion $\|\boldsymbol{\rho}\|_\infty \leq 1/2$. For our specific choice of \mathbf{h} we can upper bound, by applying Lemma 3.4.5:

$$\mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}) \leq \frac{1}{2} \delta^2 \|\mathbf{w}\|_1.$$

First we note that for this choice of parameters we can upper bound, using Corollary 3.4.10:

$$\|\tilde{\mathbf{f}}\|_\infty \leq \|\tilde{\mathbf{f}}\|_p \leq \frac{1}{10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \cdot \log \|\mathbf{w}\|_1}. \quad (3.59)$$

By Lemma 3.4.15, for each edge $e \in E$ we either have

$$\max\{|\rho_e^+|, |\rho_e^-|\} < C_\infty = \frac{1}{2\delta\sqrt{2}\|\mathbf{w}\|_1} \leq 1/2$$

by our assumption on δ , in which case we are done, or the low stretch property

$$\left| \frac{w_e^+ \rho_e^+}{s_e^+} \right| + \left| \frac{w_e^- \rho_e^-}{s_e^-} \right| \leq 6\gamma$$

holds, where by Corollary 3.4.11

$$\gamma = \delta^2 \|\mathbf{w}\|_1 \cdot 32\sqrt{6} \cdot \log \|\mathbf{w}\|_1.$$

This implies that

$$|\rho_e^+| = \left| \rho_e^+ s_e^+ \cdot \frac{\rho_e^+}{s_e^+} \right|^{1/2} = \left(|\tilde{f}_e| \cdot \left| \frac{\rho_e^+}{s_e^+} \right| \right)^{1/2} \leq \left(\frac{1}{10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \cdot \log \|\mathbf{w}\|_1} \cdot 6\gamma \right)^{1/2} < \frac{1}{2}.$$

The argument for ρ_e^- is symmetric. Therefore we have that $\|\boldsymbol{\rho}\|_\infty \leq 1/2$.

2. Now let us upper bound the ℓ_1 norm of the perturbation $\Delta\mathbf{h}$. Per Lemma 3.4.6, our step produced by solving the regularized objective yields

$$\Delta\mathbf{h} = -R_p(\tilde{\mathbf{f}})^{p-1}.$$

Therefore we can control

$$\|\Delta\mathbf{h}\|_1 = R_p \|\tilde{\mathbf{f}}\|_{p-1}^{p-1} \leq R_p \left(\|\mathbf{w}\|_1^{\frac{1}{p-1} - \frac{1}{p}} \cdot \|\tilde{\mathbf{f}}\|_p \right)^{p-1} = R_p \|\mathbf{w}\|_1^{1/p} \cdot \|\tilde{\mathbf{f}}\|_p^{p-1}.$$

Using (3.59) we further upper bound this by

$$\begin{aligned} \|\Delta \mathbf{h}\|_1 &\leq p \cdot (10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \cdot \log \|\mathbf{w}\|_1)^{p+1} \cdot \|\mathbf{w}\|_1^{1/p} \cdot \left(\frac{1}{10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \cdot \log \|\mathbf{w}\|_1} \right)^{p-1} \\ &\leq p \cdot \|\mathbf{w}\|_1^{1/p} \cdot (10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \cdot \log \|\mathbf{w}\|_1)^2. \end{aligned}$$

3. Finally, let us upper bound the demand routed by $\tilde{\mathbf{f}}$. For the demand perturbation, by Lemma 3.4.6 one has that after optimizing the regularized objective the resulting flow $\tilde{\mathbf{f}}$ routes a demand $\tilde{\mathbf{d}}$ such that

$$\|\tilde{\mathbf{d}}\|_1 \leq \left(\frac{6 \|\mathbf{w}\|_1 \cdot \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s})}{R_\star} \right)^{1/2} \leq \left(\frac{3\delta^2 \|\mathbf{w}\|_1^2}{R_\star} \right)^{1/2} = 1,$$

which gives us what we needed. □

Definition 3.4.22 (Weight balancing procedure). *Given a flow \mathbf{f} that is μ -central with respect to weights \mathbf{w} and with slacks \mathbf{s} , let $S \subseteq E$ be the set of edges that are not balanced (Invariant 3.4.14). The weight balancing procedure consists of computing new weights \mathbf{w}' such that*

- For each $e \in S$: If $w_e^+ \leq w_e^-$ then $w_e'^+ = 96 \cdot \delta^4 \|\mathbf{w}\|_1^2$, $w_e'^- = w_e^-$, while if $w_e^+ > w_e^-$ then $w_e'^+ = w_e^+$, $w_e'^- = 96 \cdot \delta^4 \|\mathbf{w}\|_1^2$.
- For each $e \notin S$ we set $w_e'^+ = w_e^+$, $w_e'^- = w_e^-$.

Additionally, we compute a flow \mathbf{f}' with slacks $\mathbf{s}' > \mathbf{0}$ such that

$$\frac{w_e'^+}{s_e'^+} - \frac{w_e'^-}{s_e'^-} = \frac{w_e^+}{s_e^+} - \frac{w_e^-}{s_e^-}.$$

Lemma 3.4.23 (Weight balancing lemma). *Given \mathbf{f} , \mathbf{w} , \mathbf{s} , S and μ as in Definition 3.4.22 after applying the weight balancing procedure we get a μ -central flow \mathbf{f}' with respect to weights $\mathbf{w}' \geq \mathbf{w}$ and with slacks \mathbf{s}' such that all edges satisfy Invariant 3.4.14 with respect to \mathbf{w}' and \mathbf{s}' . Additionally*

$$\|\mathbf{w}' - \mathbf{w}\|_1 \leq 96 \cdot |S| \cdot \delta^4 \|\mathbf{w}\|_1^2,$$

and

$$\|\mathbf{d}' - \mathbf{d}\|_1 \leq |S|,$$

and $\|\mathbf{d}' - \mathbf{d}\|_1 \leq |S|$ where \mathbf{d} is the demand routed by \mathbf{f} and \mathbf{d}' the demand routed by \mathbf{f}' .

Proof. Let $e \in S$ and without loss of generality $w_e^+ < 96 \cdot \delta^4 \|\mathbf{w}\|_1^2$ and $w_e^- > \delta \|\mathbf{w}\|_1$. First of all, we have $w_e'^+ = 96 \cdot \delta^4 \|\mathbf{w}\|_1^2 > w_e^+$ and the weight increase $w_e'^+ - w_e^+$ is at most $96 \cdot \delta^4 \|\mathbf{w}\|_1^2$. Furthermore, let us look at the function $g(t) = \frac{w_e'^+}{1-t} - \frac{w_e'^-}{t}$ for $t \in (0, 1)$. This is a continuous increasing function with $\lim_{t \rightarrow 0} g(t) = -\infty$ and $\lim_{t \rightarrow 1} g(t) = \infty$. Therefore there exists a unique $f_e' \in (0, 1)$ such that

$$\frac{w_e'^+}{1-f_e'} - \frac{w_e'^-}{f_e'} = \frac{w_e^+}{1-f_e} - \frac{w_e^-}{f_e}.$$

As $f_e, f'_e \in (0, 1)$, the demand perturbation $|f'_e - f_e|$ on this edge is at most 1. Putting everything together, we get that $\|\mathbf{w}' - \mathbf{w}\|_1 \leq 96 \cdot |S| \cdot \delta^4 \|\mathbf{w}\|_1^2$ and $\|\mathbf{d}' - \mathbf{d}\|_1 \leq |S|$. \square

Lemma 3.4.24 (Progress lemma). *Given a central instance with parameter μ and weights \mathbf{w} , we can obtain a new central instance with parameter $\mu/(1 + \delta)$ and weights $\mathbf{w}''' + \Delta\mathbf{w} \geq \mathbf{w}$ with*

$$\delta \geq m^{-1/4}/10,$$

such that

$$\begin{aligned} \|\mathbf{w}''' - \mathbf{w}\|_1 &\leq \|\Delta\mathbf{w}\|_1 \\ &+ (\delta^2 \|\mathbf{w} + \Delta\mathbf{w}\|_1)^{5/2} \cdot 6 \cdot 10^4 \cdot \log \|\mathbf{w} + \Delta\mathbf{w}\|_1 \\ &+ m^{-10} \\ &+ (10^6 \cdot \delta^2 \|\mathbf{w} + \Delta\mathbf{w}\|_1 \cdot \log \|\mathbf{w} + \Delta\mathbf{w}\|_1)^2 \cdot p \cdot \|\mathbf{w} + \Delta\mathbf{w}\|_1^{1/p}. \end{aligned}$$

where $\Delta\mathbf{w} \geq \mathbf{0}$ is the weight increase caused by applying the procedure described in Definition 3.4.22 on weights \mathbf{w} . Furthermore, the demand perturbation is $\tilde{\mathbf{d}} + \Delta\tilde{\mathbf{d}}$, where

$$\|\tilde{\mathbf{d}}\|_1 \leq 1.$$

and $\Delta\tilde{\mathbf{d}}$ is the demand perturbation caused by applying the procedure described in Definition 3.4.22 on weights \mathbf{w} .

Proof. We first apply the weight balancing procedure as described in Definition 3.4.22 to get new weights $\mathbf{w} + \Delta\mathbf{w} \geq \mathbf{w}$.

We will apply the residual correction step with the flow $\tilde{\mathbf{f}}$ guaranteed by Lemma 3.4.21, which guarantees that the corresponding flow yields a congestion $\|\boldsymbol{\rho}\|_\infty \leq 1/2$. As this step is obtained for a perturbed residual $\nabla F_\mu^{\mathbf{w} + \Delta\mathbf{w}}(\mathbf{x}) - \mathbf{C}^\top \mathbf{h} = -\mathbf{C}^\top (\mathbf{h} + \Delta\mathbf{h})$, it means that after executing it we obtain an instance with weights \mathbf{w}' and slacks \mathbf{s}' such that

$$\mathcal{E}_{\mathbf{w}', \mathbf{s}'}(\mathbf{h}' + \Delta\mathbf{h}) \leq 1/4,$$

where $\nabla F_\mu^{\mathbf{w}'}(\mathbf{x}) = -\mathbf{C}^\top \mathbf{h}'$. This follows from applying Lemma 3.4.18. Therefore, using Lemma 3.4.19 we can obtain a new point where

$$\nabla F_\mu^{\mathbf{w}''}(\mathbf{x}'') = \mathbf{C}^\top \Delta\mathbf{h},$$

for a slight change in weights from \mathbf{w}' to \mathbf{w}'' . Finally, applying Lemma 3.4.20 we can establish exact centrality

$$\nabla F_\mu^{\mathbf{w}'''}(\mathbf{x}'') = \mathbf{0},$$

while slightly increasing weights from \mathbf{w}'' to \mathbf{w}''' .

Having described the method, let us first bound the weight change it causes. By Lemma 3.4.17 we have that the weight increase $\|\mathbf{w}' - (\mathbf{w} + \Delta\mathbf{w})\|_1$ caused by performing the perturbations described in Definition 3.4.16 is upper bounded by

$$192\sqrt{2} \cdot \gamma \cdot (\delta^2 \|\mathbf{w} + \Delta\mathbf{w}\|_1)^{3/2},$$

where

$$\gamma = \delta^2 \|\mathbf{w} + \Delta\mathbf{w}\|_1 \cdot 32\sqrt{6} \cdot \log \|\mathbf{w} + \Delta\mathbf{w}\|_1.$$

Furthermore, restoring the condition that $\nabla F_\mu^{\mathbf{w}'''}(\mathbf{x}'') = \mathbf{C}^\top \Delta\mathbf{h}$ is done while increasing the ℓ_1 norm

of the weight vector by $\|\mathbf{w}'' - \mathbf{w}'\|_1 \leq m^{-10}$ and restoring exact centrality via Lemma 3.4.6 costs us a further weight increase $\|\mathbf{w}''' - \mathbf{w}''\|_1$ that is upper bounded by

$$\|\Delta \mathbf{h}\|_1 \leq p \cdot \|\mathbf{w} + \Delta \mathbf{w}\|_1^{1/p} \cdot (10^6 \cdot \delta^2 \|\mathbf{w} + \Delta \mathbf{w}\|_1 \log \|\mathbf{w} + \Delta \mathbf{w}\|_1)^2,$$

where we used the guarantee from Lemma 3.4.21. Therefore we conclude that

$$\begin{aligned} \|\mathbf{w}''' - \mathbf{w}\|_1 &= \|\Delta \mathbf{w}\|_1 + \|\mathbf{w}' - (\mathbf{w} + \Delta \mathbf{w})\|_1 + \|\mathbf{w}'' - \mathbf{w}'\|_1 + \|\mathbf{w}''' - \mathbf{w}''\|_1 \\ &\leq \|\Delta \mathbf{w}\|_1 \\ &\quad + (\delta^2 \|\mathbf{w} + \Delta \mathbf{w}\|_1)^{5/2} \cdot 6 \cdot 10^4 \cdot \log \|\mathbf{w} + \Delta \mathbf{w}\|_1 \\ &\quad + m^{-10} \\ &\quad + (10^6 \cdot \delta^2 \|\mathbf{w} + \Delta \mathbf{w}\|_1 \cdot \log \|\mathbf{w}\|_1)^2 \cdot p \cdot \|\mathbf{w} + \Delta \mathbf{w}\|_1^{1/p}. \end{aligned}$$

Finally, the two steps in which the flow demand changes is the weight balancing procedure, in which the change in demand is $\Delta \tilde{\mathbf{d}}$, and the step when we restrict the regularized problem solution to the non-augmented graph. In the latter case, the change in demand $\tilde{\mathbf{d}}$ can be upper bounded by Lemma 3.4.21 by at most 1 in ℓ_1 norm. Thus the demand will be perturbed by $\|\mathbf{d}' - \mathbf{d}\|_1 = \|\tilde{\mathbf{d}} + \Delta \tilde{\mathbf{d}}\|_1 \leq 1 + \|\Delta \tilde{\mathbf{d}}\|_1$. This concludes the proof. \square

Lemma 3.4.24 is the main workhorse of the improved algorithm. It shows that we can make large progress within the interior point method, while paying for some demand perturbation and for some slight increase in $\|\mathbf{w}\|_1$. In order to guarantee sufficient progress, all we are left to do is to ensure that we can set an appropriate δ such that the sum of weights never increases beyond $O(m)$. This is a mere consequence of the result given above.

Lemma 3.4.25. *Suppose we have a μ -central instance with weights $\mathbf{w} \geq \mathbf{1}$, where $\|\mathbf{w}\|_1 \leq 2m + 1$ and $\mu = m^{O(1)}$. Let $\varepsilon = m^{-O(1)}$, and let $\delta = m^{-(3/8+o(1))}$. In $\tilde{O}(\delta^{-1})$ iterations of the procedure described in Lemma 3.4.24 we obtain an instance with duality gap at most ε with a total demand perturbation of $\tilde{O}(\delta^{-1})$.*

Proof. We perform a sequence of iterations as described in Lemma 3.4.24.

We will argue that within $T = \tilde{O}(\delta^{-1} \log(m\mu\varepsilon^{-1}))$ iterations of the procedure described in Lemma 3.4.24 the barrier weights will always satisfy $\|\mathbf{w}\|_1 \leq 3m$. We need to take into account the total increase guaranteed by Lemma 3.4.24, together with the possible weight increases caused by the weight balancing procedure.

First let us bound the total number of balancing operations, since each of these can increase a single weight by a significant amount. First, from Definition 3.4.22 we see that once an edge gets balanced it will never become unbalanced again, as weights are monotonic.

Furthermore, we see that such an operation can only occur when the largest of the two paired weights is at least $\delta \|\mathbf{w}\|_1 \geq 2m\delta$, since we maintain $\mathbf{w} \geq \mathbf{1}$. Under the invariant that throughout the entire algorithm $\|\mathbf{w}\|_1 \leq 3m$, we therefore see that this can only happen at most $(3m)/(2m\delta) = 3/(2\delta)$ times. Invoking Lemma 3.4.23 we therefore get that the total weight change caused by these operations is at most

$$\frac{3}{2\delta} \cdot 96 \cdot \delta^4 \cdot (3m)^2 = 1296 \cdot \delta^3 m^2.$$

In addition, we incur weight increases due to the progress steps; per Lemma 3.4.24, within each of the T iterations, $\|\mathbf{w}\|_1$ increases by at most

$$(\delta^2 (3m))^{5/2} \cdot 6 \cdot 10^4 \cdot \log 3m + m^{-10} + 10^{12} \cdot p \delta^4 (3m)^{2+1/p} \cdot \log^2(3m).$$

Therefore the total weight increase over T iterations is at most

$$\tilde{O}\left(\left(\delta^4 m^{5/2} + p\delta^3 m^{2+1/p}\right) \log(m\mu\varepsilon^{-1})\right)$$

which is $o(m)$ as long as

$$\delta \leq 1/\max\left(m^{3/8+o(1)}, p^{1/3}m^{(1+1/p)/3+o(1)}\right) = 1/m^{3/8+o(1)}.$$

Thus this specific choice of δ insures that the invariant $\|\mathbf{w}\|_1 \leq 3m$ is satisfied throughout the entire run of the algorithm.

Finally, we bound the total perturbation in demand suffered by the flow we maintain. From Lemma 3.4.18 each progress step perturbs the demand by at most 1 in ℓ_1 norm. Furthermore, the weight balancing operations may perturb it by an additional $\frac{3}{2\delta}$ overall. Summing up we obtain the desired claim. \square

Combining with the repairing procedure elaborated in [42], which we show how to adapt to our present setting in Section 3.5, we obtain the main theorem.

Theorem 3.4.26. *Given a directed graph $G(V, E, \mathbf{c})$ with m arcs and n vertices, such that $\|\mathbf{c}\|_\infty \leq W$, and a demand vector $\mathbf{d} \in \mathbb{Z}^n$, in $m^{11/8+o(1)} \log W$ time we can obtain a flow \mathbf{f} which routes \mathbf{d} in G while satisfying the capacity constraints $\mathbf{0} \leq \mathbf{f} \leq \mathbf{1}$ and minimizing the cost $\sum_{e \in E} c_e f_e$, or certifies that no such flow exists.*

Proof. Before proceeding, we note that Lemma 3.4.25 assumes that the initial centrality parameter μ is bounded by a polynomial in m . This may not be exactly true as the initial centering (Lemma 9.1.2) is done while setting a parameter μ which is upper bounded by $m^{O(1)}W$. As Lemma 3.4.25 which bounds the number of iterations of our interior point method assumes that initially $\mu = m^{O(1)}$, we need to enforce the property that the method in the lemma is only called on instances where $W = m^{O(1)}$.

Similarly to [42], we enforce this property by employing the scaling technique of [71]. Thus, we reduce the problem to solving $O(\log W)$ instances of the problem, where costs are polynomially bounded. Without doing so, the running time of our algorithm would have depended super-linearly in $\log W$.

For each of the $O(\log W)$ instances, we proceed as follows. First, we apply Lemma 9.1.2 to obtain a μ -central solution for $\mu = m^{O(1)}$. Then, we apply Lemma 3.4.25 with $\varepsilon = m^{-3}$ to obtain a flow \mathbf{f} which routes a perturbed demand \mathbf{d}' where $\|\mathbf{d}' - \mathbf{d}\|_1 \leq m^{3/8+o(1)}$, and is close to optimal, in the sense that it is accompanied by dual variables which certify a duality gap of $O(m^{-3})$. Note that each of the $m^{3/8+o(1)}$ iterations of the algorithm from Lemma 3.4.25 requires solving the regularized objective (3.18) to high precision, which due to our choice of parameter p (Corollary 3.4.9) can be done in $m^{1+o(1)}$ time. Finally, we apply Lemma 3.5.4 to repair this flow in $m^{11/8+o(1)}$ time. Hence the total running time is $m^{11/8+o(1)}$.

We can certify infeasibility (the case where the demand can not be routed while satisfying the capacity constraints), by looking at the arcs used by the returned optimal solution. The initialization procedure assumes that a demand satisfying flow can be routed in the graph. This is done by adding to the graph $O(m)$ arcs with high costs (see Lemma 9.1.1). Using Lemma 9.1.1 we see that simply inspecting the optimal solution returned for this graph we can decide if the routing is infeasible in G . \square

BALANCEWEIGHTS($G; \mathbf{w}, \mathbf{f}$)

1. For each $e \in E$:
2. If $w_e^+ < 96 \cdot \delta^4 \|\mathbf{w}\|_1^2$ and $w_e^- > \delta \|\mathbf{w}\|_1$ then $w_e'^+ \leftarrow 96 \cdot \delta^4 \|\mathbf{w}\|_1^2$, $w_e'^- \leftarrow w_e^-$.
3. If $w_e^+ > \delta \|\mathbf{w}\|_1$ and $w_e^- < 96 \cdot \delta^4 \|\mathbf{w}\|_1^2$ then $w_e'^+ \leftarrow w_e^+$, $w_e'^- \leftarrow 96 \cdot \delta^4 \|\mathbf{w}\|_1^2$.
4. Set $\mathbf{0} < \mathbf{f}' < \mathbf{1}$ such that $\frac{w_e'^+}{1-f'} - \frac{w_e'^-}{f'} = \frac{w_e^+}{1-f} - \frac{w_e^-}{f}$.
5. Return \mathbf{w}', \mathbf{f}' .

Figure 3-4: Weight balancing procedure

3.5 Repairing the Flow

In this section we show that the flow produced in Section 3.4 can be repaired in such a way that we obtain an optimal flow for the original problem. This can be done by applying the combinatorial fixing techniques used in [42]. Let us first recall a few useful definitions concerning the bipartite perfect \mathbf{b} -matching problem.

Bipartite Perfect \mathbf{b} -matching. For a given weighted bipartite graph $G = (V, E)$ with $V = V_1 \cup V_2$ where V_1 and V_2 are the two sets of bipartition and a *demand vector* $\mathbf{b} \in \mathbb{R}_+^V$, a *perfect \mathbf{b} -matching* is a vector $\mathbf{x} \in \mathbb{R}_+^E$ such that $\sum_{e \in E(v)} x_e = b_v$ for all $v \in V$. A perfect \mathbf{b} -matching is a generalization of perfect matching; in the particular case where all \mathbf{b} 's equal 1, integer \mathbf{b} -matchings (or $\mathbf{1}$ -matchings) are exactly perfect matchings. We require that $\mathbf{b}(V_1) = \mathbf{b}(V_2)$ as otherwise they trivially do not exist.

The *weighted bipartite perfect \mathbf{b} -matching problem* is defined as follows: given a weighted bipartite graph $G = (V, E, \mathbf{c})$, return a perfect \mathbf{b} -matching in G that has minimum weight, or conclude that there is no perfect \mathbf{b} -matching in G . The dual problem to the weighted perfect bipartite \mathbf{b} -matching is a *\mathbf{b} -vertex packing problem* where we want to find a vector $\mathbf{y} \in \mathbb{R}^V$ satisfying the following LP

$$\begin{aligned} \max \quad & \sum_{v \in V} y_v b_v, \\ & y_u + y_v \leq w_{uv} \quad \forall u, v \in E. \end{aligned}$$

We employ the following result which follows from a direct application of Lemmas 35, 37, and Theorem 40 in [42].

Theorem 3.5.1. *Consider an instance $G = (V, E, \mathbf{c})$ of the weighted perfect bipartite \mathbf{b} -matching problem, where $\|\mathbf{b} - \hat{\mathbf{b}}\|_1 \leq P$, and $\|\mathbf{b}\|_1 = O(m)$. Given a feasible primal-dual pair of variables (\mathbf{x}, \mathbf{y}) with duality gap at most m^{-2} , in $O(Pm + m \log n)$ time we can compute an optimal primal-dual pair $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ to the perfect $\hat{\mathbf{b}}$ -matching problem, or conclude that no such matching exists.*

In order to apply this theorem we need to convert our instance into a \mathbf{b} -matching instance with small duality gap. The output we receive from Lemma 3.4.25 is a flow \mathbf{f} which routes a slightly different demand \mathbf{d} from the one we originally intended. Furthermore \mathbf{f} satisfies the centrality condition $\mathbf{C}^\top \left(\frac{\mathbf{w}^-}{\mathbf{f}} - \frac{\mathbf{w}^+}{1-\mathbf{f}} \right) = \mathbf{C}^\top \frac{\mathbf{c}}{\mu}$ for a small value of μ .

From Flow to Bipartite b -Matching. We convert this solution into a corresponding bipartite b -matching problem in a new bipartite graph $G'(V_1 \cup V_2, E', \mathbf{c}')$, defined as follows:

1. for each vertex $v \in V$, create the corresponding vertex in V_1 ;
2. for each arc $e \in E$, create a vertex $v_e \in V_2$;
3. for each arc $e = (u, v) \in E$ create edges (u, v_e) with cost c_e and (v, v_e) with cost 0.

Let \mathbf{d} be the demand routed by \mathbf{f} . For the bipartite graph G' , we define the \mathbf{b} vector as follows:

1. for each $v \in V$ set for the corresponding vertex v in V_1 : $b_v = |E^+(v)| + d_v$.
2. for each $v_2 \in V_2$ set $b_{v_2} = 1$.

We can easily verify that an optimal primal-dual solution to the \mathbf{b} -matching problem in G' maps to an optimal primal-dual solution to the minimum cost flow problem in G .

Lemma 3.5.2. *Let (\mathbf{x}, \mathbf{y}) be a pair of feasible primal-dual variables for the \mathbf{b} -matching problem in the graph G' , constructed according to the rules defined above. These can be mapped to a pair (\mathbf{f}, \mathbf{z}) of optimal feasible primal-dual variables for the minimum cost flow problem in G . Furthermore, this mapping can be constructed in linear time.*

Proof. Let \mathbf{x} be the optimal \mathbf{b} -matching, and let \mathbf{y} be its dual vector certifying optimality.

We construct \mathbf{f} as follows: for each arc $e = (u, v) \in G$ we look at the corresponding gadget in G' consisting of vertices $u, v \in V_1$ and $v_e \in V_2$ and set $f_e = x_{u, v_e}$.

First we verify that \mathbf{f} is feasible. For each arc $e = (u, v) \in G$, the corresponding vertex $v_e \in V_2$ has $b_{v_e} = 1$, and therefore $x_{u, v_e} + x_{v, v_e} = 1$. As both are non-negative it means that in G , $0 \leq f_e \leq 1$.

Next we verify that \mathbf{f} indeed routes the demand it is supposed to route, i.e. \mathbf{d} . By definition we have that for each $v \in V_1$,

$$\sum_{(v, v_e) \in E'} x_{v, v_e} = b_v = |E^+(v)| + d_v.$$

Therefore the demand at vertex v can be written as

$$\begin{aligned} \sum_{e=(u,v) \in E} f_e - \sum_{e'=(v,u) \in E} f_{e'} &= \sum_{e=(u,v) \in E} (1 - f_e) - |E^+(v)| + \sum_{e'=(v,u) \in E} f_{e'} \\ &= \sum_{(v, v_e) \in E'} x_{v, v_e} - |E^+(v)| \\ &= d_v. \end{aligned}$$

Finally we certify optimality for \mathbf{f} by exhibiting feasible dual variables which satisfy complementary slackness. For the LP formulation defined in Section 3.3.1 we require for each arc $e \in E$ two non-negative dual variables z_e^- and z_e^+ such that summing up along each cycle

$$z_e^+ - z_e^- + c_e$$

with the appropriate sign depending on the orientation of the encountered arcs, we obtain exactly 0. For each $e = (u, v) \in E$, we define them as

$$\begin{aligned} z_e^- &= c_e - y_u - y_{v_e}, \\ z_e^+ &= -y_v - y_{v_e}. \end{aligned}$$

Feasibility is obvious since $z_e^+, z_e^- \geq 0$ because \mathbf{y} is a feasible vector and hence for each edge the sum of the y 's of its vertices is upper bounded by the cost, and summing up $z_e^- - z_e^+$ along any cycle we obtain exactly the sum of the costs c_e along that cycle, summed up with the appropriate sign.

We finally certify optimality for \mathbf{f} by verifying that $(\mathbf{f}, (\mathbf{z}^-; \mathbf{z}^+))$ satisfy complementary slackness. We have that for each arc $(u, v) \in E$,

$$\begin{aligned} f_e z_e^- &= x_{uv_e}(c_e - y_u - y_{v_e}) = 0 \\ (1 - f_e) z_e^+ &= x_{vv_e}(-y_v - y_{v_e}) = 0, \end{aligned}$$

where we used the fact that (\mathbf{x}, \mathbf{y}) satisfy complementary slackness. Hence \mathbf{f} is optimal in G . \square

Our goal will be to obtain an optimal primal-dual solution for a $\widehat{\mathbf{b}}$ -matching problem which carries over to an optimal solution to the minimum cost flow problem in G with the original demand $\widehat{\mathbf{d}}$.

In order to do so, we first convert our flow instance with small duality gap to a \mathbf{b} -matching instance with the same duality gap. Afterwards, we fix the resulting \mathbf{b} -matching to optimality and we convert it to an optimal $\widehat{\mathbf{b}}$ -matching which maps to a flow $\widehat{\mathbf{f}}$ routing the original demand $\widehat{\mathbf{d}}$ in G , via Theorem 3.5.1.

In order to do so, we use the μ -central solution produced by the interior point method in Section 3.4, with $\mu = O(m^{-10})$. Recalling that our current solution satisfies

$$\mathbf{C}^\top \left(\frac{\mu \mathbf{w}^-}{\mathbf{f}} - \frac{\mu \mathbf{w}^+}{\mathbf{1} - \mathbf{f}} \right) = \mathbf{C}^\top \mathbf{c},$$

we set primal-dual variables (\mathbf{x}, \mathbf{y}) for G' as follows:

1. for each arc $e = (u, v) \in E$, set $x_{u, v_e} = f_e$ and $x_{v, v_e} = 1 - f_e$.
2. pick an arbitrary vertex $v_0 \in V$ and set its corresponding $y_{v_0} = 0$; then perform a breadth-first search in G (ignoring orientations) starting from v , and for each vertex visited for the first time when traversing an arc $e = (u, v)$ (in any direction) such that $y_u - y_v = \frac{\mu w_e^+}{1 - f_e} - \frac{\mu w_e^-}{f_e} - c_e$;
3. for all $v_e \in V_2$ where $e = (u, v) \in E$, set $y_{v_e} = -y_u - \frac{\mu w_e^-}{f_e} + c_e = -y_v - \frac{\mu w_e^-}{1 - f_e}$.

Having defined G' together with its demand vector \mathbf{b} and corresponding primal-dual solution (\mathbf{x}, \mathbf{y}) , let us show that indeed this is a feasible solution and it has low duality gap.

Lemma 3.5.3. *The primal-dual solution (\mathbf{x}, \mathbf{y}) constructed above is feasible and has duality gap $\mu \|\mathbf{w}\|_1$.*

Proof. First we check feasibility. Primal feasibility holds trivially by construction. For dual feasibility, consider for each arc $e = (u, v) \in E$ the two edges appearing in its corresponding gadget in G' , (u, v_e) and (v, v_e) . For the former we have

$$y_u + y_{v_e} = -\frac{\mu w_e^-}{f_e} + c_e \leq c_{u, v_e}$$

and for the latter

$$y_v + y_{v_e} = -\frac{\mu w_e^+}{1 - f_e} \leq c_{v, v_e} = 0.$$

Therefore the dual constraints of the \mathbf{b} -matching problem are also satisfied. Finally we write the duality gap via

$$\begin{aligned}
\sum_{e=(u,v) \in E'} x_e(c_e - y_u - y_v) &= \sum_{e=(u,v) \in E} (f_e(c_e - y_u - y_{v_e}) + (1 - f_e)(0 - y_v - y_{v_e})) \\
&= \sum_{e=(u,v) \in E} \left(f_e \cdot \frac{\mu w_e^-}{f_e} + (1 - f_e) \cdot \frac{\mu w_e^+}{1 - f_e} \right) \\
&= \sum_{e=(u,v) \in E} \mu(w_e^- + w_e^+) \\
&= \mu \|\mathbf{w}\|_1.
\end{aligned}$$

□

Now we are ready to state the main result of this section.

Lemma 3.5.4. *Given a target integral demand $\widehat{\mathbf{d}}$, a flow \mathbf{f} satisfying μ -centrality with weights \mathbf{w} such that $\|\mathbf{w}\|_1 = O(m)$, $\mu = O(m^{-10})$, and \mathbf{f} routes a demand \mathbf{d} we can obtain in time $\tilde{O}(\|\mathbf{d} - \widehat{\mathbf{d}}\|_1 \cdot m)$ a flow $\widehat{\mathbf{f}}$ which is optimal among all flows routing $\widehat{\mathbf{d}}$.*

Proof. Using Lemma 3.5.3 we can convert the flow instance into a \mathbf{b} -matching instance with duality gap $O(m^{-2})$. Now we can invoke Theorem 3.5.1 to convert the matching into an optimal $\widehat{\mathbf{b}}$ -matching where $\widehat{\mathbf{b}}$ corresponds to the demand $\widehat{\mathbf{d}}$ that the flow is supposed to route in G . We see that by construction $\|\widehat{\mathbf{b}} - \widehat{\mathbf{d}}\|_1 = \|\mathbf{d} - \widehat{\mathbf{d}}\|_1$. Therefore obtaining the corresponding optimal $\widehat{\mathbf{b}}$ -matching is done in time $O((\|\mathbf{d} - \widehat{\mathbf{d}}\|_1 + \log n)m)$. Finally applying Lemma 3.5.2 we obtain the optimal flow $\widehat{\mathbf{f}}$ in G . □

3.6 Improving the Running Time to $m^{4/3+o(1)} \log W$

The running time of the algorithm we presented above hits a barrier of $m^{11/8+o(1)} \log W$. The key bottleneck there is the post-processing we do on the residual error of the (intermediate) solutions we obtain after performing each progress step. Indeed, while the length of our steps is dictated by the ℓ_∞ -norm of our step size $\|\boldsymbol{\rho}\|_\infty$ (which is, in a sense optimal), it is unclear how to ensure that the energy required to route a flow that fixes the corresponding residual error is sufficiently small, without overly increasing the weights of the constraint barriers. In fact, the extent of weight perturbations necessary for this post-processing step are exactly what determines the $m^{11/8+o(1)} \log W$ running time. All the other weight perturbations, which are caused by the regularization terms, are much milder and would lead to the desired $m^{4/3+o(1)} \log W$ running time.

After the first version of [10] was posted, Liu and Sidford [114] published a preprint that obtains an improved running time for the unit-capacity maximum flow. The main technique introduced in that paper boils down exactly to avoiding the aforementioned bottleneck. Roughly, instead of advancing from a central point to the next one via a progress step followed by a sequence of residual correction steps, they instead directly solve the optimization problem which lands them at the next central point.

To do so, one must guarantee that this optimization problem is well-conditioned at all times, in the sense that the objective has a Hessian which always stays within a constant factor from the one at the origin. While such a Hessian stability condition is not true in general, Liu and Sidford [114] modify the logarithmic barriers they use by extending them with quadratics outside the region

where they would be naturally well-behaved. The resulting new objective function can be efficiently optimized by slightly extending the mixed ℓ_2 - ℓ_p solver of Kyng et al [103] (see Appendix 9.1.4).

Incorporating this idea in our framework yields a the desired running time improvement of our unit-capacity minimum cost flow algorithm as well. Most of the details, particularly those involving preconditioning and weight perturbations carry over from the previous sections. In fact, the only change to the algorithm needed is to replace the regularized Newton step (cf. line 3 in Figure 3-3) with solving a regularized problem which directly involves the logarithmic barrier. (Furthermore, lines 7 and 8 are no longer required.)

Method Overview. Let us specify the ideal optimization problem which we would solve in order to advance along the central path. Suppose we have a μ -central instance i.e. we have a flow $\mathbf{f} = \mathbf{f}_0 + \mathbf{C}\mathbf{x}$ which satisfies

$$\mathbf{C}^\top \left(\frac{\mathbf{w}^+}{\mathbf{1} - \mathbf{f}_0 - \mathbf{C}\mathbf{x}} - \frac{\mathbf{w}^-}{\mathbf{f}_0 + \mathbf{C}\mathbf{x}} \right) = -\frac{\mathbf{C}^\top \mathbf{c}}{\mu}. \quad (3.60)$$

as it optimizes the convex objective

$$\min_{\mathbf{x}} F_\mu^w(\mathbf{x}) = \min_{\mathbf{x}} \frac{1}{\mu} \langle \mathbf{c}, \mathbf{C}\mathbf{x} \rangle - \sum_{e \in E} (w_e^+ \log(\mathbf{1} - \mathbf{f}_0 - \mathbf{C}\mathbf{x})_e + w_e^- \log(\mathbf{f}_0 + \mathbf{C}\mathbf{x})_e).$$

Our goal is to design an optimization procedure which enables us to augment \mathbf{f} with a circulation $\mathbf{C}\tilde{\mathbf{x}}$ in order to obtain an optimizer for

$$\begin{aligned} \min_{\mathbf{x}'} F_{\mu/(1+\delta)}^w(\mathbf{x}') &= \min_{\mathbf{x}'} \frac{1+\delta}{\mu} \langle \mathbf{c}, \mathbf{C}\mathbf{x}' \rangle - \sum_{e \in E} (w_e^+ \log(\mathbf{1} - \mathbf{f}_0 - \mathbf{C}\mathbf{x}')_e + w_e^- \log(\mathbf{f}_0 + \mathbf{C}\mathbf{x}')_e) \\ &= \min_{\tilde{\mathbf{x}}} \frac{1+\delta}{\mu} \langle \mathbf{c}, \mathbf{f} + \mathbf{C}\tilde{\mathbf{x}} \rangle - \sum_{e \in E} (w_e^+ \log(\mathbf{1} - \mathbf{f} - \mathbf{C}\tilde{\mathbf{x}})_e + w_e^- \log(\mathbf{f} + \mathbf{C}\tilde{\mathbf{x}})_e). \end{aligned}$$

We can equivalently rewrite this optimization procedure, after adding a constant term to it, as

$$\min_{\tilde{\mathbf{x}}} \Psi_{\mu/(1+\delta)}^{\mathbf{w}, \mathbf{f}}(\tilde{\mathbf{x}}),$$

where

$$\begin{aligned} \Psi_{\mu/(1+\delta)}^{\mathbf{w}, \mathbf{f}}(\tilde{\mathbf{x}}) &:= F_{\mu/(1+\delta)}^w(\mathbf{x} + \tilde{\mathbf{x}}) - \frac{1+\delta}{\mu} \langle \mathbf{c}, \mathbf{f} \rangle + \sum_{e \in E} (w_e^+ \log(1 - f_e) + w_e^- \log f_e) \\ &= \frac{1+\delta}{\mu} \langle \mathbf{c}, \mathbf{C}\tilde{\mathbf{x}} \rangle - \sum_{e \in E} \left(w_e^+ \log \left(\mathbf{1} - \frac{\mathbf{C}\tilde{\mathbf{x}}}{\mathbf{1} - \mathbf{f}} \right)_e + w_e^- \log \left(\mathbf{1} + \frac{\mathbf{C}\tilde{\mathbf{x}}}{\mathbf{f}} \right)_e \right). \end{aligned}$$

Optimizing this function is hard to do in general. However, assuming its optimal augmenting circulation $\mathbf{C}\tilde{\mathbf{x}}$ satisfies

$$\left\| \frac{\mathbf{C}\tilde{\mathbf{x}}}{\min\{\mathbf{f}, \mathbf{1} - \mathbf{f}\}} \right\|_\infty \leq \frac{1}{10}, \quad (3.61)$$

which in other words, says that augmenting \mathbf{f} with $\mathbf{C}\tilde{\mathbf{x}}$ will not come close to breaking the feasibility constraints, we can instead solve a well-conditioned objective obtained by replacing the log's with a better behaved function $\tilde{\log}$ satisfying $\log(1+t) = \tilde{\log}(1+t)$ whenever $|t| \leq 1/10$.

More precisely, we use the definition from [114] which we reproduce below:

$$\widetilde{\log}(1+t) = \begin{cases} \log(1+t), & \text{if } t \in [-\theta, \theta], \\ \log(1+\theta) + (t-\theta) \cdot \log'(1+\theta) + (t-\theta)^2 \cdot \frac{1}{2} \log''(1+\theta), & \text{if } t > \theta, \\ \log(1-\theta) + (t+\theta) \cdot \log'(1-\theta) + (t+\theta)^2 \cdot \frac{1}{2} \log''(1-\theta), & \text{if } t < -\theta, \end{cases}$$

where we set $\theta = 1/10$. We can also easily verify that on the boundary of the interval $[-\theta, \theta]$ the first and second order derivatives exactly match those of $\log(1+t)$. The essential feature is that outside this range, the second derivative stays constant, whereas in the case of $\log(1+t)$ it changes very fast.

Using this, we can define the minimization problem

$$\min_{\tilde{\mathbf{x}}} \widetilde{\Psi}_{\mu/(1+\delta)}^{w,\mathbf{f}}(\tilde{\mathbf{x}}),$$

where

$$\widetilde{\Psi}_{\mu/(1+\delta)}^{w,\mathbf{f}}(\tilde{\mathbf{x}}) := \frac{1+\delta}{\mu} \langle \mathbf{c}, \mathbf{C}\tilde{\mathbf{x}} \rangle - \sum_{e \in E} \left(w_e^+ \widetilde{\log} \left(\mathbf{1} - \frac{\mathbf{C}\tilde{\mathbf{x}}}{\mathbf{1} - \mathbf{f}} \right)_e + w_e^- \widetilde{\log} \left(\mathbf{1} + \frac{\mathbf{C}\tilde{\mathbf{x}}}{\mathbf{f}} \right)_e \right). \quad (3.62)$$

Observation 3.6.1. *If the minimizer $\tilde{\mathbf{x}}$ of $\widetilde{\Psi}_{\mu/(1+\delta)}^{w,\mathbf{f}}(\tilde{\mathbf{x}})$ satisfies the low-congestion condition from (3.61), then it also minimizes $\Psi_{\mu/(1+\delta)}^{w,\mathbf{f}}(\tilde{\mathbf{x}})$.*

Due to the fact that the second order derivatives of $\widetilde{\log}$ are bounded, $\widetilde{\Psi}_{\mu/(1+\delta)}^{w,\mathbf{f}}(\tilde{\mathbf{x}})$ is well-conditioned and, as a matter of fact can easily be minimized by using a small number of calls to a routine which minimizes quadratics over the set of circulations. As specified in [114] this can be easily done by using fast Laplacian system solvers. However, the main difficulty that arises is ensuring that the minimizer $\tilde{\mathbf{x}}$ satisfies the condition from (3.61).

Enforcing this property is non-trivial, and requires regularizing the function $\widetilde{\Psi}_{\mu/(1+\delta)}^{w,\mathbf{f}}$ in an identical manner to the way we did it in the analysis from Section 3.4. Most of the results we used there carry over, after performing some minor modifications in the analysis.

Regularizing the Objective. In order to enforce the required property, we add two regularization terms to our optimization step. Just like in Section 3.4.1 we augment the graph G to G_\star , which has a cycle basis \mathbf{C}_\star , and in this graph we write down the equivalent concave maximization problem for (3.62), which we regularize with two extra terms. We slightly abuse notation by making $\widetilde{\Psi}_{\mu/(1+\delta)}^{w,\mathbf{f}}$ act on an element in the circulation space of G_\star , with the meaning that the linear and logarithmic terms only act on edges in E :

$$\max_{\tilde{\mathbf{f}}_\star \in \mathbf{C}_\star \tilde{\mathbf{x}}} -\widetilde{\Psi}_{\mu/(1+\delta)}^{w,\mathbf{f}}(\tilde{\mathbf{x}}) - \frac{R_\star}{2} \sum_{e \in E'} (\tilde{f}_\star)_e^2 - \frac{R_p}{p} \sum_{e \in E \cup E'} (\tilde{f}_\star)_e^p. \quad (3.63)$$

Writing $\tilde{\mathbf{f}}_\star = \tilde{\mathbf{f}} + \tilde{\mathbf{f}}'$ where $\tilde{\mathbf{f}}$ is the restriction of $\tilde{\mathbf{f}}_\star$ to the edges E of G , and $\tilde{\mathbf{f}}'$ is the restriction

to the augmenting edges E' , and using the centrality condition (3.60), we can further write this as

$$\begin{aligned} \max_{\tilde{\mathbf{f}}_\star = \mathbf{C}_\star \tilde{\mathbf{x}}} (1 + \delta) \left\langle \frac{\mathbf{w}^+}{\mathbf{1} - \mathbf{f}} - \frac{\mathbf{w}^-}{\mathbf{f}}, \tilde{\mathbf{f}} \right\rangle + \sum_{e \in E} \left(w_e^+ \widetilde{\log} \left(1 - \frac{\tilde{f}_e}{1 - f_e} \right) + w_e^- \widetilde{\log} \left(1 + \frac{\tilde{f}_e}{f_e} \right) \right) \\ - \frac{R_\star}{2} \sum_{e \in E'} (\tilde{f}'_e)^2 - \frac{R_p}{p} \sum_{e \in E \cup E'} (\tilde{f}_\star)_e^p. \end{aligned} \quad (3.64)$$

Note that this optimization problem can be solved efficiently using Lemma 9.1.7. The key reason is that all the terms except for the one involving p powers of the flow f have a second order derivative which is either constant, or which stays bounded within a small multiplicative factor from the one at 0. Similarly to before, the solver from Lemma 9.1.7 yields a high-accuracy, yet inexact solution. We can, however, assume we obtain an exact solution by performing minor perturbations to our problem, in a manner similar to the one discussed in Section 9.1.3.

We can now write the first order optimality condition for the objective in (3.63), thus providing an analogue of Lemma 3.4.6. Before doing so, we give the following helper lemma, which will be the main driver of the results in this section. It intuitively states that the optimal solution to (3.64) can be thought of as the solution to a regularized Newton step as the one in (3.18), but where the coefficients of the quadratic $\frac{1}{2} \sum_{e \in E} \tilde{f}'_e \cdot \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right)$ have been slightly perturbed by a small multiplicative constant.

Lemma 3.6.2 (Optimality with average Hessian). *Let $\tilde{\mathbf{f}}_\star = \mathbf{C}_\star \tilde{\mathbf{x}}_\star$ be the optimizer of (3.64), and let $\tilde{\mathbf{f}}_\star = \tilde{\mathbf{f}} + \tilde{\mathbf{f}}'$, where the two components are supported on E and E' , respectively. Then there exists a vector $\boldsymbol{\alpha} = (\boldsymbol{\alpha}^+; \boldsymbol{\alpha}^-)$, $(1 + \theta)^{-2} \cdot \mathbf{1} \leq \boldsymbol{\alpha} \leq (1 - \theta)^{-2} \cdot \mathbf{1}$, which can be explicitly computed, such that for any circulation \mathbf{g} in \mathbf{C}_\star ,*

$$\left\langle \mathbf{g}, \begin{bmatrix} \delta \left(\frac{\mathbf{w}^+}{\mathbf{1} - \mathbf{f}} - \frac{\mathbf{w}^-}{\mathbf{f}} \right) - \tilde{\mathbf{f}} \left(\frac{\boldsymbol{\alpha}^+ \mathbf{w}^+}{(\mathbf{1} - \mathbf{f})^2} + \frac{\boldsymbol{\alpha}^- \mathbf{w}^-}{\mathbf{f}^2} \right) - R_p \cdot (\tilde{\mathbf{f}})^{p-1} \\ -R_\star \cdot \tilde{\mathbf{f}}' - R_p \cdot (\tilde{\mathbf{f}}')^{p-1} \end{bmatrix} \right\rangle = 0.$$

Proof. We write the first order optimality condition for (3.64). We have that for any circulation \mathbf{g} in \mathbf{C}_\star :

$$\left\langle \mathbf{g}, \begin{bmatrix} (1 + \delta) \left(\frac{\mathbf{w}^+}{\mathbf{1} - \mathbf{f}} - \frac{\mathbf{w}^-}{\mathbf{f}} \right) + \mathbf{w}^+ \cdot \nabla_{\tilde{\mathbf{f}}} \widetilde{\log} \left(1 - \frac{\tilde{\mathbf{f}}}{\mathbf{1} - \mathbf{f}} \right) + \mathbf{w}^- \cdot \nabla_{\tilde{\mathbf{f}}} \widetilde{\log} \left(1 + \frac{\tilde{\mathbf{f}}}{\mathbf{f}} \right) - R_p \cdot (\tilde{\mathbf{f}})^{p-1} \\ -R_\star \cdot \tilde{\mathbf{f}}' - R_p \cdot (\tilde{\mathbf{f}}')^{p-1} \end{bmatrix} \right\rangle = 0.$$

Next we use the expansion

$$\frac{d}{dx} \widetilde{\log} \left(1 + \frac{x}{a} \right) = \frac{1}{a} \cdot \widetilde{\log}' \left(1 + \frac{x}{a} \right) = \frac{1}{a} \cdot \left(1 + \frac{x}{a} \int_0^1 \widetilde{\log}'' \left(1 + t \cdot \frac{x}{a} \right) dt \right).$$

From the definition of $\widetilde{\log}$ we know that its second order derivative is always stable which enables us to bound

$$-\frac{1}{(1 - \theta)^2} \leq \int_0^1 \widetilde{\log}'' \left(1 + t \cdot \frac{x}{a} \right) dt \leq -\frac{1}{(1 + \theta)^2},$$

hence for some $\alpha \in [(1 + \theta)^{-2}, (1 - \theta)^{-2}]$, one has

$$\frac{d}{dx} \widetilde{\log} \left(1 + \frac{x}{a} \right) = \frac{1}{a} \cdot \widetilde{\log}' \left(1 + \frac{x}{a} \right) = \frac{1}{a} - \alpha \cdot \frac{x}{a^2}.$$

Therefore there exists a vector $\boldsymbol{\alpha} = (\boldsymbol{\alpha}^+; \boldsymbol{\alpha}^-)$ satisfying $(1 + \theta)^{-2} \cdot \mathbf{1} \leq \boldsymbol{\alpha} \leq (1 - \theta)^{-2} \cdot \mathbf{1}$ such that

$$\begin{aligned} & \left\langle \mathbf{g}, \left[\begin{array}{c} (1 + \delta) \left(\frac{\mathbf{w}^+}{\mathbf{1} - \mathbf{f}} - \frac{\mathbf{w}^-}{\mathbf{f}} \right) + \mathbf{w}^+ \cdot \nabla_{\tilde{\mathbf{f}}} \widetilde{\log} \left(\mathbf{1} - \frac{\tilde{\mathbf{f}}}{\mathbf{1} - \mathbf{f}} \right) + \mathbf{w}^- \cdot \nabla_{\tilde{\mathbf{f}}} \widetilde{\log} \left(\mathbf{1} + \frac{\tilde{\mathbf{f}}}{\mathbf{f}} \right) - R_p \cdot (\tilde{\mathbf{f}})^{p-1} \\ -R_{\star} \cdot \tilde{\mathbf{f}}' - R_p \cdot (\tilde{\mathbf{f}}')^{p-1} \end{array} \right] \right\rangle \\ &= \left\langle \mathbf{g}, \left[\begin{array}{c} \delta \left(\frac{\mathbf{w}^+}{\mathbf{1} - \mathbf{f}} - \frac{\mathbf{w}^-}{\mathbf{f}} \right) - \tilde{\mathbf{f}} \left(\frac{\boldsymbol{\alpha}^+ \mathbf{w}^+}{(\mathbf{1} - \mathbf{f})^2} + \frac{\boldsymbol{\alpha}^- \mathbf{w}^-}{\mathbf{f}^2} \right) - R_p \cdot (\tilde{\mathbf{f}})^{p-1} \\ -R_{\star} \cdot \tilde{\mathbf{f}}' - R_p \cdot (\tilde{\mathbf{f}}')^{p-1} \end{array} \right] \right\rangle = 0, \end{aligned}$$

which is what we needed. □

Lemma 3.6.2 enables us to use an optimality condition very similar to the one we had before in Section 3.4. As a matter of fact, all the remaining statements are nothing but "robust" versions of those we previously used. Essential here are new versions of Lemma 3.4.6 and Lemma 3.4.7 which accommodate the extra multiplicative factors on resistances. Roughly, our goal is to provide upper bounds on $\|\tilde{\mathbf{f}}\|_{\infty}$ and $\|\tilde{\mathbf{f}}/\min\{\mathbf{1} - \mathbf{f}, \mathbf{f}\}\|_{\infty}$, which together will imply that the condition from (3.61) is satisfied.

After proving that this is the case, we will show how to advance to the next point on the central path – the regularization terms on (3.64) will require us to increase the weights \mathbf{w} in order to obtain optimality for the non-regularized objective i.e. $\nabla_{\tilde{\Psi}_{\mu/(1+\delta)}^{\mathbf{w}, \mathbf{f}}}(\tilde{\mathbf{x}}) = 0$. Nevertheless, this procedure is essentially identical to the one we previously used in Lemma 3.4.20.

3.6.1 Bounding Congestion

Here we show that the congestion condition from (3.61) is satisfied, and hence the minimizer of (3.64) also minimizes the expression after replacing $\widetilde{\log}$ with \log . The proofs are almost identical to those from Section 3.4. For consistency we will use the slack notation

$$\mathbf{s}^- = \mathbf{f}, \quad \mathbf{s}^+ = \mathbf{1} - \mathbf{f},$$

and use the shorthand notation for the residual

$$\mathbf{h} = \delta \left(\frac{\mathbf{w}^+}{\mathbf{s}^+} - \frac{\mathbf{w}^-}{\mathbf{s}^-} \right).$$

We first give a short lemma providing an optimality condition for the restriction of the flow $\tilde{\mathbf{f}}_{\star}$ computed by (3.64) to the edges E of the original graph G .

Lemma 3.6.3 (Optimality in the non-augmented graph). *Let $\tilde{\mathbf{f}}_{\star} = \mathbf{C}_{\star} \tilde{\mathbf{x}}_{\star}$ be the optimizer of (3.63) and let $\tilde{\mathbf{f}}$ be its restriction to the edges of G . Let $\tilde{\mathbf{d}}$ be the demand routed by $\tilde{\mathbf{f}}$ in G . Then there exists a vector $\boldsymbol{\alpha} = (\boldsymbol{\alpha}^+; \boldsymbol{\alpha}^-) \in \mathbb{R}^{2m}$, $\frac{1}{(1+\theta)^2} \cdot \mathbf{1} \leq \boldsymbol{\alpha} \leq \frac{1}{(1-\theta)^2} \cdot \mathbf{1}$, which can be explicitly computed, such that*

$$\mathbf{C}^{\top} \cdot \left(\frac{\boldsymbol{\alpha}^+ \mathbf{w}^+}{(\mathbf{s}^+)^2} + \frac{\boldsymbol{\alpha}^- \mathbf{w}^-}{(\mathbf{s}^-)^2} \right) \cdot \tilde{\mathbf{f}} = \mathbf{C}^{\top} (\mathbf{h} + \Delta \mathbf{h}) \quad (3.65)$$

where $\Delta \mathbf{h} = -R_p(\tilde{\mathbf{f}})^{p-1}$, and

$$\begin{aligned}\|\tilde{\mathbf{d}}\|_1 &\leq 3 \left(\frac{\|\mathbf{w}\|_1 \cdot \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s})}{R_\star} \right)^{1/2}, \\ \|\tilde{\mathbf{f}}_\star\|_p &\leq \left(\frac{p \cdot \frac{3}{2} \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s})}{R_p} \right)^{1/p}.\end{aligned}$$

The proof can be found in Appendix 9.1.5.

Next we provide a guarantee enforced by the component of the regularizer involving augmenting edges.

Lemma 3.6.4. *Let $\tilde{\mathbf{f}}_\star$ be the solution of the regularized objective and $\tilde{\mathbf{f}}$ its restriction on G , and suppose that $\|\mathbf{w}\|_1 \geq 3$. Then there exists a vector $\boldsymbol{\alpha} = (\boldsymbol{\alpha}^+; \boldsymbol{\alpha}^-) \in \mathbb{R}^{2m}$ such that for all edges $e \in E$:*

$$\left| \left(\frac{\alpha_e^+ w_e^+}{(s_e^+)^2} + \frac{\alpha_e^- w_e^-}{(s_e^-)^2} + R_p \cdot \tilde{f}_e^{p-2} \right) \tilde{f}_e - h_e \right| \leq \hat{\gamma}, \quad (3.66)$$

where

$$\hat{\gamma} = \left(R_\star + R_p \cdot \|\tilde{\mathbf{f}}_\star\|_\infty^{p-2} \right)^{1/2} \cdot \left\| \frac{\mathbf{h}}{\sqrt{(\mathbf{w}^+ + \mathbf{w}^-) \left(\frac{\alpha^+ \mathbf{w}^+}{(s^+)^2} + \frac{\alpha^- \mathbf{w}^-}{(s^-)^2} \right)}} \right\|_\infty \cdot 32 \log \|\mathbf{w}\|_1.$$

Furthermore, this implies that

$$\left(\frac{\alpha_e^+ w_e^+}{(s_e^+)^2} + \frac{\alpha_e^- w_e^-}{(s_e^-)^2} \right) \cdot |\tilde{f}_e| \leq |h_e| + \hat{\gamma}. \quad (3.67)$$

The proof is essentially identical to that of Lemma 3.4.7. We discuss it at the end of Appendix 9.1.2.

Combining Lemmas 3.6.3 and 3.6.4 we can finally prove the main statement of this section. This will be true for as long as the condition we specified in Definition 3.4.13 holds. This condition will be later enforced by performing the balancing operation we also specified in Definition 3.4.22.

In order to obtain the desired congestion bound, we set our regularization parameters identically to Definition 3.4.8:

$$\begin{aligned}p &= \min \left\{ k \in 2\mathbb{Z} : k \geq (\log m)^{1/3} \right\}, \\ R_p &= p \cdot (10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \cdot \log \|\mathbf{w}\|_1)^{p+1}, \\ R_\star &= 3 \cdot \delta^2 \|\mathbf{w}\|_1^2.\end{aligned}$$

Lemma 3.6.5 (Congestion bound). *Suppose that all edges $e \in E$ are balanced, per Definition 3.4.13, and $\delta > 10 \cdot \|\mathbf{w}\|_1^{-1/2}$. Then the restriction $\tilde{\mathbf{f}}$ of the flow $\tilde{\mathbf{f}}_\star$ computed via (3.64) satisfies the low-congestion condition (3.61).*

Proof. Lemma 3.4.15 proves that if all edges are balanced, for any flow $\hat{\mathbf{f}}$ in G such that for all $e \in E$

$$\left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) \cdot |\hat{f}_e| \leq |h_e| + \hat{\gamma},$$

where $\mathbf{h} = \delta \left(\frac{\mathbf{w}^+}{s^+} - \frac{\mathbf{w}^-}{s^-} \right)$, then for each $e \in E$ it is true that $\max\{|\widehat{\rho}_e^+|, |\widehat{\rho}_e^-|\} < C_\infty$ or $\left| \frac{w_e^+ \widehat{\rho}_e^+}{s_e^+} \right| + \left| \frac{w_e^- \widehat{\rho}_e^-}{s_e^-} \right| \leq 6\gamma$, where $\widehat{\rho}_e^+ = \frac{\widehat{f}_e}{s_e^+}$, $\widehat{\rho}_e^- = \frac{-\widehat{f}_e}{s_e^-}$, and $C_\infty = \frac{1}{2\delta\sqrt{2}\|\mathbf{w}\|_1}$. We apply this statement for

$$\left| \widehat{f}_e \right| := \frac{\frac{\alpha_e^+ w_e^+}{(s_e^+)^2} + \frac{\alpha_e^- w_e^-}{(s_e^-)^2}}{\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2}} \left| \widetilde{f}_e \right| \geq \frac{1}{(1+\theta)^2} \left| \widetilde{f}_e \right|,$$

which also implies that $|\rho_e^+| = \left| \frac{\widetilde{f}_e}{s_e^+} \right| \leq (1+\theta)^2 |\widehat{\rho}_e^+|$, $|\rho_e^-| = \left| \frac{\widetilde{f}_e}{s_e^-} \right| \leq (1+\theta)^2 |\widehat{\rho}_e^-|$. Note that we now have

$$\left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) \left| \widehat{f}_e \right| = \left(\frac{\alpha_e^+ w_e^+}{(s_e^+)^2} + \frac{\alpha_e^- w_e^-}{(s_e^-)^2} \right) \cdot \left| \widetilde{f}_e \right| \leq |h_e| + \widehat{\gamma},$$

and thus we get that for all $e \in E$ at least one of

$$\max\{|\rho_e^+|, |\rho_e^-|\} \leq (1+\theta)^2 \cdot \max\{|\widehat{\rho}_e^+|, |\widehat{\rho}_e^-|\} < (1+\theta)^2 \cdot C_\infty$$

and

$$\left| \frac{w_e^+ \rho_e^+}{s_e^+} \right| + \left| \frac{w_e^- \rho_e^-}{s_e^-} \right| \leq (1+\theta)^2 \cdot \left(\left| \frac{w_e^+ \widehat{\rho}_e^+}{s_e^+} \right| + \left| \frac{w_e^- \widehat{\rho}_e^-}{s_e^-} \right| \right) \leq 6 \cdot (1+\theta)^2 \cdot \gamma$$

is true. Using the fact that $\delta > 10 \cdot \|\mathbf{w}\|_1^{-1/2}$, the former becomes $\max\{|\rho_e^+|, |\rho_e^-|\} = \frac{(1+\theta)^2}{2\delta\sqrt{2}\|\mathbf{w}\|_1} < \frac{1}{20}$. For all the remaining edges, we have

$$\left| \frac{w_e^+ \rho_e^+}{s_e^+} \right| + \left| \frac{w_e^- \rho_e^-}{s_e^-} \right| \leq 6 \cdot (1+\theta)^2 \cdot \gamma = 6 \cdot (1+\theta)^2 \cdot \delta^2 \|\mathbf{w}\|_1 \cdot 32\sqrt{6} \cdot \log \|\mathbf{w}\|_1,$$

which, combined with

$$\begin{aligned} \left| \widetilde{f}_e \right| &\leq \|\widetilde{\mathbf{f}}\|_p \leq \left(\frac{p \cdot \frac{3}{2} \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s})}{R_p} \right)^{1/p} \\ &\leq \left(\frac{p \cdot \frac{3}{4} \delta^2 \|\mathbf{w}\|_1}{p (10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \log \|\mathbf{w}\|_1)^{p+1}} \right)^{1/p} \\ &\leq \frac{1}{10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \log \|\mathbf{w}\|_1} \end{aligned}$$

gives

$$\begin{aligned} |\rho_e^+| &= \left(\left| \widetilde{f}_e \right| \cdot \left| \frac{\rho_e^+}{s_e^+} \right| \right)^{1/2} \\ &\leq \left(\frac{1}{10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \log \|\mathbf{w}\|_1} \cdot 6 \cdot (1+\theta)^2 \cdot \delta^2 \|\mathbf{w}\|_1 \cdot 32\sqrt{6} \log \|\mathbf{w}\|_1 \right)^{1/2} \\ &< \frac{1}{20}. \end{aligned}$$

Symmetrically, we get that $|\rho_e^-| < \frac{1}{20}$ and so we are done. \square

3.6.2 Making Progress

Here we show how to use the flow obtained from optimizing (3.64) in order to achieve centrality for a new parameter $\mu/(1 + \delta)$, at the expense of slightly increasing weights from \mathbf{w} to some \mathbf{w}' .

Lemma 3.6.6 (Almost-centrality after executing step). *Let $\tilde{\mathbf{f}}$ be the optimizer of (3.64) and let $\tilde{\mathbf{f}}$ be its restriction to the edges of G . Then*

$$\mathbf{C}^\top \left((1 + \delta) \left(\frac{\mathbf{w}^+}{\mathbf{1} - \mathbf{f}} - \frac{\mathbf{w}^-}{\mathbf{f}} \right) - \left(\frac{\mathbf{w}^+}{\mathbf{1} - \mathbf{f} - \tilde{\mathbf{f}}} - \frac{\mathbf{w}^-}{\mathbf{f} + \tilde{\mathbf{f}}} \right) \right) = \mathbf{C}^\top \cdot R_p \cdot (\tilde{\mathbf{f}})^{p-1}.$$

Proof. Writing the optimality condition for (3.64), and using Lemma 3.6.5, we obtain

$$\mathbf{C}^\top \left((1 + \delta) \left(\frac{\mathbf{w}^+}{\mathbf{1} - \mathbf{f}} - \frac{\mathbf{w}^-}{\mathbf{f}} \right) + \mathbf{w}^+ \cdot \nabla_{\tilde{\mathbf{f}}} \log \left(\mathbf{1} - \frac{\tilde{\mathbf{f}}}{\mathbf{1} - \mathbf{f}} \right) + \mathbf{w}^- \cdot \nabla_{\tilde{\mathbf{f}}} \log \left(\mathbf{1} + \frac{\tilde{\mathbf{f}}}{\mathbf{f}} \right) - R_p \cdot (\tilde{\mathbf{f}})^{p-1} \right) = 0,$$

which we can rewrite equivalently as

$$\mathbf{C}^\top \left((1 + \delta) \left(\frac{\mathbf{w}^+}{\mathbf{1} - \mathbf{f}} - \frac{\mathbf{w}^-}{\mathbf{f}} \right) - \left(\frac{\mathbf{w}^+}{\mathbf{1} - \mathbf{f} - \tilde{\mathbf{f}}} - \frac{\mathbf{w}^-}{\mathbf{f} + \tilde{\mathbf{f}}} \right) \right) = \mathbf{C}^\top \cdot R_p \cdot (\tilde{\mathbf{f}})^{p-1},$$

which is what we needed. □

As we can see, the regularization terms have two effects. One is that the update $\tilde{\mathbf{f}}$ is not exactly a circulation, so this will account for some change in the routed demand. The other effect is that after augmenting the current flow \mathbf{f} with $\tilde{\mathbf{f}}$ we do not obtain a central solution, as shown in Lemma 3.6.6. We proceed to fix this manually by slightly increasing the weights \mathbf{w} . This follows from applying Lemma 3.4.20 to the residual $\Delta \mathbf{h} = -R_p(\tilde{\mathbf{f}})^{p-1}$, which produces a central solution with a new set of weights $\mathbf{w}' \geq \mathbf{w}$ such that $\|\mathbf{w}' - \mathbf{w}\|_1 \leq R_p \cdot \sum_{e \in E} |\tilde{f}_e|^{p-1}$.

We are now ready to characterize the amount of progress we make in a single iteration of the method previously described.

Lemma 3.6.7 (Progress lemma). *Given a μ -central instance, i.e. a flow \mathbf{f} and balanced weights \mathbf{w} such that*

$$\mathbf{C}^\top \left(\frac{\mathbf{w}^+}{\mathbf{1} - \mathbf{f}} - \frac{\mathbf{w}^-}{\mathbf{f}} \right) = -\frac{\mathbf{C}^\top \mathbf{c}}{\mu},$$

in the time require to solve (3.64) we can obtain a $\mu/(1 + \delta)$ -central instance, i.e. a flow $\mathbf{f} + \tilde{\mathbf{f}}$ and weights $\mathbf{w}' \geq \mathbf{w}$, such that

$$\mathbf{C}^\top \left(\frac{\mathbf{w}^+}{\mathbf{1} - \mathbf{f} - \tilde{\mathbf{f}}} - \frac{\mathbf{w}^-}{\mathbf{f} + \tilde{\mathbf{f}}} \right) = -(1 + \delta) \frac{\mathbf{C}^\top \mathbf{c}}{\mu},$$

where

$$\|\mathbf{w}' - \mathbf{w}\|_1 \leq p \cdot 10^{12} \cdot \delta^4 \|\mathbf{w}\|_1^{2+1/p} \cdot \log^2 \|\mathbf{w}\|_1,$$

and $\tilde{\mathbf{f}}$ routes a demand $\tilde{\mathbf{d}}$ such that

$$\|\tilde{\mathbf{d}}\|_1 \leq 3/2.$$

Proof. Using Lemma 3.6.3 and Lemma 3.6.6, after solving (3.64) we obtain an augmenting flow $\tilde{\mathbf{f}}$ routing a demand $\tilde{\mathbf{d}}$ with

$$\|\tilde{\mathbf{d}}\|_1 \leq 3 \left(\frac{\|\mathbf{w}\|_1 \cdot \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s})}{R_\star} \right)^{1/2} \leq 3 \left(\frac{\|\mathbf{w}\|_1 \cdot \frac{1}{2} \delta^2 \|\mathbf{w}\|_1}{3\delta^2 \|\mathbf{w}\|_1^2} \right)^{1/2} < 3/2,$$

such that

$$\mathbf{C}^\top \left((1 + \delta) \left(\frac{\mathbf{w}^+}{\mathbf{1} - \mathbf{f}} - \frac{\mathbf{w}^-}{\mathbf{f}} \right) - \left(\frac{\mathbf{w}^+}{\mathbf{1} - \mathbf{f} - \tilde{\mathbf{f}}} - \frac{\mathbf{w}^-}{\mathbf{f} + \tilde{\mathbf{f}}} \right) \right) = \mathbf{C}^\top \cdot R_p \cdot (\tilde{\mathbf{f}})^{p-1}.$$

We then increase the weights \mathbf{w} to \mathbf{w}' to make the right-hand side of this identity equal to $\mathbf{0}$. Per Lemma 3.4.20, this increases the weights to \mathbf{w}' such that

$$\|\mathbf{w}' - \mathbf{w}\|_1 = R_p \cdot \sum_{e \in E} |\tilde{f}_e|^{p-1} \leq R_p \left(m^{\frac{1}{p-1} - \frac{1}{p}} \cdot \|\tilde{\mathbf{f}}\|_p \right)^{p-1} \leq R_p \cdot \|\mathbf{w}\|_1^{1/p} \cdot \|\tilde{\mathbf{f}}\|_p^{p-1},$$

where we used the fact that $\|\mathbf{w}\|_1 \geq m$. From Lemma 3.6.3, we have that for our specific choice of regularization parameters,

$$\|\tilde{\mathbf{f}}\|_p \leq \left(\frac{p \cdot \frac{3}{2} \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s})}{R_p} \right)^{1/p} \leq \frac{1}{10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \log \|\mathbf{w}\|_1}.$$

Therefore

$$\begin{aligned} \|\mathbf{w}' - \mathbf{w}\|_1 &\leq \left(p \cdot (10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \cdot \log \|\mathbf{w}\|_1)^{p+1} \right) \cdot \|\mathbf{w}\|_1^{1/p} \cdot \left(\frac{1}{10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \log \|\mathbf{w}\|_1} \right)^{p-1} \\ &= p \cdot (10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \cdot \log \|\mathbf{w}\|_1)^2 \cdot \|\mathbf{w}\|_1^{1/p} = p \cdot 10^{12} \cdot \delta^4 \|\mathbf{w}\|_1^{2+1/p} \cdot \log^2 \|\mathbf{w}\|_1. \end{aligned}$$

□

3.6.3 Wrapping Up

We can now give the main statement of this section, which follows from running the interior point method, based on the guarantee provided by Lemma 3.6.7.

Lemma 3.6.8. *Suppose we have a μ -central instance with weights $\mathbf{w} \geq \mathbf{1}$, where $\|\mathbf{w}\|_1 \leq 2m + 1$ and $\mu = m^{O(1)}$. Let $\varepsilon = m^{-O(1)}$, and let $\delta = m^{-(1/3+o(1))}$. In time dominated by $\tilde{O}(\delta^{-1})$ iterations of the procedure described in Lemma 3.6.7 we obtain an instance with duality gap at most ε with a total demand perturbation of $\tilde{O}(\delta^{-1})$.*

Proof. We perform a sequence of iterations as described in Lemma 3.6.7. These are interspersed with calls to the weight balancing procedure (Lemma 3.4.23), required in order to maintain the invariant needed by Lemma 3.6.5.

We will argue that within $T = \tilde{O}(\delta^{-1} \log(m\mu\varepsilon^{-1}))$ iterations the barrier weights will always satisfy $\|\mathbf{w}\|_1 \leq 3m$. We need to take into account the total increase guaranteed by Lemma 3.6.7 together with the possible weight increases caused by the weight balancing procedure.

First let us bound the total number of balancing operations, since each of these can increase a single weight by a significant amount. First, from Definition 3.4.22 we see that once an edge gets balanced it will never become unbalanced again, as weights are monotonic.

Furthermore, we see that such an operation can only occur when the largest of the two paired weights is at least $\delta \|\mathbf{w}\|_1 \geq 2m\delta$, since we maintain $\mathbf{w} \geq \mathbf{1}$. Under the invariant that throughout the entire algorithm $\|\mathbf{w}\|_1 \leq 3m$, we therefore see that this can only happen at most $(3m)/(2m\delta) = 3/(2\delta)$ times. Invoking Lemma 3.4.23 we therefore get that the total weight change caused by these operations is at most

$$\frac{3}{2\delta} \cdot 96 \cdot \delta^4 \cdot (3m)^2 = 1296 \cdot \delta^3 m^2.$$

In addition, we incur weight increases due to the progress steps; per Lemma 3.6.7, within each of the T iterations, $\|\mathbf{w}\|_1$ increases by at most

$$p \cdot 10^{12} \cdot \delta^4 (3m)^{2+1/p} \cdot \log^2(3m).$$

Therefore the total weight increase over T iterations is at most

$$\tilde{O} \left(\left(p \delta^3 m^{2+1/p} \right) \log(m \mu \varepsilon^{-1}) \right)$$

which is $o(m)$ as long as

$$\delta \leq 1 / \left(p^{1/3} m^{(1+1/p)/3+o(1)} \right) = 1/m^{1/3+o(1)}.$$

Thus this specific choice of δ insures that the invariant $\|\mathbf{w}\|_1 \leq 3m$ was satisfied throughout the entire run of the algorithm.

Finally, we bound the total perturbation in demand suffered by the flow we maintain. Each of the $\tilde{O}(\delta^{-1})$ iterations perturbs the demand by at most $3/2$ in ℓ_1 norm, and the weight balancing operations may perturb the demand by an additional $\frac{3}{2\delta}$ overall. Summing up we obtain the desired claim. \square

This enables us to obtain a running time of $m^{4/3+o(1)}$ for minimum cost flow in unit-capacity graphs. The proof is identical to that of Theorem 3.4.26, we use scaling to obtain a logarithmic dependence in W , and resort to the fixing procedure from Section 3.5 to repair the demand perturbation. The time required to implement each iteration of the interior point method is dominated by the time required by one call to the solver in Theorem 9.1.7, which is $m^{1+o(1)}$ by our choice of parameters.

Theorem 3.6.9. *Given a directed graph $G(V, E, \mathbf{c})$ with m arcs and n vertices, such that $\|\mathbf{c}\|_\infty \leq W$, and a demand vector $\mathbf{d} \in \mathbb{Z}^n$, in $m^{4/3+o(1)} \log W$ time we can obtain a flow \mathbf{f} which routes \mathbf{d} in G while satisfying the capacity constraints $\mathbf{0} \leq \mathbf{f} \leq \mathbf{1}$ and minimizing the cost $\sum_{e \in E} c_e f_e$, or certifies that no such flow exists.*

Chapter 4

Faster Sparse Minimum Cost Flow by Electrical Flow Localization

4.1 Introduction

In the last decade, continuous optimization has proved to be an invaluable tool for designing graph algorithms, often leading to significant improvements over the best known combinatorial algorithms. This has been particularly true in the context of flow problems—arguably, some of the most prominent graph problems [44, 39, 105, 118, 152, 97, 106, 136, 117, 42, 151, 150, 153, 115, 114, 10, 26, 25]. Indeed, these developments have brought a host of remarkable improvements in a variety of regimes, such as when seeking only approximate solutions, or when the underlying graph is dense. However, most of these improvements did not fully address the challenge of seeking *exact* solutions in *sparse* graphs. Fortunately, the improvements for that regime eventually emerged [118, 117, 42, 115, 114, 10]. They still suffered though from an important shortcoming: they all had a polynomial running time dependency in the graph’s capacities, and hence—in contrast to the classical combinatorial algorithms—they did not yield efficient algorithms in the presence of arbitrary capacities. Recently, Gao, Liu and Peng [73] have finally changed this state of affairs by providing the first *exact* maximum flow algorithm to break the $\tilde{O}\left(m^{3/2} \log^{O(1)} U\right)$ barrier for sparse graphs with *general* capacities (bounded by U). Their approach, however, crucially relies on a preconditioning technique that is specific to the maximum flow problem and, in particular, having an s - t demand—rather than a general one. As a result, the corresponding improvement held only for that particular problem.

In this chapter, we demonstrate how to circumvent these limitations and provide the first algorithm that breaks the $\tilde{O}\left(m^{3/2} \log^{O(1)}(U + W)\right)$ barrier for the *minimum cost flow* problem in *sparse* graphs with general demands, capacities (bounded by U), and costs (bounded by W). This algorithm runs in time $\tilde{O}\left(m^{3/2-1/762} \log(U + W)\right)$.

4.1.1 Previous work

In 2013, Mądry [118] presented the first running time improvement to the maximum flow problem since the $\tilde{O}(m\sqrt{n} \log U)$ algorithm of [75] in the regime of sparse graphs with small capacities. To this end, he presented an algorithm that runs in time $\tilde{O}\left(m^{10/7} \text{poly}(U)\right)$, where U is a bound on edge capacities, breaking past the $\tilde{O}\left(m^{3/2}\right)$ running time barrier that has for decades resisted improvement attempts. The main idea in that work was to use an interior point method with an improved number of iterations guarantee that was delivered via use of an adaptive re-weighting of the central path and careful perturbations of the problem instance. Building on this framework,

a series of subsequent works [117, 115, 114] has brought the runtime of sparse max flow down to $\tilde{O}(m^{4/3}\text{poly}(U))$. (With the most recent of these works crucially relying on nearly-linear time ℓ_p flows [103].) In parallel [42, 10], the running time of the more general minimum cost flow problem was reduced to $\tilde{O}(m^{4/3}\text{poly}(U)\log W)$, where W is a bound on edge costs.

However, even though these algorithms offer a significant improvement when U is relatively small, the question of whether there exists an algorithm faster than $\tilde{O}(m^{3/2}\log^{O(1)}U)$ for sparse graphs with general capacities remained open. In fact, a polynomial dependence on capacities or costs seems inherent in the central path re-weighting technique used in all the aforementioned works. Recently, [73] finally made progress on this question by developing an algorithm for the maximum flow problem that runs in time $\tilde{O}(m^{3/2-1/328}\log U)$. The source of improvement here was different from previous works, in the sense that it was not based on decreasing the number of iterations of the interior point method. Instead, it was based on devising a data structure to solve the dynamically changing Laplacian system required by the interior point method in sublinear time per iteration.

The new approach put forth by [73], despite being quite different to the prior ones, still leaned on the preconditioning approach of [117], as well as on other properties that are specific to the maximum flow problem. For this reason, this improvement did not extend to the minimum cost flow problem with general capacities, for which the fastest known runtime was still $\tilde{O}(m\log(U+W) + n^{1.5}\log^2(U+W))$ [25] and $\tilde{O}(m^{3/2}\log^{O(1)}(U+W))$ [44] in the sparse regime.

4.1.2 Our result

In this work, we give an algorithm for the minimum cost flow problem with a running time of $\tilde{O}(m^{3/2-1/762}\log(U+W))$. This is the first improvement for sparse graphs with general capacities over [44], which runs in time $\tilde{O}(m^{3/2}\log^{O(1)}(U+W))$. Specifically, we prove that:

Theorem 4.1.1. *Given a graph $G(V, E)$ with edge costs $\mathbf{c} \in \mathbb{Z}_{[-W, W]}^m$, a demand $\mathbf{d} \in \mathbb{R}^n$, and capacities $\mathbf{u} \in \mathbb{Z}_{[0, U]}^m$, there exists an algorithm that with high probability runs in time $\tilde{O}(m^{3/2-1/762}\log(U+W))$ and returns a flow $\mathbf{f} \in [0, \mathbf{u}]$ in G such that \mathbf{f} routes the demand \mathbf{d} and the cost $\langle \mathbf{c}, \mathbf{f} \rangle$ is minimized.*

4.1.3 High level overview of our approach

As we build on the approach presented in [73], we first briefly overview some of the key ideas introduced there that will also be relevant for our discussion. The maximum flow interior point method by [117] works by, repeatedly over $\tilde{O}(\sqrt{m})$ steps, taking an electrical flow step that is a multiple of

$$\tilde{\mathbf{f}} = \mathbf{R}^{-1}\mathbf{B}\mathbf{L}^+\mathbf{B}^\top\mathbf{1}_{st},$$

where $\mathbf{L} = \mathbf{B}^\top\mathbf{R}^{-1}\mathbf{B}$ is a Laplacian matrix and \mathbf{r} are resistances that change per step. However, $\tilde{\mathbf{f}}$ has m entries and takes $\tilde{O}(m)$ to compute, which gives the standard $\tilde{O}(m^{3/2})$ bound. To go beyond this, [73] show that it suffices to compute $\tilde{\mathbf{f}}$ for only a *sublinear* number of high-congestion entries of $\tilde{\mathbf{f}}$, where congestion is defined as $\rho = \sqrt{\mathbf{r}}\tilde{\mathbf{f}}$. By known linear sketching results, these edges can be detected by computing the inner product $\langle \mathbf{q}, \rho \rangle$ for a small number of randomly chosen vectors $\mathbf{q} \in \mathbb{R}^m$. Crucially, given a vertex subset $C \subseteq V$ of sublinear size that contains s and t , this inner product can be equivalently written as the following sublinear-sized inner product

$$\langle \mathbf{q}, \rho \rangle = \left\langle \pi^C \left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}} \right), \mathbf{S}C^+\mathbf{d} \right\rangle, \quad (4.1)$$

where $SC := SC(G, C)$ is the *Schur complement* of G onto C , \mathbf{d} is equal to $\mathbf{B}^\top \mathbf{1}_{st}$, and $\pi^C(\cdot)$ is a *demand projection* onto C . Therefore, the problem is reduced to maintaining two quantities: $\pi^C\left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}}\right)$ and $(SC(G, C))^+ \mathbf{d}$ in sublinear time per operation. The latter is computed by using the dynamic Schur complement data structure of [51], and the former can be maintained by a careful use of random walks.

We now describe our approach. Instead of using the interior point method formulation of [117] which only applies to the maximum flow problem, we use the one by [10] for the, more general, minimum cost flow problem.

There are now several obstacles to making this approach work by maintaining the quantity $\langle \mathbf{q}, \boldsymbol{\rho} \rangle$:

Preconditioning A significant difference between [117] and [10] is that while the former is able to guarantee that the magnitude of the electrical potentials computed in each step is inversely proportional to the duality gap, meaning that a large duality gap implies potential embeddings of low stretch, no such preconditioning method is known for minimum cost flow. In fact, [10] used demand perturbations to show that a *weaker* bound on the potentials can be achieved, which was still sufficient for their purposes. Unfortunately, this bound is not strong enough to be used in the analysis of [73].

In order to alleviate this issue, we completely remove preconditioning from the picture by only requiring a bound on the *energy* of the electrical potentials (instead of their magnitude). In particular, given an approximate demand projection $\tilde{\pi}^C\left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}}\right)$, identity (4.1) is used to detect congested edges. In [73], there is a uniform upper bound on the entries of the potential embedding $\boldsymbol{\phi} = SC^+ \mathbf{d}$ because of preconditioning, thus the error in (4.1) can be bounded by

$$\left\| \tilde{\pi}^C\left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}}\right) - \pi^C\left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}}\right) \right\|_1 \|\boldsymbol{\phi}\|_\infty.$$

As we do not have a good bound on $\|\boldsymbol{\phi}\|_\infty$, we instead use an alternative upper bound on the error:

$$\sqrt{\mathcal{E}_r\left(\tilde{\pi}^C\left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}}\right) - \pi^C\left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}}\right)\right)} \sqrt{E_r(\boldsymbol{\phi})},$$

where $\mathcal{E}_r(\cdot)$ gives the energy to route a demand with resistances \mathbf{r} , and $E_r(\cdot)$ gives the energy of a potential embedding with resistances \mathbf{r} . As the standard interior point method step satisfies $E_r(\boldsymbol{\phi}) \leq 1$, all our efforts focus on ensuring that

$$\sqrt{\mathcal{E}_r\left(\tilde{\pi}^C\left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}}\right) - \pi^C\left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}}\right)\right)} \leq \varepsilon \tag{4.2}$$

for some error parameter ε . One issue is the fact that the energy depends on the current resistances, therefore even if at some point the error of the demand projection is low, after a few iterations it might increase because of resistance changes. We deal with this issue by taking the stability of resistances along the central path into account. This allows us to upper bound how much this error increases after a number of iterations. The resistance stability lemma is a generalization of the one used in [73].

Unfortunately, even though (4.2) seems like the right type of guarantee, it is unclear how to ensure that it is always true. Specifically, it involves efficiently computing the hitting probabilities from some vertex v to C in an appropriate norm, which ends up being non-trivial. Instead, we show

that the following *weaker* error bound can be ensured with high probability:

$$\left| \left\langle \tilde{\pi}^C \left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}} \right) - \pi^C \left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}} \right), \phi \right\rangle \right| \leq \varepsilon, \quad (4.3)$$

where ϕ is a *fixed* potential vector with $E_r(\phi) \leq 1$. Interestingly, this guarantee is still sufficient for our purposes.

Costs and general demand There is a fundamental obstacle to using the approach of [73] once edge costs are introduced. In particular, for the maximum flow problem, the demand pushed by the electrical flow in each iteration is an s - t demand, so—up to scaling—it is always constant. In minimum cost flow on the other hand, the augmenting flow is a multiple of $\mathbf{c} - \mathbf{R}^{-1} \mathbf{B} \mathbf{L}^+ \mathbf{B}^\top \mathbf{c}$. Here it is not possible to locate a sublinear number of congested edges just by looking at the electrical flow term $\mathbf{R}^{-1} \mathbf{B} \mathbf{L}^+ \mathbf{B}^\top \mathbf{c}$, as there might be significant cancellations with \mathbf{c} . We instead use the following equivalent form: $\frac{\frac{1}{s^+} - \frac{1}{s^-}}{\mathbf{r}} - \mathbf{R}^{-1} \mathbf{B} \mathbf{L}^+ \mathbf{B}^\top \frac{\frac{1}{s^+} - \frac{1}{s^-}}{\mathbf{r}}$, which allows us to ignore the first term because it is small and concentrate on the electrical flow term. One issue that arises is the fact that the demand vector $\mathbf{B}^\top \frac{\frac{1}{s^+} - \frac{1}{s^-}}{\mathbf{r}}$ now depends on slacks, and as a result changes throughout the interior point method. This issue can be handled relatively easily.

A more significant issue concerns the vertex sparsifier. In fact, the vertex sparsifier framework around which [73] is based only accepts demands that are supported on the vertex set C of the sparsifier. As $|C|$ is sublinear in n , this only captures demands with sublinear support, one such example being max flow with support 2. However, our demand vector $\mathbf{B}^\top \frac{\frac{1}{s^+} - \frac{1}{s^-}}{\mathbf{r}}$ in general will be supported on n vertices. Even though it might seem impossible to get around this issue, we show that the special structure of C allows us to push the demand to a small number of vertices. More specifically, we show that if one projects all of the demand onto C , the flow induced by this new demand will not differ much from the one with the original demand. Concretely, given a Laplacian system $\mathbf{L}\phi = \mathbf{d}$, we decompose it into two systems $\mathbf{L}\phi^{(1)} = \pi^C(\mathbf{d})$ and $\mathbf{L}\phi^{(2)} = \mathbf{d} - \pi^C(\mathbf{d})$, where $\pi^C(\mathbf{d})$ is the projection of \mathbf{d} onto C . Intuitively, the latter system computes the electrical flow to push all demands to C , and the former to serve this C -supported demand. We show that, as long as C is a *congestion reduction subset* (as it is also the case in [73]), $\phi^{(2)}$ has negligible contribution in the electrical flow, thus it can be ignored. More specifically, in Section 4.4.1 we prove the following lemma:

Lemma 4.4.3. *Consider a graph $G(V, E)$ with resistances \mathbf{r} and Laplacian \mathbf{L} , a β -congestion reduction subset C , and a demand $\mathbf{d} = \delta \mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}}$ for some $\delta > 0$ and $\mathbf{q} \in [-1, 1]^m$. Then, the potential embedding defined as*

$$\phi = \mathbf{L}^+ (\mathbf{d} - \pi^C(\mathbf{d}))$$

has congestion $\delta \cdot \tilde{O}(1/\beta^2)$, i.e. $\left\| \frac{\mathbf{B}\phi}{\sqrt{\mathbf{r}}} \right\|_\infty \leq \delta \cdot \tilde{O}(1/\beta^2)$.

Now, for computing $\phi^{(1)}$, we need to get an approximate estimate of $\pi^C(\mathbf{d})$. Even though the most natural approach would be to try to maintain $\pi^C(\mathbf{d})$ under vertex insertions to C , this approach has issues related to the fact that our error guarantee is based on a *fixed* potential vector. In particular, if we used an estimate of $\pi^C(\mathbf{d})$, then the potential vector in (4.3) would depend on the randomness of this estimate, and as a result the high probability guarantee would not work.

Instead, we show that it is not even necessary to maintain $\pi^C(\mathbf{d})$ very accurately. In fact, it suffices to *exactly* compute it only every few iterations of the algorithm, and use this estimate for the

calculation. What allows us to do this is the following lemma, which bounds the change of $\pi^C(\mathbf{d})$ measured in energy, after a sequence of vertex insertions and resistance changes.

Lemma 4.4.12. *Consider a graph $G(V, E)$ with resistances $\mathbf{r}^0, \mathbf{q}^0 \in [-1, 1]^m$, a β -congestion reduction subset C^0 , and a fixed sequence of updates, where the i -th update $i \in \{0, T-1\}$ is of the following form:*

- **ADDTERMINAL**(v^i): Set $C^{i+1} = C^i \cup \{v^i\}$ for some $v^i \in V \setminus C^i$, $q_e^{i+1} = q_e^i, r_e^{i+1} = r_e^i$
- **UPDATE**($e^i, \mathbf{q}, \mathbf{r}$): Set $C^{i+1} = C^i$, $q_e^{i+1} = q_e, r_e^{i+1} = r_e$, where $e^i \in E(C^i)$

Then, with high probability,

$$\sqrt{\mathcal{E}_{\mathbf{r}^T} \left(\pi^{C^0, \mathbf{r}^0} \left(\mathbf{B}^\top \frac{\mathbf{q}_S^0}{\sqrt{\mathbf{r}^0}} \right) - \pi^{C^T, \mathbf{r}^T} \left(\mathbf{B}^\top \frac{\mathbf{q}_S^T}{\sqrt{\mathbf{r}^T}} \right) \right)} \leq \tilde{O} \left(\max_{i \in \{0, \dots, T-1\}} \left\| \frac{\mathbf{r}^T}{\mathbf{r}^i} \right\|_\infty^{1/2} \beta^{-2} \right) \cdot T.$$

If we call this demand projection estimate π_{old} , the quantity that we would like to maintain (4.1) now becomes

$$\langle \mathbf{q}, \boldsymbol{\rho} \rangle \approx \left\langle \pi^C \left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}} \right), SC^+ \pi_{old} \right\rangle.$$

Therefore all that's left is to efficiently maintain approximations to demand projections of the form $\pi^C \left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}} \right)$.

Bounding demand projections. An important component for showing that demand projections can be updated efficiently is bounding the magnitude of an entry $\pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{r_e}} \right)$ of the projection, for some fixed edge $e = (u, w)$. This is apparent in the following identity which shows how a demand projection changes after inserting a vertex:

$$\pi^{C \cup \{v\}}(\mathbf{d}) = \pi^C(\mathbf{d}) + \pi_v^{C \cup \{v\}}(\mathbf{d}) \cdot (\mathbf{1}_v - \pi^C(\mathbf{1}_v)). \quad (4.4)$$

In [73] this projection entry is upper bounded by $(p_v^{C \cup \{v\}}(u) + p_v^{C \cup \{v\}}(w)) \cdot \frac{1}{\sqrt{r_e}}$, where $p_v^{C \cup \{v\}}(u)$ is the probability that a random walk starting at u hits v before C . This bound can be very bad as r_e can be arbitrarily small, although in the particular case of max flow it is possible to show that such low-resistance edges cannot get congested and thus are not of interest.

In order to overcome this issue, we provide a different bound, which in contrast works best when r_e is small.

Lemma 4.4.6. *Consider a graph $G(V, E)$ with resistances \mathbf{r} and a subset of vertices $C \subseteq V$. For any vertex $v \in V \setminus C$ we have that*

$$\left| \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{\mathbf{r}}} \right) \right| \leq (p_v^{C \cup \{v\}}(u) + p_v^{C \cup \{v\}}(w)) \cdot \frac{\sqrt{r_e}}{R_{eff}(v, e)}.$$

Here $R_{eff}(v, e)$ is the effective resistance between v and e . In fact, together with the other upper bound mentioned above, this implies that

$$\left| \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{\mathbf{r}}} \right) \right| \leq (p_v^{C \cup \{v\}}(u) + p_v^{C \cup \{v\}}(w)) \cdot \frac{1}{\sqrt{R_{eff}(v, e)}},$$

which no longer depends on the value of the resistance r_e .

As we will see, it suffices to approximate $\pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{r}} \right)$ up to additive accuracy roughly $\hat{\varepsilon} \cdot (p_v^{C \cup \{v\}}(u) + p_v^{C \cup \{v\}}(w)) / \sqrt{R_{\text{eff}}(C, v)}$ for some error parameter $\hat{\varepsilon} > 0$. Thus, Lemma 4.4.6 immediately implies that for any edge e such that $R_{\text{eff}}(v, e) \mathbf{g} R_{\text{eff}}(C, v)$, this term is small enough to begin with, and thus can be ignored.

Important edges. In order to ensure that the demand projection can be updated efficiently, we focus only on the demand coming from a special set of edges, which we call *important*. These are the edges that are close (in effective resistance metric) to C relative to their own resistance r_e . In fact, the farther an edge is from C in this sense, the smaller its worst-case congestion, and so non-important edges do not influence the set of congested edges that we are looking for. At a high level, this is because parts of the graph that are very far in the potential embedding have minimal interactions with each other.

Definition 4.1.2 (Important edges). *An edge $e \in E$ is called ε -important (or just important) if $R_{\text{eff}}(C, e) \leq r_e / \varepsilon^2$.*

Based on the above discussion, we seek to find congested edges *only* among important edges.

Lemma 4.4.9 (Localization lemma). *Let ϕ^* be any solution of*

$$\mathbf{L}\phi^* = \delta \cdot \pi^C \left(\mathbf{B}^\top \frac{\mathbf{p}}{\sqrt{\mathbf{r}}} \right),$$

where \mathbf{r} are any resistances, $\mathbf{p} \in [-1, 1]^m$, and $C \subseteq V$. Then, for any $e \in E$ that is not ε -important we have $\left| \frac{\mathbf{B}\phi^*}{\sqrt{\mathbf{r}}} \right|_e \leq 12\varepsilon$.

One issue is that the set of important edges changes whenever C changes. However, we show that, because of the stability of resistances along the central path, the set of important edges only needs to be updated once every few iterations.

4.2 Preliminaries

Given $x, y \in \mathbb{R}$ and $\alpha \in \mathbb{R}_{\geq 1}$, we say that x and y α -approximate each other and write $x \approx_\alpha y$ if $\alpha^{-1} \leq x/y \leq \alpha$.

Definition 4.2.1 (Schur complement). *Given a graph $G(V, E)$ with Laplacian $\mathbf{L} \in \mathbb{R}^{n \times n}$ and a vertex subset $C \subseteq V$ as well as $F = V \setminus C$, $SC(G, C) := \mathbf{L}_{CC} - \mathbf{L}_{CF} \mathbf{L}_{FF}^{-1} \mathbf{L}_{FC}$ (or just SC) is called the Schur complement of G onto C .*

Fact 4.2.2 (Cholesky factorization). *Given a matrix $\mathbf{L} \in \mathbb{R}^{n \times n}$, a subset $C \subseteq [n]$, and $F = [n] \setminus C$, we have*

$$\mathbf{L}^+ = \begin{pmatrix} \mathbf{I} & -\mathbf{L}_{FF}^{-1} \mathbf{L}_{FC} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{L}_{FF}^{-1} & \mathbf{0} \\ \mathbf{0} & SC(\mathbf{L}, C)^+ \end{pmatrix} \begin{pmatrix} \mathbf{I} & \mathbf{0} \\ -\mathbf{L}_{CF} \mathbf{L}_{FF}^{-1} & \mathbf{I} \end{pmatrix}.$$

4.2.1 Random walks

Definition 4.2.3 (Hitting probabilities). *Consider a graph $G(V, E)$ with resistances \mathbf{r} . For any $u, v \in V$, $C \subseteq V$, we denote by $p_v^{C, \mathbf{r}}(u)$ the probability that for random walk that starts from u*

and uses edges with probability proportional to $\frac{1}{r}$, the first vertex of C to be visited is v . When not ambiguous, we will use the notation $p_v^C(u)$.

Definition 4.2.4 (Demand projection). Consider a graph $G(V, E)$ and a demand vector \mathbf{d} . For any $v \in V$, $C \subseteq V$, we define $\pi_v^{C,r}(\mathbf{d}) = \sum_{u \in V} d_u p_v^{C,r}(u)$ and call the resulting vector $\boldsymbol{\pi}^{C,r}(\mathbf{d}) \in \mathbb{R}^n$ the demand projection of \mathbf{d} onto C . When not ambiguous, we will use the notation $\boldsymbol{\pi}^C(\mathbf{d})$.

For convenience, when we write $\boldsymbol{\pi}^C(\mathbf{d})$ we might also refer to the restriction of this vector to C . This will be clear from the context, and, as $\pi_v^C(\mathbf{d}) = 0$ for any $v \notin C$, no ambiguity is introduced.

Fact 4.2.5 ([73]). Given a graph $G(V, E)$ with Laplacian \mathbf{L} , a vertex subset $C \subseteq V$, and $\mathbf{d} \in \mathbb{R}^n$, we have

$$\boldsymbol{\pi}^C(\mathbf{d}) = \mathbf{d}_C - \mathbf{L}_{CF} \mathbf{L}_{FF}^{-1} \mathbf{d}_F.$$

Additionally,

$$[\mathbf{L}^+ \mathbf{d}]_C = \mathbf{S}C^+ \boldsymbol{\pi}^C(\mathbf{d}),$$

where $\mathbf{S}C$ is the Schur complement of G onto C .

An important property of the demand projection is that the energy required to route it is upper bounded by the energy required to route the original demand. The proof can be found in Section 9.2.2.

Lemma 4.2.6. Let \mathbf{d} be a demand vector, let \mathbf{r} be resistances, and let $C \subseteq V$ be a subset of vertices. Then

$$\mathcal{E}_r(\boldsymbol{\pi}^C(\mathbf{d})) \leq \mathcal{E}_r(\mathbf{d}).$$

The following lemma relates the effective resistance between a vertex and a vertex set, to the energy to route a particular demand, based on a demand projection.

Lemma 4.2.7 (Effective resistance and hitting probabilities). Given a graph $G(V, E)$ with resistances \mathbf{r} , any vertex set $A \subseteq V$ and vertex $u \in V \setminus A$, we have $R_{eff}(u, A) = \mathcal{E}_r(\mathbf{1}_u - \boldsymbol{\pi}^A(\mathbf{1}_u))$.

Proof. Let \mathbf{L} be the Laplacian of G with resistances \mathbf{r} and $F = V \setminus A$. We first prove that

$$\mathcal{E}_r(\mathbf{1}_u - \boldsymbol{\pi}^A(\mathbf{1}_u)) = \mathbf{1}_u^\top \mathbf{L}_{FF}^{-1} \mathbf{1}_u.$$

This is because

$$\begin{aligned} & \mathcal{E}_r(\mathbf{1}_u - \boldsymbol{\pi}^A(\mathbf{1}_u)) \\ &= \langle \mathbf{1}_u - \boldsymbol{\pi}^A(\mathbf{1}_u), \mathbf{L}^+(\mathbf{1}_u - \boldsymbol{\pi}^A(\mathbf{1}_u)) \rangle \\ &= \left\langle \mathbf{1}_u - \begin{pmatrix} \mathbf{0} & \mathbf{0} \\ -\mathbf{L}_{AF} \mathbf{L}_{FF}^{-1} & \mathbf{I} \end{pmatrix} \mathbf{1}_u, \mathbf{L}^+ \left(\mathbf{1}_u - \begin{pmatrix} \mathbf{0} & \mathbf{0} \\ -\mathbf{L}_{AF} \mathbf{L}_{FF}^{-1} & \mathbf{I} \end{pmatrix} \mathbf{1}_u \right) \right\rangle \\ &= \left\langle \begin{pmatrix} \mathbf{I} & \mathbf{0} \\ -\mathbf{L}_{AF} \mathbf{L}_{FF}^{-1} & \mathbf{I} \end{pmatrix} \left(\mathbf{1}_u + \begin{pmatrix} \mathbf{0} \\ \mathbf{L}_{AF} \mathbf{L}_{FF}^{-1} \mathbf{1}_u \end{pmatrix} \right), \right. \\ & \quad \left. \begin{pmatrix} \mathbf{L}_{FF}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{S}C(\mathbf{L}, A)^+ \end{pmatrix} \begin{pmatrix} \mathbf{I} & \mathbf{0} \\ -\mathbf{L}_{AF} \mathbf{L}_{FF}^{-1} & \mathbf{I} \end{pmatrix} \left(\mathbf{1}_u + \begin{pmatrix} \mathbf{0} \\ \mathbf{L}_{AF} \mathbf{L}_{FF}^{-1} \mathbf{1}_u \end{pmatrix} \right) \right\rangle \\ &= \left\langle \mathbf{1}_u, \begin{pmatrix} \mathbf{L}_{FF}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{S}C(\mathbf{L}, A)^+ \end{pmatrix} \mathbf{1}_u \right\rangle \\ &= \langle \mathbf{1}_u, \mathbf{L}_{FF}^{-1} \mathbf{1}_u \rangle. \end{aligned}$$

On the other hand, note that $R_{eff}(u, A) = \widehat{R}_{eff}(u, \widehat{a})$, where \widehat{R} are the effective resistances in a graph \widehat{G} that results after contracting A to a new vertex \widehat{a} . It is easy to see that the Laplacian of this new graph is

$$\widehat{\mathbf{L}} = \begin{pmatrix} \mathbf{L}_{FF} & \mathbf{L}_{FA}\mathbf{1} \\ \mathbf{1}^\top \mathbf{L}_{AF} & \mathbf{1}^\top \mathbf{L}_{FA}\mathbf{1} \end{pmatrix}.$$

We look at the system $\widehat{\mathbf{L}} \begin{pmatrix} x \\ a \end{pmatrix} = \mathbf{1}_u - \mathbf{1}_{\widehat{a}}$, where a is a scalar. The solution is given by

$$\mathbf{x} = \mathbf{L}_{FF}^{-1}(\mathbf{1}_u - a \cdot \mathbf{L}_{FA}\mathbf{1}).$$

However, as $\mathbf{1} \in \ker(\widehat{\mathbf{L}})$ by the fact that it is a Laplacian, we can assume that $a = 0$ by shifting. Therefore $\mathbf{x} = \mathbf{L}_{FF}^{-1}\mathbf{1}_u$, and so we can conclude that

$$\begin{aligned} R_{eff}(u, A) &= \langle \mathbf{1}_u - \mathbf{1}_{\widehat{a}}, \widehat{\mathbf{L}}^+(\mathbf{1}_u - \mathbf{1}_{\widehat{a}}) \rangle \\ &= \langle \mathbf{1}_u, \mathbf{L}_{FF}^{-1}\mathbf{1}_u \rangle. \end{aligned}$$

So we have proved that $R_{eff}(u, A) = \mathcal{E}_r(\mathbf{1}_u - \boldsymbol{\pi}^A(\mathbf{1}_u))$ and we are done. \square

Finally, the following lemma relates the effective resistance between a vertex and a vertex set, to the effective resistance between vertices.

Lemma 4.2.8. *Given a graph $G(V, E)$ with resistances \mathbf{r} , any vertex set $A \subseteq V$ and vertex $u \in V \setminus A$, we have*

$$\frac{1}{|A|} \cdot \min_{v \in A} R_{eff}(u, v) \leq R_{eff}(u, A) \leq \min_{v \in A} R_{eff}(u, v).$$

Proof. Let \mathbf{L} be the Laplacian of G with resistances \mathbf{r} , and note that $R_{eff}(u, v) = \left\| \mathbf{L}^{+/2}(\mathbf{1}_u - \mathbf{1}_v) \right\|_2^2$ and, by Lemma 4.2.7, $R_{eff}(u, A) = \left\| \mathbf{L}^{+/2}(\mathbf{1}_u - \boldsymbol{\pi}^A(\mathbf{1}_u)) \right\|_2^2$. Expanding the latter, we have

$$\begin{aligned} R_{eff}(u, A) &= \left\| \sum_{v \in A} \pi_v^A(\mathbf{1}_u) \cdot \mathbf{L}^{+/2}(\mathbf{1}_u - \mathbf{1}_v) \right\|_2^2 \\ &= \sum_{v \in A} (\pi_v^A(\mathbf{1}_u))^2 \left\| \mathbf{L}^{+/2}(\mathbf{1}_u - \mathbf{1}_v) \right\|_2^2 + \sum_{v \in A} \sum_{\substack{v' \in A \\ v' \neq v}} \pi_v^A(\mathbf{1}_u) \pi_{v'}^A(\mathbf{1}_u) \langle \mathbf{1}_u - \mathbf{1}_{v'}, \mathbf{L}^+(\mathbf{1}_u - \mathbf{1}_v) \rangle. \end{aligned}$$

Now, note that $\pi_v^A(\mathbf{1}_u), \pi_{v'}^A(\mathbf{1}_u) \geq 0$. Additionally, let $\boldsymbol{\phi} = \mathbf{L}^+(\mathbf{1}_u - \mathbf{1}_v)$ be the potential embedding that induces a 1-unit electrical flow from v to u . As the potential embedding stretches between $\boldsymbol{\phi}_v$

and ϕ_u , we have that $\phi_{v'} \leq \phi_u$, so $\langle \mathbf{1}_u - \mathbf{1}_{v'}, \mathbf{L}^+(\mathbf{1}_u - \mathbf{1}_{v'}) \rangle = \phi_u - \phi_{v'} \geq 0$. Therefore,

$$\begin{aligned} R_{eff}(u, A) &\geq \sum_{v \in A} (\pi_v^A(\mathbf{1}_u))^2 \left\| \mathbf{L}^{+/2}(\mathbf{1}_u - \mathbf{1}_v) \right\|_2^2 \\ &\geq \frac{1}{|A|} \sum_{v \in A} \pi_v^A(\mathbf{1}_u) \cdot R_{eff}(u, v) \\ &\geq \frac{1}{|A|} \min_{v \in A} R_{eff}(u, v), \end{aligned}$$

where we used the Cauchy-Schwarz inequality and the fact that $\sum_{v \in A} \pi_v^A(\mathbf{1}_u) = 1$. \square

4.3 Interior Point Method with Dynamic Data Structures

The goal of this section is to show that, given a data structure for approximating electrical flows in sublinear time, we can execute the min cost flow interior point method with total runtime faster than $\tilde{O}(m^{3/2})$.

4.3.1 LP formulation and background

We present the interior point method setup that we will use, which is from [10]. Our goal is to solve the following minimum cost flow linear program:

$$\begin{aligned} &\min \langle \mathbf{c}, \mathbf{C}\mathbf{x} \rangle \\ &\mathbf{0} \leq \mathbf{f}^0 + \mathbf{C}\mathbf{x} \leq \mathbf{u}, \end{aligned}$$

where \mathbf{f}^0 is a flow with $\mathbf{B}^\top \mathbf{f}^0 = \mathbf{d}$ and \mathbf{C} is an $m \times (m - n + 1)$ matrix whose image is the set of circulations in G .

In order to use an interior point method, the following log barrier objective is defined:

$$\min_{\mathbf{x}} F_\mu(\mathbf{x}) = \left\langle \frac{\mathbf{c}}{\mu}, \mathbf{C}\mathbf{x} \right\rangle - \sum_{e \in E} (\log(\mathbf{f}^0 + \mathbf{C}\mathbf{x})_e + \log(\mathbf{u} - (\mathbf{f}^0 + \mathbf{C}\mathbf{x}))_e). \quad (4.5)$$

For any parameter $\mu > 0$, the optimality condition of (4.5) is called the *centrality condition* and is given by

$$\mathbf{C}^\top \left(\frac{\mathbf{c}}{\mu} + \frac{1}{\mathbf{s}^+} - \frac{1}{\mathbf{s}^-} \right) = \mathbf{0}, \quad (4.6)$$

where $\mathbf{f} = \mathbf{f}^0 + \mathbf{C}\mathbf{x}$, and $\mathbf{s}^+ = \mathbf{u} - \mathbf{f}$, $\mathbf{s}^- = \mathbf{f}$ are called the *positive* and *negative* slacks of \mathbf{f} respectively.

This leads us to the following definitions.

Definition 4.3.1 (μ -central flow). *Given a minimum cost flow instance with costs \mathbf{c} , demands \mathbf{d} and capacities \mathbf{u} , as well as a parameter $\mu > 0$, we will say that a flow \mathbf{f} (and its corresponding slacks \mathbf{s} and resistances \mathbf{r}) is μ -central if $\mathbf{B}^\top \mathbf{f} = \mathbf{d}$, $\mathbf{s} > \mathbf{0}$, and it satisfies the centrality condition*

(4.6), i.e.

$$\mathbf{C}^\top \left(\frac{\mathbf{c}}{\mu} + \frac{1}{\mathbf{s}^+} - \frac{1}{\mathbf{s}^-} \right) = \mathbf{0}.$$

Additionally, we will denote such flow by $\mathbf{f}(\mu)$ (and its corresponding slacks and resistances by $\mathbf{s}(\mu)$ and $\mathbf{r}(\mu)$, respectively).

Definition 4.3.2 ((μ, α) -central flow). *Given parameters $\mu > 0$ and $\alpha \geq 1$, we will say that a flow \mathbf{f} with resistances $\mathbf{r} > \mathbf{0}$ is (μ, α) -central if $\mathbf{r} \approx_\alpha \mathbf{r}(\mu)$. We will also call its corresponding slacks \mathbf{s} and resistances \mathbf{r} (μ, α) -central.*

Given a μ -central flow \mathbf{f} and some step size $\delta > 0$, the standard (Newton) step to obtain an approximately $\mu/(1 + \delta)$ -central flow $\mathbf{f}' = \mathbf{f} + \tilde{\mathbf{f}}$ is given by

$$\begin{aligned} \tilde{\mathbf{f}} &= -\frac{\delta}{\mu} \frac{\mathbf{c}}{\mathbf{r}} + \frac{\delta}{\mu} \mathbf{R}^{-1} \mathbf{B} \mathbf{L}^+ \mathbf{B}^\top \frac{\mathbf{c}}{\mathbf{r}} \\ &= \delta \frac{\frac{1}{\mathbf{s}^+} - \frac{1}{\mathbf{s}^-}}{\mathbf{r}} - \delta \mathbf{R}^{-1} \mathbf{B} \mathbf{L}^+ \mathbf{B}^\top \frac{\frac{1}{\mathbf{s}^+} - \frac{1}{\mathbf{s}^-}}{\mathbf{r}} \\ &= \delta \cdot g(\mathbf{s}) - \delta \mathbf{R}^{-1} \mathbf{B} \mathbf{L}^+ \mathbf{B}^\top g(\mathbf{s}) \end{aligned}$$

where $\mathbf{r} = \frac{1}{(\mathbf{s}^+)^2} + \frac{1}{(\mathbf{s}^-)^2}$ and we have denoted $g(\mathbf{s}) = \frac{\frac{1}{\mathbf{s}^+} - \frac{1}{\mathbf{s}^-}}{\mathbf{r}}$.

Fact 4.3.3. *Using known scaling arguments, can assume that costs and capacities are bounded by $\text{poly}(m)$, while only incurring an extra logarithmic dependence in the largest network parameter [71].*

We also use the fact that the resistances in the interior point method are never too large, which is proved in Appendix 9.2.2.

Fact 4.3.4. *For any $\mu \in (1/\text{poly}(m), \text{poly}(m))$, we have $\|\mathbf{r}(\mu)\|_\infty \leq m^{\tilde{O}(\log m)}$.*

4.3.2 Making progress with approximate electrical flows

The following lemma shows that we can make k steps of the interior point method by computing $O(k^4)$ $(1 + O(k^{-6}))$ -approximate electrical flows. The proof is essentially the same as in [73], but we provide it for completeness in Appendix 9.2.3.

Lemma 4.3.5. *Let $\mathbf{f}^1, \dots, \mathbf{f}^{T+1}$ be flows with slacks \mathbf{s}^t and resistances \mathbf{r}^t for $t \in [T + 1]$, where $T = \frac{k}{\varepsilon_{\text{step}}}$ for some $k \leq \sqrt{m}/10$ and $\varepsilon_{\text{step}} = 10^{-5}k^{-3}$, such that*

- \mathbf{f}^1 is $(\mu, 1 + \varepsilon_{\text{solve}}/8)$ -central for $\varepsilon_{\text{solve}} = 10^{-5}k^{-3}$
- For all $t \in [T]$ and $e \in E$, $f_e^{t+1} = \begin{cases} f_e(\mu) + \varepsilon_{\text{step}} \sum_{i=1}^t \tilde{f}_e^i & \text{if } \exists i \in [t] : \tilde{f}_e^i \neq 0 \\ f_e^1 & \text{otherwise} \end{cases}$, where

$$\tilde{\mathbf{f}}^{*t} = \delta g(\mathbf{s}^t) - \delta (\mathbf{R}^t)^{-1} \mathbf{B} \left(\mathbf{B}^\top (\mathbf{R}^t)^{-1} \mathbf{B} \right)^+ \mathbf{B}^\top g(\mathbf{s}^t)$$

for $\delta = \frac{1}{\sqrt{m}}$ and

$$\left\| \sqrt{\mathbf{r}^t} \left(\tilde{\mathbf{f}}^{*t} - \tilde{\mathbf{f}}^t \right) \right\|_\infty \leq \varepsilon$$

for $\varepsilon = 10^{-6}k^{-6}$.

Then, setting $\varepsilon_{\text{step}} = \varepsilon_{\text{solve}} = 10^{-5}k^{-3}$ and $\varepsilon = 10^{-6}k^{-6}$ we get that $\mathbf{s}^{T+1} \approx_{1.1} \mathbf{s} \left(\mu / (1 + \varepsilon_{\text{step}} \delta)^{k\varepsilon_{\text{step}}^{-1}} \right)$.

From now and for the rest of Section 4.3 we fix the values of $\varepsilon_{\text{step}}, \varepsilon_{\text{solve}}, \varepsilon$ based on this lemma. Using this lemma together with the following recentering procedure also used in [73], we can exactly compute a $\left(\mu / (1 + \varepsilon_{\text{step}} / \sqrt{m})^{k\varepsilon_{\text{step}}^{-1}} \right)$ -central flow.

Lemma 4.3.6. *Given a flow \mathbf{f} with slacks \mathbf{s} such that $\mathbf{s} \approx_{1.1} \mathbf{s}(\mu)$ for some $\mu > 0$, we can compute $\mathbf{f}(\mu)$ in $\tilde{O}(m)$.*

4.3.3 The LOCATOR data structure

From the previous lemma it becomes obvious that the only thing left is to maintain in sublinear time an approximation to

$$\delta g(\mathbf{s}^t) - \delta (\mathbf{R}^t)^{-1} \mathbf{B} (\mathbf{B}^\top (\mathbf{R}^t)^{-1} \mathbf{B})^+ \mathbf{B}^\top g(\mathbf{s}^t).$$

for $\delta = 1/\sqrt{m}$. This is the job of the $(\alpha, \beta, \varepsilon)$ -LOCATOR data structure, which computes all the entries of this vector that have magnitude $\geq \varepsilon$. We note that the guarantees of this data structure are similar to the ones in [73], but our locator requires an extra parameter α which is a measure of how much resistances can deviate before a full recomputation has to be made.

Definition 4.3.7 ($(\alpha, \beta, \varepsilon)$ -LOCATOR). *An $(\alpha, \beta, \varepsilon)$ -LOCATOR is a data structure that maintains valid slacks \mathbf{s} and resistances \mathbf{r} , and can support the following operations against oblivious adversaries with high probability:*

- INITIALIZE(\mathbf{f}): Set $\mathbf{s}^+ = \mathbf{u} - \mathbf{f}$, $\mathbf{s}^- = \mathbf{f}$, $\mathbf{r} = \frac{1}{(s^+)^2} + \frac{1}{(s^-)^2}$.
- UPDATE(e, \mathbf{f}): Set $s_e^+ = u_e - f_e$, $s_e^- = f_e$, $r_e = \frac{1}{(s_e^+)^2} + \frac{1}{(s_e^-)^2}$. Works under the condition that

$$r_e^{\max} / \alpha \leq r_e \leq \alpha \cdot r_e^{\min},$$

where r_e^{\max} and r_e^{\min} are the maximum and minimum resistance values that edge e has had since the last call to BATCHUPDATE.

- BATCHUPDATE(Z, \mathbf{f}): Set $s_e^+ = u_e - f_e$, $s_e^- = f_e$, $r_e = \frac{1}{(s_e^+)^2} + \frac{1}{(s_e^-)^2}$ for all $e \in Z$.
- SOLVE(): Let

$$\tilde{\mathbf{f}}^* = \delta g(\mathbf{s}) - \delta \mathbf{R}^{-1} \mathbf{B} (\mathbf{B}^\top \mathbf{R}^{-1} \mathbf{B})^+ \mathbf{B}^\top g(\mathbf{s}), \quad (4.7)$$

where $\delta = \frac{1}{\sqrt{m}}$. Returns an edge set Z of size $\tilde{O}(\varepsilon^{-2})$ that with high probability contains all e such that $\sqrt{r_e} \left| \tilde{f}_e^* \right| \geq \varepsilon$.

The data structure works as long as the total number of calls to UPDATE, plus the sum of $|Z|$ for all calls to BATCHUPDATE is $O(\beta m)$.

In Section 4.4 we will prove the following lemma, which constructs an $(\alpha, \beta, \varepsilon)$ -LOCATOR and outlines its runtime guarantees:

Lemma 4.3.8 (Efficient $(\alpha, \beta, \varepsilon)$ -LOCATOR). *For any graph $G(V, E)$ and parameters $\alpha \geq 1$, $\beta \in (0, 1)$, $\varepsilon \geq \tilde{\Omega}(\beta^{-2}m^{-1/2})$, and $\hat{\varepsilon} \in \left(\tilde{\Omega}(\beta^{-2}m^{-1/2}), \varepsilon\right)$, there exists an $(\alpha, \beta, \varepsilon)$ -LOCATOR for G with the following runtimes per operation:*

- INITIALIZE(\mathbf{f}): $\tilde{O}\left(m \cdot (\hat{\varepsilon}^{-4}\beta^{-8} + \hat{\varepsilon}^{-2}\varepsilon^{-2}\alpha^2\beta^{-4})\right)$.
- UPDATE(e, \mathbf{f}): $\tilde{O}\left(m \cdot \frac{\hat{\varepsilon}\alpha^{1/2}}{\varepsilon^3} + \hat{\varepsilon}^{-4}\varepsilon^{-2}\beta^{-8} + \hat{\varepsilon}^{-2}\varepsilon^{-4}\alpha^2\beta^{-6}\right)$ amortized.
- BATCHUPDATE(Z, \mathbf{f}): $\tilde{O}\left(m \cdot \frac{1}{\varepsilon^2} + |Z| \cdot \frac{1}{\varepsilon^2\beta^2}\right)$.
- SOLVE(): $\tilde{O}\left(\beta m \cdot \frac{1}{\varepsilon^2}\right)$.

Note that even though a LOCATOR computes a set that contains all ε -congested edges, it does not return the actual flow values. The reason for that is that it only works against oblivious adversaries, and allowing (randomized) flow values to affect future updates constitutes an adaptive adversary. As in [73], we resolve this by sanitizing the outputs through a different data structure called CHECKER, which computes the flow values and works against semi-adaptive adversaries. As the definition and implementation of CHECKER is orthogonal to our contribution and also does not affect the final runtime, we defer the discussion to Appendix 9.2.6. To simplify the presentation in this section, we instead define the following idealized version of it, called PERFECTCHECKER.

Definition 4.3.9 (ε -PERFECTCHECKER). *For any error $\varepsilon > 0$, an ε -PERFECTCHECKER is an oracle that given a graph $G(V, E)$, slacks \mathbf{s} , resistances \mathbf{r} , supports the following operations:*

- UPDATE(e, \mathbf{f}): Set $s_e^+ = u_e - f_e$, $s_e^- = f_e$, $r_e = \frac{1}{(s_e^+)^2} + \frac{1}{(s_e^-)^2}$.
- CHECK(e): Compute a flow value \tilde{f}_e such that $\sqrt{r_e} \left| \tilde{f}_e - \tilde{f}_e^* \right| \leq \varepsilon$, where

$$\tilde{\mathbf{f}}^* = \delta g(\mathbf{s}) - \delta \mathbf{R}^{-1} \mathbf{B} (\mathbf{B}^\top \mathbf{R}^{-1} \mathbf{B})^+ \mathbf{B}^\top g(\mathbf{s}),$$

with $\delta = 1/\sqrt{m}$. If $\sqrt{r_e} \left| \tilde{f}_e \right| < \varepsilon/2$ return 0, otherwise return \tilde{f}_e .

4.3.4 The minimum cost flow algorithm

Now, we will show how the data structure defined in Section 4.3.3 can be used to make progress along the central path. The main lemma that analyzes the performance of the minimum cost flow algorithm given access to an $(\alpha, \beta, \varepsilon)$ -LOCATOR is Lemma 4.3.10. Also, the skeleton of the algorithm is described in Algorithm 1.

Lemma 4.3.10 (MINCOSTFLOW). *Let \mathcal{L} be an $(\alpha, \beta, \varepsilon)$ -LOCATOR, \mathbf{f} be a μ -central flow where $\mu = \text{poly}(m)$, and $k \in [m^{1/316}]$, $\beta \geq \tilde{\Omega}(k^3/m^{1/4})$, $\hat{T} \in \left[\tilde{O}(m^{1/2}/k)\right]$ be some parameters. There is an algorithm that with high probability computes a μ' -central flow \mathbf{f}' , where $\mu' \leq m^{-10}$. Additionally, the algorithm runs in time $\tilde{O}(m^{3/2}/k)$, plus*

- $\tilde{O}(k^3\beta^{-1/2})$ calls to \mathcal{L} .INITIALIZE,
- $\tilde{O}(m^{1/2}k^3)$ calls to \mathcal{L} .SOLVE,

Algorithm 1 Minimum Cost Flow

```

1: procedure MINCOSTFLOW( $G, \mathbf{c}, \mathbf{d}, \mathbf{u}$ )
2:    $\bar{\mathbf{f}}, \mu = \text{INITIALIZE}(G, \mathbf{c}, \mathbf{d}, \mathbf{u})$  ▷ Lemma 4.3.12.  $\bar{\mathbf{f}}$  is  $\mu$ -central at all times.
3:    $i = 0$ 
4:   while  $\mu \geq m^{-10}$  do
5:     if  $i$  is a multiple of  $\lfloor \varepsilon_{\text{solve}} \sqrt{\beta m} / k \rfloor$  then ▷ Re-initialize when  $|C|$  exceeds  $O(\beta m)$ .
6:        $\mathcal{L} = \text{LOCATOR.INITIALIZE}(\bar{\mathbf{f}})$  with error  $\varepsilon/2$ 
7:     if  $i$  is a multiple of  $\lfloor \varepsilon_{\text{solve}} \sqrt{\beta_{\text{CHECKER}} m} / k \rfloor$  then
8:        $\mathcal{C}^i = \text{CHECKER.INITIALIZE}(\bar{\mathbf{f}}, \varepsilon, \beta_{\text{CHECKER}})$  for  $i \in [k\varepsilon_{\text{step}}^{-1}]$ 
9:     if  $i$  is a multiple of  $\lfloor 0.5\alpha^{1/4} / k - 1 \rfloor$  then ▷ Update important edges when  $\mathcal{L}.\mathbf{r}^0$  expires
10:       $\mathcal{L}.\text{BATCHUPDATE}(\emptyset)$ 
11:       $\bar{\mathbf{f}}, \mu = \text{MULTISTEP}(\bar{\mathbf{f}}, \mu)$ 
12:      if  $i$  is a multiple of  $T$  then
13:         $Z = \emptyset$ 
14:        for  $e \in E$  do
15:           $\bar{s}_e^+ = u_e - \bar{f}_e, \bar{s}_e^- = \bar{f}_e$ 
16:          if  $\bar{s}_e^+ \not\approx_{\varepsilon_{\text{solve}}/16} \mathcal{L}.s_e^+$  or  $\bar{s}_e^- \not\approx_{\varepsilon_{\text{solve}}/16} \mathcal{L}.s_e^-$  then
17:             $\mathcal{C}^i.\text{UPDATE}(e, \bar{\mathbf{f}})$  for  $i \in [k\varepsilon_{\text{step}}^{-1}]$ 
18:             $Z = Z \cup \{e\}$ 
19:           $\mathcal{L}.\text{BATCHUPDATE}(Z, \bar{\mathbf{f}})$ 
20:        else
21:          for  $e \in E$  do
22:             $\bar{s}_e^+ = u_e - \bar{f}_e, \bar{s}_e^- = \bar{f}_e$ 
23:            if  $\bar{s}_e^+ \not\approx_{\varepsilon_{\text{solve}}/8} \mathcal{L}.s_e^+$  or  $\bar{s}_e^- \not\approx_{\varepsilon_{\text{solve}}/8} \mathcal{L}.s_e^-$  then
24:               $\mathcal{C}^i.\text{UPDATE}(e, \bar{\mathbf{f}})$  for  $i \in [k\varepsilon_{\text{step}}^{-1}]$ 
25:               $\mathcal{L}.\text{UPDATE}(e, \bar{\mathbf{f}})$ 
26:           $i = i + 1$ 
27:   return ROUND( $G, \mathbf{c}, \mathbf{d}, \mathbf{u}, \bar{\mathbf{f}}$ ) ▷ Lemma 4.3.13

```

- $\tilde{O}\left(m^{1/2}\left(k^6\hat{T} + k^{15}\right)\right)$ calls to $\mathcal{L}.\text{UPDATE}$,
- $\tilde{O}\left(m^{1/2}\alpha^{-1/4}\right)$ calls to $\mathcal{L}.\text{BATCHUPDATE}(\emptyset)$, and
- $\tilde{O}\left(m^{1/2}k^{-1}\hat{T}^{-1}\right)$ calls to $\mathcal{L}.\text{BATCHUPDATE}(Z, \bar{\mathbf{f}})$ for some $Z \neq \emptyset, \bar{\mathbf{f}}$. Additionally, the sum of $|Z|$ over all such calls is $\tilde{O}\left(mk^3\beta^{1/2}\right)$.

The proof appears in Appendix 9.2.3. Its main ingredient is the following lemma, which easily follows from Lemma 4.3.5 and essentially shows how k steps of the interior point method can be performed in $\tilde{O}(m)$ instead of $\tilde{O}(mk)$. Its proof appears in Appendix 9.2.3.

Lemma 4.3.11 (MULTISTEP). *Let $k \in \{1, \dots, \sqrt{m}/10\}$. We are given $\mathbf{f}(\mu)$, an $(\alpha, \beta, \varepsilon/2)$ -LOCATOR \mathcal{L} , and an ε -PERFECTCHECKER \mathcal{C} , such that*

- $\mathcal{L}.\mathbf{r} = \mathcal{C}.\mathbf{r}$ are $(\mu, 1 + \varepsilon_{\text{solve}}/8)$ -central resistances, and
- $\mathcal{L}.\mathbf{r}^0$ are $(\mu^0, 1 + \varepsilon_{\text{solve}}/8)$ -central resistances, where $\mu^0 \leq \mu \cdot (1 + \varepsilon_{\text{step}}/\sqrt{m})^{\hat{T}}$ and $\hat{T} = (0.5\alpha^{1/4} - k)\varepsilon_{\text{step}}^{-1}$. Additionally, for any resistances $\hat{\mathbf{r}}$ that \mathcal{L} had at any point since the last call to $\mathcal{L}.\text{BATCHUPDATE}$, $\hat{\mathbf{r}}$ are $(\hat{\mu}, 1.1)$ -central for some $\hat{\mu} \in [\mu, \mu^0]$.

Then, there is an algorithm that with high probability computes $\mathbf{f}(\mu')$, where $\mu' = \mu/(1 + \varepsilon_{\text{step}}/\sqrt{m})^{k\varepsilon_{\text{step}}^{-1}}$. The algorithm runs in time $\tilde{O}(m)$, plus $O(k^{16})$ calls to $\mathcal{L}.\text{UPDATE}$, $O(k^4)$ calls to $\mathcal{L}.\text{SOLVE}$, and $O(k^{16})$ calls to $\mathcal{C}.\text{UPDATE}$ and $\mathcal{C}.\text{CHECK}$. Additionally, $\mathcal{L}.\mathbf{r}$ and \mathcal{C} are unmodified.

Algorithm 2 MultiStep

```

1: procedure MULTISTEP( $\mathbf{f}, \mu$ )    ▷ Makes equivalent progress to  $k$  interior point method steps
2:    $\hat{\mathbf{r}} = \mathcal{L}.\mathbf{r}$                                 ▷ Save resistances to restore later
3:   for  $i = 1, \dots, k\varepsilon_{\text{step}}^{-1}$  do
4:      $Z = \mathcal{L}.\text{SOLVE}()$ 
5:     for  $e \in Z$  do                                ▷  $Z$ : Set of edges with sufficiently changed flow
6:        $\tilde{\mathbf{f}}_e = \mathcal{C}^i.\text{CHECK}(e)$ 
7:       if  $\tilde{\mathbf{f}}_e \neq 0$  then
8:          $\mathbf{f}_e = \mathbf{f}_e + \varepsilon_{\text{step}}\tilde{\mathbf{f}}_e$ 
9:          $\mathcal{L}.\text{UPDATE}(e, \mathbf{f})$ 
10:         $\mathcal{C}^j.\text{TEMPORARYUPDATE}(e, \mathbf{f})$  for  $j \in [i + 1, k\varepsilon_{\text{step}}^{-1}]$ 
11:    $\mu = \mu/(1 + \varepsilon_{\text{step}}/\sqrt{m})^{k\varepsilon_{\text{step}}^{-1}}$ 
12:    $\mathbf{f} = \text{RECENTER}(\mathbf{f}, \mu)$                                 ▷ Lemma 4.3.6
13:   for  $e \in E$  do
14:     if  $\mathcal{L}.r_e \neq \hat{\mathbf{r}}_e$  then
15:        $\mathcal{L}.\text{UPDATE}(e, \hat{\mathbf{r}})$                                 ▷ Return LOCATOR resistances to their original state
16:   Call  $\mathcal{C}^i.\text{ROLLBACK}()$  to undo all TEMPORARYUPDATES for all  $\mathcal{C}^i$ 
17:   return  $\mathbf{f}, \mu$ 

```

4.3.5 Proof of Theorem 4.1.1

Correctness. First of all, we apply capacity and cost scaling [71] to make sure that $\|\mathbf{c}\|_\infty, \|\mathbf{u}\|_\infty = \text{poly}(m)$. These incur an extra factor of $\log(U + W)$ in the runtime.

We first get an initial solution to the interior point method by using the following lemma:

Lemma 4.3.12 (Interior point method initialization, Appendix A in [10]). *Given a min cost flow instance $\mathcal{I} = (G(V, E), \mathbf{c}, \mathbf{d}, \mathbf{u})$, there exists an algorithm that runs in time $O(m)$ and produces a new min cost flow instance $\mathcal{I}' = (G'(V', E'), \mathbf{c}', \mathbf{d}', \mathbf{u}')$, where $|V'| = O(|V|)$ and $|E'| = O(|E|)$, as well as a flow \mathbf{f} such that*

- \mathbf{f} is μ -central for \mathcal{I}' for some $\mu = \Theta(\|\mathbf{c}\|_2)$
- Given an optimal solution for \mathcal{I}' , an optimal minimum cost flow solution for \mathcal{I} can be computed in $O(m)$

Therefore we now have a $\text{poly}(m)$ -central solution for an instance \mathcal{I} . We can now apply Lemma 4.3.10 to get a μ' -central solution with $\mu' \leq m^{-10}$. Then we can apply the following lemma to round the solution, which follows from Lemma 5.4 in [10].

Lemma 4.3.13 (Interior point method rounding). *Given a min cost flow instance \mathcal{I} and a μ -central flow \mathbf{f} for $\mu \leq m^{-10}$, there is an algorithm that runs in time $\tilde{O}(m)$ and returns an optimal integral flow.*

By Lemma 4.3.12, this solution can be turned into an exact solution for the original instance. As Lemma 4.3.10 succeeds with high probability, the whole algorithm does too.

Runtime. To determine the final runtime, we analyze each operation in Algorithm 1 separately.

The INITIALIZE (Lemma 4.3.12) and ROUND (Lemma 4.3.13) operations take time $\tilde{O}(m)$. Now, the runtime of Lemma 4.3.10 is $\tilde{O}(m^{3/2}/k)$ plus the runtime incurred because of calls to the locator \mathcal{L} . We will use the runtimes per operation from Lemma 4.3.8.

\mathcal{L} .SOLVE: This operation is run $\tilde{O}(m^{1/2}k^3)$ times, and each of these costs $\tilde{O}\left(\frac{\beta m}{\varepsilon^2}\right) = \tilde{O}(mk^{12}\beta)$. Therefore in total $\tilde{O}(m^{3/2}k^{15}\beta)$.

We pick β by $m^{3/2}k^{15}\beta \leq m^{3/2}/k$ as $\beta = k^{-16}$, so the runtime is $\tilde{O}(m^{3/2}/k)$. Note that this satisfies the constraint $\beta \geq \tilde{\Omega}(k^3/m^{1/4})$ as long as $k \leq \tilde{\Omega}(m^{1/76})$.

\mathcal{L} .BATCHUPDATE: This is run $\tilde{O}(m^{1/2}/\alpha^{1/4})$ times with empty arguments, each of which takes time $\tilde{O}(m/\varepsilon^2) = \tilde{O}(mk^{12})$. The total runtime because of these is $\tilde{O}(m^{3/2}k^{12}\alpha^{-1/4})$. As we need this to be below $\tilde{O}(m^{3/2}/k)$, we set $\alpha = k^{52}$.

This operation is also run $\tilde{O}(m^{1/2}k^{-1}\hat{T}^{-1})$ times with some non-empty argument Z , each of which takes time $\tilde{O}(m/\varepsilon^2 + |Z|/(\varepsilon^2\beta^2)) = \tilde{O}(mk^{12} + k^{44}|Z|)$. As by Lemma 4.3.10 the total sum of $|Z|$ over all calls is $\tilde{O}(mk^3\beta^{1/2}) = \tilde{O}(mk^{-5})$, we get a runtime of

$$\tilde{O}\left(m^{1/2}k^{-1}\hat{T}^{-1} \cdot mk^{12} + k^{44} \cdot mk^{-5}\right) = \tilde{O}\left(m^{3/2}k^{11}\hat{T}^{-1} + mk^{39}\right).$$

In order to set the first term to be at most $\tilde{O}(m^{3/2}/k)$, we set $\hat{T} = k^{12}$.

Therefore the total runtime of this operation is $\tilde{O}(m^{3/2}/k + mk^{39})$.

L.UPDATE: This is run $\tilde{O}\left(m^{1/2}\left(k^{6\hat{T}} + k^{15}\right)\right) = \tilde{O}\left(m^{1/2}k^{18}\right)$ times and the amortized cost per operation is

$$\begin{aligned} & \tilde{O}\left(m \cdot \frac{\hat{\varepsilon}\alpha^{1/2}}{\varepsilon^3} + \hat{\varepsilon}^{-4}\varepsilon^{-2}\beta^{-8} + \hat{\varepsilon}^{-2}\varepsilon^{-4}\alpha^2\beta^{-6}\right) \\ &= \tilde{O}\left(m \cdot k^{44}\hat{\varepsilon} + k^{140}\hat{\varepsilon}^{-4} + k^{224}\hat{\varepsilon}^{-2}\right), \end{aligned}$$

so in total

$$m^{3/2}k^{62}\hat{\varepsilon} + m^{1/2}k^{158}\hat{\varepsilon}^{-4} + m^{1/2}k^{242}\hat{\varepsilon}^{-2}.$$

As we need the first term to be $\tilde{O}\left(m^{3/2}/k\right)$, we set $\hat{\varepsilon} = k^{-63}$. Therefore the total runtime is

$$\tilde{O}\left(m^{3/2}/k + m^{1/2}k^{410} + m^{1/2}k^{368}\right) = \tilde{O}\left(m^{3/2}/k + m^{1/2}k^{410}\right).$$

L.INITIALIZE: This is run $\tilde{O}\left(k^3\beta^{-1/2}\right) = k^{11}$ times in total, and the runtime for each run is

$$\tilde{O}\left(m \cdot \left(\hat{\varepsilon}^{-4}\beta^{-8} + \hat{\varepsilon}^{-2}\varepsilon^{-2}\alpha^2\beta^{-4}\right)\right) = \tilde{O}\left(m \cdot \left(k^{380} + k^{306}\right)\right) = \tilde{O}\left(m \cdot k^{380}\right),$$

so in total $\tilde{O}\left(mk^{380}\right)$.

Therefore, for the whole algorithm, we get $\tilde{O}\left(m^{3/2}/k + m^{1/2}k^{410} + mk^{380}\right)$ which after balancing gives $k = m^{1/762}$.

4.4 An Efficient $(\alpha, \beta, \varepsilon)$ -LOCATOR

In this section we will show how to implement an $(\alpha, \beta, \varepsilon)$ -LOCATOR, as defined in Definition 4.3.7. In order to maintain the approximate electrical flow \mathbf{f} required by Lemma 4.3.8 we will keep a vertex sparsifier in the form of a sparsified Schur complement onto some vertex set C . As in [73], we choose C to be a *congestion reduction subset*.

Definition 4.4.1 (Congestion reduction subset [73]). *Given a graph $G(V, E)$ with resistances \mathbf{r} and any parameter $\beta \in (0, 1)$, a vertex subset $C \subseteq V$ is called a β -congestion reduction subset (or just congestion reduction subset) if:*

- $|C| \leq O(\beta m)$
- For any $u \in V$, a random walk starting from u that visits $\tilde{\Omega}(\beta^{-1} \log n)$ distinct vertices hits C with high probability
- If we generate $\deg(u)$ random walks from each $u \in V \setminus C$, the expected number of these that hit some fixed $v \in V \setminus C$ before C is $\tilde{O}(1/\beta^2)$. Concretely:

$$\sum_{u \in V} \deg(u) \cdot p_v^{C \cup \{v\}}(u) \leq \tilde{O}(1/\beta^2). \quad (4.8)$$

The following lemma shows that such a vertex subset can be constructed efficiently:

Lemma 4.4.2 (Construction of congestion reduction subset [73]). *Given a graph $G(V, E)$ with resistances \mathbf{r} and a parameter $\beta \in (0, 1)$, there is an algorithm that generates a β -congestion reduction subset in time $\tilde{O}(m/\beta^2)$.*

Intuitively, (4.8) says that “not too many” random walks go through a given vertex before reaching C . This property is crucial for ensuring that when inserting a new vertex into C , the data structure will not have to change too much. As we will see in Section 4.4.1, this property plays an even more central role when general demands are introduced, as it allows us to show that the demands outside C can be pushed to C . Additionally, in Section 4.4.2 we will use it to show that edges that are too far from C in effective resistance metric are not *important*, in the sense that neither can they get congested, nor can their demand congest anything else.

4.4.1 Moving demands to the sparsifier

The goal of this section is to show that if C is a congestion reduction subset, then any demand of the form $\mathbf{d} = \mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}}$ for some $\mathbf{q} \in [-1, 1]^m$ can be approximated by $\pi^C(\mathbf{d})$, i.e. its demand projection onto C (Definition 4.2.4). This allows us to move all demands to the sublinear-sized C and thus enables us to work with the Schur complement of G onto C .

Lemma 4.4.3. *Consider a graph $G(V, E)$ with resistances \mathbf{r} and Laplacian \mathbf{L} , a β -congestion reduction subset C , and a demand $\mathbf{d} = \delta \mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}}$ for some $\delta > 0$ and $\mathbf{q} \in [-1, 1]^m$. Then, the potential embedding defined as*

$$\phi = \mathbf{L}^+ (\mathbf{d} - \pi^C(\mathbf{d}))$$

has congestion $\delta \cdot \tilde{O}(1/\beta^2)$, i.e. $\left\| \frac{\mathbf{B}\phi}{\sqrt{\mathbf{r}}} \right\|_\infty \leq \delta \cdot \tilde{O}(1/\beta^2)$.

We first prove a restricted version of the lemma where \mathbf{d} is an $s - t$ demand. Then, Lemma 4.4.3 follows trivially by applying (4.8).

Lemma 4.4.4. *Consider a graph $G(V, E)$ with resistances \mathbf{r} and Laplacian \mathbf{L} , a β -congestion reduction subset C , and a demand $\mathbf{d} = \delta \mathbf{B}^\top \frac{\mathbf{1}_{st}}{\sqrt{\mathbf{r}}}$ for some $\delta > 0$ and $(s, t) \in E \setminus E(C)$. Then, for the potential embedding defined as*

$$\phi = \mathbf{L}^+ (\mathbf{d} - \pi^C(\mathbf{d}))$$

it follows that for any $e = (u, v) \in E$ we have

$$\left| \frac{(\mathbf{B}\phi)_e}{\sqrt{r_e}} \right| \leq 2\delta \cdot \left(p_u^{C \cup \{u\}}(s) + p_v^{C \cup \{v\}}(s) + p_u^{C \cup \{u\}}(t) + p_v^{C \cup \{v\}}(t) \right).$$

The proof of Lemma 4.4.4 appears in Appendix 9.2.4.

4.4.2 ε -Important edges

In this section we will show that the effect of edges that are “far” from the congestion reduction subset C is negligible, as both their congestion and the congestion incurred because of their demands are small. More specifically, given a demand \mathbf{d} supported on C with energy ≤ 1 , i.e. $\mathcal{E}_r(\mathbf{d}) \leq 1$, the

congestion $\rho = \mathbf{R}^{-1/2} \mathbf{B} \mathbf{L}^+ \mathbf{d}_C$ that it induces satisfies:

$$\begin{aligned}
|\rho_e| &= \left| \left\langle \mathbf{1}_e, \mathbf{R}^{-1/2} \mathbf{B} \mathbf{L}^+ \mathbf{d} \right\rangle \right| \\
&= \left| \left\langle \mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{\mathbf{r}}}, \mathbf{L}^+ \mathbf{d} \right\rangle \right| \\
&= \left| \left\langle \pi^C \left(\mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{\mathbf{r}}} \right), \mathbf{S} \mathbf{C}^+ \mathbf{d}_C \right\rangle \right| \\
&\leq \sqrt{\mathcal{E}_r \left(\pi^C \left(\mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{\mathbf{r}}} \right) \right)} \mathcal{E}_r(\mathbf{d}) \\
&\leq \sqrt{\mathcal{E}_r \left(\pi^C \left(\mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{\mathbf{r}}} \right) \right)}.
\end{aligned}$$

For the last equality we used Fact 4.2.5, for the first inequality we applied Cauchy-Schwarz, and for the second one we used the upper bound on the energy required to route \mathbf{d} . Therefore, if we bound the energy of the projection of $\mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{\mathbf{r}}}$ onto C , we can also bound the congestion of e . This is done in the following lemma, whose proof appears in Appendix 9.2.4.

Lemma 4.4.5. *Consider a graph $G(V, E)$ with resistances \mathbf{r} and $C \subseteq V$. Then, for all $e \in E \setminus E(C)$ we have*

$$\sqrt{\mathcal{E}_r \left(\pi^C \left(\mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{r_e}} \right) \right)} \leq 6 \cdot \sqrt{\frac{r_e}{R_{\text{eff}}(C, e)}}.$$

This is the consequence of the following lemma, which bounds the magnitude of the projection on a specific vertex, based on its effective resistance distance from e , as well as hitting probabilities from e to C . The proof appears in Appendix 9.2.5.

Lemma 4.4.6. *Consider a graph $G(V, E)$ with resistances \mathbf{r} and a subset of vertices $C \subseteq V$. For any vertex $v \in V \setminus C$ we have that*

$$\left| \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{\mathbf{r}}} \right) \right| \leq (p_v^{C \cup \{v\}}(u) + p_v^{C \cup \{v\}}(w)) \cdot \frac{\sqrt{r_e}}{R_{\text{eff}}(v, e)}.$$

This lemma is complementary to the more immediate property

$$\left| \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{\mathbf{r}}} \right) \right| \leq (p_v^{C \cup \{v\}}(u) + p_v^{C \cup \{v\}}(w)) \cdot \frac{1}{\sqrt{r_e}},$$

and they are both used in Section 4.5 in order to estimate demand projections. In fact, the just by multiplying these two, we get the following lemma, which is nice because it doesn't depend on r_e :

Lemma 4.4.7. *Consider a graph $G(V, E)$ with resistances \mathbf{r} and a subset of vertices $C \subseteq V$. For any vertex $v \in V \setminus C$ we have that*

$$\left| \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{\mathbf{r}}} \right) \right| \leq (p_v^{C \cup \{v\}}(u) + p_v^{C \cup \{v\}}(w)) \cdot \frac{1}{\sqrt{R_{\text{eff}}(v, e)}}.$$

By the previous discussion, Lemma 4.4.6 implies that if $|\rho_e| \geq \varepsilon$, then $R_{eff}(C, e) \leq r_e \cdot \frac{36}{\varepsilon^2}$. This motivates the following definition of ε -important edges.

Definition 4.4.8 (Important edges). *An edge $e \in E$ is called ε -important (or just important) if $R_{eff}(C, e) \leq r_e/\varepsilon^2$.*

Now it is time for the main lemma of this section, which uses Lemma 4.4.5 to show that if our goal is to detect edges with congestion $\geq \varepsilon$, it is sufficient to restrict to computing demand projections of $\Omega(\varepsilon)$ -important edges. Its proof appears in Appendix 9.2.4.

Lemma 4.4.9 (Localization lemma). *Let ϕ^* be any solution of*

$$\mathbf{L}\phi^* = \delta \cdot \pi^C \left(\mathbf{B}^\top \frac{\mathbf{p}}{\sqrt{\mathbf{r}}} \right),$$

where \mathbf{r} are any resistances, $\mathbf{p} \in [-1, 1]^m$, and $C \subseteq V$. Then, for any $e \in E$ that is not ε -important we have $\left| \frac{\mathbf{B}\phi^*}{\sqrt{\mathbf{r}}} \right|_e \leq 6\varepsilon$.

4.4.3 Proving Lemma 4.3.8

Before moving to the description of how LOCATOR works and its proof, we will provide a lemma which bounds how fast a demand projection changes.

We will use the following observation, which states that if our congestion reduction subset C contains an βm -sized uniformly random edge subset, then with high probability, effective resistance neighborhoods that are disjoint from C only have $\tilde{O}(\beta^{-1})$ edges. Note that this will be true throughout the algorithm as long as the resistances do not depend on the randomness of C . This is true, as resistance updates are only ever given as inputs to LOCATOR.

Lemma 4.4.10 (Few edges in a small neighborhood). *Let $\beta \in (0, 1)$ be a parameter and C be a vertex set which contains a subset of βm edges sampled at random. Then with high probability, for any $v \in V \setminus C$ we have that $|N_E(v, R_{eff}(C, v)/2)| \leq 10\beta^{-1} \ln m$, where*

$$N_E(v, R) := \{e \in E \mid R_{eff}(e, v) \leq R\}.$$

Proof. Suppose that for some vertex $v \in V \setminus C$, $|N_E(v, R_{eff}(C, v)/2)| \geq 10\beta^{-1} \ln m$. Since by construction C contains a random edge subset of size βm , with high probability $N_E(v, R_{eff}(C, v)/2) \cap C \neq \emptyset$, so there exists $u \in C$ such that $R_{eff}(u, v) \leq R_{eff}(C, v)/2$. This is a contradiction since $u \in C$ implies $R_{eff}(C, v) \leq R_{eff}(u, v)$. Union bounding over all v yields the claim. \square

Using this fact, we can now show that the change of the demand projection (measured in energy) is quite mild. The proof of the following lemma can be found in Appendix 9.2.4.

Lemma 4.4.11 (Projection change). *Consider a graph $G(V, E)$ with resistances \mathbf{r} , $\mathbf{q} \in [-1, 1]^m$, and a β -congestion reduction subset C . Then, with high probability,*

$$\sqrt{\mathcal{E}_{\mathbf{r}} \left(\pi^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}} \right) - \pi^C \left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}} \right) \right)} \leq \tilde{O}(\beta^{-2}).$$

The above lemma can be applied over multiple vertex insertions and resistance changes, to bound the overall energy change. This is shown in the following lemma, which is proved in Appendix 9.2.4:

Lemma 4.4.12. Consider a graph $G(V, E)$ with resistances $\mathbf{r}^0, \mathbf{q}^0 \in [-1, 1]^m$, a β -congestion reduction subset C^0 , and a fixed sequence of updates, where the i -th update $i \in \{0, \dots, T-1\}$ is of the following form:

- **ADDTERMINAL**(v^i): Set $C^{i+1} = C^i \cup \{v^i\}$ for some $v^i \in V \setminus C^i$, $q_e^{i+1} = q_e^i, r_e^{i+1} = r_e^i$
- **UPDATE**($e^i, \mathbf{q}, \mathbf{r}$): Set $C^{i+1} = C^i$, $q_e^{i+1} = q_e, r_e^{i+1} = r_e$, where $e^i \in E(C^i)$

Then, with high probability,

$$\sqrt{\mathcal{E}_{r^T} \left(\pi^{C^0, r^0} \left(\mathbf{B}^\top \frac{\mathbf{q}_S^0}{\sqrt{\mathbf{r}^0}} \right) - \pi^{C^T, r^T} \left(\mathbf{B}^\top \frac{\mathbf{q}_S^T}{\sqrt{\mathbf{r}^T}} \right) \right)} \leq \tilde{O} \left(\max_{i \in \{0, \dots, T-1\}} \left\| \frac{\mathbf{r}^T}{\mathbf{r}^i} \right\|_\infty^{1/2} \beta^{-2} \right) \cdot T.$$

We are now ready to describe the LOCATOR data structure. We will give an outline here, and defer the full proof to Appendix 9.2.4. The goal of an $(\alpha, \beta, \varepsilon)$ -LOCATOR is, given some flow \mathbf{f} with slacks \mathbf{s} and resistances \mathbf{r} , to compute all $e \in E$ such that $\sqrt{r_e} \left| \tilde{f}_e^* \right| \geq \varepsilon$, where

$$\tilde{\mathbf{f}}^* = \delta g(\mathbf{s}) - \delta \mathbf{R}^{-1} \mathbf{B} \mathbf{L}^+ \mathbf{B}^\top g(\mathbf{s})$$

($\mathbf{L} = \mathbf{B}^\top \mathbf{R}^{-1} \mathbf{B}$), where $\delta = 1/\sqrt{m}$.

If we set $\rho_e^* = \sqrt{r_e} \tilde{f}_e^*$, we can equivalently write

$$\boldsymbol{\rho}^* = \delta \sqrt{\mathbf{r}} g(\mathbf{s}) - \delta \mathbf{R}^{-1/2} \mathbf{B} \mathbf{L}^+ \mathbf{B}^\top g(\mathbf{s}),$$

and require to find all the entries of $\boldsymbol{\rho}^*$ with magnitude at least ε . As $\delta \left\| \sqrt{\mathbf{r}} g(\mathbf{s}) \right\|_\infty \leq \delta \leq \varepsilon/100$, we can concentrate on the second term, and denote

$$\boldsymbol{\rho}'^* = \delta \mathbf{R}^{-1/2} \mathbf{B} \mathbf{L}^+ \mathbf{B}^\top g(\mathbf{s})$$

for convenience.

First, we use Lemma 4.4.3 to show that

$$\delta \left\| \mathbf{R}^{-1/2} \mathbf{B} \mathbf{L}^+ (g(\mathbf{s}) - \pi^C(g(\mathbf{s}))) \right\|_\infty \leq \delta \cdot \tilde{O}(\beta^{-2}) \leq \varepsilon/100.$$

Now, let's set $\boldsymbol{\pi}_{old} = \pi^{C^0}(g(\mathbf{s}^0))$, where C^0 was the vertex set of the sparsifier and \mathbf{s}^0 the slacks after the last call to BATCHUPDATE. As we will be calling BATCHUPDATE at least every T calls to UPDATE for some $T \geq 1$, Lemma 4.4.12 implies that

$$\delta \left\| \mathbf{R}^{-1/2} \mathbf{B} \mathbf{L}^+ (\pi^C(g(\mathbf{s})) - \boldsymbol{\pi}_{old}) \right\|_\infty \leq \delta \cdot \tilde{O}(\alpha \beta^{-2}) T \leq \varepsilon/100,$$

as long as $T \leq \varepsilon \sqrt{m} / \tilde{O}(\alpha \beta^{-2})$.

Importantly, we will never be *removing* vertices from C , so $C^0 \subseteq C$. This implies that it suffices to find the large entries of

$$\delta \mathbf{R}^{-1/2} \mathbf{B} \mathbf{L}^+ \boldsymbol{\pi}_{old}.$$

Now, note that for any edge e that was *not* $\varepsilon/(100\alpha)$ -important for C^0 and corresponding

resistances \mathbf{r}^0 , we have

$$\begin{aligned}
& \delta \left| \mathbf{R}^{-1/2} \mathbf{B} \mathbf{L}^+ \boldsymbol{\pi}_{old} \right|_e \\
& \leq \delta \sqrt{\mathcal{E}_r(\boldsymbol{\pi}^{C^0}(\mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{\mathbf{r}}}))} \sqrt{\mathcal{E}_r(\boldsymbol{\pi}_{old})} \\
& \leq \delta \cdot \sqrt{\alpha} \frac{\varepsilon}{100\alpha} \cdot \sqrt{2\alpha m} \\
& = \varepsilon/50,
\end{aligned}$$

where we used Lemma 4.4.5 and the fact that $\mathcal{E}_r(\boldsymbol{\pi}^{C^0}(g(\mathbf{s}^0))) \leq 2\mathcal{E}_r(g(\mathbf{s}^0)) \leq 2\alpha m$. Therefore it suffices to approximate

$$\delta \mathbf{I}_S \mathbf{R}^{-1/2} \mathbf{B} \mathbf{L}^+ \boldsymbol{\pi}_{old},$$

where S was the set of $\frac{\varepsilon}{100\alpha}$ -important edges last computed during the last call to BATCHUPDATE.

Now, we will use the sketching lemma from (Lemma 5.1, [73] v2), which shows that in order to find all $\Omega(\varepsilon)$ large entries of this vector, it suffices to compute the inner products

$$\delta \left\langle \boldsymbol{\pi}^C \left(\mathbf{B}^\top \frac{\mathbf{q}_S^i}{\sqrt{\mathbf{r}}} \right), \mathbf{S} \mathbf{C}^+ \boldsymbol{\pi}_{old} \right\rangle$$

for $i \in [\tilde{O}(\varepsilon^{-2})]$ up to $O(\varepsilon)$ accuracy. Here $\mathbf{S} \mathbf{C}$ is the Schur complement onto C .

Based on this, there are two types of quantities that we will maintain:

- $\tilde{O}(1/\varepsilon^2)$ approximate demand projections $\tilde{\pi}^C \left(\mathbf{B}^\top \frac{\mathbf{q}_S^i}{\sqrt{\mathbf{r}}} \right)$, and
- an approximate Schur complement $\widetilde{\mathbf{S} \mathbf{C}}$ of G onto C .

For the latter, we will directly use the dynamic Schur complement data structure DYNAMICSC that was also used by [73] and is based on [51]. For completeness, we present this data structure in Appendix 9.2.1.

For the former, we will need $\tilde{O}(1/\varepsilon^2)$ data structures for maintaining demand projections onto C , under vertex insertions to C . The guarantees of each such a data structure, that we call an $(\alpha, \beta, \varepsilon)$ -DEMANDPROJECTOR, are as follows.

Definition 4.4.13 ($(\alpha, \beta, \varepsilon)$ -DEMANDPROJECTOR). *Let $\hat{\varepsilon} \in (0, \varepsilon)$ be a tradeoff parameter. Given a graph $G(V, E)$, resistances \mathbf{r} , and a vector $\mathbf{q} \in [-1, 1]^m$, an $(\alpha, \beta, \varepsilon)$ -DEMANDPROJECTOR is a data structure that maintains a vertex subset $C \subseteq V$ and an approximation to the demand projection $\boldsymbol{\pi}^C \left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}} \right)$, with high probability under oblivious adversaries. The following operations are supported:*

- **INITIALIZE**($C, \mathbf{r}, \mathbf{q}, S, \mathcal{P}$): *Initialize the data structure in order to maintain an approximation of $\boldsymbol{\pi}^C \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right)$, where $C \subseteq V$ is a β -congestion reduction subset, \mathbf{r} are resistances, $\mathbf{q} \in [-1, 1]^m$, and $S \subseteq E$ is a subset of γ -important edges. $\mathcal{P} = \{\mathcal{P}^{u,e,i} \mid u \in V, e \in E, u \in e, i \in [h]\}$ for some $h \in \mathbb{Z}_{\geq 1}$, is a set of independent random walks from u to C for any u .*
- **ADDTERMINAL**($v, \tilde{R}_{eff}(C, v)$): *Insert v into C . Also, $\tilde{R}_{eff}(C, v)$ is an estimate of $R_{eff}(C, v)$ such that $\tilde{R}_{eff}(C, v) \approx_2 R_{eff}(C, v)$. Returns an estimate*

$$\tilde{\pi}_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}} \right)$$

for the demand projection of \mathbf{q} onto $C \cup \{v\}$ at coordinate v such that

$$\left| \tilde{\pi}_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}} \right) - \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}} \right) \right| \leq \frac{\hat{\varepsilon}}{\sqrt{R_{eff}(v, C)}}.$$

- **UPDATE**($e, \mathbf{r}', \mathbf{q}'$): Set $r_e = r'_e$ and $q_e = q'_e$, where $e \in E(C)$, and $q'_e \in [-1, 1]$. Furthermore, r'_e satisfies the inequality $r_e^{\max}/\alpha \leq r'_e \leq \alpha \cdot r_e^{\min}$, where r_e^{\min} and r_e^{\max} represent the minimum, respectively the maximum values that the resistance of e has had since the last call to **INITIALIZE**.
- **OUTPUT**(\cdot): Output $\tilde{\pi}^C \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right)$ such that after $T \leq n^{O(1)}$ calls to **ADDTERMINAL**, for any fixed vector ϕ , $E_{\mathbf{r}}(\phi) \leq 1$, with high probability

$$\left| \left\langle \tilde{\pi}^C \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) - \pi^C \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right), \phi \right\rangle \right| \leq \hat{\varepsilon} \cdot \sqrt{\alpha} \cdot T.$$

We will implement such a data structure in Section 4.5, where we will prove the following lemma:

Lemma 4.4.14 (Demand projection data structure). *For any graph $G(V, E)$ and parameters $\hat{\varepsilon} \in (0, \varepsilon)$, $\beta \in (0, 1)$, there exists an $(\alpha, \beta, \varepsilon)$ -DEMANDPROJECTOR for G which, given access to $h = \tilde{\Theta}(\hat{\varepsilon}^{-4}\beta^{-6} + \hat{\varepsilon}^{-2}\beta^{-2}\gamma^{-2})$ precomputed independent random walks from u to C for each $e \in E$, $u \in e$, has the following runtimes per operation:*

- **INITIALIZE**: $\tilde{O}(m)$.
- **ADDTERMINAL**: $\tilde{O}(\hat{\varepsilon}^{-4}\beta^{-8} + \hat{\varepsilon}^{-2}\beta^{-6}\gamma^{-2})$.
- **UPDATE**: $O(1)$.
- **OUTPUT**: $O(\beta m + T)$, where T is the number of calls made to **ADDTERMINAL** after the last call to **INITIALIZE**.

Now we describe the way we will use the **DEMANDPROJECTORS** and **DYNAMICSC** to get an $(\alpha, \beta, \varepsilon)$ -LOCATOR \mathcal{L} .

Algorithm 3 LOCATOR \mathcal{L} .INITIALIZE

- 1: **procedure** \mathcal{L} .INITIALIZE(\mathbf{f})
 - 2: $\mathbf{s}^+ = \mathbf{u} - \mathbf{f}$, $\mathbf{s}^- = \mathbf{f}$, $\mathbf{r} = \frac{1}{(\mathbf{s}^+)^2} + \frac{1}{(\mathbf{s}^-)^2}$
 - 3: $\mathbf{Q} =$ Sketching matrix produced by (Lemma 5.1, [73] v2)
 - 4: DYNAMICSC = DYNAMICSC.INITIALIZE($\mathbf{G}, \emptyset, \mathbf{r}, \varepsilon, \beta$)
 - 5: $C =$ DYNAMICSC. C ▷ β -congestion reduction subset
 - 6: Estimate $\tilde{R}_{eff}(C, e) \approx_4 R_{eff}(C, e)$ using Lemma 9.2.3
 - 7: $S = \left\{ e \in E \mid \tilde{R}_{eff}(C, e) \leq r_e \cdot \left(\frac{100\alpha}{\varepsilon} \right)^2 \right\}$
 - 8: $h = \tilde{\Theta}(\hat{\varepsilon}^{-4}\beta^{-6} + \hat{\varepsilon}^{-2}\varepsilon^{-2}\alpha^2\beta^{-2})$
 - 9: Sample walks $\mathcal{P}^{u, e, i}$ from u to C for $e \in E \setminus E(C)$, $u \in e$, $i \in [h]$ (Lemma 5.15, [73] v2)
 - 10: $\text{DP}^i =$ DEMANDPROJECTOR.INITIALIZE($C, \mathbf{r}, \mathbf{q}^i, S, \mathcal{P}$) for all rows \mathbf{q}^i of \mathbf{Q}
 - 11: \mathcal{L} .BATCHUPDATE(\emptyset)
-

\mathcal{L} .INITIALIZE: Every time \mathcal{L} .INITIALIZE is called, we first generate a β -congestion reduction subset C based on Lemma 4.4.2 (takes time $\tilde{O}(m/\beta^2)$), then a sketching matrix \mathbf{Q} and its rows \mathbf{q}^i for

$i \in [\tilde{O}(1/\varepsilon^2)]$ as in (Lemma 5.1, [73] v2) (takes time $\tilde{O}(m/\varepsilon^2)$), and finally random walks $\mathcal{P}^{u,e,i}$ from u to C for each $u \in V$, $e \in E \setminus E(C)$ with $u \in e$, and $i \in [h]$, where $h = \tilde{O}(\hat{\varepsilon}^{-4}\beta^{-6} + \hat{\varepsilon}^{-2}\varepsilon^{-2}\alpha^2\beta^{-2})$ as in (Lemma 5.15, [73] v2) (takes time $\tilde{O}(h/\beta^2)$ for each (u, e)).

We also compute $\tilde{R}_{eff}(C, u) \approx_2 R_{eff}(C, u)$ for all $u \in V$ as described in Lemma 9.2.3 so that we can let S be a subset of $\varepsilon/(100\alpha)$ -important edges that contains all $\varepsilon/(200\alpha)$ -important edges. This takes time $\tilde{O}(m)$. Then, we call `DYNAMICSC.INITIALIZE`($G, C, \mathbf{r}, O(\varepsilon), \beta$) (from Appendix 9.2.1) to initialize the dynamic Schur complement onto C , with error tolerance $O(\varepsilon)$, which takes time $\tilde{O}\left(m \cdot \frac{1}{\varepsilon^4\beta^4}\right)$, as well as `DEMANDPROJECTOR.INITIALIZE`($C, \mathbf{r}, \mathbf{q}, S, \mathcal{P}$) for the $\tilde{O}(1/\varepsilon^2)$ `DEMANDPROJECTORS`, i.e. one for each $\mathbf{q} \in \{\mathbf{q}^i \mid i \in [\tilde{O}(1/\varepsilon^2)]\}$. Also, we compute

$$\boldsymbol{\pi}^{old} = \boldsymbol{\pi}^C \left(\mathbf{B}^\top g(\mathbf{s}) \right),$$

which takes $\tilde{O}(m)$ as in `DEMANDPROJECTOR.INITIALIZE`. All of this takes $\tilde{O}\left(m \cdot \left(\frac{1}{\varepsilon^4\beta^8} + \frac{\alpha^2}{\varepsilon^2\varepsilon^2\beta^4}\right)\right)$.

Algorithm 4 `LOCATOR` $\mathcal{L}.$ `UPDATE` and $\mathcal{L}.$ `BATCHUPDATE`

```

1: procedure UPDATE( $e = (u, w), \mathbf{f}$ )
2:    $s_e^+ = u_e - f_e, s_e^- = f_e, r_e = \frac{1}{(s_e^+)^2} + \frac{1}{(s_e^-)^2}$ 
3:    $\tilde{R}_{eff}(C, u) = \text{DYNAMICSC.ADDTERMINAL}(u)$ 
4:    $\tilde{R}_{eff}(C \cup \{u\}, w) = \text{DYNAMICSC.ADDTERMINAL}(w)$ 
5:    $C = C \cup \{u, w\}$ 
6:   for  $i = 1, \dots, \tilde{O}(1/\varepsilon^2)$  do
7:      $\text{DP}^i.\text{ADDTERMINAL}(u, \tilde{R}_{eff}(C, u))$ 
8:      $\text{DP}^i.\text{ADDTERMINAL}(w, \tilde{R}_{eff}(C \cup \{u\}, w))$ 
9:   DYNAMICSC.UPDATE( $e, r_e$ )
10:  for  $i = 1 \dots \tilde{O}(1/\varepsilon^2)$  do
11:     $\text{DP}^i.\text{UPDATE}(e, \mathbf{r}, \mathbf{q}^i)$ 
12: procedure BATCHUPDATE( $Z, \mathbf{f}$ )
13:   $\mathbf{s}^+ = \mathbf{u} - \mathbf{f}, \mathbf{s}^- = \mathbf{f}, \mathbf{r} = \frac{1}{(\mathbf{s}^+)^2} + \frac{1}{(\mathbf{s}^-)^2}$ 
14:  Estimate  $\tilde{R}_{eff}(C, e) \approx_4 R_{eff}(C, e)$  using Lemma 9.2.3
15:   $S = \left\{ e \in E \mid \tilde{R}_{eff}(C, e) \leq r_e \cdot \left(\frac{100\alpha}{\varepsilon}\right)^2 \right\}$   $\triangleright \frac{\varepsilon}{100\alpha}$ -important edges
16:  for  $e = (u, w) \in Z$  do
17:    DYNAMICSC.ADDTERMINAL( $u$ )
18:    DYNAMICSC.ADDTERMINAL( $w$ )
19:     $C = C \cup \{u, w\}$ 
20:    DYNAMICSC.UPDATE( $e, r_e$ )
21:  for  $i = [\tilde{O}(1/\varepsilon^2)]$  do
22:     $\text{DP}^i.\text{INITIALIZE}(C, \mathbf{r}, \mathbf{q}^i, S, \mathcal{P})$ 
23:   $\boldsymbol{\pi}_{old} = \frac{1}{\sqrt{m}} \cdot \boldsymbol{\pi}^C \left( \mathbf{B}^\top \frac{\frac{1}{\mathbf{s}^+} - \frac{1}{\mathbf{s}^-}}{\mathbf{r}} \right)$   $\triangleright$  Compute exactly using Laplacian solve

```

$\mathcal{L}.$ UPDATE: Now, whenever $\mathcal{L}.$ UPDATE is called on an edge e , either $e \in E(C)$ or $e \notin E(C)$. In the first case we simply call `UPDATE` on `DYNAMICSC` and all `DEMANDPROJECTORS`.

In the second case, we first call `DYNAMICSC.ADDTERMINAL` on one endpoint v of e . After

doing this we can also get an estimate $\widetilde{R}_{eff}(C, v) \approx_2 R_{eff}(C, v)$ by looking at the edges between C and v in the sparsified Schur complement. By the guarantees of the expander decomposition used inside DYNAMICSC [73], the number of expanders containing v , amortized over all calls to DYNAMICSC.ADDTERMINAL, is $O(\text{poly log}(n))$. As the sparsified Schur complement contains $\widetilde{O}(1/\varepsilon^2)$ neighbors of v from each expander, the amortized number of neighbors of v in the sparsified Schur complement is $\widetilde{O}(1/\varepsilon^2)$, and the amortized runtime to generate them (by random sampling) is $\widetilde{O}(1/\varepsilon^2)$.

Given the resistances r_1, \dots, r_l of these edges, setting $\widetilde{R}_{eff}(C, v) = \left(\sum_{i=1}^l r_i^{-1} \right)^{-1}$ we guarantee that $\widetilde{R}_{eff}(C, v) \approx_{1+O(\varepsilon)} R_{eff}(C, v)$, by the fact that DYNAMICSC maintains an $(1 + O(\varepsilon))$ -approximate sparsifier of the Schur complement. Then, we call ADDTERMINAL($v, \widetilde{R}_{eff}(C, v)$) on all DEMANDPROJECTORS.

After repeating the same process for the other endpoint of e , we finally call UPDATE on DYNAMICSC and all DEMANDPROJECTORS. This takes time $\widetilde{O}\left(\frac{1}{\varepsilon^2\beta^2}\right)$ because of the Schur complement and amortized $\widetilde{O}\left(m \cdot \frac{\widehat{\varepsilon}\alpha^{1/2}}{\varepsilon} + \frac{1}{\varepsilon^4\beta^8} + \frac{\alpha^2}{\varepsilon^2\varepsilon^2\beta^{-6}}\right)$ for each of the demand projectors, so the total amortized runtime is $\widetilde{O}\left(m \cdot \frac{\widehat{\varepsilon}\alpha^{1/2}}{\varepsilon^3} + \frac{1}{\varepsilon^4\varepsilon^2\beta^8} + \frac{\alpha^2}{\varepsilon^2\varepsilon^4\beta^6}\right)$.

\mathcal{L} .BATCHUPDATE: When \mathcal{L} .BATCHUPDATE is called on a set of edges Z , we add them one by one in the DYNAMICSC data structure following the same process as in \mathcal{L} .UPDATE. For the demand projectors, we first manually insert the endpoints of these edges into C and then re-initialize all DEMANDPROJECTORS, by calling INITIALIZE with a new subset S of $\frac{\varepsilon}{200\alpha}$ -important edges that contains all $\frac{\varepsilon}{100\alpha}$ -important edges. Such a set can be computed by estimating $R_{eff}(C, u)$ for all $u \in V \setminus C$ up to a constant factor and, by Lemma 9.2.3, takes time $\widetilde{O}(m)$. Also, we compute

$$\pi^{old} = \pi^C \left(\mathbf{B}^\top g(s) \right),$$

which takes $\widetilde{O}(m)$ as in DEMANDPROJECTOR.INITIALIZE. The total runtime of this is $\widetilde{O}(m/\varepsilon^2 + |Z|/(\beta^2\varepsilon^2))$.

Algorithm 5 LOCATOR \mathcal{L} .SOLVE

```

1: procedure SOLVE()
2:    $\widetilde{SC} = \text{DYNAMICSC}.\widetilde{SC}()$ 
3:    $\phi_{old} = \widetilde{SC}^+ \pi_{old}$ 
4:    $\mathbf{v} = \mathbf{0}$ 
5:   for  $i = 1, \dots, \widetilde{O}(1/\varepsilon^2)$  do
6:      $\widetilde{\pi}^i = \text{DP}^i.\text{OUTPUT}()$ 
7:      $v_i = \langle \widetilde{\pi}^i, \phi_{old} \rangle$ 
8:    $Z = \text{RECOVER}(\mathbf{v}, \varepsilon/100)$  ▷ Recovers all  $\varepsilon/2$ -congested edges (Lemma 5.1, [73] v2)
9:   return  $Z$ 

```

\mathcal{L} .SOLVE: When \mathcal{L} .SOLVE is called, we set $\widetilde{SC} = \text{DYNAMICSC}.\widetilde{SC}()$, call OUTPUT on all DEMANDPROJECTORS to obtain vectors $\widetilde{\pi}^i$ which are estimators for $\pi^C(\mathbf{B}^\top \frac{\mathbf{q}_s^i}{\sqrt{r}})$ in the sense of Definition 4.4.13. Then we compute $v_i = \langle \widetilde{\pi}^i, \widetilde{SC}^+ \pi_{old} \rangle$ where π_{old} is the demand projection that was computed exactly the last time BATCHUPDATE was called. These computed terms represent an approximation to the update in $(\mathbf{Q}\rho)_i$ between two consecutive calls of \mathcal{L} .SOLVE. As we will show in the appendix, $\langle \widetilde{\pi}^i, \widetilde{SC}^+ \pi_{old} \rangle$ is an ε -additive approximation of $\langle \pi^C(\mathbf{B}^\top \frac{\mathbf{q}_s^i}{\sqrt{r}}), \mathbf{L}^+ \pi^C(\mathbf{B}^\top g(s)) \rangle$ for

all $i \in \left[\tilde{O}(1/\varepsilon^2) \right]$. The key fact that makes this approximation feasible is that although updates to the demand projection are hard to approximate with few samples, when hitting them with the deterministic vector $\boldsymbol{\pi}_{old}$, the resulting inner products strongly concentrate. The runtime of this is $\tilde{O}(\beta m/\varepsilon^2)$.

Using these computed values with the ℓ_2 heavy hitter data structure (Lemma 5.1, [73] v2) we get all edges with congestion more than ε . The total runtime is $\tilde{O}(\beta m/\varepsilon^2)$.

4.5 The Demand Projection Data Structure

The main goal of this section is to construct an $(\alpha, \beta, \varepsilon)$ -DEMANDPROJECTOR, as defined in Definition 4.4.13, and thus prove Lemma 4.4.14. The most important operation that needs to be implemented in order to prove Lemma 4.4.14 is to maintain the demand projection after inserting a vertex $v \in V \setminus C$ to C . In order to do this, we use the following identity from [73]:

$$\boldsymbol{\pi}^{C \cup \{v\}}(\mathbf{d}) = \boldsymbol{\pi}^C(\mathbf{d}) + \pi_v^{C \cup \{v\}}(\mathbf{d}) \cdot (\mathbf{1}_v - \boldsymbol{\pi}^C(\mathbf{1}_v)), \quad (4.9)$$

where \mathbf{d} is any demand (in our case, we have $\mathbf{d} = \mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}}$ for some $S \subseteq E$). For this, we need to compute approximations to $\pi_v^{C \cup \{v\}}\left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}}\right)$ and $\boldsymbol{\pi}^C(\mathbf{1}_v)$.

In Section 4.5.1, we will show that if S is a subset of γ -important edges, we can efficiently estimate $\pi_v^{C \cup \{v\}}\left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}}\right)$ up to additive accuracy $\frac{\hat{\varepsilon}}{R_{eff}(C, v)}$ by sampling random walks to C starting only from edges with relatively high resistance. For the remaining edges, the γ -importance property will imply that we are not losing much by ignoring them.

Then, in Section 4.5.2 we will show how to approximate $\mathbf{1}_v - \boldsymbol{\pi}^C(\mathbf{1}_v)$. This is equivalent to estimating the hitting probabilities from v to C . The guarantee that we would ideally like to get is on the error to route

$$\mathcal{E}_r(\tilde{\boldsymbol{\pi}}^C(\mathbf{1}_v) - \boldsymbol{\pi}^C(\mathbf{1}_v)) \leq \hat{\varepsilon}^2 R_{eff}(C, v). \quad (4.10)$$

Note that this is not possible to do efficiently for general C . For example, suppose that the hitting distribution is uniform. In this case, $\Omega(|C|)$ random walks are required to get a bound similar to (4.10). However, it might still be possible to guarantee it by using the structure of C , and this would simplify some parts of our analysis. Instead, we are going to work with the following weaker approximation bound: For any fixed potential vector $\boldsymbol{\phi} \in \mathbb{R}^n$ with $E_r(\boldsymbol{\phi}) \leq 1$, we have w.h.p.

$$|\langle \tilde{\boldsymbol{\pi}}^C(\mathbf{1}_v) - \boldsymbol{\pi}^C(\mathbf{1}_v), \boldsymbol{\phi} \rangle| \leq \hat{\varepsilon} \sqrt{R_{eff}(C, v)}. \quad (4.11)$$

Now, using these estimation lemmas, we will bound how our demand projection degrades when inserting a new vertex into C . This is stated in the following lemma and proved in Appendix 9.2.5.

Lemma 4.5.1 (Inserting a new vertex to C). *Consider a graph $G(V, E)$ with resistances \mathbf{r} , $\mathbf{q} \in [-1, 1]^m$, a β -congestion reduction subset C , and $v \in V \setminus C$. We also suppose that we have an estimate of the $C-v$ effective resistance such that $\tilde{R}_{eff}(C, v) \approx_2 R_{eff}(C, v)$, as well as to independent random walks $\mathcal{P}^{u, e, i}$ for each $u \in V \setminus C$, $e \in E \setminus E(C)$ with $u \in e$, $i \in [h]$, where each random walk starts from u and ends at C .*

If we let S be a subset of γ -important edges for $\gamma > 0$, then for any error parameter $\hat{\varepsilon} > 0$ we

can compute $\tilde{\pi}_v^{\text{CU}\{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) \in \mathbb{R}$ and $\tilde{\pi}^{\text{CU}\{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) \in \mathbb{R} \in \mathbb{R}^n$ such that with high probability

$$\left| \tilde{\pi}_v^{\text{CU}\{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) - \tilde{\pi}^{\text{CU}\{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) \right| \leq \frac{\hat{\varepsilon}}{\sqrt{R_{\text{eff}}(C, v)}},$$

as long as $h = \tilde{\Omega} (\hat{\varepsilon}^{-4} \beta^{-6} + \hat{\varepsilon}^{-2} \beta^{-2} \gamma^{-2})$. Furthermore, for any fixed ϕ , $E_{\mathbf{r}}(\phi) \leq 1$, after T insertions after the last call to INITIALIZE, with high probability

$$\left| \left\langle \tilde{\pi}^{\text{CU}\{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) - \pi^{\text{CU}\{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right), \phi \right\rangle \right| \leq \hat{\varepsilon} T,$$

as long as $h = \tilde{\Omega} (\hat{\varepsilon}^{-2} \beta^{-4} \gamma^{-2})$.

Algorithm 6 DEMANDPROJECTOR DP.ADDTERMINAL

```

1: procedure DP.ADDTERMINAL( $v, \tilde{R}_{\text{eff}}(C, v)$ )
2:   if  $v \in C$  then
3:     return
4:    $t = t + 1$ 
5:    $\tilde{\pi}_v^{\text{CU}\{v\}} \left( \mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) = 0$ 
6:   for  $u, e \in S, i$  such that  $\mathcal{P}^{u,e,i} \ni v$  and  $\tilde{R}_{\text{eff}}(C, v) \leq \frac{1}{(\min\{\hat{\varepsilon}/\tilde{O}(\beta^{-2}), \gamma/4\})^2 r_e}$  do
7:     if  $e = (u, *)$  then
8:        $\tilde{\pi}_v^{\text{CU}\{v\}} \left( \mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) = \tilde{\pi}_v^{\text{CU}\{v\}} \left( \mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) + \frac{1}{h} \frac{q_e}{\sqrt{r_e}}$ 
9:     else
10:       $\tilde{\pi}_v^{\text{CU}\{v\}} \left( \mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) = \tilde{\pi}_v^{\text{CU}\{v\}} \left( \mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) - \frac{1}{h} \frac{q_e}{\sqrt{r_e}}$ 
11:     Shortcut  $\mathcal{P}^{u,e,i}$  at  $v$ 
12:      $h' = \tilde{O} (\hat{\varepsilon}^{-2} \beta^{-4} \gamma^{-2})$ 
13:      $\tilde{\pi}^C(\mathbf{1}_v) = \mathbf{0}$ 
14:     for  $i = 1, \dots, h'$  do
15:       Run random walk from  $v$  to  $C$  with probabilities prop. to  $\mathbf{r}^{-1}$ , let  $u$  be the last vertex
16:        $\tilde{\pi}_u^C(\mathbf{1}_v) = \tilde{\pi}_u^C(\mathbf{1}_v) + \frac{1}{h'}$ 
17:        $\tilde{\pi}^{\text{CU}\{v\}} \left( \mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) = \tilde{\pi}^C \left( \mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) + \tilde{\pi}_v^{\text{CU}\{v\}} \left( \mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) \cdot (\mathbf{1}_v - \tilde{\pi}^C(\mathbf{1}_v))$ 
18:        $C = C \cup \{v\}, F = F \setminus \{v\}$ 

```

4.5.1 Estimating $\tilde{\pi}_v^{\text{CU}\{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right)$

There is a straightforward algorithm to estimate $\tilde{\pi}_v^{\text{CU}\{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right)$. For each edge $e = (u, w) \in E \setminus E(C)$, sample a number of random walks from u and w until they hit $C \cup \{v\}$. Then, add to the estimate $\frac{q_e}{\sqrt{r_e}}$ times the fraction of the random walks starting from u that contain v , minus $\frac{q_e}{\sqrt{r_e}}$ times the fraction of the random walks starting from w that contain v . [73] uses this sampling method together with the following concentration bound, to get a good estimate if the resistances of all congested edges are sufficiently large.

Lemma 4.5.2 (Concentration inequality 1 [73]). *Let $S = X_1 + \dots + X_n$ be the sum of n independent random variables. The range of X_i is $\{0, a_i\}$ for $a_i \in [-M, M]$. Let t, E be positive numbers such*

that $t \leq E$ and $\sum_{i=1}^n |\mathbb{E}[X_i]| \leq E$. Then

$$\Pr [|S - \mathbb{E}[S]| > t] \leq 2 \exp\left(-\frac{t^2}{6EM}\right).$$

Unfortunately, in our setting there is no reason to expect these resistances to be large, so the variance of this estimate might be too high. We have already introduced the concept of important edges in order to alleviate this problem, and proved that we only need to look at important edges. Even if all edges of which the demand projection is estimated are important (i.e. close to C), however, v can still be far from C . This is an issue, since we don't directly estimate projections onto C , but instead estimate the projection onto $C \cup \{v\}$ and then from v onto C .

Intuitively, however, if v is far from C , it should also be far from the set of important edges, so the insertion of v should not affect their demand projection too much. As the distance upper bound between an important edge and C is relative to the scale of the resistance of that edge, this statement needs to be more fine-grained in order to take the resistances of important edges into account.

More concretely, in the following lemma, which is proved in Appendix 9.2.5, we show that if we only compute demand projection estimates for edges e such that $r_e \geq c^2 R_{eff}(C, v)$ for some appropriately chosen $c > 0$, then we can guarantee a good bound on the number of random walks we need to sample.

For the remaining edges, we will show that the energy of their contributions to the projection is negligible, so that we can reach to our desired statement in Lemma 4.5.4.

Lemma 4.5.3. *Consider a graph $G(V, E)$ with resistances $\mathbf{r}, \mathbf{q} \in [-1, 1]^n$, a β -congestion reduction subset C , as well as $v \in V \setminus C$. If for some $c > 0$ we are given a set of edges*

$$S' \subseteq \left\{ e \in E \setminus E(C) \mid R_{eff}(C, v) \leq \frac{1}{c^2} r_e \right\},$$

then for any $\delta'_1 > 0$ we can compute $\tilde{\pi}_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_{S'}}{\sqrt{\mathbf{r}}} \right) \in \mathbb{R}$ such that with high probability

$$\left| \tilde{\pi}_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_{S'}}{\sqrt{\mathbf{r}}} \right) - \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_{S'}}{\sqrt{\mathbf{r}}} \right) \right| \leq \frac{\delta'_1}{\beta c \sqrt{R_{eff}(C, v)}}.$$

The algorithm requires access to $\tilde{O} \left(\delta_1'^{-2} \log n \log \frac{1}{\beta} \right)$ independent random walks from u to C for each $u \in V \setminus C$ and $e \in E \setminus E(C)$ with $u \in e$.

This leads us to the desired statement for this section, whose proof appears in Appendix 9.2.5.

Lemma 4.5.4 (Estimating $\pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}} \right)$). *Consider a graph $G(V, E)$ with resistances $\mathbf{r}, \mathbf{q} \in [-1, 1]^n$, a β -congestion reduction subset C , as well as $v \in V \setminus C$. If we are given a set S of γ -important edges for some $\gamma \in (0, 1)$ and an estimate $\tilde{R}_{eff}(C, v) \approx_2 R_{eff}(C, v)$, then for any $\delta_1 \in (0, 1)$ we can compute $\tilde{\pi}_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) \in \mathbb{R}$ such that with high probability*

$$\left| \tilde{\pi}_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) - \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) \right| \leq \frac{\delta_1}{\sqrt{R_{eff}(C, v)}}. \quad (4.12)$$

The algorithm requires $\tilde{O}(\delta_1^{-4}\beta^{-6} + \delta_1^{-2}\beta^{-2}\gamma^{-2})$ independent random walks from u to C for each $u \in V \setminus C$ and $e \in E \setminus E(C)$ with $u \in e$.

Additionally, we have

$$\left| \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) \right| \leq \frac{1}{\gamma \sqrt{R_{\text{eff}}(C, v)}} \cdot \tilde{O} \left(\frac{1}{\beta^2} \right).$$

4.5.2 Estimating $\pi^C(\mathbf{1}_v)$

In contrast to the quantity $\pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}} \right)$, where there are cancellations between its two components $\pi_v^{C \cup \{v\}} \left(\sum_{e=(u,w) \in E} \frac{q_e}{\sqrt{r_e}} \mathbf{1}_u \right)$ and $\pi_v^{C \cup \{v\}} \left(\sum_{e=(u,w) \in E} -\frac{q_e}{\sqrt{r_e}} \mathbf{1}_w \right)$ (as $\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}}$ sums up to $\mathbf{0}$), in $\pi^C(\mathbf{1}_v)$ there are no cancellations. The goal is to simply estimate the hitting probabilities from v to the vertices of C , which can be done by sampling a number of random walks from v to C .

As discussed before, even though ideally we would like to have an error bound of the form $\sqrt{\mathcal{E}_r(\tilde{\pi}^C(\mathbf{1}_v) - \pi^C(\mathbf{1}_v))} \leq \delta_2 \sqrt{R_{\text{eff}}(C, v)}$, our analysis is only able to guarantee that for any fixed potential vector ϕ with $E_r(\phi) \leq 1$, with high probability $|\langle \phi, \tilde{\pi}^C(\mathbf{1}_v) - \pi^C(\mathbf{1}_v) \rangle| \leq \delta_2 \sqrt{R_{\text{eff}}(C, v)}$. However, this is still sufficient for our purposes.

In Appendix 9.2.5 we prove the following general concentration inequality, which basically states that we can estimate the desired hitting probabilities as long as we have a bound on the ℓ_2 norm of the potentials ϕ weighted by the hitting probabilities.

Lemma 4.5.5 (Concentration inequality 2). *Let π be a probability distribution over $[n]$ and $\tilde{\pi}$ an empirical distribution of Z samples from π . For any $\bar{\phi} \in \mathbb{R}^n$ with $\|\bar{\phi}\|_{\pi, 2}^2 \leq \mathcal{V}$, we have*

$$\Pr \left[|\langle \tilde{\pi} - \pi, \bar{\phi} \rangle| > t \right] \leq \frac{1}{n^{100}} + \tilde{O}(\log(n \cdot \mathcal{V}/t)) \exp \left(-\frac{Zt^2}{\tilde{O}(\mathcal{V} \log^2 n)} \right).$$

We will apply it for $\bar{\phi} = \phi - \phi_v \cdot \mathbf{1}$, and it is important to note that $\mathcal{E}_r(\bar{\phi}) = \mathcal{E}_r(\phi)$. In order to get a bound on $\|\bar{\phi}\|_{\pi^C(\mathbf{1}_v), 2}^2$, we use the following lemma, which is proved in Appendix 9.2.5.

Lemma 4.5.6 (Bounding the second moment of potentials). *For any graph G , resistances \mathbf{r} , potentials ϕ with $E_r(\phi) \leq 1$, $C \subseteq V$ and $v \in V \setminus C$ we have $\|\phi - \phi_v \mathbf{1}\|_{\pi^C(\mathbf{1}_v), 2}^2 \leq 8 \cdot R_{\text{eff}}(C, v)$.*

To give some intuition on this, consider the case when $V = C \cup \{v\} = \{1, \dots, k\} \cup \{v\}$, and there are edges e_1, \dots, e_k between C and v , one for each vertex of C . Then, we have $\pi_i^C(\mathbf{1}_v) = (r_{e_i})^{-1} / \sum_{i=1}^k (r_{e_i})^{-1}$, and so

$$\|\bar{\phi}\|_{\pi^C(\mathbf{1}_v), 2}^2 = \sum_{i=1}^k \frac{(\phi_i - \phi_v)^2}{r_{e_i}} \cdot \left(\sum_{i=1}^k (r_{e_i})^{-1} \right)^{-1} \leq \mathcal{E}_r(\bar{\phi}) \cdot R_{\text{eff}}(C, v) \leq R_{\text{eff}}(C, v).$$

We finally arrive at the desired statement about estimating $\pi^C(\mathbf{1}_v)$.

Lemma 4.5.7 (Estimating $\pi^C(\mathbf{1}_v)$). *Consider a graph $G(V, E)$ with resistances \mathbf{r} , a β -congestion reduction subset C , as well as $v \in V \setminus C$. Then, for any $\delta_2 > 0$, we can compute $\tilde{\pi}^C(\mathbf{1}_v) \in \mathbb{R}^n$ such*

that with high probability

$$|\langle \phi, \tilde{\pi}^C(\mathbf{1}_v) - \pi^C(\mathbf{1}_v) \rangle| \leq \delta_2 \cdot \sqrt{R_{eff}(C, v)}, \quad (4.13)$$

where $\phi \in \mathbb{R}^n$ is a fixed vector with $E_r(\phi) \leq 1$. The algorithm computes $\tilde{O}\left(\frac{\log n}{\delta_2^2}\right)$ random walks from v to C .

Proof. Because both $\tilde{\pi}^C(\mathbf{1}_v)$ and $\pi^C(\mathbf{1}_v)$ are probability distributions, the quantity (4.13) doesn't change when a multiple of $\mathbf{1}$ is added to ϕ , and so we can replace it by $\bar{\phi} = \phi - \phi_v \mathbf{1}$.

Now, $\tilde{\pi}^C(\mathbf{1}_v)$ will be defined as the empirical hitting distribution that results from sampling Z random walks from v to C . Directly applying the concentration bound in Lemma 4.5.5 and setting $Z = \tilde{O}\left(\frac{\log n}{\delta_2^2}\right)$, together with the fact that $\|\bar{\phi}\|_{\pi^C(\mathbf{1}_v), 2}^2 \leq 8 \cdot R_{eff}(C, v)$ by Lemma 4.5.6 and $\log \log R_{eff}(C, v) \leq O(\log \log n)$, we get

$$\Pr \left[|\langle \tilde{\pi}^C(\mathbf{1}_v) - \pi^C(\mathbf{1}_v), \bar{\phi} \rangle| > \delta_2 \cdot \sqrt{R_{eff}(C, v)} \right] < \frac{1}{n^{10}}.$$

□

4.5.3 Proof of Lemma 4.4.14

We are now ready for the proof of Lemma 4.4.14.

Proof of Lemma 4.4.14. Let DP be a demand projection data structure. We analyze its operations one by one.

Algorithm 7 DEMANDPROJECTOR DP.INITIALIZE

- 1: **procedure** DP.INITIALIZE(C, r, q, S, \mathcal{P})
 - 2: Initialize C, r, q, S, \mathcal{P}
 - 3: $F = V \setminus C$
 - 4: $h = \tilde{O}(\hat{\varepsilon}^{-4} \beta^{-6} + \hat{\varepsilon}^{-2} \beta^{-4} \gamma^{-2})$ ▷ #random walks for each pair $u \in V, e \in E$ with $u \in e$
 - 5: $t = 0$ ▷ #calls to ADDTERMINAL since last call to UPDATEFULL
 - 6: $\phi = \mathbf{L}_{FF}^+ \left[\mathbf{B}^\top \frac{q_S}{\sqrt{r}} \right]_F$
 - 7: $\tilde{\pi}^C \left(\mathbf{B}^\top \frac{q_S}{\sqrt{r}} \right) = \left[\mathbf{B}^\top \frac{q_S}{\sqrt{r}} \right]_C - \mathbf{L}_{CF} \phi$
-

DP.INITIALIZE(C, r, q, S, \mathcal{P}): We initialize the values of C, r, q, S, \mathcal{P} . Then we exactly compute the demand projection, i.e. $\tilde{\pi}^C \left(\mathbf{B}^\top \frac{q_S}{\sqrt{r}} \right) = \pi^C \left(\mathbf{B}^\top \frac{q_S}{\sqrt{r}} \right)$, which takes time $\tilde{O}(m)$ as shown in [73].

More specifically, we have $\pi^C \left(\mathbf{B}^\top \frac{q_S}{\sqrt{r}} \right) = (\mathbf{I} \quad \mathbf{L}_{CF} \mathbf{L}_{FF}^{-1}) \mathbf{B}^\top \frac{q_S}{\sqrt{r}}$ which only requires applying the operators \mathbf{L}_{FF}^{-1} and \mathbf{L}_{CF} .

DP.ADDTERMINAL($v, \tilde{R}_{eff}(C, v)$): We will serve this operation by applying Lemma 4.5.1. It is important to note that the error guarantee for the OUTPUT procedure increases with every call to ADDTERMINAL, so in general we have a bounded budget for the number of calls to this procedure before having to call again INITIALIZE.

We apply Lemma 4.5.1 to obtain $\tilde{\pi}_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{q_S}{\sqrt{r}} \right)$, and update the estimate $\tilde{\pi}_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{q_S}{\sqrt{r}} \right)$. The former can be achieved with $h = \tilde{O}(\hat{\varepsilon}^{-4} \beta^{-6} + \hat{\varepsilon}^{-2} \beta^{-2} \gamma^{-2})$ random walks. Note that these random

walks are already stored in \mathcal{P} , so accessing each of them takes time $\tilde{O}(1)$. Using the congestion reduction property of C , we see that the running time of the procedure, which is dominated by shortcutting the random walks is $\tilde{O}(h\beta^{-2})$, which gives the claimed bound. The latter can be achieved with $h' = \tilde{O}(\hat{\varepsilon}^{-2}\beta^{-4}\gamma^{-2})$ fresh random walks. Due to the congestion reduction property, simulating each of these requires $\tilde{O}(\beta^{-2})$ time.

Algorithm 8 DEMANDPROJECTOR DP.UPDATE and DP.OUTPUT

```

1: procedure DP.UPDATE( $e, \mathbf{r}', \mathbf{q}'$ )
2:   if  $e \in S$  then
3:      $\tilde{\pi}^C \left( \mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) = \tilde{\pi}^C \left( \mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) + \left( \frac{q'_e}{\sqrt{r'_e}} - \frac{q_e}{\sqrt{r_e}} \right) \cdot \mathbf{B}^\top \mathbf{1}_e$ 
4:      $q_e = q'_e, r_e = r'_e$ 
5: procedure DP.OUTPUT()
6:   return  $\tilde{\pi}^C \left( \mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right)$ 

```

DP.UPDATE($e, \mathbf{r}', \mathbf{q}'$): We update the values of r_e, q_e . We also update the projection, by noting that since $e \in E(C)$,

$$\pi^C \left(\mathbf{B}^\top \frac{\mathbf{q}'}{\sqrt{\mathbf{r}'}} \right) = \pi^C \left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}} \right) + \left(\frac{q'_e}{\sqrt{r'_e}} - \frac{q_e}{\sqrt{r_e}} \right) \mathbf{B}^\top \mathbf{1}_e,$$

so we change $\tilde{\pi}^C \left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}} \right)$ by the same amount, which takes time $O(1)$ and does not introduce any additional error in our estimate.

DP.OUTPUT(\cdot): We output our estimate $\tilde{\pi}^C \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right)$. Per Lemma 4.5.1 we see that each of the previous T calls to ADDTERMINAL add an error to our estimate of at most $\hat{\varepsilon}$ in the sense that if Δ^t were the true change in the demand projection at the t^{th} insertion, and $\tilde{\Delta}^t$ were the update made to our estimate, then

$$\left| \left\langle \tilde{\Delta}^t - \Delta^t, \phi \right\rangle \right| \leq \hat{\varepsilon},$$

w.h.p. for any fixed ϕ such that $E_{\mathbf{r}^t}(\phi) \leq 1$, where \mathbf{r}^t represents the resistances when t^{th} call to ADDTERMINAL is made. Equivalently, for any nonzero ϕ ,

$$\frac{1}{\sqrt{E_{\mathbf{r}^t}(\phi)}} \left| \left\langle \tilde{\Delta}^t - \Delta^t, \phi \right\rangle \right| \leq \hat{\varepsilon},$$

By the invariant satisfied by the resistances passed as parameters to the ADDTERMINAL routine, we have that $\mathbf{r}^t \leq \alpha \cdot \mathbf{r}^T$ for all t . Therefore $1/E_{\mathbf{r}^t}(\phi) \leq \alpha/E_{\mathbf{r}^t}(\phi)$. So we have that

$$\frac{1}{\sqrt{E_{\mathbf{r}^T}(\phi)}} \left| \left\langle \tilde{\Delta}^t - \Delta^t, \phi \right\rangle \right| \leq \hat{\varepsilon} \cdot \sqrt{\alpha}.$$

Summing up over T insertions, we obtain the desired error bound. Furthermore, note that returning the estimate takes time proportional to $|C|$, which is $\tilde{O}(\beta m + T)$. □

Chapter 5

Decomposable Submodular Function Minimization via Maximum Flow

5.1 Introduction

A significant amount of work has been dedicated to the study of submodular functions. While this topic has garnered a lot of excitement from the theory community due to its the multiple connections to diverse algorithmic areas [116, 78], on the practical side minimizing submodular functions has been intensively used to model discrete problems in machine learning. MAP inference in Markov Random Fields [99], image segmentation [8, 147], clustering [126], corpus extraction problems [110] are just a few success stories of submodular minimization.

Polynomial time algorithms for this problem have been known ever since the 80's [78], and they have seen major running time improvements in more recent years [143, 84, 61, 132, 107, 35]. However, the massive scale of the problems that use submodular minimization nowadays drives the need for further developments.

One great advantage offered by the submodular functions that occur in practice is that they are *structured*. For example, in many common cases (hypergraph cuts [165], covering functions [158], MAP inference [60, 99, 166]) these can be decomposed into sums of simple submodular functions defined on small subsets. For these instances, prior work [90, 131, 56, 57, 94] has focused on providing efficient algorithms in the regime where the functions in the decomposition admit fast optimization oracles.

Notably, many of these recent developments have leveraged a mix of ideas coming from both discrete and continuous optimization. In particular, Ene et al. [57] present algorithms for decomposable function minimization that are based on both continuous methods (i.e. gradient descent) and discrete algorithms, as the authors employ a version of the preflow-push algorithm for maximum flow [76]. As this work was paralleled by multiple improvements to the running time for maximum flow [118, 117, 114, 95, 25, 73], most of which stemmed from innovations in convex optimization, it seemed plausible that the same new optimization techniques could be helpful for improving the running times of other fundamental problems in combinatorial optimization, including submodular function minimization. In this context, a particularly intriguing question emerged:

Can we leverage the techniques used to obtain faster algorithms for maximum flow to provide faster algorithms for submodular function minimization?

We answer this question in the affirmative, by showing how to solve decomposable submodular function minimization using *black-box access* to any routine that can compute the maximum flow in

a capacitated directed graph. To compare the running times, in the case where all the functions in the decomposition act on $O(1)$ elements of the ground set and are polynomially bounded (such as the case of a hypergraph cut function, with $O(1)$ sized hyperedges), our algorithm has – up to polylogarithmic factors – the same running time as that of computing maximum flow in a sparse graph with $O(|V|)$ vertices, and polynomial integral capacities [75, 25, 73].

As it turns out, to achieve this it is not sufficient to directly use off-the-shelf maximum flow algorithms. Instead, our approach is based on solving submodular minimization in the more general parametric setting, where we further parametrize the problem with an additional time-dependent penalty term on the elements in the set, and want to simultaneously solve *all* the problems in this family. In turn, our reduction requires solving the parametric minimum cut problem, which has been intensely studied in the classical graph theoretic literature [72, 122, 160, 77]. In this setting, which is essentially a particular case of parametric submodular minimization, the capacities of certain arcs in the graph evolve in a monotonic fashion.

While some of the existing work on parametric cuts and flows does provide efficient algorithms via reductions to maximum flow [160], the type of parametric capacities it supports does not cover the requirements for our more general scenario. Therefore, we develop a new efficient algorithm for computing parametric cuts under a broad range of parametric capacities. Our algorithm is nearly optimal from the perspective of weakly-polynomial time algorithms, since its running time matches (up to polylogarithmic factors involving certain parameters) that of the fastest maximum flow algorithm in a directed graph with integer capacities. In addition, our reduction also provides novel improvements in several other regimes, involving the strongly polynomial case, and that of planar graphs, both of which may be of independent interest.

5.1.1 Our Results

In this chapter we establish further connections between discrete and continuous optimization to provide an efficient algorithm for solving the decomposable submodular function minimization problem in the more general parametric setting. Our algorithm is at its core based on a continuous optimization method, but whose progress steps are driven by a new combinatorial algorithm we devise for the parametric cut problem. In this sense, our approach leverages the paradigm of combinatorial preconditioning from scientific computing literature [156, 17, 102, 162].

To properly state our main result, we need to introduce some notation. Let $V = \{1, \dots, n\}$ and let $F : 2^V \rightarrow \mathbb{N}$ a submodular set function with the special property that

$$F(S) = \sum_{i=1}^r F_i(S), \quad \text{for all } S \subseteq V,$$

where each $F_i : 2^V \rightarrow \mathbb{N}$ is a submodular set function acting on a subset $V_i \subseteq V$ of elements, in the sense that $F_i(S) = F_i(S \cap V_i)$ for all $S \subseteq V$. Let EO_i be the time required to evaluate $F_i(S)$ for any $S \subseteq V_i$, and let \mathcal{O}_i be the time required to minimize $F_i(S) + w(S)$ over V_i , where w is any linear function, and suppose that $\max_{S \subseteq V} F(S) = n^{O(1)}$. Furthermore, for each $i \in V$ let $\psi_i : \mathbb{R} \rightarrow \mathbb{R}$ be a strictly convex function satisfying $n^{-O(1)} \leq |\psi_i''(x)| \leq n^{O(1)}$ and $|\psi_i'(0)| \leq n^{O(1)}$.

Then our main theorem is the following.

Theorem 5.1.1. *There is an algorithm which for all $\lambda \in \mathbb{R}$ simultaneously optimizes the objective*

$$\min_{S \subseteq V} F(S) + \sum_{i \in S} \psi_i'(\lambda)$$

by returning a vector x such that for any $\lambda \in \mathbb{R}$ the set $S^\lambda = \{u : x_u \geq \lambda\}$ satisfies

$$F(S^\lambda) + \sum_{u \in S^\lambda} \psi'(u) \leq \min_{S \subseteq V} F(S) + \sum_{u \in S} \psi'(u) + \varepsilon.$$

Furthermore, if $T_{\max\text{flow}}(n, m)$ is the time required to compute the maximum flow in a directed graph with polynomially bounded integral capacities, then our algorithm runs in time

$$\tilde{O}\left(\max_i |V_i|^2 \left(\sum_{i=1}^r |V_i|^2 \mathcal{O}_i + T_{\max\text{flow}}\left(n, n + \sum_{i=1}^r |V_i|^2\right)\right) \log \frac{1}{\varepsilon}\right).$$

To better understand this result, let us consider the case where each submodular function in the decomposition acts on a small number of elements, i.e. $|V_i| = O(1)$. In this case we have the following corollary:

Corollary 5.1.1. *If each function F_i in the decomposition acts on $O(1)$ elements of the ground set, then we can solve the parametric submodular minimization problem to ε precision in time*

$$\tilde{O}\left(T_{\max\text{flow}}(n, n + r) \log \frac{1}{\varepsilon}\right).$$

While our statements concern the parametric setting, it is easy to use them to recover the solution to the standard submodular minimization problem. Simply by letting $\psi'_i(t) = t$ for all i , and thresholding the returned vector at 0 we obtain the desired result. Using Goldberg-Rao [75] or the current state of the art maximum flow advancements [25, 73, 37], we see that this significantly improves over all the previous algorithms for decomposable submodular minimization, in the regime where all sets V_i are small. Following [57] it has remained widely open whether algorithms improving upon the $\tilde{O}(\min\{n^2, nr\} \log^{O(1)}(1/\varepsilon))$ running time exist, and it has been conjectured that faster running times could be obtained by leveraging the newer techniques for graph algorithms based on interior point methods.

Using [37], we obtain a running time of $O((n + r)^{1+o(1)} \log 1/\varepsilon)$.

The crucial subroutine our algorithm is based on is a novel efficient algorithm for solving the parametric cut problem using a maximum flow oracle. We give an overview of our reduction and its additional applications in Section 5.3, and describe it in detail in Appendix 9.3.3.

5.1.2 Previous Work

Related Works on Submodular Minimization Submodular function minimization is a classical problem in combinatorial optimization, which goes back to the seminal work of Edmonds [54]. The first polynomial-time algorithm was obtained by Grötschel et al. [78] using the ellipsoid method. This was followed by a plethora of improvements, among which the more recent ones [43, 91] leveraged related techniques. On a different front, there has been significant work dedicated to obtaining strongly polynomial time algorithms for this problem [61, 84, 85, 132, 143, 107, 43, 91].

For the more structured regime of *decomposable submodular function minimization*, algorithms based on both discrete and continuous methods have been developed. Kolmogorov [100] has shown that this problem reduces to computing maximum submodular flows, and gave an algorithm for this problem based on augmenting paths. This was followed by further algorithms based on discrete methods [8, 60]. The continuous methods are based on convex optimization on the submodular base polytope, which is also used here. Notably, Stobbe and Krause [158] tackled this problem using gradient descent, Nishihara et al. [131] used alternating projections to obtain an algorithm with

linear convergence, Ene and Nguyen [56] achieved an improved algorithm with a linear convergence rate based on accelerated coordinate descent, while Ene et al. provided further improvements both via gradient descent and combinatorial techniques [57].

Related Works on Parametric Min Cut The seminal work of Gallo et al. [72] studied the generalization of the maximum flow problem where some edge-capacities, instead of being fixed, are allowed to be (possibly different) monotonic functions of a single parameter. They showed how to modify certain versions of the push-relabel algorithm for ordinary maximum flow to the parametric problem with the same asymptotic time complexity. In particular, using the Goldberg-Tarjan max-flow algorithm [76] they gave an $O(nm \log(n^2/m))$ time bound for the parametric version. Their algorithm can compute the min cuts either when a set of parameter values are given [79] or the capacity functions are all affine functions of the parameter λ .

Several other max-flow algorithms were also shown to fit into their framework (see e.g., [77]) though all requiring $\Omega(mn)$ time in the worst case. Further generalizations of the parametric min-cut problems have also been considered [122, 77]. When all parameterized capacities are equal to the parameter λ , Tarjan et al. [160] give a divide and conquer approach that can use any maximum flow algorithm as a black box and is a factor $\min\{n, \log(nU)\}$ worse.

5.2 Background and Preliminaries

5.2.1 Submodular Set Functions and Convex Analysis

Let V be a finite ground set of size n , and we assume w.l.o.g. that $V = \{1, \dots, n\}$. A set function $F : 2^V \rightarrow \mathbb{R}$ is submodular if $F(A) + F(B) \geq F(A \cup B) + F(A \cap B)$ for any two sets $A, B \subseteq V$. We are concerned with minimizing submodular set functions of the form $F = \sum_{i=1}^r F_i$, where each F_i is a submodular set function:

$$\min_{A \subseteq V} F(A) = \min_{A \subseteq V} \sum_{i=1}^r F_i(A).$$

For the rest of the chapter we will assume that F_i are non-negative, integral, and that $\max_{S \subseteq V} F(S) \leq F_{\max}$. The non-negativity constraint holds without loss of generality, as we can simply shift each F_i by a constant until it becomes non-negative. For rational functions that are represented on bounded bit precision, the integrality can be enforced simply by scaling them, at the expense of increasing F_{\max} . As we will see, some of our subroutines depend on the magnitude of F_{\max} , so we will generally assume that this is polynomially bounded.

As in previous works [131, 56, 57, 94], in this chapter we are concerned with the regime where each function F_i in the decomposition acts on few elements of the ground set V . More precisely for each $i \in \{1, \dots, r\}$ there is a *small* set $V_i \subseteq V$ such that $F_i(A) = F_i(A \cap V_i)$ for all $A \subseteq V$. We assume w.l.o.g. that $F_i(\emptyset) = F_i(V_i)$, which we discuss in more detail in Section 9.3.4. The running time of our algorithm depends on $\max_{1 \leq i \leq r} |V_i|$. This assumption is important as, furthermore, the final running time of our algorithm depends on (i) the time \mathcal{O}_i to optimize functions of the form $F_i(S) + w(S)$ over V_i , where w is a linear function and (ii) the time EO_i to evaluate F_i for subsets of V_i . In the case where $|V_i| = O(1)$, this is also constant time.

Given an arbitrary vector $w \in \mathbb{R}^n$ and a subset $A \subseteq V$, we use the notation $w(A) = \sum_{i \in A} w_i$.

Definition 5.2.1. *Given a submodular set function $F : 2^V \rightarrow \mathbb{R}$, such that $F(\emptyset) = 0$, its submodular*

base polytope $B(F)$ is defined as follows:

$$B(F) = \{w \in \mathbb{R}^n : w(A) \leq F(A) \text{ for all } A \subseteq V, w(V) = F(V)\} .$$

Definition 5.2.2. Given a submodular set function $F : 2^V \rightarrow \mathbb{R}$, $F(\emptyset) = 0$, its Lovász extension $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is defined over $[0, 1]^n$ as the convex closure of F . However, it will be more convenient to consider its extension of \mathbb{R}^n , given by

$$f(x) = \int_0^\infty F(\{i : x_i \geq t\}) dt + \int_{-\infty}^0 (F(\{i : x_i \geq t\}) - F(V)) dt .$$

Fact 5.2.1. It is well known [15] that the Lovász extension of a submodular set function F can be equivalently characterized in terms of its submodular base polytope $B(F)$. More precisely, if $F(\emptyset) = 0$, then:

$$f(x) = \max_{w \in B(F)} \langle w, x \rangle .$$

For parametric submodular function minimization we consider a family of functions parameterized by $\alpha \in \mathbb{R}$:

$$F_\alpha(A) = F(A) + \sum_{i \in A} \psi'_i(\alpha) , \quad (5.1)$$

where $\psi_j : \mathbb{R} \rightarrow \mathbb{R}$ are strictly convex differentiable functions, satisfying $\lim_{\alpha \rightarrow -\infty} \psi'_i(\alpha) = -\infty$ and $\lim_{\alpha \rightarrow \infty} \psi'_i(\alpha) = \infty$, for all i . A common example is $\psi'_j(\alpha) = \alpha$, which imposes an ℓ_1 penalty on the size of the set A . It is shown in [36, 15] that minimizing $F_\alpha(A)$ for the entire range of scalars α amounts to minimizing a regularized version of the Lovász extension.

Lemma 5.2.3. Let F_α be the family of parameterized submodular set functions defined as in (5.1), where ψ_i are strictly convex functions. Let f be the Lovász extension of F , and consider the optimization problem

$$\min_{x \in \mathbb{R}^n} f(x) + \sum_{i \in V} \psi_i(x_i) . \quad (5.2)$$

Let $A^\alpha = \arg \min_{A \subseteq V} F_\alpha(A)$, and let x^* be the minimizer of (5.2). Then

$$A^\alpha = \{i : x_i^* \geq \alpha\} . \quad (5.3)$$

For completeness we reproduce the proof of Lemma 5.2.3 in Section 9.3.2.

Via convex duality one can prove that minimizing (5.2) is equivalent to a dual optimization problem on the submodular base polytope $B(F)$:

$$\min_{w \in B(F)} \sum_{i \in V} \psi_i^*(-w_i) , \quad (5.4)$$

where ψ_i^* is the Fenchel dual of ψ_i .

Definition 5.2.4 (Fenchel dual). Let $g : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$ be a convex function. Its Fenchel dual or convex conjugate $g^* : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$ is defined as

$$g^*(w) = \sup_x \langle w, x \rangle - g(x) .$$

We will refer to (5.2) as the primal problem and (5.4) as the dual problem. The algorithm described in this chapter will focus on optimizing (5.4) while strongly leveraging the decomposable

structure of F . We assume that all functions ψ_i have “nice” second derivatives, which will play an important role in the algorithm, since this will also ensure that the minimizers of (5.2) and (5.4) are unique.

Assumption 5.2.2. *The function ψ_i is L -smooth and σ -strongly convex for all $i \in V$. Equivalently for each i , its second derivative satisfies $0 < \sigma \leq \psi_i''(x) \leq L$, for all $x \in \mathbb{R}$. Furthermore, $|\psi_i'(0)| \leq n^{O(1)}$, for all $i \in V$.*

This condition also helps us ensure that we can efficiently convert between the primal and dual spaces onto which the ψ_i and its Fenchel dual ψ_i^* act. Also, whenever it is convenient, we will use the notation $\psi(x) = \sum_{i \in V} \psi_i(x_i)$, $\psi^*(y) = \sum_{i \in V} \psi_i^*(y_i)$.

5.2.2 Overview of Approach

Decomposable Submodular Minimization Here we provide an overview of our approach for minimizing decomposable submodular functions. Our approach for the parametric setting yields a strictly stronger result without sacrificing running time, so we will focus on this more general problem.

Our approach is based on minimizing a convex function on the submodular base polytope $B(F)$. As it has been seen in previous works [15], in order to solve the parametric problem (5.1), it suffices to solve the dual problem (5.4), which is a convex optimization problem over $B(F)$. For convenience let us denote by $h(w) = \sum_{i \in V} \psi_i^*(-w_i)$, so that our objective becomes computing $\min_{w \in B(F)} h(w)$.

We use an iterative method, which maintains a point $w \in B(F)$ and updates it in such a way that the objective value improves significantly in each step. To do so, we find a polytope P such that

$$w + \frac{1}{\alpha} \cdot P \subseteq B(F) \subseteq w + P \tag{5.5}$$

and such that we can efficiently minimize ψ over $w + \frac{1}{\alpha} \cdot P$. If we can find the minimizer w' over $w + \frac{1}{\alpha} \cdot P$, then moving our iterate to w' also guarantees that

$$h(w') - h(w^*) \leq \left(1 - \frac{1}{\alpha}\right) (h(w) - h(w^*)),$$

where w^* is the minimizer of h over $B(F)$. This is true due to the convexity of h . Indeed, let $\tilde{w} = w + t(w^* - w)$ where $t = \max\{t \leq 1 : w + t(w^* - w) \in w + \frac{1}{\alpha}P\}$; in other words \tilde{w} represents the furthest point on the segment connecting w and w^* such that \tilde{w} still lies inside the small polytope $w + \frac{1}{\alpha}P$. Due to the sandwiching property of the polytopes (5.5), we have that $t \geq 1/\alpha$. Hence, using the convexity of h , we obtain that

$$h(\tilde{w}) - h(w^*) = h(w + t(w^* - w)) - h(w^*) \leq (1 - t)(h(w) - h(w^*)) \leq (1 - 1/\alpha)(h(w) - h(w^*)).$$

Since w' minimizes h over $w + \frac{1}{\alpha} \cdot B(F)$, we must have $h(w') \leq h(\tilde{w})$, and we obtain the desired progress in function value. Thus iterating $\tilde{O}(\alpha)$ times we obtain a high precision solution, which we then convert back to a combinatorial solution to the original problem using some careful error analysis.

More importantly, we need to address the question of finding a polytope P satisfying (5.5).

To do so, for each i , we define the “residual” submodular functions $F'_i(A) = F_i(A) - w_i(A)$ for all $A \subseteq V$, where $w_i \in B(F_i)$ such that $\sum_{i=1}^r w_i = w$. The existence of such a decomposition of $w \in B(\sum_{i=1}^r F_i)$ is well-known, and goes back to Edmonds [53]. Very importantly, we note that since F_i were non-negative, F'_i remain non-negative submodular set functions.

It is known [45] that non-negative submodular set functions can be approximated by graph cuts. Following the proof from [45], for each i we construct a graph on $O(|V_i|)$ vertices whose cuts approximate the value of F'_i within a factor of $O(|V_i|^2)$. Combining all these graphs into a single one, we obtain a weighted directed graph on $O(|V|)$ vertices and $O(|V| + \sum_{i=1}^r |V_i|^2)$ arcs such that its cut function G approximates F' within a factor of $O(\max_i |V_i|^2)$.

Crucially, we can show that if G is the cut function which approximates F' , then we also have that

$$\frac{1}{\max_i |V_i|^2} \cdot B(G) \subseteq B(F'_i) \subseteq B(G),$$

and therefore it suffices to implement a routine that minimizes h over $w + \frac{1}{\max_i |V_i|^2} \cdot B(G)$ in order to obtain an algorithm that terminates in $\tilde{O}(\max_i |V_i|^2)$ such iterations.

To implement this routine, we devise a new combinatorial algorithm for solving the parametric flow problem, with general parameterized capacities. By comparison to previous literature, our algorithm efficiently leverages a maximum flow oracle on a sequence of graphs obtained via contracting edges, and whose running time is up to polylogarithmic factors equal to that of computing a maximum flow in a capacitated directed graph with $O(|V|)$ vertices and $O(|V| + \sum_{i=1}^r |V_i|^2)$ arcs.

Following this, we convert the combinatorial solution to the parametric flow problem into a solution to its corresponding dual problem on the submodular base polytope, which returns the new iterate w' .

Throughout the algorithm we need to control the errors introduced by the fact that both the solution we receive for the parametric flow problem and the one we return as an approximate minimizer of h over $B(F)$ are *approximate*, but these are easily tolerable since our main routines return high precision solutions.

5.3 Parametric Min s, t -Cut

In the general parametric min s, t -cut problem [72], the capacities of the source's outgoing edges (s, v) are (possibly different) nonnegative real nondecreasing functions of a parameter $\lambda \in \mathbb{D}$, where $\mathbb{D} \subseteq \mathbb{R}$ is some domain, whereas the capacities of the sink's incoming edges vt are nonincreasing functions of λ . The goal is to compute the representation of the *cut function* $\kappa : \mathbb{D} \rightarrow \mathbb{R}$ such that $\kappa(\lambda)$ equals the capacity of the minimum s, t -cut in G_λ obtained from G by evaluating the parameterized capacity functions at λ . It is known that κ consists of $O(n)$ pieces, where κ equals the parameterized capacity of some fixed cut in G .

More formally, let $\lambda_{\min} \in \mathbb{D}$ be such that the minimal min s, t -cuts of $G_{\lambda_{\min}}$ and $G_{\lambda'}$ are equal for all $\lambda' \in \mathbb{D}$, $\lambda' < \lambda_{\min}$. Similarly, let $\lambda_{\max} \in \mathbb{D}$ be such that the minimal min s, t -cuts of $G_{\lambda_{\max}}$ and $G_{\lambda'}$ are equal for all $\lambda' \in \mathbb{D}$ with $\lambda' > \lambda_{\max}$. We will consider λ_{\min} and λ_{\max} inputs to our problem. Then, there exist $O(n)$ *breakpoints* $\Lambda = \{\lambda_1, \dots, \lambda_k\}$, $\lambda_{\min} = \lambda_0 < \lambda_1 < \dots < \lambda_k$ and an embedding of vertices $\tau : V \rightarrow \Lambda \cup \{\lambda_{\min}, \infty\}$ such that for all $i = 0, \dots, k-1$, $\lambda' \in [\lambda_i, \lambda_{i+1}) \cap \mathbb{D}$, $\kappa(\lambda')$ equals the capacity of the cut $S(\lambda_i) = \{v \in V : \tau(v) \leq \lambda_i\}$ in $G_{\lambda'}$, and also $S(\lambda_k)$ is a min s, t -cut of $G_{\lambda_{\max}}$.

Motivated by our submodular minimization application, our algorithm in the most general setting solves the ε -approximate parametric min s, t -cut problem.

Definition 5.3.1 (ε -approximate parametric min s, t -cut). *Let Λ , τ , and $S : \mathbb{D} \rightarrow 2^V$ be as defined above. A pair (Λ, τ) is called an ε -approximate parametric min s, t -cut of G if:*

1. For $i = 0, \dots, k-1$, $S(\lambda_i)$ is a min s, t -cut of $G_{\lambda'}$ for all $\lambda' \in [\lambda_i, \lambda_{i+1} - \varepsilon) \cap \mathbb{D}$.

2. $S(\lambda_k)$ is a min s, t -cut of $G_{\lambda_{\max}}$.
3. For $i = 0, \dots, k - 1$, $S(\lambda_i) \subsetneq S(\lambda_{i+1})$.

We prove that an ε -approximate parametric min s, t -cut yields breakpoints within ε additive error wrt. to the breakpoints of the exact parametric min s, t -cut. Our algorithm solves the above problem assuming only constant-time black-box access to the capacity functions.

Theorem 5.3.2. *Let $R = \lambda_{\max} - \lambda_{\min}$ be an integral multiple of $\varepsilon > 0$. Let $T_{\max\text{flow}}(n', m') = \Omega(m' + n')$ be a convex function bounding the time needed to compute maximum flow in a graph with n' vertices and m' edges obtained from G_λ by edge/vertex deletions and/or edge contractions (with merging parallel edges by summing their capacities) for any $\lambda = \lambda_{\min} + \ell\varepsilon$ and any integer $\ell \in [0, R/\varepsilon]$. Then, ε -approximate parametric min s, t -cut in G can be computed in $O(T_{\max\text{flow}}(n, m \log n) \cdot \log \frac{R}{\varepsilon} \cdot \log n)$ time.*

The algorithm is recursive. In order to ensure uniqueness of the minimum cuts considered, it always computes cuts with *minimal* s -side.

Roughly speaking, given initial guesses $\lambda_{\min}, \lambda_{\max}$ the algorithm finds, using $O(\log((\lambda_{\max} - \lambda_{\min})/\varepsilon))$ maximum flow computations, the most balanced split λ_1, λ_2 of the domain such that (1) the s -sides for all the min-cuts of $G_{\lambda'}$ for $\lambda' > \lambda_2$ have size at least $n/2$, (2) the t -sides of all the min-cuts of $G_{\lambda'}$ for $\lambda' < \lambda_1$ have size at least $n/2$, (3) $\lambda_2 - \lambda_1 \leq \varepsilon$. We then recurse on the intervals $[\lambda_{\min}, \lambda_1]$ and $[\lambda_2, \lambda_{\max}]$ on *minors* of G with at least $n/2$ vertices contracted. Even though the contraction requires merging parallel edges in order to have at most $m + n$ (as opposed to $2m$) edges in the recursive calls, we are able to guarantee that the capacity functions in the recursive calls are all obtained by shifting the original capacity functions by a real number, and thus can be evaluated in constant time as well. Since the number of vertices decreases by a factor of two in every recursive call, one can prove that for each level of the recursion tree, the sum of numbers of vertices in the calls at that level is $O(n)$, whereas the sum of sizes of edge sets is $O(m + n \log n)$.

We show that the ε -approximate algorithm can be turned into an exact algorithm in two important special cases. First of all, if the capacity functions are low-degree (say, at most 4) polynomials with *integer coefficients* in $[-U, U]$, then one can compute parametric min s, t -cut only $O(\text{polylog}\{n, U\})$ factors slower than best known max-flow algorithm for integer capacities [73, 25, 75]. Second, we can solve the *discrete domain* case, i.e., when \mathbb{D} has finite size ℓ with only $O(\text{polylog}\{n, \ell\})$ multiplicative overhead wrt. the respective maximum flow algorithm.

Moreover, since our reduction runs maximum flow computations only on *minors* of the input graph G , it also yields very efficient parametric min s, t -cut algorithms for planar graphs. In particular, since near-optimal strongly-polynomial s, t -max flow algorithms for planar graphs are known [22, 58], we obtain near-optimal algorithms for the integer polynomial capacity functions (as above) and discrete domains. What is perhaps more surprising, using our reduction we can even obtain a *strongly polynomial* exact parametric min s, t -cut algorithm for planar graphs with linear capacity functions with real coefficients. The algorithm runs in $\tilde{O}(n^{1.21875})$ time and constitutes the only known subquadratic strongly polynomial parametric min- s, t -cut algorithm.

The details of our parametric min s, t -cut algorithm and its applications are covered in Appendix 9.3.3. It should be noted that the idea of using cuts contraction is not new and appeared previously in [160]. Compared to [160], our reduction provably handles more general parameterized capacity functions. As it does not operate on any auxiliary networks that may not preserve structural properties of G , but merely on minors of G , it proves much more robust in important special cases such as planar graphs. Finally, we believe that our reduction is also more natural and operates on the breakpoints of the cut function directly, whereas the reduction of [160] operates on so-called balanced flows.

5.4 Parametric Decomposable Submodular Minimization via Base Polytope Approximations

5.4.1 Algorithm Overview

In Algorithm 9 we give the description of our the main routine.

Algorithm 9 Parametric Decomposable Submodular Function Minimization

- 1: **Input:** ε : error tolerance ▷ Returns ε -optimal solution of $\min_{w \in B(F)} \psi^*(-w)$
 - 2: Set $w^{0,i} = 0$ for $i \in [r]$ ▷ Initialize a feasible solution
 - 3: Set $T = \alpha \log \frac{\psi^*(-w^0) + \psi(0)}{\varepsilon}$
 - 4: **for** $t = 1 \dots T$ **do**
 - 5: Set $G_i(V_i, E_i, c_i) = \text{GRAPHAPPROX}(w^{t-1,i})$, for $i \in [r]$, and construct $G(V, E, c)$ by combining the graphs G_i
 - 6: Set $\phi(x) := \psi(x) + \langle \sum_{i=1}^r w^{t-1,i}, x \rangle$
 - 7: Set $\tilde{w} = \text{FINDMINCUTS}(G(V, E, c), \phi, \frac{1}{3L})$ ▷ Find parametric min s, t -cuts
 - 8: Round all the entries of \tilde{w} to the nearest integer
 - 9: Decompose $\tilde{w} = \sum_{i=1}^r \tilde{w}^i$, with $\tilde{w}^i \in B(G_i)$ using Lemma 5.4.3.
 - 10: Set $w^{t,i} = w^{t-1,i} + \tilde{w}^i$, for all $i \in [r]$.
 - 11: **return** $\sum_{i=1}^r w^{T,i}$
-

5.4.2 Removing Assumptions

In Section 5.2 we assumed that for all i , $F_i(\emptyset) = F_i(V_i) = 0$ and $F_i(S) \geq 0$ for all S . These assumptions hold without loss of generality. In Section 9.3.4 we show how to preprocess the input such that these assumptions are valid.

5.4.3 From Parametric Minimum Cut to Cut Base Polytope Optimization

In this section, we will focus on the problem of minimizing a convex function over the base polytope of the s, t -cut function of a graph $G(V \cup \{s, t\}, E, c)$. We define the s, t -cut function $G : 2^V \rightarrow \mathbb{R}$ as $G(S) = c^+(S \cup \{s\})$ for $S \subseteq V$, and let $B(G)$ be the base polytope of G . Now, the parametric min s, t -cut problem can be written as

$$\min_{S \subseteq V} G(S) + \sum_{u \in S} \phi'_u(\lambda), \quad (5.6)$$

where the capacity of an edge (u, t) at time λ is $c_{ut} + \phi'_u(\lambda)$ and ϕ is a function satisfying Assumption 5.2.2. In particular, our goal in this section is, given solutions to (5.6) for all λ , to solve the following dual problem:

$$\min_{w \in B(G)} \phi^*(-w). \quad (5.7)$$

Definition 5.4.1 (W -restricted function). *A submodular function $F : 2^V \rightarrow \mathbb{R}_{\geq 0}$ is called W -restricted if $F(S) = F(S \cap W)$ for all $S \subseteq V$, where $W \subseteq V$.*

As the cut functions $G(S)$ that we will be concerned with will be decomposable, i.e. $G(S) = \sum_{i=1}^r G_i(S)$ for all $S \subseteq V$, we introduce the following notion of a *decomposition* of some $w \in B(G)$ into a sum of $w^i \in B(G_i)$.

Definition 5.4.2 (*F*-decomposition). *Let $F : 2^V \rightarrow \mathbb{R}_{\geq 0}$ with $|V| = n$ be a submodular function that is decomposable, i.e. $F(S) = \sum_{i=1}^r F_i(S)$ for all $S \subseteq V$, where $F_i : 2^V \rightarrow \mathbb{R}_{\geq 0}$ are submodular functions. Then for any $w \in B(F)$ there exist vectors $w^1, \dots, w^r \in \mathbb{R}^n$, where $w^i \in B(F_i)$ for all $i \in [r]$ and $\sum_{i=1}^r w^i = w$. We call the sequence of vectors w^1, \dots, w^r an (F_1, \dots, F_r) -decomposition of w , or just an *F*-decomposition of w if the F_i 's are clear from context.*

What follows is the main lemma of this section, whose full proof appears in Appendix 9.3.5.

Lemma 5.4.3 (From parametric min-cut to cut base polytope optimization). *Consider a graph $G(V, E, c \geq 0)$ and a function $\phi(x) = \sum_{u \in V} \phi_u(x_u)$ that satisfies Assumption 5.2.2. Additionally, let $G(S) = c^+(S)$ for all $S \subseteq V$ be the cut function associated with the graph, and suppose it is decomposable as $G(S) = \sum_{i=1}^r G_i(S)$ where $G_i : 2^V \rightarrow \mathbb{Z}_{\geq 0}$ are V_i -restricted cut functions defined as $G_i(S) = c^{i+}(S)$ that correspond to graphs $G_i(V, E, c^i \geq 0)$, and $c = \sum_{i=1}^r c^i$.*

We define an extended vertex set $V' = V \cup \{s, t\}$ with edge set $E' = E \cup \bigcup_{u \in V} (s, u) \cup \bigcup_{u \in V} (u, t)$, the parametric capacity of an edge $(u, v) \in E'$ as

$$c_\lambda(u, v) = \begin{cases} \max\{0, \phi'_u(-\lambda)\} & \text{if } u \in V, v = t \\ \max\{0, -\phi'_u(-\lambda)\} & \text{if } u = s, v \in V \\ c_{uv} & \text{otherwise} \end{cases}$$

and let (Λ, τ) be a $\frac{1}{3L}$ -approximate parametric min s, t -cut of $G'(V', E', c_\lambda)$.

There exists an algorithm that, given (Λ, τ) , outputs $\tilde{w}^ = \operatorname{argmin}_{w \in B(G)} \phi^*(-w)$ and a G -decomposition $\tilde{w}^{*1}, \dots, \tilde{w}^{*r}$ of \tilde{w}^* . The running time of this algorithm is $O\left(n + \sum_{i=1}^r |V_i|^2\right)$.*

5.4.4 Dual Progress Analysis in One Step

Lemma 5.4.4 (Dual progress in one step). *Let $F : 2^V \rightarrow \mathbb{Z}_{\geq 0}$ be a submodular function that is separable, i.e. $F(S) = \sum_{i=1}^r F_i(S)$ for all $S \subseteq V$, where $F_i : 2^V \rightarrow \mathbb{Z}_{\geq 0}$ are V_i -restricted submodular functions with $F_i(\emptyset) = 0$. Additionally, let $\psi : \mathbb{R}^n \rightarrow \mathbb{R}$ be a function that satisfies Assumption 5.2.2, where $|V| = n$.*

Given a feasible dual solution $w \in B(F)$ and an F -decomposition $w^1, \dots, w^r \in \mathbb{Z}^n$ of w , there is an algorithm that outputs a vector $w' \in \mathbb{Z}^n$, along with an (F_1, \dots, F_r) -decomposition $w'^1, \dots, w'^r \in \mathbb{Z}^n$ of w' , such that

$$\psi^*(-w') - \psi^*(-w^*) \leq \left(1 - \frac{1}{\alpha}\right) (\psi^*(-w) - \psi^*(-w^*))$$

where $\alpha = \max_{i \in [r]} \{|V_i|^2/4\}$ and $w^* = \operatorname{argmin}_{w^* \in B(F)} \psi^*(-w^*)$ is the dual optimum. The running time of the algorithm is

$$\tilde{O} \left(\sum_{i=1}^r |V_i|^2 \mathcal{O}_i + T_{\max\text{flow}} \left(n, n + \sum_{i=1}^r |V_i|^2 \right) \right).$$

The full proof of Lemma 5.4.4 appears in Appendix 9.3.5.

5.4.5 Main Theorem

Proof of Theorem 5.1.1. We repeatedly apply Lemma 5.4.4, and let w^0, \dots, w^T be the iterates after $T = \alpha \log \frac{\psi^*(-w^0) - \psi^*(-w^*)}{\zeta}$ iterations. We have

$$\psi^*(-w^T) - \psi^*(-w^*) \leq \left(1 - \frac{1}{\alpha}\right) (\psi^*(-w^{T-1}) - \psi^*(-w^*)).$$

Applying induction over T steps we obtain that

$$\psi^*(-w^T) - \psi^*(-w^*) \leq \left(1 - \frac{1}{\alpha}\right)^T (\psi^*(-w^0) - \psi^*(-w^*)) \leq e^{-T/\alpha} (\psi^*(-w^0) - \psi^*(-w^*)) \leq \zeta.$$

We have obtained a high precision solution to the objective (5.4). Finally, setting $\frac{1}{\zeta} = \operatorname{poly}(\frac{L}{\sigma} n F_{\max}/\varepsilon) = n^{O(1)}/\varepsilon$ and applying Corollary 9.3.1 to convert from this solution to the actual sets, we obtain the desired solution.

As Assumption 5.2.2 implies $\psi^*(-w^0) - \psi^*(-w^*) = n^{O(1)}$, the total running time is

$$\tilde{O} \left(\max_i |V_i|^2 \left(\sum_{i=1}^r |V_i|^2 \mathcal{O}_i + T_{\max\text{flow}} \left(n, n + \sum_{i=1}^r |V_i|^2 \right) \right) \log \frac{1}{\varepsilon} \right).$$

□

Acknowledgements

KA was supported in part by the NSF grants CCF-1553428 and CNS-1815221. AK, AM and PS were supported in part by ERC CoG project TUGbOAT no 772346. We are grateful to Michael B. Cohen for discussions on the potential of second order methods for submodular minimization, which motivated parts of this project. We thank Alina Ene for pointing us to the relevant material from [15].

Part II

Optimization Under Sparsity Constraints

Chapter 6

Sparse Convex Optimization via Adaptively Regularized Hard Thresholding

6.1 Introduction

Sparse Convex Optimization is the problem of optimizing a convex objective, while constraining the sparsity of the solution (its number of non-zero entries). Variants and special cases of this problem have been studied for many years, and there have been countless applications in Machine Learning, Signal Processing, and Statistics. In Machine Learning it is used to regularize models by enforcing parameter sparsity, since a sparse set of parameters often leads to better model generalization. Furthermore, in a lot of large scale applications the number of parameters of a trained model is a significant factor in computational efficiency, thus improved sparsity can lead to improved time and memory performance. In applied statistics, a single extra feature translates to a real cost from increasing the number of samples. In compressed sensing, finding a sparse solution to a Linear Regression problem can be used to significantly reduce the sample size for the recovery of a target signal. In the context of these applications, decreasing sparsity by even a small amount while not increasing the accuracy can have a significant impact.

6.1.1 Sparse Optimization

Given a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ and any s^* -sparse (*unknown*) target solution x^* , the Sparse Optimization problem is to find an s -sparse solution x , i.e. a solution with at most s non-zero entries, such that $f(x) \leq f(x^*) + \varepsilon$ and $s \leq s^* \gamma$, where $\varepsilon > 0$ is a desired accuracy and $\gamma \geq 1$ is an approximation factor for the target sparsity. Even if f is a convex function, the sparsity constraint makes this problem non-convex, and it has been shown that it is an intractable problem, even when $\gamma = O(2^{\log^{1-\delta} n})$ and f is the Linear Regression objective [127, 63]. However, this worst-case behavior is not observed in practice, and so a large body of work has been devoted to the analysis of algorithms under the assumption that the *restricted condition number* $\kappa_{s+s^*} = \frac{\rho_{s+s^*}^+}{\rho_{s+s^*}^-}$ (or just $\kappa = \frac{\rho^+}{\rho^-}$) of f is bounded [127, 146, 170, 16, 113, 89, 169, 148, 149, 89, 154]. Note: Here, $\rho_{s+s^*}^+$ is the maximum smoothness constant of any restriction of f on an $(s + s^*)$ -sparse subset of coordinates and $\rho_{s+s^*}^-$ is the minimum strong convexity constant of any restriction of f on an $(s + s^*)$ -sparse subset of coordinates.

The first algorithm for this problem, often called *Orthogonal Matching Pursuit (OMP)* or *Greedy*,

was analyzed by [127] for Linear Regression, and subsequently for general f by [146], obtaining the guarantee that the sparsity of the returned solution is $O\left(s^* \kappa \log \frac{f(x^0) - f(x^*)}{\varepsilon}\right)$.¹ In applications where having low sparsity is crucial, the dependence of sparsity on the required accuracy ε is undesirable. The question of whether this dependence can be removed was answered positively [146, 89] giving a sparsity guarantee of $O(s^* \kappa^2)$. As remarked in [146], this bound sacrifices the linear dependence on κ , while removing the dependence on ε and $f(x^0) - f(x^*)$.

Since then, there has been some work on improving these results by introducing non-trivial assumptions, such as the target solution x^* being close to globally optimal. More specifically, [170] defines the *Restricted Gradient Optimal Constant (RGOC) at level s* , ζ_s (or just ζ) as the ℓ_2 norm of the top- s elements in $\nabla f(x^*)$ and analyzes an algorithm that gives sparsity $s = O(s^* \kappa \log(s^* \kappa))$, and such that $f(x) \leq f(x^*) + O(\zeta^2/\rho^-)$. [154] strengthens this bound to $f(x) \leq f(x^*) + O(\zeta^2/\rho^+)$ with sparsity $s = O(s^* \kappa \log \kappa)$. However, this means that $f(x)$ might be much larger than $f(x^*) + \varepsilon$ in general. To the best of our knowledge, no improvement has been made over the $O\left(s^* \min\left\{\kappa \frac{f(x^0) - f(x^*)}{\varepsilon}, \kappa^2\right\}\right)$ bound in the general case.

Related work. Sparse convex optimization is closely related to the problem of optimizing a convex function under a rank constraint, which is a very general optimization problem encompassing matrix completion, robust principal component analysis, and others. Close analogues of OMP and OMPR that give guarantees on the rank of the solution based on the condition number have been analyzed for this setting [144, 13].

Another line of work studies a maximization version of the sparse convex optimization problem as well as its generalizations for matroid constraints [6, 55, 38].

6.1.2 Sparse Solution and Support Recovery

Often, as is the case in compressed sensing, one needs a guarantee on the closeness of the solution x to the target solution x^* in absolute terms, rather than in terms of the value of f . The goal is usually either to recover (a superset of) the target support, or to ensure that the returned solution is close to the target solution in ℓ_2 norm. The results for this problem either assume a constant upper bound on the *Restricted Isometry Property (RIP)* constant $\delta_r := \frac{\kappa_r - 1}{\kappa_r + 1}$ for some r (RIP-based recovery), or that x^* is close to being a global optimum (RIP-free recovery). This problem has been extensively studied and is an active research area in the vast compressed sensing literature. See also the survey by [21].

In the seminal papers of [33, 32, 46, 31] it was shown that for the Linear Regression problem when $\delta_{2s^*} < \sqrt{2} - 1 \approx 0.41$, the LASSO algorithm [161] can recover a solution with $\|x - x^*\|_2^2 \leq C f(x^*)$, where C is a constant depending only on δ_{2s^*} and $f(x^*) = \frac{1}{2} \|Ax^* - b\|_2^2$ is the error of the target solution². Since then, a multitude of results of similar flavor have appeared, either giving related guarantees for the LASSO algorithm while improving the RIP upper bound [67, 30, 64, 29, 124, 7] which culminate in a bound of $\delta_{2s^*} < 0.6248$, or showing that similar guarantees can be obtained by greedy algorithms under more restricted RIP conditions, but that are typically faster than LASSO [130, 129, 128, 19, 87, 65, 66]. See also the comprehensive surveys by [68, 125].

[128] presents a greedy algorithm called *CoSaMP* and shows that for Linear Regression it achieves a bound in the form of [31] while having a more efficient implementation. Their method works for the more restricted RIP upper bound of $\delta_{2s^*} < 0.025$, or $\delta_{4s^*} < 0.4782$ as improved by [68]. [19] proves that another greedy algorithm called *Iterative Hard Thresholding (IHT)* achieves a similar

¹Even though [127] states a less general result, this is what is implicitly proven.

² $f(x^*)$ is also commonly denoted as $\frac{1}{2} \|\eta\|_2^2$, where $Ax^* = b + \eta$, i.e. η is the measurement *noise*.

bound to that of CoSaMP for Linear Regression, with the condition $\delta_{3s^*} < 0.067$, which is improved to $\delta_{2s^*} < \frac{1}{3}$ by [87] and to $\delta_{3s^*} < 0.5774$ by [65].

The RIP-free line of research has shown that strong results can be achieved without a RIP upper bound, given that the target solution is sufficiently close to being a global optimum. These results typically require that s is significantly larger than s^* . In particular, [170] shows that if ζ is the RGOC of f it can be guaranteed that $\|x - x^*\|_2 \leq 2\sqrt{6}\frac{\zeta}{\rho^-}$ (or $(1 + \sqrt{6})\frac{\zeta}{\rho^-}$ with a slightly tighter analysis). [154] strengthens this bound to $(1 + \sqrt{1 + \frac{5}{\kappa}})\frac{\zeta}{\rho^-}$. Furthermore, it has been shown that as long as a ‘‘Signal-to-Noise’’ condition holds, one can actually recover a superset of the target support. Typically the condition is a lower bound on $|x_{\min}^*|$, the minimum magnitude non-zero entry of the target solution. Different lower bounds that have been devised include $\Omega\left(\frac{\sqrt{s+s^*}\|\nabla f(x^*)\|_\infty}{\rho_{s+s^*}^-}\right)$ [89], which was later improved to $\Omega\left(\sqrt{\frac{f(x^*)-f(\bar{x}^*)}{\rho_{2s}^-}}\right)$, where \bar{x}^* is an optimal s -sparse solution [169]. Finally, [154] improves the sparsity bound to $O(s^*\kappa \log(s^*\kappa))$ in the statistical setting and [149] shows that the sparsity can be brought down to $s = s^* + O(\kappa^2)$ if a stronger lower bound of $\Omega\left(\sqrt{\kappa}\frac{\zeta}{\rho}\right)$ is assumed.

6.1.3 Our Work

In this work we present a new algorithm called *Adaptively Regularized Hard Thresholding (ARHT)*, that closes the longstanding gap between the $O\left(s^*\kappa\frac{f(x^0)-f(x^*)}{\varepsilon}\right)$ and $O(s^*\kappa^2)$ bounds by getting a sparsity of $O(s^*\kappa)$ and thus achieving the best of both worlds. As [63] shows that for a general class of algorithms (including greedy algorithms like OMP, IHT as well as LASSO) the linear dependence on κ is necessary even for the special case of Sparse Regression, our result is tight for this class of algorithms. In Section 6.5.1 we briefly describe this example and also state a conjecture that it can be turned into an inapproximability result in Conjecture 6.5.1. Furthermore, in Section 6.5.2 we show that the $O(s^*\kappa^2)$ sparsity bound is tight for OMPR, thus highlighting the importance of regularization in our method. Our algorithm is efficient, as it requires roughly $O\left(s \log^3 \frac{f(x^0)-f(x^*)}{\varepsilon}\right)$ iterations, each of which includes one function minimization in a restricted support of size s and is simple to describe and implement. Furthermore, it directly implies non-trivial results in the area of compressed sensing.

We also provide a new analysis of OMPR [87] and show that under the condition that $s > s^*\frac{\kappa^2}{4}$, or equivalently under the RIP condition $\delta_{s+s^*} < \frac{2\sqrt{\frac{s}{s^*}-1}}{2\sqrt{\frac{s}{s^*}+1}}$, it is possible to approximately minimize the function f up to some error depending on the RIP constant and the closeness of x^* to global optimality. More specifically, we show that for any $\varepsilon > 0$ OMPR returns a solution x such that

$$f(x) \leq f(x^*) + \varepsilon + C_1(f(x^*) - f(x^{\text{opt}})),$$

where x^{opt} is the globally optimal solution, as well as

$$\|x - x^*\|_2^2 \leq \varepsilon + C_2(f(x^*) - f(x^{\text{opt}})),$$

where C_1, C_2 are constants that only depend on $\frac{s}{s^*}$ and δ_{s+s^*} . An important feature of our approach is that it provides a meaningful tradeoff between the RIP constant upper bound and the sparsity of the solution, even when the sparsity s is arbitrarily close to s^* . In other words, one can relax the RIP condition at the expense of increasing the sparsity of the returned solution. Furthermore, our

Table 6.1: Compressed sensing tradeoffs implied by Theorem 6.4.2: Sparsity vs RIP condition

s	RIP CONDITION
s^*	$\delta_{2s^*} < 0.33$
$2s^*$	$\delta_{3s^*} < 0.47$
$3s^*$	$\delta_{4s^*} < 0.55$
$30s^*$	$\delta_{31s^*} < 0.83$

analysis applies to general functions with bounded RIP constant.

Experiments with real data suggest that ARHT and a variant of OMPR which we call *Exhaustive Local Search* achieve promising performance in recovering sparse solutions.

6.1.4 Comparison to Previous Work

In this section we compare our results with previous work on sparse optimization, solution, and support recovery.

Sparse Optimization

Our Algorithm 15 (ARHT) returns a solution with $s = O(s^* \kappa)$ without any additional assumptions, thus significantly improving over the bound $O\left(s^* \min\left\{\kappa \frac{f(x^0) - f(x^*)}{\varepsilon}, \kappa^2\right\}\right)$ that was known in previous work. This proves that neither any dependence on the required solution accuracy ε , nor the quadratic dependence on the condition number κ is necessary. Furthermore, no assumption on the function or the target solution is required to achieve this bound. Importantly, previous results imply that our bound is tight up to constants for a general class of algorithms, including Greedy-type algorithms and LASSO [63].

Sparse Solution Recovery

In Corollary 6.3.11, we show that the improved guarantees of Theorem 6.3.1 immediately imply that ARHT gives a bound of $\|x - x^*\|_2 \leq (2 + \theta) \frac{\zeta}{\rho}$ for any $\theta > 0$, where ζ is the Restricted Gradient Optimal Constant. This improves the constant factor in front of the corresponding results of [170, 154].

As we saw, our Theorem 6.4.2 directly implies that OMPR gives an upper bound on $\|x - x^*\|_2^2$ of the same form as the RIP-based bounds in previous work, under the condition $\delta_{s+s^*} < \frac{2\sqrt{\frac{s}{s^*}-1}}{2\sqrt{\frac{s}{s^*}+1}}$. While previous results either concentrate on the case $s = s^*$, or sgs^* , our analysis provides a way to trade off increased sparsity for a more relaxed RIP bound, allowing for a whole family of RIP conditions where s is arbitrarily close to s^* . Specifically, if we set $s = s^*$ our work implies recovery for $\delta_{2s^*} < \frac{1}{3} \approx 0.33$, which matches the best known bound for any greedy algorithm [87], although it is a stricter condition than the $\delta_{2s^*} < 0.62$ required by LASSO [68]. Table 6.1 contains a few such RIP bounds implied by our analysis and shows that it readily surpasses the bounds for Subspace Pursuit $\delta_{3s^*} < 0.35$, CoSaMP $\delta_{4s^*} < 0.48$, and OMP $\delta_{31s^*} < 0.33$ [87, 170]. Another important feature compared to previous work is that all our guarantees are not restricted to Linear Regression and are true for any function f , as long as it satisfies the required RIP condition, which makes the result more general.

Sparse Support Recovery

Corollary 6.3.12 shows that as a direct consequence of our work, the condition $|x_{\min}^*| > \frac{\zeta}{\rho^-}$ suffices for our algorithm to recover a superset of the support with size $s = O(s^* \kappa)$. Compared to [89], we improve both the size of the superset, as well as the condition, since $\sqrt{s} \frac{\|\nabla f(x^*)\|_\infty}{\rho^-} \geq \sqrt{\frac{s}{s^*}} \frac{\zeta}{\rho^-} = \Omega\left(\frac{\zeta}{\rho^-}\right)$. Compared to [149], the bounds on the superset size are incomparable in general, but our $|x_{\min}^*|$ condition is more relaxed, since $\sqrt{\kappa} \frac{\zeta}{\rho^-} = \Omega\left(\frac{\zeta}{\rho^-}\right)$. Finally, [169] works under the condition $|x_{\min}^*| > \sqrt{\frac{2(f(x^*) - \min_{\|x\|_0 \leq s} f(x))}{\rho^-}}$, which is more relaxed since this quantity is always in $\left[\frac{1}{\sqrt{\kappa}} \frac{\zeta}{\rho^-}, \frac{\zeta}{\rho^-}\right]$, but uses a larger superset size of $O(s^* \kappa^2)$ instead of $O(s^* \kappa)$. Although not explicitly stated, [170, 154] also give a similar lower bound of $\sqrt{1 + \frac{10}{\kappa}} \frac{\zeta}{\rho^-}$ which we improve by a constant factor.

6.1.5 Runtime discussion

ARHT has the advantage of being very simple to implement in practice. The runtime of Algorithm 15 (ARHT) is comparable to that of the most efficient greedy algorithms (e.g. OMP/OMPR), as it requires a single function minimization per iteration. On the other hand, Algorithm 13 (Exhaustive Local Search) is less efficient, as it requires $O(n)$ function minimizations in each iteration, although in practice one might be able to speed it up by exploiting the fact that the problems being solved in each iteration are very closely related.

6.1.6 Naming Conventions

The algorithm that we call *Orthogonal Matching Pursuit (OMP)*, is also known as “Greedy” [127], “Fully Corrective Forward Greedy Selection” or just “Forward Selection”. What we call *Orthogonal Matching Pursuit with Replacement (OMPR)* [87] is also known by various other names. It is referenced in [146] as a simpler variant of their “Fully Corrective Forward Greedy Selection with Replacement” algorithm, or just Forward Selection with Replacement, or “Partial Hard Thresholding with parameter $r = 1$ (PHT(r) where $r = 1$)” [88] which is a generalization of Iterative Hard Thresholding. Finally, what we call *Exhaustive Local Search* is essentially a variant of “Orthogonal Least Squares” that includes replacement steps, and also appears in [146] as “Fully Corrective Forward Greedy Selection with Replacement”, or just “Forward Stepwise Selection with Replacement”. See also [20] regarding naming conventions.

Remark: Most of the results in the literature either only apply to, or are only presented for the Linear Regression problem. Since all of our results apply to general function minimization, we present them as such.

6.2 Preliminaries

We will make use of $\tilde{\kappa}_s = \rho_2^+ / \rho_s^-$ which is at most κ_s as long as $s \geq 2$. The following lemma stems from the definitions of ρ_2^+, ρ_1^+ and can be used to relate ρ_2^+ with ρ_1^+

Lemma 6.2.1. *For any function f that has the RSC property at sparsity level ≥ 2 and RSS constants ρ_1^+, ρ_2^+ at sparsity levels 1 and 2 respectively, we have $\rho_2^+ \leq 2\rho_1^+$.*

Proof. For any $x, y \in \mathbb{R}^n$ such that $|\text{supp}(y - x)| \leq 2$, We will prove that

$$f(y) \leq f(x) + \langle \nabla f(x), y - x \rangle + \frac{2\rho_1^+}{2} \|y - x\|_2^2.$$

Let $y = x + \alpha \vec{1}_i + \beta \vec{1}_j$ for some $i, j \in [n]$ and $\alpha, \beta \in \mathbb{R}$. We assume $i \neq j$ and since otherwise the claim already follows from RSS at sparsity level 1. We apply the RSS property with sparsity level 1 to get the inequalities

$$f(x + 2\alpha \vec{1}_i) \leq f(x) + 2\langle \nabla f(x), \alpha \vec{1}_i \rangle + 4\frac{\rho_1^+}{2} \|\alpha \vec{1}_i\|_2^2$$

and

$$f(x + 2\beta \vec{1}_j) \leq f(x) + 2\langle \nabla f(x), \beta \vec{1}_j \rangle + 4\frac{\rho_1^+}{2} \|\beta \vec{1}_j\|_2^2.$$

Now, by using convexity (more precisely restricted convexity at sparsity level 2 that is implied by RSC) we have

$$\begin{aligned} f(y) &= f(x + \alpha \vec{1}_i + \beta \vec{1}_j) \\ &\leq \frac{1}{2} \left(f(x + 2\alpha \vec{1}_i) + f(x + 2\beta \vec{1}_j) \right) \\ &\leq f(x) + \langle \nabla f(x), \alpha \vec{1}_i + \beta \vec{1}_j \rangle + \frac{2\rho_1^+}{2} \|\alpha \vec{1}_i + \beta \vec{1}_j\|_2^2 \\ &= f(x) + \langle \nabla f(x), y - x \rangle + \frac{2\rho_1^+}{2} \|y - x\|_2^2. \end{aligned}$$

□

Definition 6.2.2 (Restricted Gradient Optimal Constant (RGOCC)). *Given a differentiable function f and a “target” solution x^* , the Restricted Gradient Optimal Constant [170] at sparsity level s is the minimum $\zeta_s \in \mathbb{R}_+$ such that*

$$|\langle \nabla f(x^*), y \rangle| \leq \zeta_s \|y\|_2$$

for all s -sparse y . Setting $y = \nabla_S f(x^*)$ for some set S with $|S| \leq s$, this implies that $\zeta_s \geq \|\nabla_S f(x^*)\|$. An alternative definition is that ζ_s is the ℓ_2 norm of the s elements of $\nabla f(x^*)$ with highest absolute value.

Definition 6.2.3 (S -restricted minimizer). *Given $f : \mathbb{R}^n \rightarrow \mathbb{R}$, $x^* \in \mathbb{R}^n$, and $S \subseteq [n]$, we will call x^* an S -restricted minimizer of f if $\text{supp}(x^*) \subseteq S$ and for all x such that $\text{supp}(x) \subseteq S$ we have $f(x^*) \leq f(x)$.*

In Lemma 6.2.4 we state a standard martingale concentration inequality that we will use. See also [40] for more on martingales.

Lemma 6.2.4 (Martingale concentration inequality [40]). *Let $Y_0 = 0, Y_1, \dots, Y_n$ be a martingale with respect to the sequence i_1, \dots, i_n such that*

$$\text{Var}(Y_k \mid i_1, \dots, i_{k-1}) \leq \sigma^2$$

and

$$Y_{k-1} - Y_k \leq M,$$

for all $k \in [n]$, then for any $\lambda > 0$,

$$\Pr [Y_n \leq -\lambda] \leq e^{-\lambda^2 / (2(n\sigma^2 + M\lambda/3))}.$$

6.2.1 Algorithms

ℓ_1 optimization (LASSO)

The LASSO approach is to relax the ℓ_0 constraint into an ℓ_1 one, thus solving the following optimization problem:

$$\min_x f(x) + \lambda \|x\|_1, \quad (6.1)$$

for some parameter $\lambda > 0$.

Iterative Hard Thresholding (IHT):

[19] define the hard thresholding operator $H_r(x)$ as

$$[H_r(x)]_i = \begin{cases} x_i & \text{if } |x_i| \text{ is one of the } r \text{ entries of } x \\ & \text{with largest magnitude} \\ 0 & \text{otherwise} \end{cases}.$$

Using this, the algorithm is described in Algorithm 10.

Algorithm 10 Iterative Hard Thresholding (IHT)

```

1: function IHT( $s, T$ )
2:   function to be minimized  $f : \mathbb{R}^n \rightarrow \mathbb{R}$ 
3:   number of iterations  $T$ 
4:   output sparsity  $s$ 
5:    $S^0 \leftarrow \emptyset$ 
6:    $x^0 \leftarrow \vec{0}$ 
7:   for  $t = 0 \dots T - 1$  do
8:      $x^{t+1} \leftarrow H_s(x^t - \eta \nabla f(x^t))$ 
9:   return  $x^T$ 

```

Orthogonal Matching Pursuit (Greedy/OMP/Fwd stepwise selection)

The algorithm is described in Algorithm 11.

Orthogonal Matching Pursuit with Replacement (Local search/OMPR/Fwd stepwise selection with replacement steps)

The algorithm is described in Algorithm 12.

Exhaustive Local Search

The algorithm in this section is similar to OMPR, in that it iteratively inserts a new element in the support while removing one from it at the same time. While, as in OMPR, the element to be

Algorithm 11 Greedy/OMP/Fwd stepwise selection

```
1: function greedy( $s$ )
2:   function to be minimized  $f : \mathbb{R}^n \rightarrow \mathbb{R}$ 
3:   output sparsity  $s$ 
4:    $x^0 \leftarrow \vec{0}$ 
5:   for  $t = 0 \dots s - 1$  do
6:      $i \leftarrow \operatorname{argmax}\{|\nabla_i f(x^t)| \mid i \in [n] \setminus S^t\}$ 
7:      $S^{t+1} \leftarrow S^t \cup \{i\}$ 
8:      $x^{t+1} \leftarrow \operatorname{argmin}\{f(x) \mid \operatorname{supp}(x) \subseteq S^{t+1}\}$ 
9:   return  $x^s$ 
```

Algorithm 12 Orthogonal Matching Pursuit with Replacement

```
1: function OMPR( $s$ )
2:   function to be minimized  $f : \mathbb{R}^n \rightarrow \mathbb{R}$ 
3:   output sparsity  $s$ 
4:    $S^0 \leftarrow [s]$ 
5:    $x^0 \leftarrow \operatorname{argmin}\{f(x) \mid \operatorname{supp}(x) \subseteq S^0\}$ 
6:    $t \leftarrow 0$ 
7:   while true do
8:      $i \leftarrow \operatorname{argmax}\{|\nabla_i f(x^t)| \mid i \in [n] \setminus S^t\}$ 
9:      $j \leftarrow \operatorname{argmin}\{|x_j^t| \mid j \in S^t\}$ 
10:     $S^{t+1} \leftarrow S^t \cup \{i\} \setminus \{j\}$ 
11:     $x^{t+1} \leftarrow \operatorname{argmin}\{f(x) \mid \operatorname{supp}(x) \subseteq S^{t+1}\}$ 
12:    if  $f(x^{t+1}) \geq f(x^t)$  then
13:      break
14:     $t \leftarrow t + 1$ 
15:   $T \leftarrow t$ 
16:  return  $x^T$ 
```

removed is the minimum magnitude entry, the one to be inserted is chosen to be the one which results in the maximum decrease in the value of the objective. It is described in Algorithm 13.

Remark 6.2.5. *In the following sections, we will denote the minimization objective by f , the RSS and RSC parameters ρ_2^+ and $\rho_{s+s^*}^-$ by ρ^+ and ρ^- respectively, as well as $\kappa = \frac{\rho_{s+s^*}^+}{\rho_{s+s^*}^-}$ and $\tilde{\kappa} = \frac{\rho_2^+}{\rho_{s+s^*}^-}$. Note that the use of ρ_2^+ instead of ρ_1^+ used in some works is not restrictive. As shown in Lemma 6.2.1, $\rho_2^+ \leq 2\rho_1^+$ and so in all the bounds involving $\tilde{\kappa}$, it can be replaced by $2\frac{\rho_1^+}{\rho_{s+s^*}^-}$, thus only losing a factor of 2. Furthermore, we state our results in terms of $\tilde{\kappa}$ as opposed to κ . This is always more general since $\tilde{\kappa} \leq \kappa$.*

In order to simplify the exposition, we will assume that $\min_x f(x) = 0$. This property can be ensured by adding a constant to f .

When no additional context is provided, we denote current solution by x and the target solution x^ , with respective support sets S and S^* and sparsities $s = |S|$ and $s^* = |S^*|$.*

Algorithm 13 Exhaustive Local Search

```
1: function to be minimized  $f : \mathbb{R}^n \rightarrow \mathbb{R}$ 
2: target sparsity  $s$ 
3: number of iterations  $T$ 
4:  $S^0 \leftarrow [s]$ 
5:  $x^0 \leftarrow \operatorname{argmin}\{f(x) \mid \operatorname{supp}(x) \subseteq S^0\}$ 
6: for  $t = 0 \dots T - 1$  do
7:    $j \leftarrow \operatorname{argmin}_{j \in S^t} x_j^2$ 
8:    $i \leftarrow \operatorname{argmin}_{i \in [n] \setminus S^t} \left\{ \min_{x : \operatorname{supp}(x) \subseteq S^t \cup \{i\} \setminus \{j\}} f(x) \right\}$ 
9:    $S^{t+1} \leftarrow S^t \cup \{i\} \setminus \{j\}$ 
10:   $x^{t+1} \leftarrow \operatorname{argmin}\{f(x) \mid \operatorname{supp}(x) \subseteq S^{t+1}\}$ 
11:  if  $f(x^{t+1}) \geq f(x^t)$  then
12:    return  $x^t$ 
13: return  $x^T$ 
```

6.3 Adaptively Regularized Hard Thresholding (ARHT)

6.3.1 Overview and Main Theorem

Our algorithm is essentially a hard thresholding algorithm (and more specifically OMPR, also known as PHT(1)) with the crucial novelty that it is applied on an adaptively regularized objective function. Hard thresholding algorithms maintain a solution x supported on $S \subseteq [n]$, which they iteratively update by inserting new elements into the support set S and removing the same number of elements from it, in order to preserve the sparsity of x . More specifically, OMPR makes one insertion and one removal in each iteration. In order to evaluate the element i to be *inserted* into S , OMPR uses the fact that, because of smoothness, $\frac{(\nabla_i f(x))^2}{2\rho_2^+}$ is a lower bound on the decrease of $f(x)$ caused by inserting i into the support, and therefore picks i to maximize $|\nabla_i f(x)|$. Similarly, in order to evaluate the element j to be *removed* from S , OMPR uses the fact that $\frac{\rho_2^+}{2} x_j^2$ upper bounds the increase of $f(x)$ caused by setting $x_j = 0$, and therefore picks j to minimize $|x_j|$. However, the real worth of j might be as small as $\frac{\rho_2^-}{2} x_j^2$, so the upper bound can be loose by a factor of $\frac{\rho_2^+}{\rho_2^-} \geq \frac{\rho_2^+}{\rho_{s+s^*}^-} = \tilde{\kappa}$.

We eliminate this discrepancy by running the algorithm on the regularized function $g(z) := f(z) + \frac{\rho_2^+}{2} \|z\|_2^2$. As the restricted condition number of g is now $O(1)$, the real worth of a removal candidate j matches the upper bound up to a constant factor.

However, even though g is now well conditioned, the analysis can only guarantee the quality of the solution in terms of the *original* objective f if the regularization is *not* applied on elements S^* , i.e. $\frac{\rho_2^+}{2} \|x_{R \setminus S^*}\|_2^2$ for some sufficiently large $R \subseteq [n]$; if this is the case, a solution with sparsity $O(s^* \tilde{\kappa})$ can be recovered. Unfortunately, there is no way of knowing a priori which elements not to regularize, as this is equivalent to finding the target solution. As a result, the algorithm can get trapped in local minima, which are defined as states in which one iteration of the algorithm does not decrease $g(x)$, even though x is a suboptimal solution in terms of f (i.e. $f(x) > f(x^*)$).

The main contribution of this work is to characterize such local minima and devise a procedure that is able to successfully escape them, thus allowing x to converge to a desired solution for the original objective.

The core algorithm is presented in Algorithm 14. The full algorithm additionally requires some

standard routines like binary search and is presented in Algorithm 15.

Algorithm 14 Adaptively Regularized Hard Thresholding core routine

```

1: function ARHT_CORE( $s, \widehat{f}^*, \varepsilon$ )
2:   function to be minimized  $f : \mathbb{R}^n \rightarrow \mathbb{R}$ 
3:   target sparsity  $s$ 
4:   target value  $\widehat{f}^*$  (current guess for the optimal value)
5:   target error  $\varepsilon$ 
6:   Define  $g_R(x) := f(x) + \frac{\rho_2^+}{2} \|x_R\|_2^2$  for all  $R \subseteq [n]$ .
7:    $R^0 \leftarrow [n]$ 
8:    $S^0 \leftarrow [s]$ 
9:    $x^0 \leftarrow \operatorname{argmin}_{\operatorname{supp}(x) \subseteq S^0} g_{R^0}(x)$ 
10:   $T = 2s \log \frac{f(\vec{0}) - \min_x f(x)}{\varepsilon}$  (number of iterations)
11:  for  $t = 0 \dots T - 1$  do
12:    if  $\min_{\operatorname{supp}(x) \subseteq S^t} f(x) \leq \widehat{f}^*$  then
13:      return  $\operatorname{argmin}_{\operatorname{supp}(x) \subseteq S^t} f(x)$ 
14:     $i \leftarrow \operatorname{argmax}_{i \in [n]} |\nabla_i g_{R^t}(x^t)|$ 
15:     $j \leftarrow \operatorname{argmin}_{j \in S^t} |x_j|$ 
16:     $S^{t+1} \leftarrow S^t \cup \{i\} \setminus \{j\}$ 
17:     $x^{t+1} \leftarrow \operatorname{argmin}_{\operatorname{supp}(x) \subseteq S^{t+1}} g_{R^t}(x)$ 
18:    if  $g_{R^t}(x^t) - g_{R^t}(x^{t+1}) < \frac{1}{s} (g_{R^t}(x^t) - \widehat{f}^*)$  then
19:       $S^{t+1} \leftarrow S^t$ 
20:      Sample  $i \in R^t$  proportional to  $(x_i^t)^2$ 
21:       $R^{t+1} \leftarrow R^t \setminus \{i\}$ 
22:       $x^{t+1} \leftarrow \operatorname{argmin}_{\operatorname{supp}(x) \subseteq S^{t+1}} g_{R^{t+1}}(x)$ 
23:  return  $x^T$ 

```

In the following, we will let \widehat{f}^* denote a guess on the target value $f(x^*)$. Also, x^0 will denote the initial solution, which is an S^0 -restricted minimizer for an arbitrary set $S^0 \subseteq [n]$ with $|S^0| = s$. In Algorithm 14, S^0 is defined explicitly as $[s]$, however in practice one might want to pick a better initial set (e.g. returned by running OMP).

It should be noted that even though the value ρ_2^+ is used by the algorithm to define the regularizer, exact knowledge of ρ_2^+ is not required, and an upper bound U on it can be used. Of course, the final sparsity and runtime bound will then depend on $U/\rho_{s+s^*}^-$ instead of $\rho_2^+/\rho_{s+s^*}^-$. One such upper bound is $2\rho_1^+$. For linear and logistic regression where A is the $(\# \text{ examples}) \times (\# \text{ features})$ data matrix, ρ_1^+ is the maximum ℓ_2 norm of a column of the data matrix A (so $\rho_1^+ = 1$ if the columns of A are normalized). More generally, getting such a bound would depend on the specific function we are trying to minimize. For example, if we are trying to minimize $\sum_{i=1}^m L(Ax)_i$ where A is a data matrix with normalized columns and $L : \mathbb{R} \rightarrow \mathbb{R}$ is a 1-smooth loss function, then $\rho_1^+ \leq 1$ so 2 is a good upper bound for ρ_2^+ . Another option is to run the algorithm for different candidate values for

Algorithm 15 Adaptively Regularized Hard Thresholding

```

1: function ARHT_ROBUST( $s, \widehat{f}^*, \varepsilon$ )
2:   function to be minimized  $f : \mathbb{R}^n \rightarrow \mathbb{R}$ 
3:    $x^{\text{ret}} \leftarrow \vec{0}$ 
4:   for  $z = 1 \dots 5 \log \left( 6n \log \frac{f(\vec{0})}{\varepsilon} \right)$  do
5:      $x \leftarrow \text{ARHT\_core}(s, \widehat{f}^*, \varepsilon)$ 
6:     if  $f(x) < f(x^{\text{ret}})$  then
7:        $x^{\text{ret}} \leftarrow x$ 
8:   return  $x^{\text{ret}}$ 
9: function ARHT( $s, \varepsilon$ )
10:  function to be minimized  $f : \mathbb{R}^n \rightarrow \mathbb{R}$ 
11:  target sparsity  $s$ 
12:  target error  $\varepsilon$ 
13:   $l \leftarrow \vec{0}$ 
14:   $r \leftarrow f(\vec{0})$ 
15:   $b \leftarrow \vec{0}$ 
16:  while  $r - l > \varepsilon$  do
17:     $m \leftarrow \frac{l+r}{2}$ 
18:     $x \leftarrow \text{ARHT\_robust}(s, m, \varepsilon/3)$ 
19:    if  $f(x) > m + \varepsilon/3$  then
20:       $l \leftarrow m$ 
21:    else
22:       $b \leftarrow x$ 
23:       $r \leftarrow f(x)$ 
24:  return  $b$ 

```

ρ_2^+ until we get the desired performance.

We are now ready for the main result of this section. It basically states that for any solution x^* with sparsity s^* , a solution x with $f(x) \leq f(x^*) + \varepsilon$ and sparsity $O(\kappa s^*)$ can be recovered by Algorithm 15. In comparison, previously known analyses could guarantee either a sparsity $O(\kappa s^* \log \frac{f(\vec{0})}{\varepsilon})$ (OMP) or $O(\kappa^2 s^*)$ (OMPR, IHT, LASSO). It is useful to note here that, because of the constant factor in front of κs^* (which is less than 20), this result cannot be directly used to obtain compressed sensing results with sparsity close to s^* , but is instead useful in the asymptotic regime (with relaxed sparsity sgs^*).

Theorem 6.3.1. *Given a function f and an (unknown) s^* -sparse solution x^* , with probability at least $1 - \frac{1}{n}$ Algorithm 15 returns an s -sparse solution x with $f(x) \leq f(x^*) + \varepsilon$, as long as $s \geq s^* \max\{4\tilde{\kappa} + 7, 12\tilde{\kappa} + 6\}$. The number of iterations is $O\left(s \log^2 \frac{f(\vec{0})}{\varepsilon} \log\left(n \log \frac{f(\vec{0})}{\varepsilon}\right)\right)$.*

The following corollary that bounds the total runtime can be immediately extracted. Note that in practice the total runtime heavily depends on the choice of f , and it can often be improved for various special cases (e.g. linear regression).

Corollary 6.3.2 (Theorem 6.3.1 runtime). *If we denote by G the time needed to compute ∇f and by M the time to minimize f in a restricted subset of $[n]$ of size s , the total runtime of Algorithm 15 is $O\left((G + M)s \log^2 \frac{f(\vec{0})}{\varepsilon} \log\left(n \log \frac{f(\vec{0})}{\varepsilon}\right)\right)$. If gradient descent is used for the implementation of the*

inner optimization problem, then $M = O\left(G\tilde{\kappa} \log \frac{f(\vec{0})}{\varepsilon}\right)$ and so the total runtime can be bounded by $O\left(Gs\tilde{\kappa} \log^3 \frac{f(\vec{0})}{\varepsilon} \log\left(n \log \frac{f(\vec{0})}{\varepsilon}\right)\right)$.

Before proving the above theorem, we provide the main components that are needed for its proof. It is important to split the iterations of Algorithm 14 into two categories: Those that make enough progress, i.e. for which the condition in Line 19 of Algorithm 14 is *false*, and those that don't, i.e. for which the condition in Line 19 is *true*. We call the former *Type 1* iterations and the latter *Type 2* iterations. Intuitively, Type 1 iterations signify that $g(x)$ is decreasing at a sufficient rate to achieve the desired convergence, while Type 2 iterations indicate a local minimum that should be dealt with. Our argument consists of two steps: Showing that as long as there are enough Type 1 iterations, a desired solution will be obtained (Lemma 6.3.3), and bounding the total number of Type 2 iterations with constant probability (Lemma 6.3.4).

Lemma 6.3.3 (Convergence rate). *If Algorithm 14 executes at least $T_1 = s \log \frac{g(x^0) - \hat{f}^*}{\varepsilon}$ Type 1 iterations, then $f(x^T) \leq \hat{f}^* + \varepsilon$.*

The proof of this lemma can be found in Appendix 9.4.1.

Lemma 6.3.4 (Bounding Type 2 iterations). *If $s \geq s^* \max\{4\tilde{\kappa} + 7, 12\tilde{\kappa} + 6\}$ and $\hat{f}^* \geq f(x^*)$, then with probability at least 0.2 the number of Type 2 iterations is at most $(s^* - 1)(4\tilde{\kappa} + 6)$.*

The proof of this lemma appears in Section 6.3.2. These lemmas can now be directly used to obtain the following lemma, which states the performance guarantee of the ARHT core routine (Algorithm 14).

Lemma 6.3.5 (Algorithm 14 guarantee). *If $s \geq s^* \max\{4\tilde{\kappa} + 7, 12\tilde{\kappa} + 6\}$ and $\hat{f}^* \geq f(x^*)$, with probability at least 0.2 $\text{ARHT_core}(s, \hat{f}^*, \varepsilon)$ returns an s -sparse solution x such that $f(x) \leq \hat{f}^* + \varepsilon$.*

Proof. By Lemma 6.3.4, with probability at least 0.2 there will be at most $(s^* - 1)(4\tilde{\kappa} + 6)$ Type 2 iterations. This means that the number of Type 1 iterations is at least

$$T - (s^* - 1)(4\tilde{\kappa} + 6) \geq s \log \frac{f(\vec{0})}{\varepsilon} \geq s \log \frac{g^0(x^0) - \hat{f}^*}{\varepsilon},$$

where the latter inequality follows from the fact that $f(\vec{0}) = g^0(\vec{0}) \geq g^0(x^0)$ and $\hat{f}^* \geq f(x^*) \geq 0$. Lemma 6.3.3 then implies that $f(x^T) \leq \hat{f}^* + \varepsilon$. \square

In other words, as long as $\hat{f}^* \geq f(x^*)$, a solution of value $\leq \hat{f}^* + \varepsilon$ will be found. As the value $f(x^*)$ is not known a priori, \hat{f}^* is just an estimate for it. We perform binary search over \hat{f}^* , as described in Algorithm 15. The reason we need this estimate of $f(x^*)$ is line 19 of Algorithm 14. As our algorithm is randomized, we use this estimate to decide whether there was enough progress in one iteration (Type 1 iteration), or not (Type 2 iteration, in which case we perform the randomization step). If \hat{f}^* is smaller than $f(x^*)$, the algorithm might mistake a Type 1 iteration for a Type 2 iteration, thus performing the randomization step even though the algorithm makes enough progress. On the other hand, if \hat{f}^* is much larger than $f(x^*)$, the algorithm might terminate with a suboptimal solution of value much larger than $f(x^*)$.

The probability of success in the previous lemma can be boosted by repeating multiple times. Combining these arguments will lead us to the proof of Theorem 6.3.1. First, we turn the result of Lemma 6.3.5 into a high probability result by repeating multiple times:

Lemma 6.3.6. *If $s \geq s^* \max\{4\tilde{\kappa} + 7, 12\tilde{\kappa} + 6\}$ and $\hat{f}^* \geq f(x^*)$, $\text{ARHT_robust}(s, \hat{f}^*, \varepsilon)$ returns an s -sparse solution x such that $f(x) \leq \hat{f}^* + \varepsilon$ with probability at least $1 - \frac{1}{6n \log \frac{f(\bar{0})}{\varepsilon}}$.*

Proof. From Lemma 6.3.5, the probability that a given call to ARHT_core fails is at most 0.8. Since this random experiment is executed $5 \log \left(6n \log \frac{f(\bar{0})}{\varepsilon}\right)$ times independently, the probability that it never succeeds is at most $(0.8)^{5 \log \left(6n \log \frac{f(\bar{0})}{\varepsilon}\right)} < \frac{1}{6n \log \frac{f(\bar{0})}{\varepsilon}}$, therefore the statement follows. \square

Lemma 6.3.7. *If $s \geq s^* \max\{4\tilde{\kappa} + 7, 12\tilde{\kappa} + 6\}$, $\text{ARHT}(s, \varepsilon)$ (in Algorithm 15) returns an s -sparse solution x such that $f(x) \leq f(x^*) + \varepsilon$. The algorithm succeeds with probability at least $1 - \frac{1}{n}$. and the number of calls to ARHT_robust is $\leq 6 \log \frac{f(\bar{0})}{\varepsilon}$.*

Proof. First we will bound the number of calls to ARHT_robust . Let L_k be the equal to $r - l$ before the k -th iteration in Line 21 of Algorithm 15. Then either $L_{k+1} = L_k/2$ (Line 25) or $L_{k+1} \leq L_k/2 + \varepsilon/3 < 5L_k/6$ (Line 28). Therefore in any case we have $L_{k+1} < 5L_k/6$ which implies that after $T = 6 \log \frac{f(\bar{0})}{\varepsilon}$ iterations we will have $r - l \leq \varepsilon$.

Now let us compute the probability that all the calls to ARHT_robust are successful. The number of such calls is at most $6 \log \frac{f(\bar{0})}{\varepsilon}$ and we know each one of them independently fails with probability less than $\frac{1}{6n \log \frac{f(\bar{0})}{\varepsilon}}$, so by a union bound the probability that at least one call fails is less than $\frac{1}{n}$.

To prove correctness, note that by Lemma 6.3.6, for each $r \geq f(x^*)$ we have $f(\text{ARHT_robust}(s, r, \varepsilon/3)) \leq r + \varepsilon/3$. After Line 20 of Algorithm 15, we will have $l = 0 \leq f(x^*)$. In the while construct, it is always true that $f(x^*) \geq l$. This is initially true, as we saw. For each m chosen in Line 22 and x in Line 23, note that if $f(x) > m + \varepsilon/3$, then by Lemma 6.3.6 $f(x^*) > m$ and so the invariant that $f(x^*) \geq l$ stays true. On the other hand, it is always true that $f(b) \leq r$. Initially this is so because $f(\bar{0}) = r$, and when we decrease r to some $f(x)$ we also update $b = x$. This implies that in the end of the algorithm the returned solution will have the required property, since we will have $f(b) \leq r \leq l + \varepsilon \leq f(x^*) + \varepsilon$. \square

The proof Theorem 6.3.1 now easily follows.

Proof of Theorem 6.3.1. Lemma 6.3.7 already establishes the correctness of the algorithm with probability at least $1 - \frac{1}{n}$. For the runtime, note that ARHT_core takes $O\left(s \log \frac{f(\bar{0})}{\varepsilon}\right)$ iterations, ARHT_robust takes $O\left(\log\left(n \log \frac{f(\bar{0})}{\varepsilon}\right)\right)$ iterations, and ARHT takes $O\left(\log \frac{f(\bar{0})}{\varepsilon}\right)$ iterations. In conclusion, the total number of iterations is $O\left(s \log^2 \frac{f(\bar{0})}{\varepsilon} \log\left(n \log \frac{f(\bar{0})}{\varepsilon}\right)\right)$, each of which requires a constant number of minimizations of f . \blacksquare

6.3.2 Bounding Type 2 Iterations

When x has significant ℓ_2^2 mass in the target support, the regularization term $\frac{\rho_2^+}{2} \|x\|_2^2$ might penalize the target solution too much, leading to a Type 2 iteration. In this case, we use random sampling to detect an element in the optimal support and unregularize it. This procedure escapes all local minima, thus leading to a bound in the total number of Type 2 iterations.

More concretely, we show that if at some iteration of the algorithm the value of $g(x)$ does not decrease sufficiently (Type 2 iteration), then roughly at least a $\frac{1}{\kappa}$ -fraction of the ℓ_2^2 mass of x lies

in the target support S^* . We exploit this property by sampling an element i proportional to x_i^2 and removing its corresponding term from the regularizer (*unregularizing* it). We show that with constant probability this will happen at most $O(s^*\tilde{\kappa})$ times, as after that all the elements in S^* will have been unregularized.

When referring to the t -th iteration of Algorithm 14, we let x^t be the current solution with support set S^t and $R^t \subseteq [n]$ the current regularization set as defined in the algorithm. For ease of notation, we will drop the subscript of the regularizer, i.e. $\Phi^t(z) := \frac{\rho_2^+}{2} \|z_{R^t}\|_2^2$ and of the regularized function, i.e. $g^t(z) := f(z) + \Phi^t(z)$. Note that by definition of the algorithm x^t is an S^t -restricted minimizer of g^t .

Let $(\rho_2^+)'$ and $(\rho_{s+s^*}^-)'$ be RSS and RSC parameters of g^t . We start with a lemma that relates $(\rho_2^+)'$ to ρ_2^+ and $(\rho_{s+s^*}^-)'$ to $\rho_{s+s^*}^-$, and is proved in Appendix 9.4.1.

Lemma 6.3.8 (RSC, RSS of regularized function). $(\rho_2^+)' \leq 2\rho_2^+$ and $(\rho_{s+s^*}^-)' \geq \rho_{s+s^*}^-$

This states that the restricted smoothness and strong convexity constants of the regularized function are always within a constant factor of those of the original function, and thus we can make our statements in terms of the RSC, RSS of the original function. Next, we present a lemma that establishes a lower bound on the progress $g^t(x^t) - g^{t+1}(x^{t+1})$ in one iteration. This will be helpful in order to diagnose the cause of having insufficient progress in one iteration.

Lemma 6.3.9 (ARHT Progress Lemma). *If $\hat{f}^* \geq f(x^*)$, for the progress $g^t(x^t) - g^t(x^{t+1})$ in Line 19 of Algorithm 14 it holds that*

$$\begin{aligned} & g^t(x^t) - g^t(x^{t+1}) \\ & \geq \frac{\rho^-}{2|S^* \setminus S^t| \rho^+} \left(f(x^t) - f(x^*) + \langle \nabla_{S^t \setminus S^*} \Phi^t(x^t), x_{S^t \setminus S^*}^t \rangle - \frac{1}{2\rho^-} \|\nabla_{S^t \cap S^*} \Phi^t(x^t)\|_2^2 \right) - \rho^+(x_j^t)^2. \end{aligned}$$

Proof. The proof will proceed as follows: We first use the smoothness of g^t to get a lower bound on the progress in one step, $g^t(x^t) - g^t(x^{t+1})$. This lower bound will depend on $\|\nabla_{S^* \setminus S^t} g^t(x^t)\|_2^2$, which is the norm of the gradient of g^t restricted to the set $S^* \setminus S^t$, as well as $(x_j^t)^2$, where j is the position of the minimum-magnitude entry of x^t . Then, we use the strong convexity of g^t to relate $\|\nabla_{S^* \setminus S^t} g^t(x^t)\|_2^2$ to the difference in function value $f(x^t) - f(x^*)$, plus some terms that come from the regularizer.

First of all, since the condition in Line 12 (“if $\min_{\text{supp}(x) \subseteq S^t} f(x) \leq \hat{f}^*$ ”) was not triggered, we have that $\min_{\text{supp}(x) \subseteq S^t} f(x) > \hat{f}^* \geq f(x^*)$ and so $S^* \setminus S^t \neq \emptyset$. By Lemma 6.3.8 we have that $(\rho^+)' \leq 2\rho^+$, therefore the decrease in g^t that is achieved is

$$\begin{aligned} & g^t(x^t) - g^t(x^{t+1}) \\ & \geq \max_{\eta \in \mathbb{R}} \left\{ g^t(x^t) - g^t(x^t + \eta \vec{1}_i - x_j^t \vec{1}_j) \right\} \\ & \geq \max_{\eta \in \mathbb{R}} \left\{ -\langle \nabla g^t(x^t), \eta \vec{1}_i - x_j^t \vec{1}_j \rangle - \rho^+ \eta^2 - \rho^+(x_j^t)^2 \right\} := B. \end{aligned}$$

Note that, as defined by the algorithm, x^t is an S^t -restricted minimizer of g^t and since $j \in S^t$, we

have $\nabla_j g^t(x^t) = 0$. Therefore

$$\begin{aligned}
B &= \max_{\eta \in \mathbb{R}} \{-\langle \nabla g^t(x^t), \eta \vec{1}_i \rangle - \rho^+ \eta^2 - \rho^+(x_j^t)^2\} \\
&= \frac{[\nabla_i g^t(x^t)]^2}{4\rho^+} - \rho^+(x_j^t)^2 \\
&\geq \max_{k \in S^* \setminus S} \frac{[\nabla_k g^t(x^t)]^2}{4\rho^+} - \rho^+(x_j^t)^2 \\
&\geq \frac{\|\nabla_{S^* \setminus S^t} g^t(x^t)\|_2^2}{4|S^* \setminus S^t| \rho^+} - \rho^+(x_j^t)^2,
\end{aligned} \tag{6.2}$$

where we used the fact that i was picked to maximize $|\nabla_k g^t(x^t)|$. Now we would like to relate this to $g^t(x^t) - f(x^*)$ (and not $g^t(x^t) - g^t(x^*)$). By applying the Restricted Strong Convexity property,

$$\begin{aligned}
f(x^*) - f(x^t) &\geq \langle \nabla f(x^t), x^* - x^t \rangle + \frac{\rho^-}{2} \|x^t - x^*\|_2^2 \\
&\geq \langle \nabla f(x^t), x^* - x^t \rangle + \frac{\rho^-}{2} \|x_{S^* \setminus S^t}^*\|_2^2 + \frac{\rho^-}{2} \|(x^t - x^*)_{S^t \cap S^*}\|_2^2.
\end{aligned}$$

Now note that $f(x^t) = g^t(x^t) - \Phi^t(x^t)$, $\nabla_{S^t} g^t(x^t) = \vec{0}$ (since x^t is an S^t -restricted minimizer of g^t), and $\nabla \Phi^t(x^t) = \nabla_{S^t} \Phi^t(x^t)$ therefore

$$\begin{aligned}
&\langle \nabla f(x^t), x^* - x^t \rangle \\
&= \langle \nabla g^t(x^t), x^* - x^t \rangle - \langle \nabla \Phi^t(x^t), x^* - x^t \rangle \\
&= \langle \nabla_{S^* \setminus S^t} g^t(x^t), x_{S^* \setminus S^t}^* \rangle + \langle \nabla_{S^t \setminus S^*} \Phi^t(x^t), x_{S^t \setminus S^*}^t \rangle + \langle \nabla_{S^t \cap S^*} \Phi^t(x^t), (x^t - x^*)_{S^t \cap S^*} \rangle.
\end{aligned}$$

Plugging this into the previous inequality, we get

$$\begin{aligned}
f(x^*) - f(x^t) &\geq \langle \nabla_{S^* \setminus S^t} g^t(x^t), x_{S^* \setminus S^t}^* \rangle + \frac{\rho^-}{2} \|x_{S^* \setminus S^t}^*\|_2^2 + \langle \nabla_{S^t \setminus S^*} \Phi^t(x^t), x_{S^t \setminus S^*}^t \rangle \\
&\quad + \langle \nabla_{S^t \cap S^*} \Phi^t(x^t), (x^t - x^*)_{S^t \cap S^*} \rangle + \frac{\rho^-}{2} \|(x^t - x^*)_{S^t \cap S^*}\|_2^2 \\
&\geq -\frac{1}{2\rho^-} \|\nabla_{S^* \setminus S^t} g^t(x^t)\|_2^2 + \langle \nabla_{S^t \setminus S^*} \Phi^t(x^t), x_{S^t \setminus S^*}^t \rangle - \frac{1}{2\rho^-} \|\nabla_{S^t \cap S^*} \Phi^t(x^t)\|_2^2,
\end{aligned}$$

where we twice used the inequality $\langle u, v \rangle + \frac{\lambda}{2} \|v\|_2^2 \geq -\frac{1}{2\lambda} \|u\|_2^2$ for any $\lambda > 0$. This inequality is derived by expanding $\frac{1}{2} \left\| \frac{1}{\sqrt{\lambda}} u + \sqrt{\lambda} v \right\|_2^2 \geq 0$. So plugging in $\|\nabla_{S^* \setminus S^t} g^t(x^t)\|_2^2$ into (6.2),

$$\begin{aligned}
B &\geq \frac{\rho^-}{2|S^* \setminus S^t| \rho^+} \left(f(x^t) - f(x^*) + \langle \nabla_{S^t \setminus S^*} \Phi^t(x^t), x_{S^t \setminus S^*}^t \rangle - \frac{1}{2\rho^-} \|\nabla_{S^t \cap S^*} \Phi^t(x^t)\|_2^2 \right) - \rho^+(x_j^t)^2.
\end{aligned}$$

□

Let $R \subseteq [n]$ be the set of currently regularized elements. The following invariant is a crucial

ingredient for bringing the sparsity from $O(s^*\tilde{\kappa}^2)$ down to $O(s^*\tilde{\kappa})$, and we intend to enforce it at all times. It essentially states that there will always be enough elements in the current solution that are being regularized.

Invariant 6.3.10.

$$|R \cap S| \geq s^* \max\{1, 8\tilde{\kappa}\}$$

To give some intuition on this, ARHT owes its improved $\tilde{\kappa}$ dependence on the regularizer $\frac{\rho^+}{2} \|x\|_2^2$. However, during the algorithm, some elements are being unregularized. Our analysis requires that the current solution support always contains $\Omega(s^*\tilde{\kappa})$ regularized elements, which is what Invariant 6.3.10 states.

We can now proceed to show that, with constant probability, Algorithm 14 will only have $O(s^*\tilde{\kappa})$ Type 2 iterations, which is the goal of this section.

Proof of Lemma 6.3.4. The idea of the proof is to use the progress bound in Lemma 6.3.9 to obtain a necessary condition under which the progress (i.e. the decrease of $g^t(x^t)$) is not sufficient in one iteration (thus, we have a Type 2 iteration). For our choice of regularizer, this condition implies that at least an $\Omega(\frac{1}{\tilde{\kappa}})$ fraction of the ℓ_2^2 mass of x^t lies in the optimal support set S^* , which means that we can find an element in S^* with decent probability by appropriately sampling an element of x^t . We finally apply a probabilistic analysis over all iterations, to show that if each sampled element is unregularized, with constant probability the total number of Type 2 iterations cannot exceed $\Theta(\tilde{\kappa}s^*)$.

We first observe some useful properties of our regularizer, which can be verified by simple substitution. The definition of $\Phi^t(x^t)$ implies that

$$\begin{aligned} \langle \nabla_{S^t \setminus S^*} \Phi^t(x^t), x_{S^t \setminus S^*}^t \rangle &= \rho^+ \langle x_{(R^t \cap S^t) \setminus S^*}^t, x_{S^t \setminus S^*}^t \rangle \\ &= \rho^+ \sum_{i \in (R^t \cap S^t) \setminus S^*} x_i^2 \\ &= \rho^+ \sum_{i \in R^t \setminus S^*} x_i^2 \\ &= \rho^+ \left\| x_{R^t \setminus S^*}^t \right\|_2^2, \end{aligned} \tag{6.3}$$

where the second-to-last equality follows because x^t is 0 outside of S^t . For the same reason, we also have

$$\left\| \nabla_{S^t \cap S^*} \Phi^t(x^t) \right\|_2^2 = (\rho^+)^2 \left\| x_{R^t \cap S^*}^t \right\|_2^2. \tag{6.4}$$

By combining (6.3) and (6.4) we get that

$$\Phi^t(x^t) = \frac{1}{2} \langle \nabla_{S^t \setminus S^*} \Phi^t(x^t), x_{S^t \setminus S^*}^t \rangle + \frac{1}{2\rho^+} \left\| \nabla_{S^t \cap S^*} \Phi^t(x^t) \right\|_2^2. \tag{6.5}$$

Equations (6.3), (6.4), and (6.5) will be used later on. Now, before the first iteration we have $|R^0 \cap S^0| = |S^0| = s$. Since in each Type 2 iteration we have $|R^{t+1}| = |R^t| - 1$,

$$|R^t \cap S^t| \geq s - [\text{number of Type 2 iterations up to } t].$$

This implies that for the first $(s^* - 1)(4\tilde{\kappa} + 6)$ Type 2 iterations,

$$|R^t \cap S^t| \geq s - (s^* - 1)(4\tilde{\kappa} + 6) \geq s^* \max\{1, 8\tilde{\kappa}\}, \quad (6.6)$$

since $s \geq s^* \max\{4\tilde{\kappa} + 7, 12\tilde{\kappa} + 6\}$. From this it follows that

$$\begin{aligned} |(R^t \cap S^t) \setminus S^*| &= |R^t \cap S^t| - |R^t \cap S^t \cap S^*| \\ &\geq s^* \max\{1, 8\tilde{\kappa}\} - |S^t \cap S^*| \\ &\geq |S^* \setminus S^t| 8\tilde{\kappa} \\ &= |S^* \setminus S^t| 8 \frac{\rho^+}{\rho^-} \end{aligned}$$

and so

$$\begin{aligned} (x_j^t)^2 &\leq \frac{1}{|(R^t \cap S^t) \setminus S^*|} \left\| x_{(R^t \cap S^t) \setminus S^*}^t \right\|_2^2 \\ &\leq \frac{\rho^-}{8|S^* \setminus S^t| \rho^+} \left\| x_{R^t \setminus S^*}^t \right\|_2^2 \\ &= \frac{\rho^-}{8|S^* \setminus S^t| (\rho^+)^2} \langle \nabla_{S^t \setminus S^*} \Phi^t(x^t), x_{S^t \setminus S^*}^t \rangle, \end{aligned}$$

where $j \in S^t$ is the element that the algorithm removes from S^t , and we used (6.3). Combining this inequality with the statement of Lemma 6.3.9 we have

$$\begin{aligned} &g^t(x^t) - g^t(x^{t+1}) \\ &\geq \frac{\rho^-}{2|S^* \setminus S^t| \rho^+} \left(f(x^t) - f(x^*) + \langle \nabla_{S^t \setminus S^*} \Phi^t(x^t), x_{S^t \setminus S^*}^t \rangle - \frac{1}{2\rho^-} \|\nabla_{S^t \cap S^*} \Phi^t(x^t)\|_2^2 \right) - \rho^+ (x_j^t)^2 \quad (6.7) \\ &\geq \frac{\rho^-}{2|S^* \setminus S^t| \rho^+} \left(f(x^t) - f(x^*) + \frac{3}{4} \langle \nabla_{S^t \setminus S^*} \Phi^t(x^t), x_{S^t \setminus S^*}^t \rangle - \frac{1}{2\rho^-} \|\nabla_{S^t \cap S^*} \Phi^t(x^t)\|_2^2 \right). \end{aligned}$$

By definition of a Type 2 iteration,

$$\begin{aligned} g^t(x^t) - g^t(x^{t+1}) &< \frac{1}{s} \left(g^t(x^t) - \hat{f}^* \right) \\ &\leq \frac{\rho^-}{2|S^* \setminus S^t| \rho^+} (g^t(x^t) - f(x^*)) \quad , \quad (6.8) \\ &= \frac{\rho^-}{2|S^* \setminus S^t| \rho^+} (f(x^t) - f(x^*) + \Phi^t(x^t)) \end{aligned}$$

where we used the fact that $s \geq 2s^*\tilde{\kappa} \geq 2|S^* \setminus S^t|\tilde{\kappa}$ and $f(x^*) \leq \hat{f}^*$. Combining inequalities (6.7) and (6.8) we get

$$\Phi^t(x^t) > \frac{3}{4} \langle \nabla_{S^t \setminus S^*} \Phi^t(x^t), x_{S^t \setminus S^*}^t \rangle - \frac{1}{2\rho^-} \|\nabla_{S^t \cap S^*} \Phi^t(x^t)\|_2^2,$$

or equivalently, by replacing $\Phi^t(x^t)$ from (6.5),

$$\frac{1}{2} \left(\frac{1}{\rho^-} + \frac{1}{\rho^+} \right) \|\nabla_{S^t \cap S^*} \Phi^t(x^t)\|_2^2 > \frac{1}{4} \langle \nabla_{S^t \setminus S^*} \Phi^t(x^t), x_{S^t \setminus S^*}^t \rangle.$$

Further applying (6.3) and (6.4), we equivalently get

$$2(1 + \tilde{\kappa}) \|x_{R^t \cap S^*}^t\|_2^2 > \|x_{R^t \setminus S^*}^t\|_2^2. \quad (6.9)$$

Now, note that in Lines 21-22 the algorithm picks an element $i \in R^t$ with probability proportional to $(x_i^t)^2$ and unregularizes it, i.e. sets $R^{t+1} \leftarrow R^t \setminus \{i\}$. We denote this probability distribution over $i \in R^t$ by \mathcal{D} . From what we have established already in (6.9), we can lower bound the probability that i lies in the target support:

$$\begin{aligned} \Pr_{i \sim \mathcal{D}}[i \in S^*] &= \frac{\|x_{R^t \cap S^*}^t\|_2^2}{\|x_{R^t \cap S^*}^t\|_2^2 + \|x_{R^t \setminus S^*}^t\|_2^2} \\ &> \frac{\frac{1}{2(1+\tilde{\kappa})}}{1 + \frac{1}{2(1+\tilde{\kappa})}} \\ &= \frac{1}{2\tilde{\kappa} + 3} \\ &:= p. \end{aligned} \quad (6.10)$$

Note that this event can happen at most once for each $i \in S^*$ during the whole execution of the algorithm, since each element can only be removed once from the set of regularized elements.

We will prove that with constant probability the number of Type 2 steps will be at most $(s^* - 1)(4\tilde{\kappa} + 6) := b$. For $1 \leq k \leq b$, we define the following random variables:

- $i_k \in [n]$ is the index picked in the k -th Type 2 iteration, or \perp if there are less than k Type 2 iterations.
- q_k is the probability of picking an index in the optimal support in the k -th Type 2 iteration (i.e. $i_k \in S^*$):

$$q_k = \begin{cases} \|x_{R^{t_k} \cap S^*}^{t_k}\|_2^2 / \|x_{R^{t_k}}^{t_k}\|_2^2 & \text{if } i_k \neq \perp \\ 0 & \text{otherwise} \end{cases},$$

where $t_k \in [T]$ is the index of the k -th Type 2 iteration within all iterations of the algorithm. Note that, by (6.10), $q_k > 0$ implies $q_k \geq p$.

- X_k is 1 if the index picked in the k -th Type 2 step was in the optimal support:

$$X_k = \begin{cases} 1 & \text{with probability } q_k \\ 0 & \text{otherwise} \end{cases}$$

Our goal is to upper bound $\Pr \left[\sum_{k=1}^b X_k \leq s^* - 1 \right]$. This automatically implies the same upper bound on the probability that there will be more than b Type 2 iterations.

We define another sequence of random variables Y_0, \dots, Y_b , where $Y_0 = 0$, and

$$Y_k = \begin{cases} Y_{k-1} + \frac{p}{q_k} - p & \text{if } X_k = 1 \\ Y_{k-1} - p & \text{if } X_k = 0 \end{cases},$$

for $k \in [b]$. Since if $q_k > 0$ we have $\frac{p}{q_k} \leq 1$, it is immediate that

$$Y_k - Y_{k-1} \leq X_k - p$$

and so $Y_b \leq \sum_{k=1}^b X_k - bp$. Furthermore,

$$\begin{aligned} \mathbb{E}[Y_k \mid i_1, \dots, i_{k-1}] &= Y_{k-1} + q_k \left(\frac{p}{q_k} - p \right) - (1 - q_k)p \\ &= Y_{k-1}, \end{aligned}$$

meaning that Y_0, \dots, Y_b is a martingale with respect to i_1, \dots, i_b . We will apply the inequality from Lemma 6.2.4. We compute a bound on the differences

$$\begin{aligned} Y_{k-1} - Y_k &= \begin{cases} p - \frac{p}{q_k} & \text{if } X_k = 1 \\ p & \text{if } X_k = 0 \end{cases} \\ &\leq p \end{aligned} \tag{6.11}$$

and the variance

$$\begin{aligned} \text{Var}(Y_k \mid i_1, \dots, i_{k-1}) &= \mathbb{E} \left[(Y_k - \mathbb{E}[Y_k \mid i_1, \dots, i_{k-1}])^2 \mid i_1, \dots, i_{k-1} \right] \\ &= \mathbb{E} \left[(Y_k - Y_{k-1})^2 \mid i_1, \dots, i_{k-1} \right] \\ &= q_k \cdot \left(p - \frac{p}{q_k} \right)^2 + (1 - q_k) \cdot p^2 \\ &= q_k \cdot p^2 \left(1 - \frac{2}{q_k} + \frac{1}{q_k^2} \right) + (1 - q_k) \cdot p^2 \\ &= p^2 \left(\frac{1}{q_k} - 1 \right) \\ &\leq p, \end{aligned}$$

where we used (6.11) along with the fact that $q_k \geq p$. Using the concentration inequality from Lemma 6.2.4 we obtain

$$\begin{aligned} \Pr \left[\sum_{k=1}^b X_k \leq s^* - 1 \right] &\leq \Pr [Y_b \leq s^* - 1 - b \cdot p] \\ &\leq e^{-(bp - s^* + 1)^2 / (2(b \cdot p + p \cdot (bp - s^* + 1)/3))} \\ &= e^{-(s^* - 1) / (2(2 + p/3))} \\ &\leq e^{-1 / (2(2 + 1/9))} \\ &< 0.8, \end{aligned}$$

where we used the fact that $bp = 2(s^* - 1)$, $s^* \geq 2$ (otherwise the problem is trivial), and $p = \frac{1}{2\tilde{\kappa} + 3} \leq \frac{1}{3}$. Therefore we conclude that the probability that we have not unregularized the whole set S^* after b steps is at most 0.8. Since we can only have a Type 2 step if there is a regularized element in S^* (this is immediate e.g. from (6.10)), this implies that with probability at least 0.2 the number of Type 2 steps is at most $b = (s^* - 1)(4\tilde{\kappa} + 6)$.

■

6.3.3 Corollaries

As the first corollary of Theorem 6.3.1, we show that it directly implies solution recovery bounds similar to those of [170], while also improving the recovery bound by a constant factor.

Corollary 6.3.11 (Solution recovery). *Given a function f and an (unknown) s^* -sparse solution x^* , such that the Restricted Gradient Optimal Constant at sparsity level s is ζ , i.e.*

$$|\langle \nabla f(x^*), y \rangle| \leq \zeta \|y\|_2 ,$$

for all s -sparse y and as long as

$$s \geq s^* \max \{4\tilde{\kappa} + 7, 12\tilde{\kappa} + 6\} ,$$

Algorithm 15 ensures that

$$f(x) \leq f(x^*) + \varepsilon$$

and

$$\|x - x^*\|_2 \leq \frac{\zeta}{\rho^-} \left(1 + \sqrt{1 + 2\varepsilon \frac{\rho^-}{\zeta^2}} \right) .$$

For any $\theta > 0$ and $\varepsilon \leq \frac{\zeta^2}{\rho^-} \theta (1 + \frac{\theta}{2})$, this implies that

$$\|x - x^*\|_2 \leq (2 + \theta) \frac{\zeta}{\rho^-} .$$

Proof. By strong convexity we have

$$\begin{aligned} \varepsilon &\geq f(x) - f(x^*) \\ &\geq \langle \nabla f(x^*), x - x^* \rangle + \frac{\rho^-}{2} \|x - x^*\|_2^2 \\ &\geq -\zeta \|x - x^*\|_2 + \frac{\rho^-}{2} \|x - x^*\|_2^2 , \end{aligned}$$

therefore

$$\frac{\rho^-}{2} \|x - x^*\|_2^2 - \zeta \|x - x^*\|_2 - \varepsilon \leq 0 ,$$

looking at which as a quadratic polynomial in $\|x - x^*\|_2$, it follows that

$$\begin{aligned}\|x - x^*\|_2 &\leq \frac{\zeta + \sqrt{\zeta^2 + 2\varepsilon\rho^-}}{\rho^-} \\ &= \frac{\zeta}{\rho^-} \left(1 + \sqrt{1 + 2\varepsilon\frac{\rho^-}{\zeta^2}} \right) \\ &= (2 + \theta) \frac{\zeta}{\rho^-},\end{aligned}$$

by setting $\varepsilon = \frac{\zeta^2}{\rho^-} (\theta + \frac{1}{2}\theta^2)$. □

The next corollary shows that our Theorem 6.3.1 can be also used to obtain support recovery results under a ‘‘Signal-to-Noise’’ condition given as a lower bound to $|x_{\min}^*|$.

Corollary 6.3.12 (Support recovery). *As long as*

$$s \geq s^* \max \{4\tilde{\kappa} + 7, 12\tilde{\kappa} + 6\}$$

and $|x_{\min}^*| > \frac{\zeta}{\rho^-}$, Algorithm 15 with $\varepsilon < -\frac{1}{2\rho^-}\zeta^2 + \frac{\rho^-}{2}(x_{\min}^*)^2$ returns a solution x with support S such that

$$S^* \subseteq S.$$

Proof. Let us suppose that $S^* \setminus S^t \neq \emptyset$. By restricted strong convexity we have

$$\begin{aligned}-\frac{1}{2\rho^-}\zeta^2 + \frac{\rho^-}{2}(x_{\min}^*)^2 &> \varepsilon \\ &\geq f(x) - f(x^*) \\ &\geq \langle \nabla f(x^*), x - x^* \rangle + \frac{\rho^-}{2} \|x - x^*\|_2^2 \\ &\geq \langle \nabla f(x^*), x \rangle + \frac{\rho^-}{2} \|x_{S^t \setminus S^*}\|_2^2 + \frac{\rho^-}{2} \|x_{S^* \setminus S^t}^*\|_2^2 \\ &\geq -\frac{1}{2\rho^-} \|\nabla_{S^t \setminus S^*} f(x^*)\|_2^2 + \frac{\rho^-}{2} \|x_{S^* \setminus S^t}^*\|_2^2 \\ &\geq -\frac{1}{2\rho^-}\zeta^2 + \frac{\rho^-}{2}(x_{\min}^*)^2\end{aligned}$$

a contradiction. Here we used the fact that by local optimality $\nabla_{S^*} f(x^*) = \vec{0}$, the inequality $\langle u, v \rangle + \frac{\lambda}{2} \|v\|_2^2 \geq -\frac{1}{2\lambda} \|u\|_2^2$ for any vectors u, v and scalar $\lambda > 0$, and the fact that $\|\nabla_{S^t \setminus S^*} f(x^*)\|_2^2 \leq \zeta^2$ by Definition 6.2.2. Therefore $S^* \subseteq S^t$. □

6.4 Analysis of Orthogonal Matching Pursuit with Replacement (OMPR)

6.4.1 Overview and Main Theorem

The OMPR algorithm was first described (under a different name) in [146]. It is an extension of OMP but after each iteration some element is removed from S^t so that the sparsity remains the

same. The algorithm description is in Algorithm 12.

For each iteration t of Algorithm 12, we will define a solution

$$\tilde{x}^t = \underset{\text{supp}(x) \subseteq S^t \cup S^*}{\text{argmin}} f(x)$$

to be the optimal solution supported on $S^t \cup S^*$. Furthermore, we let \bar{x}^* be the optimal $(s + s^*)$ -sparse solution, i.e.

$$\bar{x}^* = \underset{|\text{supp}(x)| \leq s + s^*}{\text{argmin}} f(x).$$

By definition, the following chain of inequalities holds

$$\min_{x \in \mathbb{R}^n} f(x) \leq f(\bar{x}^*) \leq f(\tilde{x}^t) \leq \min\{f(x^t), f(x^*)\}.$$

We will denote $\mu = \sqrt{\frac{s^*}{s}}$. The following technical lemma is important for our approach, and roughly states that if there is significant ℓ_2 norm difference between x^t and x^* , at least one of x^t, x^* is significantly larger than \tilde{x}^t in function value. Its importance lies on the fact that instead of directly applying strong convexity between x^t and x^* , it gets a tighter bound by making use of \tilde{x}^t .

Lemma 6.4.1. *For any function f with RSC constant ρ^- at sparsity level $s + s^*$ and any two solutions x^t, x^* with respective supports S^t, S^* and sparsity levels s, s^* , we have that*

$$\left(\sqrt{f(x^t) - f(\tilde{x}^t)} + \sqrt{f(x^*) - f(\tilde{x}^t)} \right)^2 \geq \frac{\rho^-}{2} \left(\|x_{S^* \setminus S^t}^*\|_2^2 + \|x_{S^t \setminus S^*}^t\|_2^2 \right).$$

The proof can be found in Appendix 9.4.1.

The following theorem is the main result of this section. Its strength lies in its generality, and various useful corollaries can be directly extracted from it. It can be seen as a careful and general analysis of OMP, and, in contrast to Theorem 6.3.1, the interesting part is not the asymptotic sparsity bound (which is $O(\kappa^2 s^*)$ and is known), but the constant factor in front of it, which allows its use in recovering a solution with sparsity close to s^* , under a RIP bound. It roughly states that for any solution x^* that is s^* -sparse, OMP can be used to obtain a solution x

- with sparsity $\leq \kappa^2 s^*$ that comes arbitrarily close to x^* in function value, i.e. $f(x) \leq f(x^*) + \varepsilon$
- with sparsity $\leq \kappa^2 s^*/4$ that approximates x^* in function value, and the approximation depends on how close x^* is to being a global optimum.

The latter can be used to obtain compressed sensing results, as it gives sparsity very close to s^* given that upper bounds on κ (equivalently, RIP upper bounds) are met. In comparison with previously known results, our work is the first to obtain compressed sensing RIP bounds for general functions f and for a wide range of sparsity levels from s^* to much larger than that.

Theorem 6.4.2. *Given a function f , an (unknown) s^* -sparse solution x^* , a desired solution sparsity level s , and error parameters $\varepsilon > 0$ and $0 < \theta < 1$, Algorithm 12 returns an s -sparse solution x such that*

- If $\tilde{\kappa} \sqrt{\frac{s^*}{s}} \leq 1$, then

$$f(x) \leq f(x^*) + \varepsilon$$

in $O\left(\sqrt{ss^*} \log \frac{f(x^0) - f(x^*)}{\varepsilon}\right)$ iterations.

- If $1 < \tilde{\kappa} \sqrt{\frac{s^*}{s}} < 2 - \theta$, then

$$f(x) \leq f(x^*) + B$$

where

$$B = \varepsilon + \frac{4(1 - \theta) \left(\tilde{\kappa} \sqrt{\frac{s^*}{s}} - 1\right)}{\left(2 - \tilde{\kappa} \sqrt{\frac{s^*}{s}} - \theta\right)^2} (f(x^*) - f(\bar{x}^*))$$

in $O\left(\frac{\sqrt{ss^*}}{\theta} \log \frac{f(x^0) - f(x^*)}{B}\right)$ iterations.

6.4.2 Progress Lemma and Theorem Proof

The main ingredient needed to prove Theorem 6.4.2 is the following lemma, which bounds the progress of Algorithm 12 in one iteration.

Lemma 6.4.3 (OMPR Progress Lemma). *We can bound the progress of one step of the algorithm by distinguishing the following three cases:*

- If $\mu\tilde{\kappa} \leq 1$, then

$$f(x^{t+1}) - f(x^*) \leq (f(x^t) - f(x^*)) \left(1 - \frac{\mu}{|S^* \setminus S^t|}\right)$$

- If $\mu\tilde{\kappa} > 1$ and $f(x^*) = f(\tilde{x}^t)$, then

$$f(x^{t+1}) - f(x^*) \leq (f(x^t) - f(x^*)) \cdot \left(1 - \frac{\mu}{|S^* \setminus S^t|} (2 - \mu\tilde{\kappa})\right)$$

- If $\mu\tilde{\kappa} > 1$ and $f(x^*) > f(\tilde{x}^t)$, then

$$f(x^{t+1}) - f(x^*) \leq (f(x^t) - f(x^*)) \cdot \left(1 - \frac{\mu}{|S^* \setminus S^t|} \left(2 - \mu\tilde{\kappa} - \frac{2(\mu\tilde{\kappa} - 1)}{\sqrt{\frac{f(x^t) - f(\tilde{x}^t)}{f(x^*) - f(\tilde{x}^t)} - 1}}\right)\right)$$

Proof. The proof will proceed as follows: We will use the smoothness and strong convexity of f to get a bound on the progress of one step of the algorithm in decreasing f , based on $f(x^t) - f(x^*)$. This progress will be offset by the ℓ_2 norm of x^* restricted to $S^* \setminus S^t$. We use Lemma 6.4.1 to upper bound this norm by a quantity that depends on $f(x^t) - f(\tilde{x}^t)$ and $f(x^*) - f(\tilde{x}^t)$, where \tilde{x}^t is the optimal solution in the joint support $S^t \cup S^*$. Finally, by a careful case analysis based on the value of μ , we obtain the three bullet points in the lemma statement.

First of all, if $S^* \subseteq S^t$ then, since x^t is an S^t -restricted minimizer, we have $f(x^t) \leq f(x^*)$ and we are done. So suppose otherwise, i.e. $S^* \setminus S^t \neq \emptyset$ and $f(x^t) > f(x^*)$. Let $i = \operatorname{argmax}_{i \notin S^t} |\nabla_i f(x^t)|$

and $j = \operatorname{argmin}_{j \in S^t} |x_j^t|$. By definition of OMPR (Algorithm 12) and restricted smoothness of f , we

have

$$\begin{aligned}
f(x^{t+1}) &\leq \min_{\eta \in \mathbb{R}} f(x^t + \eta \vec{1}_i - x_j^t \vec{1}_j) \\
&\leq \min_{\eta \in \mathbb{R}} f(x^t) + \langle \nabla f(x^t), \eta \vec{1}_i - x_j^t \vec{1}_j \rangle + \frac{\rho^+}{2} \left\| \eta \vec{1}_i - x_j^t \vec{1}_j \right\|_2^2 \\
&= \min_{\eta \in \mathbb{R}} f(x^t) + \eta \nabla_i f(x^t) + \frac{\rho^+}{2} \eta^2 + \frac{\rho^+}{2} (x_j^t)^2 \\
&= f(x^t) - \frac{(\nabla_i f(x^t))^2}{2\rho^+} + \frac{\rho^+}{2} (x_j^t)^2 \\
&\leq f(x^t) - \frac{\|\nabla_{S^* \setminus S^t} f(x^t)\|_2^2}{2\rho^+ |S^* \setminus S^t|} + \frac{\rho^+}{2|S^t \setminus S^*|} \|x_{S^t \setminus S^*}^t\|_2^2,
\end{aligned} \tag{6.12}$$

where the second to last equality follows from the fact that $\nabla_j f(x^t) = \vec{0}$, as x^t is an S^t -restricted minimizer of f , and the last inequality since

$$(x_j^t)^2 = \min_{j \in S^t \setminus S^*} (x_j^t)^2 \leq \frac{\|x_{S^t \setminus S^*}^t\|_2^2}{|S^t \setminus S^*|}.$$

Re-arranging (6.12), we get

$$|S^* \setminus S^t| (f(x^t) - f(x^{t+1})) \geq \frac{\|\nabla_{S^* \setminus S^t} f(x^t)\|_2^2}{2\rho^+} - \frac{\rho^+ |S^* \setminus S^t|}{2 |S^t \setminus S^*|} \|x_{S^t \setminus S^*}^t\|_2^2. \tag{6.13}$$

On the other hand, by restricted strong convexity of f ,

$$\begin{aligned}
f(x^*) - f(x^t) &\geq \langle \nabla f(x^t), x^* - x^t \rangle + \frac{\rho^-}{2} \|x^* - x^t\|_2^2 \\
&= \langle \nabla_{S^* \setminus S^t} f(x^t), x_{S^* \setminus S^t}^* \rangle + \frac{\rho^-}{2} \|x^* - x^t\|_2^2 \\
&\geq \langle \nabla_{S^* \setminus S^t} f(x^t), x_{S^* \setminus S^t}^* \rangle + \frac{\rho^-}{2} \|x_{S^* \setminus S^t}^*\|_2^2 + \frac{\rho^-}{2} \|x_{S^t \setminus S^*}^t\|_2^2 \\
&\geq \langle \nabla_{S^* \setminus S^t} f(x^t), x_{S^* \setminus S^t}^* \rangle + \frac{\mu\rho^+}{2} \|x_{S^* \setminus S^t}^*\|_2^2 \\
&\quad + \frac{\rho^- - \mu\rho^+}{2} \|x_{S^* \setminus S^t}^*\|_2^2 + \frac{\rho^-}{2} \|x_{S^t \setminus S^*}^t\|_2^2 \\
&\geq -\frac{1}{2\mu\rho^+} \|\nabla_{S^* \setminus S^t} f(x^t)\|_2^2 + \frac{\rho^- - \mu\rho^+}{2} \|x_{S^* \setminus S^t}^*\|_2^2 + \frac{\rho^-}{2} \|x_{S^t \setminus S^*}^t\|_2^2,
\end{aligned} \tag{6.14}$$

where the first equality follows from the fact that $\nabla_{S^t} f(x^t) = \vec{0}$ as x^t is an S^t -restricted minimizer of f and the last inequality from using the fact that $\langle u, v \rangle + \frac{\lambda}{2} \|v\|_2^2 \geq -\frac{1}{2\lambda} \|u\|_2^2$ for any $\lambda > 0$.

Re-arranging (6.14), we get

$$\frac{1}{2\mu\rho^+} \|\nabla_{S^* \setminus S^t} f(x^t)\|_2^2 \geq f(x^t) - f(x^*) - \frac{\mu\rho^+ - \rho^-}{2} \|x_{S^* \setminus S^t}^*\|_2^2 + \frac{\rho^-}{2} \|x_{S^t \setminus S^*}^t\|_2^2.$$

By substituting this into (6.13),

$$\begin{aligned} & |S^* \setminus S^t| (f(x^t) - f(x^{t+1})) \\ & \geq \mu (f(x^t) - f(x^*)) - \frac{\mu^2 \rho^+ - \mu \rho^-}{2} \left\| x_{S^* \setminus S^t}^* \right\|_2^2 + \frac{\mu \rho^-}{2} \left\| x_{S^t \setminus S^*}^t \right\|_2^2 - \frac{\rho^+ |S^* \setminus S^t|}{2 |S^t \setminus S^*|} \left\| x_{S^t \setminus S^*}^t \right\|_2^2. \end{aligned}$$

Note that by our choice of μ and since $s^* \leq s$,

$$\mu^2 \rho^+ = \rho^+ \frac{s^*}{s} \geq \rho^+ \frac{s^* - |S^* \cap S^t|}{s - |S^* \cap S^t|} = \rho^+ \frac{|S^* \setminus S^t|}{|S^t \setminus S^*|}$$

and so

$$\begin{aligned} & \mu (f(x^t) - f(x^*)) - \frac{\mu^2 \rho^+ - \mu \rho^-}{2} \left\| x_{S^* \setminus S^t}^* \right\|_2^2 + \frac{\mu \rho^-}{2} \left\| x_{S^t \setminus S^*}^t \right\|_2^2 - \frac{\rho^+ |S^* \setminus S^t|}{2 |S^t \setminus S^*|} \left\| x_{S^t \setminus S^*}^t \right\|_2^2 \\ & \geq \mu (f(x^t) - f(x^*)) - \frac{\mu}{2} (\mu \rho^+ - \rho^-) \left(\left\| x_{S^* \setminus S^t}^* \right\|_2^2 + \left\| x_{S^t \setminus S^*}^t \right\|_2^2 \right), \end{aligned}$$

concluding that

$$|S^* \setminus S^t| (f(x^t) - f(x^{t+1})) \geq \mu (f(x^t) - f(x^*)) - \frac{\mu}{2} (\mu \rho^+ - \rho^-) \left(\left\| x_{S^* \setminus S^t}^* \right\|_2^2 + \left\| x_{S^t \setminus S^*}^t \right\|_2^2 \right).$$

For $\mu \tilde{\kappa} \leq 1 \Leftrightarrow \mu \rho^+ - \rho^- \leq 0$, this automatically implies that

$$\begin{aligned} f(x^t) - f(x^{t+1}) & \geq \frac{\mu}{|S^* \setminus S^t|} (f(x^t) - f(x^*)) \\ \Leftrightarrow f(x^{t+1}) - f(x^*) & \leq \left(1 - \frac{\mu}{|S^* \setminus S^t|} \right) (f(x^t) - f(x^*)). \end{aligned}$$

On the other hand, if $\mu \tilde{\kappa} > 1$ we have

$$\begin{aligned} & |S^* \setminus S^t| (f(x^t) - f(x^{t+1})) \\ & \geq \mu (f(x^t) - f(x^*)) - \frac{\mu}{2} (\mu \rho^+ - \rho^-) \left(\left\| x_{S^* \setminus S^t}^* \right\|_2^2 + \left\| x_{S^t \setminus S^*}^t \right\|_2^2 \right) \\ & \geq \mu (f(x^t) - f(x^*)) - \mu (\mu \tilde{\kappa} - 1) \left(\sqrt{f(x^t) - f(\tilde{x}^t)} + \sqrt{f(x^*) - f(\tilde{x}^t)} \right)^2, \end{aligned}$$

where we used Lemma 6.4.1. If $f(x^*) = f(\tilde{x}^t)$ it is immediate that

$$f(x^{t+1}) - f(x^*) \leq \left(1 - \frac{\mu}{|S^* \setminus S^t|} (2 - \mu \tilde{\kappa}) \right) (f(x^t) - f(x^*)),$$

so let us from now on assume that $f(x^*) > f(\tilde{x}^t)$ and set $a = f(x^t) - f(\tilde{x}^t)$, $a' = f(x^{t+1}) - f(\tilde{x}^t)$, and $b = f(x^*) - f(\tilde{x}^t)$. From what we have concluded before

$$|S^* \setminus S^t| (a - a') \geq \mu (a - b) - \mu (\mu \tilde{\kappa} - 1) \left(\sqrt{a} + \sqrt{b} \right)^2,$$

or equivalently

$$\begin{aligned}
a' - b &\leq \left(1 - \frac{\mu}{|S^* \setminus S^t|}\right) (a - b) + \frac{\mu}{|S^* \setminus S^t|} (\mu\tilde{\kappa} - 1) (\sqrt{a} + \sqrt{b})^2 \\
&= (a - b) \left(1 - \frac{\mu}{|S^* \setminus S^t|} \left(1 - (\mu\tilde{\kappa} - 1) \frac{(\sqrt{a} + \sqrt{b})^2}{a - b}\right)\right) \\
&= (a - b) \left(1 - \frac{\mu}{|S^* \setminus S^t|} \left(1 - (\mu\tilde{\kappa} - 1) \frac{\sqrt{\frac{a}{b}} + 1}{\sqrt{\frac{a}{b}} - 1}\right)\right) \\
&= (a - b) \left(1 - \frac{\mu}{|S^* \setminus S^t|} \left(1 - (\mu\tilde{\kappa} - 1) \left(1 + \frac{2}{\sqrt{\frac{a}{b}} - 1}\right)\right)\right) \\
&= (a - b) \left(1 - \frac{\mu}{|S^* \setminus S^t|} \left(2 - \mu\tilde{\kappa} - \frac{2(\mu\tilde{\kappa} - 1)}{\sqrt{\frac{a}{b}} - 1}\right)\right),
\end{aligned}$$

Replacing back a, a', b , the desired statement follows:

$$f(x^{t+1}) - f(x^*) \leq (f(x^t) - f(x^*)) \cdot \left(1 - \frac{\mu}{|S^* \setminus S^t|} \left(2 - \mu\tilde{\kappa} - \frac{2(\mu\tilde{\kappa} - 1)}{\sqrt{\frac{f(x^t) - f(\tilde{x}^t)}{f(x^*) - f(\tilde{x}^t)}} - 1}\right)\right).$$

□

The proof of Theorem 6.4.2 now follows by appropriately applying Lemma 6.4.3.

Proof of Theorem 6.4.2. We will directly apply Lemma 6.4.3 over the course of multiple iterations. For the second bullet of the theorem statement, the progress bound of Lemma 6.4.3 also depends on $(f(x^t) - f(\tilde{x}^t))/(f(x^*) - f(\tilde{x}^t))$. We show that, as long as $f(x)$ is significantly larger than $f(x^*)$, this can be lower bounded by a sufficiently large quantity, leading to fast convergence.

Case 1: $\mu\tilde{\kappa} \leq 1$.

By Lemma 6.4.3, we have

$$\begin{aligned}
f(x^T) - f(x^*) &\leq (f(x^{T-1}) - f(x^*)) \left(1 - \frac{\mu}{|S^* \setminus S^{T-1}|}\right) \\
&\leq (f(x^{T-1}) - f(x^*)) \left(1 - \frac{\mu}{s^*}\right) \\
&\leq (f(x^{T-1}) - f(x^*)) e^{-\frac{\mu}{s^*}} \\
&\leq \dots \\
&\leq (f(x^0) - f(x^*)) e^{-T \frac{\mu}{s^*}} \\
&\leq \varepsilon,
\end{aligned}$$

for our choice of $T = O\left(\sqrt{ss^*} \log \frac{f(x^0) - f(x^*)}{\varepsilon}\right)$ and replacing $\mu = \sqrt{\frac{s^*}{s}}$.

Case 2: $\mu\tilde{\kappa} > 1$.

Let \mathcal{A} be the set of $0 \leq t \leq T - 1$ such that $f(x^*) = f(\tilde{x}^t)$ and \mathcal{B} the set of $0 \leq t \leq T - 1$ such

that $f(x^*) > f(\tilde{x}^t)$. By Lemma 6.4.3, for $t \in \mathcal{A}$ we then have

$$\begin{aligned} f(x^{t+1}) - f(x^*) &\leq (f(x^t) - f(x^*)) \left(1 - \frac{\mu}{|S^* \setminus S^t|} (2 - \mu\tilde{\kappa}) \right) \\ &\leq (f(x^t) - f(x^*)) \left(1 - \frac{\mu}{s^*} (2 - \mu\tilde{\kappa}) \right). \end{aligned}$$

We now consider the case $t \in \mathcal{B}$. By Lemma 6.4.3,

$$f(x^{t+1}) - f(x^*) \leq (f(x^t) - f(x^*)) \cdot \left(1 - \frac{\mu}{|S^* \setminus S^t|} \left(2 - \mu\tilde{\kappa} - \frac{2(\mu\tilde{\kappa} - 1)}{\sqrt{\frac{f(x^t) - f(\tilde{x}^t)}{f(x^*) - f(\tilde{x}^t)} - 1}} \right) \right). \quad (6.15)$$

Let us suppose that the theorem statement is not true. This implies

$$\begin{aligned} f(x^t) - f(x^*) &\geq f(x^T) - f(x^*) \\ &> \varepsilon + \frac{4(1 - \theta)(\mu\tilde{\kappa} - 1)}{(2 - \mu\tilde{\kappa} - \theta)^2} (f(x^*) - f(\tilde{x}^*)) \\ &\geq \varepsilon + \frac{4(1 - \theta)(\mu\tilde{\kappa} - 1)}{(2 - \mu\tilde{\kappa} - \theta)^2} (f(x^*) - f(\tilde{x}^t)) \\ &\geq \frac{4(1 - \theta)(\mu\tilde{\kappa} - 1)}{(2 - \mu\tilde{\kappa} - \theta)^2} (f(x^*) - f(\tilde{x}^t)), \end{aligned} \quad (6.16)$$

for all $0 \leq t \leq T$. Therefore

$$\begin{aligned} f(x^t) - f(\tilde{x}^t) &> \left(\frac{4(1 - \theta)(\mu\tilde{\kappa} - 1)}{(2 - \mu\tilde{\kappa} - \theta)^2} + 1 \right) (f(x^*) - f(\tilde{x}^t)) \\ &= \left(\frac{4(1 - \theta)(\mu\tilde{\kappa} - 1) + 4 + (\mu\tilde{\kappa} + \theta)^2 - 4(\mu\tilde{\kappa} + \theta)}{(2 - \mu\tilde{\kappa} - \theta)^2} \right) \cdot (f(x^*) - f(\tilde{x}^t)) \\ &= \frac{(\mu\tilde{\kappa} - \theta)^2}{(2 - \mu\tilde{\kappa} - \theta)^2} (f(x^*) - f(\tilde{x}^t)), \end{aligned}$$

or equivalently for all $t \in \mathcal{B}$

$$\sqrt{\frac{f(x^t) - f(\tilde{x}^t)}{f(x^*) - f(\tilde{x}^t)}} - 1 > \frac{\mu\tilde{\kappa} - \theta}{2 - \mu\tilde{\kappa} - \theta} - 1 = \frac{2(\mu\tilde{\kappa} - 1)}{2 - \mu\tilde{\kappa} - \theta}.$$

Replacing this into (6.15), we get that for any $t \in \mathcal{B}$

$$\begin{aligned} f(x^{t+1}) - f(x^*) &\leq (f(x^t) - f(x^*)) \cdot \left(1 - \frac{\mu}{|S^* \setminus S^t|} \left(2 - \mu\tilde{\kappa} - \frac{2(\mu\tilde{\kappa} - 1)}{\sqrt{\frac{f(x^t) - f(\tilde{x}^t)}{f(x^*) - f(\tilde{x}^t)} - 1}} \right) \right) \\ &\leq (f(x^t) - f(x^*)) \left(1 - \frac{\mu}{|S^* \setminus S^t|} \theta \right) \end{aligned}$$

and so combining it with the case $t \in \mathcal{A}$ and using the fact that $\mu\tilde{\kappa} < 2 - \theta \Leftrightarrow \theta < 2 - \mu\tilde{\kappa}$,

$$\begin{aligned}
f(x^T) - f(x^*) &\leq (f(x^{T-1}) - f(x^*)) \left(1 - \frac{\mu}{|S^* \setminus S^{T-1}|} \min\{2 - \mu\tilde{\kappa}, \theta\}\right) \\
&\leq (f(x^{T-1}) - f(x^*)) \left(1 - \frac{\mu}{s^*} \theta\right) \\
&\leq (f(x^{T-1}) - f(x^*)) e^{-\frac{\mu}{s^*} \theta} \\
&\leq \dots \\
&\leq (f(x^0) - f(x^*)) e^{-T \frac{\mu}{s^*} \theta} \\
&= \varepsilon + \frac{4(1 - \theta)(\mu\tilde{\kappa} - 1)}{(2 - \mu\tilde{\kappa} - \theta)^2} (f(x^*) - f(\bar{x}^*)),
\end{aligned}$$

where the last equality follows by our choice of

$$T = \frac{\sqrt{ss^*}}{\theta} \log \frac{f(x^0) - f(x^*)}{B}$$

and replacing $\mu = \sqrt{\frac{s^*}{s}}$. This is a contradiction. ■

6.4.3 Corollaries of Theorem 6.4.2

The first corollary states that in the “noiseless” case (i.e. when the target solution is globally optimal), the returned solution can reach arbitrarily close to the target solution. For $s = s^*$, it gives the same condition of $\tilde{\kappa} < 2$ (or $\delta < 1/3$) as in [87], while working for any function f . For $s > s^*$, it additionally gives a tradeoff between the sparsity and the RIP bound required for the algorithm.

Corollary 6.4.4 (Noiseless case). *If $\tilde{\kappa} \sqrt{\frac{s^*}{s}} < 2$ and x^* is a globally optimal solution, i.e. $f(x^*) = \min_z f(z)$, Algorithm 12 returns a solution with*

$$f(x) \leq f(x^*) + \varepsilon$$

in $O\left(\frac{\sqrt{ss^*}}{2 - \tilde{\kappa} \sqrt{\frac{s^*}{s}}} \log \frac{f(x^0) - f(x^*)}{\varepsilon}\right)$ iterations.

Proof. We apply Theorem 6.4.2 with $\theta = \frac{1}{2} \left(2 - \tilde{\kappa} \sqrt{\frac{s^*}{s}}\right)$. □

The following result is in the usual form of sparse recovery results, which provide a bound on $\|x - x^*\|_2$ given a RIP constant upper bound. It provides a tradeoff between the RIP constant and the sparsity of the returned solution. Similar results can be found e.g. in [31] using LASSO, albeit they only apply to case of linear regression ($f(x) = \frac{1}{2} \|Ax - b\|_2^2$) and do not offer a sparsity tradeoff.

Corollary 6.4.5 (ℓ_2 solution recovery). *Given any parameters $\varepsilon > 0$ and $0 < \theta < 1$, the returned solution x of Algorithm 12 will satisfy*

$$\|x - x^*\|_2^2 \leq \varepsilon + C \left(f(x^*) - \min_z f(z)\right)$$

as long as

$$\delta_{s+s^*} < \frac{(2-\theta)\sqrt{\frac{s}{s^*}} - 1}{(2-\theta)\sqrt{\frac{s}{s^*}} + 1},$$

where C is a constant that depends only on θ , δ_{s+s^*} , and $\frac{s}{s^*}$.

Proof. By triangle inequality and strong convexity, and letting \tilde{x}^* be the optimal solution in $\text{supp}(x) \cup \text{supp}(x^*)$, we have

$$\begin{aligned} \|x - x^*\|_2^2 &\leq 2 (\|x - \tilde{x}^*\|_2^2 + \|x^* - \tilde{x}^*\|_2^2) \\ &\leq \frac{4}{1 - \delta_{s+s^*}} (f(x) - f(\tilde{x}^*) + f(x^*) - f(\tilde{x}^*)) \\ &= \frac{4}{1 - \delta_{s+s^*}} (f(x) - f(x^*) + 2(f(x^*) - f(\tilde{x}^*))) \\ &= \frac{4}{1 - \delta_{s+s^*}} \left(f(x) - f(x^*) + 2(f(x^*) - \min_z f(z)) \right). \end{aligned}$$

Now, by applying Theorem 6.4.2 for some error tolerance $\hat{\varepsilon} > 0$, we get

$$f(x) \leq f(x^*) + \hat{\varepsilon} + \widehat{C}(f(x^*) - f(\bar{x}^*)) \leq f(x^*) + \hat{\varepsilon} + \widehat{C}(f(x^*) - \min_z f(z)),$$

for some $\widehat{C} > 0$, and so we conclude that

$$\|x - x^*\|_2^2 \leq \frac{4\hat{\varepsilon}}{1 - \delta_{s+s^*}} + \frac{4(\widehat{C} + 2)}{1 - \delta_{s+s^*}} (f(x^*) - \min_z f(z)).$$

The statement follows by setting $\varepsilon = \frac{4\hat{\varepsilon}}{1 - \delta_{s+s^*}}$ and $C = \frac{4(\widehat{C} + 2)}{1 - \delta_{s+s^*}}$. \square

In particular, for $s = s^*$, the above lemma implies recovery under the condition $\delta_{2s^*} < \frac{1}{3}$.

6.5 Lower Bounds

6.5.1 $\Omega(s^*\kappa)$ Lower Bound due to [63]

In Appendix B of [63] a matrix $A \in \mathbb{R}^{m \times n}$ and a vector $b \in \mathbb{R}^m$ are constructed and let us define $f(x) = \frac{1}{2} \|Ax - b\|_2^2$. If we let $\overline{S^*} = \{1, \dots, n-2\}$ and $S^* = \{n-1, n\}$, then f has the property that

$$\min_{\text{supp}(x) \subseteq S^*} f(x) = \min_{\text{supp}(x) \subseteq \overline{S^*}} f(x) = 0,$$

but for any $S \subset \overline{S^*}$,

$$\min_{\text{supp}(x) \subseteq S} f(x) > 0.$$

Furthermore, for any $S \subset \overline{S^*}$ and $x = \underset{\text{supp}(x) \subseteq S}{\text{argmin}} f(x)$, it is true that

$$\max_{i \in S^*} |\nabla_i f(x)| < \min_{i \in S^* \setminus S} |\nabla_i f(x)|.$$

This means that for any algorithm with an OMP-like criterion like Orthogonal Matching Pursuit, Orthogonal Matching Pursuit with Replacement, Iterative Hard Thresholding, and Partial Hard Thresholding, if the initial solution does not have an intersection with S^* , then it will never have, therefore implying that the sparsity returned by the algorithm is $|S| = n - 2 = \Omega(n)$. As for this construction $\kappa = \frac{\rho_n^+}{\rho_n} = O(n)$, there exists a constant c such that the sparsity of the returned solution cannot be less than $cs^*\kappa$, since $s^*\kappa = O(n) = O(|S|)$. Therefore none of these algorithms can improve the bound $O(s^*\kappa)$ of Theorem 6.3.1 by more than a constant factor. This example also applies to ARHT and Exhaustive Local Search.

It seems difficult to get past this example and achieve sparsity $s = O(s^*\kappa^{1-\delta})$ for some $\delta > 0$. We conjecture that there might be a way to turn the above example into an inapproximability result:

Conjecture 6.5.1. *For any $\delta > 0$, there is no polynomial time algorithm that given a matrix $A \in \mathbb{R}^{m \times n}$, a vector $b \in \mathbb{R}^m$, a target sparsity $s^* \geq 1$, and a desired accuracy $\varepsilon > 0$, returns an $s = O(s^*\kappa_{s+s^*}^{1-\delta})$ -sparse solution x such that $\|Ax - b\|_2^2 \leq \min_{\|x^*\|_0 \leq s^*} \|Ax^* - b\|_2^2 + \varepsilon$, if such a solution exists.*

6.5.2 $\Omega(s^*\kappa^2)$ Lower Bound for OMPR

The following lemma shows that, without regularization, OMPR requires sparsity $\Omega(s^*\kappa^2)$ in general, and therefore the sparsity upper bound is tight. We assume that the algorithm is run for a fixed T iterations, even when the solution stops improving, for a clearer presentation.

Lemma 6.5.2. *There is a function $f(x) = \frac{1}{2} \|Ax - b\|_2^2$ where $A \in \mathbb{R}^{n \times n}$ and $b \in \mathbb{R}^n$ and a target solution x^* of f with sparsity s^* , as well as a set $S \subseteq [n]$ with $|S| = \Theta(s^*\kappa^2)$ such that OMPR initialized with support set S returns a solution x with $f(x) = f(x^*) + \Theta(s^*\kappa^2)$.*

Proof. Without loss of generality we assume that κ is an even integer and set $n = s^*(1 + \kappa + \kappa^2)$. We then partition $[n]$ into three intervals $I_1 = [1, s^*]$, $I_2 = [s^* + 1, s^*(1 + \kappa)]$, $I_3 = [s^*(1 + \kappa) + 1, s^*(1 + \kappa + \kappa^2)]$. We define the diagonal matrix $A \in \mathbb{R}^{n \times n}$ such that

$$A_{ii} = \begin{cases} 1 & \text{if } i \in I_1 \\ \sqrt{\kappa} & \text{if } i \in I_2 \\ 1 & \text{if } i \in I_3 \end{cases}$$

and vector $b \in \mathbb{R}^n$ such that

$$b_i = \begin{cases} \kappa\sqrt{1 - 4\delta} & \text{if } i \in I_1 \\ \sqrt{\kappa}\sqrt{1 - 2\delta} & \text{if } i \in I_2 \\ 1 & \text{if } i \in I_3 \end{cases},$$

where $\delta > 0$ is a sufficiently small scalar used to avoid ties in the steps of the algorithm. The target solution is defined as

$$x_i^* = \begin{cases} \kappa(1 - 4\delta) & \text{if } i \in I_1 \\ 0 & \text{if } i \in I_2 \cup I_3 \end{cases}$$

and its value is $f(x^*) = s^*\kappa^2(1 - \delta)$. Now consider any initial support set $S^0 \subset I_3$ such that

$|S^0| = s^* \kappa^2 / 2$. The initial solution will then be

$$x_i^0 = \begin{cases} 0 & \text{if } i \in I_1 \cup I_2 \cup I_3 \setminus S^0 \\ 1 & \text{if } i \in S^0 \end{cases}$$

and its value $f(x^0) = s^* \kappa^2 (\frac{5}{4} - 3\delta) = f(x^*) + \Theta(s^* \kappa^2)$. The gradient at x^0 is

$$\nabla_i f(x^0) = \begin{cases} -\kappa \sqrt{1 - 4\delta} & \text{if } i \in I_1 \\ -\kappa \sqrt{1 - 2\delta} & \text{if } i \in I_2 \\ -1 & \text{if } i \in I_3 \setminus S^0 \\ 0 & \text{if } i \in S^0 \end{cases},$$

therefore the algorithm will pick $S^1 = S^0 \cup \{i^0\} \setminus \{j^0\}$ for some $i^0 \in I_2$ and some $j^0 \in S^0$, since the gradient entries in I_2 have the largest magnitude among those in $[n]$. The new solution will be

$$x_i^1 = \begin{cases} 0 & \text{if } i \in I_1 \cup I_2 \cup I_3 \setminus S^1 \\ \sqrt{1 - 2\delta} & \text{if } i = i^0 \\ 1 & \text{if } i \in S^1 \setminus \{i^0\} \end{cases}$$

with value $f(x^1) = s^* \kappa^2 (\frac{5}{4} - 3\delta) - \frac{1}{2}(\kappa(1 - 2\delta) - 1)$ and gradient

$$\nabla_i f(x^1) = \begin{cases} -\kappa \sqrt{1 - 4\delta} & \text{if } i \in I_1 \\ -\kappa \sqrt{1 - 2\delta} & \text{if } i \in I_2 \setminus S^1 \\ -1 & \text{if } i \in I_3 \setminus S^1 \\ 0 & \text{if } i \in S^1 \end{cases}$$

and therefore the algorithm will pick $S^2 = S^1 \cup \{i^1\} \setminus \{i^0\}$ for some $i^1 \in I_2$. i^0 will be the one to be removed from S^1 because x_{i^0} has the smallest magnitude out of all entries in S^1 . Continuing this process, the algorithm will always have $S^t \cap I_2 = 1$ and $S^t \cap I_3 = |S^t| - 1$, and so $f(x^t) = s^* \kappa^2 (\frac{5}{4} - 3\delta) - \frac{1}{2}(\kappa(1 - 2\delta) - 1) = f(x^*) + \Theta(s^* \kappa^2)$ for $t \geq 1$. \square

6.6 Experiments

6.6.1 Overview

In this section we evaluate the training performance of different algorithms in the tasks of Linear Regression and Logistic Regression. More specifically, for each algorithm we are interested in how the *loss* over the training set (the quality of the solution) evolves as a function of the *sparcity* of the solution, i.e. the number of non-zeros.

The algorithms that we will consider are *LASSO*, *Orthogonal Matching Pursuit (OMP)*, *Orthogonal Matching Pursuit with Replacement (OMPR)*, *Adaptively Regularized Hard Thresholding (ARHT)* (Algorithm 15), and *Exhaustive Local Search* (Algorithm 13). We run our experiments on publicly available regression and binary classification data sets, out of which we have presented those on which the algorithms have significantly different performance between each other. In some of the other data sets that we tested, we observed that all algorithms had similar performance. The results are presented in Figures 6-1, 6-2, 6-3, 6-4. Also, in Figure 6-5 we present a runtime comparison

between ARHT and Exhaustive Local Search in the year and census datasets. Another relevant class of algorithms that we considered was ℓ_p *Approximate Message Passing* algorithms [48, 171]. Brief experiments showed its performance in terms of sparsity for $p \leq 0.5$ to be promising (on par with OMPR and ARHT although these had much faster runtimes), however a detailed comparison is left for future work.

In both types of objectives (linear and logistic) we include an intercept term, which is present in all solutions (i.e. it is always counted as +1 in the sparsity of the solution). For consistency, all greedy algorithms (OMPR, ARHT, Exhaustive Local Search) are initialized with the OMP solution of the same sparsity.

We note that this section is not supposed to be a conclusive experimental evaluation of the aforementioned algorithms, but rather a preliminary set of experiments that gives partial evidence for their performance. A more extensive future experimental evaluation should focus on implementing runtime-optimized versions of these algorithms and generating sparsity vs loss and loss vs runtime plots for a larger collection of real datasets and with more datapoints and features.

These experiments suggest that Exhaustive Local Search outperforms the other algorithms. However, ARHT also has promising performance and it might be preferred because of better computational efficiency. As a general conclusion, however, both Exhaustive Local Search and ARHT give sparser solutions than other known methods in all the examples we tested. We leave a comprehensive comparison for future work. As a limitation, we observe that ARHT has inconsistent performance in some cases, oscillating between the Exhaustive Local Search and OMPR solutions.

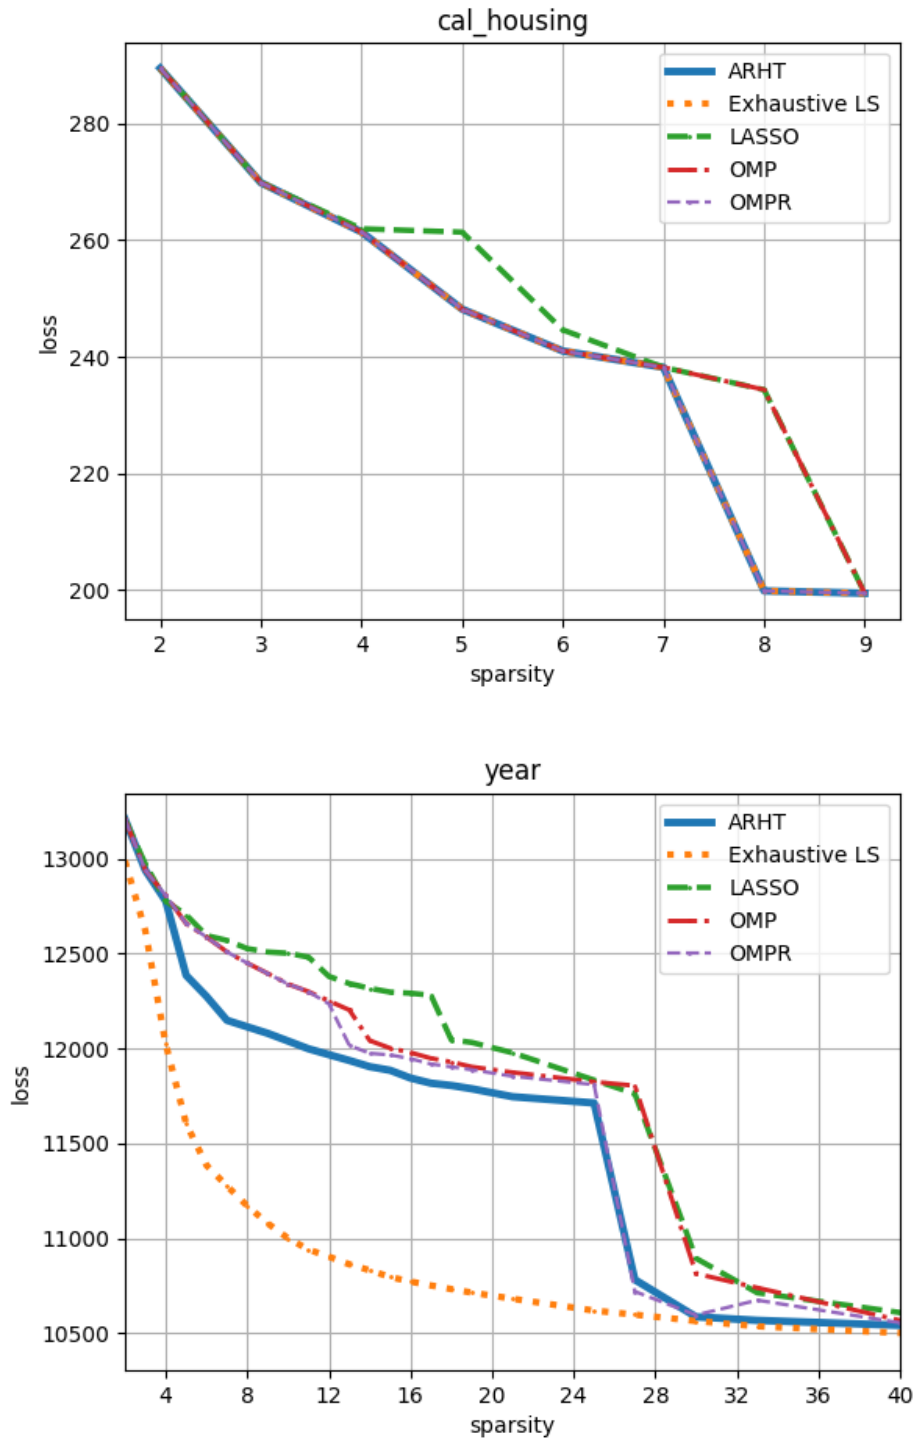


Figure 6-1: Comparison of different algorithms in the Regression data sets *cal_housing* and *year* using the Linear Regression loss.

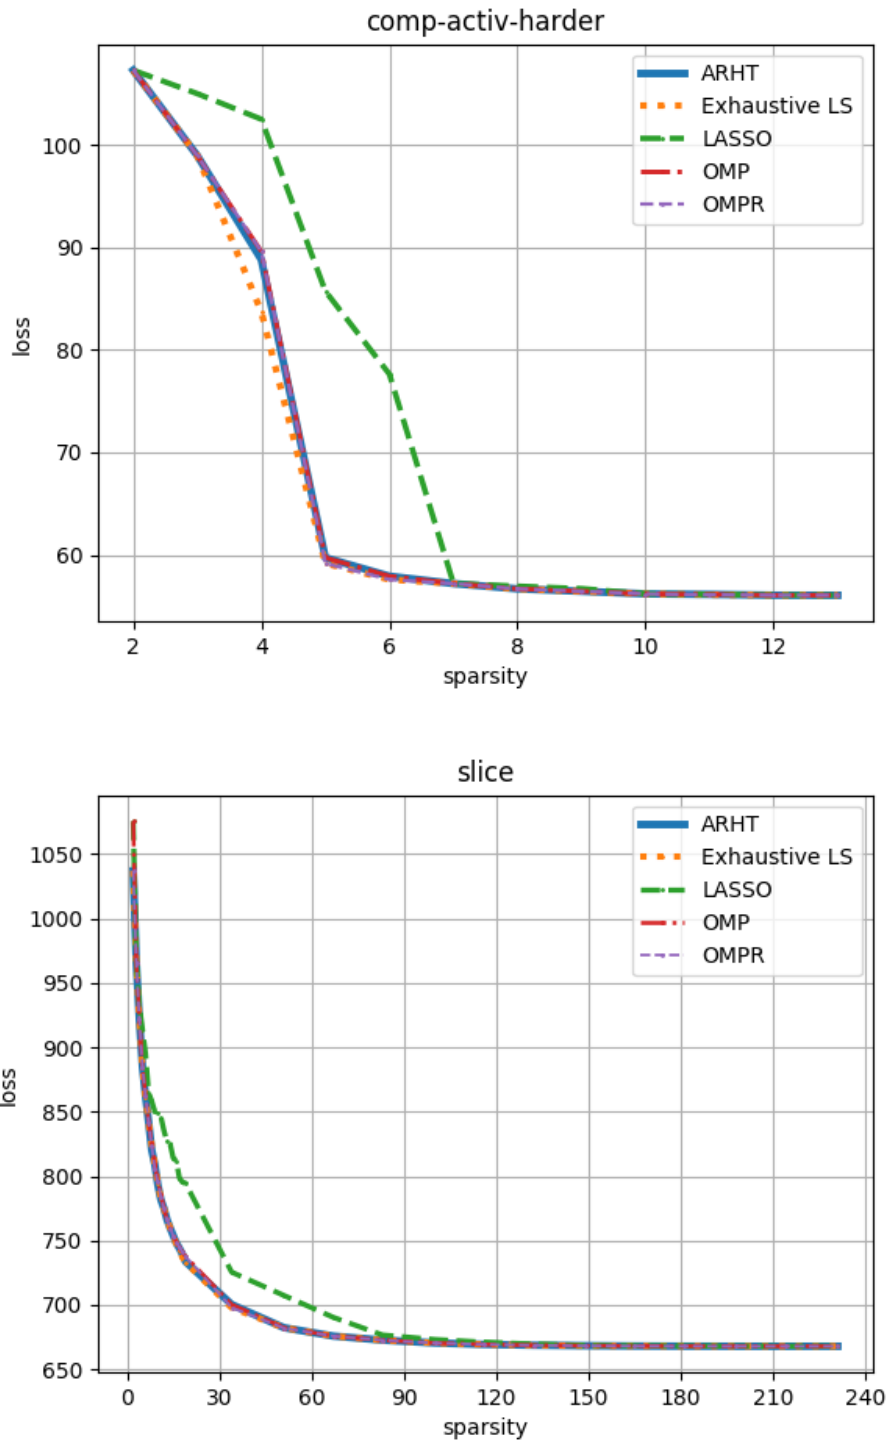


Figure 6-2: Comparison of different algorithms in the Regression data sets *comp-activ-harder* and *slice* using the Linear Regression loss.

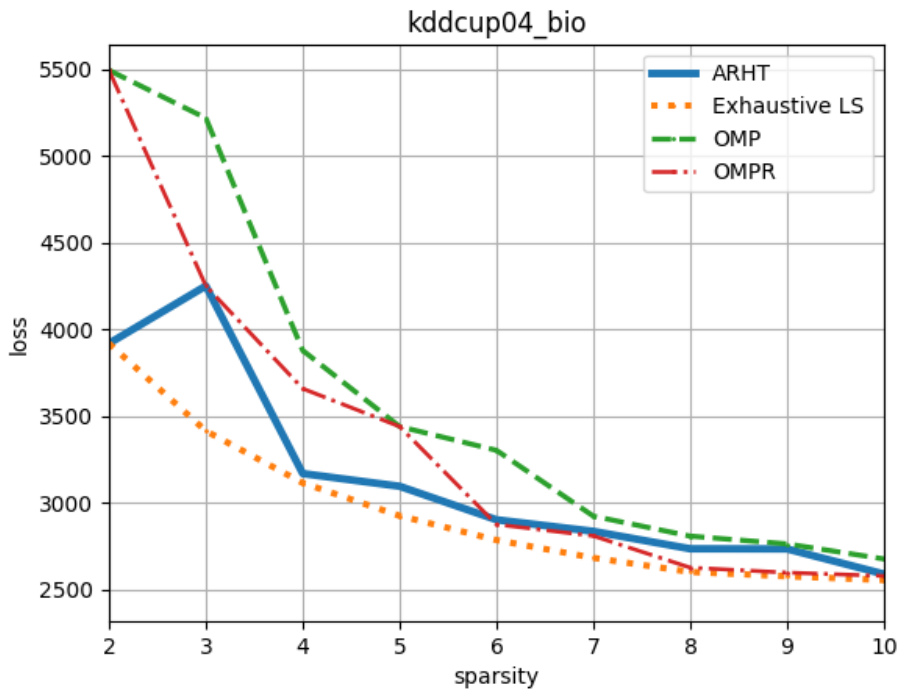
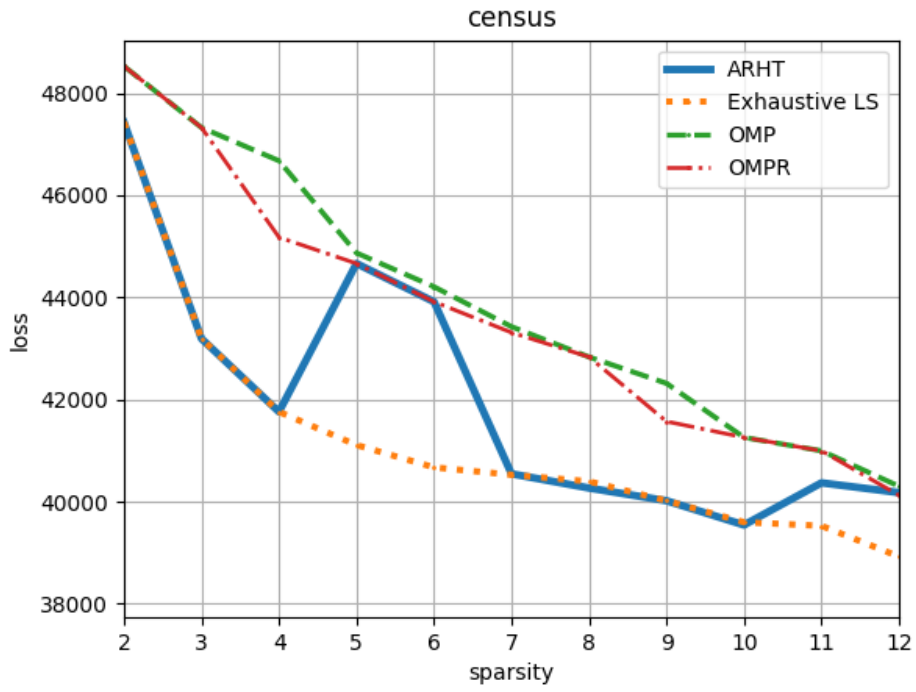


Figure 6-3: Comparison of different algorithms in the Binary classification data sets *census* and *kddcup04_bio* using the Logistic Regression loss.

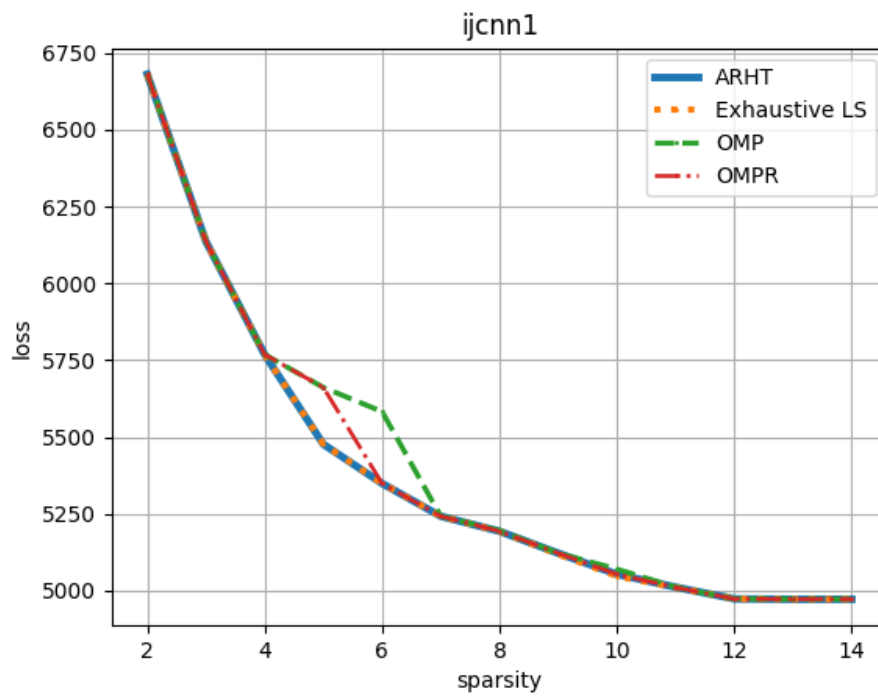



Figure 6-4: Comparison of different algorithms in the Binary classification data sets *letter* and *ijcnn1* using the Logistic Regression loss.

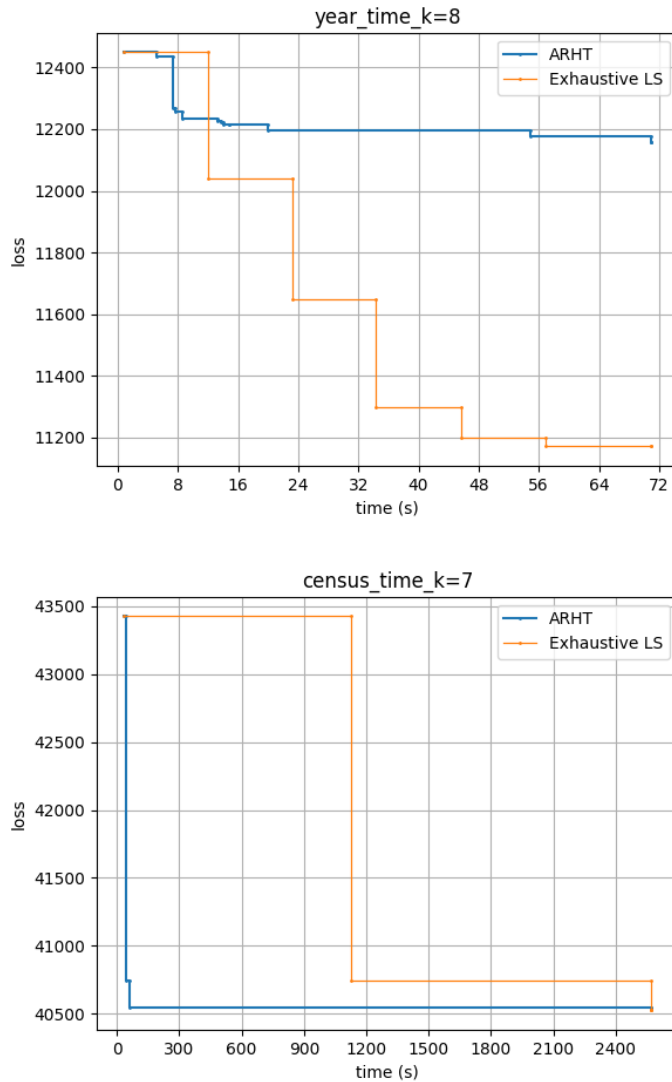


Figure 6-5: Comparison of the loss of a solution vs time elapsed to compute it, between ARHT and Exhaustive Local Search. The first experiment is on the *year* dataset (90 total features) with fixed sparsity 8 and the second is on the *census* dataset (401 total features) with fixed sparsity 7. We note that in the first case Exhaustive Local Search returns significantly sparser solutions without too significant time overhead compared to ARHT. However, in the second case ARHT computes a solution of similar loss to that of Exhaustive Local Search, but around 40 times faster. We attribute this to the fact that the Exhaustive Local Search has an extra n factor in the runtime, so it doesn't scale as well as ARHT as the number of features increases. Note: The reason there are “steps” in the plot is that the solution is improved at discrete time steps, i.e. whenever an insertion and removal from the solution support improves the solution.

For experimental evaluation we used well known and publicly available data sets. Their names and basic properties are outlined in Table 6.2.

Table 6.2: Data sets used for experimental evaluation. The columns are the data set name, the number of examples m , and the number of features n . The data sets can be downloaded here.

NAME	n	d	PROBLEM
KDDCUP04_BIO	145750	74	BINARY
CAL_HOUSING	20639	8	REGRESSION
CENSUS	299284	401	BINARY
COMP-ACTIV-HARDER	8191	12	REGRESSION
IJCNN1	24995	22	BINARY
LETTER	20000	16	BINARY
SLICE	53500	384	REGRESSION
YEAR	463715	90	REGRESSION

6.6.2 Setup Details

Basic Definitions

The two quantities that take part in our experiments are the *sparsity* and the *loss* of a particular solution. We have already defined and discussed the former at length. The latter refers to the training loss for the problems of Linear Regression and Logistic Regression. We let m denote the number of examples and n the number of features in each example.

In the *Linear Regression* task we are given the data set (A, b) , where $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$. The columns of A correspond to features and the rows to examples. The (ℓ_2 *Linear Regression*) *loss* of a solution $x \in \mathbb{R}^n$ is defined as $\ell_2_loss(x) = \frac{1}{2} \|Ax - b\|_2^2$.

In the *Logistic Regression* task we are given the data set (A, b) , where $A \in \mathbb{R}^{m \times n}$, $b \in \{0, 1\}^m$. The columns of A correspond to features and the rows to examples. The (*Logistic Regression*) *loss* of a solution $x \in \mathbb{R}^n$ is defined as

$$\text{logistic_loss}(x) = \sum_{i \in [m]} (-b_i \log \sigma(Ax)_i - (1 - b_i) \log(1 - \sigma(Ax)_i)),$$

where $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ defined as $\sigma(t) = \frac{1}{1 + e^{-t}}$ is the sigmoid function.

Data Pre-processing

We apply a very basic form of pre-processing to the data. More specifically, we use one-hot encoding to turn categorical features into numerical ones. Then, we discard any examples with missing data so that all the entries of A are defined. We also augment the matrix A with an extra all-ones column (i.e. $\vec{1}$) in order to encode the constant (y -intercept) term into A , and we scale all the columns of A so that their ℓ_2 norm is 1. Finally, for the case of ARHT we further augment A in order to encode the regularizer as well. We do this by adding an identity matrix as extra rows. In other words, $A \leftarrow \begin{pmatrix} A \\ I \end{pmatrix}$ and $b \leftarrow \begin{pmatrix} b \\ \vec{0} \end{pmatrix}$.

6.6.3 Implementation Details

The code has been implemented in *python3*, with libraries *numpy*, *sklearn*, and *scipy*.

Inner Optimization Problem

All the algorithms except for LASSO rely on an inner optimization routine in a restricted subset of coordinates in each step. The inner optimization problem consists of solving a standard Linear Regression or Logistic Regression problem using only a submatrix of A defined by a subset of s of its columns. For that, we use *LinearRegression* and *LogisticRegression* from *sklearn.linear_model*. For Logistic Regression we used an LBFGS solver with 1000 iterations.

Overall Algorithm

The LASSO solver we used is *Lasso* from *sklearn.linear_model* with 1000 iterations. As LASSO is not tuned in terms of a required sparsity s , but rather in terms of the regularization parameter α , for each sparsity level we applied binary search on α in order to find a parameter α that gives the required sparsity.

For ARHT, we used a fixed number of 20 iterations at Line 5 of Algorithm 15. In Line 19 of Algorithm 14 we slightly weaken the progress condition to

$$g_{R^t}(x^t) - g_{R^t}(x^{t+1}) \geq \frac{10^{-3}}{s} (g_{R^t}(x^t) - \text{opt}) . \quad (6.17)$$

Furthermore, we do not perform a fixed number of iterations. Instead, we use a stopping criterion: If the progress condition (6.17) is not met and at least half the elements in x^t have already been unregularized, i.e. $|S^t \setminus R^t| \geq \frac{1}{2} |S^t|$, then we stop. If a desirable solution has not been found, it means that this might be an unsuccessful run, and early termination can be used to detect such runs early and re-start, thus improving the runtime. The routine which samples an index i proportional to x_i^2 was implemented by a standard sampling method that uses binary search on i and flips a random coin at each step. This requires computation of interval sums of x_i^2 , which is done by computing partial sums.

Chapter 7

Iterative Hard Thresholding with Adaptive Regularization: Sparser Solutions Without Sacrificing Runtime

7.1 Introduction

Sparse optimization is the task of optimizing a function f over s -sparse vectors, i.e. those with at most s non-zero entries. Examples of such optimization problems arise in machine learning, with the goal to make models smaller for efficiency, generalization, or interpretability reasons, and compressed sensing, where the goal is to recover an s -sparse signal from a small number of measurements. A closely related problem is *low rank optimization*, where the sparsity constraint is instead placed on the spectrum of the solution (which is a matrix). This problem is central in matrix factorization, recommender systems, robust principal components analysis, among other tasks. More generally, *structured sparsity* constraints have the goal of capturing the special structure of a particular task by restricting the set of solutions to those that are “simple” in an appropriate sense. Examples include group sparsity, tree- and graph-structured sparsity. For more on generalized sparsity measures see e.g. [142].

Among the huge number of algorithms that have been developed for the sparse optimization problems, three stand out as the most popular ones:

- The **LASSO** [161], which works by relaxing the ℓ_0 (sparsity) constraint to an ℓ_1 constraint, thus convexifying the problem.
- **Orthogonal matching pursuit** (OMP) [135], which works by building the solution greedily in an incremental fashion.
- **Iterative hard thresholding** (IHT) [19], which performs projected gradient descent on the set of sparse solutions.

Among these, IHT is generally the most efficient, since it has essentially no overhead over plain gradient descent, making it the tool of choice for large-scale applications.

7.1.1 Iterative Hard Thresholding (IHT)

Consider the *sparse convex optimization* problem

$$\min_{\|\mathbf{x}\|_0 \leq s} f(\mathbf{x}), \quad (7.1)$$

where f is convex and $\|\mathbf{x}\|_0$ is the number of non-zero entries in the vector \mathbf{x} , i.e. the sparsity of \mathbf{x} . IHT works by repeatedly performing the following iteration

$$\mathbf{x}^{t+1} = H_{s'}(\mathbf{x}^t - \eta \cdot \nabla f(\mathbf{x}^t)), \quad (7.2)$$

where $H_{s'}$ is the *hard thresholding operator* that zeroes out all but the top s' entries, for some (potentially relaxed) sparsity level s' , and $\eta > 0$ is the step size.

As (7.1) is known to be NP-hard [127] and even hard to approximate [63], an extra assumption needs to be made for the performance of the algorithm to be theoretically evaluated in a meaningful way. The most common assumption is that the (*restricted*) *condition number* of f is bounded by κ (or the *restricted isometry property* constant is bounded by δ [31]), but other assumptions have been studied as well, such as incoherence [47] and weak supermodularity [109]. The performance is then measured in terms of the sparsity s' of the returned solution, as well as its error (value of f).

As it is known [89], IHT is guaranteed to return an $s' = O(s\kappa^2)$ -sparse solution \mathbf{x} with $f(\mathbf{x}) \leq f(\mathbf{x}^*) + \varepsilon$. In fact, as we show in Section 9.5.5, the κ^2 factor cannot be improved in the analysis. Recently, [14] presented an algorithm called ARHT, which improves the sparsity to $s' = O(s\kappa)$. However, their algorithm is much less efficient than IHT, for many reasons. So the question emerges:

Is there a sparse convex optimization algorithm that returns $O(s\kappa)$ -sparse solutions, but whose runtime efficiency is comparable to IHT?

The main contribution of our work is to show that this goal can be achieved, and done so by a surprisingly simple tweak to IHT.

7.1.2 Reconciling Sparsity and Efficiency: Regularized IHT

Our main result is the following theorem, which states that running IHT on an *adaptively regularized* objective function returns $O(s\kappa)$ -sparse solutions that are ε -optimal in function value, while having no significant runtime overhead over plain gradient descent.

Theorem 7.1.1 (Regularized IHT). *Let $f \in \mathbb{R}^n \rightarrow \mathbb{R}$ be a convex function that is β -smooth and α -strongly convex¹, with condition number $\kappa = \beta/\alpha$, and \mathbf{x}^* be an (unknown) s -sparse solution. Then, running Algorithm 16 with $\eta = (2\beta)^{-1}$ and $c = s'/(4T)$ for*

$$T = O\left(\kappa \log \frac{f(\mathbf{x}^0) + (\beta/2) \|\mathbf{x}^0\|_2^2 - f(\mathbf{x}^*)}{\varepsilon}\right)$$

iterations starting from an arbitrary $s' = O(s\kappa)$ -sparse solution \mathbf{x}^0 , the algorithm returns an s' -sparse solution \mathbf{x}^T such that $f(\mathbf{x}^T) \leq f(\mathbf{x}^) + \varepsilon$. Furthermore, each iteration requires $O(1)$ evaluations of f , ∇f , and $O(n)$ additional time.*

¹The theorem also holds if the smoothness and strong convexity constants are replaced by $(s' + s)$ -restricted smoothness and strong convexity constants.

To achieve this result, we significantly refine and generalize the *adaptive regularization* technique of [14]. This refined version fixes many of the shortcomings of the original, by (i) not requiring *re-optimization* in every iteration (a relic of OMP-style algorithms), (ii) taking $\tilde{O}(\kappa)$ instead of $\tilde{O}(s\kappa)$ iterations, (iii) being *deterministic*, (iv) not requiring knowledge of the optimal function value $f(x^*)$ thus avoiding the overhead of an outer binary search, and (v) being more easily generalizable to other settings, like low rank minimization.

In short, our main idea is to run IHT on a regularized function

$$g(\mathbf{x}) = f(\mathbf{x}) + (\beta/2) \|\mathbf{x}\|_{\mathbf{w},2}^2 ,$$

where $\|\mathbf{x}\|_{\mathbf{w},2}^2 = \sum_{i=1}^n w_i x_i^2$ and \mathbf{w} are non-negative weights. These weights change dynamically during the algorithm, in a way that depends on the value of \mathbf{x} . The effect is that now the IHT step will instead be given by

$$\mathbf{x}^{t+1} = H_{s'} \left((\mathbf{1} - 0.5\mathbf{w}^t) \mathbf{x}^t - \eta \cdot \nabla f(\mathbf{x}^t) \right) ,$$

which is almost the same as (7.2), except that it has an extra term that biases the solution towards $\mathbf{0}$. Additionally, in each step the weights \mathbf{w}^t are updated based on the current solution as

$$w_i^{t+1} = \left(w_i^t \cdot \left(\mathbf{1} - c \cdot \frac{w_i^t (x_i^t)^2}{\|\mathbf{x}^t\|_{\mathbf{w}^t,2}^2} \right) \right)_{\geq 1/2}$$

for some parameter $c > 0$, where $(\cdot)_{\geq 1/2}$ denotes zeroing out all the entries that are $< 1/2$ and keeping the others intact.

In Section 7.2, we will go over the central ideas of our refined adaptive regularization technique, and also explain how it can be extended to deal with more general sparsity measures.

7.1.3 Beyond Sparsity: Low Rank Optimization

As discussed, our new techniques transfer to the problem of minimizing a convex function under a rank constraint. In particular, we prove the following theorem:

Theorem 7.1.2 (Adaptive Regularization for Low Rank Optimization). *Let $f \in \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ be a convex function with condition number κ and consider the low rank minimization problem*

$$\min_{\text{rank}(\mathbf{A}) \leq r} f(\mathbf{A}) . \tag{7.3}$$

For any error parameter $\varepsilon > 0$, there exists a polynomial time algorithm that returns a matrix \mathbf{A} with $\text{rank}(\mathbf{A}) \leq O\left(r \left(\kappa + \log \frac{f(\mathbf{O}) - f(\mathbf{A}^)}{\varepsilon}\right)\right)$ and $f(\mathbf{A}) \leq f(\mathbf{A}^*) + \varepsilon$, where \mathbf{O} is the all-zero matrix and \mathbf{A}^* is any rank- r matrix.*

This result can be compared to the Greedy algorithm of [13], which works by incrementally adding a rank-1 component to the solution and achieves rank $O(r\kappa \log \frac{f(\mathbf{O}) - f(\mathbf{A}^*)}{\varepsilon})$, as well as their Local Search algorithm, which works by simultaneously adding a rank-1 component and removing another, and achieves rank $O(r\kappa^2)$. In contrast, our Theorem 7.1.2 returns a solution with rank $O\left(r \left(\kappa + \log \frac{f(\mathbf{O}) - f(\mathbf{A}^*)}{\varepsilon}\right)\right)$.

7.1.4 Related Work

The sparse optimization and compressed sensing literature has a wealth of different algorithms and analyses. Examples include the seminal paper of [31] on recovery with LASSO and followup works [64], the CoSaMP algorithm [128], orthogonal matching pursuit and variants [127, 146, 87, 14] iterative hard thresholding [19, 89], hard thresholding pursuit [65, 169, 148, 149], partial hard thresholding [88], and message passing algorithms [48]. For a survey on compressed sensing, see [21, 68].

A family of algorithms that is closely related to IHT are Frank-Wolfe (FW) methods [69], which have been used for dealing with generalized sparsity constraints [86]. The basic version can be viewed as a variant of OMP without re-optimization in each iteration. Block-FW methods are more resemblant of IHT without the projection step, see e.g. [5] for an application to the low rank minimization problem.

[112] presented an interesting connection between hard and soft thresholding algorithms by studying a concavity property of the thresholding operator, and proposed new thresholding operators.

Recently it has been shown [138] that IHT can be guaranteed to work for sparse optimization of *non-convex* functions, under appropriate assumptions. In particular, [138] studies a stochastic version of IHT for sparse deep learning problems, from both a theoretical and practical standpoint.

7.2 The Adaptive Regularization Method

Consider the sparse optimization problem

$$\min_{\|\mathbf{x}\|_0 \leq s} f(\mathbf{x}) \tag{7.4}$$

on a convex function f with condition number at most κ , and an optimal solution \mathbf{x}^* that is supported on the set of indices $S^* \subseteq [n]$.

The main hurdle towards solving this problem is that it is NP hard. Therefore, it is common to relax it by a factor depending on κ . In fact, IHT requires relaxing the sparsity constraint by a factor of $O(\kappa^2)$ (i.e. $\|\mathbf{x}\|_0 \leq O(s\kappa^2)$), in order to return a near-optimal solution. Also, the κ^2 factor is tight for IHT (see Appendix 9.5.5).

Remark We state all our results in terms of the condition number κ , even though the statements can be strengthened to depend on the *restricted* condition number $\kappa_{s'+s}$, specifically the condition number restricted on $(s' + s)$ -sparse directions. We state our results in this weaker form for clarity of presentation.

7.2.1 Regularized IHT

Perhaps surprisingly, there is a way to *regularize* the objective by a weighted ℓ_2 norm so that running IHT on the new objective will only require relaxing the sparsity by $O(\kappa)$:

$$\min_{\|\mathbf{x}\|_0 \leq s} f(\mathbf{x}) + (\beta/2) \|\mathbf{x}\|_{\mathbf{w},2}^2. \tag{7.5}$$

One way to do this is by setting the weights \mathbf{w} to be 1 everywhere except in the indices from S^* , where it is set to 0. An inquisitive reader will protest that this is not a very useful statement, since it requires knowledge of S^* , which was our goal to begin with. In fact, we could just as easily have used the regularizer $(\beta/2) \|\mathbf{x} - \mathbf{x}^*\|_2^2$, thus penalizing everything that is far from the optimum!

7.2.2 Learning Weights

Our main contribution is to show that the optimal weights \mathbf{w} can in fact be *learned* in the duration of the algorithm². More precisely, consider running IHT starting from the setting of $\mathbf{w} = \mathbf{1}$. The regularized objective (7.5) is now $O(1)$ -conditioned, which is great news. On the other hand, (7.5) is not what we set out to minimize. In other words, even though this approach might work great for minimizing (7.5), it might (and generally will) fail to achieve sufficient decrease in (7.4)—one could view this as the algorithm getting trapped in a local minimum.

Our main technical tool is to characterize these local minima, by showing that they can only manifest themselves if the current solution \mathbf{x} satisfies the following condition:

$$\|\mathbf{x}_{S^*}\|_{\mathbf{w},2}^2 \geq \Omega(\kappa^{-1}) \|\mathbf{x}\|_{\mathbf{w},2}^2. \quad (7.6)$$

In words, this means that a significant fraction of the mass of the current solution lies in the support S^* of the optimal solution. Interestingly, this gives us enough information based on which to update the regularization weights \mathbf{w} in a way that the sum of weights in S^* drops fast enough compared to the total sum of weights. This implies that the vector \mathbf{w} moves in a direction that correlates with the direction of the optimal weight vector.

These are the core ideas needed to bring the sparsity overhead of IHT from $O(\kappa^2)$ down to $O(\kappa)$.

7.2.3 Beyond Sparsity: Learning Subspaces

One can summarize the approach of the previous section in the following more general way: If we know that the optimal solution \mathbf{x}^* lies in a particular low-dimensional subspace (in our case this was the span of $\mathbf{1}_i$ for all $i \in S^*$), then we can define a regularization term that penalizes all the solutions based on their distance to that subspace. Of course, this subspace is unknown to us, but we can try to adaptively modify the regularization term every time the algorithm gets stuck, just as we did in the previous section.

More concretely, given a collection \mathcal{A} of unit vectors from \mathbb{R}^n (commonly called *atoms*), we define the following problem:

$$\min_{\text{rank}_{\mathcal{A}}(\mathbf{x}) \leq r} f(\mathbf{x}), \quad (7.7)$$

where $\text{rank}_{\mathcal{A}}(\mathbf{x})$ is the smallest number of vectors from \mathcal{A} such that \mathbf{x} can be written as their linear combination. We can pick $\mathcal{A} = \{\mathbf{1}_1, \mathbf{1}_2, \dots, \mathbf{1}_n\}$ to obtain the sparse optimization problem, $\mathcal{A} = \{\text{vec}(\mathbf{u}\mathbf{v}^\top) \mid \|\mathbf{u}\|_2 = \|\mathbf{v}\|_2 = 1\}$ for the low rank minimization problem, and other choices of \mathcal{A} can capture more sophisticated problem constraints such as graph structure. Defining an IHT variant for these more general settings is usually straightforward, although the analysis for even obtaining a rank overhead of $O(\kappa^2)$ does not trivially follow and depends on the structure of \mathcal{A} .

So, how would a regularizer look in this more general setting? Given our above discussion, it is fairly simple to deduce it. Consider a decomposition of \mathbf{x} as the sum of rank-1 components from $\mathcal{A} = \{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_{|\mathcal{A}|}\}$:

$$\mathbf{x} = \sum_{i \in S} \mathbf{a}_i,$$

²The idea of adaptively learning regularization weights looks on the surface similar to adaptive gradient algorithms such as AdaGrad [50]. An important difference is that these algorithms regularize the function around the current solution, while we regularize it around the origin. Still, this is a potentially intriguing connection that deserves to be investigated further.

where $\text{rank}_{\mathcal{A}}(\mathbf{a}_i) = 1$, and let $L^* = \text{span}(\{\mathbf{a}_i \mid i \in S^*\})$ be a low-dimensional subspace that contains the optimal solution and L_{\perp}^* is its complement. We can then define the regularizer

$$\Phi^*(\mathbf{x}) = (\beta/2) \sum_{i \in S} \left\| \mathbf{\Pi}_{L_{\perp}^*} \mathbf{a}_i \right\|_2^2,$$

where $\mathbf{\Pi}_{L_{\perp}^*}$ is the orthogonal projection onto the subspace perpendicular to L^* —in other words $\left\| \mathbf{\Pi}_{L_{\perp}^*} \mathbf{a}_i \right\|_2$ is the ℓ_2 distance from \mathbf{a}_i to L^* . An equivalent but slightly more concise way is to write:

$$\Phi^*(\mathbf{x}) = (\beta/2) \left\langle \mathbf{\Pi}_{L_{\perp}^*}, \sum_{i \in S} \mathbf{a}_i \mathbf{a}_i^{\top} \right\rangle.$$

Then, we can replace the unknown projection matrix $\mathbf{\Pi}_{L_{\perp}^*}$ by a weight matrix \mathbf{W} initialized at \mathbf{I} , and proceed by adaptively modifying \mathbf{W} as we did in the previous section.

It should be noted that the full analysis of this framework is not automatic for general \mathcal{A} , and there are several technical challenges that arise depending on the choice of \mathcal{A} . In particular, it does not directly apply to the low rank minimization case, and we end up using a different choice of regularizer. However, the discussion in this section should serve as a basic framework for improving the IHT analysis in more general settings, as in particular it did to motivate the low rank optimization analysis that we will present in Section 7.4.

7.3 Sparse Optimization Using Regularized IHT

The main result of this section is an efficient algorithm for sparse optimization of convex functions that, even though is a slight modification of IHT, improves the sparsity by an $O(\kappa)$ factor, where κ is the condition number. The regularized IHT algorithm is presented in Algorithm 16 and its analysis is in Theorem 7.1.1, whose proof can be found in Appendix 9.5.2.

Algorithm 16 Regularized IHT

\mathbf{x}^0 : initial s' -sparse solution
 $\mathbf{w}^0 = \mathbf{1}$: initial regularization weights
 η : step size, T : #iterations
 c : weight step size
for $t = 0 \dots T - 1$ **do**
 $\mathbf{x}^{t+1} = H_{s'}((\mathbf{1} - 0.5\mathbf{w}^t)\mathbf{x}^t - \eta \cdot \nabla f(\mathbf{x}^t))$
 $\mathbf{w}^{t+1} = \left(\mathbf{w}^t - c \cdot (\mathbf{w}^t \mathbf{x}^t)^2 / \|\mathbf{x}^t\|_{\mathbf{w}^t, 2}^2 \right)_{\geq 1/2}$
 if $f(\mathbf{x}^{t+1}) + (4\eta)^{-1} \|\mathbf{x}^{t+1}\|_{\mathbf{w}^{t+1}, 2}^2 > f(\mathbf{x}^t) + (4\eta)^{-1} \|\mathbf{x}^t\|_{\mathbf{w}^{t+1}, 2}^2$ **then**
 $\mathbf{x}^{t+1} = \mathbf{x}^t$ ▷ In practice there is no need to perform this step.

Theorem 7.1.1 (Regularized IHT). *Let $f \in \mathbb{R}^n \rightarrow \mathbb{R}$ be a convex function that is β -smooth and α -strongly convex, with condition number $\kappa = \beta/\alpha$, and \mathbf{x}^* be an (unknown) s -sparse solution. Then, running Algorithm 16 with $\eta = (2\beta)^{-1}$ and $c = s'/(4T)$ for*

$$T = O \left(\kappa \log \frac{f(\mathbf{x}^0) + (\beta/2) \|\mathbf{x}^0\|_2^2 - f(\mathbf{x}^*)}{\varepsilon} \right)$$

iterations starting from an arbitrary $s' = O(s\kappa)$ -sparse solution \mathbf{x}^0 , the algorithm returns an s' -sparse solution \mathbf{x}^T such that $f(\mathbf{x}^T) \leq f(\mathbf{x}^*) + \varepsilon$. Furthermore, each iteration requires $O(1)$ evaluations of f , ∇f , and $O(n)$ additional time.

The main ingredient for proving Theorem 7.1.1 is Lemma 7.3.1, which states that each step of the algorithm either makes substantial (multiplicative) progress in an appropriately regularized function $f(\mathbf{x}) + (\beta/2) \|\mathbf{x}\|_{\mathbf{w},2}^2$, or a significant fraction of the mass of \mathbf{x}^2 lies in S^* , which is the support of the target solution. This latter condition allows us to adapt the weights \mathbf{w} in order to obtain a new regularization function that penalizes the target solution less. The proof of the lemma can be found in Appendix 9.5.3.

Lemma 7.3.1 (Regularized IHT step progress). *Let $f \in \mathbb{R}^n \rightarrow \mathbb{R}$ be a convex function that is β -smooth and α -strongly convex, $\kappa = \beta/\alpha$ be its condition number, and \mathbf{x}^* be any s -sparse solution. Given any s' -sparse solution $\mathbf{x} \in \mathbb{R}^n$ where*

$$s' \geq (128\kappa + 2)s$$

and a weight vector $\mathbf{w} \in (\{0\} \cup [1/2, 1])^n$ such that $\|\mathbf{w}\|_1 \geq n - s'/2$, we make the following update:

$$\mathbf{x}' = H_{s'} \left((\mathbf{1} - 0.5\mathbf{w})\mathbf{x} - (2\beta)^{-1} \nabla f(\mathbf{x}) \right) .$$

Then, at least one of the following two conditions holds:

- Updating \mathbf{x} makes regularized progress:

$$g(\mathbf{x}') \leq g(\mathbf{x}) - (16\kappa)^{-1} (g(\mathbf{x}) - f(\mathbf{x}^*)) ,$$

where

$$g(\mathbf{x}) := f(\mathbf{x}) + (\beta/2) \|\mathbf{x}\|_{\mathbf{w},2}^2$$

is the ℓ_2 -regularized version of f with weights given by \mathbf{w} . Note: The regularized progress statement is true as long as \mathbf{x} is suboptimal, i.e. $g(\mathbf{x}) > f(\mathbf{x}^*)$. Otherwise, we just have $g(\mathbf{x}') \leq g(\mathbf{x})$.

- \mathbf{x} is significantly correlated to the optimal support $S^* := \text{supp}(\mathbf{x}^*)$:

$$\|\mathbf{x}_{S^*}\|_{\mathbf{w},2}^2 \geq (4\kappa + 6)^{-1} \|\mathbf{x}\|_{\mathbf{w},2}^2 ,$$

and the regularization term restricted to S^* is non-negligible:

$$(\beta/2) \|\mathbf{x}_{S^*}\|_{\mathbf{w},2}^2 \geq (8\kappa + 8)^{-1} (g(\mathbf{x}) - f(\mathbf{x}^*)) .$$

Comparison to ARHT. The ARHT algorithm of [14] is also able to achieve a sparsity bound of $O(s\kappa)$. However, their algorithm is not practically desirable for a variety of reasons.

- First of all, it follows the OMP (more accurately, OMP with Removals) paradigm, which makes local changes to the support of the solution by inserting or removing a single element of the support, and then *fully re-optimizing* the function on its restriction to this support. Even though the support will generally be very small compared to the ambient dimension n , this is still a significant runtime overhead. In contrast, regularized IHT does not require re-optimization.

Additionally, the fact that in the ARHT only one new element is added at a time leads to an iteration count that scales with $s\kappa$, instead of the κ of regularized IHT. This is a significant speedup, since both algorithms have to evaluate the gradient in each iteration. Therefore, regularized IHT will require $O(s)$ times fewer gradient evaluations.

- When faced with the non-progress condition, in which the regularized function value does not decrease sufficiently, ARHT moves by selecting a random index i with probability proportional to x_i^2 , and proceeds to *unregularize* this element, i.e. remove it from the sum of regularization terms. Instead, our algorithm is completely deterministic. This is achieved by allowing a *weighted* regularization term, and gradually reducing the regularization weights instead of dropping terms.
- ARHT requires knowledge of the optimal function value $f(\mathbf{x}^*)$. The reason is that in each iteration they need to gauge whether enough progress was made in reducing the value of the regularized function g , compared to how far it is from the optimal function value. If so, they would perform the unregularization step. In contrast, our analysis does not require these two cases (updates to \mathbf{x} or \mathbf{w}) to be exclusive, and in fact simultaneously updates both, regardless of how much progress was made in g . Thus, our algorithm avoids the expensive overhead of an outer binary search over the optimal value $f(\mathbf{x}^*)$.

For all these reasons, as well as its striking simplicity, we believe that regularized IHT can prove to be a useful practical sparse optimization tool.

7.4 Low Rank Optimization Using Regularized Local Search

In this section we present a regularized local search algorithm for low rank optimization of convex functions, that returns an ε -optimal solution with rank $O\left(r\left(\kappa + \log \frac{f(\mathbf{O}) - f(\mathbf{A}^*)}{\varepsilon}\right)\right)$, where r is the target rank. The algorithm is based on the Local Search algorithm of [13], but also uses adaptive regularization, which leads to a lot new technical hurdles that are addressed in the analysis. This is presented in Theorem 7.1.2 and proved in Appendix 9.5.4.

Theorem 7.1.2 (Adaptive Regularization for Low Rank Optimization). *Let $f \in \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ be a convex function with condition number κ and consider the low rank minimization problem*

$$\min_{\text{rank}(\mathbf{A}) \leq r} f(\mathbf{A}). \quad (7.8)$$

For any error parameter $\varepsilon > 0$, there exists a polynomial time algorithm that returns a matrix \mathbf{A} with $\text{rank}(\mathbf{A}) \leq O\left(r\left(\kappa + \log \frac{f(\mathbf{O}) - f(\mathbf{A}^)}{\varepsilon}\right)\right)$ and $f(\mathbf{A}) \leq f(\mathbf{A}^*) + \varepsilon$, where \mathbf{O} is the all-zero matrix and \mathbf{A}^* is any rank- r matrix.*

Discussion about ε dependence. Some of the technical issues in the rank case have to do with operator *non-commutativity* and thus pose no issue in the sparsity case. In particular, the extra $\log \frac{f(\mathbf{O}) - f(\mathbf{A}^*)}{\varepsilon}$ dependence in the rank comes exactly because of these issues. However, we think that it should be possible to completely remove this dependence in the future by a more careful analysis.

Discussion about computational efficiency. We note that the goal of this section is to show an improved rank bound, and not to argue about the computational efficiency of such an algorithm. It might be possible to derive an efficient algorithm by transforming the proof in Theorem 7.1.2 into

a proof for a matrix IHT algorithm, which might be significantly more efficient, as it will not require solving linear systems in each iteration. Still, there are a lot of remaining issues to be tackled, as currently the algorithm requires computing multiple singular value decompositions and orthogonal projections in each iteration. Therefore working on a computationally efficient algorithm that can guarantee a rank of $O(r\kappa)$ is a very interesting direction for future research.

Matrix regularizer Getting back into the main ingredients of Theorem 7.1.2, we describe the choice of our regularizer. As we are working over general rectangular matrices, we use *two* regularizers, one for the left singular vectors and one for the right singular vectors of \mathbf{A} . Concretely, given two weight matrices \mathbf{Y} , \mathbf{W} such that $\mathbf{O} \preceq \mathbf{Y} \preceq \mathbf{I}$, $\mathbf{O} \preceq \mathbf{W} \preceq \mathbf{I}$, we define

$$\Phi(\mathbf{A}) = (\beta/4) \left(\langle \mathbf{W}, \mathbf{A}\mathbf{A}^\top \rangle + \langle \mathbf{Y}, \mathbf{A}^\top \mathbf{A} \rangle \right),$$

where β is a bound on the smoothness of f . The gradient of the regularized function is

$$\nabla g(\mathbf{A}) = \nabla f(\mathbf{A}) + (\beta/2) (\mathbf{W}\mathbf{A} + \mathbf{A}\mathbf{Y}),$$

and the new solution $\bar{\mathbf{A}}$ is defined as

$$\bar{\mathbf{A}} = H_{s'-1}(\mathbf{A}) - \eta H_1(\nabla g(\mathbf{A})),$$

where we remind that the thresholding operator $H_r : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{m \times n}$ that is used in the algorithm returns the top r components of the singular value decomposition of a matrix, i.e. given $\mathbf{M} = \sum_{i=1}^k \lambda_i \mathbf{u}_i \mathbf{v}_i^\top$, where $\lambda_1 \geq \dots \geq \lambda_k$ are the singular values and $r \leq k$, $H_r(\mathbf{M}) = \sum_{i=1}^r \lambda_i \mathbf{u}_i \mathbf{v}_i^\top$. In other words, we drop the bottom rank-1 component of \mathbf{A} and add the top rank-1 component of the gradient.

After taking a step, we re-optimize over matrices with the current left and right singular space, also known as performing a fully corrective step, as in [144, 13]. To do this, we first compute the SVD $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ of $\bar{\mathbf{A}}$ and then solve the optimization problem $\min_{\mathbf{A}=\mathbf{U}\mathbf{X}\mathbf{V}^\top} g^t(\mathbf{A})$. For simplicity we assume that this optimization problem can be solved exactly, but the analysis can be modified to account for the case when we have an approximate solution and we are only given a bound on the norm of the gradient (projected onto the relevant subspace), i.e. $\|\mathbf{\Pi}_{\text{im}(\mathbf{U})} \nabla g^t(\mathbf{A}) \mathbf{\Pi}_{\text{im}(\mathbf{V})}\|_F$.

Whenever there is not enough progress, we make the following updates on the weight matrices \mathbf{W} and \mathbf{Y} :

$$\begin{aligned} \mathbf{W}' &= \mathbf{W} - \mathbf{W}\mathbf{A}\mathbf{A}^\top \mathbf{W} / \langle \mathbf{W}, \mathbf{A}\mathbf{A}^\top \rangle \\ \mathbf{Y}' &= \mathbf{Y} - \mathbf{Y}\mathbf{A}^\top \mathbf{A} \mathbf{Y} / \langle \mathbf{Y}, \mathbf{A}^\top \mathbf{A} \rangle. \end{aligned}$$

The full algorithm description is in Algorithm 25, Appendix 9.5.4. In the algorithm description we assume that $f(\mathbf{A}^*)$ is known. This assumption can be removed by performing binary search over this value, as in [14].

7.5 Experiments

Introduction. In this section we present numerical experiments in order to compare the performance of IHT and regularized IHT (Algorithm 16) in training sparse linear models. In particular, we will look at the tasks of linear regression and logistic regression using both real and synthetic data.

In the former, we are given a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, where each row represents an example and each column a feature, and a vector $\mathbf{b} \in \mathbb{R}^m$ that represents the ground truth outputs, and our objective is to minimize the ℓ_2 loss

$$(1/2) \|\mathbf{Ax} - \mathbf{b}\|_2^2 .$$

In logistic regression, \mathbf{b} has binary instead of real entries, and our objective is to minimize the logistic loss

$$\underbrace{- \sum_{i=1}^m (b_i \log \sigma(\mathbf{Ax})_i + (1 - b_i) \log(1 - \sigma(\mathbf{Ax})_i))}_{l(\mathbf{x})} ,$$

where $\sigma(z) = (1 + e^{-z})^{-1}$ is the sigmoid function. As is common, we look at the *regularized* logistic regression objective:

$$l(\mathbf{x}) + (\rho/2) \|\mathbf{x}\|_2^2 ,$$

for some $\rho > 0$. For our experiments we use $\rho = 0.1$.

Preprocessing and choice of parameters. The only preprocessing we perform is to center the columns of \mathbf{A} , i.e. we subtract the mean of each column from each entry of the column, and then scale the columns to unit ℓ_2 norm. This ensures that for any sparsity parameter $s' \in [n]$, the function f is s' -smooth when restricted to s' -sparse directions, or in other words the s' -restricted smoothness constant of f is at most s' . Thus we set our smoothness estimate to $\beta := s'$. Our smoothness estimate β influences the (regularized) IHT algorithm in two ways. First, as the step size of the algorithm is given by $1/\beta$, a value of β that is too large can slow down the algorithm, or even get it stuck to a local minimum. Second, the strength of the regularization term in regularized IHT should be close to the $(s + s')$ -restricted smoothness constant, as shown in the analysis of Theorem 7.1.1.

Even though having a perfectly accurate estimate of the smoothness constant is not necessary, a more accurate estimate improves the performance of the algorithm. In fact, the estimate $1/s'$ for the step size is generally too conservative. When used in practice, one should either tune this parameter or use a variable/adaptive step size to achieve the best results.

For the weight step size of regularized IHT, we set the weight step size to $c = s'/T$, but we also experiment with how changing c affects the performance of the algorithm. The downside of this setting is that it requires knowing the number of iterations a priori. However, in practice one could tune c and then run the algorithm for $O(s'/c)$ iterations. Note that ideally, based on the theoretical analysis, T would be proportional to the restricted condition number of f , however this quantity is hard to compute in general. Another idea to avoid this in practice could be to let c be a variable step size.

Implementation. Both the IHT and regularized IHT algorithms are incredibly simple, and can be described in a few lines of python code, as can be seen in Figure 9-1, Appendix 9.5.1. Note that in comparison to Algorithm 16 we do not perform the conditional assignment. All the experiments were run on a single 2.6GHz Intel Core i7 core of a 2019 MacBook Pro with 16GB DDR4 RAM using Python 3.9.10.

7.5.1 Real data

We first experiment with real data, specifically the *year* regression dataset from UCI [49] and the *rcv1* binary classification dataset [108], which have been previously used in the literature. Performance on other datasets was similar. In Figure 7-1 (a)-(b) we have a comparison between the error of the solution returned by IHT and regularized IHT for a fixed sparsity level. Specifically, if we let \mathbf{x}^{**} be the (dense) global minimizer of f , we plot the logarithm of the (normalized) *excess loss* $(f(\mathbf{x}) - f(\mathbf{x}^{**}))/f(\mathbf{0})$ against the number of iterations. Note that $f(\mathbf{x}^{**})$ will typically be considerably lower than the loss of the sparse optimum $f(\mathbf{x}^*)$. In order to make a fair comparison, for each algorithm we pick the best fixed step size of the form $2^i/s$ for integer $i \geq 0$, where s is the fixed sparsity level. The best step sizes of IHT and regularized IHT end up being $2/s, 4/s$ respectively for the linear regression example, and $8/s, 16/s$ respectively for the logistic regression example.

We notice that initially regularized IHT has a much higher error than IHT, but after some iterations it is lower than IHT. This phenomenon is to be expected, because the algorithm runs on a regularized function, and so tries to keep not just $f(\mathbf{x})$ but also $\|\mathbf{x}\|_2^2$ small. After some iterations, when the algorithm has learnt regularization weights that are closer to the optimal ones, it converges to sparser solutions than IHT (equivalently, lower error solutions with the same sparsity, which is what is shown in the plot).

In Figure 7-1 (c) we compare IHT and regularized IHT for different sparsity levels on the *year* dataset. If e_1 and e_2 are the excess errors of IHT and regularized IHT respectively, we plot e_2/e_1 , which is the relative excess error of regularized IHT with respect to that of IHT. We notice a reduction of up to 40% on the excess error. In Figure 7-1 (d) we examine the effect of the choice of the weight step size c . We conclude that c can give a tradeoff between runtime and accuracy, as setting it to a large value will lead to faster weight decay and thus resemble IHT, while a small value of c will lead to slow weight decrease, which will lead to more iterations but also potentially recover an improved solution. Here we can see an interesting tradeoff between the number of iterations and the error of the solution that is eventually returned. In particular, the *larger* c is, the *faster* the degradation of regularization weights. Thus, for $c \rightarrow \infty$, the algorithm tends to be the same as IHT. On the other hand, with *smaller* values of c , one can get an improved error rate, but at the cost of a larger number of iterations. This is because the regularization weights decrease slowly, and so in the early iterations of the algorithm the regularization term will account for a significant fraction of the objective function value.

7.5.2 Synthetic data

We now turn to synthetically generated linear regression instances. The first result presented in Figure 7-2 (a) is the hard IHT instance that we derived in our lower bound in Appendix 9.5.5. This experiment shows that there exist examples where, with bad initialization, IHT cannot decrease the objective at all (i.e. is stuck at a local minimum), while regularized IHT with the same initialization manages to reduce the loss by more than 70%.

The second result is a result in the well known setting of sparse signal recovery from linear measurements. We generate a matrix \mathbf{A} with entries that are sampled i.i.d. from the standard normal distribution, an s -sparse signal \mathbf{x} again with entries sampled i.i.d. from the standard normal distribution, and an observed vector $\mathbf{b} := \mathbf{A}\mathbf{x}$. The goal is to recover \mathbf{x} by minimizing the objective

$$f(\mathbf{x}) = (1/2) \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 .$$

In Figure 7-2 (b), we plot the normalized value of this objective, after running both IHT and regularized IHT for the same number of iterations. Here we pick the best step size per instance,

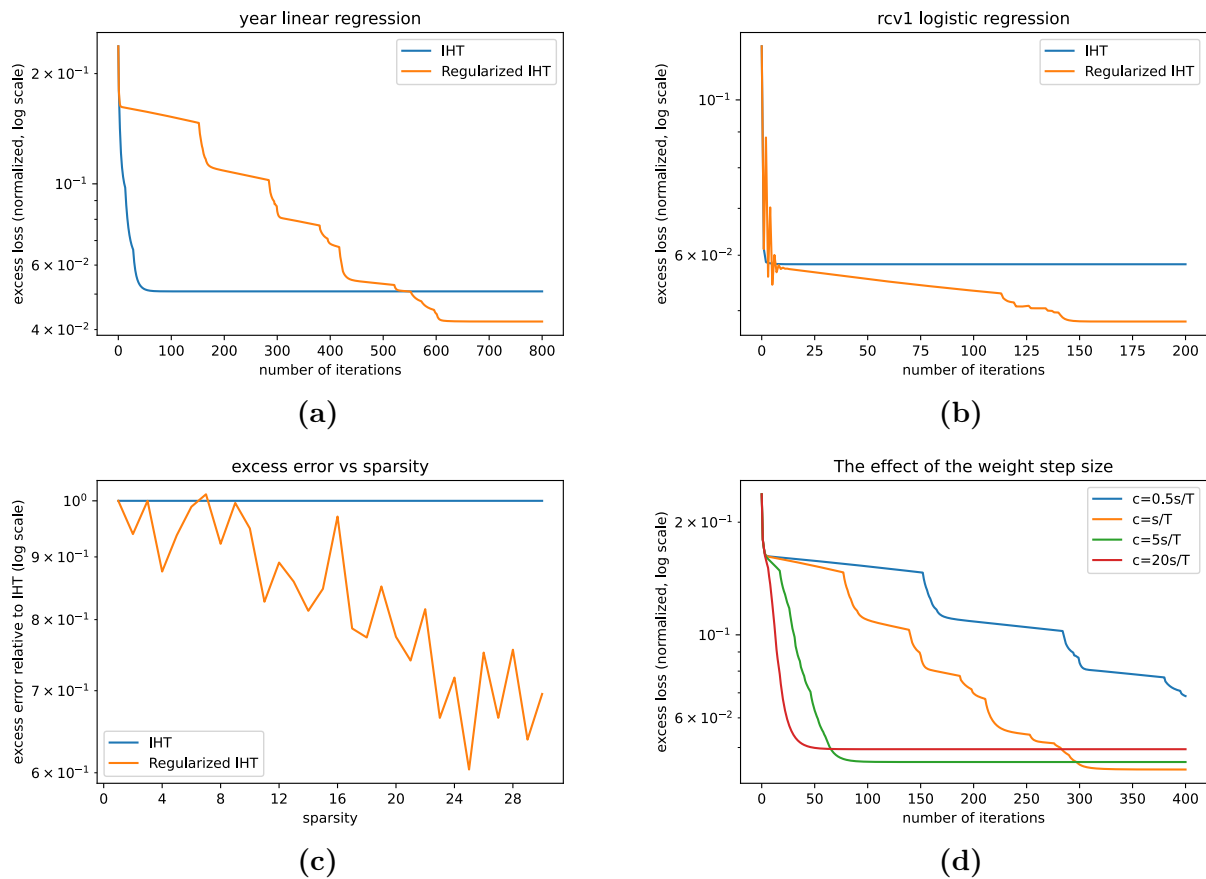


Figure 7-1: **(a)-(b)**: IHT vs Regularized IHT performance on the year and rcv1 datasets with fixed sparsity levels $s = 11$ and $s = 10$ respectively. On the x axis we have number of iterations and on the y axis we have the normalized excess loss (compared to the dense global optimum), in log scale. The excess loss of regularized IHT is less than that of IHT, specifically 17.3% and 17.2% respectively less in the two experiments. **(c)**: Excess error of regularized IHT relative to IHT in the *year* dataset, where sparsity values range from 1 to 30. Both algorithms are run for $T = 800$ iterations. **(d)**: Error rate vs number of iterations of regularized IHT on the *year* dataset with fixed sparsity $s = 11$ and step size $\eta = 4/s$, using different values for the weight step size c .




Figure 7-2: **(a)**: A demonstration of a 80% decrease in loss by using regularized IHT instead of IHT on the hard instance for IHT presented in Section 9.5.5. We have generated the data with a condition number of $\kappa = 20$, and a planted sparse solution with sparsity $s = 2$. The dimension is $n = 842$. It can be observed that, for the given initialization vector, IHT never makes any progress on decreasing the error. In contrast, regularized IHT is able to decrease it by almost a factor of 5. **(b)**: Sparse signal recovery, where \mathbf{A} is an 100×800 measurement matrix, the sparsity level ranges from 1 to 100, and each algorithm is run for 240 iterations. Bands of 1 standard error are shown, after running each data point 20 times independently.

starting from $\eta = 1/s$ and increasing in multiples of 1.2. Also, for each fixed value of s and algorithm, we run the experiments 20 times in order to account for the variance. The results show a superiority in the performance of regularized IHT for the sparse signal recovery task.

Chapter 8

Local Search Algorithms for Rank-Constrained Convex Optimization

8.1 Introduction

Given a real-valued convex function $R : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ on real matrices and a parameter $r^* \in \mathbb{N}$, the *rank-constrained convex optimization* problem consists of finding a matrix $A \in \mathbb{R}^{m \times n}$ that minimizes $R(A)$ among all matrices of rank at most r^* :

$$\min_{\text{rank}(A) \leq r^*} R(A) \tag{8.1}$$

Even though R is convex, the rank constraint makes this problem non-convex. Furthermore, it is known that this problem is NP-hard and even hard to approximate ([127, 63]).

In this work, we propose efficient greedy and local search algorithms for this problem. Our contribution is twofold:

1. We provide theoretical analyses that bound the rank and objective value of the solutions returned by the two algorithms in terms of the *rank-restricted condition number*, which is the natural generalization of the condition number for low rank subspaces. The results are significantly stronger than previous known bounds for this problem.
2. We experimentally demonstrate that, after careful performance adjustments, the proposed general-purpose greedy and local search algorithms have superior performance to other methods, even for some of those that are tailored to a particular problem. Thus, these algorithms can be considered as a general tool for rank-constrained convex optimization and a viable alternative to methods that use convex relaxations or alternating minimization.

The rank-restricted condition number Similarly to the work in sparse convex optimization, a restricted condition number quantity has been introduced as a reasonable assumption on R . If we let ρ_r^+ be the maximum smoothness bound and ρ_r^- be the minimum strong convexity bound only along rank- r directions of R (these are called rank-restricted smoothness and strong convexity respectively), the rank-restricted condition number is defined as $\kappa_r = \frac{\rho_r^+}{\rho_r^-}$. If this quantity is bounded, one can efficiently find a solution A with $R(A) \leq R(A^*) + \varepsilon$ and $\text{rank } r = O(r^* \cdot \kappa_{r+r^*} \frac{R(\mathbf{0})}{\varepsilon})$ using a greedy algorithm ([144]). However, this is not an ideal bound since the rank scales linearly with $\frac{R(\mathbf{0})}{\varepsilon}$, which can be particularly high in practice. Inspired by the analogous literature on sparse convex

optimization by [127, 146, 170, 89] and more recently [14], one would hope to achieve a logarithmic dependence or no dependence at all on $\frac{R(\mathbf{0})}{\varepsilon}$. In this chapter we achieve this goal by providing an improved analysis showing that the greedy algorithm of [144] in fact returns a matrix of rank of $r = O(r^* \cdot \kappa_{r+r^*} \log \frac{R(\mathbf{0})}{\varepsilon})$. We also provide a new local search algorithm together with an analysis guaranteeing a rank of $r = O(r^* \cdot \kappa_{r+r^*}^2)$. Apart from significantly improving upon previous work on rank-restricted convex optimization, these results directly generalize a lot of work in sparse convex optimization, e.g. [127, 146, 89]. Our algorithms and theorem statements can be found in Section 8.2.

Runtime improvements Even though the rank bound guaranteed by our theoretical analyses is adequate, the algorithm runtimes leave much to be desired. In particular, both the greedy algorithm of [144] and our local search algorithm have to solve an optimization problem in each iteration in order to find the best possible linear combination of features added so far. Even for the case that $R(A) = \frac{1}{2} \sum_{(i,j) \in \Omega} (M - A)_{ij}^2$, this requires solving a least squares problem on $|\Omega|$ examples and r^2 variables. For practical implementations of these algorithms, we circumvent this issue by solving a related optimization problem that is usually much smaller. This instead requires solving n least squares problems with *total* number of examples $|\Omega|$, each on r variables. This not only reduces the size of the problem by a factor of r , but also allows for a straightforward distributed implementation. Interestingly, our theoretical analyses still hold for these variants. We propose an additional heuristic that reduces the runtime even more drastically, which is to only run a few (less than 10) iterations of the algorithm used for solving the inner optimization problem. Experimental results show that this modification not only does not significantly worsen results, but for machine learning applications also acts as a regularization method that can dramatically improve generalization. These matters, as well as additional improvements for making the local search algorithm more practical, are addressed in Section 8.2.3.

Roadmap In Section 8.2, we provide the descriptions and theoretical results for the algorithms used, along with several modifications to boost performance. In Section 8.3, we evaluate the proposed greedy and local search algorithms on optimization problems like robust PCA. Then, in Section 8.4 we evaluate their generalization performance in machine learning problems like matrix completion.

8.2 Algorithms & Theoretical Guarantees

In Sections 8.2.1 and 8.2.2 we state and provide theoretical performance guarantees for the basic greedy and local search algorithms respectively. Then in Section 8.2.3 we state the algorithmic adjustments that we propose in order to make the algorithms efficient in terms of runtime and generalization performance. A discussion regarding the tightness of the theoretical analysis is deferred to Appendix 9.6.

When the dimension is clear from context, we will denote the all-ones vector by $\mathbf{1}$, and the vector that is 0 everywhere and 1 at position i by $\mathbf{1}_i$. Given a matrix A , we denote by $\text{im}(A)$ its column span. One notion that we will find useful is that of *singular value thresholding*. More specifically, given a rank- k matrix $A \in \mathbb{R}^{m \times n}$ with SVD $\sum_{i=1}^k \sigma_i u^i v^{i\top}$ such that $\sigma_1 \geq \dots \geq \sigma_k$, as well as an integer parameter $r \geq 1$, we define $H_r(A) = \sum_{i=1}^r \sigma_i u^i v^{i\top}$ to be the operator that truncates to the r highest singular values of A .

8.2.1 Greedy

Algorithm 17 (Greedy) was first introduced in [144] as the GECO algorithm. It works by iteratively adding a rank-1 matrix to the current solution. This matrix is chosen as the rank-1 matrix that best approximates the gradient, i.e. the pair of singular vectors corresponding to the maximum singular value of the gradient. In each iteration, an additional procedure is run to optimize the combination of previously chosen singular vectors.

In [144] guarantee on the rank of the solution returned by the algorithm is $r^* \kappa_{r+r^*} \frac{R(\mathbf{0})}{\varepsilon}$. The main bottleneck in order to improve on the $\frac{R(\mathbf{0})}{\varepsilon}$ factor is the fact that the analysis is done in terms of the squared nuclear norm of the optimal solution. As the worst-case discrepancy between the squared nuclear norm and the rank is $R(\mathbf{0})/\varepsilon$, their bounds inherit this factor. Our analysis works directly with the rank, in the spirit of sparse optimization results (e.g. [144, 89, 14]). A challenge compared to these works is the need for a suitable notion of “intersection” between two sets of vectors. The main technical contribution of this work is to show that the orthogonal projection of one set of vectors into the span of the other is such a notion, and, based on this, to define a decomposition of the optimal solution that is used in the analysis.

Algorithm 17 Greedy

```

1: procedure GREEDY( $r \in \mathbb{N}$  : target rank)
2:   function to be minimized  $R : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ 
3:    $U \in \mathbb{R}^{m \times 0}$  ▷ Initially rank is zero
4:    $V \in \mathbb{R}^{n \times 0}$ 
5:   for  $t = 0 \dots r - 1$  do
6:      $\sigma uv^\top \leftarrow H_1(\nabla R(UV^\top))$  ▷ Max singular value  $\sigma$  and corresp. singular vectors  $u, v$ 
7:      $U \leftarrow \begin{pmatrix} U & u \end{pmatrix}$  ▷ Append new vectors as columns
8:      $V \leftarrow \begin{pmatrix} V & v \end{pmatrix}$ 
9:      $U, V \leftarrow \text{OPTIMIZE}(U, V)$ 
10:  return  $UV^\top$ 
11: procedure OPTIMIZE( $U \in \mathbb{R}^{m \times r}, V \in \mathbb{R}^{n \times r}$ )
12:   $X \leftarrow \underset{X \in \mathbb{R}^{r \times r}}{\text{argmin}} R(UXV^\top)$ 
13:  return  $UX, V$ 

```

Theorem 8.2.1 (Algorithm 17 (greedy) analysis). *Let A^* be any fixed optimal solution of (8.1) for some function R and rank bound r^* , and let $\varepsilon > 0$ be an error parameter. For any integer $r \geq 2r^* \cdot \kappa_{r+r^*} \log \frac{R(\mathbf{0}) - R(A^*)}{\varepsilon}$, if we let $A = \text{GREEDY}(r)$ be the solution returned by Algorithm 17, then $R(A) \leq R(A^*) + \varepsilon$. The number of iterations is r .*

The proof of Theorem 8.2.1 can be found in Appendix 9.6.

8.2.2 Local Search

One drawback of Algorithm 1 is that it increases the rank in each iteration. Algorithm 2 is a modification of Algorithm 1, in which the rank is truncated in each iteration. The advantage of Algorithm 2 compared to Algorithm 1 is that it is able to make progress without increasing the rank of A , while Algorithm 1 necessarily increases the rank in each iteration. More specifically, because of the greedy nature of Algorithm 1, some rank-1 components that have been added to A might become obsolete or have reduced benefit after a number of iterations. Algorithm 2 is able to identify such candidates and remove them, thus allowing it to continue making progress.

Algorithm 18 Local Search

```

1: procedure LOCAL_SEARCH( $r \in \mathbb{N}$  : target rank)
2:   function to be minimized  $R : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ 
3:    $U \leftarrow \mathbf{0}_{m \times r}$  ▷ Initialize with all-zero solution
4:    $V \leftarrow \mathbf{0}_{n \times r}$ 
5:   for  $t = 0 \dots L - 1$  do ▷ Run for  $L$  iterations
6:      $\sigma v v^\top \leftarrow H_1(\nabla R(UV^\top))$  ▷ Max singular value  $\sigma$  and corresp. singular vectors  $u, v$ 
7:      $U, V \leftarrow \text{TRUNCATE}(U, V)$  ▷ Reduce rank of  $UV^\top$  by one
8:      $U \leftarrow \begin{pmatrix} U & u \end{pmatrix}$  ▷ Append new vectors as columns
9:      $V \leftarrow \begin{pmatrix} V & v \end{pmatrix}$ 
10:     $U, V \leftarrow \text{OPTIMIZE}(U, V)$ 
11:  return  $UV^\top$ 
12: procedure TRUNCATE( $U \in \mathbb{R}^{m \times r}, V \in \mathbb{R}^{n \times r}$ )
13:   $U \Sigma V^\top \leftarrow \text{SVD}(H_{r-1}(UV^\top))$  ▷ Keep all but minimum singular value
14:  return  $U \Sigma, V$ 

```

Theorem 8.2.2 (Algorithm 18 (local search) analysis). *Let A^* be any fixed optimal solution of (8.1) for some function R and rank bound r^* , and let $\varepsilon > 0$ be an error parameter. For any integer $r \geq r^* \cdot (1 + 8\kappa_{r+r^*}^2)$, if we let $A = \text{LOCAL_SEARCH}(r)$ be the solution returned by Algorithm 18, then $R(A) \leq R(A^*) + \varepsilon$. The number of iterations is $O\left(r^* \kappa_{r+r^*} \log \frac{R(\mathbf{0}) - R(A^*)}{\varepsilon}\right)$.*

The proof of Theorem 8.2.2 can be found in Appendix 9.6.

8.2.3 Algorithmic adjustments

Inner optimization problem The inner optimization problem that is used in both greedy and local search is:

$$\min_{X \in \mathbb{R}^{r \times r}} R(UXV^\top). \quad (8.2)$$

It essentially finds the choice of matrices U' and V' , with columns in the column span of U and V respectively, that minimizes $R(U'V'^\top)$. We, however, consider the following problem instead:

$$\min_{V \in \mathbb{R}^{n \times r}} R(UV^\top). \quad (8.3)$$

Note that the solution recovered from (8.3) will never have worse objective value than the one recovered from (8.2), and that nothing in the analysis of the algorithms breaks. Importantly, (8.3) can usually be solved much more efficiently than (8.2). As an example, consider the following objective that appears in matrix completion: $R(A) = \frac{1}{2} \sum_{(i,j) \in \Omega} (M - A)_{ij}^2$ for some $\Omega \subseteq [m] \times [n]$. If

we let $\Pi_\Omega(\cdot)$ be an operator that zeroes out all positions in the matrix that are not in Ω , we have $\nabla R(A) = -\Pi_\Omega(M - A)$. The optimality condition of (8.2) now is $U^\top \Pi_\Omega(M - UXV^\top)V = \mathbf{0}$ and that of (8.3) is $U^\top \Pi_\Omega(M - UV^\top) = \mathbf{0}$. The former corresponds to a least squares linear regression problem with $|\Omega|$ examples and r^2 variables, while the latter can be decomposed into n independent systems $U^\top \left(\sum_{i:(i,j) \in \Omega} \mathbf{1}_i \mathbf{1}_i^\top \right) UV^j = U^\top \Pi_\Omega(M \mathbf{1}_j)$, where the variable is V^j which is the j -th column of V . The j -th of these systems corresponds to a least squares linear regression problem with

$|\{i : (i, j) \in \Omega\}|$ examples and r variables. Note that the total number of examples in all systems is $\sum_{j \in [n]} |\{i : (i, j) \in \Omega\}| = |\Omega|$. The choice of V here as the variable to be optimized is arbitrary.

In particular, as can be seen in Algorithm 19, in practice we alternate between optimizing U and V in each iteration. It is worthy of mention that the OPTIMIZE_FAST procedure is basically the same as one step of the popular alternating minimization procedure for solving low rank problems. As a matter of fact, when our proposed algorithms are viewed from this lens, they can be seen as alternating minimization interleaved with rank-1 insertion and/or removal steps.

Algorithm 19 Fast inner Optimization

```

1: procedure OPTIMIZE_FAST( $U \in \mathbb{R}^{m \times r}, V \in \mathbb{R}^{n \times r}, t \in \mathbb{N}$  : iteration index of algorithm)
2:   if  $t \bmod 2 = 0$  then
3:      $X \leftarrow \underset{X \in \mathbb{R}^{m \times r}}{\operatorname{argmin}} R(XV^\top)$ 
4:     return  $X, V$ 
5:   else
6:      $X \leftarrow \underset{X \in \mathbb{R}^{n \times r}}{\operatorname{argmin}} R(UX^\top)$ 
7:     return  $U, X$ 

```

Singular value decomposition As modern methods for computing the top entries of a singular value decomposition scale very well even for large sparse matrices ([120, 159, 163]), the “insertion” step of greedy and local search, in which the top entry of the SVD of the gradient is determined, is quite fast in practice. However, these methods are not suited for computing the *smallest* singular values and corresponding singular vectors, a step required for the local search algorithm that we propose. Therefore, in our practical implementations we opt to perform the alternative step of directly removing one pair of vectors from the representation UV^\top . A simple approach is to go over all r possible removals and pick the one that increases the objective by the least amount. A variation of this approach has been used by [144]. However, a much faster approach is to just pick the pair of vectors $U\mathbf{1}_i, V\mathbf{1}_i$ that minimizes $\|U\mathbf{1}_i\|_2 \|V\mathbf{1}_i\|_2$. This is the approach that we use, as can be seen in Algorithm 20.

Algorithm 20 Fast rank reduction

```

1: procedure TRUNCATE_FAST( $U \in \mathbb{R}^{m \times r}, V \in \mathbb{R}^{n \times r}$ )
2:    $i \leftarrow \underset{i \in [r]}{\operatorname{argmin}} \|U\mathbf{1}_i\|_2 \|V\mathbf{1}_i\|_2$ 
3:   return  $(U_{[m],[1,i-1]} \quad U_{[m],[i+1,r]}), (V_{[n],[1,i-1]} \quad V_{[n],[i+1,r]})$  ▷ Remove column  $i$ 

```

After the previous discussion, we are ready to state the fast versions of Algorithm 17 and Algorithm 18 that we use for our experiments. These are Algorithm 8.2.3 and Algorithm 21. Notice that we initialize Algorithm 21 with the solution of Algorithm 8.2.3 and we run it until the value of $R(\cdot)$ stops decreasing rather than for a fixed number of iterations.

Algorithm 8.2.3 (Fast Greedy). *The Fast Greedy algorithm is defined identically as Algorithm 17, with the only difference that it uses the OPTIMIZE_FAST routine as opposed to the OPTIMIZE routine.*

Algorithm 21 Fast Local Search

```
1: procedure FAST_LOCAL_SEARCH( $r \in \mathbb{N}$  : target rank)
2:   function to be minimized  $R : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ 
3:    $U, V \leftarrow$  solution returned by FAST_GREEDY( $r$ )
4:   do
5:      $U_{\text{prev}}, V_{\text{prev}} \leftarrow U, V$ 
6:      $\sigma uv^\top \leftarrow H_1(\nabla R(UV^\top))$   $\triangleright$  Max singular value  $\sigma$  and corresp. singular vectors  $u, v$ 
7:      $U, V \leftarrow$  TRUNCATE_FAST( $U, V$ )  $\triangleright$  Reduce rank of  $UV^\top$  by one
8:      $U \leftarrow \begin{pmatrix} U & u \end{pmatrix}$   $\triangleright$  Append new vectors as columns
9:      $V \leftarrow \begin{pmatrix} V & v \end{pmatrix}$ 
10:     $U, V \leftarrow$  OPTIMIZE_FAST( $U, V, t$ )
11:  while  $R(UV^\top) < R(U_{\text{prev}}V_{\text{prev}}^\top)$ 
12:  return  $U_{\text{prev}}V_{\text{prev}}^\top$ 
```

8.3 Optimization Applications

An immediate application of the above algorithms is in the problem of *low rank matrix recovery*. Given any convex distance measure between matrices $d : \mathbb{R}_{m \times n} \times \mathbb{R}_{m \times n} \rightarrow \mathbb{R}_{\geq 0}$, the goal is to find a low-rank matrix A that matches a target matrix M as well as possible in terms of d : $\min_{\text{rank}(A) \leq r^*} d(M, A)$

This problem captures a lot of different applications, some of which we go over in the following sections.

8.3.1 Low-rank approximation on observed set

A particular case of interest is when $d(M, A)$ is the Frobenius norm of $M - A$, but only applied to entries belonging to some set Ω . In other words, $d(M, A) = \frac{1}{2} \|\Pi_\Omega(M - A)\|_F^2$. We have compared our Fast Greedy and Fast Local Search algorithms with the SoftImpute algorithm of [121] as implemented by [141], on the same experiments as in [121]. We have solved the inner optimization problem required by our algorithms by the LSQR algorithm [134]. More specifically, $M = UV^\top + \eta \in \mathbb{R}^{100 \times 100}$, where η is some noise vector. We let every entry of U, V, η be i.i.d. normal with mean 0 and the entries of Ω are chosen i.i.d. uniformly at random over the set $[100] \times [100]$. The experiments have three parameters: The true rank r^* (of UV^\top), the percentage of observed entries $p = |\Omega|/10^4$, and the signal-to-noise ratio SNR. We measure the normalized MSE, i.e. $\|\Pi_\Omega(M - A)\|_F^2 / \|\Pi_\Omega(M)\|_F^2$. The results can be seen in Figure 8-1, where it is illustrated that Fast Local Search sometimes returns significantly more accurate and lower-rank solutions than Fast Greedy, and Fast Greedy generally returns significantly more accurate and lower-rank solutions than SoftImpute.

8.3.2 Robust principal component analysis (RPCA)

The robust PCA paradigm asks one to decompose a given matrix M as $L + S$, where L is a low-rank matrix and S is a sparse matrix. This is useful for applications with outliers where directly computing the principal components of M is significantly affected by them. For a comprehensive survey on Robust PCA survey one can look at [24]. The following optimization problem encodes the above-stated requirements:

$$\min_{\text{rank}(L) \leq r^*} \|M - L\|_0 \quad (8.4)$$

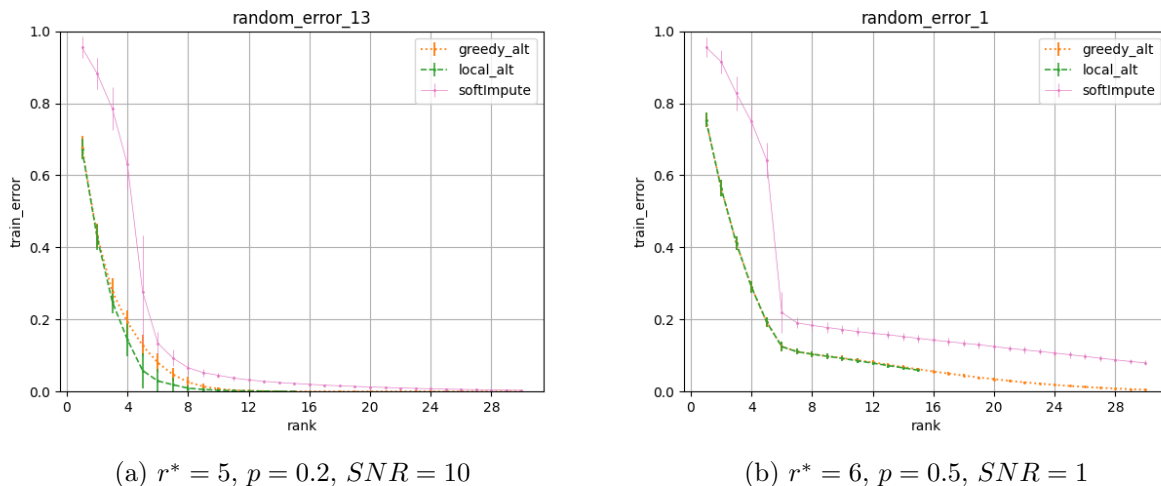


Figure 8-1: Objective value error vs rank in the problem of Section 8.3.1.

where $\|X\|_0$ is the sparsity (i.e. number of non-zeros) of X . As neither the rank constraint or the ℓ_0 function are convex, [34] replaced them by their usual convex relaxations, i.e. the nuclear norm $\|\cdot\|_*$ and ℓ_1 norm respectively. However, we opt to only relax the ℓ_0 function but not the rank constraint, leaving us with the problem:

$$\min_{\text{rank}(L) \leq r^*} \|M - L\|_1 \quad (8.5)$$

In order to make the objective differentiable and thus more well-behaved, we further replace the ℓ_1 norm by the Huber loss $H_\delta(x) = \begin{cases} x^2/2 & \text{if } |x| \leq \delta \\ \delta|x| - \delta^2/2 & \text{otherwise} \end{cases}$, thus getting: $\min_{\text{rank}(L) \leq r^*} \sum_{ij} H_\delta(M - L)_{ij}$. This is a problem on which we can directly apply our algorithms. We solve the inner optimization problem by applying 10 L-BFGS iterations.

In Figure 8-2 one can see an example of foreground-background separation from video using robust PCA. The video is from the BMC 2012 dataset [164]. In this problem, the low-rank part corresponds to the background and the sparse part to the foreground. We compare three algorithms: Our Fast Greedy algorithm, standard PCA with 1 component (the choice of 1 was picked to get the best outcome), and the standard Principal Component Pursuit (PCP) algorithm ([34]), as implemented in [111], where we tuned the regularization parameter λ to achieve the best result. We find that Fast Greedy has the best performance out of the three algorithms in this sample task.

8.4 Machine Learning Applications

8.4.1 Regularization techniques

In the previous section we showed that our proposed algorithms bring down different optimization objectives aggressively. However, in applications where the goal is to obtain a low generalization error, regularization is needed. We considered two different kinds of regularization. The first method is to run the inner optimization algorithm for less iterations, usually 2-3. Usually this is straightforward since an iterative method is used. For example, in the case $R(A) = \frac{1}{2} \|\Pi_\Omega(M - A)\|_F^2$ the inner optimization is a least squares linear regression problem that we solve using the LSQR algorithm. The second one is to add an ℓ_2 regularizer to the objective function. However, this option did not



Figure 8-2: Foreground-background separation from video. From left to right: Fast Greedy with rank=3 and Huber loss with $\delta = 20$. Standard PCA with rank=1. Principal Component Pursuit (PCP) with $\lambda = 0.002$. Both PCA and PCP have visible "shadows" in the foreground that appear as "smudges" in the background. These are less obvious in a still frame but more apparent in a video.

provide a substantial performance boost in our experiments, and so we have not implemented it.

8.4.2 Matrix Completion with Random Noise

In this section we evaluate our algorithms on the task of recovering a low rank matrix UV^\top after observing $\Pi_\Omega(UV^\top + \eta)$, i.e. a fraction of its entries with added noise. As in Section 8.3.1, we use the setting of [121] and compare with the SoftImpute method. The evaluation metric is the normalized MSE, defined as $(\sum_{(i,j) \notin \Omega} (UV^\top - A)_{ij}^2) / (\sum_{(i,j) \notin \Omega} (UV^\top)_{ij}^2)$, where A is the predicted matrix

and UV^\top the true low rank matrix. A few example plots can be seen in Figure 8-3 and a table of results in Table 8.1. We have implemented the Fast Greedy and Fast Local Search algorithms with 3 inner optimization iterations. In the first few iterations there is a spike in the relative MSE of the algorithms that use the OPTIMIZE_FAST routine. We attribute this to the aggressive alternating minimization steps of this procedure and conjecture that adding a regularization term to the objective might smoothen the spike. However, the Fast Local Search algorithm still gives the best overall performance in terms of how well it approximates the true low rank matrix UV^\top , and in particular with a very small rank—practically the same as the true underlying rank.

8.4.3 Recommender Systems

In this section we compare our algorithms on the task of movie recommendation on the Movielens datasets [80]. In order to evaluate the algorithms, we perform random 80%-20% train-test splits that are the same for all algorithms and measure the mean RMSE in the test set. If we let $\Omega \subseteq [m] \times [n]$ be the set of user-movie pairs in the training set, we assume that the true user-movie matrix is low rank, and thus pose (8.1) with $R(A) = \frac{1}{2} \|\Pi_\Omega(M - A)\|_F^2$. We make the following slight modification

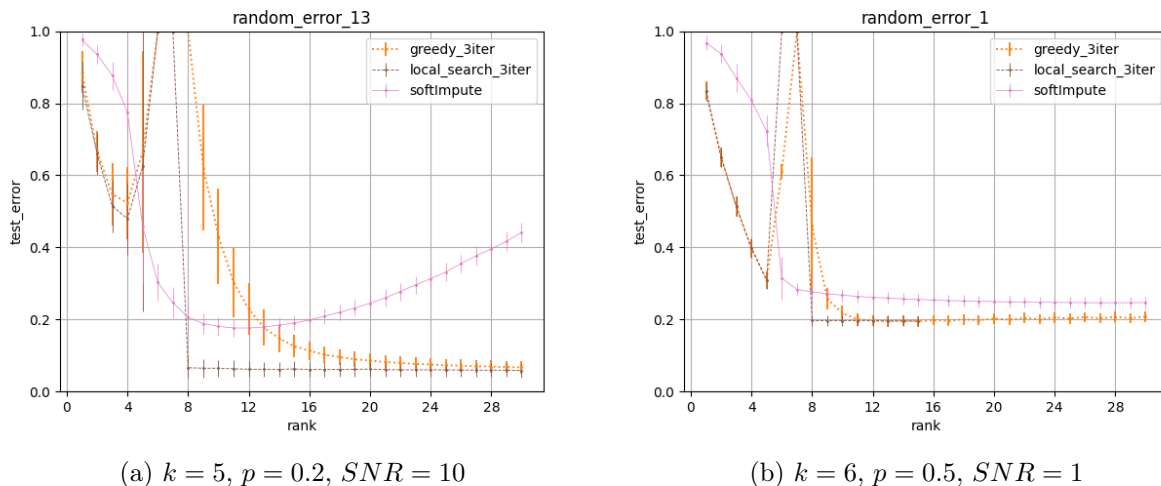


Figure 8-3: Test error vs rank in the matrix completion problem of Section 8.4.2. Bands of ± 1 standard error are shown. Note that SoftImpute starts to overfit for ranks larger than 12 in (a). The “jumps” at around rank 5-7 happen because of overshooting (taking too large a step) during the insertion the rank-1 component in both Fast Greedy and Fast Local Search. More specifically, these implementations only apply 3 iterations of the inner optimization step, which in some cases are too few to amend the overshooting. However, after a few more iterations of the algorithm the overshooting issue is resolved (i.e. the algorithm has had enough iterations to scale down the rank-1 component that caused the overshooting).

Algorithm	random_error_13	random_error_1	random_error_2
SoftImpute ([121])	0.1759/10	0.2465/28	0.2284/30
Fast Greedy (Algorithm 8.2.3)	0.0673/30	0.1948/13	0.1826/21
Fast Local Search (Algorithm 21)	0.0613/14	0.1952/15	0.1811/15

Table 8.1: Lowest test error for any rank in the matrix completion problem of Section 8.4.2, and associated rank returned by each algorithm. In the form error/rank.

in order to take into account the range of the ratings $[1, 5]$: We clip the entries of A between 1 and 5 when computing $\nabla R(A)$ in Algorithm 8.2.3 and Algorithm 21. In other words, instead of $\Pi_{\Omega}(A - M)$ we compute the gradient as $\Pi_{\Omega}(\text{clip}(A, 1, 5) - M)$. This is similar to replacing our objective by a Huber loss, with the difference that we only do so in the steps that we mentioned and not the inner optimization step, mainly for runtime efficiency reasons.

The results can be seen in Table 8.2. We do not compare with Fast Local Search, as we found that it only provides an advantage for small ranks (< 30), and otherwise matches Fast Greedy. For the inner optimization steps we have used the LSQR algorithm with 2 iterations in the 100K and 1M datasets, and with 3 iterations in the 10M dataset. Note that even though the SVD algorithm by [101] as implemented by [83] (with no user/movie bias terms) is a highly tuned algorithm for recommender systems that was one of the top solutions in the famous Netflix prize, it has comparable performance to our general-purpose Algorithm 8.2.3.

Finally, Table 8.3 demonstrates the speedup achieved by our algorithms over the basic greedy implementation. It should be noted that the speedup compared to the basic greedy of [144] (Algorithm 17) is larger as rank increases, since the fast algorithms scale linearly with rank, but the

Algorithm	MovieLens 100K	MovieLens 1M	MovieLens 10M
NMF ([104])	0.9659	0.9166	0.8960
SoftImpute	1.0106	0.9599	0.957
Alternating Minimization	0.9355	0.8732	0.8410
SVD ([101])	0.9533	0.8743	0.8315
Fast Greedy (Algorithm 8.2.3)	0.9451	0.8714	0.8330

Table 8.2: Mean RMSE and standard error among 5 random splits for 100K and 1M with standard errors < 0.01 , and 3 random splits for 10M with standard errors < 0.001 . The rank of the prediction is set to 100 except for NMF where it is 15 and Fast Greedy in the 10M dataset where it is chosen to be 35 by cross-validation. Alternating Minimization is a well known algorithm (e.g. [157]) that alternatively minimizes the left and right subspace, and also uses Frobenius norm regularization. For SoftImpute and Alternating Minimization we have found the best choice of parameters by performing a grid search over the rank and the multiplier of the regularization term. We have found the best choice of parameters by performing a grid search over the rank and the multiplier of the regularization term. We ran 20 iterations of Alternating Minimization in each case.

basic greedy scales quadratically.

Algorithm	Figure 8-3 (a)	MovieLens 100K	MovieLens 1M
SoftImpute	10.6	9.4	40.6
Alternating Minimization	18.9	252.0	1141.4
Greedy ([144])	18.8	418.4	4087.3
Fast Greedy	10.2	43.4	244.2
Fast Local Search	10.8	46.1	263.0

Table 8.3: Runtimes (in seconds) of different algorithms for fitting a rank=30 solution in various experiments. Code written in python and tested on an Intel Skylake CPU with 16 vCPUs.

It is important to note that our goal here is not to be competitive with the best known algorithms for matrix completion, but rather to propose a general yet practically applicable method for rank-constrained convex optimization. For a recent survey on the best performing algorithms in the MovieLens datasets see [139]. It should be noted that a lot of these algorithms have significant performance boost compared to our methods because they use additional features (meta information about each user, movie, timestamp of a rating, etc.) or stronger models (user/movie biases, "implicit" ratings). A runtime comparison with these recent approaches is an interesting avenue for future work. As a rule of thumb, however, Fast Greedy has roughly the same runtime as SVD ([101]) in each iteration, i.e. $O(|\Omega|r)$, where Ω is the set of observable elements and r is the rank. As some better performing approaches have been reported to be much slower than SVD (e.g. SVD++ is reported to be 50-100x slower than SVD in the MovieLens 100K and 1M datasets ([83]), this might also suggest a runtime advantage of our approach compared to some better performing methods.

8.5 Conclusions

We presented simple algorithms with strong theoretical guarantees for optimizing a convex function under a rank constraint. Although the basic versions of these algorithms have appeared before, through a series of careful runtime, optimization, and generalization performance improvements

that we introduced, we have managed to reshape the performance of these algorithms in all fronts. Via our experimental validation on a host of practical problems such as low-rank matrix recovery with missing data, robust principal component analysis, and recommender systems, we have shown that the performance in terms of the solution quality matches or exceeds other widely used and even specialized solutions, thus making the argument that our Fast Greedy and Fast Local Search routines can be regarded as strong and practical general purpose tools for rank-constrained convex optimization. Interesting directions for further research include the exploration of different kinds of regularization and tuning for machine learning applications, as well as a competitive implementation and extensive runtime comparison of our algorithms.

Bibliography

- [1] Deeksha Adil, Rasmus Kyng, Richard Peng, and Sushant Sachdeva. “Iterative Refinement for ℓ_p -norm Regression”. In: *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM. 2019, pp. 1405–1424.
- [2] Deeksha Adil and Sushant Sachdeva. “Faster p -norm minimizing flows, via smoothed q -norm problems”. In: *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM. 2020, pp. 892–910.
- [3] Pankaj K. Agarwal, Micha Sharir, and Sivan Toledo. “Applications of Parametric Searching in Geometric Optimization”. In: *J. Algorithms* 17.3 (1994), pp. 292–318. DOI: 10.1006/jagm.1994.1038. URL: <https://doi.org/10.1006/jagm.1994.1038>.
- [4] Ravindra K Ahuja, Thomas L Magnanti, and James B Orlin. *Network Flows*. Tech. rep. Alfred P. Sloan School of Management, Cambridge, MA, 1988.
- [5] Zeyuan Allen-Zhu, Elad Hazan, Wei Hu, and Yuanzhi Li. “Linear convergence of a frank-wolfe type algorithm over trace-norm balls”. In: *Advances in Neural Information Processing Systems* 30 (2017).
- [6] Jason Altschuler, Aditya Bhaskara, Gang Fu, Vahab Mirrokni, Afshin Rostamizadeh, and Morteza Zadimoghaddam. “Greedy Column Subset Selection: New Bounds and Distributed Algorithms”. In: *International Conference on Machine Learning*. 2016, pp. 2539–2548.
- [7] Joel Andersson and Jan-Olov Strömberg. “On the theorem of uniform recovery of random sampling matrices”. In: *IEEE Transactions on Information Theory* 60.3 (2014), pp. 1700–1710.
- [8] Chetan Arora, Subhashis Banerjee, Prem Kalra, and SN Maheshwari. “Generic cuts: An efficient algorithm for optimal inference in higher order MRF-MAP”. In: *European Conference on Computer Vision*. Springer. 2012, pp. 17–30.
- [9] Kyriakos Axiotis, Adam Karczmarz, Anish Mukherjee, Piotr Sankowski, and Adrian Vladu. “Decomposable submodular function minimization via maximum flow”. In: *International Conference on Machine Learning*. PMLR. 2021, pp. 446–456.
- [10] Kyriakos Axiotis, Aleksander Mądry, and Adrian Vladu. “Circulation control for faster minimum cost flow in unit-capacity graphs”. In: *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2020, pp. 93–104.
- [11] Kyriakos Axiotis, Aleksander Mądry, and Adrian Vladu. “Faster sparse minimum cost flow by electrical flow localization”. In: *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2022, pp. 528–539.
- [12] Kyriakos Axiotis and Maxim Sviridenko. “Iterative Hard Thresholding with Adaptive Regularization: Sparser Solutions Without Sacrificing Runtime”. In: *International Conference on Machine Learning* (2022).

- [13] Kyriakos Axiotis and Maxim Sviridenko. “Local Search Algorithms for Rank-Constrained Convex Optimization”. In: *International Conference on Learning Representations*. 2021.
- [14] Kyriakos Axiotis and Maxim Sviridenko. “Sparse Convex Optimization via Adaptively Regularized Hard Thresholding”. In: *Journal of Machine Learning Research* 22 (2021), pp. 1–47.
- [15] Francis Bach et al. “Learning with submodular functions: A convex optimization perspective”. In: *Foundations and Trends® in Machine Learning* 6.2-3 (2013), pp. 145–373.
- [16] Sohail Bahmani, Bhiksha Raj, and Petros T Boufounos. “Greedy sparsity-constrained optimization”. In: *Journal of Machine Learning Research* 14.Mar (2013), pp. 807–841.
- [17] Marshall Bern, John R Gilbert, Bruce Hendrickson, Nhat Nguyen, and Sivan Toledo. “Support-graph preconditioners”. In: *SIAM Journal on Matrix Analysis and Applications* 27.4 (2006), pp. 930–951.
- [18] Aaron Bernstein, Jan van den Brand, Maximilian Probst Gutenberg, Danupon Nanongkai, Thatchaphol Saranurak, Aaron Sidford, and He Sun. “Fully-dynamic graph sparsifiers against an adaptive adversary”. In: *arXiv preprint arXiv:2004.08432* (2020).
- [19] Thomas Blumensath and Mike E Davies. “Iterative hard thresholding for compressed sensing”. In: *Applied and computational harmonic analysis* 27.3 (2009), pp. 265–274.
- [20] Thomas Blumensath and Mike E Davies. *On the difference between orthogonal matching pursuit and orthogonal least squares*. Tech. rep. 2007.
- [21] Holger Boche, Robert Calderbank, Gitta Kutyniok, and Jan Vybiral. “A survey of compressed sensing”. In: *Compressed sensing and its applications*. Springer, 2015, pp. 1–39.
- [22] Glencora Borradaile and Philip N. Klein. “An $O(n \log n)$ algorithm for maximum st -flow in a directed planar graph”. In: *J. ACM* 56.2 (2009), 9:1–9:30. DOI: 10.1145/1502793.1502798. URL: <https://doi.org/10.1145/1502793.1502798>.
- [23] Jonathan Borwein and Adrian S Lewis. *Convex analysis and nonlinear optimization: theory and examples*. Springer Science & Business Media, 2010.
- [24] Thierry Bouwmans, Sajid Javed, Hongyang Zhang, Zhouchen Lin, and Ricardo Otazo. “On the applications of robust PCA in image and video processing”. In: *Proceedings of the IEEE* 106.8 (2018), pp. 1427–1457.
- [25] Jan van den Brand, Yin Tat Lee, Yang P Liu, Thatchaphol Saranurak, Aaron Sidford, Zhao Song, and Di Wang. “Minimum cost flows, mdps, and ℓ_1 -regression in nearly linear time for dense instances”. In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 2021, pp. 859–869.
- [26] Jan van den Brand, Yin-Tat Lee, Danupon Nanongkai, Richard Peng, Thatchaphol Saranurak, Aaron Sidford, Zhao Song, and Di Wang. “Bipartite matching in nearly-linear time on moderately dense graphs”. In: *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2020, pp. 919–930.
- [27] Sébastien Bubeck, Michael B Cohen, Yin Tat Lee, James R Lee, and Aleksander Mądry. “K-server via multiscale entropic regularization”. In: *Proceedings of the 50th annual ACM SIGACT symposium on theory of computing*. 2018, pp. 3–16.
- [28] Sébastien Bubeck, Michael B Cohen, Yin Tat Lee, and Yuanzhi Li. “An homotopy method for ℓ_p regression provably beyond self-concordance and in input-sparsity time”. In: *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. 2018, pp. 1130–1137.

- [29] T Tony Cai, Lie Wang, and Guangwu Xu. “New bounds for restricted isometry constants”. In: *IEEE Transactions on Information Theory* 56.9 (2010), pp. 4388–4394.
- [30] T Tony Cai, Lie Wang, and Guangwu Xu. “Shifting inequality and recovery of sparse signals”. In: *IEEE Transactions on Signal processing* 58.3 (2009), pp. 1300–1308.
- [31] Emmanuel J Candes. “The restricted isometry property and its implications for compressed sensing”. In: *Comptes rendus mathématique* 346.9-10 (2008), pp. 589–592.
- [32] Emmanuel J Candes, Justin K Romberg, and Terence Tao. “Stable signal recovery from incomplete and inaccurate measurements”. In: *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences* 59.8 (2006), pp. 1207–1223.
- [33] Emmanuel J Candes and Terence Tao. “Decoding by linear programming”. In: *IEEE transactions on information theory* 51.12 (2005), pp. 4203–4215.
- [34] Emmanuel J Candès, Xiaodong Li, Yi Ma, and John Wright. “Robust principal component analysis?” In: *Journal of the ACM (JACM)* 58.3 (2011), pp. 1–37.
- [35] Deeparnab Chakrabarty, Yin Tat Lee, Aaron Sidford, and Sam Chiu-wai Wong. “Subquadratic submodular function minimization”. In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. 2017, pp. 1220–1231.
- [36] Antonin Chambolle and Jérôme Darbon. “On total variation minimization and surface evolution using parametric maximum flows”. In: *International journal of computer vision* 84.3 (2009), p. 288.
- [37] Li Chen, Rasmus Kyng, Yang P Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. “Maximum flow and minimum-cost flow in almost-linear time”. In: *arXiv preprint arXiv:2203.00671* (2022).
- [38] Lin Chen, Moran Feldman, and Amin Karbasi. “Weakly Submodular Maximization Beyond Cardinality Constraints: Does Randomization Help Greedy?” In: *International Conference on Machine Learning*. 2018, pp. 804–813.
- [39] Paul Christiano, Jonathan A Kelner, Aleksander Mądry, Daniel A Spielman, and Shang-Hua Teng. “Electrical flows, Laplacian systems, and faster approximation of maximum flow in undirected graphs”. In: *Proceedings of the 43rd Annual ACM Symposium on Theory of computing*. 2011, pp. 273–282.
- [40] Fan Chung and Linyuan Lu. “Concentration inequalities and martingale inequalities: a survey”. In: *Internet Mathematics* 3.1 (2006), pp. 79–127.
- [41] Michael B Cohen, Rasmus Kyng, Gary L Miller, Jakub W Pachocki, Richard Peng, Anup B Rao, and Shen Chen Xu. “Solving SDD linear systems in nearly $m \log^{1/2} n$ time”. In: *Proceedings of the 46th Annual ACM Symposium on Theory of computing*. 2014, pp. 343–352.
- [42] Michael B Cohen, Aleksander Mądry, Piotr Sankowski, and Adrian Vladu. “Negative-Weight Shortest Paths and Unit Capacity Minimum Cost Flow in $\tilde{O}(m^{10/7} \log W)$ Time”. In: *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM. 2017, pp. 752–771.
- [43] Daniel Dadush, László A Végh, and Giacomo Zambelli. “Geometric rescaling algorithms for submodular function minimization”. In: *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM. 2018, pp. 832–848.

- [44] Samuel I Daitch and Daniel A Spielman. “Faster approximate lossy generalized flow via interior point algorithms”. In: *Proceedings of the 40th Annual ACM Symposium on Theory of computing*. 2008, pp. 451–460.
- [45] Nikhil R Devanur, Shaddin Dughmi, Roy Schwartz, Ankit Sharma, and Mohit Singh. “On the approximation of submodular functions”. In: *arXiv preprint arXiv:1304.4948* (2013).
- [46] David L Donoho. “Compressed sensing”. In: *IEEE Transactions on information theory* 52.4 (2006), pp. 1289–1306.
- [47] David L Donoho and Michael Elad. “Optimally sparse representation in general (nonorthogonal) dictionaries via l_1 minimization”. In: *Proceedings of the National Academy of Sciences* 100.5 (2003), pp. 2197–2202.
- [48] David L Donoho, Arian Maleki, and Andrea Montanari. “Message-passing algorithms for compressed sensing”. In: *Proceedings of the National Academy of Sciences* 106.45 (2009), pp. 18914–18919.
- [49] Dheeru Dua and Casey Graff. *UCI Machine Learning Repository*. 2017. URL: <http://archive.ics.uci.edu/ml>.
- [50] John Duchi, Elad Hazan, and Yoram Singer. “Adaptive subgradient methods for online learning and stochastic optimization.” In: *Journal of machine learning research* 12.7 (2011).
- [51] David Durfee, Yu Gao, Gramoz Goranci, and Richard Peng. “Fully dynamic spectral vertex sparsifiers and applications”. In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. 2019, pp. 914–925.
- [52] David Durfee, Kevin A Lai, and Saurabh Sawlani. “ ℓ_1 Regression using Lewis Weights Preconditioning and Stochastic Gradient Descent”. In: *Conference on Learning Theory*. 2018, pp. 1626–1656.
- [53] Jack Edmonds. “Matroids and the greedy algorithm”. In: *Mathematical programming* 1.1 (1971), pp. 127–136.
- [54] Jack Edmonds. “Submodular functions, matroids, and certain polyhedra”. In: *Combinatorial Structures and Their Applications*. 1970, pp. 69–87.
- [55] Ethan Elenberg, Alexandros G Dimakis, Moran Feldman, and Amin Karbasi. “Streaming weak submodularity: Interpreting neural networks on the fly”. In: *Advances in Neural Information Processing Systems*. 2017, pp. 4044–4054.
- [56] Alina Ene and Huy Nguyen. “Random coordinate descent methods for minimizing decomposable submodular functions”. In: *International Conference on Machine Learning*. PMLR. 2015, pp. 787–795.
- [57] Alina Ene, Huy L Nguyen, and László A Végh. “Decomposable Submodular Function Minimization: Discrete and Continuous”. In: *NIPS*. 2017.
- [58] Jeff Erickson. “Maximum Flows and Parametric Shortest Paths in Planar Graphs”. In: *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010*. Ed. by Moses Charikar. SIAM, 2010, pp. 794–804. DOI: 10.1137/1.9781611973075.65.
- [59] Steve Fisk. “A note on weyl’s inequality”. In: (1996).
- [60] Alexander Fix, Thorsten Joachims, Sung Min Park, and Ramin Zabih. “Structured learning of sum-of-submodular higher order energy functions”. In: *Proceedings of the IEEE International Conference on Computer Vision*. 2013, pp. 3104–3111.

- [61] 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.
- [62] L. R. Ford and D. R. Fulkerson. *Flows in Networks*. Princeton University Press, 1962. ISBN: 9780691625393.
- [63] Dean Foster, Howard Karloff, and Justin Thaler. “Variable selection is hard”. In: *Conference on Learning Theory*. 2015, pp. 696–709.
- [64] Simon Foucart. “A note on guaranteed sparse recovery via ℓ_1 -minimization”. In: *Applied and Computational Harmonic Analysis* 29.1 (2010), pp. 97–103.
- [65] Simon Foucart. “Hard thresholding pursuit: an algorithm for compressive sensing”. In: *SIAM Journal on Numerical Analysis* 49.6 (2011), pp. 2543–2563.
- [66] Simon Foucart. “Sparse recovery algorithms: sufficient conditions in terms of restricted isometry constants”. In: *Approximation Theory XIII: San Antonio 2010*. Springer, 2012, pp. 65–77.
- [67] Simon Foucart and Ming-Jun Lai. “Sparsest solutions of underdetermined linear systems via ℓ_q -minimization for $0 < q \leq 1$ ”. In: *Applied and Computational Harmonic Analysis* 26.3 (2009), pp. 395–407.
- [68] Simon Foucart and Holger Rauhut. “A mathematical introduction to compressive sensing”. In: *Bull. Am. Math* 54 (2017), pp. 151–165.
- [69] Marguerite Frank, Philip Wolfe, et al. “An algorithm for quadratic programming”. In: *Naval research logistics quarterly* 3.1-2 (1956), pp. 95–110.
- [70] Satoru Fujishige. “Lexicographically optimal base of a polymatroid with respect to a weight vector”. In: *Mathematics of Operations Research* 5.2 (1980), pp. 186–196.
- [71] Harold N Gabow. “Scaling algorithms for network problems”. In: *24th Annual Symposium on Foundations of Computer Science (SFCS 1983)*. IEEE. 1983, pp. 248–258.
- [72] Giorgio Gallo, Michael D. Grigoriadis, and Robert Endre Tarjan. “A Fast Parametric Maximum Flow Algorithm and Applications”. In: *SIAM J. Comput.* 18.1 (1989), pp. 30–55. DOI: 10.1137/0218003. URL: <https://doi.org/10.1137/0218003>.
- [73] Yu Gao, Yang P Liu, and Richard Peng. “Fully dynamic electrical flows: Sparse maxflow faster than goldberg-rao”. In: *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2022, pp. 516–527.
- [74] Andrew V Goldberg, Sagi Hed, Haim Kaplan, and Robert E Tarjan. “Minimum-cost flows in unit-capacity networks”. In: *Theory of Computing Systems* 61.4 (2017), pp. 987–1010.
- [75] Andrew V Goldberg and Satish Rao. “Beyond the flow decomposition barrier”. In: *Journal of the ACM (JACM)* 45.5 (1998), pp. 783–797.
- [76] 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.
- [77] Frieda Granot, S. Thomas McCormick, Maurice Queyranne, and Fabio Tardella. “Structural and algorithmic properties for parametric minimum cuts”. In: *Math. Program.* 135.1-2 (2012), pp. 337–367. DOI: 10.1007/s10107-011-0463-1.
- [78] Martin Grötschel, László Lovász, and Alexander Schrijver. “The ellipsoid method and its consequences in combinatorial optimization”. In: *Combinatorica* 1.2 (1981), pp. 169–197.

- [79] Dan Gusfield and Charles U. Martel. “A Fast Algorithm for the Generalized Parametric Minimum Cut Problem and Applications”. In: *Algorithmica* 7.5&6 (1992), pp. 499–519. DOI: 10.1007/BF01758775.
- [80] F Maxwell Harper and Joseph A Konstan. “The movielens datasets: History and context”. In: *Acm transactions on interactive intelligent systems (tiis)* 5.4 (2015), pp. 1–19.
- [81] Nicholas J Higham. *Accuracy and stability of numerical algorithms*. Vol. 80. SIAM, 2002.
- [82] R. A. Horn and C.R. Johnson. *Topics in Matrix Analysis*. 1991.
- [83] Nicolas Hug. “Surprise: A Python library for recommender systems”. In: *Journal of Open Source Software* 5.52 (2020), p. 2174. DOI: 10.21105/joss.02174. URL: <https://doi.org/10.21105/joss.02174>.
- [84] Satoru Iwata. “A faster scaling algorithm for minimizing submodular functions”. In: *SIAM Journal on Computing* 32.4 (2003), pp. 833–840.
- [85] 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.
- [86] Martin Jaggi. “Revisiting Frank-Wolfe: Projection-free sparse convex optimization”. In: *International Conference on Machine Learning*. PMLR. 2013, pp. 427–435.
- [87] Prateek Jain, Ambuj Tewari, and Inderjit S Dhillon. “Orthogonal matching pursuit with replacement”. In: *Advances in neural information processing systems*. 2011, pp. 1215–1223.
- [88] Prateek Jain, Ambuj Tewari, and Inderjit S Dhillon. “Partial hard thresholding”. In: *IEEE Transactions on Information Theory* 63.5 (2017), pp. 3029–3038.
- [89] Prateek Jain, Ambuj Tewari, and Purushottam Kar. “On iterative hard thresholding methods for high-dimensional m-estimation”. In: *Advances in Neural Information Processing Systems*. 2014, pp. 685–693.
- [90] Stefanie Jegelka, Francis Bach, and Suvrit Sra. “Reflection methods for user-friendly submodular optimization”. In: *Advances in Neural Information Processing Systems* 26 (2013).
- [91] Haotian Jiang, Yin Tat Lee, Zhao Song, and Sam Chiu-wai Wong. “An improved cutting plane method for convex optimization, convex-concave games, and its applications”. In: *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*. 2020, pp. 944–953.
- [92] Sham M Kakade, Shai Shalev-Shwartz, and Ambuj Tewari. “Regularization techniques for learning with matrices”. In: *The Journal of Machine Learning Research* 13.1 (2012), pp. 1865–1890.
- [93] Adam Karczmarz and Piotr Sankowski. “A Deterministic Parallel APSP Algorithm and its Applications”. In: *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 255–272. DOI: 10.1137/1.9781611976465.17. URL: <https://epubs.siam.org/doi/abs/10.1137/1.9781611976465.17>.
- [94] Senanayak Sesh Kumar Karri, Francis Bach, and Thomas Pock. “Fast decomposable submodular function minimization using constrained total variation”. In: *Advances in Neural Information Processing Systems* 32 (2019).
- [95] Tarun Kathuria, Yang P Liu, and Aaron Sidford. “Unit Capacity Maxflow in Almost $O(m^{4/3})$ Time”. In: *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2020, pp. 119–130.

- [96] Jonathan Kelner and Petar Maymounkov. “Electric routing and concurrent flow cutting”. In: *International Symposium on Algorithms and Computation*. Springer. 2009, pp. 792–801.
- [97] Jonathan A Kelner, Yin Tat Lee, Lorenzo Orecchia, and Aaron Sidford. “An almost-linear-time algorithm for approximate max flow in undirected graphs, and its multicommodity generalizations”. In: *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete algorithms*. SIAM. 2014, pp. 217–226.
- [98] Jonathan A Kelner, Lorenzo Orecchia, Aaron Sidford, and Zeyuan Allen Zhu. “A simple, combinatorial algorithm for solving SDD systems in nearly-linear time”. In: *Proceedings of the 45th Annual ACM Symposium on Theory of Computing*. 2013, pp. 911–920.
- [99] Pushmeet Kohli, Philip HS Torr, et al. “Robust higher order potentials for enforcing label consistency”. In: *International Journal of Computer Vision* 82.3 (2009), pp. 302–324.
- [100] Vladimir Kolmogorov. “Minimizing a sum of submodular functions”. In: *Discrete Applied Mathematics* 160.15 (2012), pp. 2246–2258.
- [101] Yehuda Koren, Robert Bell, and Chris Volinsky. “Matrix factorization techniques for recommender systems”. In: *Computer* 42.8 (2009), pp. 30–37.
- [102] Ioannis Koutis, Gary L Miller, and David Tolliver. “Combinatorial preconditioners and multilevel solvers for problems in computer vision and image processing”. In: *International Symposium on Visual Computing*. Springer. 2009, pp. 1067–1078.
- [103] Rasmus Kyng, Richard Peng, Sushant Sachdeva, and Di Wang. “Flows in almost linear time via adaptive preconditioning”. In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. 2019, pp. 902–913.
- [104] Daniel D Lee and H Sebastian Seung. “Algorithms for non-negative matrix factorization”. In: *Advances in neural information processing systems*. 2001, pp. 556–562.
- [105] Yin Tat Lee, Satish Rao, and Nikhil Srivastava. “A new approach to computing maximum flows using electrical flows”. In: *Proceedings of the 45th Annual ACM Symposium on Theory of Computing*. 2013, pp. 755–764.
- [106] Yin Tat Lee and Aaron Sidford. “Path finding methods for linear programming: Solving linear programs in $O(\sqrt{\text{rank}})$ iterations and faster algorithms for maximum flow”. In: *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*. IEEE. 2014, pp. 424–433.
- [107] Yin Tat Lee, Aaron Sidford, and Sam Chiu-wai Wong. “A faster cutting plane method and its implications for combinatorial and convex optimization”. In: *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*. IEEE. 2015, pp. 1049–1065.
- [108] David D Lewis, Yiming Yang, Tony Russell-Rose, and Fan Li. “Rcv1: A new benchmark collection for text categorization research”. In: *Journal of machine learning research* 5.Apr (2004), pp. 361–397.
- [109] Edo Liberty and Maxim Sviridenko. “Greedy minimization of weakly supermodular set functions”. In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik. 2017.
- [110] Hui Lin and Jeff Bilmes. “Optimal selection of limited vocabulary speech corpora”. In: *Twelfth Annual Conference of the International Speech Communication Association*. 2011.
- [111] Zhouchen Lin, Minming Chen, and Yi Ma. “The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices”. In: *arXiv preprint arXiv:1009.5055* (2010).

- [112] Haoyang Liu and Rina Foygel Barber. “Between hard and soft thresholding: optimal iterative thresholding algorithms”. In: *Information and Inference: A Journal of the IMA* 9.4 (2020), pp. 899–933.
- [113] Ji Liu, Jieping Ye, and Ryohei Fujimaki. “Forward-backward greedy algorithms for general convex smooth functions over a cardinality constraint”. In: *International Conference on Machine Learning*. 2014, pp. 503–511.
- [114] Yang P Liu and Aaron Sidford. “Faster Divergence Maximization for Faster Maximum Flow”. In: *arXiv:2003.08929* (2020).
- [115] Yang P Liu and Aaron Sidford. “Faster Energy Maximization for Faster Maximum Flow”. In: *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*. 2020.
- [116] László Lovász. “Submodular functions and convexity”. In: *Mathematical programming the state of the art*. Springer, 1983, pp. 235–257.
- [117] Aleksander Mądry. “Computing maximum flow with augmenting electrical flows”. In: *57th Annual Symposium on Foundations of Computer Science*. IEEE. 2016, pp. 593–602.
- [118] Aleksander Mądry. “Navigating central path with electrical flows: From flows to matchings, and back”. In: *54th Annual Symposium on Foundations of Computer Science*. IEEE. 2013, pp. 253–262.
- [119] K. Mahler. “An inequality for the discriminant of a polynomial.” In: *Michigan Math. J.* 11.3 (Sept. 1964), pp. 257–262. DOI: 10.1307/mmj/1028999140. URL: <https://doi.org/10.1307/mmj/1028999140>.
- [120] Per-Gunnar Martinsson, Vladimir Rokhlin, and Mark Tygert. “A randomized algorithm for the decomposition of matrices”. In: *Applied and Computational Harmonic Analysis* 30.1 (2011), pp. 47–68.
- [121] Rahul Mazumder, Trevor Hastie, and Robert Tibshirani. “Spectral regularization algorithms for learning large incomplete matrices”. In: *The Journal of Machine Learning Research* 11 (2010), pp. 2287–2322.
- [122] S. Thomas McCormick. “Fast Algorithms for Parametric Scheduling Come From Extensions to Parametric Maximum Flow”. In: *Oper. Res.* 47.5 (1999), pp. 744–756. DOI: 10.1287/opre.47.5.744.
- [123] Nimrod Megiddo. “Applying Parallel Computation Algorithms in the Design of Serial Algorithms”. In: *J. ACM* 30.4 (1983), pp. 852–865. DOI: 10.1145/2157.322410. URL: <https://doi.org/10.1145/2157.322410>.
- [124] Qun Mo and Song Li. “New bounds on the restricted isometry constant δ_{2k} ”. In: *Applied and Computational Harmonic Analysis* 31.3 (2011), pp. 460–468.
- [125] Ahmad Mousavi, Mehdi Rezaee, and Ramin Ayanzadeh. “A survey on compressive sensing: classical results and recent advancements”. In: *Journal of Mathematical Modeling* 8.3 (2020), pp. 309–344.
- [126] Mukund Narasimhan and Jeff A Bilmes. “Local Search for Balanced Submodular Clusterings.” In: *IJCAI*. 2007, pp. 981–986.
- [127] Balas Kausik Natarajan. “Sparse approximate solutions to linear systems”. In: *SIAM journal on computing* 24.2 (1995), pp. 227–234.

- [128] Deanna Needell and Joel A Tropp. “CoSaMP: Iterative signal recovery from incomplete and inaccurate samples”. In: *Applied and computational harmonic analysis* 26.3 (2009), pp. 301–321.
- [129] Deanna Needell and Roman Vershynin. “Signal recovery from incomplete and inaccurate measurements via regularized orthogonal matching pursuit”. In: *IEEE Journal of selected topics in signal processing* 4.2 (2010), pp. 310–316.
- [130] Deanna Needell and Roman Vershynin. “Uniform uncertainty principle and signal recovery via regularized orthogonal matching pursuit”. In: *Foundations of computational mathematics* 9.3 (2009), pp. 317–334.
- [131] Robert Nishihara, Stefanie Jegelka, and Michael I Jordan. “On the Convergence Rate of Decomposable Submodular Function Minimization”. In: *Advances in Neural Information Processing Systems* 27 (2014), pp. 640–648.
- [132] James B Orlin. “A faster strongly polynomial time algorithm for submodular function minimization”. In: *Mathematical Programming* 118.2 (2009), pp. 237–251.
- [133] James B Orlin. “A polynomial time primal network simplex algorithm for minimum cost flows”. In: *Mathematical Programming* 78.2 (1997), pp. 109–129.
- [134] Christopher C Paige and Michael A Saunders. “LSQR: An algorithm for sparse linear equations and sparse least squares”. In: *ACM Transactions on Mathematical Software (TOMS)* 8.1 (1982), pp. 43–71.
- [135] Yagyensh Chandra Pati, Ramin Rezaifar, and Perinkulam Sambamurthy Krishnaprasad. “Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition”. In: *Proceedings of 27th Asilomar conference on signals, systems and computers*. IEEE. 1993, pp. 40–44.
- [136] Richard Peng. “Approximate undirected maximum flows in $O(m \text{ polylog } n)$ time”. In: *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete algorithms*. SIAM. 2016, pp. 1862–1867.
- [137] Richard Peng and Daniel A Spielman. “An efficient parallel solver for SDD linear systems”. In: *Proceedings of the 46th Annual ACM Symposium on Theory of computing*. 2014, pp. 333–342.
- [138] Alexandra Peste, Eugenia Iofinova, Adrian Vladu, and Dan Alistarh. “AC/DC: Alternating Compressed/DeCompressed Training of Deep Neural Networks”. In: *Thirty-Fifth Conference on Neural Information Processing Systems*. 2021.
- [139] Steffen Rendle, Li Zhang, and Yehuda Koren. “On the difficulty of evaluating baselines: A study on recommender systems”. In: *arXiv preprint arXiv:1905.01395* (2019).
- [140] R Tyrrell Rockafellar. *Convex analysis*. Vol. 36. Princeton University Press, 1970.
- [141] Alex Rubinsteyn and Sergey Feldman. *fancyimpute: Version 0.0.16*. Version v0.0.16. May 2016. DOI: 10.5281/zenodo.51773. URL: <https://doi.org/10.5281/zenodo.51773>.
- [142] Ludwig Schmidt. “Algorithms above the noise floor”. PhD thesis. Massachusetts Institute of Technology, Cambridge, USA, 2018. URL: <http://hdl.handle.net/1721.1/118098>.
- [143] 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.
- [144] Shai Shalev-Shwartz, Alon Gonen, and Ohad Shamir. “Large-scale convex minimization with a low-rank constraint”. In: *Proceedings of the 28th International Conference on International Conference on Machine Learning*. 2011, pp. 329–336.

- [145] Shai Shalev-Shwartz and Yoram Singer. “Convex repeated games and Fenchel duality”. In: *NIPS*. Vol. 6. 2006, pp. 1265–1272.
- [146] Shai Shalev-Shwartz, Nathan Srebro, and Tong Zhang. “Trading accuracy for sparsity in optimization problems with sparsity constraints”. In: *SIAM Journal on Optimization* 20.6 (2010), pp. 2807–2832.
- [147] Ishant Shanu, Chetan Arora, and Parag Singla. “Min norm point algorithm for higher order MRF-MAP inference”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2016, pp. 5365–5374.
- [148] Jie Shen and Ping Li. “On the iteration complexity of support recovery via hard thresholding pursuit”. In: *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org. 2017, pp. 3115–3124.
- [149] Jie Shen and Ping Li. “Partial hard thresholding: Towards a principled analysis of support recovery”. In: *Advances in Neural Information Processing Systems*. 2017, pp. 3124–3134.
- [150] Jonah Sherman. “Area-convexity, ℓ_∞ regularization, and undirected multicommodity flow”. In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. 2017, pp. 452–460.
- [151] Jonah Sherman. “Generalized preconditioning and undirected minimum-cost flow”. In: *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM. 2017, pp. 772–780.
- [152] Jonah Sherman. “Nearly maximum flows in nearly linear time”. In: *54th Annual Symposium on Foundations of Computer Science*. IEEE. 2013, pp. 263–269.
- [153] Aaron Sidford and Kevin Tian. “Coordinate methods for accelerating ℓ_∞ regression and faster approximate maximum flow”. In: *59th Annual Symposium on Foundations of Computer Science*. IEEE. 2018, pp. 922–933.
- [154] Raghav Somani, Chirag Gupta, Prateek Jain, and Praneeth Netrapalli. “Support recovery for orthogonal matching pursuit: upper and lower bounds”. In: *Advances in Neural Information Processing Systems*. 2018, pp. 10814–10824.
- [155] Daniel A Spielman and Shang-Hua Teng. “Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems”. In: *Proceedings of the 36th annual ACM symposium on Theory of computing*. 2004, pp. 81–90.
- [156] Daniel A Spielman and Shang-Hua Teng. “Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time”. In: *Journal of the ACM (JACM)* 51.3 (2004), pp. 385–463.
- [157] Nathan Srebro, Jason Rennie, and Tommi Jaakkola. “Maximum-margin matrix factorization”. In: *Advances in neural information processing systems* 17 (2004), pp. 1329–1336.
- [158] Peter Stobbe and Andreas Krause. “Efficient minimization of decomposable submodular functions”. In: *Advances in Neural Information Processing Systems* 23 (2010).
- [159] Arthur Szlam, Yuval Kluger, and Mark Tygert. “An implementation of a randomized algorithm for principal component analysis”. In: *arXiv preprint arXiv:1412.3510* (2014).
- [160] Robert Endre Tarjan, Julie Ward, Bin Zhang, Yunhong Zhou, and Jia Mao. “Balancing Applied to Maximum Network Flow Problems”. In: *Algorithms - ESA 2006, 14th Annual European Symposium, Zurich, Switzerland, September 11-13, 2006, Proceedings*. 2006, pp. 612–623. DOI: 10.1007/11841036_55. URL: https://doi.org/10.1007/11841036%5C_55.

- [161] Robert Tibshirani. “Regression shrinkage and selection via the lasso”. In: *Journal of the Royal Statistical Society: Series B (Methodological)* 58.1 (1996), pp. 267–288.
- [162] Sivan Toledo and Haim Avron. “Combinatorial preconditioners”. In: *Combinatorial Scientific Computing* (2010), pp. 69–93.
- [163] Andrew Tulloch. *Fast Randomized SVD*. <https://research.fb.com/blog/2014/09/fast-randomized-svd/>. 2014.
- [164] Antoine Vacavant, Thierry Chateau, Alexis Wilhelm, and Laurent Lequière. “A benchmark dataset for outdoor foreground/background extraction”. In: *Asian Conference on Computer Vision*. Springer. 2012, pp. 291–300.
- [165] Nate Veldt, Austin R Benson, and Jon Kleinberg. “Minimizing Localized Ratio Cut Objectives in Hypergraphs”. In: *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2020, pp. 1708–1718.
- [166] Sara Vicente, Vladimir Kolmogorov, and Carsten Rother. “Joint optimization of segmentation and appearance models”. In: *2009 IEEE 12th International Conference on Computer Vision*. IEEE. 2009, pp. 755–762.
- [167] James Hardy Wilkinson. *Rounding errors in algebraic processes*. Courier Corporation, 1994.
- [168] Chee-Keng Yap et al. *Fundamental problems of algorithmic algebra*. Vol. 49. Oxford University Press Oxford, 2000.
- [169] Xiaotong Yuan, Ping Li, and Tong Zhang. “Exact recovery of hard thresholding pursuit”. In: *Advances in Neural Information Processing Systems*. 2016, pp. 3558–3566.
- [170] Tong Zhang. “Sparse recovery with orthogonal matching pursuit under RIP”. In: *IEEE Transactions on Information Theory* 57.9 (2011), pp. 6215–6221.
- [171] Le Zheng, Arian Maleki, Haolei Weng, Xiaodong Wang, and Teng Long. “Does ℓ_p -Minimization Outperform ℓ_1 -Minimization?” In: *IEEE Transactions on Information Theory* 63.11 (2017), pp. 6896–6935.

Chapter 9

Appendix

9.1 Appendix for Chapter 3

9.1.1 Initializing the Interior Point Method

We require finding an initial \mathbf{f}_0 for which the solution can be easily brought to centrality. To do so, we modify the original graph by adding a set of $O(m)$ arcs with high cost $c_\infty = (m + 1)\|\mathbf{c}\|_\infty$ such that the flow which pushes exactly $1/2$ units on every arc in the modified graph routes the original demand \mathbf{d} .

The consequence of adding such these edges is that, while the solution to the original problem remains unchanged, since it will never be beneficial to use an arc with cost c_∞ in the optimal solution. Meanwhile initializing $\mathbf{f}_0 = \mathbf{1}/2$ ensures that the flow on each arc is equally far from the upper and the lower barrier, and therefore their contributions to the gradient will cancel. This will make centering trivially easy.

First let us show that such an augmentation is indeed possible.

Lemma 9.1.1. *Let $G = (V, E, \mathbf{c})$ be a directed graph with $|E| = m$ arcs with unit capacity, $|V| = n$ vertices, costs on arcs $\mathbf{c} \geq 0$, and let $\mathbf{d} \in \mathbb{Z}^n$ be a demand vector $\sum_{i=1}^n d_i = 0$.*

Then there exists a graph $G' = (V', E', \mathbf{c}')$ with at most $2m$ unit-capacity arcs, and a demand vector \mathbf{d}' such that the flow $\mathbf{f}' = \mathbf{1}/2$ routes the demand \mathbf{d}' in G' .

Furthermore if $(\mathbf{f}')^$ is a flow in G' with $\mathbf{0} \leq (\mathbf{f}')^* \leq \mathbf{1}$ and which routes \mathbf{d}' such that the cost $\langle \mathbf{c}', \mathbf{f}' \rangle$ is minimized, then one can convert it in $O(m)$ time into a flow \mathbf{f}^* which routes \mathbf{d} in G , such that $\mathbf{0} \leq \mathbf{f}^* \leq \mathbf{1}$ and the cost $\langle \mathbf{c}, \mathbf{f} \rangle$ is minimized, or certify that no such flow exists.*

Proof. First we construct the graph G' . Let $V' = V \cup \{v_0\}$, where v_0 is a new vertex. Initialize $S = \emptyset$. For each vertex $v \in V$, let $\ell(v) = d_v - \frac{1}{2}(|E^-(v)| - |E^+(v)|)$, representing the excess flow at vertex v after routing $\mathbf{f} = \mathbf{1}/2$ on each arc. Let $c_\infty = (m + 1)\|\mathbf{c}\|_\infty$. For each v where $\ell(v) > 0$, create $2\ell(v)$ arcs (v, v_0) with cost c_∞ , and add them to S . Similarly, for each v where $\ell(v) < 0$ create $-2\ell(v)$ arcs (v_0, v) with cost c_∞ and add them to S . Note that $2\ell(v)$ is an integer, since \mathbf{d} is integral. Let $E' = E \cup S$ and \mathbf{c}' be the corresponding cost vector where arcs in E preserve their original cost \mathbf{c} , while those in S have cost c_∞ . Let $d'_v = d_v$ for all $v \in V$ and $d'_{v_0} = 0$.

Now let $(\mathbf{f}')^*$ be the minimum cost flow in G' which satisfies capacity constraints. If $(\mathbf{f}')^*$ is not supported on any arcs in S , then $(\mathbf{f}')^*$ is also a feasible flow in G . Furthermore it must be the optimal flow in G , since otherwise $(\mathbf{f}')^*$ would not have been optimal in G' . If $(\mathbf{f}')^*$ has nonzero flow on some arc in S , then we must conclude that it is impossible to route \mathbf{d} in G while satisfying capacity constraints.

Suppose it were possible to do so using a flow \mathbf{f}^* . Then, consider the circulation $\mathbf{g} = \mathbf{f}^* - (\mathbf{f}')^*$. Now convert \mathbf{g} into a *minimal* circulation $\tilde{\mathbf{g}}$ which preserves the flows on S as follows: while \mathbf{g} contains a cycle without any arcs in S , decrease the value of all flows along that cycle by the minimum value of the flow along it. This operation always zeroes out the flow on at least one edge, so the process must finish. Furthermore note that the cost of any such cycle must be non-positive, since otherwise we contradict the optimality of \mathbf{f}^* .

At this point we are left with a circulation $\tilde{\mathbf{g}}$ which does not contain any cycles supported only in E , and whose cost is at least $\langle \mathbf{c}, \mathbf{g} \rangle \geq 0$; the latter inequality follows from the optimality of $(\mathbf{f}')^*$.

Using the fact that $\tilde{\mathbf{g}}$ does not contain cycles supported in E , we get that the restriction to the arcs in E , $\tilde{\mathbf{g}}|_E$ satisfies $\|\tilde{\mathbf{g}}|_E\|_1 \leq m \cdot \|\tilde{\mathbf{g}}|_S\|_1$, where $\tilde{\mathbf{g}}|_S$ is the corresponding restriction to S . Therefore $|\langle \mathbf{c}, \tilde{\mathbf{g}}|_E \rangle| \leq \|\mathbf{c}\|_\infty \cdot \|\tilde{\mathbf{g}}|_E\|_1 \leq m \cdot \|\mathbf{c}\|_\infty \cdot \|\tilde{\mathbf{g}}|_S\|_1$. Since by definition $\langle \mathbf{c}, \tilde{\mathbf{g}}|_S \rangle = -c_\infty \cdot \|\tilde{\mathbf{g}}|_S\|_1$, we conclude that $\langle \mathbf{c}, \tilde{\mathbf{g}} \rangle = \langle \mathbf{c}, \tilde{\mathbf{g}}|_E \rangle + \langle \mathbf{c}, \tilde{\mathbf{g}}|_S \rangle \leq m \cdot \|\mathbf{c}\|_\infty \cdot \|\tilde{\mathbf{g}}|_S\|_1 - c_\infty \cdot \|\tilde{\mathbf{g}}|_S\|_1 = \|\tilde{\mathbf{g}}|_S\|_1 \cdot (m\|\mathbf{c}\|_\infty - c_\infty) < 0$. This yields a contradiction, so a feasible \mathbf{f}^* can not possibly exist. \square

At this point, using the reduction given above, we can assume without loss of generality that the flow $\mathbf{f} = 1/2 \cdot \mathbf{1}$ routes the demand \mathbf{d} .

Lemma 9.1.2. *If $\mathbf{f} = 1/2 \cdot \mathbf{1}$ routes the input demand \mathbf{d} , then in the time dominated by $O(\log \log m)$ residual correction steps, we can produce a solution $\mathbf{f} = \mathbf{C}\mathbf{x}$ and a set of weights $\mathbf{w} \geq \mathbf{1}$ such that $\|\mathbf{w}\|_1 \leq 2m + 1$ and $\nabla F_\mu^{\mathbf{w}}(\mathbf{x}) = \mathbf{0}$, for $\mu \leq 2\|\mathbf{c}\|_2$.*

Proof. For this flow $\mathbf{f} = \mathbf{C}\mathbf{x}$ we have that $\mathbf{s}^+ = \mathbf{s}^- = \mathbf{1}/2$ and the corresponding residual satisfies $\nabla F_\mu^{\mathbf{1}}(\mathbf{x}) = \frac{\mathbf{C}^\top \mathbf{c}}{\mu}$. Per Definition 3.3.3 we can certify an upper bound on $\mathcal{E}_{1,s}(\nabla F_\mu^{\mathbf{1}}(\mathbf{x}))$ using $\mathbf{y} = (\mathbf{0}; \mathbf{c}/\mu)$ which shows that

$$\mathcal{E}_{1,s}(\nabla F_\mu^{\mathbf{1}}(\mathbf{x})) \leq \frac{1}{2} \cdot \langle (\mathbf{s}^+; \mathbf{s}^-)^2, (\mathbf{0}; \mathbf{c}/\mu)^2 \rangle = \frac{1}{2} \cdot \frac{1}{4\mu^2} \cdot \|\mathbf{c}\|_2^2.$$

Therefore setting $\mu = \|\mathbf{c}\|_2$ we have that the energy is at most $1/8$, which fulfills the conditions required by Corollary 3.3.7 and Lemma 3.3.8 to produce in the time dominated by $O(\log \log m)$ residual correction steps a solution \mathbf{x} and a weight vector $\mathbf{w} \geq \mathbf{1}$ such that $\|\mathbf{w} - \mathbf{1}\|_1 \leq 2m^{-9} \leq 1$ for which $\nabla F_{\mu'}^{\mathbf{w}}(\mathbf{x}) = \mathbf{0}$, where $\mu' \leq \|\mathbf{c}\|_2(1 + m^{-10}) \leq 2\|\mathbf{c}\|_2$. \square

9.1.2 Preconditioning Proof

In this section we provide the proof for Lemma 3.4.7.

Proof of Lemma 3.4.7. Let $\tilde{\mathbf{f}}_\star = \tilde{\mathbf{f}} + \tilde{\mathbf{f}}'$. Writing the optimality condition for (3.18) we obtain

$$\mathbf{C}_\star^\top (\mathbf{r}\tilde{\mathbf{f}}_\star - \mathbf{h}) = \mathbf{0}, \quad (9.1)$$

where

$$\mathbf{r} = \begin{cases} \frac{\mathbf{w}^+}{(s^+)^2} + \frac{\mathbf{w}^-}{(s^-)^2} + R_p \cdot (\tilde{\mathbf{f}})^{p-2} & \text{for edges in } E, \\ R_\star + R_p \cdot (\tilde{\mathbf{f}}')^{p-2} & \text{for edges in } E'. \end{cases}$$

Since \mathbf{C}_\star is a cycle basis for G_\star , the condition (9.1) implies that along any cycle in G_\star the (appropriately signed) sum of weights $\mathbf{r}\tilde{\mathbf{f}}_\star - \mathbf{h}$ is exactly 0. This means that these weights are determined by an embedding of the vertices of G_\star onto the line, c.f. Lemma 9.1.3. Hence there exists a vector

$\phi \in \mathbb{R}^{|\mathcal{V} \cup \{v_\star\}|}$ such that for every edge $(u, v) \in E \cup E'$:

$$r_e(\tilde{f}_\star)_e - h_e = \phi_u - \phi_v. \quad (9.2)$$

Also, to shorten notation, let us define

$$\bar{w} = w^+ + w^-. \quad (9.3)$$

Next we will prove that the coordinates of ϕ must lie within a short interval. The intuition here relies on the fact that the preconditioning edges make the graph G_\star have good enough expansion; in turn, using an argument similar to [96] which was subsequently employed under various forms in [97, 118, 42] we argue that good expansion means that all the vertices are close to each other in the potential embedding.

Given a scalar t , let S_t be the set of edges (u, v) for which $t \in (\min\{\phi_u, \phi_v\}, \max\{\phi_u, \phi_v\})$. Our proof proceeds by lower bounding for each t :

$$\sum_{(u,v) \in S_t} \frac{\bar{w}_{uv}}{|\phi_u - \phi_v|} \geq \frac{\left(\sum_{(u,v) \in S_t} \frac{\sqrt{\bar{w}_{uv}}}{\sqrt{r_{uv}}} \right)^2}{\sum_{(u,v) \in S_t} \frac{|\phi_u - \phi_v|}{r_{uv}}}, \quad (9.4)$$

which we obtained using Cauchy-Schwarz.

Next we observe that the quantity in the denominator is upper bounded by

$$\sum_{(u,v) \in S_t} \frac{|\phi_u - \phi_v|}{r_{uv}} \leq \sum_{(u,v) \in S_t} \frac{|h_{uv}|}{r_{uv}}. \quad (9.5)$$

We can see why this is true by using the fact that \tilde{f}_\star is a circulation, therefore along any cut S_t one has

$$\sum_{\substack{(u,v) \in S_t \\ \phi_u > \phi_v}} (\tilde{f}_\star)_{uv} = \sum_{\substack{(u,v) \in S_t \\ \phi_u \leq \phi_v}} (\tilde{f}_\star)_{uv},$$

i.e. the sum of the values of flows going from left to right is equal to the sum of values of flows going from right to left in the embedding.

By substituting $(\tilde{f}_\star)_{uv}$ with the value determined from (9.2) we equivalently obtain that

$$\sum_{\substack{(u,v) \in S_t \\ \phi_u > \phi_v}} \frac{h_{uv} + |\phi_u - \phi_v|}{r_{uv}} = \sum_{\substack{(u,v) \in S_t \\ \phi_u \leq \phi_v}} \frac{h_{uv} - |\phi_u - \phi_v|}{r_{uv}},$$

and by rearranging

$$\sum_{(u,v) \in S_t} \frac{|\phi_u - \phi_v|}{r_{uv}} = \sum_{\substack{(u,v) \in S_t \\ \phi_u > \phi_v}} \frac{-h_{uv}}{r_{uv}} + \sum_{\substack{(u,v) \in S_t \\ \phi_u \leq \phi_v}} \frac{h_{uv}}{r_{uv}} \leq \sum_{(u,v) \in S_t} \frac{|h_{uv}|}{r_{uv}}.$$

Therefore, plugging (9.5) into (9.4), we obtain

$$\sum_{(u,v) \in S_t} \frac{\bar{w}_{uv}}{|\phi_u - \phi_v|} \geq \frac{\left(\sum_{(u,v) \in S_t} \frac{\sqrt{\bar{w}_{uv}}}{\sqrt{r_{uv}}} \right)^2}{\sum_{(u,v) \in S_t} \frac{|h_{uv}|}{r_{uv}}} \geq \frac{\sum_{(u,v) \in S_t} \frac{\sqrt{\bar{w}_{uv}}}{\sqrt{r_{uv}}}}{\max_{(u,v) \in S_t} \frac{|h_{uv}|/r_{uv}}{\sqrt{\bar{w}_{uv}}/\sqrt{r_{uv}}}} = \left(\sum_{(u,v) \in S_t} \frac{\sqrt{\bar{w}_{uv}}}{\sqrt{r_{uv}}} \right) \cdot \frac{1}{\|\mathbf{h}\bar{\mathbf{w}}^{-1/2}\mathbf{r}^{-1/2}\|_\infty}. \quad (9.6)$$

At this point we can prove that all the values in ϕ lie within a small interval. In order to do so we will crucially use the augmenting edges, which endow G_\star with better expansion properties. Suppose w.l.o.g. that $\phi_{v_\star} = 0$. For every edge $e = (u, v) \in E \cup E'$, let $m_e = \min\{\phi_u, \phi_v\}$, $M_e = \max\{\phi_u, \phi_v\}$. For all edges in $e = (u, v_\star) \in E'$ (of which multiple copies can occur) we define

$$\bar{w}_e = 1.$$

Hence we may define for $t \geq 0$:

$$F(t) = \sum_{\substack{(u,v) \in E \cup E' \\ M_{uv} \geq t}} \bar{w}_{uv} \cdot \frac{M_{uv} - \max\{m_{uv}, t\}}{|\phi_u - \phi_v|}, \quad (9.7)$$

which represents the sum of weighted fractions of edges that are on the right side of the cut at position t on the real line.

Our goal is to show that since $F(t)$ decreases very fast, we do not need to increase t very much before we run out of edges i.e. $F(t)$ becomes 0. Indeed, (9.6) offers a lower bound on the instantaneous decrease of $F(t)$, as t increases. The reason is that all the augmenting edges (v_\star, u) for $\phi_u > t$ appear in the cut. This also means that $F(t) > 0$ if $S_t \neq \emptyset$ and $F(t) = 0$ otherwise.

Intuitively (9.6) states that when slightly increasing t , $F(t)$ must decrease by a constant factor, scaled by the change in t . To formalize this we simply use the fact that

$$\sum_{(u,v) \in S_t} \frac{\sqrt{\bar{w}_{uv}}}{\sqrt{r_{uv}}} = \sum_{(u,v) \in S_t} \frac{\bar{w}_{uv}}{\sqrt{\bar{w}_{uv}r_{uv}}} \geq \sum_{(u,v_\star) \in S_t} \frac{\bar{w}_{uv_\star}}{\sqrt{\bar{w}_{uv_\star}r_{uv_\star}}} \geq \frac{1}{\max_{e \in E'} \sqrt{\bar{w}_e r_e}} \cdot \frac{F(t)}{2}, \quad (9.8)$$

which follows from the inequality

$$F(t) \leq \sum_{e \in S_t} \bar{w}_e \leq 2 \cdot \sum_{e \in S_t \cap E'} \bar{w}_e.$$

The latter is ensured by the fact that by definition each vertex $v \in V$ is incident to at least $\sum_{e \in E: e \sim v} \bar{w}_e$ augmenting edges in E' . Furthermore, even if $F(t)$ is very small but S_t is still nonempty, we can use the lower bound

$$\sum_{(u,v) \in S_t} \frac{\sqrt{\bar{w}_{uv}}}{\sqrt{r_{uv}}} = \sum_{(u,v) \in S_t} \frac{\bar{w}_{uv}}{\sqrt{\bar{w}_{uv}r_{uv}}} \geq \frac{|S_t|}{\max_{e \in E'} \sqrt{\bar{w}_e r_e}} \geq \frac{1}{\max_{e \in E'} \sqrt{\bar{w}_e r_e}}. \quad (9.9)$$

Combining (9.6), (9.8), and (9.9) we obtain that if S_t is nonempty, or equivalently $F(t) > 0$, then

$$\sum_{(u,v) \in S_t} \frac{\bar{w}_{uv}}{|\phi_u - \phi_v|} \geq \frac{1}{\|\mathbf{h}\bar{\mathbf{w}}^{-1/2}\mathbf{r}^{-1/2}\|_\infty \cdot \max_{e \in E'} \sqrt{\bar{w}_e r_e}} \cdot \max\{F(t)/2, 1\}. \quad (9.10)$$

As a matter of fact, this tells us that F must decrease very fast, since from the definition (9.7) we have that F is a continuous function, differentiable almost everywhere, such that at all points $t \geq 0$ where it is differentiable it satisfies

$$\frac{d}{dt}F(t) = - \sum_{\substack{(u,v) \in E \cup E' \\ M_{uv} \geq t}} \frac{\bar{w}_{uv}}{|\phi_u - \phi_v|} \leq - \sum_{(u,v) \in S_t} \frac{\bar{w}_{uv}}{|\phi_u - \phi_v|}.$$

Thus using (9.10) and Lemma 9.1.4 we obtain that $F(T) = 0$ for

$$\begin{aligned} T &= \|\mathbf{h}\bar{\mathbf{w}}^{-1/2}\mathbf{r}^{-1/2}\|_\infty \cdot \max_{e \in E'} \sqrt{\bar{w}_e r_e} \cdot (1 + 2 \log F(0)) \\ &\leq \|\mathbf{h}\bar{\mathbf{w}}^{-1/2}\mathbf{r}^{-1/2}\|_\infty \cdot \max_{e \in E'} \sqrt{\bar{w}_e r_e} \cdot (1 + 2 \log(4\|\mathbf{w}\|_1)), \end{aligned} \quad (9.11)$$

where the last inequality follows from accounting the weights of the augmenting edges in E' .

Using the identical argument for vertices embedded to the left of v_* , we conclude that (9.11) yields an upper bound on $M_e - m_e$, and therefore, for all edges $e \in E$

$$\begin{aligned} |r_e \tilde{f}_e - h_e| &\leq \max_{e \in E'} \sqrt{\bar{w}_e r_e} \cdot \|\mathbf{h}\bar{\mathbf{w}}^{-1/2}\mathbf{r}^{-1/2}\|_\infty \cdot 2(1 + 2 \log(4\|\mathbf{w}\|_1)) \\ &\leq \max_{e \in E'} \sqrt{\bar{w}_e r_e} \cdot \|\mathbf{h}\bar{\mathbf{w}}^{-1/2}\mathbf{r}^{-1/2}\|_\infty \cdot 32 \log \|\mathbf{w}\|_1. \end{aligned}$$

Now we use Lemma 3.4.6 to obtain that for all $(u, v_*) = e \in E'$:

$$\bar{w}_e r_e = R_* + R_p \cdot (\tilde{f}'_e)^{p-2} \leq R_* + R_p \cdot \|\tilde{\mathbf{f}}'\|_\infty^{p-2},$$

which ensures that

$$\left| \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} + R_p (\tilde{f}'_e)^{p-2} \right) \tilde{f}_e - h_e \right| \leq \left(R_* + R_p \cdot \|\tilde{\mathbf{f}}'\|_\infty^{p-2} \right)^{1/2} \cdot \left\| \frac{\mathbf{h}}{\sqrt{\bar{\mathbf{w}}\mathbf{r}}} \right\|_\infty \cdot 32 \log \|\mathbf{w}\|_1, \quad (9.12)$$

and therefore that

$$\left| \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) \tilde{f}_e - h_e \right| \leq \left| R_p \cdot (\tilde{f}'_e)^{p-1} \right| + \left(R_* + R_p \cdot \|\tilde{\mathbf{f}}'\|_\infty^{p-2} \right)^{1/2} \quad (9.13)$$

$$\cdot \left\| \frac{\mathbf{h}}{\sqrt{(\mathbf{w}^+ + \mathbf{w}^-) \left(\frac{w^+}{(s^+)^2} + \frac{w^-}{(s^-)^2} \right)}} \right\|_\infty \cdot 32 \log \|\mathbf{w}\|_1, \quad (9.14)$$

which is what we needed.

We can furthermore establish a similar upper bound on

$$\begin{aligned}
\left| \frac{w_e^+}{(s_e^+)^2} \cdot \tilde{f}_e \right| &\leq \left| \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} + R_p \cdot (\tilde{f}_e)^{p-2} \right) \tilde{f}_e \right| \\
&\leq |h_e| + \max_{e \in E'} \sqrt{\bar{w}_e r_e} \cdot \|\mathbf{h} \bar{\mathbf{w}}^{-1/2} \mathbf{r}^{-1/2}\|_\infty \cdot 32 \log \|\mathbf{w}\|_1 \\
&\leq |h_e| + \left(R_\star + R_p \cdot \|\tilde{\mathbf{f}}_\star\|_\infty^{p-2} \right)^{1/2} \\
&\quad \cdot \left\| \frac{\mathbf{h}}{\sqrt{(\mathbf{w}^+ + \mathbf{w}^-) \left(\frac{\mathbf{w}^+}{(s^+)^2} + \frac{\mathbf{w}^-}{(s^-)^2} \right)}} \right\|_\infty \cdot 32 \log \|\mathbf{w}\|_1,
\end{aligned}$$

and an identical upper bound on $\left| \frac{w_e^-}{(s_e^-)^2} \cdot \tilde{f}_e \right|$.

□

Lemma 9.1.3. *Let G be a graph with n vertices and m edges, and \mathbf{C} be a matrix that encodes a cycle basis for G , in the sense that for any \mathbf{x} , $\mathbf{C}\mathbf{x}$ is a circulation in G and for any circulation \mathbf{f} in G , there exists a vector \mathbf{x} such that $\mathbf{f} = \mathbf{C}\mathbf{x}$. Suppose that $\mathbf{y} \in \mathbb{R}^m$ is a vector such that $\mathbf{C}^\top \mathbf{y} = 0$. Then there exists a vector $\phi \in \mathbb{R}^n$ such that for all $(u, v) \in E$ one has that $y_{uv} = \phi_u - \phi_v$.*

Proof. By definition, the image of \mathbf{C} is the space of circulations in G . Therefore the kernel of \mathbf{C}^\top is orthogonal to the space of circulations in G , and therefore so is the vector \mathbf{y} . Now consider the incidence matrix $\mathbf{B} \in \mathbb{R}^{m \times n}$, which is constructed as follows: for each edge $(u, v) \in E$, add a row in with two nonzero entries, +1 at position u , and -1 at position v . One can easily see that $\ker \mathbf{B}^\top$ is exactly the space of circulations, as the \mathbf{B}^\top operator acts on flows by returning the vector of demands that they route. Hence \mathbf{y} lies in the image of \mathbf{B} , i.e. $\mathbf{y} = \mathbf{B}\phi$. By the definition of \mathbf{B} , the conclusion follows. □

Lemma 9.1.4. *Let $F : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ be a continuous function, differentiable almost everywhere, such that $F(0) \geq 0$. Suppose that for all $t \geq 0$ where $F(t) > 0$ and $dF(t)/dt$ exists, we have*

$$\frac{d}{dt} F(t) \leq -\frac{1}{\alpha} \max \left\{ \frac{F(t)}{2}, 1 \right\},$$

for some $\alpha > 0$. Then

$$F(\alpha(1 + 2 \log F(0))) = 0.$$

Proof. Let $T > 0$ be any point for which $F(T) > 0$. From the hypothesis we know that the instantaneous decrease in F at all points $t \in [0, T]$ is at least dt/α . Hence we have that:

$$F(T) \leq \min_{0 \leq t \leq T} F(t) - \frac{T-t}{\alpha}.$$

Furthermore, using the fact that the instantaneous decrease in $F(t)$ is at least $F(t)/(2\alpha)$, we solve the corresponding ODE to obtain that for all t ,

$$F(t) \leq F(0) \exp \left(-\frac{t}{2\alpha} \right).$$

Combining the two inequalities, and setting $t = 2\alpha \log F(0)$, we get that

$$F(T) \leq 1 - \frac{T - 2\alpha \log F(0)}{\alpha} = (1 + 2 \log F(0)) - \frac{T}{\alpha}.$$

which implies that $T \leq \alpha(1 + 2 \log F(0))$, since $F(T) \geq 0$. \square

Finally, we discuss how to adapt the proof of Lemma 3.4.7 to obtain the Lemma 3.6.4.

Proof of Lemma 3.6.4. This proof is identical to that of Lemma 3.4.7. The main difference consists of including the α factors in the definition of \mathbf{r} for edges in e , more specifically, we let:

$$\mathbf{r} = \begin{cases} \frac{\alpha^+ w^+}{(s^+)^2} + \frac{\alpha^- w^-}{(s^-)^2} + R_p \cdot (\tilde{\mathbf{f}})^{p-2} & \text{for edges in } E, \\ R_x + R_p \cdot (\tilde{\mathbf{f}})^{p-2} & \text{for edges in } E'. \end{cases}$$

Using this new definition for \mathbf{r} , the remaining proof carries over by following the exact same steps as before. \square

9.1.3 Solving the Mixed Objective

Invoking the Solver

Producing a high precision solution to the regularized objective in (3.18) can be done efficiently in our particular setting, where we aim to optimize a mixed ℓ_2 - ℓ_p objective in the space of circulations. To this extent we use the following result from [103], also restated and improved in [2].

Theorem 9.1.5. *For any $p \geq 2$, given weights $\mathbf{r} \in \mathbb{R}_{\geq 0}^{|E|}$, and a cost vector $\mathbf{g} \in \mathbb{R}^{|E|}$ define the function defined over circulations \mathbf{f} in G :*

$$\text{val}(\mathbf{f}) = \sum_e g_e f_e + r_e f_e^2 + |f_e|^p.$$

Given any circulation \mathbf{f} for which all the parameters are bounded by $2^{(\log n)^{O(1)}}$ we can compute a circulation $\tilde{\mathbf{f}}$ such that

$$\text{val}(\tilde{\mathbf{f}}) - OPT \leq \frac{1}{2^{(\log m)^{O(1)}}} (\text{val}(\mathbf{f}) - OPT) + \frac{1}{2^{(\log m)^{O(1)}}}$$

in $2^{O(p^{3/2})} m^{1+O(1/\sqrt{p})}$ time.

For large values of p this solves the regularized objective defined in (3.18) to high precision in almost linear time $O(m^{1+o(1)})$, which is comparably fast to the time required to minimize a convex quadratic function in the space of circulations via fast Laplacian solvers.

Discussion on Solver Errors

Throughout Chapter 3 we assume that the solutions to the regularized objective are exact. This is not exactly true due to the approximate nature of the solver specified in Theorem 9.1.5. Instead, we argue that the entire analysis we showed carries over even if the solver returns a solution which carries some small error.

Consider the equivalent problem of minimizing the negative of the objective $\Phi(\mathbf{x})$ stated in (3.18). We verify that at a point \mathbf{x} such that the current flow $\tilde{\mathbf{f}} = \mathbf{C}_\star \mathbf{x}$, its Hessian is

$$\nabla^2 \Phi(\mathbf{x}) = \mathbf{C}_\star^\top \mathbf{D} \mathbf{C}_\star$$

where \mathbf{D} is a diagonal matrix whose entries are defined such that

$$(\mathbf{D})_{e,e} = \begin{cases} \frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} + (p-1)R_p(\tilde{f}_e)^{p-2} & \text{if } e \in E, \\ R_\star + R_p(p-1)(\tilde{f}_e)^{p-2} & \text{if } e \in E'. \end{cases}$$

Our choice of regularization parameters yield an upper bound on $\|\tilde{\mathbf{f}}\|_\infty$ as we showed in the proof of Lemma 3.4.6 (Equation (3.43)), which together with our choice of regularization parameters (Section 3.4.2) and the invariants that $\mathbf{1} \leq \mathbf{w}$ and $\|\mathbf{w}\|_1 = O(m)$ ensure that $R_p(\tilde{f}_e)^{p-2} \leq \tilde{O}((\delta^2 m)^3) = o(m)$. Assuming that we always maintain all our slacks large enough i.e. for all e : $\tau^{-1} \leq f_e \leq 1 - \tau^{-1}$, for $\tau = m^{O(1)}$, we see that \mathbf{D} is always well-conditioned in the sense that all of its diagonal entries are between τ^{-2} and $M = O(m\tau^2)$. We will discuss later how to maintain this property.¹

Under this assumption we can therefore use basic tools from convex analysis to argue that after optimizing Φ to high precision, the gradient of the returned solution is small. Let \mathbf{x} be the solution returned by the high-precision solver such that $\Phi(\mathbf{x}) \leq \Phi(\mathbf{x}^*) + \varepsilon$. Letting $\nabla \Phi(\mathbf{x}) = \mathbf{C}_\star^\top \mathbf{h}$ for some vector \mathbf{h} , we can write

$$\begin{aligned} \|\nabla \Phi(\mathbf{x})\|_{\mathbf{C}_\star^\top \mathbf{C}_\star}^* &= \left\| \int_0^1 \nabla^2 \Phi((1-t)\mathbf{x}^* + t\mathbf{x})(\mathbf{x} - \mathbf{x}^*) dt \right\|_{\mathbf{C}_\star^\top \mathbf{C}_\star}^* \\ &\leq \max_{t \in [0,1]} \left\| \nabla^2 \Phi((1-t)\mathbf{x}^* + t\mathbf{x}) \right\|_{\|\cdot\|_{\mathbf{C}_\star^\top \mathbf{C}_\star} \rightarrow \|\cdot\|_{\mathbf{C}_\star^\top \mathbf{C}_\star}^*} \cdot \|\mathbf{x} - \mathbf{x}^*\|_{\mathbf{C}_\star^\top \mathbf{C}_\star} \\ &\leq \max_{\substack{\mathbf{D} \text{ diagonal} \\ \mathbf{D}_{ii} \in [M^{-1}, M]}} \left\| \mathbf{C}_\star^\top \mathbf{D} \mathbf{C}_\star \right\|_{\|\cdot\|_{\mathbf{C}_\star^\top \mathbf{C}_\star} \rightarrow \|\cdot\|_{\mathbf{C}_\star^\top \mathbf{C}_\star}^*} \cdot \|\mathbf{x} - \mathbf{x}^*\|_{\mathbf{C}_\star^\top \mathbf{C}_\star}. \end{aligned}$$

Here we use the notation

$$\begin{aligned} \|\mathbf{x}\|_{\mathbf{C}_\star^\top \mathbf{C}_\star} &= \|\mathbf{C}_\star \mathbf{x}\|_2, \\ \|\mathbf{f}\|_{\mathbf{C}_\star^\top \mathbf{C}_\star}^* &= \max_{\mathbf{x}: \|\mathbf{x}\|_{\mathbf{C}_\star^\top \mathbf{C}_\star} \leq 1} \langle \mathbf{f}, \mathbf{x} \rangle, \end{aligned}$$

and

$$\|\mathbf{A}\|_{\|\cdot\|_{\mathbf{C}_\star^\top \mathbf{C}_\star} \rightarrow \|\cdot\|_{\mathbf{C}_\star^\top \mathbf{C}_\star}^*} = \max_{\mathbf{x}: \|\mathbf{x}\|_{\mathbf{C}_\star^\top \mathbf{C}_\star} \leq 1} \|\mathbf{A}\mathbf{x}\|_{\mathbf{C}_\star^\top \mathbf{C}_\star}^*.$$

Using these definitions we can further write

$$\begin{aligned} \|\mathbf{C}_\star^\top \mathbf{D} \mathbf{C}_\star\|_{\|\cdot\|_{\mathbf{C}_\star^\top \mathbf{C}_\star} \rightarrow \|\cdot\|_{\mathbf{C}_\star^\top \mathbf{C}_\star}^*} &= \max_{\mathbf{x}: \|\mathbf{x}\|_{\mathbf{C}_\star^\top \mathbf{C}_\star} \leq 1} \|\mathbf{C}_\star^\top \mathbf{D} \mathbf{C}_\star \mathbf{x}\|_{\mathbf{C}_\star^\top \mathbf{C}_\star}^* \\ &= \max_{\mathbf{x}: \|\mathbf{C}_\star \mathbf{x}\|_2 \leq 1} \max_{\mathbf{y}: \|\mathbf{C}_\star \mathbf{y}\|_2 \leq 1} \langle \mathbf{C}_\star^\top \mathbf{D} \mathbf{C}_\star \mathbf{x}, \mathbf{y} \rangle \\ &= \max_{\mathbf{f}, \mathbf{g}: \|\mathbf{f}\|_2 \leq 1, \|\mathbf{g}\|_2 \leq 1} \langle \mathbf{D} \mathbf{f}, \mathbf{g} \rangle \\ &\leq \max_i \mathbf{D}_{ii}. \end{aligned}$$

¹Alternatively, we can enforce lower and upper bounds on the entries of \mathbf{D} by adding an additional small quadratic regularizer to the edges in E .

Next we bound the distance $\|\mathbf{x} - \mathbf{x}^*\|_{\mathbf{C}_\star^\top \mathbf{C}_\star}$. To do so, we Taylor expand:

$$\begin{aligned} \Phi(\mathbf{x}) - \Phi(\mathbf{x}^*) &= \frac{1}{2}(\mathbf{x} - \mathbf{x}^*)^\top \left(\int_0^1 \nabla^2 \Phi((1-t)\mathbf{x}^* + t\mathbf{x}) dt \right) (\mathbf{x} - \mathbf{x}^*) \\ &= \frac{1}{2} \left(\left\| \int_0^1 \nabla^2 \Phi((1-t)\mathbf{x}^* + t\mathbf{x}) dt \right\|_{\|\cdot\|_{\mathbf{C}_\star^\top \mathbf{C}_\star} \rightarrow \|\cdot\|_{\mathbf{C}_\star^\top \mathbf{C}_\star}} \|\mathbf{x} - \mathbf{x}^*\|_{\mathbf{C}_\star^\top \mathbf{C}_\star} \right)^2 \\ &\geq \frac{1}{2} \left(\min_i \mathbf{D}_{ii} \cdot \|\mathbf{x} - \mathbf{x}^*\|_{\mathbf{C}_\star^\top \mathbf{C}_\star} \right)^2. \end{aligned}$$

Since the solution we obtain satisfies $\Phi(\mathbf{x}) - \Phi(\mathbf{x}^*) \leq \varepsilon$, we thus infer that

$$\|\mathbf{x} - \mathbf{x}^*\|_{\mathbf{C}_\star^\top \mathbf{C}_\star} \leq \frac{\sqrt{2\varepsilon}}{\min_i \mathbf{D}_{ii}}.$$

Plugging into the upper bound on the gradient norm, we obtain that:

$$\|\nabla \Phi(\mathbf{x})\|_{\mathbf{C}_\star^\top \mathbf{C}_\star}^* \leq \frac{\max_i \mathbf{D}_{ii}}{\min_i \mathbf{D}_{ii}} \cdot \sqrt{2\varepsilon}.$$

Since we have that $\nabla \Phi(\mathbf{x}) = \mathbf{C}_\star^\top \mathbf{h}$, this implies an upper bound on the norm of \mathbf{h} . Indeed we have that $\|\mathbf{C}_\star^\top \mathbf{h}\|_{\mathbf{C}_\star^\top \mathbf{C}_\star}^* = \max_{\mathbf{y}: \|\mathbf{C}_\star \mathbf{y}\| \leq 1} \langle \mathbf{C}_\star^\top \mathbf{h}, \mathbf{y} \rangle = \max_{\mathbf{y}: \|\mathbf{C}_\star \mathbf{y}\| \leq 1} \langle \mathbf{h}, \mathbf{C}_\star \mathbf{y} \rangle = \|\mathbf{h}\|_2$. Therefore the same upper bound also holds for $\|\mathbf{h}\|_2$. Hence under the assumption that all slacks are lower bounded by a small polynomial, we can obtain that $\|\mathbf{h}\|_2 \leq M^{-1}$ by setting ε to an appropriately small polynomial.

Therefore after (approximately) solving the regularized objective (3.18) we can set its gradient exactly to 0 only by slightly perturbing the linear term. This extra error, in turn, gets passed to Lemma 3.4.20 which removes the error by slightly modifying the weight vector \mathbf{w} . Although this operation may enable some weights to decrease by a tiny amount below 1, this can be prevented simply by uniformly upscaling all the weights. The centrality property will be then achieved for a slightly smaller parameter μ and the weight increase will be upper bounded by an arbitrary inverse polynomial in m thus negligible.

Finally, Lemma 9.1.6 ensures that as a matter of fact all slacks are always polynomially lower bounded, and so our claim holds. Furthermore, all calls to the solver in Theorem 9.1.5 are made on instances where all parameters are well-conditioned, as required.

Lemma 9.1.6 (Slack lower bound). *When invoking the mixed objective solver during the procedure described in Theorem 3.4.26, at all times we have that $\min \{s_e^+, s_e^-\} \geq 1/m^{O(1)}$ for all edges $e \in E$.*

Proof. We note that the stretch condition from Lemma 3.4.7 holds for any \mathbf{h} . We apply it once for $\mathbf{h} = \frac{\mathbf{w}^+}{s^+} - \frac{\mathbf{w}^-}{s^-}$ and once for $\mathbf{h} = \frac{\mathbf{c}}{\mu}$. For the former, since

$$\left\| \frac{\left| \frac{\mathbf{w}^+}{s^+} - \frac{\mathbf{w}^-}{s^-} \right|}{\sqrt{(\mathbf{w}^+ + \mathbf{w}^-) \left(\frac{\mathbf{w}^+}{(s^+)^2} + \frac{\mathbf{w}^-}{(s^-)^2} \right)}} \right\|_\infty \leq 1,$$

we get

$$\left\| \left(\frac{\mathbf{w}^+}{(s^+)^2} + \frac{\mathbf{w}^-}{(s^-)^2} + R_p \cdot \tilde{\mathbf{f}}^{p-2} \right) \tilde{\mathbf{f}} - \left(\frac{\mathbf{w}^+}{s^+} - \frac{\mathbf{w}^-}{s^-} \right) \right\|_\infty = O(\delta \|\mathbf{w}\|_1 \cdot \log \|\mathbf{w}\|_1) = O(m), \quad (9.15)$$

and for the latter, since $\left\| \frac{\frac{c}{\mu}}{\sqrt{(w^+ + w^-) \left(\frac{w^+}{(s^+)^2} + \frac{w^-}{(s^-)^2} \right)}} \right\|_\infty \leq \frac{W \cdot m^4}{2}$, we get

$$\left\| \left(\frac{\mathbf{w}^+}{(s^+)^2} + \frac{\mathbf{w}^-}{(s^-)^2} + R_p \cdot \tilde{\mathbf{f}}^{p-2} \right) \tilde{\mathbf{f}} - \frac{\mathbf{c}}{\mu} \right\|_\infty = O(\delta \|\mathbf{w}\|_1 \cdot W \cdot m^4 \cdot \log \|\mathbf{w}\|_1) = m^{O(1)}. \quad (9.16)$$

Combining (9.15) and (9.16) and using the triangle inequality we get that

$$\left\| \frac{\mathbf{w}^+}{s^+} - \frac{\mathbf{w}^-}{s^-} \right\|_\infty \leq \left\| \frac{\mathbf{c}}{\mu} \right\|_\infty + m^{O(1)} = m^{O(1)}.$$

Since $\max\{s_e^+, s_e^-\} \geq \frac{1}{2}$, using the triangle inequality and the fact that $1 \leq w_e^+, w_e^- \leq O(m)$ in the above inequality implies that $\min\{s_e^+, s_e^-\} \geq 1/m^{O(1)}$. \square

9.1.4 Strengthening the Mixed Objective Solver

In this section we prove that the ℓ_2 - ℓ_p solver from [103] can be extended to handle a broader class of optimization problems on graphs. This will be useful in order to solve the optimization problem required by the improved method we present in Section 3.6. We next state the main lemma we prove in this section.

Lemma 9.1.7. *Let a graph $G = (V, E)$ with m edges, and a family of functions $\{g_e\}_{e \in E}$, $g_e : \mathbb{R} \rightarrow \mathbb{R}$ such that each function satisfies $\left| \ln \frac{g_e''(x)}{g_e''(0)} \right| \leq \alpha$, for all $x \in \mathbb{R}$. Let the function defined over circulations \mathbf{f} in G :*

$$\tilde{\text{val}}(\mathbf{f}) = \sum_e g_e(f_e) + |f_e|^p.$$

We can compute a circulation $\tilde{\mathbf{f}}$ such that

$$\tilde{\text{val}}(\tilde{\mathbf{f}}) - OPT \leq \frac{1}{2^{(\log m)^{O(1)}}} (\tilde{\text{val}}(\mathbf{f}) - OPT) + \frac{1}{2^{(\log m)^{O(1)}}}$$

in $2^{O(\alpha + p^3/2)} m^{1+O(1/\sqrt{p})}$ time.

The proof closely follows the lines of the iterative refinement proofs from the original paper [1]. Intuitively, since g has bounded second order derivatives, it is always well approximated by a quadratic. Therefore the solver from Theorem 9.1.5 can be used iteratively to improve the error of the current solution.

Iterative refinement is based on the following basic statement.

Lemma 9.1.8. *Let a linear subspace $\mathcal{X} \subseteq \mathbb{R}^m$, let $h, k : \mathcal{X} \rightarrow \mathbb{R}$ be convex twice-differentiable functions, and let $\mathbf{x}, \mathbf{x}^* \in \mathbb{R}^m$ such that \mathbf{x}^* minimizes h . Suppose that k satisfies*

$$k(c\boldsymbol{\delta}) \leq c^2 k(\boldsymbol{\delta}),$$

for all $c \in \mathbb{R}$, $\boldsymbol{\delta} \in \mathcal{X}$, such that

$$k(\boldsymbol{\delta}) \leq h(\mathbf{x} + \boldsymbol{\delta}) - h(\mathbf{x}) - \langle \nabla h(\mathbf{x}), \boldsymbol{\delta} \rangle \leq \beta \cdot k(\boldsymbol{\delta}),$$

for any $\boldsymbol{\delta} \in \mathcal{X}$. Then letting

$$\boldsymbol{\delta}^* = \arg \min_{\boldsymbol{\delta} \in \mathcal{X}} \langle \nabla h(\mathbf{x}), \boldsymbol{\delta} \rangle + k(\boldsymbol{\delta})$$

and $\boldsymbol{\delta}^\sharp \in \mathcal{X}$ such that

$$\langle \nabla h(\mathbf{x}), \boldsymbol{\delta}^\sharp \rangle + k(\boldsymbol{\delta}^\sharp) \leq \frac{1}{\gamma} (\langle \nabla h(\mathbf{x}), \boldsymbol{\delta}^* \rangle + k(\boldsymbol{\delta}^*)) + \varepsilon,$$

one has that

$$h(\mathbf{x}) - h(\mathbf{x} + \boldsymbol{\delta}^\sharp) \geq \frac{1}{\beta\gamma} (h(\mathbf{x}) - h(\mathbf{x}^*)) - \frac{\varepsilon}{\beta}.$$

Proof. Using the hypothesis, we have that

$$\begin{aligned} h(\mathbf{x} + \boldsymbol{\delta}^\sharp/\beta) &\leq h(\mathbf{x}) + \langle \nabla h(\mathbf{x}), \boldsymbol{\delta}^\sharp/\beta \rangle + \beta \cdot k(\boldsymbol{\delta}^\sharp/\beta) \\ &\leq h(\mathbf{x}) + \frac{1}{\beta} \left(\langle \nabla h(\mathbf{x}), \boldsymbol{\delta}^\sharp \rangle + k(\boldsymbol{\delta}^\sharp) \right), \end{aligned}$$

where we used the right hand side of the sandwiching inequality from the hypothesis, and the fact that $k(\boldsymbol{\delta}^\sharp/\beta) \leq k(\boldsymbol{\delta}^\sharp)/\beta^2$.

Next we plug in the relation between $\boldsymbol{\delta}^*$ and $\boldsymbol{\delta}^\sharp$ to obtain:

$$\begin{aligned} h(\mathbf{x} + \boldsymbol{\delta}^\sharp/\beta) &\leq h(\mathbf{x}) + \frac{1}{\beta\gamma} (\langle \nabla h(\mathbf{x}), \boldsymbol{\delta}^* \rangle + k(\boldsymbol{\delta}^*)) + \frac{\varepsilon}{\beta} \\ &\leq h(\mathbf{x}) + \frac{1}{\beta\gamma} (\langle \nabla h(\mathbf{x}), \mathbf{x}^* - \mathbf{x} \rangle + k(\mathbf{x}^* - \mathbf{x})) + \frac{\varepsilon}{\beta}, \end{aligned}$$

where the latter inequality follows from the fact that $\boldsymbol{\delta}^*$ minimizes $\langle \nabla h(\mathbf{x}), \boldsymbol{\delta}^* \rangle + k(\boldsymbol{\delta}^*)$, so plugging in $\mathbf{x}^* - \mathbf{x}$ as an argument can only increase the sum of these two terms. Finally, we plug in the left hand side of the sandwiching inequality, i.e.

$$\langle \nabla h(\mathbf{x}), \mathbf{x}^* - \mathbf{x} \rangle + k(\mathbf{x}^* - \mathbf{x}) \leq h(\mathbf{x}^*) - h(\mathbf{x}),$$

which combined with the previous inequality yields,

$$h(\mathbf{x}) - h(\mathbf{x} + \boldsymbol{\delta}^\sharp/\beta) \geq \frac{1}{\beta\gamma} (h(\mathbf{x}) - h(\mathbf{x}^*)) - \frac{\varepsilon}{\beta},$$

which is what we needed. \square

We apply Lemma 9.1.8 for a custom choice of k , which is dictated by the specific family of functions $\{g_e\}_{e \in E}$.

In order to do so, we also require a sandwiching inequality for the $\sum f_e^p$ term. We use the following inequality from [103].

Lemma 9.1.9 ([103], p. 11). *Let $x, \delta \in \mathbb{R}$ and $p \geq 2$. Then*

$$2^{-O(p)}(x^{p-2}\delta^2 + \delta^p) \leq (x + \delta)^p - x^p - px^{p-1} \leq 2^{O(p)}(x^{p-2}\delta^2 + \delta^p).$$

Using Lemmas 9.1.8 and 9.1.9, we can now define an appropriate function k for each circulation \mathbf{f} . In order to do so, we use the following simple lemma.

Lemma 9.1.10. *Let $\widetilde{val}(\mathbf{f})$ specified as in Lemma 9.1.7, and let a circulation \mathbf{f} . Then one has that for any circulation $\boldsymbol{\delta}$:*

$$\begin{aligned} \sum_e \left(\frac{e^{-\alpha} g_e''(0)}{2} \delta_e^2 + 2^{-O(p)} (f_e^{p-2} \delta_e^2 + \delta_e^p) \right) &\leq \widetilde{val}(\mathbf{f} + \boldsymbol{\delta}) - \widetilde{val}(\mathbf{f}) - \langle \nabla \widetilde{val}(\mathbf{f}), \boldsymbol{\delta} \rangle \\ &\leq \sum_e \left(\frac{e^\alpha g_e''(0)}{2} \delta_e^2 + 2^{O(p)} (f_e^{p-2} \delta_e^2 + \delta_e^p) \right). \end{aligned}$$

Proof. We require lower bounding and upper bounding the terms of order higher than 1, after expanding $\widetilde{val}(\mathbf{f} + \boldsymbol{\delta})$ around $\widetilde{val}(\mathbf{f})$. To do so, we first notice that by the hypothesis in Lemma 9.1.7 we have

$$\frac{e^{-\alpha} g_e''(0)}{2} \delta^2 \leq g_e(x + \delta) - g_e(x) - g_e'(x) \delta \leq \frac{e^\alpha g_e''(0)}{2} \delta^2.$$

Similarly, we use Lemma 9.1.9 to lower and upper bound the higher order terms of f_e^p for each e . Combining, we obtain the desired claim. \square

Now we can prove that we can decrease $\widetilde{val}(\mathbf{f}) - OPT$ very fast, which in turn enables us to prove the main lemma in this section.

Proof of Lemma 9.1.7. We use Lemma 9.1.8 where we define the functions k_e based on Lemma 9.1.10. More precisely, for each $e \in E$ we let

$$k(\boldsymbol{\delta}) = \sum_e k_e(\delta_e),$$

where

$$k_e(\delta_e) = \frac{e^{-\alpha} g_e''(0)}{2} \delta_e^2 + 2^{-O(p)} (f_e^{p-2} \delta_e^2 + \delta_e^p).$$

Using Lemma 9.1.10 we verify that

$$k(\boldsymbol{\delta}) \leq \widetilde{val}(\mathbf{f} + \boldsymbol{\delta}) - \widetilde{val}(\mathbf{f}) - \langle \nabla \widetilde{val}(\mathbf{f}), \boldsymbol{\delta} \rangle \leq \max\{e^{2\alpha}, 2^{O(p)}\} \cdot k(\boldsymbol{\delta}).$$

We use the solver from Theorem 9.1.5 to approximately minimize $k(\boldsymbol{\delta})$ plus the corresponding linear term. Then, for our specific setting, we can apply Lemma 9.1.8 with

$$\begin{aligned} \gamma &= 1 + \frac{1}{2^{(\log m)^{O(1)}}}, \\ \varepsilon &= \frac{1}{2^{(\log m)^{O(1)}}}, \\ \beta &= \max\{e^{2\alpha}, 2^{O(p)}\}, \end{aligned}$$

to get that the newly obtained iterate \mathbf{f}' satisfies

$$\widetilde{val}(\mathbf{f}) - \widetilde{val}(\mathbf{f}') \geq \frac{1}{\beta\gamma} (\widetilde{val}(\mathbf{f}) - OPT) - \frac{\varepsilon}{\beta}.$$

Therefore executing this step for T iterations, we obtain \mathbf{f}_T such that

$$\begin{aligned}\widetilde{\text{val}}(\mathbf{f}_T) - OPT &\leq \left(1 - \frac{1}{\beta\gamma}\right)^T (\widetilde{\text{val}}(\mathbf{f}) - OPT) + \frac{\varepsilon}{\beta} \left(\sum_{t=0}^{T-1} \left(1 - \frac{1}{\beta\gamma}\right)^t\right) \\ &\leq \left(1 - \frac{1}{\beta\gamma}\right)^T (\widetilde{\text{val}}(\mathbf{f}) - OPT) + \varepsilon\gamma.\end{aligned}$$

Hence setting $T = \beta\gamma(\log m)^{O(1)}$ we make

$$\widetilde{\text{val}}(\mathbf{f}_T) - OPT \leq \frac{1}{2^{(\log m)^{O(1)}}} (\widetilde{\text{val}}(\mathbf{f}) - OPT) + \frac{1}{2^{(\log m)^{O(1)}}}.$$

Hence we require $T = (e^{2\alpha} + 2^{O(p)}) \log^{O(1)} m$ iterations to obtain the target accuracy. Together with the running time guarantee from Theorem 9.1.5 we obtain the claim. \square

9.1.5 Deferred Proofs

Proof of Lemma 3.3.1

Proof. Writing first order optimality conditions for (3.3), we have

$$\nabla F_\mu^w(\mathbf{x}) = \frac{\mathbf{C}^\top \mathbf{c}}{\mu} + \mathbf{C}^\top \left(\frac{\mathbf{w}^+}{\mathbf{s}^+} - \frac{\mathbf{w}^-}{\mathbf{s}^-} \right).$$

Setting this to $\mathbf{0}$ yields the conclusion. Now letting $\mathbf{y} = \mu \cdot \mathbf{w}/\mathbf{s}$, we see that $\mathbf{y} \geq \mathbf{0}$ since both weights \mathbf{w} and slacks \mathbf{s} are non-negative, and $\mathbf{C}^\top (\mathbf{y}^+ - \mathbf{y}^-) = -\mathbf{C}^\top \mathbf{c}$, from which we conclude that \mathbf{y} is a feasible dual vector. Finally, we can write the duality gap as

$$\begin{aligned}\langle \mathbf{c}, \mathbf{C}\mathbf{x} \rangle + \langle \mathbf{1} - \mathbf{f}_0, \mathbf{y}^+ \rangle + \langle \mathbf{f}_0, \mathbf{y}^- \rangle &= -\langle \mathbf{C}^\top (\mathbf{y}^+ - \mathbf{y}^-), \mathbf{x} \rangle + \langle \mathbf{1} - \mathbf{f}_0, \mathbf{y}^+ \rangle + \langle \mathbf{f}_0, \mathbf{y}^- \rangle \\ &= \langle \mathbf{1} - \mathbf{f}_0 - \mathbf{C}\mathbf{x}, \mathbf{y}^+ \rangle + \langle \mathbf{f}_0 + \mathbf{C}\mathbf{x}, \mathbf{y}^- \rangle = \langle \mathbf{s}^+, \mu \cdot \mathbf{w}^+ / \mathbf{s}^+ \rangle + \langle \mathbf{s}^-, \mu \cdot \mathbf{w}^- / \mathbf{s}^- \rangle \\ &= \mu \|\mathbf{w}\|_1.\end{aligned}$$

\square

Proof of Lemma 3.3.4

Proof. We write

$$\begin{aligned}\mathcal{E}_{w,s}(\mathbf{h}) &= \min_{\tilde{\mathbf{y}}: \mathbf{C}^\top(\tilde{\mathbf{y}}+\mathbf{h})=\mathbf{0}} \frac{1}{2} \sum_{e \in E} \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right)^{-1} (\tilde{y}_e)^2 \\ &= \min_{\tilde{\mathbf{y}}} \max_{\tilde{\mathbf{x}}} \frac{1}{2} \sum_{e \in E} \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right)^{-1} (\tilde{y}_e)^2 - \langle \tilde{\mathbf{x}}, \mathbf{C}^\top(\tilde{\mathbf{y}}+\mathbf{h}) \rangle \\ &= \max_{\tilde{\mathbf{x}}} \min_{\tilde{\mathbf{y}}} \frac{1}{2} \sum_{e \in E} \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right)^{-1} (\tilde{y}_e)^2 - \langle \mathbf{C}\tilde{\mathbf{x}}, \tilde{\mathbf{y}}+\mathbf{h} \rangle\end{aligned}$$

The solution to the inner optimization problem is $\tilde{\mathbf{y}} = \left(\frac{\mathbf{w}^+}{(s^+)^2} + \frac{\mathbf{w}^-}{(s^-)^2} \right) \mathbf{C}\tilde{\mathbf{x}}$, therefore by substituting we get

$$\begin{aligned} \mathcal{E}_{\mathbf{w},s}(\mathbf{h}) &= \max_{\tilde{\mathbf{x}}} \langle \mathbf{h}, \mathbf{C}\tilde{\mathbf{x}} \rangle - \frac{1}{2} \sum_{e \in E} \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) (\mathbf{C}\tilde{\mathbf{x}})_e^2 \\ &= \max_{\tilde{\mathbf{f}} = \mathbf{C}\tilde{\mathbf{x}}} \langle \mathbf{h}, \tilde{\mathbf{f}} \rangle - \frac{1}{2} \sum_{e \in E} \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) (\tilde{f}_e)^2, \end{aligned}$$

proving (3.9). Now the first order optimality condition of this problem is given by

$$\mathbf{C}^\top \left(\frac{\mathbf{w}^+}{(s^+)^2} + \frac{\mathbf{w}^-}{(s^-)^2} \right) \cdot (\mathbf{C}\tilde{\mathbf{x}}) = \mathbf{C}^\top \mathbf{h}$$

or equivalently, by setting $\tilde{\mathbf{f}} = \mathbf{C}\tilde{\mathbf{x}}$,

$$\mathbf{C}^\top \left(\frac{\mathbf{w}^+}{(s^+)^2} + \frac{\mathbf{w}^-}{(s^-)^2} \right) \cdot \tilde{\mathbf{f}} = \mathbf{C}^\top \mathbf{h}.$$

If $\tilde{\mathbf{x}}, \tilde{\mathbf{f}}$ are solutions to the above linear system, this implies

$$\begin{aligned} \mathcal{E}_{\mathbf{w},s}(\mathbf{h}) &= \langle \mathbf{h}, \mathbf{C}\tilde{\mathbf{x}} \rangle - \frac{1}{2} \sum_{e \in E} \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) (\mathbf{C}\tilde{\mathbf{x}})_e^2 \\ &= \langle \mathbf{C}^\top \mathbf{h}, \tilde{\mathbf{x}} \rangle - \frac{1}{2} \sum_{e \in E} \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) (\mathbf{C}\tilde{\mathbf{x}})_e^2 \\ &= \frac{1}{2} \sum_{e \in E} \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) (\mathbf{C}\tilde{\mathbf{x}})_e^2 \\ &= \frac{1}{2} \sum_{e \in E} \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) (\tilde{f}_e)^2. \end{aligned}$$

The equations in terms of $\boldsymbol{\rho}$ follow by simple substitution. □

Proof of Lemma 3.3.6

Proof. We explicitly write the new residual after performing the update. We have:

$$\begin{aligned} -\mathbf{C}^\top \mathbf{h}' &:= \nabla F_\mu^{\mathbf{w}}(\mathbf{x}') = \mathbf{C}^\top \left(\frac{\mathbf{w}^+}{(s^+)' } - \frac{\mathbf{w}^-}{(s^-)' } + \frac{\mathbf{c}}{\mu} \right) \\ &= \mathbf{C}^\top \left(\frac{\mathbf{w}^+}{s^+} - \frac{\mathbf{w}^-}{s^-} + \frac{\mathbf{c}}{\mu} \right) + \mathbf{C}^\top \left(\frac{\mathbf{w}^+}{(s^+)' } - \frac{\mathbf{w}^+}{s^+} - \frac{\mathbf{w}^-}{(s^-)' } + \frac{\mathbf{w}^-}{s^-} \right) \\ &\stackrel{(1)}{=} -\mathbf{C}^\top \left(\frac{\mathbf{w}^+ \boldsymbol{\rho}^+}{s^+} - \frac{\mathbf{w}^- \boldsymbol{\rho}^-}{s^-} \right) + \mathbf{C}^\top \left(\frac{\mathbf{w}^+ \boldsymbol{\rho}^+}{s^+(1-\boldsymbol{\rho}^+)} - \frac{\mathbf{w}^- \boldsymbol{\rho}^-}{s^-(1-\boldsymbol{\rho}^-)} \right) \\ &= \mathbf{C}^\top \left(\frac{\mathbf{w}^+ (\boldsymbol{\rho}^+)^2}{s^+(1-\boldsymbol{\rho}^+)} - \frac{\mathbf{w}^- (\boldsymbol{\rho}^-)^2}{s^-(1-\boldsymbol{\rho}^-)} \right) \\ &= \mathbf{C}^\top \left(\frac{\mathbf{w}^+ (\boldsymbol{\rho}^+)^2}{(s^+)' } - \frac{\mathbf{w}^- (\boldsymbol{\rho}^-)^2}{(s^-)' } \right), \end{aligned}$$

where (1) follows from Equations (3.11) and (3.12) and from writing

$$\begin{aligned} (\mathbf{s}^+)' &= \mathbf{s}^+ - \mathbf{C}\tilde{\mathbf{x}} = \mathbf{s}^+ \left(\mathbf{1} - \frac{\mathbf{C}\tilde{\mathbf{x}}}{\mathbf{s}^+} \right) = \mathbf{s} (\mathbf{1} - \boldsymbol{\rho}^+) \\ (\mathbf{s}^-)' &= \mathbf{s}^- + \mathbf{C}\tilde{\mathbf{x}} = \mathbf{s}^- \left(\mathbf{1} + \frac{\mathbf{C}\tilde{\mathbf{x}}}{\mathbf{s}^-} \right) = \mathbf{s} (\mathbf{1} - \boldsymbol{\rho}^-) . \end{aligned}$$

Now we can upper bound the energy required to route the new residual with resistances determined by $(\mathbf{w}, \mathbf{s}')$, by substituting

$$\tilde{\mathbf{y}} = \frac{\mathbf{w}^+(\boldsymbol{\rho}^+)^2}{(\mathbf{s}^+)' } - \frac{\mathbf{w}^-(\boldsymbol{\rho}^-)^2}{(\mathbf{s}^-)' }$$

into Definition 3.3.3, after noting that

$$\mathbf{C}^\top \tilde{\mathbf{y}} = \mathbf{C}^\top \left(\frac{\mathbf{w}^+(\boldsymbol{\rho}^+)^2}{(\mathbf{s}^+)' } - \frac{\mathbf{w}^-(\boldsymbol{\rho}^-)^2}{(\mathbf{s}^-)' } \right) = -\mathbf{C}^\top \mathbf{h}' .$$

We thus obtain

$$\begin{aligned} \mathcal{E}_{\mathbf{w}, \mathbf{s}'}(\mathbf{h}') &\leq \frac{1}{2} \sum_{e \in E} \frac{\left(\frac{w_e^+(\rho_e^+)^2}{(s_e^+)' } - \frac{w_e^-(\rho_e^-)^2}{(s_e^-)' } \right)^2}{\frac{w_e^+}{(s_e^+)'2} + \frac{w_e^-}{(s_e^-)'2}} \leq \frac{1}{2} \sum_{e \in E} \frac{\left(\frac{w_e^+(\rho_e^+)^2}{(s_e^+)' } \right)^2 + \left(\frac{w_e^-(\rho_e^-)^2}{(s_e^-)' } \right)^2}{\frac{w_e^+}{(s_e^+)'2} + \frac{w_e^-}{(s_e^-)'2}} \\ &\leq \frac{1}{2} \sum_{e \in E} \frac{\left(\frac{w_e^+(\rho_e^+)^2}{(s_e^+)' } \right)^2}{\frac{w_e^+}{(s_e^+)'2}} + \frac{\left(\frac{w_e^-(\rho_e^-)^2}{(s_e^-)' } \right)^2}{\frac{w_e^-}{(s_e^-)'2}} = \frac{1}{2} \sum_{e \in E} (w_e^+(\rho_e^+)^4 + w_e^-(\rho_e^-)^4) . \end{aligned}$$

□

Proof of Corollary 3.3.7

Proof. From Lemma 3.3.4 we have that

$$\mathcal{E}_{\mathbf{w}, \mathbf{s}}(\mathbf{h}) = \frac{1}{2} \sum_{e \in E} (w_e^+(\rho_e^+)^2 + w_e^-(\rho_e^-)^2) .$$

Since $\mathbf{w} \geq 1$ we have that

$$\|\boldsymbol{\rho}\|_\infty^2 \leq 2 \cdot \mathcal{E}_{\mathbf{w}, \mathbf{s}}(\mathbf{h}) .$$

Using Lemma 3.3.6 we upper bound

$$\mathcal{E}_{\mathbf{w}, \mathbf{s}'}(\mathbf{h}') \leq \frac{1}{2} \sum_{e \in E} (w_e^+(\rho_e^+)^4 + w_e^-(\rho_e^-)^4) \leq \frac{1}{2} \|\boldsymbol{\rho}\|_\infty^2 \cdot \sum_{e \in E} (w_e^+(\rho_e^+)^2 + w_e^-(\rho_e^-)^2) \leq 2 \cdot \mathcal{E}_{\mathbf{w}, \mathbf{s}}(\mathbf{h})^2 .$$

□

Proof Lemma 3.3.8

Proof. By Lemma 3.3.4 we know that there exists a vector $\boldsymbol{\rho}$ such that

$$\mathbf{C}^\top \mathbf{h} = \mathbf{C}^\top \left(\frac{\mathbf{w}^+ \boldsymbol{\rho}^+}{\mathbf{s}^+} - \frac{\mathbf{w}^- \boldsymbol{\rho}^-}{\mathbf{s}^-} \right)$$

and

$$\mathcal{E}_{\mathbf{w},s}(\mathbf{h}) = \frac{1}{2} \sum_{e \in E} (w_e^+(\rho_e^+)^2 + w_e^-(\rho_e^-)^2) \leq \varepsilon.$$

Therefore $\|\boldsymbol{\rho}\|_\infty \leq \sqrt{2\varepsilon}$. Now let $\mathbf{w}' = \frac{\mathbf{w}(1+\boldsymbol{\rho})}{1-\|\boldsymbol{\rho}\|_\infty}$ and $\mu' = \frac{\mu}{1-\|\boldsymbol{\rho}\|_\infty}$. We have that

$$\begin{aligned} \nabla F_{\mu'}^{\mathbf{w}'}(\mathbf{x}) &= \mathbf{C}^\top \left(\frac{(\mathbf{w}^+)'}{s^+} - \frac{(\mathbf{w}^-)'}{s^-} - \frac{\mathbf{c}}{\mu'} \right) \\ &= \frac{1}{1-\|\boldsymbol{\rho}\|_\infty} \cdot \left(\mathbf{C}^\top \left(\frac{\mathbf{w}^+}{s^+} - \frac{\mathbf{w}^-}{s^-} - \frac{\mathbf{c}}{\mu} \right) + \mathbf{C}^\top \left(\frac{\mathbf{w}^+ \boldsymbol{\rho}^+}{s^+} - \frac{\mathbf{w}^- \boldsymbol{\rho}^-}{s^-} \right) \right) \\ &= \frac{1}{1-\|\boldsymbol{\rho}\|_\infty} \cdot \left(-\mathbf{C}^\top \mathbf{h} + \mathbf{C}^\top \mathbf{h} \right) \\ &= \mathbf{0}. \end{aligned}$$

Furthermore, by this construction we see that $\mathbf{w} \leq \mathbf{w}'$, and that

$$\mathbf{w}' \leq \mathbf{w} \cdot \frac{1 + \|\boldsymbol{\rho}\|_\infty}{1 - \|\boldsymbol{\rho}\|_\infty} \leq \mathbf{w} \cdot \frac{1 + \sqrt{2\varepsilon}}{1 - \sqrt{2\varepsilon}} \leq \mathbf{w} \cdot (1 + 4\sqrt{\varepsilon}),$$

whenever $\varepsilon \leq 1/100$. Therefore the total increase in weight is at most $4\sqrt{\varepsilon}\|\mathbf{w}\|_1$. Furthermore the loss of duality gap is determined by $\mu' \leq \mu(1 + 2\sqrt{\varepsilon})$. \square

Proof of Lemma 3.3.9

Proof. We first perform $O(\log \log \|\mathbf{w}\|_1)$ vanilla residual correction steps as described in Definition 3.3.5 to obtain a new solution $\mathbf{f}' = \mathbf{f}_0 + \mathbf{C}\mathbf{x}'$ with residual $\nabla F_\mu^{\mathbf{w}}(\mathbf{x}') = -\mathbf{C}^\top \mathbf{g}'$ and low energy $\mathcal{E}_{\mathbf{w},s'}(\mathbf{g}') \leq \|\mathbf{w}\|_1^{-22}/16$. Then we apply the perfect correction step from Lemma 3.3.8 to eliminate the residual $-\mathbf{C}^\top \mathbf{g}'$ by obtaining a new set of weights $\mathbf{w}' \geq \mathbf{w}$ such that $\|\mathbf{w}' - \mathbf{w}\|_1 \leq \|\mathbf{w}\|^{-10}$ and $\nabla F_{\mu'}^{\mathbf{w}'}(\mathbf{x}') = \mathbf{0}$ where $\mu' \leq \mu(1 + \frac{1}{2}\|\mathbf{w}\|_1^{-11})$. \square

Proof of Lemma 3.3.10

Proof. Since \mathbf{f} is μ -central, we can write:

$$\begin{aligned} \nabla F_\mu^{\mathbf{w}}(\mathbf{x}) &= \mathbf{C}^\top \left(\frac{\mathbf{w}^+}{s^+} - \frac{\mathbf{w}^-}{s^-} + \frac{\mathbf{c}}{\mu} \right) = (1 + \delta) \mathbf{C}^\top \left(\frac{\mathbf{w}^+}{s^+} - \frac{\mathbf{w}^-}{s^-} + \frac{\mathbf{c}}{\mu} \right) - \delta \mathbf{C}^\top \left(\frac{\mathbf{w}^+}{s^+} - \frac{\mathbf{w}^-}{s^-} \right) \\ &= -\delta \mathbf{C}^\top \left(\frac{\mathbf{w}^+}{s^+} - \frac{\mathbf{w}^-}{s^-} \right), \end{aligned}$$

and so we can set $\mathbf{h}' = \delta \left(\frac{\mathbf{w}^+}{s^+} - \frac{\mathbf{w}^-}{s^-} \right)$. Using Definition 3.3.3 we can upper bound the energy required to route this residual by exhibiting the solution $\tilde{\mathbf{y}} = -\delta \left(\frac{\mathbf{w}^+}{s^+} - \frac{\mathbf{w}^-}{s^-} \right)$ after noting that

$$\mathbf{C}^\top (\tilde{\mathbf{y}} + \mathbf{h}') = \mathbf{0}:$$

$$\begin{aligned} \mathcal{E}_{\mathbf{w},s}(\mathbf{h}') &\leq \frac{1}{2} \delta^2 \sum_{e \in E} \frac{\left(\frac{w_e^+}{s_e^+} - \frac{w_e^-}{s_e^-} \right)^2}{\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2}} \leq \frac{1}{2} \delta^2 \sum_{e \in E} \frac{\left(\frac{w_e^+}{s_e^+} \right)^2 + \left(\frac{w_e^-}{s_e^-} \right)^2}{\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2}} \leq \frac{1}{2} \delta^2 \left(\sum_{e \in E} \frac{\left(\frac{w_e^+}{s_e^+} \right)^2}{\frac{w_e^+}{(s_e^+)^2}} + \sum_{e \in E} \frac{\left(\frac{w_e^-}{s_e^-} \right)^2}{\frac{w_e^-}{(s_e^-)^2}} \right) \\ &= \frac{1}{2} \delta^2 \|\mathbf{w}\|_1, \end{aligned}$$

which is at most $1/4$ as long as $\delta \leq \frac{1}{(2\|\mathbf{w}\|_1)^{1/2}}$. \square

Proof of Lemma 3.3.11

Proof. Given a μ -central flow and using Lemma 3.3.10 we see that setting $\delta = \frac{1}{(2\|\mathbf{w}\|_1)^{1/2}}$ and $\mu' = \mu(1 - \delta)$ we obtain $\mathcal{E}_{\mathbf{w},s}(\mathbf{h}) \leq 1/4$, where $-\mathbf{C}^\top \mathbf{h} = \nabla F_{\mu'}^{\mathbf{w}}(\mathbf{x})$. Hence applying Corollary 3.3.7 for $O(\log \log m)$ iterations we obtain a new flow $\mathbf{f}' = \mathbf{f}_0 + \mathbf{C} \mathbf{x}'$ with slacks \mathbf{s}' and residual $-\mathbf{C}^\top \mathbf{h}' = \nabla F_{\mu'}^{\mathbf{w}}(\mathbf{x}')$ such that $\mathcal{E}_{\mathbf{w},s'}(\mathbf{h}') \leq m^{-20}/4$.

Finally, applying the perfect correction step from Lemma 3.3.8 we obtain a new set of weights $\mathbf{w}' \geq \mathbf{w}$, such that $\mathbf{w}' \leq \mathbf{w}(1 + m^{-10})$ and $\nabla F_{\mu''}^{\mathbf{w}'}(\mathbf{x}') = \mathbf{0}$ for $\mu'' \leq \mu'(1 + m^{-10})$. In other words, \mathbf{f}' is μ'' -central with respect to \mathbf{w}' .

Since the increase in weights is very small, iterating this procedure for $O(m^{1/2} \log m)$ steps maintains the invariant that $\|\mathbf{w}\|_1 \leq 2m + 1$. Furthermore, in each iteration the parameter μ gets scaled down by a factor of $1 + \frac{1}{(4\|\mathbf{w}\|_1)^{1/2}} \geq 1 + \frac{1}{3m^{1/2}}$, after which it gets slightly scaled up by at most $1 + m^{-10}$ due to the perfect correction step. Hence within $O(m^{1/2} \log \frac{m\mu^0}{\varepsilon}) = O(m^{1/2} \log m)$ iterations, the parameter μ gets scaled down by a factor of $\Omega(m\mu^0/\varepsilon)$, which thus implies that the final duality gap will be $O\left(\frac{3m \cdot \mu^0}{m\mu^0/\varepsilon}\right) = O(\varepsilon)$. \square

Proof of Lemma 3.4.5

Proof. For the first inequality, we use the test vector $\tilde{\mathbf{y}} = -\mathbf{h}$ into Definition 3.3.3. For the second one, we have

$$\begin{aligned} \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}) &= \frac{1}{2} \sum_{e \in E} \frac{\delta^2 \left(\frac{w_e^+}{s_e^+} - \frac{w_e^-}{s_e^-} \right)^2}{\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2}} \leq \frac{1}{2} \delta^2 \sum_{e \in E} \frac{\left(\frac{w_e^+}{s_e^+} \right)^2 + \left(\frac{w_e^-}{s_e^-} \right)^2}{\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2}} \\ &\leq \frac{1}{2} \delta^2 \sum_{e \in E} \left(\frac{\left(\frac{w_e^+}{s_e^+} \right)^2}{\frac{w_e^+}{(s_e^+)^2}} + \frac{\left(\frac{w_e^-}{s_e^-} \right)^2}{\frac{w_e^-}{(s_e^-)^2}} \right) = \frac{1}{2} \delta^2 \sum_{e \in E} (w_e^+ + w_e^-) \\ &= \frac{1}{2} \delta^2 \|\mathbf{w}\|_1. \end{aligned}$$

\square

Proof of Corollary 3.4.10

Proof. We can upper bound using Lemma 3.4.6:

$$\begin{aligned} \|\tilde{\mathbf{f}}_\star\|_p &\leq \left(\frac{p \cdot \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s})}{R_p} \right)^{1/p} \leq \left(\frac{p \cdot \frac{1}{2} \delta^2 \|\mathbf{w}\|_1}{p \cdot (10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \cdot \log \|\mathbf{w}\|_1)^{p+1}} \right)^{1/p} \\ &\leq \frac{1}{10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \cdot \log \|\mathbf{w}\|_1}. \end{aligned} \quad (9.17)$$

□

Proof of Corollary 3.4.11

Proof. We have

$$\hat{\gamma} = \left(R_\star + R_p \cdot \|\tilde{\mathbf{f}}_\star\|_\infty^{p-2} \right)^{1/2} \cdot \left\| \frac{\mathbf{h}}{\sqrt{(\mathbf{w}^+ + \mathbf{w}^-) \left(\frac{\mathbf{w}^+}{(s^+)^2} + \frac{\mathbf{w}^-}{(s^-)^2} \right)}} \right\|_\infty \cdot 32 \log \|\mathbf{w}\|_1.$$

In particular, for sufficiently large m and since $\|\tilde{\mathbf{f}}_\star\|_\infty \leq \|\tilde{\mathbf{f}}_\star\|_p$ we have

$$\begin{aligned} R_p \|\tilde{\mathbf{f}}_\star\|_\infty^{p-2} &\leq \frac{p (10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \cdot \log \|\mathbf{w}\|_1)^{p+1}}{(10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \cdot \log \|\mathbf{w}\|_1)^{p-2}} \leq p \cdot (10^6 \cdot \delta^2 \|\mathbf{w}\|_1 \cdot \log \|\mathbf{w}\|_1)^3 \\ &= (\delta^2 \|\mathbf{w}\|_1)^3 \cdot p \cdot (10^6 \cdot \log \|\mathbf{w}\|_1)^3 < 3\delta^2 \|\mathbf{w}\|_1^2 = R_\star, \end{aligned} \quad (9.18)$$

where the last inequality follows from the assumption on δ . Furthermore, we bound the term under the ℓ_∞ norm as:

$$\begin{aligned} \frac{|h_e|}{\sqrt{(w_e^+ + w_e^-) \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right)}} &= \delta \frac{\left| \frac{w_e^+}{s_e^+} - \frac{w_e^-}{s_e^-} \right|}{\sqrt{(w_e^+ + w_e^-) \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right)}} \\ &= \delta \left(\frac{\left(\frac{w_e^+}{s_e^+} - \frac{w_e^-}{s_e^-} \right)^2}{(w_e^+ + w_e^-) \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right)} \right)^{1/2} \\ &\leq \delta \left(\frac{1}{w_e^+ + w_e^-} \left(\frac{(w_e^+/s_e^+)^2}{w_e^+/(s_e^+)^2} + \frac{(w_e^-/s_e^-)^2}{w_e^-/(s_e^-)^2} \right) \right)^{1/2} \\ &= \delta. \end{aligned} \quad (9.19)$$

Combining (9.18) and (9.19), we get the upper bound

$$\hat{\gamma} \leq (2 \cdot R_\star)^{1/2} \cdot \delta \cdot 32 \log \|\mathbf{w}\|_1 = \delta^2 \|\mathbf{w}\|_1 \cdot 32\sqrt{6} \cdot \log \|\mathbf{w}\|_1 = \gamma.$$

□

Proof of Lemma 3.4.15

Proof. We restate (3.45) in terms of $\gamma \geq \hat{\gamma}$:

$$\left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) \cdot |\tilde{f}_e| \leq |h_e| + \gamma.$$

More specifically we will use the following, which the above implies for the setting of $\mathbf{h} = \delta \left(\frac{w^+}{s^+} - \frac{w^-}{s^-} \right)$:

$$\left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) \cdot |\tilde{f}_e| \leq \frac{w_e^+ \delta}{s_e^+} + \frac{w_e^- \delta}{s_e^-} + \gamma$$

or equivalently since ρ_e^+ and ρ_e^- have opposite signs:

$$\left| \frac{w_e^+ \rho_e^+}{s_e^+} \right| + \left| \frac{w_e^- \rho_e^-}{s_e^-} \right| \leq \frac{w_e^+ \delta}{s_e^+} + \frac{w_e^- \delta}{s_e^-} + \gamma. \quad (9.20)$$

Assume without loss of generality that $s_e^+ \leq s_e^-$, and so $|\rho_e^+| \geq |\rho_e^-|$. Now, for the sake of contradiction we suppose that $|\rho_e^+| \geq C_\infty$ and $\left| \frac{w_e^+ \rho_e^+}{s_e^+} \right| + \left| \frac{w_e^- \rho_e^-}{s_e^-} \right| > 6\gamma$. We consider the two cases of Definition 3.4.13:

(1) $\max\{w_e^+, w_e^-\} \leq \delta \|\mathbf{w}\|_1$: Since $0 < s_e^+ \leq s_e^- < 1$ and $s_e^+ + s_e^- = 1$, we have that $s_e^- \geq \frac{1}{2}$. Therefore $\frac{w_e^- \delta}{s_e^-} \leq 2 \cdot \delta^2 \|\mathbf{w}\|_1 \leq 2 \cdot \gamma$. We conclude that

$$\left| \frac{w_e^+ \rho_e^+}{s_e^+} \right| + \left| \frac{w_e^- \rho_e^-}{s_e^-} \right| \leq \frac{w_e^+ \delta}{s_e^+} + \frac{w_e^- \delta}{s_e^-} + \gamma \leq \frac{1}{2} \left| \frac{w_e^+ \rho_e^+}{s_e^+} \right| + 3\gamma,$$

since $|\rho_e^+| \geq C_\infty = \frac{1}{2\delta\sqrt{2}\|\mathbf{w}\|_1} \geq 2\delta$ by our assumption on δ . Thus

$$\left| \frac{w_e^+ \rho_e^+}{s_e^+} \right| + \left| \frac{w_e^- \rho_e^-}{s_e^-} \right| \leq 6\gamma,$$

a contradiction.

(2) $\min\{w_e^+, w_e^-\} \geq 96 \cdot \delta^4 \|\mathbf{w}\|_1^2$: We will first prove that $|\tilde{f}_e| < 2\delta$. Suppose to the contrary. We have that $|\rho_e^-| = \frac{|\tilde{f}_e|}{s_e^-} \geq \left| \frac{\tilde{f}_e}{s_e^-} \right| \geq 2\delta$, so

$$\left| \frac{w_e^+ \rho_e^+}{s_e^+} \right| + \left| \frac{w_e^- \rho_e^-}{s_e^-} \right| \leq \frac{w_e^+ \delta}{s_e^+} + \frac{w_e^- \delta}{s_e^-} + \gamma \leq \frac{1}{2} \left| \frac{w_e^+ \rho_e^+}{s_e^+} \right| + \frac{1}{2} \left| \frac{w_e^- \rho_e^-}{s_e^-} \right| + \gamma$$

and thus

$$\left| \frac{w_e^+ \rho_e^+}{s_e^+} \right| + \left| \frac{w_e^- \rho_e^-}{s_e^-} \right| \leq 2\gamma,$$

a contradiction, therefore $|\tilde{f}_e| < 2\delta$. This immediately implies that $2\delta > |\tilde{f}_e| = |\rho_e^+| s_e^+ \geq C_\infty s_e^+$, and so $s_e^+ \leq 2\delta/C_\infty$. Therefore we get

$$\left| \frac{w_e^+ \rho_e^+}{s_e^+} \right| \geq \frac{(96 \cdot \delta^4 \|\mathbf{w}\|_1^2) C_\infty}{2\delta/C_\infty} = 48 \cdot \delta^3 \|\mathbf{w}\|_1^2 \cdot C_\infty^2.$$

Since $C_\infty^2 = \frac{1}{8\delta^2 \|\mathbf{w}\|_1}$, we conclude that

$$\left| \frac{w_e^+ \rho_e^+}{s_e^+} \right| \geq 6 \cdot \delta \|\mathbf{w}\|_1 \geq 3 \frac{w_e^- \delta}{s_e^-},$$

where we used the fact that $\|\mathbf{w}\|_\infty \leq \|\mathbf{w}\|_1$ and the fact that $s_e^- \geq \frac{1}{2}$. Therefore

$$\left| \frac{w_e^+ \rho_e^+}{s_e^+} \right| + \left| \frac{w_e^- \rho_e^-}{s_e^-} \right| \leq \frac{w_e^+ \delta}{s_e^+} + \frac{w_e^- \delta}{s_e^-} + \gamma \leq \frac{1}{2} \left| \frac{w_e^+ \rho_e^+}{s_e^+} \right| + \frac{1}{3} \left| \frac{w_e^+ \rho_e^+}{s_e^+} \right| + \gamma.$$

Therefore we conclude that

$$\left| \frac{w_e^+ \rho_e^+}{s_e^+} \right| + \left| \frac{w_e^- \rho_e^-}{s_e^-} \right| \leq 6\gamma,$$

again a contradiction. □

Proof of Lemma 3.4.18

Proof. Let us analyze the new residual after performing the update described in (3.47-3.48). Just like in the standard correction step which we analyzed in Section 3.3.4, here we can write:

$$\begin{aligned} \nabla F_\mu^w(\mathbf{x}') &= \frac{\mathbf{C}^\top \mathbf{c}}{\mu} + \mathbf{C}^\top \left(\frac{\mathbf{w}^+}{(s^+)' } - \frac{\mathbf{w}^-}{(s^-)' } \right) \\ &= \nabla F_\mu^w(\mathbf{x}) + \mathbf{C}^\top \left(\frac{\mathbf{w}^+}{(s^+)' } - \frac{\mathbf{w}^-}{(s^-)' } \right) - \mathbf{C}^\top \left(\frac{\mathbf{w}^+}{s^+} - \frac{\mathbf{w}^-}{s^-} \right) \\ &= \mathbf{C}^\top \Delta \mathbf{h} - \mathbf{C}^\top \left(\frac{\mathbf{w}^+ \rho^+}{s^+} - \frac{\mathbf{w}^- \rho^-}{s^-} \right) + \mathbf{C}^\top \left(\frac{\mathbf{w}^+}{(s^+)' } - \frac{\mathbf{w}^-}{(s^-)' } \right) - \mathbf{C}^\top \left(\frac{\mathbf{w}^+}{s^+} - \frac{\mathbf{w}^-}{s^-} \right) \\ &= \mathbf{C}^\top \Delta \mathbf{h} - \mathbf{C}^\top \left(\frac{\mathbf{w}^+ (1 + \rho^+)}{s^+} - \frac{\mathbf{w}^- (1 + \rho^-)}{s^-} \right) + \mathbf{C}^\top \left(\frac{\mathbf{w}^+}{(s^+)' } - \frac{\mathbf{w}^-}{(s^-)' } \right) \\ &= \mathbf{C}^\top \Delta \mathbf{h} - \mathbf{C}^\top \left(\frac{\mathbf{w}^+ (1 + \rho^+) (1 - \rho^+)}{(s^+)' } - \frac{\mathbf{w}^- (1 + \rho^-) (1 - \rho^-)}{(s^-)' } \right) + \mathbf{C}^\top \left(\frac{\mathbf{w}^+}{(s^+)' } - \frac{\mathbf{w}^-}{(s^-)' } \right) \\ &= \mathbf{C}^\top \Delta \mathbf{h} + \mathbf{C}^\top \left(\frac{\mathbf{w}^+ (\rho^+)^2}{(s^+)' } - \frac{\mathbf{w}^- (\rho^-)^2}{(s^-)' } \right). \end{aligned}$$

Therefore

$$\nabla F_\mu^w(\mathbf{x}') - \mathbf{C}^\top \Delta \mathbf{h} = \mathbf{C}^\top \left(\frac{\mathbf{w}^+ (\rho^+)^2}{(s^+)' } - \frac{\mathbf{w}^- (\rho^-)^2}{(s^-)' } \right).$$

Next we show that modifying the weights from \mathbf{w} to \mathbf{w}' sets a subset of these entries to 0, which will enable us to correct this new perturbed residual. Using the weight update described in (3.53-3.54) we obtain

$$\begin{aligned}
& -\mathbf{C}^\top (\mathbf{g} + \Delta \mathbf{h}) \\
& = \nabla F_\mu^{\mathbf{w}'}(\mathbf{x}') - \mathbf{C}^\top \Delta \mathbf{h} \\
& = \mathbf{C}^\top \frac{\mathbf{c}}{\mu} - \mathbf{C}^\top \Delta \mathbf{h} + \mathbf{C}^\top \left(\frac{(\mathbf{w}^+)'}{(\mathbf{s}^+)' } - \frac{(\mathbf{w}^-)' }{(\mathbf{s}^-)' } \right) \\
& = \nabla F_\mu^{\mathbf{w}}(\mathbf{x}') - \mathbf{C}^\top \Delta \mathbf{h} + \mathbf{C}^\top \left(\frac{(\mathbf{w}^+)' - \mathbf{w}^+}{(\mathbf{s}^+)' } - \frac{(\mathbf{w}^-)' - \mathbf{w}^-}{(\mathbf{s}^-)' } \right) \\
& = \nabla F_\mu^{\mathbf{w}}(\mathbf{x}') - \mathbf{C}^\top \Delta \mathbf{h} + \mathbf{C}^\top \left(\frac{\mathbf{w}^- (\rho^-)^2}{(\mathbf{s}^-)' } \cdot \mathbf{1}_{|\rho^-| \geq C_\infty} - \frac{\mathbf{w}^+ (\rho^+)^2}{(\mathbf{s}^+)' } \cdot \mathbf{1}_{|\rho^+| \geq C_\infty} \right) \\
& = \mathbf{C}^\top \left(\left(\frac{\mathbf{w}^+ (\rho^+)^2}{(\mathbf{s}^+)' } \cdot \mathbf{1}_{|\rho^+| < C_\infty} \right) - \left(\frac{\mathbf{w}^- (\rho^-)^2}{(\mathbf{s}^-)' } \cdot \mathbf{1}_{|\rho^-| < C_\infty} \right) \right).
\end{aligned}$$

Finally, using Lemma 3.3.4 we certify an upper bound on

$$\begin{aligned}
\mathcal{E}_{\mathbf{w}', \mathbf{s}'}(\mathbf{g} + \Delta \mathbf{h}) & \leq \mathcal{E}^{\max}(\mathbf{g} + \Delta \mathbf{h}, \mathbf{w}', \mathbf{s}') \\
& = \frac{1}{2} \sum_{e \in E} \frac{\left(\left(\frac{\mathbf{w}^+ (\rho^+)^2}{(\mathbf{s}^+)' } \cdot \mathbf{1}_{|\rho^+| < C_\infty} \right) - \left(\frac{\mathbf{w}^- (\rho^-)^2}{(\mathbf{s}^-)' } \cdot \mathbf{1}_{|\rho^-| < C_\infty} \right) \right)_e^2}{\frac{w_e^+}{(s_e^+)^{1/2}} + \frac{w_e^-}{(s_e^-)^{1/2}}} \\
& \leq \frac{1}{2} \sum_{\substack{e \in E \\ |\rho_e^+| < C_\infty}} \frac{\left(\frac{w_e^+ (\rho_e^+)^2}{(s_e^+)' } \right)^2}{\frac{w_e^+}{(s_e^+)^{1/2}}} + \frac{1}{2} \sum_{\substack{e \in E \\ |\rho_e^-| < C_\infty}} \frac{\left(\frac{w_e^- (\rho_e^-)^2}{(s_e^-)' } \right)^2}{\frac{w_e^-}{(s_e^-)^{1/2}}} \\
& \leq \frac{1}{2} \sum_{\substack{e \in E \\ |\rho_e^+| < C_\infty}} w_e^+ (\rho_e^+)^4 + \frac{1}{2} \sum_{\substack{e \in E \\ |\rho_e^-| < C_\infty}} w_e^- (\rho_e^-)^4 \\
& \leq \frac{1}{2} C_\infty^2 \cdot \sum_{e \in E} (w_e^+ (\rho_e^+)^2 + w_e^- (\rho_e^-)^2) \\
& = \frac{1}{2} C_\infty^2 \cdot \sum_{e \in E} (\tilde{f}_e)^2 \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) \\
& \leq \frac{1}{2} C_\infty^2 \cdot 8 \cdot \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}) \\
& \leq \frac{2\delta^2 \|\mathbf{w}\|_1}{8\delta^2 \|\mathbf{w}\|_1} \\
& = \frac{1}{4},
\end{aligned}$$

where we used the fact that $\mathbf{w}' \geq \mathbf{w}$, (3.24), and Lemma 3.4.5. \square

Proof of Lemma 3.6.3

Proof. Let $\tilde{\mathbf{f}}_\star = \tilde{\mathbf{f}} + \tilde{\mathbf{f}}'$ where $\tilde{\mathbf{f}}'$ is the restriction of $\tilde{\mathbf{f}}_\star$ to the edges incident to v_\star . To prove (3.65), note that from Lemma 3.6.2 there exists a vector $\boldsymbol{\alpha} = (\boldsymbol{\alpha}^+; \boldsymbol{\alpha}^-)$, $(1 + \theta)^{-2} \cdot \mathbf{1} \leq \boldsymbol{\alpha} \leq (1 - \theta)^{-2} \cdot \mathbf{1}$,

such that for any circulation $\mathbf{g} = \mathbf{C}_\star \mathbf{z}_\star$ in G_\star ,

$$\left\langle \mathbf{g}, \begin{bmatrix} \mathbf{h} - \tilde{\mathbf{f}} \left(\frac{\alpha^+ w^+}{(s^+)^2} + \frac{\alpha^- w^-}{(s^-)^2} \right) - R_p \cdot (\tilde{\mathbf{f}})^{p-1} \\ -R_\star \cdot \tilde{\mathbf{f}}' - R_p \cdot (\tilde{\mathbf{f}}')^{p-1} \end{bmatrix} \right\rangle = 0. \quad (9.21)$$

Restricting ourselves to circulations supported only in the non-preconditioned graph G , one has that for any circulation in $\mathbf{g}' = \mathbf{C} \mathbf{z}$ in G :

$$\left\langle \mathbf{z}, \mathbf{C}^\top \left(\mathbf{h} + \Delta \mathbf{h} - \tilde{\mathbf{f}} \left(\frac{\alpha^+ w^+}{(s^+)^2} + \frac{\alpha^- w^-}{(s^-)^2} \right) \right) \right\rangle = 0$$

Since this holds for any test vector \mathbf{z} , it must be that the second term in the inner product is $\mathbf{0}$. Rearranging, it yields the identity from (3.65). Next, we notice that (9.21) is the optimality condition of the following objective:

$$\begin{aligned} \max_{\tilde{\mathbf{f}}_\star = \mathbf{C}_\star \tilde{\mathbf{x}}} \left\langle \mathbf{h}, \tilde{\mathbf{f}} \right\rangle - \frac{1}{2} \sum_{e \in E} (\tilde{f}_e)^2 \cdot \left(\frac{\alpha_e^+ w_e^+}{(s_e^+)^2} + \frac{\alpha_e^- w_e^-}{(s_e^-)^2} \right) \\ - \frac{R_\star}{2} \sum_{e \in E'} (\tilde{f}'_e)^2 - \frac{R_p}{p} \sum_{e \in E \cup E'} (\tilde{f}_\star)_e^p. \end{aligned} \quad (9.22)$$

Let us proceed to bound the norm of the demand routed by $\tilde{\mathbf{f}}$. Consider the value of the objective in (9.22) after truncating it to only the first two terms, which we can write as:

$$\left\langle \mathbf{h}, \tilde{\mathbf{f}} \right\rangle - \frac{1}{2} \sum_{e \in E} (\tilde{f}_e)^2 \cdot \left(\frac{\alpha_e^+ w_e^+}{(s_e^+)^2} + \frac{\alpha_e^- w_e^-}{(s_e^-)^2} \right) \quad (9.23)$$

$$\leq \sum_{e \in E} h_e \cdot \tilde{f}_e - \frac{1}{3} \sum_{e \in E} (\tilde{f}_e)^2 \cdot \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right) \quad (9.24)$$

$$\leq \frac{3}{4} \sum_{e \in E} h_e^2 \cdot \left(\frac{w_e^+}{(s_e^+)^2} + \frac{w_e^-}{(s_e^-)^2} \right)^{-1} \quad (9.25)$$

$$= \frac{3}{2} \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}), \quad (9.26)$$

where we used $\alpha \geq \frac{1}{(1+\theta)^2} \cdot \mathbf{1} > \frac{2}{3} \cdot \mathbf{1}$ and the fact that $\langle \mathbf{a}, \mathbf{b} \rangle \leq \frac{1}{2} \|\mathbf{a}\|^2 + \frac{1}{2} \|\mathbf{b}\|^2$.

Note that the value of the regularized objective (9.22) is at least 0 since we can always substitute $\tilde{\mathbf{x}} = \mathbf{0}$ and obtain exactly 0. By re-arranging,

$$\frac{R_\star}{2} \sum_{e \in E'} (\tilde{f}'_e)^2 \quad (9.27)$$

$$\leq \left\langle \mathbf{h}, \tilde{\mathbf{f}} \right\rangle - \frac{1}{2} \sum_{e \in E} (\tilde{f}_e)^2 \cdot \left(\frac{\alpha_e^+ w_e^+}{(s_e^+)^2} + \frac{\alpha_e^- w_e^-}{(s_e^-)^2} \right) - \frac{R_p}{p} \sum_{e \in E \cup E'} (\tilde{f}_\star)_e^p \quad (9.28)$$

$$\leq \left\langle \mathbf{h}, \tilde{\mathbf{f}} \right\rangle - \frac{1}{2} \sum_{e \in E} (\tilde{f}_e)^2 \cdot \left(\frac{\alpha_e^+ w_e^+}{(s_e^+)^2} + \frac{\alpha_e^- w_e^-}{(s_e^-)^2} \right) \quad (9.29)$$

$$\leq \frac{3}{2} \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}), \quad (9.30)$$

where we also used the fact that the last term of (9.28) is non-positive and (9.26). Therefore (9.30) enables us to upper bound

$$\sum_{e \in E'} (\tilde{f}'_e)^2 \leq \frac{3}{R_\star} \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}), \quad (9.31)$$

which implies that

$$\sum_{e \in E'} |\tilde{f}'_e| \leq |E'|^{1/2} \cdot \left(\sum_{e \in E'} (\tilde{f}'_e)^2 \right)^{1/2} \quad (9.32)$$

$$\leq \left(3 \|\mathbf{w}\|_1 \cdot \frac{3}{R_\star} \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}) \right)^{1/2} \quad (9.33)$$

$$= 3 \left(\frac{\|\mathbf{w}\|_1 \cdot \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s})}{R_\star} \right)^{1/2}, \quad (9.34)$$

a quantity that upper bounds the demand perturbation. Using a similar argument we can upper bound $\|\tilde{\mathbf{f}}_\star\|_p$. We have

$$\frac{R_p}{p} \sum_{e \in E \cup E'} (\tilde{f}_\star)_e^p \quad (9.35)$$

$$\leq \langle \mathbf{h}, \tilde{\mathbf{f}} \rangle - \frac{1}{2} \sum_{e \in E} (\tilde{f}_e)^2 \cdot \left(\frac{\alpha_e^+ w_e^+}{(s_e^+)^2} + \frac{\alpha_e^- w_e^-}{(s_e^-)^2} \right) - \frac{R_\star}{2} \sum_{e \in E'} (\tilde{f}'_e)^2 \quad (9.36)$$

$$\leq \langle \mathbf{h}, \tilde{\mathbf{f}} \rangle - \frac{1}{2} \sum_{e \in E} (\tilde{f}_e)^2 \cdot \left(\frac{\alpha_e^+ w_e^+}{(s_e^+)^2} + \frac{\alpha_e^- w_e^-}{(s_e^-)^2} \right) \quad (9.37)$$

$$\leq \frac{3}{2} \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s}), \quad (9.38)$$

thus concluding that

$$\|\tilde{\mathbf{f}}_\star\|_p = \left(\sum_{e \in E \cup E'} (\tilde{f}_\star)_e^p \right)^{1/p} \leq \left(\frac{p \cdot \frac{3}{2} \mathcal{E}^{\max}(\mathbf{h}, \mathbf{w}, \mathbf{s})}{R_p} \right)^{1/p}. \quad (9.39)$$

□

9.2 Appendix for Chapter 4

9.2.1 Maintaining the Schur Complement

Following the scheme from [73] we maintain a dynamic Schur complement of the graph onto a subset of terminals C . The approach follows rather directly from [73] and leverages the recent work of [18] to dynamically maintain an edge sparsifier of the Schur complement of the graph onto C . Compared to [73] we do not require a parameter that depends on the adaptivity of the adversary. In addition, when adding a vertex to C we also return a $(1 + \varepsilon)$ -approximation of the effective resistance $R_{eff}(v, C)$, which gets returned by the function call.

Lemma 9.2.1 (DYNAMICSC (Theorem 4, [73])). *There is a DYNAMICSC data structure supporting*

the following operations with the given runtimes against oblivious adversaries, for constants $0 < \beta, \varepsilon < 1$:

- **INITIALIZE**($G, C^{(init)}, \mathbf{r}, \varepsilon, \beta$): Initializes a graph G with resistances \mathbf{r} and a set of safe terminals $C^{(safe)}$. Sets the terminal set $C = C^{(safe)} \cup C^{(init)}$. Runtime: $\tilde{O}(m\beta^{-4}\varepsilon^{-4})$.
- **ADDTERMINAL**($v \in V(G)$): Returns $\tilde{R}_{eff}(C, v) \approx_2 R_{eff}(C, v)$ and adds v as a terminal. Runtime: Amortized $\tilde{O}(\beta^{-2}\varepsilon^{-2})$.
- **TEMPORARYADDTERMINALS**($\Delta C \subseteq V(G)$): Adds all vertices in the set ΔC as (temporary) terminals. Runtime: Worst case $\tilde{O}(K^2\beta^{-4}\varepsilon^{-4})$, where K is the total number of terminals added by all of the **TEMPORARYADDTERMINALS** operations that have not been rolled back using **ROLLBACK**. All **TEMPORARYADDTERMINALS** operations should be rolled back before the next call to **ADDTERMINALS**.
- **UPDATE**(e, r): Under the guarantee that both endpoints of e are terminals, updates $r_e = r$. Runtime: Worst case $\tilde{O}(1)$.
- $\widetilde{SC}()$: Returns a spectral sparsifier $\widetilde{SC} \approx_{1+\varepsilon} SC(G, C)$ (with respect to resistances \mathbf{r}) with $\tilde{O}(|C|\varepsilon^{-2})$ edges. Runtime: Worst case $\tilde{O}((\beta m + (K\beta^{-2}\varepsilon^{-2})^2)\varepsilon^{-2})$ where K is the total number of terminals added by all of the **TEMPORARYADDTERMINALS** operations that have not been rolled back.
- **ROLLBACK**(): Rolls back the last **UPDATE**, **ADDTERMINALS**, or **TEMPORARYADDTERMINALS** if it exists. The runtime is the same as the original operation.

Finally, all calls return valid outputs with high probability. The size of C should always be $O(\beta m)$.

This data structure is analyzed in detail in [73]. Additionally, let us show that an approximation to $R_{eff}(v, C)$ can be efficiently computed along with the **ADDTERMINAL** operation. To get an estimate we simply inspect the neighbors of v in the sparsified Schur complement of $C \cup \{v\}$ and compute the inverse of the sum of their inverses. This is indeed a $1 + O(\varepsilon)$ -approximation, as effective resistances are preserved within a $1 + O(\varepsilon)$ factor in the sparsifier.

To show that this operation takes little amortized time, we note that by the proof appearing in [73, Lemma 6.2], vertex v appears in amortized $\tilde{O}(1)$ expanders maintained dynamically. As the dynamic sparsifier keeps $\tilde{O}(\varepsilon^{-2})$ neighbors of v from each expander, the number of neighbors to inspect with each call is $\tilde{O}(\varepsilon^{-2})$, which also bounds the time necessary to approximate the resistance.

9.2.2 Auxiliary Lemmas

Lemma 4.2.6. *Let \mathbf{d} be a demand vector, let \mathbf{r} be resistances, and let $C \subseteq V$ be a subset of vertices. Then*

$$\mathcal{E}_r(\boldsymbol{\pi}^C(\mathbf{d})) \leq \mathcal{E}_r(\mathbf{d}).$$

Proof. Letting $F = V \setminus C$, and \mathbf{L} be the Laplacian of the underlying graph, we can write

$$\boldsymbol{\pi}^C(\mathbf{d}) = \mathbf{d}_C - \mathbf{L}_{CF} \mathbf{L}_{FF}^{-1} \mathbf{d}_F.$$

By factoring \mathbf{L}^+ as

$$\mathbf{L}^+ = \begin{bmatrix} I & 0 \\ -\mathbf{L}_{FF}^{-1} \mathbf{L}_{FC} & I \end{bmatrix} \begin{bmatrix} SC(\mathbf{L}, C)^+ & 0 \\ 0 & \mathbf{L}_{FF}^{-1} \end{bmatrix} \begin{bmatrix} I & -\mathbf{L}_{CF} \mathbf{L}_{FF}^{-1} \\ 0 & I \end{bmatrix}$$

we can write

$$\mathcal{E}_r(\mathbf{d}) = \mathbf{d}^\top \mathbf{L}^+ \mathbf{d} = \begin{bmatrix} \boldsymbol{\pi}^C(\mathbf{d}) \\ \mathbf{d}_F \end{bmatrix}^\top \begin{bmatrix} SC(\mathbf{L}, C)^+ & 0 \\ 0 & \mathbf{L}_{FF}^{-1} \end{bmatrix} \begin{bmatrix} \boldsymbol{\pi}^C(\mathbf{d}) \\ \mathbf{d}_F \end{bmatrix} = \|\boldsymbol{\pi}^C(\mathbf{d})\|_{SC(\mathbf{L}, C)^+}^2 + \|\mathbf{d}_F\|_{\mathbf{L}_{FF}^{-1}}^2.$$

Furthermore, we can use the same factorization to write

$$\mathcal{E}_r(\boldsymbol{\pi}^C(\mathbf{d})) = \begin{bmatrix} \boldsymbol{\pi}^C(\mathbf{d}) \\ 0 \end{bmatrix}^\top \begin{bmatrix} SC(\mathbf{L}, C)^+ & 0 \\ 0 & \mathbf{L}_{FF}^{-1} \end{bmatrix} \begin{bmatrix} \boldsymbol{\pi}^C(\mathbf{d}) \\ 0 \end{bmatrix} = \|\boldsymbol{\pi}^C(\mathbf{d})\|_{SC(\mathbf{L}, C)^+}^2,$$

which proves the claim. \square

Lemma 9.2.2. *For any $\mu \in (1/\text{poly}(m), \text{poly}(m))$, we have $\|\mathbf{r}(\mu)\|_\infty \leq m^{\tilde{O}(\log m)}$.*

Proof. By Appendix A in [10], for some $\mu_0 = \Theta(\|\mathbf{c}\|_2)$, the solution $\mathbf{f} = \mathbf{u}/2$ has

$$\left\| \mathbf{C}^\top \left(\frac{\mathbf{c}}{\mu_0} + \frac{\mathbf{1}}{\mathbf{s}^+} - \frac{\mathbf{1}}{\mathbf{s}^-} \right) \right\|_{(\mathbf{C}^\top \mathbf{R} \mathbf{C})^+} \leq 1/10.$$

This implies that $\min_e \{s_e(\mu_0)^+, s_e(\mu_0)^-\} \geq \min_e u_e/4 \geq 1/4$, and so $\|\mathbf{r}(\mu_0)\|_\infty \leq O(1)$. Additionally, $\|\mathbf{c}\|_\infty \in [1, \text{poly}(m)]$, so $\mu_0 = \Theta(\text{poly}(m))$.

Now, for any integer $i \geq 0$ we let $\mu_{i+1} = \mu_i \cdot (1 - 1/\sqrt{m})^{\sqrt{m}/10}$. By Lemma 9.2.7 we have that $\mathbf{r}(\mu_{i+1}) \approx_{m^2} \mathbf{r}(\mu_i)$, and so

$$\mathbf{r} \left(\frac{1}{\text{poly}(m)} \right) = \mathbf{r}(\mu_{\tilde{O}(\log m)}) \leq \left(\frac{9}{100} m^2 \right)^{\tilde{O}(\log m)} \mathbf{r}(\mu_0) \leq m^{\tilde{O}(\log m)} \mathbf{r}(\mu_0) \leq m^{\tilde{O}(\log m)}.$$

\square

Lemma 9.2.3. *Given a graph $G(V, E)$ with resistances \mathbf{r} and any parameter $\varepsilon > 0$, there exists an algorithm that runs in time $\tilde{O}(m/\varepsilon^2)$ and produces a matrix $\mathbf{Q} \in \mathbb{R}^{\tilde{O}(1/\varepsilon^2) \times n}$ such that with high probability for any $u, v \in V$,*

$$R_{eff}(u, v) \approx_{1+\varepsilon} \|\mathbf{Q}\mathbf{1}_u - \mathbf{Q}\mathbf{1}_v\|_2^2$$

9.2.3 Deferred Proofs from Section 4.3

Central path stability bounds

Lemma 9.2.4 (Central path energy stability). *Consider a minimum cost flow instance on a graph $G(V, E)$. For any $\mu > 0$ and $\mu' = \mu/(1 + 1/\sqrt{m})^k$ for some $k \in (0, \sqrt{m}/10)$, we have*

$$\sum_{e \in E} \left(\frac{1}{s_e(\mu)^+ \cdot s_e(\mu')^+} + \frac{1}{s_e(\mu)^- \cdot s_e(\mu')^-} \right) (f_e(\mu') - f_e(\mu))^2 \leq 2k^2$$

Proof of Lemma 9.2.4. We let $\delta = 1/\sqrt{m}$, $\mathbf{f} = \mathbf{f}(\mu)$, $\mathbf{s} = \mathbf{s}(\mu)$, $\mathbf{r} = \mathbf{r}(\mu)$, $\mathbf{f}' = \mathbf{f}(\mu')$, $\mathbf{s}' = \mathbf{s}(\mu')$,

and $\mathbf{r}' = \mathbf{r}(\mu')$. We also set $\tilde{\mathbf{f}} = \mathbf{f}' - \mathbf{f}$. By definition of centrality we have

$$\begin{aligned} \mathbf{C}^\top \left(\frac{1}{s^-} - \frac{1}{s^+} \right) &= \mathbf{C}^\top \frac{\mathbf{c}}{\mu} \\ \mathbf{C}^\top \left(\frac{1}{s^- + \tilde{\mathbf{f}}} - \frac{1}{s^+ - \tilde{\mathbf{f}}} \right) &= \mathbf{C}^\top \frac{\mathbf{c}}{\mu'}, \end{aligned}$$

which, after subtracting, give

$$\begin{aligned} \mathbf{C}^\top \left(\frac{1}{s^- + \tilde{\mathbf{f}}} - \frac{1}{s^-} - \frac{1}{s^+ - \tilde{\mathbf{f}}} + \frac{1}{s^+} \right) &= \mathbf{C}^\top \left(\frac{\mathbf{c}}{\mu'} - \frac{\mathbf{c}}{\mu} \right) \\ \Leftrightarrow \mathbf{C}^\top \left(\left(\frac{1}{s^-(s^- + \tilde{\mathbf{f}})} + \frac{1}{s^+(s^+ - \tilde{\mathbf{f}})} \right) \tilde{\mathbf{f}} \right) &= - \left((1 + \delta)^k - 1 \right) \mathbf{C}^\top \frac{\mathbf{c}}{\mu}. \end{aligned}$$

As $\tilde{\mathbf{f}} = \mathbf{C}\mathbf{x}$ for some \mathbf{x} , after taking the inner product of both sides with \mathbf{x} we get

$$\left\langle \tilde{\mathbf{f}}, \left(\frac{1}{s^-(s^- + \tilde{\mathbf{f}})} + \frac{1}{s^+(s^+ - \tilde{\mathbf{f}})} \right) \tilde{\mathbf{f}} \right\rangle = - \left((1 + \delta)^k - 1 \right) \left\langle \frac{\mathbf{c}}{\mu}, \tilde{\mathbf{f}} \right\rangle. \quad (9.40)$$

We will now prove that $-\left\langle \frac{\mathbf{c}}{\mu}, \tilde{\mathbf{f}} \right\rangle \leq k\sqrt{m}$. First of all, by differentiating the centrality condition

$$\mathbf{C}^\top \left(\frac{\mathbf{c}}{\nu} + \frac{\mathbf{1}}{s(\nu)^+} - \frac{\mathbf{1}}{s(\nu)^-} \right) = \mathbf{0}$$

with respect to ν we get

$$\mathbf{C}^\top \left(-\frac{\mathbf{c}}{\nu^2} + \left(\frac{\mathbf{1}}{(s(\nu)^+)^2} + \frac{\mathbf{1}}{(s(\nu)^-)^2} \right) \frac{d\mathbf{f}(\nu)}{d\nu} \right) = \mathbf{0},$$

or equivalently

$$\mathbf{C}^\top \left(\mathbf{r}(\nu) \frac{d\mathbf{f}(\nu)}{d\nu} \right) = -\frac{1}{\nu} \mathbf{C}^\top \left(\frac{\mathbf{1}}{s(\nu)^+} - \frac{\mathbf{1}}{s(\nu)^-} \right).$$

If we set $g(\mathbf{s}) = \frac{\frac{1}{s^+} - \frac{1}{s^-}}{r}$, this can also be equivalently written as

$$\frac{d\mathbf{f}(\nu)}{d\nu} = -\frac{1}{\nu} \left(g(\mathbf{s}(\nu)) - (\mathbf{R}(\nu))^{-1} \mathbf{B} (\mathbf{B}^\top (\mathbf{R}(\nu)^{-1}) \mathbf{B})^+ \mathbf{B}^\top g(\mathbf{s}(\nu)) \right).$$

We have

$$\begin{aligned}
-\left\langle \frac{\mathbf{c}}{\mu}, \tilde{\mathbf{f}} \right\rangle &= -\int_{\nu=\mu}^{\mu'} \left\langle \frac{\mathbf{c}}{\mu}, d\mathbf{f}(\nu) \right\rangle \\
&= \frac{1}{\mu} \int_{\nu=\mu}^{\mu'} \left\langle \frac{\nu}{\mathbf{s}(\nu)^-} - \frac{\nu}{\mathbf{s}(\nu)^+}, \frac{1}{\nu} \left(g(\mathbf{s}(\nu)) - (\mathbf{R}(\nu))^{-1} \mathbf{B} (\mathbf{B}^\top (\mathbf{R}(\nu))^{-1} \mathbf{B})^+ \mathbf{B}^\top g(\mathbf{s}(\nu)) \right) \right\rangle d\nu \\
&= -\frac{1}{\mu} \int_{\nu=\mu}^{\mu'} \left\langle \sqrt{\mathbf{r}(\nu)} g(\mathbf{s}(\nu)), \mathbf{\Pi}_{\ker(\mathbf{B}^\top (\mathbf{R}(\nu))^{-1/2})} \sqrt{\mathbf{r}(\nu)} g(\mathbf{s}(\nu)) \right\rangle d\nu \\
&= \frac{1}{\mu} \int_{\nu=\mu}^{\mu'} \left\| \mathbf{\Pi}_{\ker(\mathbf{B}^\top (\mathbf{R}(\nu))^{-1/2})} \sqrt{\mathbf{r}(\nu)} g(\mathbf{s}(\nu)) \right\|_2^2 d\nu \\
&\leq \frac{1}{\mu} \int_{\nu=\mu}^{\mu'} \left\| \sqrt{\mathbf{r}(\nu)} g(\mathbf{s}(\nu)) \right\|_2^2 d\nu \\
&\leq \frac{1}{\mu} \int_{\nu=\mu}^{\mu'} m d\nu \\
&= m \frac{\mu - \mu'}{\mu} \\
&= m(1 - (1 + \delta)^{-k}) \\
&\leq \delta k m \\
&= k\sqrt{m},
\end{aligned}$$

where $\mathbf{\Pi}_{\ker(\mathbf{B}^\top (\mathbf{R}(\nu))^{-1/2})} = \mathbf{I} - (\mathbf{R}(\nu))^{-1/2} \mathbf{B} (\mathbf{B}^\top (\mathbf{R}(\nu))^{-1} \mathbf{B})^+ \mathbf{B}^\top (\mathbf{R}(\nu))^{-1/2}$ is the orthogonal projection onto the kernel of $\mathbf{B}^\top (\mathbf{R}(\nu))^{-1/2}$.

Plugging this into (9.40) and using the fact that $(1 + \delta)^k \leq 1 + 1.1\delta k = 1 + 1.1k/\sqrt{m}$, we get

$$\sum_{e \in E} \left(\frac{1}{s_e(\mu)^+ \cdot s_e(\mu')^+} + \frac{1}{s_e(\mu)^- \cdot s_e(\mu')^-} \right) (f_e(\mu') - f_e(\mu))^2 \leq 2k^2.$$

□

We give an auxiliary lemma which converts between different kinds of slack approximations.

Lemma 9.2.5. *We consider flows \mathbf{f}, \mathbf{f}' with slacks \mathbf{s}, \mathbf{s}' and resistances \mathbf{r}, \mathbf{r}' . Then,*

$$\max \left\{ \left| \frac{s_e'^+ - s_e^+}{s_e^+} \right|, \left| \frac{s_e'^- - s_e^-}{s_e^-} \right| \right\} \leq \sqrt{r_e} |f_e' - f_e| \leq \sqrt{2} \max \left\{ \left| \frac{s_e'^+ - s_e^+}{s_e^+} \right|, \left| \frac{s_e'^- - s_e^-}{s_e^-} \right| \right\}$$

and if $r_e \not\approx_{1+\gamma} r_e'$ for some $\gamma \in (0, 1)$, then $\sqrt{r_e} |f_e' - f_e| \geq \gamma/6$.

Proof. For the first one, note that

$$r_e = \frac{1}{(s_e^+)^2} + \frac{1}{(s_e^-)^2} \in \left[\max \left\{ \frac{1}{(s_e^+)^2}, \frac{1}{(s_e^-)^2} \right\}, 2 \max \left\{ \frac{1}{(s_e^+)^2}, \frac{1}{(s_e^-)^2} \right\} \right].$$

Together with the fact that $|f_e' - f_e| = |s_e'^+ - s_e^+| = |s_e'^- - s_e^-|$, it implies the first statement.

For the second one, without loss of generality let $s_e^+ \leq s_e^-$, so by the previous statement we have $\sqrt{r_e} |f_e' - f_e| \geq \frac{|s_e'^+ - s_e^+|}{s_e^+}$. If this is $< \gamma/6$ then $(1 - \gamma/6)s_e^+ \leq s_e'^+ \leq (1 + \gamma/6)s_e^+$, so

$s_e^+ \approx_{1+\gamma/3} s_e^+$. However, we also have that $\frac{|s_e'^- - s_e^-|}{s_e^-} \leq \frac{|s_e'^+ - s_e^+|}{s_e^+} \leq \gamma/6$, so $s_e'^- \approx_{1+\gamma/3} s_e^-$. Therefore, $r_e' = \frac{1}{(s_e'^+)^2} + \frac{1}{(s_e'^-)^2} \approx_{1+\gamma} \frac{1}{(s_e^+)^2} + \frac{1}{(s_e^-)^2} = r_e$, a contradiction. \square

The following lemma is a fine-grained explanation of how resistances can change.

Lemma 9.2.6. *Consider a minimum cost flow instance on a graph $G(V, E)$ and parameters $\mu > 0$ and $\mu' \geq \mu/(1+1/\sqrt{m})^k$, where $k \in (0, \sqrt{m}/10)$. For any $e \in E$ and $\gamma \in (0, 1)$ we let $\text{change}(e, \gamma)$ be the largest integer $t(e) \geq 0$ such that there are real numbers $\mu = \mu_1(e) > \mu_2(e) > \dots > \mu_{t(e)+1}(e) = \mu'$ with $\sqrt{r_e(\mu_i)} |f_e(\mu_{i+1}) - f_e(\mu_i)| \geq \gamma$ for all $i \in [t(e)]$.*

Then, $\sum_{e \in E} (\text{change}(e, \gamma))^2 \leq O(k^2/\gamma^2)$.

Proof of Lemma 9.2.6. First, we assume that without loss of generality, $r_e(\mu_{i+1}(e)) \approx_{(1+6\gamma)^2} r_e(\mu_i(e))$ for all $e \in E$ and $i \in [t(e)]$. If this is not true, then by continuity there exists a $\nu \in (\mu_{i+1}, \mu_i)$ such that $r_e(\mu_{i+1}(e)) \not\approx_{1+6\gamma} r_e(\nu)$ and $r_e(\nu) \not\approx_{1+6\gamma} r_e(\mu_i(e))$. By Lemma 9.2.5, this implies that $\sqrt{r_e(\mu_{i+1}(e))} |f_e(\nu) - f_e(\mu_{i+1}(e))| \geq \gamma$ and $\sqrt{r_e(\nu)} |f_e(\mu_i(e)) - f_e(\nu)| \geq \gamma$. Therefore we can break the interval (μ_{i+1}, μ_i) into (μ_{i+1}, ν) and (ν, μ_i) and make the statement stronger.

Similarly, we also assume that $r_e(\nu) \approx_{(1+6\gamma)^3} r_e(\mu_i(e))$ for all $e \in E$, $i \in [t(e)]$, and $\nu \in (\mu_{i+1}, \mu_i)$. If this is not the case, then by using the fact that $r_e(\mu_{i+1}(e)) \approx_{(1+6\gamma)^2} r_e(\mu_i(e))$, we also get that $r_e(\mu_{i+1}) \not\approx_{1+6\gamma} r_e(\nu)$, and so we can again break the interval as before and obtain a stronger statement.

Now, we look at the following integral:

$$\mathcal{E} := \int_{\nu=\mu}^{\mu'} \sum_{e \in E} r_e(\nu) \left(\frac{df_e(\nu)}{d\nu} \right)^2 |d\nu| ,$$

where $df_e(\nu)$ is the differential of the flow $f_e(\nu)$ with respect to the centrality parameter. Similarly to Lemma 9.2.4, we use the following equation that describes how the flow changes:

$$\frac{d\mathbf{f}(\nu)}{d\nu} = -\frac{1}{\nu} \left(g(\mathbf{s}(\nu)) - (\mathbf{R}(\nu))^{-1} \mathbf{B} (\mathbf{B}^\top (\mathbf{R}(\nu)^{-1}) \mathbf{B})^+ \mathbf{B}^\top g(\mathbf{s}(\nu)) \right) .$$

This implies that

$$\begin{aligned} \left\| \sqrt{\mathbf{r}(\nu)} \frac{d\mathbf{f}(\nu)}{d\nu} \right\|_2^2 &= \frac{1}{\nu^2} \left\| \sqrt{\mathbf{r}(\nu)} g(\mathbf{s}(\nu)) - (\mathbf{R}(\nu))^{-1/2} \mathbf{B} (\mathbf{B}^\top (\mathbf{R}(\nu)^{-1}) \mathbf{B})^+ \mathbf{B}^\top g(\mathbf{s}(\nu)) \right\|_2^2 \\ &\leq \frac{1}{\nu^2} \left\| \left(I - (\mathbf{R}(\nu))^{-1/2} \mathbf{B} (\mathbf{B}^\top (\mathbf{R}(\nu)^{-1}) \mathbf{B})^+ \mathbf{B}^\top (\mathbf{R}(\nu))^{-1/2} \right) \sqrt{\mathbf{r}} g(\mathbf{s}(\nu)) \right\|_2^2 \\ &\leq \frac{1}{\nu^2} \left\| \sqrt{\mathbf{r}} g(\mathbf{s}(\nu)) \right\|_2^2 \\ &\leq \frac{m}{\nu^2} , \end{aligned}$$

and so

$$\mathcal{E} \leq \int_{\nu=\mu}^{\mu'} \frac{m}{\nu^2} |d\nu| = m \left(\frac{1}{\mu'} - \frac{1}{\mu} \right) \leq m \frac{1.1\delta k}{\mu} = 1.1k\sqrt{m}/\mu . \quad (9.41)$$

On the other hand, for any $e \in E$ and $i \in [t(e)]$ we have

$$\begin{aligned}
\int_{\nu=\mu_i(e)}^{\mu_{i+1}(e)} r_e(\nu) \left(\frac{df_e(\nu)}{d\nu} \right)^2 |d\nu| &\geq \frac{r_e(\mu_i(e))}{(1+6\gamma)^3} \int_{\nu=\mu_i(e)}^{\mu_{i+1}(e)} \left(\frac{df_e(\nu)}{d\nu} \right)^2 |d\nu| \\
&\geq \frac{r_e(\mu_i(e))}{(1+6\gamma)^3} \frac{\left(\int_{\nu=\mu_i(e)}^{\mu_{i+1}(e)} \left| \frac{df_e(\nu)}{d\nu} \right| |d\nu| \right)^2}{\int_{\nu=\mu_i(e)}^{\mu_{i+1}(e)} |d\nu|} \\
&= \frac{r_e(\mu_i(e))}{(1+6\gamma)^3(\mu_i(e) - \mu_{i+1}(e))} (f(\mu_i(e)) - f(\mu_{i+1}(e)))^2 \\
&\geq \frac{\gamma^2}{36(1+6\gamma)^3(\mu_i(e) - \mu_{i+1}(e))},
\end{aligned}$$

where we used the Cauchy-Schwarz inequality. Now, note that

$$\begin{aligned}
\int_{\nu=\mu_1(e)}^{\mu_{t(e)+1}(e)} r_e(\nu) \left(\frac{df_e(\nu)}{d\nu} \right)^2 |d\nu| &\geq \sum_{i=1}^{t(e)} \frac{\gamma^2}{(1+6\gamma)^3(\mu_i(e) - \mu_{i+1}(e))} \\
&\geq \frac{\gamma^2(t(e))^2}{(1+6\gamma)^3(\mu - \mu')} \\
&\geq \frac{\gamma^2(t(e))^2 \sqrt{m}}{(1+6\gamma)^3 k \mu},
\end{aligned}$$

where remember that $t(e) = \text{change}(e, \gamma)$ and we again used Cauchy-Schwarz. Summing this up for all $e \in E$ and combining with (9.41), we get that $\sum_{e \in E} (\text{change}(e, \gamma))^2 \leq O(k^2/\gamma^2)$. \square

Lemma 9.2.7 (Central path ℓ_∞ slack stability). *Consider a minimum cost flow instance on a graph $G(V, E)$. For any $\mu > 0$ and $\mu' = \mu/(1 + 1/\sqrt{m})^k$ for some $k \in (0, \sqrt{m}/10)$, we have*

$$\mathbf{s}(\mu') \approx_{3k^2} \mathbf{s}(\mu).$$

Proof of Lemma 9.2.7. By Lemma 9.2.4, for any $e \in E$ we have that

$$\left(\frac{1}{s_e(\mu)^+ \cdot s_e(\mu')^+} + \frac{1}{s_e(\mu)^- \cdot s_e(\mu')^-} \right) (f_e(\mu') - f_e(\mu))^2 \leq 2k^2. \quad (9.42)$$

If $s_e(\mu')^+ = (1+c) \cdot s_e(\mu)^+$ for some $c \geq 0$, then

$$(f_e(\mu') - f_e(\mu))^2 = c^2 (s_e(\mu)^+)^2$$

and

$$s_e(\mu)^+ \cdot s_e(\mu')^+ = (1+c)(s_e(\mu)^+)^2,$$

so by (9.42) we have that $c \leq 3k^2$.

Similarly, if $s_e(\mu')^+ = (1+c)^{-1} \cdot s_e(\mu)^+$ for some $c \geq 0$, then

$$(f_e(\mu') - f_e(\mu))^2 = c^2 (s_e(\mu')^+)^2$$

and

$$s_e(\mu)^+ \cdot s_e(\mu')^+ = (1+c)(s_e(\mu')^+)^2,$$

so by (9.42) we have that $c \leq 3k^2$.

We have proved that $s_e(\mu')^+ \approx_{3k^2} s_e(\mu)^+$ and by symmetry we also have $s_e(\mu')^- \approx_{3k^2} s_e(\mu)^-$. \square

Proof of Lemma 4.3.5

Our goal is to keep track of how close \mathbf{f}^* remains to centrality (in ℓ_2 norm) and how close \mathbf{f} remains to \mathbf{f}^* in ℓ_∞ norm. From these two we can conclude that at all times \mathbf{f} is close in ℓ_∞ to the central flow. We first prove the following lemma, which bounds how the distance of \mathbf{f}^* to centrality (measured in energy of the residual) degrades when taking a progress step.

Lemma 9.2.8. *Let \mathbf{f}^* be a flow with slacks \mathbf{s}^* and resistances \mathbf{r}^* , and \mathbf{f} be a flow with slacks \mathbf{s} and resistances \mathbf{r} , where $\mathbf{s} \approx_{1+\varepsilon_{\text{solve}}} \mathbf{s}^*$ for some $\varepsilon_{\text{solve}} \in (0, 0.1)$. We define $\mathbf{f}'^* = \mathbf{f}^* + \varepsilon_{\text{step}} \tilde{\mathbf{f}}^*$ for some $\varepsilon_{\text{step}} \in (0, 0.1)$ (and the new slacks \mathbf{s}'^*), where*

$$\tilde{\mathbf{f}}^* = \delta g(\mathbf{s}) - \delta \mathbf{R}^{-1} \mathbf{B} (\mathbf{B}^\top \mathbf{R}^{-1} \mathbf{B})^+ \mathbf{B}^\top g(\mathbf{s}), \quad (9.43)$$

$\delta = \frac{1}{\sqrt{m}}$, and $g(\mathbf{s}) := \frac{\frac{1}{s^+} - \frac{1}{s^-}}{\mathbf{r}}$. If we let $\mathbf{h} = \frac{c}{\mu} + \frac{1}{s^{*+}} - \frac{1}{s^{*-}}$ and $\mathbf{h}' = \frac{c(1+\varepsilon_{\text{step}}\delta)}{\mu} + \frac{1}{s'^{*+}} - \frac{1}{s'^{-}}$ be the residuals of \mathbf{f}^* and \mathbf{f}'^* for some $\mu > 0$, then

$$\left\| \mathbf{C}^\top \mathbf{h}' \right\|_{\overline{\mathbf{H}}^+} \leq (1 + \varepsilon_{\text{step}} \delta) \left\| \mathbf{C}^\top \mathbf{h} \right\|_{\overline{\mathbf{H}}^+} + 5 \left\| \frac{\mathbf{r}^*}{\bar{\mathbf{r}}} \right\|_\infty^{1/2} \varepsilon_{\text{solve}} \cdot \varepsilon_{\text{step}} + 2 \left\| \frac{\mathbf{r}'^*}{\bar{\mathbf{r}}} \right\|_\infty^{1/2} \varepsilon_{\text{step}}^2,$$

where $\bar{\mathbf{r}}$ are some arbitrary resistances and $\overline{\mathbf{H}} = \mathbf{C}^\top \overline{\mathbf{R}} \mathbf{C}$.

Proof. Let $\boldsymbol{\rho}^+ = \varepsilon_{\text{step}} \tilde{\mathbf{f}}^* / \mathbf{s}^{*+}$ and $\boldsymbol{\rho}^- = -\varepsilon_{\text{step}} \tilde{\mathbf{f}}^* / \mathbf{s}^{*-}$. First of all, it is easy to see that

$$\begin{aligned} \|\boldsymbol{\rho}\|_2 &\leq \left\| \frac{\mathbf{s}}{\mathbf{s}^*} \right\|_\infty \left\| \frac{\mathbf{s}^*}{\mathbf{s}} \boldsymbol{\rho} \right\|_2 \\ &\leq \varepsilon_{\text{step}} (1 + \varepsilon_{\text{solve}}) \left\| \tilde{\mathbf{f}}^* \right\|_{\mathbf{r}, 2} \\ &= \varepsilon_{\text{step}} \delta (1 + \varepsilon_{\text{solve}}) \left\| \sqrt{\mathbf{r}} g(\mathbf{s}) - \mathbf{R}^{-1/2} \mathbf{B} (\mathbf{B}^\top \mathbf{R}^{-1} \mathbf{B})^+ \mathbf{B}^\top g(\mathbf{s}) \right\|_2 \\ &= \varepsilon_{\text{step}} \delta (1 + \varepsilon_{\text{solve}}) \left\| \left(\mathbf{I} - \mathbf{R}^{-1/2} \mathbf{B} (\mathbf{B}^\top \mathbf{R}^{-1} \mathbf{B})^+ \mathbf{B}^\top \mathbf{R}^{-1/2} \right) \sqrt{\mathbf{r}} g(\mathbf{s}) \right\|_2 \\ &\leq \varepsilon_{\text{step}} \delta (1 + \varepsilon_{\text{solve}}) \left\| \sqrt{\mathbf{r}} g(\mathbf{s}) \right\|_2 \\ &= \varepsilon_{\text{step}} \delta (1 + \varepsilon_{\text{solve}}) \left\| \frac{\frac{1}{s^+} - \frac{1}{s^-}}{\sqrt{\frac{1}{(s^+)^2} + \frac{1}{(s^-)^2}}} \right\|_2 \\ &\leq \varepsilon_{\text{step}} \delta (1 + \varepsilon_{\text{solve}}) \sqrt{m} \\ &= \varepsilon_{\text{step}} (1 + \varepsilon_{\text{solve}}). \end{aligned}$$

We bound the energy to route the residual of \mathbf{f}'^* as

$$\begin{aligned}
& \left\| \mathbf{C}^\top \mathbf{h}' \right\|_{\overline{H}^+} \\
&= \left\| \mathbf{C}^\top \mathbf{h} + \mathbf{C}^\top \left(\frac{\varepsilon_{\text{step}} \delta \mathbf{c}}{\mu} + \frac{1}{\mathbf{s}'^{*+}} - \frac{1}{\mathbf{s}^{*+}} - \frac{1}{\mathbf{s}'^{*-}} + \frac{1}{\mathbf{s}^{*-}} \right) \right\|_{\overline{H}^+} \\
&= \left\| \mathbf{C}^\top \mathbf{h} + \mathbf{C}^\top \left(\frac{\varepsilon_{\text{step}} \delta \mathbf{c}}{\mu} + \frac{\boldsymbol{\rho}^+}{\mathbf{s}'^{*+}} - \frac{\boldsymbol{\rho}^-}{\mathbf{s}'^{*-}} \right) \right\|_{\overline{H}^+} \\
&= \left\| \mathbf{C}^\top \mathbf{h} + \mathbf{C}^\top \left(\frac{\varepsilon_{\text{step}} \delta \mathbf{c}}{\mu} + \frac{\boldsymbol{\rho}^+}{\mathbf{s}^{*+}} - \frac{\boldsymbol{\rho}^-}{\mathbf{s}^{*-}} \right) + \mathbf{C}^\top \left(\frac{(\boldsymbol{\rho}^+)^2}{\mathbf{s}'^{*+}} - \frac{(\boldsymbol{\rho}^-)^2}{\mathbf{s}'^{*-}} \right) \right\|_{\overline{H}^+}.
\end{aligned}$$

Now, using (9.43) we get that $\widetilde{\mathbf{r}}\mathbf{f}^* = \delta \mathbf{r}g(\mathbf{s}) - \delta \mathbf{B}(\mathbf{B}^\top \mathbf{R}^{-1} \mathbf{B})^+ \mathbf{B}^\top g(\mathbf{s})$ and so $\mathbf{C}^\top (\widetilde{\mathbf{r}}\mathbf{f}^*) = \delta \mathbf{C}^\top (\mathbf{r}g(\mathbf{s}))$, which follows by the fact that for any i , $\mathbf{1}_i^\top \mathbf{C}^\top \mathbf{B} = (\mathbf{B}^\top \mathbf{C} \mathbf{1}_i)^\top = \mathbf{0}$, since $\mathbf{C} \mathbf{1}_i$ is a circulation by definition of \mathbf{C} . As $\widetilde{\mathbf{r}}\mathbf{f}^* = \left(\frac{1}{(\mathbf{s}^+)^2} + \frac{1}{(\mathbf{s}^-)^2} \right) \widetilde{\mathbf{f}}^* = \varepsilon_{\text{step}}^{-1} \frac{\mathbf{s}^{*+}}{(\mathbf{s}^+)^2} \boldsymbol{\rho}^+ - \varepsilon_{\text{step}}^{-1} \frac{\mathbf{s}^{*-}}{(\mathbf{s}^-)^2} \boldsymbol{\rho}^-$, we have $\varepsilon_{\text{step}} \delta \mathbf{C}^\top (\mathbf{r}g(\mathbf{s})) = \mathbf{C}^\top \left(\frac{\mathbf{s}^{*+}}{(\mathbf{s}^+)^2} \boldsymbol{\rho}^+ - \frac{\mathbf{s}^{*-}}{(\mathbf{s}^-)^2} \boldsymbol{\rho}^- \right)$ and so

$$\begin{aligned}
& \left\| \mathbf{C}^\top \mathbf{h} + \mathbf{C}^\top \left(\frac{\varepsilon_{\text{step}} \delta \mathbf{c}}{\mu} + \frac{\boldsymbol{\rho}^+}{\mathbf{s}^{*+}} - \frac{\boldsymbol{\rho}^-}{\mathbf{s}^{*-}} \right) + \mathbf{C}^\top \left(\frac{(\boldsymbol{\rho}^+)^2}{\mathbf{s}'^{*+}} - \frac{(\boldsymbol{\rho}^-)^2}{\mathbf{s}'^{*-}} \right) \right\|_{\overline{H}^+} \\
&= \left\| \mathbf{C}^\top \mathbf{h} + \mathbf{C}^\top \left(\frac{\varepsilon_{\text{step}} \delta \mathbf{c}}{\mu} + \varepsilon_{\text{step}} \delta \mathbf{r}g(\mathbf{s}) - \frac{\mathbf{s}^{*+}}{(\mathbf{s}^+)^2} \boldsymbol{\rho}^+ + \frac{\mathbf{s}^{*-}}{(\mathbf{s}^-)^2} \boldsymbol{\rho}^- + \frac{\boldsymbol{\rho}^+}{\mathbf{s}^{*+}} - \frac{\boldsymbol{\rho}^-}{\mathbf{s}^{*-}} \right) + \mathbf{C}^\top \left(\frac{(\boldsymbol{\rho}^+)^2}{\mathbf{s}'^{*+}} - \frac{(\boldsymbol{\rho}^-)^2}{\mathbf{s}'^{*-}} \right) \right\|_{\overline{H}^+} \\
&= \left\| \mathbf{C}^\top \mathbf{h} + \varepsilon_{\text{step}} \delta \mathbf{C}^\top \left(\frac{\mathbf{c}}{\mu} + \mathbf{r}g(\mathbf{s}) \right) + \mathbf{C}^\top \left(\left(\mathbf{1} - \left(\frac{\mathbf{s}^{*+}}{\mathbf{s}^+} \right)^2 \right) \frac{\boldsymbol{\rho}^+}{\mathbf{s}^{*+}} - \left(\mathbf{1} - \left(\frac{\mathbf{s}^{*-}}{\mathbf{s}^-} \right)^2 \right) \frac{\boldsymbol{\rho}^-}{\mathbf{s}^{*-}} \right) \right. \\
&\quad \left. + \mathbf{C}^\top \left(\frac{(\boldsymbol{\rho}^+)^2}{\mathbf{s}'^{*+}} - \frac{(\boldsymbol{\rho}^-)^2}{\mathbf{s}'^{*-}} \right) \right\|_{\overline{H}^+} \\
&\leq (1 + \varepsilon_{\text{step}} \delta) \left\| \mathbf{C}^\top \mathbf{h} \right\|_{\overline{H}^+} + 5 \left\| \frac{\mathbf{r}^*}{\bar{\mathbf{r}}} \right\|_{\infty}^{1/2} \varepsilon_{\text{solve}} \cdot \varepsilon_{\text{step}} + 2 \left\| \frac{\mathbf{r}'^*}{\bar{\mathbf{r}}} \right\|_{\infty}^{1/2} \varepsilon_{\text{step}}^2
\end{aligned}$$

where we have used the triangle inequality, the fact that

$$\begin{aligned}
& \varepsilon_{\text{step}} \delta \left\| \mathbf{C}^\top \left(\frac{\mathbf{c}}{\mu} + \mathbf{r}g(\mathbf{s}) \right) \right\|_{\overline{\mathbf{H}}^+} \\
&= \varepsilon_{\text{step}} \delta \left\| \mathbf{C}^\top \left(\frac{\mathbf{c}}{\mu} + \frac{1}{\mathbf{s}^+} - \frac{1}{\mathbf{s}^-} \right) \right\|_{\overline{\mathbf{H}}^+} \\
&\leq \varepsilon_{\text{step}} \delta \left\| \mathbf{C}^\top \left(\frac{\mathbf{c}}{\mu} + \frac{1}{\mathbf{s}^{*+}} - \frac{1}{\mathbf{s}^{*-}} \right) \right\|_{\overline{\mathbf{H}}^+} + \varepsilon_{\text{step}} \delta \left\| \mathbf{C}^\top \left(\frac{1}{\mathbf{s}^+} - \frac{1}{\mathbf{s}^{*+}} - \frac{1}{\mathbf{s}^-} + \frac{1}{\mathbf{s}^{*-}} \right) \right\|_{\overline{\mathbf{H}}^+} \\
&\leq \varepsilon_{\text{step}} \delta \left\| \mathbf{C}^\top \left(\frac{\mathbf{c}}{\mu} + \frac{1}{\mathbf{s}^{*+}} - \frac{1}{\mathbf{s}^{*-}} \right) \right\|_{\overline{\mathbf{H}}^+} + \varepsilon_{\text{step}} \delta \left\| \frac{\mathbf{r}^*}{\bar{\mathbf{r}}} \right\|_\infty^{1/2} \left\| \mathbf{C}^\top \left(\frac{1}{\mathbf{s}^+} - \frac{1}{\mathbf{s}^{*+}} - \frac{1}{\mathbf{s}^-} + \frac{1}{\mathbf{s}^{*-}} \right) \right\|_{(\mathbf{C}^\top \mathbf{R}^* \mathbf{C})^+} \\
&\leq \varepsilon_{\text{step}} \delta \left\| \mathbf{C}^\top \mathbf{h} \right\|_{\overline{\mathbf{H}}^+} + \varepsilon_{\text{step}} \delta \left\| \frac{\mathbf{r}^*}{\bar{\mathbf{r}}} \right\|_\infty^{1/2} \left\| \frac{\frac{1}{\mathbf{s}^+} - \frac{1}{\mathbf{s}^{*+}} - \frac{1}{\mathbf{s}^-} + \frac{1}{\mathbf{s}^{*-}}}{\left(\frac{1}{(\mathbf{s}^{*+})^2} + \frac{1}{(\mathbf{s}^{*-})^2} \right)^{1/2}} \right\|_2 \\
&\leq \varepsilon_{\text{step}} \delta \left\| \mathbf{C}^\top \mathbf{h} \right\|_{\overline{\mathbf{H}}^+} + \varepsilon_{\text{step}} \delta \left\| \frac{\mathbf{r}^*}{\bar{\mathbf{r}}} \right\|_\infty^{1/2} \left(\left\| \frac{\frac{1}{\mathbf{s}^+} - \frac{1}{\mathbf{s}^{*+}}}{\left(\frac{1}{(\mathbf{s}^{*+})^2} \right)^{1/2}} \right\|_2 + \left\| \frac{\frac{1}{\mathbf{s}^-} - \frac{1}{\mathbf{s}^{*-}}}{\left(\frac{1}{(\mathbf{s}^{*-})^2} \right)^{1/2}} \right\|_2 \right) \\
&= \varepsilon_{\text{step}} \delta \left\| \mathbf{C}^\top \mathbf{h} \right\|_{\overline{\mathbf{H}}^+} + \varepsilon_{\text{step}} \delta \left\| \frac{\mathbf{r}^*}{\bar{\mathbf{r}}} \right\|_\infty^{1/2} \left(\left\| \frac{\mathbf{s}^{*+}}{\mathbf{s}^+} - \mathbf{1} \right\|_2 + \left\| \frac{\mathbf{s}^{*-}}{\mathbf{s}^-} - \mathbf{1} \right\|_2 \right) \\
&\leq \varepsilon_{\text{step}} \delta \left\| \mathbf{C}^\top \mathbf{h} \right\|_{\overline{\mathbf{H}}^+} + \varepsilon_{\text{step}} \delta \left\| \frac{\mathbf{r}^*}{\bar{\mathbf{r}}} \right\|_\infty^{1/2} 2\varepsilon_{\text{solve}} \sqrt{m} \text{ (as } \mathbf{s} \approx_{1+\varepsilon_{\text{solve}}} \mathbf{s}^*) \\
&= \varepsilon_{\text{step}} \delta \left\| \mathbf{C}^\top \mathbf{h} \right\|_{\overline{\mathbf{H}}^+} + 2 \left\| \frac{\mathbf{r}^*}{\bar{\mathbf{r}}} \right\|_\infty^{1/2} \varepsilon_{\text{step}} \varepsilon_{\text{solve}}
\end{aligned}$$

and similarly, using $\left\| 1 - \left(\frac{\mathbf{s}^*}{\mathbf{s}} \right)^2 \right\|_\infty \leq \varepsilon_{\text{solve}}(2 + \varepsilon_{\text{solve}})$, the fact that

$$\begin{aligned}
& \left\| \mathbf{C}^\top \left(\left(\mathbf{1} - \left(\frac{\mathbf{s}^{*+}}{\mathbf{s}^+} \right)^2 \right) \frac{\boldsymbol{\rho}^+}{\mathbf{s}^{*+}} - \left(\mathbf{1} - \left(\frac{\mathbf{s}^{*-}}{\mathbf{s}^-} \right)^2 \right) \frac{\boldsymbol{\rho}^-}{\mathbf{s}^{*-}} \right) + \mathbf{C}^\top \left(\frac{(\boldsymbol{\rho}^+)^2}{\mathbf{s}^{*+}} + \frac{(\boldsymbol{\rho}^-)^2}{\mathbf{s}^{*-}} \right) \right\|_{\overline{\mathbf{H}}^+} \\
&\leq \left\| \frac{\mathbf{r}^*}{\bar{\mathbf{r}}} \right\|_\infty^{1/2} \varepsilon_{\text{solve}}(2 + \varepsilon_{\text{solve}}) \|\boldsymbol{\rho}\|_2 + \left\| \frac{\mathbf{r}^*}{\bar{\mathbf{r}}} \right\|_\infty^{1/2} \|\boldsymbol{\rho}\|_4^2 \\
&\leq \left\| \frac{\mathbf{r}^*}{\bar{\mathbf{r}}} \right\|_\infty^{1/2} \varepsilon_{\text{step}} \varepsilon_{\text{solve}}(2 + \varepsilon_{\text{solve}})(1 + \varepsilon_{\text{solve}}) + \left\| \frac{\mathbf{r}^*}{\bar{\mathbf{r}}} \right\|_\infty^{1/2} \varepsilon_{\text{step}}^2 (1 + \varepsilon_{\text{solve}})^2 \\
&\leq 3 \left\| \frac{\mathbf{r}^*}{\bar{\mathbf{r}}} \right\|_\infty^{1/2} \varepsilon_{\text{step}} \varepsilon_{\text{solve}} + 2 \left\| \frac{\mathbf{r}^*}{\bar{\mathbf{r}}} \right\|_\infty^{1/2} \varepsilon_{\text{step}}^2.
\end{aligned}$$

□

We will also use the following lemma, which is standard [10].

Lemma 9.2.9 (Small residual implies ℓ_∞ closeness). *Given a flow $\mathbf{f} = \mathbf{f}^0 + \mathbf{C}\mathbf{x}$ with slacks \mathbf{s} and resistances \mathbf{r} , if $\left\| \mathbf{C}^\top \left(\frac{\mathbf{c}}{\mu} + \frac{1}{\mathbf{s}^+} - \frac{1}{\mathbf{s}^-} \right) \right\|_{(\mathbf{C}^\top \mathbf{R} \mathbf{C})^+} \leq 1/1000$ then \mathbf{f} is $(\mu, 1.01)$ -central.*

Applying Lemma 9.2.8 for $T = \frac{k}{\varepsilon_{\text{step}}}$ iterations, we get the lemma below, which measures the closeness of \mathbf{f}^* to the central path in ℓ_2 after T iterations.

Lemma 9.2.10 (Centrality of \mathbf{f}^*). Let $\mathbf{f}^{*1}, \dots, \mathbf{f}^{*T+1}$ be flows with slacks $\mathbf{s}^{*1}, \dots, \mathbf{s}^{*T+1}$ and resistances $\mathbf{r}^{*1}, \dots, \mathbf{r}^{*T+1}$, and $\mathbf{f}^1, \dots, \mathbf{f}^{T+1}$ be flows with slacks $\mathbf{s}^1, \dots, \mathbf{s}^{T+1}$ and resistances $\mathbf{r}^1, \dots, \mathbf{r}^{T+1}$, such that $\mathbf{s}^t \approx_{1+\varepsilon_{\text{solve}}} \mathbf{s}^{*t}$ for all $t \in [T]$, where $T = \frac{k}{\varepsilon_{\text{step}}}$ for some $k \leq \sqrt{m}/10$, $\varepsilon_{\text{step}} \in (0, 0.1)$ and $\varepsilon_{\text{solve}} \in (0, 0.1)$. Additionally, we have that

- \mathbf{f}^{*1} is μ -central
- For all $t \in [T]$, $\mathbf{f}^{*t+1} = \mathbf{f}^{*t} + \varepsilon_{\text{step}} \cdot \tilde{\mathbf{f}}^t$, where

$$\left\| \sqrt{\mathbf{r}^t} \left(\tilde{\mathbf{f}}^{*t} - \tilde{\mathbf{f}}^t \right) \right\|_{\infty} \leq \varepsilon,$$

$$\tilde{\mathbf{f}}^{*t} = \delta g(\mathbf{s}^t) - \delta (\mathbf{R}^t)^{-1} \mathbf{B} \left(\mathbf{B}^{\top} (\mathbf{R}^t)^{-1} \mathbf{B} \right)^+ \mathbf{B}^{\top} g(\mathbf{s}^t)$$

$$\text{and } \delta = \frac{1}{\sqrt{m}}.$$

Then, \mathbf{f}^{*T+1} is $(\mu/(1+\varepsilon_{\text{step}}\delta)^T, 1.01)$ -central, as long as we set $\varepsilon_{\text{step}} \leq 10^{-5}k^{-3}$ and $\varepsilon_{\text{solve}} \leq 10^{-5}k^{-3}$.

Proof. For all $t \in [T+1]$, we denote the residual of \mathbf{f}^{*t} as $\mathbf{h}^t = \frac{c(1+\varepsilon_{\text{step}}\delta)^{t-1}}{\mu} + \frac{1}{\mathbf{s}^{+,*t}} - \frac{1}{\mathbf{s}^{-,*t}}$. Note that $\mathbf{C}^{\top} \mathbf{h}^1 = \mathbf{0}$ as \mathbf{f}^{*1} is μ -central.

We assume that the statement of the lemma is not true, and let \hat{T} be the smallest $t \in [T+1]$ such that \mathbf{f}^{*t} is not $(\mu/(1+\varepsilon_{\text{step}}\delta)^{t-1}, 1.01)$ -central. Obviously $\hat{T} > 1$. This means that \mathbf{f}^{*t} is $(\mu/(1+\varepsilon_{\text{step}}\delta)^{t-1}, 1.01)$ -central for all $t \in [\hat{T}-1]$, i.e. $\mathbf{s}^{*t} \approx_{1.01} \mathbf{s} (\mu/(1+\varepsilon_{\text{step}}\delta)^{t-1})$.

Also, note that by Lemma 9.2.7 about slack stability, and since $(1+\varepsilon_{\text{step}}\delta)^{|\hat{T}-t|} \leq (1+\delta)^{1.1k}$, we have $\mathbf{s} (\mu/(1+\varepsilon_{\text{step}}\delta)^{t-1}) \approx_{3.7k^2} \mathbf{s} (\mu/(1+\varepsilon_{\text{step}}\delta)^{\hat{T}-1})$ for all $t \in [T+1]$. Additionally, note that, as shown in proof of Lemma 9.2.8, we have

$$\left\| \tilde{\mathbf{f}}^{*\hat{T}-1} \right\|_{\mathbf{r}^{\hat{T}-1, \infty}} \leq \left\| \tilde{\mathbf{f}}^{*\hat{T}-1} \right\|_{\mathbf{r}^{\hat{T}-1, 2}} \leq 1,$$

so

$$\begin{aligned} \left\| \frac{\mathbf{s}^{*\hat{T}}}{\mathbf{s}^{*\hat{T}-1}} - \mathbf{1} \right\|_{\infty} &= \varepsilon_{\text{step}} \left\| \frac{\tilde{\mathbf{f}}^{*\hat{T}-1}}{\mathbf{s}^{*\hat{T}-1}} \right\|_{\infty} \\ &\leq \varepsilon_{\text{step}} (1 + \varepsilon_{\text{solve}}) \left\| \sqrt{\mathbf{r}^{\hat{T}-1}} \tilde{\mathbf{f}}^{*\hat{T}-1} \right\|_{\infty} \\ &\leq \varepsilon_{\text{step}} (1 + \varepsilon_{\text{solve}}) \left(\left\| \sqrt{\mathbf{r}^{\hat{T}-1}} \tilde{\mathbf{f}}^{*\hat{T}-1} \right\|_{\infty} + \varepsilon \right) \\ &\leq \varepsilon_{\text{step}} (1 + \varepsilon_{\text{solve}}) (1 + \varepsilon) \\ &\leq 1.3 \varepsilon_{\text{step}}. \end{aligned}$$

From this we conclude that $\mathbf{s}^{*\hat{T}} \approx_{1+2.6\varepsilon_{\text{step}}} \mathbf{s}^{*\hat{T}-1}$, and from the previous discussion we get that

$$\mathbf{s}^{*\hat{T}} \approx_{1+2.6\varepsilon_{\text{step}}} \mathbf{s}^{*\hat{T}-1} \approx_{1.01} \mathbf{s} (\mu/(1+\varepsilon_{\text{step}}\delta)^{\hat{T}-1}) \approx_{3.7k^2} \mathbf{s} (\mu/(1+\varepsilon_{\text{step}}\delta)^{t-1}) \approx_{1.01} \mathbf{s}^{*t},$$

so $\mathbf{s}^{*\hat{T}} \approx_{4k^2} \mathbf{s}^{*t}$ for all $t \in [\hat{T}-1]$.

On the other hand, if we apply Lemma 9.2.8 $\widehat{T} - 1$ times with $\bar{\mathbf{r}} = \mathbf{r}^{*\widehat{T}}$, we get

$$\begin{aligned}
& \left\| \mathbf{C}^\top \mathbf{h}^{\widehat{T}} \right\|_{(\mathbf{C}^\top \mathbf{R}^{\widehat{T}} \mathbf{C})^+} \\
&= \left\| \mathbf{C}^\top \mathbf{h}^{\widehat{T}} \right\|_{\bar{\mathbf{H}}^+} \\
&\leq (1 + \varepsilon_{\text{step}} \delta) \left\| \mathbf{C}^\top \mathbf{h}^{\widehat{T}-1} \right\|_{\bar{\mathbf{H}}^+} + 5 \left\| \frac{\mathbf{r}^{*\widehat{T}-1}}{\bar{\mathbf{r}}} \right\|_{\infty}^{1/2} \varepsilon_{\text{step}} \cdot \varepsilon_{\text{solve}} + 2 \left\| \frac{\mathbf{r}^{*\widehat{T}}}{\bar{\mathbf{r}}} \right\|_{\infty}^{1/2} \varepsilon_{\text{step}}^2 \\
&\dots \\
&\leq 5 \sum_{t=1}^{\widehat{T}-1} (1 + \varepsilon_{\text{step}} \delta)^{\widehat{T}-t-1} \left\| \frac{\mathbf{r}^{*t}}{\mathbf{r}^{*\widehat{T}}} \right\|_{\infty}^{1/2} \varepsilon_{\text{step}} \cdot \varepsilon_{\text{solve}} + 2 \sum_{t=1}^{\widehat{T}-1} (1 + \varepsilon_{\text{step}} \delta)^{\widehat{T}-t-1} \left\| \frac{\mathbf{r}^{*t+1}}{\mathbf{r}^{*\widehat{T}}} \right\|_{\infty}^{1/2} \varepsilon_{\text{step}}^2 \\
&\leq 6T \max_{t \in [\widehat{T}-1]} \left\| \frac{\mathbf{r}^{*t}}{\mathbf{r}^{*\widehat{T}}} \right\|_{\infty}^{1/2} \varepsilon_{\text{step}} \cdot \varepsilon_{\text{solve}} + 2.4T \max_{t \in [\widehat{T}-1]} \left\| \frac{\mathbf{r}^{*t+1}}{\mathbf{r}^{*\widehat{T}}} \right\|_{\infty}^{1/2} \varepsilon_{\text{step}}^2 \\
&\leq 24Tk^2 \varepsilon_{\text{step}} \cdot \varepsilon_{\text{solve}} + 10Tk^2 \varepsilon_{\text{step}}^2 \\
&= 24k^3 \varepsilon_{\text{solve}} + 10k^3 \varepsilon_{\text{step}} \\
&\leq 1/1000,
\end{aligned}$$

where we used the fact that $(1 + \varepsilon_{\text{step}} \delta)^{\widehat{T}} \leq e^{\varepsilon_{\text{step}} \delta T} = e^{\delta k} \leq 1.2$ and our setting of $\varepsilon_{\text{solve}} \leq 10^{-5} k^{-3}$ and $\varepsilon_{\text{step}} \leq 10^{-5} k^{-3}$. By Lemma 9.2.9 this implies that $\mathbf{f}^{*\widehat{T}}$ is $(\mu/(1 + \varepsilon_{\text{step}} \delta)^{\widehat{T}-1}, 1.01)$ -central, a contradiction. \square

We are now ready to prove the following lemma, which is the goal of this section:

Lemma 4.3.5. *Let $\mathbf{f}^1, \dots, \mathbf{f}^{T+1}$ be flows with slacks \mathbf{s}^t and resistances \mathbf{r}^t for $t \in [T+1]$, where $T = \frac{k}{\varepsilon_{\text{step}}}$ for some $k \leq \sqrt{m}/10$ and $\varepsilon_{\text{step}} = 10^{-5} k^{-3}$, such that*

- \mathbf{f}^1 is $(\mu, 1 + \varepsilon_{\text{solve}}/8)$ -central for $\varepsilon_{\text{solve}} = 10^{-5} k^{-3}$
- For all $t \in [T]$, $\mathbf{f}^{t+1} = \begin{cases} \mathbf{f}(\mu) + \varepsilon_{\text{step}} \sum_{i=1}^t \tilde{\mathbf{f}}^i & \text{if } \exists i \in [t] : \tilde{\mathbf{f}}^i \neq \mathbf{0} \\ \mathbf{f}^1 & \text{otherwise} \end{cases}$, where

$$\tilde{\mathbf{f}}^{*t} = \delta g(\mathbf{s}^t) - \delta (\mathbf{R}^t)^{-1} \mathbf{B} \left(\mathbf{B}^\top (\mathbf{R}^t)^{-1} \mathbf{B} \right)^+ \mathbf{B}^\top g(\mathbf{s}^t)$$

for $\delta = \frac{1}{\sqrt{m}}$ and

$$\left\| \sqrt{\mathbf{r}^t} \left(\tilde{\mathbf{f}}^{*t} - \tilde{\mathbf{f}}^t \right) \right\|_{\infty} \leq \varepsilon$$

for $\varepsilon = 10^{-6} k^{-6}$.

Then, setting $\varepsilon_{\text{step}} = \varepsilon_{\text{solve}} = 10^{-5} k^{-3}$ and $\varepsilon = 10^{-6} k^{-6}$ we get that $\mathbf{s}^{T+1} \approx_{1.1} \mathbf{s} \left(\mu / (1 + \varepsilon_{\text{step}} \delta)^{k \varepsilon_{\text{step}}^{-1}} \right)$.

Proof. We set $\mathbf{f}^{*1} = \mathbf{f}(\mu)$ and for each $t \in [T]$,

$$\mathbf{f}^{*t+1} = \mathbf{f}^{*t} + \varepsilon_{\text{step}} \tilde{\mathbf{f}}^{*t},$$

and the corresponding slacks \mathbf{s}^{*t} and resistances \mathbf{r}^{*t} . Let \hat{T} be the first $t \in [T + 1]$ such that $\mathbf{s}^{\hat{T}} \not\approx_{1+\varepsilon_{\text{solve}}} \mathbf{s}^{*\hat{T}}$. Obviously $t > 1$ as $\mathbf{s}^1 \approx_{1+\varepsilon_{\text{solve}}/8} \mathbf{s}(\mu) = \mathbf{s}^{*1}$.

Now, for all $t \in [T]$ we have

$$\left\| \sqrt{\mathbf{r}^t} \left(\tilde{\mathbf{f}}^{*t} - \tilde{\mathbf{f}}^t \right) \right\|_{\infty} \leq \varepsilon.$$

Fix some $e \in E$. If $\tilde{f}_e^t = \mathbf{0}$ for all $t \in [\hat{T} - 1]$, then we have $\sqrt{r_e^{\hat{T}}} |\tilde{f}_e^{*\hat{T}}| = \sqrt{r_e^t} |\tilde{f}_e^{*t}| \leq \varepsilon$ for all such t . This means that

$$\begin{aligned} \sqrt{r_e^{\hat{T}}} |f_e^{*\hat{T}} - f_e^{\hat{T}}| &\leq \sqrt{r_e^{\hat{T}}} |f_e^{*\hat{T}} - f_e^{*1}| + \sqrt{r_e^{\hat{T}}} |f_e^{*1} - f_e^{\hat{T}}| \\ &= \sqrt{r_e^{\hat{T}}} |f_e^{*\hat{T}} - f_e^{*1}| + \sqrt{r_e^{\hat{T}}} |f_e^{*1} - f_e^1| \\ &\leq \varepsilon_{\text{step}} \sqrt{r_e^{\hat{T}}} \sum_{t=1}^{\hat{T}-1} |\tilde{f}_e^{*t}| + \sqrt{r_e^{\hat{T}}} |f_e^{*1} - f_e^1| \\ &\leq \hat{T} \varepsilon_{\text{step}} \varepsilon + \sqrt{r_e^{\hat{T}}} |f_e^{*1} - f_e^1| \\ &\leq k\varepsilon + \sqrt{2}\varepsilon_{\text{solve}}/8 \\ &\leq \varepsilon_{\text{solve}}/2, \end{aligned}$$

as long as $\varepsilon \leq \varepsilon_{\text{solve}}/(2k) = O(1/k^4)$. In the second to last inequality we used Lemma 9.2.5.

Otherwise, there exists $t \in [\hat{T} - 1]$ such that $\tilde{f}_e^t \neq \mathbf{0}$, and by definition $f_e^{\hat{T}} = f_e(\mu) + \varepsilon_{\text{step}} \sum_{t=1}^{\hat{T}-1} \tilde{f}_e^t$, so

$$\begin{aligned} \sqrt{r_e^{*\hat{T}}} |f_e^{*\hat{T}} - f_e^{\hat{T}}| &\leq \sqrt{r_e^{*\hat{T}}} |f_e^{*1} - f_e(\mu)| + \varepsilon_{\text{step}} \sum_{t=1}^{\hat{T}-1} \sqrt{r_e^{*\hat{T}}} |\tilde{f}_e^{*t} - \tilde{f}_e^t| \\ &\leq 3k^2 \varepsilon_{\text{step}} \sum_{t=1}^{\hat{T}-1} \sqrt{r_e^{*t}} |\tilde{f}_e^{*t} - \tilde{f}_e^t| \\ &\leq 3k^2 \varepsilon_{\text{step}} (1 + \varepsilon_{\text{solve}}) \sum_{t=1}^{\hat{T}-1} \sqrt{r_e^t} |\tilde{f}_e^{*t} - \tilde{f}_e^t| \\ &\leq 3k^2 \varepsilon_{\text{step}} (1 + \varepsilon_{\text{solve}}) T \varepsilon \\ &\leq 4k^3 \varepsilon, \end{aligned}$$

where we have used Lemma 9.2.7 and the fact that $\mathbf{s}^t \approx_{1+\varepsilon_{\text{solve}}} \mathbf{s}^{*t}$ for all $t \in [\hat{T} - 1]$ which also implies that $\sqrt{\mathbf{r}^t} \approx_{1+\varepsilon_{\text{solve}}} \sqrt{\mathbf{r}^{*t}}$. Setting $\varepsilon = \frac{\varepsilon_{\text{solve}}}{8k^3} = O\left(\frac{1}{k^6}\right)$, this becomes $\leq \varepsilon_{\text{solve}}/2$.

Therefore we have proved that $\left\| \sqrt{\mathbf{r}^{*\hat{T}}} \left(\mathbf{f}^{*\hat{T}} - \mathbf{f}^{\hat{T}} \right) \right\|_{\infty} \leq \varepsilon_{\text{solve}}/2$, and so $\mathbf{s}^{*\hat{T}} \approx_{1+\varepsilon_{\text{solve}}} \mathbf{s}^{\hat{T}}$, a contradiction. Therefore we conclude that $\mathbf{s}^t \approx_{1+\varepsilon_{\text{solve}}} \mathbf{s}^{*t}$ for all $t \in [T + 1]$.

Now, as long as $\varepsilon_{\text{step}}, \varepsilon_{\text{solve}} \leq 10^{-5}k^{-3}$, we can apply Lemma 9.2.10, which guarantees that $\mathbf{s}^{*T+1} \approx_{1.01} \mathbf{s}(\mu/(1 + \varepsilon_{\text{step}}\delta)^T)$, and so $\mathbf{s}^{T+1} \approx_{1.1} \mathbf{s}(\mu/(1 + \varepsilon_{\text{step}}\delta)^T)$. Therefore we set $\varepsilon_{\text{step}} = \varepsilon_{\text{solve}} = 10^{-5}k^{-3}$ and $\varepsilon = 10^{-6}k^{-6} \leq \frac{\varepsilon_{\text{solve}}}{8k^3}$. \square

Proof of Lemma 4.3.11

Proof. We will apply Lemma 4.3.5 with \mathbf{f}^1 being the flow corresponding to the resistances $\mathcal{L}.\mathbf{r} = \mathcal{C}.\mathbf{r}$, and $T = k\varepsilon_{\text{step}}^{-1}$. Note that it is important to maintain the invariant $\mathcal{L}.\mathbf{r} = \mathcal{C}.\mathbf{r}$ throughout the algorithm so that both data structures correspond to the same electrical flow problem. For each $t \in [T]$, for the t -th iteration, Lemma 4.3.5 requires an estimate $\tilde{\mathbf{f}}^t$ such that $\left\| \sqrt{r^t} (\tilde{\mathbf{f}}^{*t} - \tilde{\mathbf{f}}^t) \right\|_{\infty} \leq \varepsilon$, where

$$\tilde{\mathbf{f}}^{*t} = \delta g(\mathbf{s}^t) - \delta (\mathbf{R}^t)^{-1} \mathbf{B} \left(\mathbf{B}^{\top} (\mathbf{R}^t)^{-1} \mathbf{B} \right)^+ \mathbf{B}^{\top} g(\mathbf{s}^t)$$

and $\delta = 1/\sqrt{m}$.

We claim that such an estimate can be computed for all t by using \mathcal{L} and \mathcal{C} . We apply the following process for each $t \in [T]$:

- Let Z be the edge set returned by $\mathcal{L}.\text{SOLVE}()$.
- Call $\mathcal{C}.\text{CHECK}(e)$ for each $e \in Z$ to obtain flow values \tilde{f}_e^t .
- Compute \mathbf{f}^t and its slacks \mathbf{s}^{t+1} and resistances \mathbf{r}^{t+1} as in Lemma 4.3.5, i.e.

$$f_e^{t+1} = \begin{cases} f_e(\mu) + \varepsilon_{\text{step}} \sum_{i=1}^t \tilde{f}_e^i & \text{if } \exists i \in [t] : \tilde{f}_e^i \neq 0 \\ f_e^1 & \text{otherwise} \end{cases}.$$

This can be computed in $O(|Z|)$ by adding either $\varepsilon_{\text{step}} \tilde{f}_e^i$ or $f_e(\mu) - f_e^1 + \varepsilon_{\text{step}} \tilde{f}_e^i$ to f_e^t for each $e \in Z$.

- Call $\mathcal{L}.\text{UPDATE}(e, \mathbf{f}^{t+1})$ and $\mathcal{C}.\text{UPDATE}(e, \mathbf{f}^{t+1})$ for all e in the support of $\tilde{\mathbf{f}}^t$. Note that $\mathcal{L}.\text{UPDATE}$ works as long as

$$r_e^{\max} / \alpha \leq r_e^{t+1} \leq \alpha \cdot r_e^{\min},$$

where r_e^{\max} , r_e^{\min} are the maximum and minimum values of $\mathcal{L}.r_e$ since the last call to $\mathcal{L}.\text{BATCHUPDATE}$. After this, we have $\mathcal{L}.\mathbf{r} = \mathcal{C}.\mathbf{r} = \mathbf{r}^{t+1}$.

In the above process, when $\mathcal{L}.\text{SOLVE}()$ is called we have $\mathcal{L}.\mathbf{r} = \mathbf{r}^t$ (for $t = 1$ this is true because $\mathcal{L}.\mathbf{r}$ are the resistances corresponding to \mathbf{f}^1). By the $(\alpha, \beta, \varepsilon/2)$ -LOCATOR guarantees in Definition 4.3.7, with high probability Z contains all the edges e such that $\sqrt{r_e^t} \left| \tilde{f}_e^{*t} \right| \geq \varepsilon/2$. Now, for each $e \in Z$, $\mathcal{C}.\text{CHECK}(e)$ returns a flow value \tilde{f}_e^t such that:

- $\sqrt{r_e^t} \left| \tilde{f}_e^t - \tilde{f}_e^{*t} \right| \leq \varepsilon$
- if $\sqrt{r_e^t} \left| \tilde{f}_e^{*t} \right| < \varepsilon/2$, then $\tilde{f}_e^t = 0$.

Therefore, the condition that

$$\sqrt{r^t} \left\| \tilde{\mathbf{f}}^t - \tilde{\mathbf{f}}^{*t} \right\|_{\infty} \leq \varepsilon$$

is satisfied. Additionally $\tilde{\mathbf{f}}^t$ is independent of the randomness of \mathcal{L} , because (the distribution of) $\tilde{\mathbf{f}}^t$ would be the same if $\mathcal{C}.\text{CHECK}$ was run for **all** edges e .

It remains to show that the LOCATOR requirement

$$\mathbf{r}^{\max} / \alpha \leq \mathbf{r}^{t+1} \leq \alpha \cdot \mathbf{r}^{\min}$$

is satisfied. Consider the minimum value of t for which this is not satisfied. By Lemma 4.3.5, we have that

$$\mathbf{s}^{\tau+1} \approx_{1.1} \mathbf{s}(\mu/(1 + \varepsilon_{\text{step}}\delta)^\tau) \quad (9.44)$$

for any $\tau \in [t]$.

Now let $\hat{\mathbf{r}}$ be the resistances of \mathcal{L} at any point since the last call to $\mathcal{L}.\text{BATCHUPDATE}$. By the lemma statement and (9.44), we know that $\hat{\mathbf{r}} \approx_{1.12} \mathbf{r}(\hat{\mu})$ for some $\hat{\mu} \in [\mu/(1 + \varepsilon_{\text{step}}\delta)^t, \mu^0]$. However, we also know that $\mu^0 \leq \mu \cdot (1 + \varepsilon_{\text{step}}\delta)^{(0.5\alpha^{1/4}-k)\varepsilon_{\text{step}}^{-1}}$ and so

$$\frac{\hat{\mu}}{\mu/(1 + \varepsilon_{\text{step}}\delta)^t} \leq (1 + \varepsilon_{\text{step}}\delta)^{0.5\alpha^{1/4}\varepsilon_{\text{step}}^{-1}} \leq (1 + \delta)^{0.5\alpha^{1/4}},$$

so by Lemma 9.2.7 we have

$$\mathbf{s}(\mu/(1 + \varepsilon_{\text{step}}\delta)^t) \approx_{0.75\alpha^{1/2}} \mathbf{s}(\hat{\mu}).$$

As $\mathbf{s}^{t+1} \approx_{1.1} \mathbf{s}(\mu/(1 + \varepsilon_{\text{step}}\delta)^t)$, we have that $\mathbf{s}^{t+1} \approx_{0.825\alpha^{1/2}} \mathbf{s}(\hat{\mu})$, and so

$$\mathbf{r}^{t+1} \approx_{0.825\alpha} \mathbf{r}(\hat{\mu}) \approx_{1.12} \hat{\mathbf{r}}.$$

This means that $\mathbf{r}^{t+1} \approx_\alpha \hat{\mathbf{r}}$ and is a contradiction.

We conclude that the requirements of \mathcal{L} are met for all t , and as a result Lemma 4.3.5 shows that $\mathbf{s}^{T+1} \approx_{1.1} \mathbf{s}(\mu/(1 + \varepsilon_{\text{step}}\delta)^{k\varepsilon_{\text{step}}^{-1}})$. By Lemma 4.3.6, we can now obtain $\mathbf{f}(\mu/(1 + \varepsilon_{\text{step}}\delta)^{k\varepsilon_{\text{step}}^{-1}})$. Finally, we return $\mathcal{L}.\mathbf{r}$ and \mathcal{C} to their original states.

Success probability. We note that all the outputs of \mathcal{C} are independent of the randomness of \mathcal{L} , and \mathcal{L} is only updated based on these outputs. As each operation of \mathcal{L} succeeds with high probability, the whole process succeeds with high probability as well.

Runtime. The recentering operation in Lemma 4.3.6 takes $\tilde{O}(m)$. Additionally, we call $\mathcal{L}.\text{SOLVE}$ $k\varepsilon_{\text{step}}^{-1} = O(k^4)$ times and, as $|Z| = O(1/\varepsilon^2)$, the total number of times $\mathcal{L}.\text{UPDATE}$, $\mathcal{C}.\text{UPDATE}$, and $\mathcal{C}.\text{CHECK}$ are called is $O(k\varepsilon_{\text{step}}^{-1}\varepsilon^{-2}) = O(k^{16})$. □

Proof of Lemma 4.3.10

Proof. Let $\delta = 1/\sqrt{m}$. Over a number of $T = \tilde{O}(m^{1/2}/k)$ iterations, we will repeatedly apply MULTISTEP (Lemma 4.3.11). We will also replace the oracle from Definition 4.3.9 by the CHECKER data structure in Section 9.2.6.

Initialization. We first initialize the LOCATOR with error $\varepsilon/2$, by calling $\mathcal{L}.\text{INITIALIZE}(\mathbf{f})$. Let \mathbf{s}^t be the slacks $\mathcal{L}.\mathbf{s}$ before the t -th iteration and \mathbf{r}^t the corresponding resistances, and \mathbf{s}^{0t} be the slacks $\mathcal{L}.\mathbf{r}^0$ before the t -th iteration and \mathbf{r}^{0t} the corresponding resistances, for $t \in [T]$. Also, we let $\mu_t = \mu/(1 + \varepsilon_{\text{step}}\delta)^{(t-1)k\varepsilon_{\text{step}}^{-1}}$. We will maintain the invariant that $\mathbf{s}^t \approx_{1+\varepsilon_{\text{solve}}/8} \mathbf{s}(\mu_t)$, which is a requirement in order to apply Lemma 4.3.11.

As in [73], we will also need to maintain $O(k^4)$ CHECKERS \mathcal{C}^i for $i \in [O(k^4)]$, so we call $\mathcal{C}^i.\text{INITIALIZE}(\mathbf{f}, \varepsilon, \beta_{\text{CHECKER}})$ for each one of these. Note that in general $\beta_{\text{CHECKER}} \neq \beta$, as the vertex sparsifiers \mathcal{L} and \mathcal{C}^i will not be on the same vertex set. As in Lemma 4.3.11, we will maintain the invariant that $\mathcal{L}.\mathbf{r} = \mathcal{C}^i.\mathbf{r}$ for all i .

Resistance updates. Assuming that all the requirements of Lemma 4.3.11 (MULTISTEP) are satisfied at the t -th iteration, that lemma computes a flow $\bar{\mathbf{f}} = \mathbf{f}(\mu_{t+1})$ with slacks $\bar{\mathbf{s}}$. In order to guarantee that $\mathbf{s}^{t+1} \approx_{1+\varepsilon_{\text{solve}}/8} \mathbf{s}(\mu_{t+1})$, we let Z be the set of edges such that either

$$s_e^{+,t} \not\approx_{1+\varepsilon_{\text{solve}}/8} \bar{s}_e^+ \text{ or } s_e^{-,t} \not\approx_{1+\varepsilon_{\text{solve}}/8} \bar{s}_e^-$$

and then call $\mathcal{L}.\text{UPDATE}(e, \bar{\mathbf{f}})$ for all $e \in Z$. This guarantees that $s_e^{+,t+1} = \bar{s}_e^+$ and $s_e^{-,t+1} = \bar{s}_e^-$ for all $e \in Z$ and so $\mathbf{s}^{t+1} \approx_{1+\varepsilon_{\text{solve}}/8} \bar{\mathbf{s}} = \mathbf{s}(\mu_{t+1})$. We also apply the same updates to the \mathcal{C}^i 's using $\mathcal{C}^i.\text{UPDATE}$, in order to ensure that they have the same resistances with \mathcal{L} .

Batched resistance updates. The number of times $\mathcal{L}.\text{UPDATE}$ is called can be quite large because of multiple edges on which error slowly accumulates. This is because in general $\Omega(m)$ resistances will be updated throughout the algorithm. As $\text{LOCATOR}.\text{UPDATE}$ is only slightly sublinear, this would lead to an $\Omega(m^{3/2})$ -time algorithm. For this reason, as in [73], we occasionally (every \hat{T} iterations for some $\hat{T} \geq 1$ to be defined later) perform batched updates by calling $\mathcal{L}.\text{BATCHUPDATE}(Z, \bar{\mathbf{f}})$, where Z is the set of edges such that either

$$s_e^{+,t} \not\approx_{1+\varepsilon_{\text{solve}}/16} \bar{s}_e^+ \text{ or } s_e^{-,t} \not\approx_{1+\varepsilon_{\text{solve}}/16} \bar{s}_e^- .$$

This again guarantees that $s_e^{+,t+1} = \bar{s}_e^+$ and $s_e^{-,t+1} = \bar{s}_e^-$ for all $e \in Z$ and so $\mathbf{s}^{t+1} \approx_{1+\varepsilon_{\text{solve}}/16} \bar{\mathbf{s}} = \mathbf{s}(\mu_{t+1})$. Note that after updating $\mathcal{L}.\mathbf{s}$ and $\mathcal{L}.\mathbf{r}$, this operation also sets $\mathcal{L}.\mathbf{r}^0 = \mathcal{L}.\mathbf{r}$. We perform the same resistance updates to the \mathcal{C}^i 's in the regular (i.e. not batched) way, using $\mathcal{C}^i.\text{UPDATE}$.

LOCATOR requirements. What is left is to ensure that the requirements of Lemma 4.3.11 are satisfied at the t -th iteration, as well as that the requirements of $\mathcal{L}.\text{UPDATE}$ and $\mathcal{L}.\text{BATCHUPDATE}$ from Definition 4.3.7 are satisfied.

The requirements are as follows:

1. Lemma 4.3.11: $\mathbf{s}^{0t} \approx_{1+\varepsilon_{\text{solve}}/8} \mathbf{s}(\mu^0)$ for some $\mu^0 \leq \mu_t \cdot (1 + \varepsilon_{\text{step}}\delta)^{(0.5\alpha^{1/4}-k)\varepsilon_{\text{step}}^{-1}}$.

Note that $\mathcal{L}.\mathbf{s}^0$ is updated every time $\mathcal{L}.\text{BATCHUPDATE}$ is called, and after the call we have $\mathcal{L}.\mathbf{s}^0 = \mathcal{L}.\mathbf{s} \approx_{1+\varepsilon_{\text{solve}}/16} \mathbf{s}(\mu^0)$ for some $\mu^0 > 0$. To ensure that it is called often enough, we call $\mathcal{L}.\text{BATCHUPDATE}(\emptyset)$ every $(0.5\alpha^{1/4}/k - 1)\varepsilon_{\text{step}}^{-1}$ iterations. Because of this, we have $\mu^0 \leq \mu_t \cdot (1 + \varepsilon_{\text{step}}\delta)^{(0.5\alpha^{1/4}/k-1)\varepsilon_{\text{step}}^{-1} \cdot k} = \mu_t \cdot (1 + \varepsilon_{\text{step}}\delta)^{(0.5\alpha^{1/4}-k)\varepsilon_{\text{step}}^{-1}}$. Additionally, for any resistances $\hat{\mathbf{r}}$ that \mathcal{L} had at any point since the last call to $\mathcal{L}.\text{BATCHUPDATE}$, it is immediate that

$$\hat{\mathbf{r}} \approx_{(1+\varepsilon_{\text{solve}}/8)^2} \mathbf{r}(\hat{\mu})$$

for some $\hat{\mu} \in [\mu_t, \mu^0]$, as this is exactly the invariant that our calls to $\mathcal{L}.\text{UPDATE}$ maintain. Therefore, $\hat{\mathbf{r}} \approx_{1.1^2} \mathbf{r}(\hat{\mu})$.

2. $\mathcal{L}.\text{UPDATE}$: $r_e^{\max}/\alpha \leq r_e^{t+1} \leq \alpha \cdot r_e^{\min}$, where r_e^{\min}, r_e^{\max} are the minimum and maximum values that $\mathcal{L}.\mathbf{r}_e$ has had since the last call to $\mathcal{L}.\text{BATCHUPDATE}$.

Let $\hat{\mathbf{r}}$ be any value of $\mathcal{L}.\mathbf{r}$ since the last call to $\mathcal{L}.\text{BATCHUPDATE}$. Because of the invariant maintained by resistance updates (including inside MULTISTEP), we have that $\hat{\mathbf{r}}$ are $(\hat{\mu}, 1.1)$ -central resistances for some $\hat{\mu}$ such that

$$\mu_{t+1} \leq \hat{\mu} \leq \mu_{t+1} \cdot (1 + \varepsilon_{\text{step}}\delta)^{(0.5\alpha^{1/4}-k)\varepsilon_{\text{step}}^{-1}} \leq \mu_{t+1} \cdot (1 + \delta)^{0.5\alpha^{1/4}} .$$

As in the previous item, we have that and \mathbf{s}^{0t} are $(\mu^0, 1 + \varepsilon_{\text{solve}}/8)$ -central. By Lemma 9.2.7 this implies

$$\mathbf{s}(\mu_{t+1}) \approx_{0.75\alpha^{1/2}} \mathbf{s}(\hat{\mu}),$$

and since $\hat{\mathbf{r}} \approx_{1.12} \mathbf{r}(\hat{\mu})$, we conclude that $\mathbf{r}(\mu_{t+1}) \approx_{\alpha} \hat{\mathbf{r}}$.

3. $\mathcal{L}.\text{BATCHUPDATE}$: Between any two successive calls to $\mathcal{L}.\text{INITIALIZE}$, the number of edges updated (number of calls to $\mathcal{L}.\text{UPDATE}$ plus the sum of $|Z|$ for all calls to $\mathcal{L}.\text{BATCHUPDATE}$) is $O(\beta m)$.

We make sure that this is satisfied by calling $\mathcal{L}.\text{INITIALIZE}(\bar{\mathbf{f}})$ every $\varepsilon_{\text{solve}}\sqrt{\beta m}/k$ iterations, where $\bar{\mathbf{f}} = \mathbf{f}(\mu_t)$, at the beginning of the t -th iteration.

Consider any two successive initializations at iterations t^{init} and t^{end} respectively. Let ℓ be the number of edges e that have potentially been updated, i.e. such that either

$$s_e(\mu_{t^{\text{init}}})^+ \not\approx_{1+\varepsilon_{\text{solve}}/16} s_e(\mu_i)^+ \text{ or } s_e(\mu_{t^{\text{init}}})^- \not\approx_{1+\varepsilon_{\text{solve}}/16} s_e(\mu_i)^-$$

for some $i \in [t^{\text{init}}, t^{\text{end}}]$. First, note that this implies that

$$\sqrt{r_e(\mu_{t^{\text{init}}})} |f_e(\mu_{t^{\text{init}}}) - f_e(\mu_i)| > \frac{\varepsilon_{\text{solve}}/16}{1 + \varepsilon_{\text{solve}}/16} > \varepsilon_{\text{solve}}/17.$$

Now, by the fact that $t^{\text{end}} - t^{\text{init}} \leq \varepsilon_{\text{solve}}\sqrt{\beta m}/k$, we have that

$$\mu_{t^{\text{init}}} \leq \mu_i \cdot (1 + \varepsilon_{\text{step}}\delta)^{k\varepsilon_{\text{step}}^{-1} \cdot \varepsilon_{\text{solve}}\sqrt{\beta m}/k} \leq \mu_i \cdot (1 + \delta)^{\varepsilon_{\text{solve}}\sqrt{\beta m}}.$$

By applying Lemma 9.2.6 with $k = \varepsilon_{\text{solve}}\sqrt{\beta m}$ and $\gamma = \varepsilon_{\text{solve}}/17$, we get that $\ell \leq O(\beta m)$. Therefore the statement follows.

We will also need to show how to implement the PERFECTCHECKER used in MULTISTEP using the \mathcal{C}^i 's, as well as how to satisfy all CHECKER requirements.

CHECKER requirements.

1. Implementing ε - PERFECTCHECKER inside MULTISTEP .

We follow almost the same procedure as in [73], other than the fact that we also need to provide some additional information to $\mathcal{C}^i.\text{SOLVE}$. Each call to $\text{PERFECTCHECKER.UPDATE}$ translates to calls to $\mathcal{C}^i.\text{TEMPORARYUPDATE}$ for all i . In addition, the i -th batch of calls to $\text{PERFECTCHECKER.CHECK}$ inside MULTISTEP (i.e. that corresponding to a single set of edges returned by $\mathcal{L}.\text{SOLVE}$) is only run on \mathcal{C}^i using $\mathcal{C}^i.\text{CHECK}$. As each call to $\mathcal{C}^i.\text{CHECK}$ is independent of previous calls to it, we can get correct outputs with high probability even when we run it multiple times (one for each edge returned by $\mathcal{L}.\text{SOLVE}$).

In order to guarantee that we have a vector $\boldsymbol{\pi}_{\text{old}}^i$ as required by $\mathcal{C}^i.\text{CHECK}$, once every k^4 calls to MULTISTEP (i.e. if t is a multiple of k^4) we compute

$$\boldsymbol{\pi}_{\text{old}}^i = \boldsymbol{\pi}^{C^{i,t}} \left(\mathbf{B}^\top g(\mathbf{s}(\mu_t)) \right)$$

for all $i \in [O(k^4)]$, where $C^{i,t}$ is the vertex set of the Schur complement data structure stored internally by the \mathcal{C}^i right before the t -th call to MULTISTEP . This can be computed in $\tilde{O}(m)$ for each i as in $\text{DEMANDPROJECTOR.INITIALIZE}$ in Lemma 4.4.14. Now, the total number

of \mathcal{C}^i .TEMPORARYUPDATES that have not been rolled back is $O(k^{16})$, and the total number of \mathcal{C}^i .UPDATES over k^4 calls to MULTISTEP by Lemma 9.2.6 is $O(k^{10}/\varepsilon_{\text{solve}}^2) = O(k^{16})$. This means that the total number of terminal insertions to $\mathcal{C}^{i,t}$ as well as resistance changes is $O(k^{16})$. By Lemma 4.4.12, if \mathcal{C}^i is the current state of the vertex set of the Schur complement of \mathcal{C}^i and \mathbf{s} are the current slacks,

$$\mathcal{E}_r \left(\boldsymbol{\pi}_{old}^i - \boldsymbol{\pi}^{\mathcal{C}^i}(g(\mathbf{s})) \right) \leq \tilde{O} \left(\alpha' \beta_{\text{CHECKER}}^{-4} \right) \cdot k^{32},$$

where α' is the largest possible multiplicative change of some $\mathcal{C}^i.r_e$ since the computation of $\boldsymbol{\pi}_{old}^i$. Furthermore, note that $\boldsymbol{\pi}_{old}^i$ is supported on \mathcal{C}^i . This is because $\mathcal{C}^{i,t} \subseteq \mathcal{C}^i$ and $\mathcal{C}^{i,t}$ does not contain any temporary terminals.

Now, as we have already proved in Lemma 4.3.11, at any point inside the t -th call to MULTISTEP, $\mathcal{C}^i.r$ are $(\hat{\mu}, 1.1)$ -central resistances for some $\hat{\mu} \in [\mu_{t+1}, \mu_t]$.

Fix $\hat{\mu} \in [\mu_{t+1}, \mu_t], \hat{\mu}' \in [\mu_{t'+1}, \mu_{t'}]$, as well as the corresponding resistances of $\mathcal{C}^i, \hat{\mathbf{r}}, \hat{\mathbf{r}}'$, where $t' \geq t$. Now, note that since we are computing $\boldsymbol{\pi}_{old}^i$ every k^4 calls to MULTISTEP, we have that

$$\frac{\hat{\mu}}{\hat{\mu}'} \leq \frac{\mu_t}{\mu_{t'+1}} \leq (1 + \varepsilon_{\text{step}}\delta)^{k\varepsilon_{\text{step}}^{-1} \cdot (t'-t+1)} \leq (1 + \delta)^{O(k^5)},$$

so Lemma 9.2.7 implies that

$$\mathbf{s}(\hat{\mu}) \approx_{O(k^{10})} \mathbf{s}(\hat{\mu}').$$

As $\hat{\mathbf{r}} \approx_{1.12} \mathbf{r}(\hat{\mu})$ and $\hat{\mathbf{r}}' \approx_{1.12} \mathbf{r}(\hat{\mu}')$, we get that $\hat{\mathbf{r}} \approx_{O(k^{20})} \hat{\mathbf{r}}'$. Therefore, $\alpha' \leq O(k^{20})$. Setting $\beta_{\text{CHECKER}} \geq \tilde{\Omega}(\alpha'^{1/4} k^8 \varepsilon^{-1/2} m^{-1/4}) = \tilde{\Omega}(k^{16}/m^{1/4})$, we get that

$$\tilde{O}(\alpha' \beta_{\text{CHECKER}}^{-4}) \cdot k^{32} \leq \varepsilon^2 m/4,$$

as required by \mathcal{C}^i .CHECK.

Finally, at the end of MULTISTEP we bring all \mathcal{C}^i to their original state before calling MULTISTEP, by calling \mathcal{C}^i .ROLLBACK. We also update all the resistances of \mathcal{L} to their original state by calling \mathcal{L} .UPDATE.

2. Between any two successive calls to \mathcal{C}^i .INITIALIZE, the total number of edges updated at any point (via \mathcal{C}^i .UPDATE or \mathcal{C}^i .TEMPORARYUPDATE that have not been rolled back) is $O(\beta_{\text{CHECKER}} m)$.

For UPDATE, we can apply a similar analysis as in the LOCATOR case to show that if we call \mathcal{C}^i .INITIALIZE every $\varepsilon_{\text{solve}} \sqrt{\beta_{\text{CHECKER}} m}/k$ iterations, the total number of updates never exceeds $O(\beta_{\text{CHECKER}} m)$. For TEMPORARYUPDATE, note that at any time there are at most $O(k^{16})$ of these that have not been rolled back (this is inside MULTISTEP). Therefore, as long as $k^{16} \leq O(\beta_{\text{CHECKER}} m) \Leftrightarrow \beta_{\text{CHECKER}} \geq \Omega(k^{16}/m)$, the requirement is met.

Output Guarantee. After the application of Lemma 4.3.11 at the last iteration, we will have $\bar{\mathbf{f}} = \mathbf{f}(\mu_{T+1})$, where $\mu_{T+1} = \mu/(1 + \varepsilon_{\text{step}}\delta)^{\tilde{O}(m^{1/2}\varepsilon_{\text{step}}^{-1})} = \mu/\text{poly}(m) \leq m^{-10}$.

Success probability. Note that all operations of LOCATOR and CHECKER work with high probability. Regarding the interaction of the randomness of these data structures and the fact that they work against oblivious adversaries, we defer to [73], where there is a detailed discussion of why this works.

In short, note that outside of MULTISTEP, all updates are deterministic (as they only depend on the central path), and in MULTISTEP the updates to LOCATOR and CHECKER only depend on outputs of a CHECKER. As each time we are getting the output from a different CHECKER, the inputs to \mathcal{C}^i .CHECK are independent of the randomness of \mathcal{C}^i , and thus succeed with high probability. Finally, note that the output of LOCATOR is only passed onto \mathcal{C}^i .CHECK, whose output is then independent of the inputs received by LOCATOR. Therefore, LOCATOR does not “leak” any randomness.

Our only deviation from [73] in CHECKER has to do with the extra input of CHECKER.CHECK (π_{old}^i). However, note that this is computed outside of MULTISTEP, and as such the only randomness it depends on is the β_{CHECKER} -congestion reduction subset \mathcal{C}^i generated when calling \mathcal{C}^i .INITIALIZE. As such, it only depends on the internal randomness of \mathcal{C}^i . As we mentioned, the output of \mathcal{C}^i .CHECK is never fed back to \mathcal{C}^i , and thus the operation works with high probability.

Runtime (except CHECKER). Each call to MULTISTEP (Lemma 4.3.11) takes time $\tilde{O}(m)$ plus $O(k^{16})$ calls to \mathcal{L} .UPDATE and $O(k^4)$ calls to \mathcal{L} .SOLVE. As the total number of iterations is $\tilde{O}(m^{1/2}/k)$, the total time because of calls to MULTISTEP is $\tilde{O}(m^{3/2}/k)$, plus $\tilde{O}(m^{1/2}k^{15})$ calls to \mathcal{L} .UPDATE and $\tilde{O}(m^{1/2}k^3)$ calls to \mathcal{L} .SOLVE.

Now, the total number of calls to \mathcal{L} .INITIALIZE is $\tilde{O}\left(\frac{m^{1/2}/k}{\varepsilon_{\text{solve}}\sqrt{\beta m/k}}\right) = \tilde{O}(k^3\beta^{-1/2})$.

The total number of calls to \mathcal{L} .BATCHUPDATE(\emptyset) is $\tilde{O}\left(\frac{m^{1/2}/k}{0.5\alpha^{1/4}/k}\right) = \tilde{O}(m^{1/2}\alpha^{-1/4})$ and the total number of calls to \mathcal{L} .BATCHUPDATE(Z, \mathbf{f}) is $\tilde{O}\left(\frac{m^{1/2}}{k\hat{T}}\right)$. Regarding the size of Z , let us focus on the calls to \mathcal{L} .BATCHUPDATE(Z, \mathbf{f}) between two successive calls to \mathcal{L} .INITIALIZE. We already showed that the sum of $|Z|$ over all calls during this interval is $O(\beta m)$. Therefore the total sum of $|Z|$ over all iterations of the algorithm is $\tilde{O}(mk^3\beta^{1/2})$.

In order to bound the number of calls to \mathcal{L} .UPDATE, we concentrate on those between two successive calls to \mathcal{L} .BATCHUPDATE(Z, \mathbf{f}) in iterations t^{old} and $t^{new} > t^{old}$. After the call to \mathcal{L} .BATCHUPDATE(Z, \mathbf{f}) in iteration t^{old} we have $\mathcal{L}.s \approx_{1+\varepsilon_{\text{solve}}/16} s(\mu_{t^{old}})$. Fix $\mu \in [\mu(t^{new}), \mu(t^{old})]$ and let ℓ be the number of $e \in E$ such that $s_e(\mu) \not\approx_{1+\varepsilon_{\text{solve}}/8} s_e^{t^{old}}$. As $s_e^{t^{old}} \approx_{1+\varepsilon_{\text{solve}}/16} s_e(\mu_{t^{old}})$ by the guarantees of \mathcal{L} .BATCHUPDATE, this implies that $s_e(\mu) \not\approx_{1+\varepsilon_{\text{solve}}/16} s_e(\mu_{t^{old}})$, and so

$$\sqrt{r_e(\mu^{t^{old}})} |f_e(\mu_{t^{old}}) - f_e(\mu)| \geq \frac{\varepsilon_{\text{solve}}/16}{1 + \varepsilon_{\text{solve}}/16} > \varepsilon_{\text{solve}}/17.$$

As

$$\mu_{t^{old}} \leq \mu \cdot (1 + \varepsilon_{\text{step}}\delta)^{k\varepsilon_{\text{step}}^{-1}\hat{T}} \leq \mu \cdot (1 + \delta)^{k\hat{T}},$$

by applying Lemma 9.2.6 with $k = (1 + \delta)^{k\hat{T}}$ and $\gamma = \varepsilon_{\text{solve}}/17$, we get that $\ell \leq O(k^2\hat{T}^2\varepsilon_{\text{solve}}^{-2})$. As there are $\tilde{O}\left(\frac{m^{1/2}}{k\hat{T}}\right)$ calls to \mathcal{L} .BATCHUPDATE(Z, \mathbf{f}), the total number of calls to \mathcal{L} .UPDATE is

$$\tilde{O}\left(m^{1/2}\hat{T}\varepsilon_{\text{solve}}^{-2}\right) \tilde{O}\left(m^{1/2}k^6\hat{T}\right).$$

We conclude that we have runtime $\tilde{O}(m^{3/2}/k)$, plus

- $\tilde{O}(k^3\beta^{-1/2})$ calls to \mathcal{L} .INITIALIZE,
- $\tilde{O}(m^{1/2}k^3)$ calls to \mathcal{L} .SOLVE,
- $\tilde{O}\left(m^{1/2}\left(k^6\hat{T} + k^{15}\right)\right)$ calls to \mathcal{L} .UPDATE,

- $\tilde{O}(m^{1/2}\alpha^{-1/4})$ calls to $\mathcal{L}.\text{BATCHUPDATE}(\emptyset)$, and
- $\tilde{O}(m^{1/2}k^{-1}\hat{T}^{-1})$ calls to $\mathcal{L}.\text{BATCHUPDATE}(Z, \mathbf{f})$.

Runtime of CHECKER. We look at each operation separately. We begin with the runtime of CHECKER.CHECK. We have

$$\tilde{O}\left(\underbrace{m^{1/2}/k}_{\# \text{ calls to MULTISTEP}} \cdot \underbrace{k^{16}}_{\# \text{ calls in each MULTISTEP}} \cdot \underbrace{(\beta_{\text{CHECKER}}m + (k^{16}\beta_{\text{CHECKER}}^{-2}\varepsilon^{-2})^2)\varepsilon^{-2}}_{\text{runtime per call}}\right).$$

To make the first term $\tilde{O}(m^{3/2}/k)$, we set $\beta_{\text{CHECKER}} = k^{-28}$. Note that this satisfies our previous requirements that $\beta_{\text{CHECKER}} \geq \tilde{\Omega}(k^{16}/m)$ and $\beta_{\text{CHECKER}} \geq \tilde{\Omega}(k^{16}/m^{1/4})$ as long as $k \leq m^{1/176}$. Therefore the total runtime because of this operation is $\tilde{O}(m^{3/2}/k + m^{1/2}k^{195})$.

For CHECKER.INITIALIZE, we have

$$\tilde{O}\left(\underbrace{k^4}_{\# \text{ CHECKERS}} \cdot \underbrace{k^3\beta_{\text{CHECKER}}^{-1/2}}_{\# \text{ times initialized}} \cdot \underbrace{m\beta_{\text{CHECKER}}^{-4}\varepsilon^{-4}}_{\text{runtime per init}}\right) = \tilde{O}(mk^{157}).$$

For CHECKER.UPDATE, similarly with the analysis of LOCATOR but noting that there are no batched updates, we have

$$\tilde{O}\left(\underbrace{k^4}_{\# \text{ CHECKERS}} \cdot \underbrace{m\varepsilon_{\text{solve}}^{-2}}_{\# \text{ calls per CHECKER}} \cdot \underbrace{\beta_{\text{CHECKER}}^{-2}\varepsilon^{-2}}_{\text{runtime per call}}\right) = \tilde{O}(mk^{78}).$$

For CHECKER.TEMPORARYUPDATE, we have

$$\tilde{O}\left(\underbrace{k^4}_{\# \text{ CHECKERS}} \cdot \underbrace{m^{1/2}/k}_{\# \text{ calls to MULTISTEP}} \cdot \underbrace{k^{16}}_{\# \text{ calls per MULTISTEP}} \cdot \underbrace{(k^{16}\beta_{\text{CHECKER}}^{-2}\varepsilon^{-2})^2}_{\text{runtime per call}}\right) = \tilde{O}(m^{1/2}k^{187}).$$

Finally, note that, by definition, computing the vectors π_{old}^i takes $\tilde{O}(m^{3/2}/k)$, as we do it once per k^4 calls to MULTISTEP and it takes $\tilde{O}(mk^4)$.

As long as $k \leq m^{1/316}$, the total runtime because of CHECKER is $\tilde{O}(m^{3/2}/k)$. \square

9.2.4 Deferred Proofs from Section 4.4

Proof of Lemma 4.4.4

We first provide a helper lemma for upper bounding escape probabilities in terms of the underlying graph's resistances.

Lemma 9.2.11 (Bounding escape probabilities). *Let a graph with resistances \mathbf{r} , and consider a random walk which at each step moves from the current vertex u to an adjacent vertex v sampled with probability proportional to $1/r_{uv}$. Let $p_u^{\{u,t\}}(s)$ represent the probability that a walk starting at s*

hits u before t . Then

$$p_u^{\{u,t\}}(s) = \frac{R_{eff}(s,t)}{R_{eff}(u,t)} \cdot p_s^{\{s,t\}}(u) \leq \frac{R_{eff}(s,t)}{R_{eff}(u,t)} \leq \frac{r_{st}}{R_{eff}(u,t)}.$$

Proof. Using standard arguments we can prove that if \mathbf{L} is the Laplacian associated with the underlying graph, then

$$p_u^{\{u,t\}}(s) = \frac{(\mathbf{1}_s - \mathbf{1}_t)^\top \mathbf{L}^+(\mathbf{1}_u - \mathbf{1}_t)}{R_{eff}(u,t)}.$$

This immediately yields the claim as we can further write it as

$$p_u^{\{u,t\}}(s) = \frac{R_{eff}(s,t)}{R_{eff}(u,t)} \cdot \frac{(\mathbf{1}_u - \mathbf{1}_t)^\top \mathbf{L}^+(\mathbf{1}_s - \mathbf{1}_t)}{R_{eff}(s,t)} = \frac{R_{eff}(s,t)}{R_{eff}(u,t)} \cdot p_s^{\{s,t\}}(u) \leq \frac{R_{eff}(s,t)}{R_{eff}(u,t)},$$

where we crucially used the symmetry of \mathbf{L} . The final inequality is due to the fact that $R_{eff}(s,t) \leq r_{st}$.

Now let us prove the claimed identity for escape probabilities. Let $\boldsymbol{\psi}$ be the vector defined by $\psi_i = p_u^{\{u,t\}}(i)$ for all $i \in V$, which clearly satisfies $\psi_u = 1$ and $\psi_t = 0$. Furthermore, for all $i \notin \{u,t\}$ we have

$$\psi_i = \sum_{j \sim i} \frac{r_{ij}^{-1}}{\sum_{k \sim i} r_{ik}^{-1}} \psi_j,$$

which can be written in short as

$$(\mathbf{L}\boldsymbol{\psi})_i = 0 \quad \text{for all } i \notin \{s,t\}.$$

Now we solve the corresponding linear system. We interpret $\boldsymbol{\psi}$ as electrical potentials corresponding to routing $1/R_{eff}(u,t)$ units of electrical flow from u to t . Indeed, by Ohm's law, this corresponds to a potential difference $\psi_u - \psi_t = 1$. Furthermore, this shows that

$$\psi_s - \psi_t = (\mathbf{1}_s - \mathbf{1}_t)^\top \mathbf{L}^+(\mathbf{1}_u - \mathbf{1}_t) \cdot \frac{1}{R_{eff}(u,t)},$$

which concludes the proof. \square

Now we are ready to prove the main statement.

Proof of Lemma 4.4.4. Note that the demand can be decomposed as $\mathbf{d} - \boldsymbol{\pi}^C(\mathbf{d}) = \mathbf{d}^1 - \mathbf{d}^2$, where $\mathbf{d}^1 = \frac{\mathbf{1}_s}{\sqrt{r_{st}}} - \boldsymbol{\pi}^C(\frac{\mathbf{1}_s}{\sqrt{r_{st}}})$ and $\mathbf{d}^2 = \frac{\mathbf{1}_t}{\sqrt{r_{st}}} - \boldsymbol{\pi}^C(\frac{\mathbf{1}_t}{\sqrt{r_{st}}})$. Now let p^1 be the probability distribution of $s - C$ random walks obtained via a random walk from s with transition probabilities proportional to inverse resistances. Similarly, let p^2 be the probability distribution of $t - C$ random walks obtained by running the same process starting from t .

Now, it is well known that an electrical flow is the sum of these random walks, i.e.

$$\mathbf{R}^{-1} \mathbf{B} \mathbf{L}^+ \mathbf{d}^1 = \frac{1}{\sqrt{r_{st}}} \cdot \mathbb{E}_{P \sim p^1} [\text{net}(P)]$$

and similarly for \mathbf{d}^2

$$\mathbf{R}^{-1} \mathbf{B} \mathbf{L}^+ \mathbf{d}^2 = \frac{1}{\sqrt{r_{st}}} \cdot \mathbb{E}_{P \sim p^2} [\text{net}(P)],$$

where $\text{net}(P) \in \mathbb{R}^m$ is a flow vector whose e -th entry is the net number of times the edge $e = (u, v)$

is used by P . Therefore we can write:

$$\begin{aligned} \left| \frac{\phi_u - \phi_v}{\sqrt{r_{uv}}} \right| &= \sqrt{r_{uv}} \left| \mathbf{R}^{-1} \mathbf{B} \mathbf{L}^+ \mathbf{d} \right|_{uv} \\ &= \sqrt{\frac{r_{uv}}{r_{st}}} \left| \mathbb{E}_{P^1 \sim p^1} [net_e(P^1)] - \mathbb{E}_{P^2 \sim p^2} [net_e(P^2)] \right|. \end{aligned}$$

Let us also subdivide e by inserting an additional vertex w in the middle (i.e. $r_{uw} = r_{vw} = r_{uv}/2$). This has no effect in the random walks, but will be slightly more convenient in terms of notation. The first expectation term can be expressed as

$$\begin{aligned} \mathbb{E}_{P^1 \sim p^1} [net_e(P^1)] &= \Pr_{P^1 \sim p^1} [P^1 \text{ visits } t \text{ before } C \cup \{w\}] \cdot \mathbb{E}_{P^1 \sim p^1} [net_e(P^1) \mid P^1 \text{ visits } t \text{ before } C \cup \{w\}] \\ &\quad + \Pr_{P^1 \sim p^1} [P^1 \text{ visits } w \text{ before } C \cup \{t\}] \cdot \mathbb{E}_{P^1 \sim p^1} [net_e(P^1) \mid P^1 \text{ visits } t \text{ before } C \cup \{t\}]. \end{aligned}$$

Now, note that

$$\mathbb{E}_{P^1 \sim p^1} [net_e(P^1) \mid P^1 \text{ visits } t \text{ before } C \cup \{w\}] = \mathbb{E}_{P^2 \sim p^2} [net_e(P^2)].$$

Additionally,

$$\begin{aligned} &\Pr_{P^1 \sim p^1} [P^1 \text{ visits } w \text{ before } C \cup \{t\}] \\ &= \Pr_{P^1 \sim p^1} [P^1 \text{ visits } w \text{ before } C] \cdot \Pr_{P^1 \sim p^1} [P^1 \text{ visits } w \text{ before } t \mid P^1 \text{ visits } w \text{ before } C] \end{aligned}$$

The first term of the product is $p_w^{C \cup \{w\}}(s)$. For the second term, we define a new graph \widehat{G} by deleting C , and denote the hitting probabilities in \widehat{G} by \widehat{p} . Then, the second term is equal to $\widehat{p}_w^{\{t,w\}}(s)$.

We have concluded that

$$\mathbb{E}_{P^1 \sim p^1} [net_e(P^1)] \leq \mathbb{E}_{P^2 \sim p^2} [net_e(P^2)] + p_w^{C \cup \{w\}}(s) \cdot \widehat{p}_w^{\{t,w\}}(s).$$

Combining this with the symmetric argument for p^2 shows that

$$\left| \mathbb{E}_{P^1 \sim p^1} [net_e(P^1)] - \mathbb{E}_{P^2 \sim p^2} [net_e(P^2)] \right| \leq p_w^{C \cup \{w\}}(s) \cdot \widehat{p}_w^{\{t,w\}}(s) + p_w^{C \cup \{w\}}(t) \cdot \widehat{p}_w^{\{s,w\}}(t).$$

Using Lemma 9.2.11 and the fact that $\widehat{R}_{eff}(w, t) \geq r_{uw}/4$ (\widehat{R}_{eff} are the effective resistances in \widehat{G}), we can bound

$$\widehat{p}_w^{\{t,w\}}(s) \leq \min\left\{1, \frac{r_{st}}{\widehat{R}_{eff}(w, t)}\right\} \leq \min\left\{1, 4 \frac{r_{st}}{r_{uv}}\right\},$$

and the same upper bound holds for $\widehat{p}_w^{\{s,w\}}(t)$. Therefore,

$$\begin{aligned} &\left| \mathbb{E}_{P^1 \sim p^1} [net_e(P^1)] - \mathbb{E}_{P^2 \sim p^2} [net_e(P^2)] \right| \\ &\leq \min\left\{1, 4 \frac{r_{st}}{r_{uv}}\right\} \left(p_w^{C \cup \{w\}}(s) + p_w^{C \cup \{w\}}(t) \right) \\ &\leq 2 \sqrt{\frac{r_{st}}{r_{uv}}} \left(p_w^{C \cup \{w\}}(s) + p_w^{C \cup \{w\}}(t) \right) \\ &\leq 2 \sqrt{\frac{r_{st}}{r_{uv}}} \left(p_u^{C \cup \{u\}}(s) + p_v^{C \cup \{v\}}(s) + p_u^{C \cup \{u\}}(t) + p_v^{C \cup \{v\}}(t) \right). \end{aligned}$$

Putting everything together, we have that

$$\left| \frac{\phi_u - \phi_v}{\sqrt{r_{uv}}} \right| \leq 2 \left(p_u^{C \cup \{u\}}(s) + p_v^{C \cup \{v\}}(s) + p_u^{C \cup \{u\}}(t) + p_v^{C \cup \{v\}}(t) \right).$$

□

Proof of Lemma 4.4.5

Proof. Let $\mathbf{d} = \mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{r}}$ and $e = (u, w)$. Note that $\mathcal{E}_r(\mathbf{d}) \leq r_e \cdot \left(\frac{1}{\sqrt{r_e}}\right)^2 = 1$, therefore the case that remains is

$$R_{eff}(C, e) \geq 36 \cdot r_e. \quad (9.45)$$

For each $v \in C$ by Lemma 4.4.6 we have that

$$|\pi_v^C(\mathbf{d})| \leq (p_v^C(u) + p_v^C(w)) \cdot \frac{\sqrt{r_e}}{R_{eff}(v, e)}.$$

Now, we would like to bound the energy of routing $\pi^C(\mathbf{d})$ by the energy to route it via w . For each $v \in C$ we let \mathbf{d}^v be the following demand:

$$\mathbf{d}^v = \pi_v^C(\mathbf{d}) \cdot (\mathbf{1}_v - \mathbf{1}_w).$$

Note that $\pi^C(\mathbf{d}) = \sum_{v \in C} \mathbf{d}^v$. We have,

$$\begin{aligned} \sqrt{\mathcal{E}_r(\pi^C(\mathbf{d}))} &= \sqrt{\mathcal{E}_r\left(\pi^C\left(\sum_{v \in C} \mathbf{d}^v\right)\right)} \\ &\leq \sum_{v \in C} \sqrt{\mathcal{E}_r(\pi^C(\mathbf{d}^v))} \\ &\leq \sum_{v \in C} (R_{eff}(v, w))^{1/2} |\pi_v^C(\mathbf{d})| \\ &\leq \sum_{v \in C} (R_{eff}(v, w))^{1/2} \cdot (p_v^C(u) + p_v^C(w)) \cdot \frac{\sqrt{r_e}}{R_{eff}(v, e)}. \end{aligned}$$

Now, note that, because R_{eff} is a metric,

$$\begin{aligned} R_{eff}(v, u) &\geq R_{eff}(v, w) - |R_{eff}(v, w) - R_{eff}(v, u)| \\ &\geq R_{eff}(v, w) - r_e \\ &\geq R_{eff}(v, w) - \frac{1}{36} R_{eff}(C, e) \\ &\geq \frac{35}{36} R_{eff}(v, w), \end{aligned}$$

where we used the triangle inequality twice, (9.45), and also the fact that $R_{eff}(v, w) \geq R_{eff}(C, e)$ because $v \in C$ and $w \in e$. Now, note also that

$$R_{eff}(v, e) \geq \frac{1}{2} \min\{R_{eff}(v, w), R_{eff}(v, u)\} \geq \frac{1}{2} \cdot \frac{35}{36} R_{eff}(v, w),$$

so

$$\begin{aligned} & 2 \sum_v (R_{eff}(v, w))^{1/2} \cdot (p_v^C(u) + p_v^C(w)) \cdot \frac{\sqrt{r_e}}{R_{eff}(v, e)} \\ & \leq 2 \sqrt{2 \cdot \frac{36}{35}} \sum_v (p_v^C(u) + p_v^C(w)) \cdot \sqrt{\frac{r_e}{R_{eff}(v, e)}} \\ & \leq 6 \cdot \sqrt{\frac{r_e}{R_{eff}(C, e)}}, \end{aligned}$$

where we used the fact that $\sum_{v \in C} (p_v^C(u) + p_v^C(w)) = 2$ and $R_{eff}(v, e) \geq R_{eff}(C, e)$ because $v \in C$. \square

Proof of Lemma 4.4.9

Proof. By definition of the fact that e is not ε -important, we have

$$R_{eff}(C, e) > r_e/\varepsilon^2.$$

Using the fact that the demand $\pi^C \left(\mathbf{B}^\top \frac{\mathbf{p}}{\sqrt{r}} \right)$ is supported on C and Lemma 4.4.5, we get

$$\begin{aligned} \left| \left\langle \mathbf{1}_e, \mathbf{R}^{-1/2} \mathbf{B} \phi^* \right\rangle \right| &= \left| \left\langle \pi^C \left(\mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{r}} \right), \phi_C^* \right\rangle \right| \\ &\leq \sqrt{\mathcal{E}_r \left(\pi^C \left(\mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{r}} \right) \right)} \cdot \sqrt{E_r(\phi^*)} \\ &= \sqrt{\mathcal{E}_r \left(\pi^C \left(\mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{r}} \right) \right)} \cdot \delta \sqrt{\mathcal{E}_r \left(\pi^C \left(\mathbf{B}^\top \frac{\mathbf{p}}{\sqrt{r}} \right) \right)} \\ &\leq 6 \sqrt{\frac{r_e}{R_{eff}(C, e)}} \cdot \delta \sqrt{m} \\ &\leq 6\varepsilon. \end{aligned}$$

\square

Proof of Lemma 4.4.11

Proof. We note that

$$\begin{aligned}
& \sqrt{\mathcal{E}_r \left(\pi^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{r}} \right) - \pi^C \left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{r}} \right) \right)} \\
&= \left| \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{r}} \right) \right| \sqrt{R_{eff}(C, v)} \\
&\leq \sum_{e=(u,w) \in E} \left| \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{q_e}{\sqrt{r_e}} \mathbf{1}_e \right) \right| \sqrt{R_{eff}(C, v)} \\
&\leq \sum_{e=(u,w) \in E} (p_v^{C \cup \{v\}}(u) + p_v^{C \cup \{v\}}(w)) \cdot \min \left\{ \sqrt{\frac{R_{eff}(C, v)}{r_e}}, \frac{\sqrt{r_e R_{eff}(C, v)}}{R_{eff}(e, v)} \right\},
\end{aligned}$$

where we used Lemma 4.4.6. For some sufficiently large c to be defined later, we partition E into X and Y , where $X = \{e \in E \mid R_{eff}(C, v) \leq c^2 \cdot r_e \text{ or } r_e R_{eff}(C, v) \leq c^2 \cdot (R_{eff}(e, v))^2\}$ and $Y = E \setminus X$. We first note that

$$\begin{aligned}
& \sum_{e=(u,w) \in X} (p_v^{C \cup \{v\}}(u) + p_v^{C \cup \{v\}}(w)) \cdot \min \left\{ \sqrt{\frac{R_{eff}(C, v)}{r_e}}, \frac{\sqrt{r_e R_{eff}(C, v)}}{R_{eff}(e, v)} \right\} \\
&\leq c \cdot \sum_{e=(u,w) \in X} (p_v^{C \cup \{v\}}(u) + p_v^{C \cup \{v\}}(w)) \\
&\leq c \cdot \tilde{O}(\beta^{-2}),
\end{aligned}$$

where the last inequality follows by the congestion reduction property.

Now, let $e = (u, w) \in Y$. We will prove that both u and w are much closer to v than C . This, in turn, will imply that their hitting probabilities on v are roughly the same, and so they mostly cancel out in the projection.

First of all, we let $R_{eff}(C, v) = c_1^2 \cdot r_e$ and $r_e R_{eff}(C, v) = c_2^2 \cdot (R_{eff}(e, v))^2$, for some $c_1, c_2 > 0$, where by definition $c_1, c_2 \geq c$. Now, we assume without loss of generality that $R_{eff}(u, v) \leq R_{eff}(w, v)$, and so

$$R_{eff}(u, v) \leq 2R_{eff}(e, v) = \frac{2}{c_2} \sqrt{r_e R_{eff}(C, v)} = \frac{2}{c_1 c_2} R_{eff}(C, v) \leq \frac{2}{c_1 c_2} (R_{eff}(C, u) + R_{eff}(u, v)),$$

so

$$R_{eff}(u, v) \leq \frac{\frac{2}{c_1 c_2}}{1 - \frac{2}{c_1 c_2}} R_{eff}(C, u) = \frac{2}{c_1 c_2 - 2} R_{eff}(C, u) \leq \frac{3}{c_1 c_2} R_{eff}(C, u). \quad (9.46)$$

Futhermore, note that

$$R_{eff}(w, v) \leq R_{eff}(u, v) + r_e \leq \left(\frac{2}{c_1 c_2} + \frac{1}{c_1^2} \right) R_{eff}(C, v) \leq \left(\frac{2}{c_1 c_2} + \frac{1}{c_1^2} \right) (R_{eff}(C, w) + R_{eff}(w, v)),$$

and so we have

$$R_{eff}(w, v) \leq \frac{\frac{2}{c_1 c_2} + \frac{1}{c_1^2}}{1 - \frac{2}{c_1 c_2} - \frac{1}{c_1^2}} R_{eff}(C, w) \leq 3 \left(\frac{1}{c_1 c_2} + \frac{1}{c_1^2} \right) R_{eff}(C, w). \quad (9.47)$$

Now, by Lemma 9.2.11 together with (9.46) we have

$$p_v^{C \cup \{v\}}(u) \geq 1 - \frac{3}{c_1 c_2}$$

and with (9.47) we have

$$p_v^{C \cup \{v\}}(w) \geq 1 - 3 \left(\frac{1}{c_1 c_2} + \frac{1}{c_1^2} \right),$$

therefore

$$\left| p_v^{C \cup \{v\}}(u) - p_v^{C \cup \{v\}}(w) \right| \leq 6 \left(\frac{1}{c_1 c_2} + \frac{1}{c_1^2} \right).$$

So,

$$\begin{aligned} \left| \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_e}{\sqrt{r_e}} \mathbf{1}_e \right) \right| \sqrt{R_{eff}(C, v)} &= \left| p_v^{C \cup \{v\}}(u) - p_v^{C \cup \{v\}}(w) \right| \sqrt{\frac{R_{eff}(C, v)}{r_e}} \\ &\leq 6 \left(\frac{1}{c_1 c_2} + \frac{1}{c_1^2} \right) \cdot c_1 \\ &= 6 \left(\frac{1}{c_2} + \frac{1}{c_1} \right) \\ &= O \left(\frac{1}{c} \right). \end{aligned}$$

Now, we will apply Lemma 4.4.10 to prove that with high probability $|Y| \leq \tilde{O}(\beta^{-1})$. The reason we can apply the lemma is that for any $e = (u, w) \in Y$, we have

$$R_{eff}(e, v) = \frac{1}{c_1 c_2} R_{eff}(C, v) \leq \frac{1}{2} R_{eff}(C, v),$$

and so $Y \subseteq N_E(v, R_{eff}(C, v)/2)$. Therefore, we get that

$$\sum_{e=(u,w) \in Y} \left| \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_e}{\sqrt{r_e}} \mathbf{1}_e \right) \right| \sqrt{R_{eff}(C, v)} \leq c^{-1} \cdot \tilde{O}(\beta^{-1}).$$

Overall, we conclude that

$$\sqrt{\mathcal{E}_r \left(\pi^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}} \right) - \pi^C \left(\mathbf{B}^\top \frac{\mathbf{q}}{\sqrt{\mathbf{r}}} \right) \right)} \leq (c + c^{-1}) \tilde{O}(\beta^{-2}) = \tilde{O}(\beta^{-2}),$$

by setting c to be a large enough constant. □

Proof of Lemma 4.4.12

Proof. We write

$$\begin{aligned}
& \pi^{C^T, r^T} \left(B^\top \frac{\mathbf{q}^T}{\sqrt{\mathbf{r}^T}} \right) - \pi^{C^0, r^0} \left(B^\top \frac{\mathbf{q}^0}{\sqrt{\mathbf{r}^0}} \right) \\
&= \underbrace{\sum_{i \text{ is an ADDTERMINAL}} \pi_{v^i}^{C^{i+1}, r^i} \left(B^\top \frac{\mathbf{q}^i}{\sqrt{\mathbf{r}^i}} \right) \cdot (\mathbf{1}_{v^i} - \pi^{C^i, r^i}(\mathbf{1}_{v^i}))}_{d_{Add}} \\
&+ \underbrace{\sum_{i \text{ is an UPDATE}} \pi^{C^i, r^i} \left(B^\top \left(\frac{\mathbf{q}^{i+1}}{\sqrt{\mathbf{r}^{i+1}}} - \frac{\mathbf{q}^i}{\sqrt{\mathbf{r}^i}} \right) \mathbf{1}_{e^i} \right)}_{d_{Upd}},
\end{aligned}$$

which implies that

$$\sqrt{\mathcal{E}_{r^T} \left(\pi^{C^T, r^T} \left(B^\top \frac{\mathbf{q}^T}{\sqrt{\mathbf{r}^T}} \right) - \pi^{C^0, r^0} \left(B^\top \frac{\mathbf{q}^0}{\sqrt{\mathbf{r}^0}} \right) \right)} \leq \sqrt{\mathcal{E}_{r^T}(d_{Add})} + \sqrt{\mathcal{E}_{r^T}(d_{Upd})}.$$

We bound each of these terms separately. For the second one, we have that

$$\begin{aligned}
& \sqrt{\mathcal{E}_{r^T}(d_{Upd})} \\
&\leq \sum_{i \text{ is an UPDATE}} \sqrt{\mathcal{E}_{r^T} \left(\pi^{C^i, r^i} \left(B^\top \left(\frac{\mathbf{q}^{i+1}}{\sqrt{\mathbf{r}^{i+1}}} - \frac{\mathbf{q}^i}{\sqrt{\mathbf{r}^i}} \right) \mathbf{1}_{e^i} \right) \right)} \\
&= \sum_{i \text{ is an UPDATE}} \sqrt{\mathcal{E}_{r^T} \left(B^\top \left(\frac{\mathbf{q}^{i+1}}{\sqrt{\mathbf{r}^{i+1}}} - \frac{\mathbf{q}^i}{\sqrt{\mathbf{r}^i}} \right) \mathbf{1}_{e^i} \right)} \\
&\leq \sum_{i \text{ is an UPDATE}} \left(\sqrt{\mathcal{E}_{r^T} \left(B^\top \frac{\mathbf{q}^{i+1}}{\sqrt{\mathbf{r}^{i+1}}} \mathbf{1}_{e^i} \right)} + \sqrt{\mathcal{E}_{r^T} \left(B^\top \frac{\mathbf{q}^i}{\sqrt{\mathbf{r}^i}} \mathbf{1}_{e^i} \right)} \right) \\
&\leq \max_i \left\| \frac{\mathbf{r}^T}{\mathbf{r}^i} \right\|_\infty^{1/2} \sum_{i \text{ is an UPDATE}} \left(\sqrt{\mathcal{E}_{r^{i+1}} \left(B^\top \frac{\mathbf{q}^{i+1}}{\sqrt{\mathbf{r}^{i+1}}} \mathbf{1}_{e^i} \right)} + \sqrt{\mathcal{E}_{r^i} \left(B^\top \frac{\mathbf{q}^i}{\sqrt{\mathbf{r}^i}} \mathbf{1}_{e^i} \right)} \right) \\
&\leq 2 \max_i \left\| \frac{\mathbf{r}^T}{\mathbf{r}^i} \right\|_\infty^{1/2} T.
\end{aligned}$$

For the first one, we have

$$\begin{aligned}
& \sqrt{\mathcal{E}_{r^T}(d_{Add})} \\
&\leq \sum_{i \text{ is an ADDTERMINAL}} \left| \pi_{v^i}^{C^{i+1}, r^i} \left(B^\top \frac{\mathbf{q}^i}{\sqrt{\mathbf{r}^i}} \right) \right| \cdot \sqrt{\mathcal{E}_{r^T}(\mathbf{1}_{v^i} - \pi^{C^i, r^i}(\mathbf{1}_{v^i}))} \\
&\leq \left\| \frac{\mathbf{r}^T}{\mathbf{r}^i} \right\|_\infty^{1/2} \sum_{i \text{ is an ADDTERMINAL}} \left| \pi_{v^i}^{C^{i+1}, r^i} \left(B^\top \frac{\mathbf{q}^i}{\sqrt{\mathbf{r}^i}} \right) \right| \cdot \sqrt{\mathcal{E}_{r^i}(\mathbf{1}_{v^i} - \pi^{C^i, r^i}(\mathbf{1}_{v^i}))} \\
&\leq \tilde{O} \left(\max_i \left\| \frac{\mathbf{r}^T}{\mathbf{r}^i} \right\|_\infty^{1/2} \beta^{-2} \right) \cdot T,
\end{aligned}$$

where in the last inequality we used Lemma 4.4.11. The desired statement now follows immediately. \square

Proof of Lemma 4.3.8

Proof. INITIALIZE(\mathbf{f}): We set $\mathbf{s}^+ = \mathbf{u} - \mathbf{f}$, $\mathbf{s}^- = \mathbf{f}$, $\mathbf{r}^0 = \mathbf{r} = \frac{1}{(s^+)^2} + \frac{1}{(s^-)^2}$.

We first initialize a β -congestion reduction subset C based on Lemma 4.4.2, which takes time $\tilde{O}(m\beta^{-2})$, and a data structure DYNAMICSC for maintaining the sparsified Schur Complement onto C , as described in Appendix 9.2.1, which takes time $\tilde{O}(m\beta^{-4}\varepsilon^{-4})$. We also set $C^0 = C$.

Then, we generate an $\tilde{O}(\varepsilon^{-2}) \times m$ sketching matrix \mathbf{Q} as in (Lemma 5.1, [73] v2), which takes time $\tilde{O}(m\varepsilon^{-2})$, and let its rows be \mathbf{q}^i for $i \in [\tilde{O}(\varepsilon^{-2})]$.

In order to compute the set of important edges, we use Lemma 9.2.3 after contracting C , which shows that we can compute all resistances of the form $R_{eff}(C, u)$ for $u \in V \setminus C$ up to a factor of 2 in $\tilde{O}(m)$. From these, we can get 4-approximate estimates of $R_{eff}(C, e)$ for $e \in E \setminus E(C)$, using the fact that

$$\min\{R_{eff}(C, u), R_{eff}(C, w)\} \approx_2 R_{eff}(C, e).$$

Then, in $O(m)$ time, we can easily compute a set of edges S such that

$$\{e \mid e \text{ is } \frac{\varepsilon\beta}{\alpha}\text{-important}\} \subseteq S \subseteq \{e \mid e \text{ is } \frac{\varepsilon\beta}{4\alpha}\text{-important}\}$$

We also need to sample the random walks that will be used inside the demand projection data structures. We use (Lemma 5.15, [73] v2) to sample $h = \tilde{O}(\hat{\varepsilon}^{-4}\beta^{-6} + \hat{\varepsilon}^{-2}\beta^{-2}\gamma^{-2})$ random walks for each $u \in V \setminus C$ and $e \in E \setminus E(C)$ with $u \in e$, where we set $\gamma = \frac{\varepsilon}{4\alpha}$ so that S is a subset of γ -important edges. Note that, by Definition 4.3.7, a γ -important edge will always remain γ -important until the LOCATOR is re-initialized, as any edge's resistive distance to C can only decrease, and its own resistance is constant. Therefore S can be assumed to always be a subset of γ -important edges.

The runtime to sample the set \mathcal{P} of these random walks is

$$\tilde{O}(mh\beta^{-2}) = \tilde{O}(m(\hat{\varepsilon}^{-4}\beta^{-8} + \hat{\varepsilon}^{-2}\beta^{-4}\gamma^{-2})) = \tilde{O}(m(\hat{\varepsilon}^{-4}\beta^{-8} + \hat{\varepsilon}^{-2}\varepsilon^{-2}\alpha^2\beta^{-4})).$$

In order to be able to detect congested edges, we will initialize $\tilde{O}(\varepsilon^{-2})$ demand projection data structures, with the guarantees from Lemma 4.4.14. We will maintain an approximation to $\boldsymbol{\pi}^C(\mathbf{B}^\top \frac{\mathbf{q}_S^i}{\sqrt{r}})$ for all $i \in \tilde{O}(\varepsilon^{-2})$, where \mathbf{q}^i are the rows of the sketching matrix that we have generated, as well as $\boldsymbol{\pi}^{old} := \boldsymbol{\pi}^{C^0, \mathbf{r}^0}(\mathbf{B}^\top \frac{\mathbf{p}^0}{\sqrt{r^0}})$, where $\mathbf{p}^0 = \sqrt{r^0}g(\mathbf{s}^0) = \frac{\frac{1}{s^+,0} - \frac{1}{s^-,0}}{\sqrt{r^0}}$.

Specifically, we call

$$\text{DP}^i.\text{INITIALIZE}(C, \mathbf{r}, \mathbf{q}^i, S, \mathcal{P})$$

for all $i \in [\tilde{O}(\varepsilon^{-2})]$, and also exactly compute $\boldsymbol{\pi}^{old}$, which can be done by calling

$$\text{DEMANDPROJECTOR.INITIALIZE}(C, \mathbf{r}, \mathbf{p}, [m], \mathcal{P}).$$

The total runtime for this operation is dominated by the random walk generation, and is

$$\tilde{O}(m(\hat{\varepsilon}^{-4}\beta^{-8} + \hat{\varepsilon}^{-2}\varepsilon^{-2}\alpha^2\beta^{-4})).$$

UPDATE(e, \mathbf{f}): We set $s_e^+ = u_e - f_e$, $s_e^- = f_e$, and $r_e = \frac{1}{(s_e^+)^2} + \frac{1}{(s_e^-)^2}$. Then, we also set $p_e = \frac{\frac{1}{s_e^+} - \frac{1}{s_e^-}}{\sqrt{r_e}}$.

We distinguish two cases:

- $e \in E(C)$:

In this case, we can simply call

$$\text{DYNAMICSC.UPDATE}(e, r_e)$$

and

$$\text{DP}^i.\text{UPDATE}(e, \mathbf{r}, \mathbf{q})$$

for all $i \in [\tilde{O}(\varepsilon^{-2})]$. Note that we can do this as DP^i was initialized with resistances \mathbf{r}^0 and $\mathbf{r}^0 \approx_\alpha \mathbf{r}$.

- $e \in E \setminus E(C)$:

We let $e = (u, w)$. We want to insert u and w into C , but for doing that DP^i 's require constant factor estimates of the resistances $R_{eff}(C, u)$ and $R_{eff}(C \cup \{u\}, w)$. In order to get these estimates, we will use **DYNAMICSC**.

We first call

$$\text{DYNAMICSC.ADDTERMINAL}(u),$$

which takes time $\tilde{O}(\beta^{-2}\varepsilon^{-2})$ and returns $\tilde{R}_{eff}(C, u) \approx_2 R_{eff}(C, u)$. Given this estimate, we can call

$$\text{DP}^i.\text{ADDTERMINAL}(u, \tilde{R}_{eff}(C, u)),$$

for all $i \in [\tilde{O}(\varepsilon^{-2})]$, each of which takes time

$$\tilde{O}(\hat{\varepsilon}^{-4}\beta^{-8} + \hat{\varepsilon}^{-2}\beta^{-6}\gamma^{-2}) = \tilde{O}(\hat{\varepsilon}^{-4}\beta^{-8} + \hat{\varepsilon}^{-2}\varepsilon^{-2}\alpha^2\beta^{-6}).$$

Now, we can set $C = C \cup \{u\}$ and repeat the same process for w .

Finally, to update the resistance, note that we now have $e \in E(C)$, so we apply the procedure from the first case.

Finally, if the total number of calls to $\text{DP}^i.\text{ADDTERMINAL}$ for some fixed i since the last call to $\mathcal{L}.\text{BATCHUPDATE}(\emptyset)$ exceeds $\frac{\varepsilon}{\hat{\varepsilon}\alpha^{1/2}}$ (note that the number of calls is actually the same for all i), we call $\mathcal{L}.\text{BATCHUPDATE}(\emptyset)$ in order to re-initialize the demand projections.

We conclude that the total runtime is

$$\tilde{O}\left(m\frac{\hat{\varepsilon}\alpha^{1/2}}{\varepsilon^3} + \hat{\varepsilon}^{-4}\varepsilon^{-2}\beta^{-8} + \hat{\varepsilon}^{-2}\varepsilon^{-4}\alpha^2\beta^{-6}\right),$$

where the first term comes from amortizing the calls to $\mathcal{L}.\text{BATCHUPDATE}(\emptyset)$, each of which, as we will see, takes $\tilde{O}(m\varepsilon^{-2})$.

BATCHUPDATE(Z, \mathbf{f}): First, for each $e \in Z$, we set $s_e^+ = u_e - f_e$, $s_e^- = f_e$, $r_e^0 = r_e = \frac{1}{(s_e^+)^2} + \frac{1}{(s_e^-)^2}$,

and $p_e^0 = p_e = \frac{\frac{1}{s_e^+} - \frac{1}{s_e^-}}{\sqrt{r_e}}$.

For each $e = (u, w) \in Z$, we call

$$\text{DYNAMICSC.ADDTERMINAL}(u)$$

and

$$\text{DYNAMICSC.ADDTERMINAL}(w)$$

(if u and w are not already in C), and then we call

$$\text{DYNAMICSC.UPDATE}(e, r_e).$$

Then, we set $C^0 = C = C \cup (\cup_{(u,w) \in Z} \{u, w\})$. Additionally, we re-compute $\boldsymbol{\pi}^{old}$ based on the new values of $C^0, \mathbf{r}^0, \mathbf{p}^0$. All of this takes time $\tilde{O}(m + |Z|\beta^{-2}\varepsilon^{-2})$.

Now, to pass these updates to the DEMANDPROJECTORS, we first have to re-compute the set of important edges S (with the newly updated resistances) as any set such that

$$\{e \mid e \text{ is } \frac{\varepsilon}{\alpha}\text{-important}\} \subseteq S \subseteq \{e \mid e \text{ is } \frac{\varepsilon}{4\alpha}\text{-important}\}.$$

As we have already argued, this takes $\tilde{O}(m)$.

Now, finally, we re-initialize all the DEMANDPROJECTORS by calling

$$\text{DP}^i.\text{INITIALIZE}(C, \mathbf{r}, \mathbf{q}, S, \mathcal{P}).$$

for all $i \in [\tilde{O}(\varepsilon^{-2})]$, where each call takes $\tilde{O}(m)$.

We conclude with a total runtime of

$$\tilde{O}(m\varepsilon^{-2} + |Z|\beta^{-2}\varepsilon^{-2}).$$

SOLVE(): This operation performs the main task of the locator, which is to detect congested edges. We will do that by using the approximate demand projections that we have been maintaining.

We remind that the congestion vector we are trying to approximate to $O(\varepsilon)$ additive accuracy is

$$\boldsymbol{\rho}^* = \delta\sqrt{\mathbf{r}}g(\mathbf{s}) - \delta\mathbf{R}^{-1/2}\mathbf{B}\mathbf{L}^+\mathbf{B}^\top g(\mathbf{s}).$$

We will first reduce the problem of finding the entries of $\boldsymbol{\rho}^*$ with magnitude $\geq \Omega(\varepsilon)$, to the problem of computing an $O(\varepsilon)$ -additive approximation to

$$v_i^* = \delta \cdot \left\langle \boldsymbol{\pi}^C \left(\mathbf{B}^\top \frac{\mathbf{q}_S^i}{\sqrt{\mathbf{r}}} \right), \widetilde{\text{SC}}^+ \boldsymbol{\pi}^{C^0, \mathbf{r}^0} \left(\mathbf{B}^\top \frac{\mathbf{p}^0}{\sqrt{\mathbf{r}^0}} \right) \right\rangle$$

for all $i \in [\tilde{O}(\varepsilon^{-2})]$, where $\widetilde{\text{SC}}$ is the approximate Schur complement maintained in DYNAMICSC. Then, we will see how to approximate v_i^* to additive accuracy $O(\varepsilon)$ using the demand projection data structures.

First of all, note that, by definition of $g(\mathbf{s}) = \frac{\frac{1}{s^+} - \frac{1}{s^-}}{\mathbf{r}}$,

$$\|\delta\sqrt{\mathbf{r}}g(\mathbf{s})\|_\infty \leq \delta \leq \varepsilon,$$

so this term can be ignored. Using Lemma 4.4.3, we get that

$$\delta \left\| \mathbf{R}^{-1/2}\mathbf{B}\mathbf{L}^+(\mathbf{B}^\top g(\mathbf{s}) - \boldsymbol{\pi}^C(\mathbf{B}^\top g(\mathbf{s}))) \right\|_\infty \leq \delta \cdot \tilde{O}(\beta^{-2}) \leq \varepsilon/2.$$

This means that the entries of the vector

$$\delta\mathbf{R}^{-1/2}\mathbf{B}\mathbf{L}^+\boldsymbol{\pi}^C(\mathbf{B}^\top g(\mathbf{s}))$$

that have magnitude $\leq \varepsilon$ do not correspond to the $\Omega(\varepsilon)$ -congested edges that we are looking for.

Now, we set $T = |C \setminus C^0|$, where C^0 was the congestion reduction subset during the last call to

BATCHUPDATE, and apply Lemma 4.4.12. This shows that

$$\sqrt{\mathcal{E}_r(\boldsymbol{\pi}^{old} - \boldsymbol{\pi}^C(\mathbf{B}^\top g(\mathbf{s})))} \leq \tilde{O}(\alpha^{1/2}\beta^{-2}) \cdot T.$$

Therefore, if we define

$$\boldsymbol{\rho} = -\delta \mathbf{R}^{-1/2} \mathbf{B} \mathbf{L}^+ \boldsymbol{\pi}^{old},$$

we conclude that

$$\|\boldsymbol{\rho} - \boldsymbol{\rho}^*\|_\infty \leq O(\varepsilon) + \delta \left\| \mathbf{R}^{-1/2} \mathbf{B} \mathbf{L}^+ \left(\boldsymbol{\pi}^{old} - \boldsymbol{\pi}^C(\mathbf{B}^\top g(\mathbf{s})) \right) \right\|_\infty \leq O(\varepsilon) + \delta T \cdot \tilde{O}(\alpha^{1/2}\beta^{-2}) \leq O(\varepsilon),$$

where we used the fact that $T = \frac{\varepsilon}{\tilde{\varepsilon}\alpha^{1/2}} \leq \frac{\varepsilon}{\delta\beta^{-2}\alpha^{1/2}}$. Therefore it suffices to estimate $\boldsymbol{\rho}$ up to $O(\varepsilon)$ -additive accuracy.

Now, note that, by definition, no edge $e \in E \setminus S$ is ε/α -important with respect to \mathbf{r}^0 and C^0 . By using Lemma 4.4.5, for each such edge we get

$$\begin{aligned} & \delta \left| \mathbf{R}^{-1/2} \mathbf{B} \mathbf{L}^+ \boldsymbol{\pi}^{old} \right|_e \\ & \leq \delta \sqrt{\mathcal{E}_r\left(\boldsymbol{\pi}^{C^0}\left(\mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{\mathbf{r}}}\right)\right)} \sqrt{\mathcal{E}_r(\boldsymbol{\pi}^{old})} \\ & \leq \delta \alpha \sqrt{\mathcal{E}_{r^0}\left(\boldsymbol{\pi}^{C^0}\left(\mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{\mathbf{r}}}\right)\right)} \sqrt{\mathcal{E}_{r^0}(\boldsymbol{\pi}^{old})} \\ & \leq \delta \alpha \cdot \frac{\varepsilon}{\alpha} \cdot O(\sqrt{m}) \\ & = O(\varepsilon), \end{aligned}$$

where we also used the fact that $\mathcal{E}_{r^0}(\boldsymbol{\pi}^{C^0, r^0}(g(\mathbf{s}^0))) \leq O(\mathcal{E}_{r^0}(g(\mathbf{s})))$. Therefore it suffices to approximate

$$\boldsymbol{\rho}' = \delta \mathbf{I}_S \mathbf{R}^{-1/2} \mathbf{B} \mathbf{L}^+ \boldsymbol{\pi}^{old}.$$

Note that here we can replace \mathbf{L}^+ by $\begin{pmatrix} -\mathbf{L}_{FF}^{-1} \mathbf{L}_{FC} \\ \mathbf{I} \end{pmatrix} \widetilde{SC}^+$ where $\widetilde{SC} \approx_{1+\varepsilon} SC$ and only lose another additive ε error, as

$$\begin{aligned} & \delta \left| \left\langle \mathbf{1}_e, \mathbf{R}^{-1/2} \mathbf{B} \mathbf{L}^+ \boldsymbol{\pi}^{old} \right\rangle - \delta \left\langle \mathbf{1}_e, \mathbf{R}^{-1/2} \mathbf{B} \begin{pmatrix} -\mathbf{L}_{FF}^{-1} \mathbf{L}_{FC} \\ \mathbf{I} \end{pmatrix} \widetilde{SC}^+ \boldsymbol{\pi}^{old} \right\rangle \right| \\ & = \delta \left| \left\langle \boldsymbol{\pi}^C \left(\mathbf{B}^\top \frac{\mathbf{1}_e}{\sqrt{\mathbf{r}}} \right), (SC^+ - \widetilde{SC}^+) \boldsymbol{\pi}^{old} \right\rangle \right| \\ & \leq O(\delta\varepsilon \cdot \sqrt{m}) \\ & \leq O(\varepsilon), \end{aligned}$$

where we used the fact that

$$(1 - \varepsilon)SC \preceq \widetilde{SC} \preceq (1 + \varepsilon)SC \Rightarrow -O(\varepsilon)\widetilde{SC}^+ \preceq SC^+ - \widetilde{SC}^+ \preceq O(\varepsilon)\widetilde{SC}^+.$$

and that

$$\begin{aligned}\sqrt{\mathcal{E}_r(\boldsymbol{\pi}^{old})} &\leq \sqrt{\mathcal{E}_r(\boldsymbol{\pi}^C(\mathbf{B}^\top g(\mathbf{s})))} + \sqrt{\mathcal{E}_r(\boldsymbol{\pi}^{old} - \boldsymbol{\pi}^C(\mathbf{B}^\top g(\mathbf{s})))} \\ &\leq O(m) + \tilde{O}(\alpha^{1/2}\beta^{-2}) \cdot T \\ &\leq O(m),\end{aligned}$$

where we used the fact that $T = \frac{\varepsilon}{\tilde{\varepsilon}\alpha^{1/2}} \leq \frac{\varepsilon}{\delta\beta^{-2}\alpha^{1/2}} \leq \frac{\sqrt{m}}{\beta^{-2}\alpha^{1/2}}$.

Now, we will use the sketching lemma (Lemma 5.1, [73] v2), which shows that in order to find all entries of

$$\mathbf{I}_S \mathbf{R}^{-1/2} \mathbf{B} \begin{pmatrix} -\mathbf{L}_{FF} \mathbf{L}_{FC} \\ \mathbf{I} \end{pmatrix} \widetilde{SC}^+ \boldsymbol{\pi}^{old}$$

with magnitude $\Omega(\varepsilon)$, it suffices to compute the inner products

$$\begin{aligned}\delta \left\langle \mathbf{B}^\top \frac{\mathbf{q}_S^i}{\sqrt{r}}, \begin{pmatrix} -\mathbf{L}_{FF}^{-1} \mathbf{L}_{FC} \\ \mathbf{I} \end{pmatrix} \widetilde{SC}^+ \boldsymbol{\pi}^{old} \right\rangle \\ = \delta \left\langle \boldsymbol{\pi}^C \left(\mathbf{B}^\top \frac{\mathbf{q}_S^i}{\sqrt{r}} \right), \widetilde{SC}^+ \boldsymbol{\pi}^{old} \right\rangle\end{aligned}$$

for $i \in [\tilde{O}(\varepsilon^{-2})]$, up to additive accuracy

$$\varepsilon \cdot \left\| \delta \mathbf{I}_S \mathbf{R}^{-1/2} \mathbf{B} \begin{pmatrix} -\mathbf{L}_{FF}^{-1} \mathbf{L}_{FC} \\ \mathbf{I} \end{pmatrix} \widetilde{SC}^+ \boldsymbol{\pi}^{old} \right\|_2^{-1} \geq \Omega(\varepsilon),$$

where we used the fact that

$$\begin{aligned}\left\| \delta \mathbf{I}_S \mathbf{R}^{-1/2} \mathbf{B} \begin{pmatrix} -\mathbf{L}_{FF}^{-1} \mathbf{L}_{FC} \\ \mathbf{I} \end{pmatrix} \widetilde{SC}^+ \boldsymbol{\pi}^{old} \right\|_2^2 \\ = \delta^2 \langle \widetilde{SC}^+ \boldsymbol{\pi}^{old}, (-\mathbf{L}_{CF} \mathbf{L}_{FF}^{-1} \quad \mathbf{I}) \mathbf{L} \begin{pmatrix} -\mathbf{L}_{FF}^{-1} \mathbf{L}_{FC} \\ \mathbf{I} \end{pmatrix} \widetilde{SC}^+ \boldsymbol{\pi}^{old} \rangle \\ = \delta^2 \left\langle \widetilde{SC}^+ \boldsymbol{\pi}^{old}, (-\mathbf{L}_{CF} \mathbf{L}_{FF}^{-1} \quad \mathbf{I}) \begin{pmatrix} \mathbf{L}_{FF} & \mathbf{L}_{FC} \\ \mathbf{L}_{CF} & \mathbf{L}_{CC} \end{pmatrix} \begin{pmatrix} -\mathbf{L}_{FF}^{-1} \mathbf{L}_{FC} \\ \mathbf{I} \end{pmatrix} \widetilde{SC}^+ \boldsymbol{\pi}^{old} \right\rangle \\ = \delta^2 \left\langle \widetilde{SC}^+ \boldsymbol{\pi}^{old}, SC \widetilde{SC}^+ \boldsymbol{\pi}^{old} \right\rangle \\ \leq 2\delta^2 \left\langle \boldsymbol{\pi}^{old}, \widetilde{SC}^+ \boldsymbol{\pi}^{old} \right\rangle \\ \leq O(\delta^2 m) \\ = O(1).\end{aligned}$$

Now, for the second part of the proof, we would like to compute \mathbf{v} such that $\|\mathbf{v} - \mathbf{v}^*\|_\infty \leq O(\varepsilon)$, where we remind that

$$\mathbf{v}^* = \left\langle \boldsymbol{\pi}^C \left(\mathbf{B}^\top \frac{\mathbf{q}_S^i}{\sqrt{r}} \right), \widetilde{SC}^+ \boldsymbol{\pi}^{old} \right\rangle.$$

Note that we already have estimates $\tilde{\boldsymbol{\pi}}^C \left(\mathbf{B}^\top \frac{\mathbf{q}_S^i}{\sqrt{r}} \right)$ given by DP^i for all $i \in [\tilde{O}(\varepsilon^{-2})]$. We obtain these estimates by calling

$$\text{DP}^i.\text{OUTPUT}()$$

each of which takes time $O(\beta m)$. By the guarantees of Definition 4.4.13, with high probability we have

$$\delta \left| \left\langle \tilde{\pi}^C \left(\mathbf{B}^\top \frac{\mathbf{q}_S^i}{\sqrt{r}} \right) - \pi^C \left(\mathbf{B}^\top \frac{\mathbf{q}_S^i}{\sqrt{r}} \right), \widetilde{SC}^+ \pi^{old} \right\rangle \right| \leq \widehat{\varepsilon} \sqrt{\alpha} T,$$

where we used the fact that

$$E_r(\delta \cdot \widetilde{SC}^+ \pi^{old}) \leq O(1).$$

Now, since by definition BATCHUPDATE is called every $\frac{\varepsilon}{\widehat{\varepsilon} \alpha^{1/2}}$ calls to UPDATE, We have $T \leq \frac{\varepsilon}{\widehat{\varepsilon} \alpha^{1/2}}$ and so

$$\delta \left| \left\langle \tilde{\pi}^C \left(\mathbf{B}^\top \frac{\mathbf{q}_S^i}{\sqrt{r}} \right) - \pi^C \left(\mathbf{B}^\top \frac{\mathbf{q}_S^i}{\sqrt{r}} \right), \widetilde{SC}^+ \pi^{old} \right\rangle \right| \leq \varepsilon.$$

This means that, running the algorithm from (Lemma 5.1, [73] v2), we can obtain an edge set of size $\widetilde{O}(\varepsilon^{-2})$ that contains all edges such that $|\rho_e^*| \geq c \cdot \varepsilon$ for some constant $c > 0$. By rescaling ε to get the right constant, we obtain all edges such that $|\rho_e^*| \geq \varepsilon/2$ with high probability. The runtime is dominated by the time to get \widetilde{SC} and apply its inverse, and is $\widetilde{O}(\beta m \varepsilon^{-2})$.

Success probability We will argue that \mathcal{L} uses DYNAMICSC and the DP^i as an oblivious adversary. First of all, note that no randomness is injected into the inputs of DYNAMICSC, as they are all coming from the inputs of \mathcal{L} .

Regarding DP^i , note that its only output is given by the call to $DP^i.OUTPUT$. However, note that its output is only used to estimate the inner product

$$\left\langle \pi^C \left(\mathbf{B}^\top \frac{\mathbf{q}_S^i}{\sqrt{r}} \right), \widetilde{SC}^+ \pi^{old} \right\rangle,$$

from which we obtain the set of congested edges and we directly return it from \mathcal{L} . Thus, it does not influence the state of \mathcal{L} , DYNAMICSC or any future inputs. □

9.2.5 Deferred Proofs from Section 4.5

Proof of Lemma 4.4.6

Proof. Let $\mathcal{P}_v(u)$ be a random walk that starts from u and stops when it hits v .

$$\begin{aligned} p_v^{C \cup \{v, w\}}(u) &= \Pr[\mathcal{P}_v(u) \cap C = \emptyset \text{ and } w \notin \mathcal{P}_v(u)] \\ &= \Pr[\mathcal{P}_v(u) \cap C = \emptyset] \cdot \Pr[w \notin \mathcal{P}_v(u) \mid \mathcal{P}_v(u) \cap C = \emptyset] \\ &= p_v^{C \cup \{v\}}(u) \cdot \Pr[w \notin \mathcal{P}_v(u) \mid \mathcal{P}_v(u) \cap C = \emptyset] \end{aligned}$$

Consider new resistances \widehat{r} , where $\widehat{r}_e = r_e$ for all $e \in E$ not incident to C and $\widehat{r}_e = \infty$ for all $e \in E$ incident to C . Also, let \widehat{p} be the hitting probability function for these new resistances. It is easy to see that

$$\Pr[w \notin \mathcal{P}_v(u) \mid \mathcal{P}_v(u) \cap C = \emptyset] = \widehat{p}_v^{\{v, w\}}(u).$$

Therefore, we have

$$p_v^{C \cup \{v, w\}}(u) = p_v^{C \cup \{v\}}(u) \cdot \widehat{p}_v^{\{v, w\}}(u).$$

Now we will bound $\widehat{p}_v^{\{v,w\}}(u)$. Let ψ be electrical potentials for pushing 1 unit of flow from v to w with resistances \widehat{r} and let f be the associated electrical flow. We have that

$$|\psi_u - \psi_w| = |f_e| \widehat{r}_e \leq \widehat{r}_e = r_e \quad (9.48)$$

(because $|f_e| \leq 1$ and e is not incident to C) and

$$|\psi_v - \psi_w| = \widehat{R}_{eff}(v, w) \geq R_{eff}(v, w) \quad (9.49)$$

Additionally, by well known facts that connect electrical potential embeddings with random walks, we have that

$$\psi_u = \psi_w + \widehat{p}_v^{\{v,w\}}(u)(\psi_v - \psi_w),$$

or equivalently

$$\widehat{p}_v^{\{v,w\}}(u) = \frac{\psi_u - \psi_w}{\psi_v - \psi_w}.$$

Using (9.48) and (9.49), this immediately implies that

$$\widehat{p}_v^{\{v,w\}}(u) \leq \frac{r_e}{R_{eff}(v, w)}.$$

So we have proved that

$$p_v^{C \cup \{v,w\}}(u) \leq p_v^{C \cup \{v\}}(u) \frac{r_e}{R_{eff}(v, w)}$$

and symmetrically

$$p_v^{C \cup \{v,u\}}(w) \leq p_v^{C \cup \{v\}}(w) \frac{r_e}{R_{eff}(v, u)}.$$

Now, let's look at $\pi_v^{C \cup \{v\}}(B^\top \mathbf{1}_e) = p_v^{C \cup \{v\}}(u) - p_v^{C \cup \{v\}}(w)$. Note that

$$p_v^{C \cup \{v\}}(u) = p_v^{C \cup \{v,w\}}(u) + p_w^{C \cup \{v,w\}}(u) p_v^{C \cup \{v\}}(w)$$

which we re-write as

$$p_v^{C \cup \{v\}}(u) - p_v^{C \cup \{v\}}(w) = p_v^{C \cup \{v,w\}}(u) - (1 - p_w^{C \cup \{v,w\}}(u)) p_v^{C \cup \{v\}}(w) \leq p_v^{C \cup \{v,w\}}(u).$$

Symmetrically,

$$p_v^{C \cup \{v\}}(w) - p_v^{C \cup \{v\}}(u) \leq p_v^{C \cup \{v,u\}}(w).$$

From these we conclude that

$$\begin{aligned} \left| \pi_v^{C \cup \{v\}}(B^\top \mathbf{1}_e) \right| &= \left| p_v^{C \cup \{v\}}(u) - p_v^{C \cup \{v\}}(w) \right| \\ &\leq \max \left\{ p_v^{C \cup \{v,w\}}(u), p_v^{C \cup \{v,u\}}(w) \right\} \\ &\leq \max \left\{ p_v^{C \cup \{v\}}(u) \cdot \frac{r_e}{R_{eff}(v, w)}, p_v^{C \cup \{v\}}(w) \cdot \frac{r_e}{R_{eff}(v, u)} \right\} \\ &\leq (p_v^{C \cup \{v\}}(u) + p_v^{C \cup \{v\}}(w)) \cdot \max \left\{ \frac{r_e}{R_{eff}(v, w)}, \frac{r_e}{R_{eff}(v, u)} \right\}, \end{aligned}$$

which, after dividing by $\sqrt{r_e}$ gives

$$\left| \pi_v^{C \cup \{v\}}(\mathbf{B}^\top \mathbf{1}_e) \right| \leq (p_v^{C \cup \{v\}}(u) + p_v^{C \cup \{v\}}(w)) \cdot \max \left\{ \frac{\sqrt{r_e}}{R_{eff}(v, w)}, \frac{\sqrt{r_e}}{R_{eff}(v, u)} \right\}.$$

□

Proof of Lemma 4.5.3

Proof. For each $u \in V \setminus C$ and $e \in S'$ with $u \in e$, we generate Z random walks $P^1(u), \dots, P^Z(u)$ from u to $C \cup \{v\}$. We set

$$\begin{aligned} \tilde{\pi}_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_{S'}}{\sqrt{\mathbf{r}}} \right) &= \sum_{e=(u,w) \in S'} \sum_{z=1}^Z \frac{1}{Z} \frac{q_e}{\sqrt{r_e}} (1_{\{v \in P^z(u)\}} - 1_{\{v \in P^z(w)\}}) \\ &= \sum_{e=(u,w) \in S'} \sum_{z=1}^Z (X_{e,u,z} - X_{e,w,z}), \end{aligned}$$

where we have set $X_{e,u,z} = \frac{1}{Z} \frac{q_e}{\sqrt{r_e}} 1_{\{v \in P^z(u)\}}$ and $X_{e,w,z} = -\frac{1}{Z} \frac{q_e}{\sqrt{r_e}} 1_{\{v \in P^z(w)\}}$.

Note that $\mathbb{E}_{P^z(u)} [X_{e,u,z}] = \frac{1}{Z} \frac{q_e}{\sqrt{r_e}} p_v^{C \cup \{v\}}(u)$ and $\mathbb{E}_{P^z(w)} [X_{e,w,z}] = -\frac{1}{Z} \frac{q_e}{\sqrt{r_e}} p_v^{C \cup \{v\}}(w)$. This implies that our estimate is unbiased, as

$$\mathbb{E} \left[\tilde{\pi}_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_{S'}}{\sqrt{\mathbf{r}}} \right) \right] = \sum_{e=(u,w) \in S'} \frac{q_e}{\sqrt{r_e}} (p_v^{C \cup \{v\}}(u) - p_v^{C \cup \{v\}}(w)) = \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_{S'}}{\sqrt{\mathbf{r}}} \right).$$

We now need to show that our estimate is concentrated around the mean. To apply the concentration bound in Lemma 4.5.2, we need the following bounds:

$$\sum_{e=(u,w) \in S'} \sum_{z=1}^Z (|\mathbb{E}[X_{e,u,z}]| + |\mathbb{E}[X_{e,w,z}]|) = \sum_{e=(u,w) \in S'} \frac{|q_e|}{\sqrt{r_e}} (p_v^{C \cup \{v\}}(u) + p_v^{C \cup \{v\}}(w)) := E$$

$$\max_{\substack{e=(u,w) \in S' \\ z \in [Z]}} \max\{|X_{e,u,z}|, |X_{e,w,z}|\} \leq \max_{e \in S'} \frac{1}{Z \sqrt{r_e}} := M.$$

So now for any $t \in [0, E]$ we have

$$\begin{aligned}
& \Pr \left[\left| \tilde{\pi}_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_{S'}}{\sqrt{r}} \right) - \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_{S'}}{\sqrt{r}} \right) \right| > t \right] \\
& \leq 2 \exp \left(-\frac{t^2}{6EM} \right) \\
& = 2 \exp \left(-\frac{Zt^2}{6 \sum_{e=(u,w) \in S'} \frac{|q_e|}{\sqrt{r_e}} (p_v^{C \cup \{v\}}(u) + p_v^{C \cup \{v\}}(w)) \max_{e \in S'} \frac{1}{\sqrt{r_e}}} \right) \\
& \leq 2 \exp \left(-\frac{Zt^2}{6 \sum_{e=(u,w) \in S'} (p_v^{C \cup \{v\}}(u) + p_v^{C \cup \{v\}}(w)) \max_{e \in S'} \frac{1}{r_e}} \right) \\
& \leq 2 \exp \left(-\frac{Zt^2 c^2 R_{eff}(C, v)}{6 \sum_{e=(u,w) \in S'} (p_v^{C \cup \{v\}}(u) + p_v^{C \cup \{v\}}(w))} \right) \\
& \leq 2 \exp \left(-Zt^2 c^2 R_{eff}(C, v) / \tilde{O}(\beta^{-2}) \right) \\
& \leq \frac{1}{n^{100}},
\end{aligned}$$

where the last inequality follows by setting $Z = \tilde{O} \left(\frac{\log n \log \frac{1}{\beta}}{\delta_1'^2} \right)$ and $t = \frac{\delta_1'}{\beta c \sqrt{R_{eff}(C, v)}}$. Note that we have used the fact that $R_{eff}(C, v) \leq r_e/c^2$ for all $e \in S'$, as well as the congestion reduction property (Definition 4.4.1)

$$\sum_{e=(u,w) \in E \setminus E(C)} (p_v^{C \cup \{v\}}(u) + p_v^{C \cup \{v\}}(w)) \leq \tilde{O}(1/\beta^2).$$

□

Proof of Lemma 4.5.4

Proof. In order to compute $\tilde{\pi}_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}} \right)$ we use Lemma 4.5.3 with demand $\mathbf{B}^\top \frac{\mathbf{q}_{S'}}{\sqrt{r}}$ and error parameter $\delta_1' > 0$, where

$$\{e \in S \mid R_{eff}(C, v) \leq r_e/(2c^2)\} \subseteq S' \subseteq \{e \in S \mid R_{eff}(C, v) \leq r_e/c^2\},$$

and $c > 0$ will be defined later. Note that such a set S' can be trivially computed given our effective resistance estimate $\tilde{R}_{eff}(C, v) \approx_2 R_{eff}(C, v)$. However, algorithmically we do not directly compute S' , but instead find its intersection with the edges from which a sampled random walk ends up at v . (Using the congestion reduction property of C , this can be done in $\tilde{O} \left(\delta_1'^{-2} \beta^{-2} \log n \log \frac{1}{\beta} \right)$ time just by going through all random walks that contain v .)

Now, Lemma 4.5.3 guarantees that

$$\left| \tilde{\pi}_v^{CU\{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_{S'}}{\sqrt{r}} \right) - \pi_v^{CU\{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_{S'}}{\sqrt{r}} \right) \right| \leq \frac{\delta'_1}{\beta c \sqrt{R_{eff}(C, v)}}$$

given access to $O(\delta'^{-2} \log n \log \frac{1}{\beta})$ random walks for each $u \in V \setminus C$, $e \in S'$ with $u \in e$.

Then, we set $\tilde{\pi}_v^{CU\{v\}}(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}}) := \tilde{\pi}_v^{CU\{v\}}(\mathbf{B}^\top \frac{\mathbf{q}_{S'}}{\sqrt{r}})$, and we have that

$$\begin{aligned} & \left| \tilde{\pi}_v^{CU\{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}} \right) - \pi_v^{CU\{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}} \right) \right| \\ & \leq \left| \tilde{\pi}_v^{CU\{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_{S'}}{\sqrt{r}} \right) - \pi_v^{CU\{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_{S'}}{\sqrt{r}} \right) \right| + \left| \pi_v^{CU\{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S - \mathbf{q}_{S'}}{\sqrt{r}} \right) \right| \\ & \leq \frac{\delta'_1}{\beta c \sqrt{R_{eff}(C, v)}} + \left| \pi_v^{CU\{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S - \mathbf{q}_{S'}}{\sqrt{r}} \right) \right|. \end{aligned} \quad (9.50)$$

Now, to bound the second term, we use Lemma 4.4.6, which gives

$$\left| \pi_v^{CU\{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S - \mathbf{q}_{S'}}{\sqrt{r}} \right) \right| \leq \sum_{e=(u,w) \in S \setminus S'} \left(p_v^{CU\{v\}}(u) + p_v^{CU\{v\}}(w) \right) \frac{\sqrt{r_e}}{R_{eff}(v, e)}.$$

Now, note that for each $e \in S \setminus S'$, e is close to C , but v is far from C , so $R_{eff}(v, e)$ should be large. Specifically, by Lemma 4.2.8 we have $R_{eff}(v, e) \geq \frac{1}{2} \min \{R_{eff}(v, u), R_{eff}(v, w)\}$, and by the triangle inequality

$$\min \{R_{eff}(v, u), R_{eff}(v, w)\} \geq R_{eff}(C, v) - \max \{R_{eff}(C, u), R_{eff}(C, w)\} \geq R_{eff}(C, v) - 2R_{eff}(C, e).$$

By the fact that e is γ -important and that

$$e \notin S' \supseteq \{e \in S \mid R_{eff}(C, v) \leq r_e / (2c^2)\},$$

we have $R_{eff}(C, e) \leq r_e / \gamma^2 \leq 2c^2 R_{eff}(C, v) / \gamma^2$, so

$$\begin{aligned} \frac{\sqrt{r_e}}{R_{eff}(v, e)} & \leq \frac{1}{1/2 - 2c^2/\gamma^2} \frac{\sqrt{r_e}}{R_{eff}(C, v)} \\ & \leq \frac{c}{1/2 - 2c^2/\gamma^2} \frac{1}{\sqrt{R_{eff}(C, v)}}. \end{aligned}$$

By using the congestion reduction property (Definition 4.4.1), we obtain

$$\left| \pi_v^{CU\{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S - \mathbf{q}_{S'}}{\sqrt{r}} \right) \right| \leq \frac{c}{1/2 - 2c^2/\gamma^2} \frac{1}{\sqrt{R_{eff}(C, v)}} \tilde{O} \left(\frac{1}{\beta^2} \right). \quad (9.51)$$

Setting $c = \min \{ \delta_1 / \tilde{O}(\beta^{-2}), \gamma/4 \}$ and $\delta'_1 = \beta c \cdot \delta_1 / 2$, (9.50) becomes

$$\left| \tilde{\pi}_v^{CU\{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}} \right) - \pi_v^{CU\{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}} \right) \right| \leq \frac{\delta_1}{\sqrt{R_{eff}(C, v)}}.$$

Also, the number of random walks needed for each valid pair (u, e) is

$$\tilde{O}\left(\delta_1^{\prime-2} \log n \log \frac{1}{\beta}\right) = \tilde{O}\left(\delta_1^{-2} \beta^{-2} c^{-2} \log n \log \frac{1}{\beta}\right) = \tilde{O}\left((\delta_1^{-4} \beta^{-6} + \delta_1^{-2} \beta^{-2} \gamma^{-2}) \log n \log \frac{1}{\beta}\right)$$

For the last part of the lemma, we let $S'' = \{e \in S \mid R_{\text{eff}}(C, v) \leq r_e / (\gamma/4)^2\}$ and write

$$\begin{aligned} & \left| \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}} \right) \right| \\ & \leq \left| \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_{S''}}{\sqrt{r}} \right) \right| + \left| \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S - \mathbf{q}_{S''}}{\sqrt{r}} \right) \right|. \end{aligned}$$

For the first term,

$$\begin{aligned} \left| \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_{S''}}{\sqrt{r}} \right) \right| & \leq \sum_{e=(u,w) \in S''} \left(p_v^{C \cup \{v\}}(u) + p_v^{C \cup \{v\}}(w) \right) \frac{1}{\sqrt{r_e}} \\ & \leq \frac{1}{(\gamma/4) \sqrt{R_{\text{eff}}(C, v)}} \tilde{O}\left(\frac{1}{\beta^2}\right), \end{aligned}$$

and for the second term we have already proved in (9.51) (after replacing c by $\gamma/4$) that

$$\left| \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S - \mathbf{q}_{S''}}{\sqrt{r}} \right) \right| \leq \frac{\gamma/4}{\sqrt{R_{\text{eff}}(C, v)}} \tilde{O}\left(\frac{1}{\beta^2}\right).$$

Putting these together, we conclude that

$$\left| \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}} \right) \right| \leq \frac{1}{\gamma \sqrt{R_{\text{eff}}(C, v)}} \cdot \tilde{O}\left(\frac{1}{\beta^2}\right).$$

□

Proof of Lemma 4.5.5

Proof. For any $I \subseteq \mathbb{R}$, we define $F_I = \{i \in [n] \mid |\bar{\phi}_i| \in I\}$. For some $0 < a < b$ to be defined later, we partition $[n]$ as

$$[n] = F_{I_0} \cup F_{I_1} \cup \dots \cup F_{I_K} \cup F_{I_{K+1}},$$

where $I_0 = [0, a)$, $I_{K+1} = [b, \infty)$, and I_1, \dots, I_K is a partition of $[a, b)$ into $K = O(\log \frac{b}{a})$ intervals such that for all $k \in [K]$ we have $\Phi_k := \max_{i \in F_k} |\bar{\phi}_i| \leq 2 \cdot \min_{i \in F_k} |\bar{\phi}_i|$.

A union bound gives

$$\Pr \left[\left| \langle \tilde{\pi} - \pi, \bar{\phi} \rangle \right| > t \right] \leq \sum_{k=0}^{K+1} \Pr \left[\left| \sum_{i \in I_k} (\tilde{\pi}_i - \pi_i) \bar{\phi}_i \right| > t / (K+2) \right].$$

We first examine I_0 and I_{K+1} separately. Note that

$$\left| \sum_{i \in F_{I_0}} (\tilde{\pi}_i - \pi_i) \bar{\phi}_i \right| \leq \|\tilde{\pi} - \pi\|_1 a \leq 2a$$

and

$$\left| \sum_{i \in F_{I_{K+1}}} (\tilde{\pi}_i - \pi_i) \bar{\phi}_i \right| \leq \sum_{i \in F_{I_{K+1}}} \tilde{\pi}_i |\bar{\phi}_i| + \sum_{i \in F_{I_{K+1}}} \pi_i |\bar{\phi}_i| \leq \frac{1}{b} \sum_{i \in F_{I_{K+1}}} |\tilde{\pi}_i - \pi_i| \bar{\phi}_i^2 \leq \frac{1}{b} \sum_{i \in F_{I_{K+1}}} \tilde{\pi}_i |\bar{\phi}_i| + \frac{\|\bar{\phi}\|_{\pi,2}^2}{b}$$

But note that by picking $b \geq \max \left\{ \frac{(K+2) \text{Var}_{\pi}(\bar{\phi})}{t}, \sqrt{\text{Var}_{\pi}(\bar{\phi}) \cdot n^{101}} \right\}$, we have $\frac{\text{Var}_{\pi}(\bar{\phi})}{b} \leq t/(K+2)$ and also for any $i \in F_{K+1}$ we have $\pi_i \leq \frac{\text{Var}_{\pi}(\bar{\phi})}{b^2} \leq \frac{1}{n^{101}}$. This means that $\Pr[\tilde{\pi}_i \neq 0] \leq \frac{1}{n^{101}}$, and so by union bound

$$\Pr \left[\sum_{i \in F_{I_{K+1}}} \tilde{\pi}_i |\bar{\phi}_i| \neq 0 \right] \leq \frac{1}{n^{100}}.$$

Now, we proceed to F_1, \dots, F_K . We draw Z samples x_1, \dots, x_Z from π . Then, we also define the following random variables for $z \in [Z]$ and $i \in [n]$:

$$X_{z,i} = \begin{cases} 1 & \text{if } x_z = i \\ 0 & \text{otherwise} \end{cases}$$

for $i \in [n]$ and

$$Y_{z,k} = \frac{1}{Z} \sum_{i \in F_k} X_{z,i} \bar{\phi}_i$$

This allows us to write $\sum_{i \in F_k} \tilde{\pi}_i \bar{\phi}_i = \sum_{z=1}^Z Y_{z,k}$.

Fix $k \in [K]$. We will apply Lemma 4.5.2 on the random variable $\sum_{z=1}^Z Y_{z,k}$. We first compute

$$\sum_{z=1}^Z |\mathbb{E}[Y_{z,k}]| = \left| \sum_{i \in F_k} \pi_i \bar{\phi}_i \right| \leq \sum_{i \in F_k} \pi_i |\bar{\phi}_i| := E_k$$

and

$$\max_{z \in [Z]} |Y_{z,k}| \leq \frac{\Phi_k}{Z} := M_k.$$

Therefore we immediately have $E_k M_k \leq \frac{2}{Z} \sum_{i \in F_k} \pi_i \bar{\phi}_i^2 \leq \frac{2}{Z} \cdot \text{Var}_\pi(\bar{\phi})$. By Lemma 4.5.2,

$$\begin{aligned} & \Pr \left[\left| \sum_{z=1}^Z Y_{z,k} - \mathbb{E} \left[\sum_{z=1}^Z Y_{z,k} \right] \right| > t/(K+2) \right] \\ & \leq 2 \exp \left(-\frac{t^2}{6E_k M_k (K+2)^2} \right) \\ & \leq 2 \exp \left(-\frac{Zt^2}{12 \cdot \text{Var}_\pi(\bar{\phi})(K+2)^2} \right). \end{aligned}$$

Summarizing, and using the fact that $K = \tilde{O}(\log(n \cdot \text{Var}_\pi(\bar{\phi})/t^2))$, we get

$$\Pr [|\langle \tilde{\pi} - \pi, \bar{\phi} \rangle| > t] \leq \tilde{O} \left(\frac{1}{n^{100}} \right) + 2\tilde{O}(\log(n \cdot \text{Var}_\pi(\bar{\phi})/t^2)) \exp \left(-\frac{Zt^2}{12 \cdot \tilde{O}(\text{Var}_\pi(\bar{\phi}) \log^2 n)} \right).$$

□

Proof of Lemma 4.5.6

Proof. Let $S_0 = \emptyset$ and for each $k \in \mathbb{N}$ let

$$S_k = \{i \in [n] \setminus S_{k-1} : \phi_i^2 \leq 2^{k+1} R_{\text{eff}}(C, v)\}.$$

Fix some $k \geq 2$. Note that $\phi_i^2 > 2^k R_{\text{eff}}(C, v)$ for all $i \in S_k$, implying $\frac{2^k R_{\text{eff}}(C, v)}{R_{\text{eff}}(S_k, v)} < E_r(\phi) \leq 1$, and so $R_{\text{eff}}(S_k, v) > 2^k R_{\text{eff}}(C, v) \geq 4R_{\text{eff}}(C, v)$. As

$$R_{\text{eff}}(C, v) \geq \frac{1}{4} \min\{R_{\text{eff}}(S_k, v), R_{\text{eff}}(C \setminus S_k, v)\} > \min\{R_{\text{eff}}(C, v), \frac{1}{4} R_{\text{eff}}(C \setminus S_k, v)\},$$

we have $R_{\text{eff}}(C \setminus S_k, v) < 4R_{\text{eff}}(C, v)$. This implies that $\left\| \pi_{S_k}^C(\mathbf{1}_v) \right\|_1 \leq \frac{R_{\text{eff}}(C \setminus S_k, v)}{R_{\text{eff}}(S_k, v)} < \frac{1}{2^{k-2}}$.

So we conclude that $\text{Var}_\pi(\phi) = \sum_{i \in S_k} \pi_i \phi_i^2 \leq \frac{1}{2^{k-2}} \cdot 2^{k+1} R_{\text{eff}}(C, v) = 8R_{\text{eff}}(C, v)$. □

Proof of Lemma 4.5.1

Proof. The first part of the statement is given by applying Lemma 4.5.4, and we see that it requires $\tilde{O}(\delta_1^{-4} \beta^{-6} + \delta_1^{-2} \beta^{-2} \gamma^{-2})$ random walks for each $u \in V \setminus C$ and $e \in E \setminus E(C)$ with $u \in e$.

For the second part we use the fact that the change in the demand projection after inserting v into C is given by

$$\pi^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) - \pi^C \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) = \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) \cdot (\mathbf{1}_v - \pi^C(\mathbf{1}_v)),$$

and therefore we can estimate this update via

$$\tilde{\pi}_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{\mathbf{r}}} \right) \cdot (\mathbf{1}_v - \tilde{\pi}^C(\mathbf{1}_v)).$$

where $\tilde{\pi}_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}} \right)$ is the estimate we computed using Lemma 4.5.4 and $\tilde{\pi}^C(\mathbf{1}_v)$ is obtained by applying Lemma 4.5.7.

Let us show that this estimation indeed introduces only a small amount of error. For any ϕ , such that $E_r(\phi) \leq 1$, we can write

$$\begin{aligned} & \left| \left\langle \tilde{\pi}_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}} \right) \cdot (\mathbf{1}_v - \tilde{\pi}^C(\mathbf{1}_v)) - \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}} \right) \cdot (\mathbf{1}_v - \pi^C(\mathbf{1}_v)), \phi \right\rangle \right| \\ & \leq \left| \left\langle \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}} \right) \cdot (\pi^C(\mathbf{1}_v) - \tilde{\pi}^C(\mathbf{1}_v)), \phi \right\rangle \right| \\ & + \left| \left\langle \left(\tilde{\pi}_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}} \right) - \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}} \right) \right) \cdot (\mathbf{1}_v - \pi^C(\mathbf{1}_v)), \phi \right\rangle \right| \\ & + \left| \left\langle \left(\tilde{\pi}_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}} \right) - \pi_v^{C \cup \{v\}} \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}} \right) \right) \cdot (\pi^C(\mathbf{1}_v) - \tilde{\pi}^C(\mathbf{1}_v)), \phi \right\rangle \right|. \end{aligned}$$

At this point we can bound these quantities using Lemmas 4.5.4 and 4.5.7. It is important to notice that they require that S is a set of γ -important edges, for some parameter γ . Our congestion reduction subset C keeps increasing due to vertex insertions. This, however, means that effective resistances between any vertex in $V \setminus C$ and C can only decrease, and therefore the set of important edges can only increase. Thus we are still in a valid position to apply these lemmas.

Using $E_r(\phi) \leq 1$, which allows us to write:

$$\langle \mathbf{1}_v - \pi^C(\mathbf{1}_v), \phi \rangle \leq \mathcal{E}_r(\mathbf{1}_v - \pi^C(\mathbf{1}_v)) = R_{eff}(v, C),$$

we can continue to upper bound the error by:

$$\begin{aligned} & \frac{\tilde{O}(\gamma^{-1}\beta^{-2})}{\sqrt{R_{eff}(C, v)}} \cdot \delta_2 \sqrt{R_{eff}(C, v)} + \frac{\delta_1}{\sqrt{R_{eff}(C, v)}} \cdot \sqrt{R_{eff}(C, v)} + \frac{\delta_1}{\sqrt{R_{eff}(C, v)}} \cdot \delta_2 \sqrt{R_{eff}(C, v)} \\ & = \delta_2 \cdot \tilde{O}(\gamma^{-1}\beta^{-2}) + \delta_1 + \delta_1 \delta_2. \end{aligned}$$

Setting $\delta_1 = \hat{\varepsilon}/2$ and $\delta_2 = \hat{\varepsilon}\beta^2\gamma/\tilde{O}(1)$, we conclude that w.h.p. each operation introduces at most $\hat{\varepsilon}$ additive error in the maintained estimate for $\left\langle \tilde{\pi}^C \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}} \right), \phi \right\rangle$.

Per Lemma 4.5.4, estimating one coordinate of the demand projection requires

$$\tilde{O}(\delta_1^{-4}\beta^{-6} + \delta_1^{-2}\beta^{-2}\gamma^{-2}) = \tilde{O}(\hat{\varepsilon}^{-4}\beta^{-6} + \hat{\varepsilon}^{-2}\beta^{-2}\gamma^{-2})$$

random walks, and estimating $\tilde{\pi}^C \left(\mathbf{B}^\top \frac{\mathbf{q}_S}{\sqrt{r}} \right)$, per Lemma 4.5.7, requires

$$\tilde{O}(\delta_2^{-2}) = \tilde{O}(\hat{\varepsilon}^{-2}\beta^{-4}\gamma^{-2})$$

random walks. This concludes the proof. □

9.2.6 The CHECKER Data Structure

Theorem 9.2.12 (Theorem 3, [73]). *There is a CHECKER data structure supporting the following operations with the given runtimes against oblivious adversaries, for parameters $0 < \beta_{\text{CHECKER}}, \varepsilon < 1$ such that $\beta_{\text{CHECKER}} \geq \tilde{\Omega}(\varepsilon^{-1/2}/m^{1/4})$.*

- INITIALIZE($\mathbf{f}, \varepsilon, \beta_{\text{CHECKER}}$): Initializes the data structure with slacks $\mathbf{s}^+ = \mathbf{u} - \mathbf{f}$, $\mathbf{s}^- = \mathbf{f}$, and resistances $\mathbf{r} = \frac{1}{(s^+)^2} + \frac{1}{(s^-)^2}$. Runtime: $\tilde{O}(m\beta_{\text{CHECKER}}^{-4}\varepsilon^{-4})$.
- UPDATE(e, \mathbf{f}'): Set $s_e^+ = u_e - f'_e$, $s_e^- = f'_e$, and $r_e = \frac{1}{(s_e^+)^2} + \frac{1}{(s_e^-)^2}$. Runtime: Amortized $\tilde{O}(\beta_{\text{CHECKER}}^{-2}\varepsilon^{-2})$.
- TEMPORARYUPDATE(e, \mathbf{f}'): Set $s_e^+ = u_e - f'_e$, $s_e^- = f'_e$, and $r_e = \frac{1}{(s_e^+)^2} + \frac{1}{(s_e^-)^2}$. Runtime: Worst case $\tilde{O}((K\beta_{\text{CHECKER}}^{-2}\varepsilon^{-2})^2)$, where K is the number of TEMPORARYUPDATES that have not been rolled back using ROLLBACK. All TEMPORARYUPDATES should be rolled back before the next call to UPDATE.
- ROLLBACK(): Rolls back the last TEMPORARYUPDATE if it exists. The runtime is the same as the original operation.
- CHECK($e, \boldsymbol{\pi}_{\text{old}}$): Returns \tilde{f}_e such that $\sqrt{r_e}|\tilde{f}_e - \tilde{f}_e^*| \leq \varepsilon$, where

$$\tilde{\mathbf{f}}^* = \delta g(\mathbf{s}) - \delta \mathbf{R}^{-1} \mathbf{B} \left(\mathbf{B}^\top \mathbf{R}^{-1} \mathbf{B} \right)^+ \mathbf{B}^\top g(\mathbf{s}),$$

for $\delta = 1/\sqrt{m}$. Additionally, a vector $\boldsymbol{\pi}_{\text{old}}$ that is supported on C such that

$$\mathcal{E}_{\mathbf{r}} \left(\boldsymbol{\pi}_{\text{old}} - \boldsymbol{\pi}^C \left(\mathbf{B}^\top g(\mathbf{s}) \right) \right) \leq \varepsilon^2 m/4$$

is provided, where C is the vertex set of the dynamic sparsifier in the DYNAMICSC that is maintained internally. Runtime: Worst case $\tilde{O}((\beta_{\text{CHECKER}} m + (K\beta_{\text{CHECKER}}^{-2}\varepsilon^{-2})^2)\varepsilon^{-2})$, where K is the number of TEMPORARYUPDATES that have not been rolled back. Additionally, the output of CHECK(e) is independent of any previous calls to CHECK.

Finally, all calls to CHECK return valid outputs with high probability. The total number of UPDATES and TEMPORARYUPDATES that have not been rolled back should always be $O(\beta_{\text{CHECKER}} m)$.

This theorem is from [73]. The only difference is in the guarantee of CHECK. We will now show how it can be implemented. Let

$$\tilde{\mathbf{f}}^* = \delta g(\mathbf{s}) - \delta \mathbf{R}^{-1} \mathbf{B} \mathbf{L}^+ \mathbf{B}^\top g(\mathbf{s}).$$

Let DYNAMICSC be the underlying Schur complement data structure. We first add the endpoints u, w of e as terminals by calling

$$\text{DYNAMICSC.TEMPORARYADDTERMINALS}(\{u, w\})$$

so that the new Schur complement is on the vertex set $C' = C \cup \{u, w\}$. Then, we set

$$\boldsymbol{\phi} = -\widetilde{SC}^+ \boldsymbol{\pi}_{\text{old}}$$

and $\tilde{f}_e = (\phi_u - \phi_w)/\sqrt{r_e}$, where \widetilde{SC} is the output of DYNAMICSC. \widetilde{SC} (\cdot). Equivalently, note that

$$\tilde{f}_e = \delta \cdot \mathbf{1}_e^\top \mathbf{R}^{-1} \mathbf{B} \mathbf{L}^+ \boldsymbol{\pi}_{\text{old}}.$$

We will show that $\sqrt{r_e}|\tilde{f}_e - \tilde{f}_e^*| \leq \varepsilon$.

We write

$$\begin{aligned} & \sqrt{r_e} \left| \tilde{f}_e - \tilde{f}_e^* \right| \\ & \leq \left\| \delta \sqrt{r} g(s) \right\|_\infty + \left\| \delta \mathbf{R}^{-1/2} \mathbf{B} \mathbf{L}^+ \left(\mathbf{B}^\top g(s) - \pi^{C'} \left(\mathbf{B}^\top g(s) \right) \right) \right\|_\infty \\ & + \left\| \delta \mathbf{R}^{-1/2} \mathbf{B} \mathbf{L}^+ \left(\pi^{C'} \left(\mathbf{B}^\top g(s) \right) - \pi^C \left(\mathbf{B}^\top g(s) \right) \right) \right\|_\infty + \left\| \delta \mathbf{R}^{-1/2} \mathbf{B} \mathbf{L}^+ \left(\pi^C \left(\mathbf{B}^\top g(s) \right) - \pi_{old} \right) \right\|_\infty. \end{aligned}$$

For the first term,

$$\left\| \delta \sqrt{r} g(s) \right\|_\infty = \left\| \delta \frac{\frac{1}{s^+} - \frac{1}{s^-}}{\sqrt{\frac{1}{(s^+)^2} + \frac{1}{(s^-)^2}}} \right\|_\infty \leq \delta \leq \varepsilon/10.$$

Now, by the fact that C' is a β_{CHECKER} -congestion reduction subset by definition in DYNAMICSC, Lemma 4.4.3 immediately implies that the second term is $\leq \delta \cdot \tilde{O}(\beta_{\text{CHECKER}}^{-2}) \leq \varepsilon/10$.

For the third term, we apply Lemma 4.4.12, which shows that

$$\begin{aligned} & \left\| \delta \mathbf{R}^{-1/2} \mathbf{B} \mathbf{L}^+ \left(\pi^{C'} \left(\mathbf{B}^\top g(s) \right) - \pi^C \left(\mathbf{B}^\top g(s) \right) \right) \right\|_\infty \\ & \leq \delta \cdot \sqrt{\mathcal{E}_r \left(\pi^{C'} \left(\mathbf{B}^\top g(s) \right) - \pi^C \left(\mathbf{B}^\top g(s) \right) \right)} \\ & \leq \delta \cdot \tilde{O}(\beta_{\text{CHECKER}}^{-2}) \\ & \leq \varepsilon/10, \end{aligned}$$

as the resistances don't change and we only have two terminal insertions from C to C' .

Finally, the fourth term is

$$\left\| \delta \mathbf{R}^{-1/2} \mathbf{B} \mathbf{L}^+ \left(\pi^C \left(\mathbf{B}^\top g(s) \right) - \pi_{old} \right) \right\|_\infty \leq \delta \cdot \sqrt{\mathcal{E}_r \left(\pi^C \left(\mathbf{B}^\top g(s) \right) - \pi_{old} \right)} \leq \delta \cdot \varepsilon \sqrt{m}/2 = \varepsilon/2.$$

We conclude that $\sqrt{r_e} \left| \tilde{f}_e - \tilde{f}_e^* \right| \leq \varepsilon$. Finally, we call DYNAMICSC.ROLLBACK to undo the terminal insertions.

The runtime of this operation is dominated by the call to DYNAMICSC. $\widetilde{SC}()$, which takes time $\tilde{O}((\beta_{\text{CHECKER}} m + (K \beta_{\text{CHECKER}}^{-2} \varepsilon^{-2})^2) \varepsilon^{-2})$.

9.3 Appendix for Chapter 5

9.3.1 Approximating Submodular Set Functions with Graph Cuts

It is shown in [45] that submodular set functions defined on a ground set of n elements can be $O(n^2)$ approximated by directed graph cuts. We state this fact as a lemma, and we include the proof for completeness in Section 9.3.1 below.

Definition 9.3.1. *Given a submodular set function $F : 2^V \rightarrow \mathbb{R}$, such that $F(\emptyset) = F(V) = 0$, and a weighted directed graph $G = (V, E, c)$, we say that the cut function of G α -approximates F if*

$$\frac{1}{\alpha} c^+(A) \leq F(A) \leq c^+(A), \text{ for all } A \subseteq V.$$

Lemma 9.3.2. *Let $V = \{1, \dots, n\}$, and let $F : V \rightarrow \mathbb{R}$ be a non-negative submodular set function, satisfying $F(\emptyset) = F(V) = 0$. Using $O(n^2)$ calls to a minimization oracle which can compute for all*

Algorithm 22 Approximate non-negative submodular function $F = F^0 - w^0$ by graph cuts, where $F^0 : 2^V \rightarrow \mathbb{R}_{\geq 0}$ is the initial submodular function and the shift vector $w^0 : V \rightarrow \mathbb{R}$ is given as input

```

1: function GRAPHAPPROX( $w^0 : V \rightarrow \mathbb{R}$ )
2:   Call GRAPHAPPROXSHIFTED( $F^0 - w^0$ )

3: function GRAPHAPPROXSHIFTED( $F : 2^V \rightarrow \mathbb{R}_{\geq 0}$ )
4:   Let  $E = \{(u, v) \in V \times V : u \neq v\}$ .
5:   for  $u, v \in V : u \neq v$  do
6:     Compute  $w_{uv} = \min_{\substack{A \subseteq V: \\ u \in A, v \notin A}} F(A)$ .
7:      $c_{uv} = w_{uv}$ .
   return  $G = (V, E, c)$ 

```

pairs $u, v \in V$

$$\min_{\substack{A \subseteq V \\ u \in A, v \notin A}} F(A)$$

one can compute a weighted directed graph $G(V, E, c)$ such that its cut function

$$c^+(A) := \sum_{\substack{(u,v) \in E: \\ u \in A, v \notin A}} c_{uv}$$

$(n^2/4)$ -approximates G . In other words, for any $A \subseteq V$ the size of the graph cut satisfies:

$$\frac{1}{n^2/4} \cdot c^+(A) \leq F(A) \leq c^+(A) .$$

Furthermore, if F takes only values that are discrete multiples of Δ , i.e. $F(A) \in \Delta \cdot \mathbb{Z}_{\geq 0}$ for all A , then all elements of c are discrete multiples of Δ .

As a consequence, we obtain a good approximation by graph cuts for decomposable submodular functions where each component in the decomposition acts on few elements, i.e., when $F_i(A) = F_i(A \cap V_i)$ for some $V_i \subseteq V$.

Lemma 9.3.3. Let $V = \{1, \dots, n\}$, and let $F_i : V_i \rightarrow \mathbb{R}$, $V_i \subseteq V$ be non-negative submodular set functions, with $F_i(\emptyset) = F_i(V_i)$, for $i = 1, \dots, r$. In the time required to compute for all pairs $u \neq v \in V$ and for all $1 \leq i \leq r$

$$\min_{\substack{A \subseteq V_i \\ u \in A, v \notin A}} F_i(A)$$

one can compute a weighted directed graph $G(V, E, c)$ such that its cut function $(M^2/4)$ -approximates $\sum_{i=1}^r F_i$, where $M = \max_{i=1, \dots, r} |V_i|$.

Proof. For each i compute the corresponding graph as in Lemma 9.3.2. Then take the union of edges over the same vertex set. \square

We showed that the function $F(A)$ is well approximated by the cut function $c^+(A)$ for the graph we constructed. Note that c^+ is only defined on internal vertices of the graph, excluding s and t . However this does not affect its submodularity. Therefore the submodular base polytopes for the two function approximate each other well.

Lemma 9.3.4. *Let F, G be two submodular functions defined over the same vertex set V such that $F(\emptyset) = G(\emptyset) = 0$, $F(V) = G(V) = 0$, and for any $A \subseteq V$, $\frac{1}{\alpha}G(A) \leq F(A) \leq G(A)$. Then their submodular base polytopes satisfy:*

$$\frac{1}{\alpha}B(G) \subseteq B(F) \subseteq B(G).$$

Proof. Let any $w \in B(F)$. Then $w(V) = F(V) = G(V)$. Furthermore for any set $A \subseteq V$, we have $w(A) \leq F(A) \leq G(A)$. Similarly for any $w \in B(G)$, we have $w(A) \leq G(A) \leq \alpha F(A)$, so $B(G) \subseteq \alpha B(F)$, which yields the claim. \square

At this point we can prove that the submodular base polytope of the cut function created in Lemma 9.3.3 approximates the submodular base polytope of the decomposable function $\sum_{i=1}^r F_i$.

Lemma 9.3.5. *Let $V = \{1, \dots, n\}$, let $F_i : V_i \rightarrow \mathbb{R}$, $V_i \subseteq V$ be non-negative submodular set functions, with $F_i(\emptyset) = F_i(V_i)$, for $i = 1, \dots, r$, and let $F = \sum_{i=1}^r F_i$. In the time required to solve $\min_{A \subseteq V_i: u \in A, v \notin A} F_i(A)$ for all $u, v \in V_i$ and all i , we can compute a weighted directed graph $G = (V, E, c)$ such that the submodular base polytope of the cut function $c^+(A)$ satisfies $\frac{1}{M^2/4}B(c^+) \subseteq B(F) \subseteq B(c^+)$, where $M = \max_{i=1, \dots, r} |V_i|$.*

Proof. The proof follows directly from applying Lemma 9.3.3, followed by Lemma 9.3.4. \square

Proof of Lemma 9.3.2

Proof. To simplify notation let us denote by

$$w_{uv} = \min_{\substack{A \subseteq V \\ u \in A, v \notin A}} F(A),$$

and let T_{uv} be the set achieving this minimum.

Consider the graph defined as follows. For every $u, v \in V$, create an arc (u, v) with weight $c_{uv} = w_{uv}$. By construction all capacities are discrete multiples of Δ .

Now we can prove the lower bound on F . We have that

$$c^+(A) = \sum_{u \in A, v \notin A} c_{uv} \leq \sum_{u \in A, v \notin A} w_{uv} = \sum_{u \in A, v \notin A} F(T_{uv}) \leq \sum_{u \in A, v \notin A} F(A) \leq (n^2/4) F(A).$$

We used the fact that $F(A)$ upper bounds c_{uv} for all $u \in A, v \notin A$. Now we prove the upper bound. For any nonempty set $A \subset V$ we can write $A = \bigcup_{u \in A} \left(\bigcap_{v \in V \setminus A} T_{uv} \right)$. By twice applying Lemma 9.3.6, we obtain that

$$F(A) \leq \sum_{u \in A} \sum_{v \notin A} F(T_{uv}) = c^+(A).$$

Additionally, we have by construction that

$$c^+(\emptyset) = c^+(V) = 0.$$

\square

Lemma 9.3.6. *Let F be a non-negative submodular set function $F : 2^V \rightarrow \mathbb{R}$, and let A_1, \dots, A_t be subsets of V . Then*

$$F\left(\bigcup_{i=1}^t A_i\right) \leq \sum_{i=1}^t F(A_i)$$

and

$$F\left(\bigcap_{i=1}^t A_i\right) \leq \sum_{i=1}^t F(A_i).$$

Proof. We prove by induction on t . If $t = 1$, both inequalities are equalities. Otherwise, suppose they hold for $t - 1$. Let $S = \bigcup_{i=1}^{t-1} A_i$. By submodularity, $F(S \cup A_t) \leq F(S) + F(A_t) - F(S \cap A_t)$. Since F is non-negative, so is $F(S \cap A_t)$, and therefore $F(S \cup A_t) \leq F(S) + F(A_t)$. Applying the induction hypothesis this concludes the first part of the proof.

Similarly, let $S = \bigcap_{i=1}^{t-1} A_i$. By submodularity, $F(S \cap A_t) \leq F(S) + F(A_t) - F(S \cup A_t)$. Since F is non-negative, so is $F(S \cup A_t)$, and therefore $F(S \cap A_t) \leq F(S) + F(A_t)$. Again, applying the induction hypothesis this concludes the second part of the proof. \square

9.3.2 Parametric Submodular Minimization via Optimization on the Base Polytope

In this section, for completeness, we provide a proof of Lemma 5.2.3, which is based on [15] (see Chapter 8). In addition, we provide error analysis for reductions between approximate solutions to the combinatorial parametric submodular minimization problem, its continuous version involving the Lovász extension, and the dual formulation on the base polytope.

Proof of Lemma 5.2.3. Given any point x , let $\beta \leq \min\{0, \min_i x_i\}$. Applying the definition of the Lovász extension, and the fundamental theorem of calculus, we can write:

$$\begin{aligned} f(x) + \sum_{i \in V} \psi_i(x_i) &= \int_0^\infty F(\{i : x_i \geq t\}) dt + \int_\beta^0 (F(\{i : x_i \geq t\}) - F(V)) dt \\ &+ \sum_{i \in V} \psi_i(\beta) + \int_\beta^\infty \sum_{i: x_i \geq t} \psi'_i(t) dt \\ &= \int_\beta^\infty \left(F(\{i : x_i \geq t\}) + \sum_{i: x_i \geq t} \psi'_i(t) \right) dt + \sum_{i \in V} \psi_i(\beta) - \beta F(V). \end{aligned}$$

Note that we crucially used the fact that the parametric term $\sum_i \psi'_i(t)$ is separable.

Next we show that if the optimal sets A^α were different from those defined in (5.3), then we could obtain a different iterate x' such that $f(x') + \sum_{i \in V} \psi_i(x'_i) \leq f(x^*) + \sum_{i \in V} \psi_i(x^*_i)$. However, since ψ is strictly convex, the minimizer of $f(x) + \sum_{i \in V} \psi_i(x_i)$ is unique. This gives a contradiction leading us to the desired conclusion.

Indeed, let $x'_i = \sup_{i \in A^\alpha} \alpha$. By the strict convexity property of ψ_i , we have that for any $\alpha > \beta$, $A^\alpha \subseteq A^\beta$, which we reprove for completeness in Lemma 9.3.7.

Using this fact, we know that if A^α are the optimizers of $F(A) + \sum_{i \in V} \psi_i(\alpha)$, then we can write:

$$\int_\beta^\infty \left(F(A^t) + \sum_{i \in A^t} \psi'_i(t) \right) dt = \int_\beta^\infty \left(F(\{i : x'_i \geq t\}) + \sum_{i: x'_i \geq t} \psi'_i(t) \right) dt.$$

Since by the optimality of A^t we have that

$$f(A^t) + \sum_{i \in A^t} \psi_i(t) \leq F(\{i : x_i \geq t\}) + \sum_{i: x_i \geq t} \psi'_i(t),$$

it means that letting $\beta = \min\{0, \min_i x'_i, \min_i x_i^*\}$,

$$\int_{\beta}^{\infty} \left(F(\{i : x'_i \geq t\}) + \sum_{i: x'_i \geq t} \psi'_i(t) \right) dt \leq \int_{\beta}^{\infty} \left(F(\{i : x_i^* \geq t\}) + \sum_{i: x_i^* \geq t} \psi'_i(t) \right) dt$$

and therefore

$$f(x') + \sum_{i \in V} \psi_i(x') \leq f(x^*) + \sum_{i \in V} \psi_i(x^*),$$

which concludes the proof. \square

Lemma 9.3.7. *Let $F : 2^V \rightarrow \mathbb{R}$ be a submodular set function, and let $\psi_i : \mathbb{R} \rightarrow \mathbb{R}$ be a family of strictly convex functions, for $i \in V$. Let $F_{\alpha}(A) = F(A) + \sum_{i \in A} \psi'_i(\alpha)$, and $A^{\alpha} = \arg \min_{A \subseteq V} F_{\alpha}(A)$. If $\alpha > \beta$, then $A^{\alpha} \subseteq A^{\beta}$.*

Proof. By optimality we have that

$$F(A^{\alpha}) + \sum_{i \in A^{\alpha}} \psi'_i(\alpha) \leq F(A^{\alpha} \cap A^{\beta}) + \sum_{i \in A^{\alpha} \cap A^{\beta}} \psi'_i(\alpha)$$

and

$$F(A^{\beta}) + \sum_{i \in A^{\beta}} \psi'_i(\beta) \leq F(A^{\alpha} \cup A^{\beta}) + \sum_{i \in A^{\alpha} \cup A^{\beta}} \psi'_i(\beta).$$

Summing up we obtain that

$$\begin{aligned} & \sum_{i \in A^{\alpha}} \psi'_i(\alpha) - \sum_{i \in A^{\alpha} \cap A^{\beta}} \psi'_i(\alpha) + \sum_{i \in A^{\beta}} \psi'_i(\beta) - \sum_{i \in A^{\alpha} \cup A^{\beta}} \psi'_i(\beta) \\ & \leq F(A^{\alpha} \cap A^{\beta}) + F(A^{\alpha} \cup A^{\beta}) - F(A^{\alpha}) - F(A^{\beta}) \\ & \leq 0, \end{aligned}$$

where we used submodularity in the last step. Therefore

$$\sum_{i \in A^{\alpha} \setminus A^{\beta}} (\psi'_i(\alpha) - \psi'_i(\beta)) \leq 0.$$

Hence we conclude that $A^{\alpha} \setminus A^{\beta} = \emptyset$, since by strict convexity we have that for all i , $\psi'_i(\alpha) > \psi'_i(\beta)$, which would make the term above strictly positive had there been any elements in the set difference. \square

Next we perform a careful error analysis to bound the total error we incur in the case where the iterate we consider is not an exact minimizer of (5.2), but has some small error in norm.

Lemma 9.3.8. *Under the conditions from Lemma 5.2.3, let $\tilde{x} \in \mathbb{R}$ be a point satisfying $\|\tilde{x} - x^*\| \leq \varepsilon$, where x^* is the minimizer of (5.2). Let the sets*

$$\tilde{A}^{\alpha} = \{i : \tilde{x}_i \geq \alpha\}.$$

If ψ_i is σ -strongly convex, for all i , and $\max_{A \subseteq V} F(A) - \min_{A' \subseteq V} F(A') \leq M$, then:

$$F(\tilde{A}^\alpha) + \sum_{i \in \tilde{A}^\alpha} \psi'_i(\alpha) \leq F(A^\alpha) + \sum_{i \in A^\alpha} \psi'_i(\alpha) + Mn^{3/2}\varepsilon + \beta\varepsilon^2/2.$$

Proof. First, using the smoothness of ψ_i we prove that

$$f(\tilde{x}) + \sum_i \psi_i(\tilde{x}_i) \leq f(x^*) + \sum_i \psi_i(x_i^*) + Mn^{3/2}\varepsilon + \beta\varepsilon^2/2.$$

To prove this we first note that f is Lipschitz, since we can use the fact that entries of the gradient of the Lovász extension consist of differences between F evaluated at different subsets of V . Hence for any x , $|\nabla_i f(x)| \leq \max_{A \subseteq V} F(A) - \min_{A' \subseteq V} F(A') \leq M$, and thus $\|\nabla f(x)\| \leq M\sqrt{n}$. Therefore

$$f(\tilde{x}) - f(x^*) \leq M\sqrt{n}\|\tilde{x} - x^*\| \leq M\sqrt{n}\varepsilon.$$

Secondly, we use the smoothness of ψ_i , to obtain that

$$\psi_i(\tilde{x}_i) \leq \psi_i(x_i^*) + \psi'_i(x_i^*)(\tilde{x}_i - x_i^*) + \frac{\beta}{2}(\tilde{x}_i - x_i^*)^2$$

Using Lemma 9.3.10 we see that $\psi'_i(x_i^*) = -w_i^*$, where w^* is the optimizer of a certain function over the base polytope $B(F)$. By the definition of $B(F)$ we have $w_i^* \leq F(\{i\}) \leq M$ and $-w_i^* + \sum_{j \neq i} w_j^* = F(V)$, so $-w_i^* \geq F(V) - \sum_{j \neq i} F(\{j\}) \geq -M(n-1)$. Thus, by applying Cauchy-Schwarz, we have

$$\sum_{i \in V} \psi'_i(x_i^*)(\tilde{x}_i - x_i^*) \leq \max_i |\psi'_i(x_i^*)| \sqrt{n} \cdot \|\tilde{x} - x^*\| \leq M(n-1)n^{1/2}\varepsilon,$$

and thus

$$\sum_{i \in V} \psi'_i(\tilde{x}_i) - \psi'_i(x_i^*) \leq M(n-1)n^{1/2}\varepsilon + \beta\varepsilon^2/2.$$

Combining with the bound on $f(\tilde{x})$, we obtain our claimed error in function value.

Now we can finalize the argument. Following the proof of Lemma 5.2.3 we write $f(\tilde{x}) + \sum_{i \in V} \psi'_i(\tilde{x}_i)$ as an integral, and similarly for x^* , to conclude that for $\beta = \min\{0, \min_i \tilde{x}_i, \min_i x_i^*\}$,

$$\int_\beta^\infty \left(F(\tilde{A}^t) + \sum_{i \in \tilde{A}^t} \psi'_i(t) \right) dt \leq \int_\beta^\infty \left(F(A^t) + \sum_{i \in A^t} \psi'_i(t) \right) dt + Mn^{3/2}\varepsilon + \beta\varepsilon^2/2.$$

Since by definition A^t minimizes $\sum_{i \in A^t} \psi'_i(t)$, we conclude that for all t ,

$$F(\tilde{A}^t) + \sum_{i \in \tilde{A}^t} \psi'_i(t) \leq F(A^t) + \sum_{i \in A^t} \psi'_i(t) + Mn^{3/2}\varepsilon + \beta\varepsilon^2/2.$$

□

We can also show that if we obtain an approximate minimizer of the dual problem (5.4) over $B(F)$, we can use it to recover an approximate minimizer of the primal problem (5.2).

Lemma 9.3.9. *Let w^* be the minimizer of the dual problem (5.4), and let x^* be the minimizer*

of the primal problem (5.2). If $w \in B(F)$ such that

$$\sum_{i \in V} \psi_i^*(-w_i) \leq \sum_{i \in V} \psi_i^*(-w_i^*) + \varepsilon,$$

then the point $x \in \mathbb{R}^n$ where $x_i = (\psi_i^*)'(-w_i)$ satisfies

$$\|x - x^*\| \leq \sqrt{\frac{2L\varepsilon}{\sigma^2}}.$$

Proof. By Lemma 9.3.10 we know that x^* and w^* are related via $x_i^* = (\psi_i^*)'(-w_i^*)$. Therefore we can write

$$|x_i - x_i^*| = |(\psi_i^*)'(-w_i) - (\psi_i^*)'(-w_i^*)| \leq \frac{1}{\sigma} |w_i - w_i^*|,$$

where in the last inequality we used the fact that ψ_i is σ -strongly convex, and hence ψ_i^* is $1/\sigma$ -smooth [145, 92]. Next we show that $|w_i - w_i^*|$ is bounded by a function of ε .

Since by assumption ψ_i is L -smooth, its dual ψ_i^* is $1/L$ -strongly convex. Therefore we have that, for all i :

$$\psi_i^*(-w_i) \geq \psi_i^*(-w_i^*) + (\psi_i^*)'(-w_i^*) \cdot (-w_i - (-w_i^*)) + \frac{\sigma}{2} (w_i^* - w_i)^2.$$

Furthermore, since w^* is an optimizer over $B(F)$, we know by first-order optimality that for any $w \in B(F)$:

$$\sum_{i \in V} (\psi_i^*)'(-w_i^*) \cdot (-w_i - (-w_i^*)) \geq 0,$$

i.e. slightly moving the point from $-w^*$ towards $-w$ can only increase function value. Thus we obtain that

$$\sum_{i \in V} \psi_i^*(-w_i) \geq \sum_{i \in V} \psi_i^*(-w_i^*) + \frac{1}{2L} \sum_{i \in V} (w_i^* - w_i)^2.$$

Combining with the hypothesis, this implies that

$$\frac{1}{2L} \sum_{i \in V} (w_i^* - w_i)^2 \leq \varepsilon,$$

and therefore

$$\|x - x^*\|^2 \leq \frac{1}{\sigma^2} \sum_{i \in V} (w_i - w_i^*)^2 \leq \frac{2L\varepsilon}{\sigma^2},$$

which implies the claimed result. \square

As a corollary of the previous lemmas, we see that an approximate solution to the dual problem (5.4) yields an approximate solution to the original parametric problem (5.1).

Corollary 9.3.1. *Let $F : 2^V \rightarrow \mathbb{R}$ be a non-negative submodular set function, and let the family of parametric problems defined in (5.1). Let $w \in B(F)$ such that*

$$\sum_{i \in V} \psi_i^*(-w_i) \leq \sum_{i \in V} \psi_i^*(-w_i^*) + \varepsilon,$$

where w^* is the true minimizer of the dual problem (5.4). Then for any α , the set

$$\tilde{A}^\alpha = \{i : \psi_i^*(-w_i) \geq \alpha\}$$

satisfies

$$F_\alpha(\tilde{A}^\alpha) \leq F_\alpha(A^\alpha) + \sqrt{\varepsilon} \cdot Mn^{3/2} \sqrt{2L/\sigma^2} + \varepsilon \cdot (L/\sigma)^2,$$

where $A^\alpha = \arg \min_{A \subseteq V} F_\alpha(A)$.

Proof. From Lemma 9.3.9 we know that the hypothesis implies that the point x where $x_i = (\psi_i^*)(-w_i)$ satisfies $\|x - x^*\| \leq \sqrt{2L\varepsilon/\sigma^2}$. Applying Lemma 9.3.8 we thus obtain that the sets constructed satisfy

$$F_\alpha(\tilde{A}^\alpha) \leq F_\alpha(A^\alpha) + Mn^{3/2} \sqrt{2L\varepsilon/\sigma^2} + L/2 \cdot (2L\varepsilon/\sigma^2),$$

which yields our claim. \square

The following helper lemma shows that we can efficiently convert between (exact) solutions to the primal and dual problems (5.2) and (5.4). Using standard techniques we can prove that these also enable us to convert between suboptimal solutions, while satisfying certain error bounds.

Lemma 9.3.10. *Let x^* be the (unique) minimizer of (5.2), and let w^* be the minimizer of (5.4). Then $w_i^* = -\psi_i'(x_i)$ and $(\psi_i^*)'(-w_i) = x_i$, for all $i \in V$.*

Proof. We use the dual characterization of f and Sion's theorem, to write

$$\min_{x \in \mathbb{R}^n} f(x) + \sum_{i \in V} \psi_i(x_i) = \min_{x \in \mathbb{R}^n} \max_{w \in B(F)} \langle w, x \rangle + \sum_{i \in V} \psi_i(x_i) = \max_{w \in B(F)} \min_{x \in \mathbb{R}^n} \langle w, x \rangle + \sum_{i \in V} \psi_i(x_i).$$

Since each ψ_i acts on a different coordinate we can write the inner minimization problem as

$$\min_{x \in \mathbb{R}^n} \sum_{i \in V} (w_i x_i + \psi_i(x_i)) = - \sum_{i \in V} \psi_i^*(-w_i),$$

where we applied the definition of the Fenchel dual. Furthermore by standard convex analysis [23, 140], as ψ_i' ranges from $-\infty$ to ∞ for each i we have that $(\psi_i^*)'(-w_i) = x_i$, and similarly $\psi_i'(x_i) = -w_i$.

Thus we can equivalently write (5.2) as

$$\max_{w \in B(F)} - \sum_{i \in V} \psi_i^*(-w_i).$$

By the previous observation, the optima are thus related via $(\psi_i^*)'(-w_i^*) = x_i^*$, and similarly $\psi_i'(x_i^*) = -w_i^*$. \square

9.3.3 Parametric s - t Cuts

In this section we show how to solve the parametric minimum cut problem by efficiently using a maximum flow oracle. In Section 5.4 we show how to convert the solution obtained by this combinatorial routine to a nearly-optimal solution to a related optimization problem on the submodular base polytope of the corresponding cut function.

In the parametric min s , t -cut problem, we are given a directed network $G = (V, E)$ with two distinguished vertices: a *source* $s \in V$, and a *sink* $t \in V$, $s \neq t$. The capacities of individual edges of G are nonnegative functions of a real parameter λ in some possibly infinite domain $\mathbb{D} \subseteq \mathbb{R}$ (as opposed to constants in the classical setting of min s , t -cut). Following [72], we assume that the capacities of edges $sv \in E$ are nondecreasing in λ and the capacities of edges $vt \in E$ are nonincreasing in λ . The capacities of all other edges of G are constant.

We denote by $c_\lambda(uv) : \mathbb{D} \rightarrow \mathbb{R}$ the capacity function of an edge $uv \in E$. Moreover, we assume that these edge capacity functions can be evaluated for arbitrary λ in constant time.

Roughly speaking, the goal of the parametric min s, t -cut problems is to compute a representation of min s, t -cut for all the possible parameters λ . Before we precisely define what this means, let us introduce some more notation and state some useful properties of (parametric) min-cuts.

Denote by $\text{cap}(G)$ the capacity of a min s, t -cut in G . Let $G_{\lambda'}$ be the graph with all the parameterized capacities replaced with the corresponding values for $\lambda = \lambda'$. For any $S, s \in S \subseteq V \setminus \{t\}$, let $c_{\lambda}(S)$ be the capacity function of S , i.e., the sum of capacity functions $c_{\lambda}(uv)$ through all uv with $u \in S$ and $v \in V \setminus S$.

Lemma 9.3.2 ([62]). *For any G , there exists a unique minimal minimum s, t -cut (S, T) with $|S|$ smallest possible, such that for any min s, t -cut (S', T') of G we have $S \subseteq S'$. Given any maximum s, t -flow f in G , such a cut can be computed from f in $O(m)$ time.*

Proof. Let G_f be the residual network associated with flow f . We let S be the set of vertices reachable from s in G_f (via edges with positive capacity). As proven by Ford and Fulkerson [62, Theorem 5.5], S defined this way does not depend on the chosen maximum flow f , and $S \subseteq S'$ holds. Clearly, given f , S can be found using any graph search algorithm. \square

Ford and Fulkerson [62, Corollary 5.4] showed that for any two min s, t -cuts $(S_1, T_1), (S_2, T_2)$ of G , $(S_1 \cap S_2, T_1 \cup T_2)$ is also a min s, t -cut of G . Gallo et al. [72, Lemma 2.8] gave the following generalization of this property to parametric min s, t -cuts.

Lemma 9.3.3 ([72]). *Let $\lambda_1 \leq \lambda_2$. For $i = 1, 2$, let $(S_{\lambda_i}, T_{\lambda_i})$ be some min s, t -cut in G_{λ_i} . Then $(S_{\lambda_1} \cap S_{\lambda_2}, T_{\lambda_1} \cup T_{\lambda_2})$ is a min s, t -cut in G_{λ_1} .*

Our algorithm will use the following crucial property of parametric minimal min s, t -cuts.

Lemma 9.3.4. *Let $\lambda_1 \leq \lambda_2$. For $i = 1, 2$, let $(S_{\lambda_i}, T_{\lambda_i})$ be the unique minimal min s, t -cut in G_{λ_i} . Then $S_{\lambda_1} \subseteq S_{\lambda_2}$.*

Proof. The uniqueness of S_{λ_1} and S_{λ_2} follows by Lemma 9.3.2 applied to G_{λ_1} and G_{λ_2} , respectively. By Lemma 9.3.3, $(S_{\lambda_1} \cap S_{\lambda_2}, T_{\lambda_1} \cup T_{\lambda_2})$ is a min s, t -cut in G_{λ_1} . By Lemma 9.3.2, we have $S_{\lambda_1} \subseteq S_{\lambda_1} \cap S_{\lambda_2}$. It follows that $S_{\lambda_1} \subseteq S_{\lambda_2}$. \square

Now given Lemma 9.3.4, we can formally state our goal in this section, which is to compute a parametric min s, t -cut defined as follows. Let $\lambda_{\min} \in \mathbb{D}$ be such that the minimal min s, t -cuts of $G_{\lambda_{\min}}$ and $G_{\lambda'}$ are equal for all $\lambda' \in \mathbb{D}$, $\lambda' < \lambda_{\min}$. Similarly, let $\lambda_{\max} \in \mathbb{D}$ be such that the minimal min s, t -cuts of $G_{\lambda_{\max}}$ and $G_{\lambda'}$ are equal for all $\lambda' \in \mathbb{D}$ with $\lambda' > \lambda_{\max}$. We will consider λ_{\min} and λ_{\max} additional inputs to our problem.

For simplicity, in the remaining part of this section we denote by S_{λ} and T_{λ} the s -side and the t -side (resp.) of the minimal min- s, t -cut of G_{λ} .

Definition 9.3.11 (Parametric min s, t -cut). *Let $\Lambda = \{\lambda_1, \dots, \lambda_k\} \subseteq \mathbb{D}$, where $k \leq n - 1$ and $\lambda_{\min} < \lambda_1 < \dots < \lambda_k \leq \lambda_{\max}$. Let $\lambda_0 = \lambda_{\min}$. Let $\tau : V \rightarrow \Lambda \cup \{\lambda_{\min}, \infty\}$ be such that $\tau(s) = \lambda_{\min}$ and $\tau(t) = \infty$. Let $S(z) = \{v \in V : \tau(v) \leq z\}$. A pair (Λ, τ) is a parametric min s, t -cut of G if:*

1. For $i = 0, \dots, k - 1$, $S(\lambda_i)$ is a minimal min s, t -cut of $G_{\lambda'}$ for all $\lambda' \in [\lambda_i, \lambda_{i+1}) \cap \mathbb{D}$.
2. $S(\lambda_k)$ is a minimal min s, t -cut of $G_{\lambda_{\max}}$.
3. For $i = 0, \dots, k - 1$, $S(\lambda_i) \subsetneq S(\lambda_{i+1})$.

It will also prove useful to define an approximate version of parametric min s, t -cut.

Definition 9.3.12 (ε -approximate parametric min s, t -cut). Let Λ , τ , and $S : \mathbb{D} \rightarrow 2^V$ be as in Definition 9.3.11. A pair (Λ, τ) is called an ε -approximate parametric min s, t -cut of G if:

1. For $i = 0, \dots, k-1$, $S(\lambda_i)$ is a minimal min s, t -cut of $G_{\lambda'}$ for all $\lambda' \in [\lambda_i, \lambda_{i+1} - \varepsilon) \cap \mathbb{D}$.
2. $S(\lambda_k)$ is a minimal min s, t -cut of $G_{\lambda_{\max}}$.
3. For $i = 0, \dots, k-1$, $S(\lambda_i) \subsetneq S(\lambda_{i+1})$.

Lemma 9.3.5. Let (Λ, τ) be the parametric min s, t -cut of G . Let $(\Lambda_\varepsilon, \tau_\varepsilon)$ be an ε -approximate parametric min s, t -cut of G . Then for all $v \in V$, $\tau(v) \leq \tau_\varepsilon(v) \leq \tau(v) + \varepsilon$.

Proof. Let $S(z) = \{v \in V : \tau(v) \leq z\}$, and $S_\varepsilon(z) = \{v \in V : \tau_\varepsilon(v) \leq z\}$. First of all, $\tau(v) = \infty$ if and only if $\tau_\varepsilon(v) = \infty$. This is because each of those is equivalent to $v \notin S_{\lambda_{\max}}$. In this case the lemma holds trivially.

So in the following let us assume that $\tau(v)$ and $\tau_\varepsilon(v)$ are both finite. We first prove $\tau_\varepsilon(v) \geq \tau(v)$. If $\tau(v) = \lambda_{\min}$ then this follows by $\tau_\varepsilon(v) \geq \lambda_{\min}$. So suppose $\tau(v) = \lambda$ for some $\lambda \in \Lambda$. Then by item (1) of Definition 9.3.11, for any $\lambda' < \lambda$, $S(\lambda')$ is a minimal min s, t -cut of $G_{\lambda'}$ and $v \notin S(\lambda')$. If we had $\tau_\varepsilon(v) < \tau(v)$, then $S_\varepsilon(\tau_\varepsilon(v))$ would be a minimal min s, t -cut of $G_{\tau_\varepsilon(v)}$ such that $v \in S_{\tau_\varepsilon(v)}$ and $\tau_\varepsilon(v) < \lambda$, a contradiction.

Now let us prove $\tau_\varepsilon(v) \leq \tau(v) + \varepsilon$. To this end, suppose $\tau_\varepsilon(v) > \tau(v) + \varepsilon$. If $\tau_\varepsilon(v) = \lambda_{\min}$, then we have $\lambda_{\min} > \tau(v) + \varepsilon \geq \lambda_{\min} + \varepsilon$, a clear contradiction. So let us assume that $\tau_\varepsilon(v) \in \Lambda_\varepsilon$ and let λ^* be the element preceding $\tau_\varepsilon(v)$ in Λ_ε , or $\lambda^* = \lambda_{\min}$ if no such element exists. We have $v \notin S_{\lambda^*}$ and S_{λ^*} is a minimal min s, t -cut in $G_{\lambda'}$ for $\lambda' = \lambda^*$ and all $\lambda' \in [\lambda^*, \tau_\varepsilon(v) - \varepsilon)$. As a result, for any $\lambda'' < \tau_\varepsilon(v) - \varepsilon$, the minimal min s, t -cut of $G_{\lambda''}$ does not contain v in the s -side. But $\tau(v) < \tau_\varepsilon(v) - \varepsilon$, $v \in S(\tau(v))$, and $S(\tau(v))$ is a minimal min s, t -cut of $G_{\tau(v)}$, a contradiction. \square

Our main result in this section is the following theorem.

Theorem 5.3.2. Let $R = \lambda_{\max} - \lambda_{\min}$ be an integral multiple of $\varepsilon > 0$. Let $T_{\max\text{flow}}(n', m') = \Omega(m' + n')$ be a convex function bounding the time needed to compute maximum flow in a graph with n' vertices and m' edges obtained from G_λ by edge/vertex deletions and/or edge contractions (with merging parallel edges by summing their capacities) for any $\lambda = \lambda_{\min} + \ell\varepsilon$ and any integer $\ell \in [0, R/\varepsilon]$. Then, ε -approximate parametric min s, t -cut in G can be computed in $O(T_{\max\text{flow}}(n, m \log n) \cdot \log \frac{R}{\varepsilon} \cdot \log n)$ time.

The rest of this section is devoted to proving Theorem 5.3.2. For a connected subset $X \subseteq V(G)$, $\{s, t\} \not\subseteq X$, let G/X denote G after merging the vertex set X into a single vertex. If the contracted vertex set X contains s (t), then the resulting vertex inherits the identity of s (t , resp.).

Lemma 9.3.6. Let λ be arbitrary and let (S_λ, T_λ) be the minimal min s, t -cut in G_λ . Then:

1. For any $\lambda' \geq \lambda$, $\text{cap}(G_{\lambda'}) = \text{cap}(G_{\lambda'}/S_\lambda)$.
2. For any $\lambda' \leq \lambda$, $\text{cap}(G_{\lambda'}) = \text{cap}(G_{\lambda'}/T_\lambda)$.

Proof. We only prove item 1, as item 2 is analogous. Since merging vertices is equivalent to connecting them with infinite capacity edges, it cannot decrease the min s, t -cut capacity, i.e., $\text{cap}(G_{\lambda'}) \leq \text{cap}(G_{\lambda'}/S_\lambda)$. On the other hand, by Lemma 9.3.4, the minimal s, t min-cut $(S_{\lambda'}, T_{\lambda'})$ in $G_{\lambda'}$ satisfies $S_\lambda \subseteq S_{\lambda'}$. Hence, the capacity of the s, t -cut $(S_{\lambda'}/S_\lambda, T_{\lambda'})$ in $G_{\lambda'}/S_\lambda$ is the same as the capacity $\text{cap}(G_{\lambda'})$ of $(S_{\lambda'}, T_{\lambda'})$ in $G_{\lambda'}$. Consequently, $\text{cap}(G_{\lambda'}) \geq \text{cap}(G_{\lambda'}/S_\lambda)$. \square

Remark 9.3.7. If (S_λ, T_λ) is a minimal min s, t -cut in G_λ , then $G[S_\lambda]$ is connected by construction (Lemma 9.3.2). However, $G[T_\lambda]$ might in general consist of several connected components if G_λ contains zero-capacity edges. In that case, we can still obtain $G_{\lambda'}/T_\lambda$ above using edge/vertex deletions and edge contractions. Namely, we contract only the connected component A of T_λ that contains t . For any other component C_i ($i = 1, \dots, q$) of T_λ , its incoming edges start in S_λ and all have capacity 0 in G_λ , and thus also in $G_{\lambda'}$ for $\lambda' < \lambda$. Consequently, removing the vertices of $\bigcup_{i=1}^q C_i$ and subsequently contracting A has the same effect on $G_{\lambda'}$ as merging the entire T_λ , i.e., $G_{\lambda'}/T_\lambda = G_{\lambda'}[V \setminus \bigcup_{i=1}^q C_i]/A$.

We use a recursive “divide-and-conquer” algorithm. The input to a recursive procedure `APXPARAMETRICMINCUT` is a graph $G = (V, E)$ with n vertices, m edges, source s and sink t , the parametric capacity function $c_\lambda : E \rightarrow \mathbb{D} \rightarrow \mathbb{R}$, and two parameters $\lambda_{\min}, \lambda_{\max}$ such that ε evenly divides $\lambda_{\max} - \lambda_{\min}$. The output of the procedure is an ε -approximate parametric min s, t -cut $(\{\lambda_1, \dots, \lambda_k\}, \tau)$ as in Definition 9.3.12. By Lemma 9.3.4, $k \leq |V(G)| - 1$.

Algorithm 23 Computing an ε -approximate parametric min s, t -cut.

```

1: Let  $s, t, \varepsilon$  be globally defined.
2: function APXPARAMETRICMINCUT( $G = (V, E), c_\lambda : E \rightarrow \mathbb{D} \rightarrow \mathbb{R}, \lambda_{\min} \in \mathbb{D}, \lambda_{\max} \in \mathbb{D}$ )
3:   if  $|V| \leq 2$  then return  $(\emptyset, \{s \rightarrow \lambda_{\min}, t \rightarrow \infty\})$ 
4:   For any  $\lambda' \in \mathbb{D}$ , let  $c_\lambda[\lambda = \lambda']$  the capacity function  $E \rightarrow \mathbb{R}$  of  $G_{\lambda'}$ 
5:    $S_{\lambda_{\min}} = \text{MINIMALMINCUT}(G, c_\lambda[\lambda = \lambda_{\min}])$ 
6:    $S_{\lambda_{\max}} = \text{MINIMALMINCUT}(G, c_\lambda[\lambda = \lambda_{\max}])$ 
7:   if  $|S_{\lambda_{\min}}| > |V|/2$  then return APXPARAMETRICMINCUT(CONTRACT( $G, c_\lambda, S_{\lambda_{\min}}$ ),  $\lambda_{\min}, \lambda_{\max}$ )
8:   if  $|S_{\lambda_{\max}}| < |V|/2$  then return APXPARAMETRICMINCUT(CONTRACT( $G, c_\lambda, V \setminus$ 
    $S_{\lambda_{\max}}$ ),  $\lambda_{\min}, \lambda_{\max}$ )
9:    $(\lambda_1, \lambda_2) := (\lambda_{\min}, \lambda_{\max})$ 
10:  while  $\lambda_2 - \lambda_1 > \varepsilon$  do
11:     $\lambda' := \lambda_1 + \lfloor (\lambda_2 - \lambda_1)/2\varepsilon \rfloor \cdot \varepsilon$ 
12:     $S_{\lambda'} = \text{MINIMALMINCUT}(G, c_\lambda[\lambda = \lambda'])$ 
13:    if  $|S_{\lambda'}| \geq |V|/2$  then
14:       $\lambda_2 := \lambda'$ 
15:    else
16:       $\lambda_1 := \lambda'$ 
17:  For  $i = 1, 2$ ,  $S_{\lambda_i} := \text{MINIMALMINCUT}(G, c_\lambda[\lambda = \lambda_i])$ 
18:   $(\Lambda_1, \tau_1) = \text{APXPARAMETRICMINCUT}(\text{CONTRACT}(G, c_\lambda, V \setminus S_{\lambda_1}), \lambda_{\min}, \lambda_1)$ 
19:   $(\Lambda_2, \tau_2) = \text{APXPARAMETRICMINCUT}(\text{CONTRACT}(G, c_\lambda, S_{\lambda_2}), \lambda_2, \lambda_{\max})$ 
20:   $\Lambda :=$  if  $|S_{\lambda_1}| = |S_{\lambda_2}|$  then  $\Lambda_1 \cup \Lambda_2$  else  $\Lambda_1 \cup \{\lambda_2\} \cup \Lambda_2$ 
21:   $\tau := \{v \in S_{\lambda_1} \rightarrow \tau_1(v), v \in V \setminus S_{\lambda_2} \rightarrow \tau_2(v), v \in S_{\lambda_2} \setminus S_{\lambda_1} \rightarrow \lambda_2\}$  return  $(\Lambda, \tau)$ 

```

The main idea of the procedure `APXPARAMETRICMINCUT` is to find the (approximately) most balanced minimal s, t -cuts S_{λ_1} and S_{λ_2} and use them to reduce the problem size in the recursive calls significantly. Specifically, we want to find such $\lambda_1 \leq \lambda_2$ that $|S_{\lambda_1}| \leq n/2$, $|S_{\lambda_2}| \geq n/2$ and $\lambda_2 - \lambda_1 = \varepsilon$.

Suppose $n > 2$ as otherwise the problem is trivial. First, we compute minimal min-cuts in $G_{\lambda_{\min}}$ and $G_{\lambda_{\max}}$. This takes two max-flow runs, i.e., $T_{\max\text{flow}}(n, m)$ time, plus $O(m)$ time by Lemma 9.3.2.

It might happen that $|S_{\lambda_{\min}}| \leq |S_{\lambda_{\max}}| < n/2$ or $n/2 < |S_{\lambda_{\min}}| \leq |S_{\lambda_{\max}}|$. In these special cases we can immediately reduce the vertex set by a factor of at least two by contracting $T_{\lambda_{\max}}$ or $S_{\lambda_{\min}}$

```

22: function MINIMALMINCUT( $G = (V, E), c : E \rightarrow \mathbb{R}$ )
23:    $f = \text{MAXFLOW}(G, s, t, c)$  return  $\{v \in V : v \text{ reachable from } s \text{ in the residual network } G_f\}$ 

24: function CONTRACT( $G = (V, E), c_\lambda : E \rightarrow \mathbb{D} \rightarrow \mathbb{R}, X \subseteq V$ ) ▷  $|X \cap \{s, t\}| = 1$ 
25:    $w^* :=$  if  $s \in X$  then  $s$  else  $t$ 
26:    $V' := V \setminus X \cup \{w^*\}$ 
27:    $E' := \emptyset$ 
28:    $c'_\lambda := E' \rightarrow \mathbb{D} \rightarrow \mathbb{R}$ 
29:   for  $uv \in E$  do
30:      $u' :=$  if  $u \in X$  then  $w^*$  else  $u$ 
31:      $v' :=$  if  $v \in X$  then  $w^*$  else  $v$ 
32:     if  $(u', v') \neq (s, t)$  then
33:       if  $u'v' \notin E'$  then
34:          $E' := E' \cup \{u'v'\}$ 
35:          $c'_\lambda(u'v') := c_\lambda(uv)$ 
36:       else
37:          $c'_\lambda(u'v') := c'_\lambda(u'v') + c_\lambda(uv)$  ▷ We add functions here.
return  $(G' = (V', E'), c'_\lambda)$ 

```

respectively, and recurse on the reduced graph. By Lemma 9.3.6 and the definition of $\lambda_{\min}, \lambda_{\max}$ this reduction does not influence the structure of parametric cuts.

So suppose $|S_{\lambda_{\min}}| \leq n/2$ and $|S_{\lambda_{\max}}| \geq n/2$. Set $\lambda_1 = \lambda_{\min}$ and $\lambda_2 = \lambda_{\max}$. So we have $|S_{\lambda_1}| \leq n/2$ and $|S_{\lambda_2}| \geq n/2$ initially. We maintain this invariant and gradually shrink the interval $[\lambda_1, \lambda_2]$ until its length gets precisely ε in a binary search-like way. We repeatedly try the pivot $\lambda' = \lambda_1 + \lfloor (\lambda_2 - \lambda_1)/2 \rfloor \cdot \varepsilon$ and compute $S_{\lambda'}$. If $|S_{\lambda'}| \geq n/2$, we set $\lambda_2 = \lambda'$, and otherwise we set $\lambda_1 = \lambda'$. Note that $\lambda_2 - \lambda_1$ remains an integer multiple of ε at all times. The whole process costs $O(\log[(\lambda_{\max} - \lambda_{\min})/\varepsilon]) = O(\log(R/\varepsilon))$ max-flow executions.

Let $G_1 = G/T_{\lambda_1}$ and $G_2 = G/S_{\lambda_2}$. Note that G_1, G_2 may contain parallel edges or a direct st edge as a result of contraction. Hence, these graphs are first preprocessed by (1) removing self-loops and direct st edges, (2) merging parallel edges by summing their cost functions. The contraction and preprocessing is performed using the procedure CONTRACT. Note that none of these preprocessing steps change the minimal cuts of $G_{1,\lambda}$ or $G_{2,\lambda}$ for any λ : the direct st edges cross *all* s, t -cuts.

Next, we recursively compute ε -approximate parametric min s, t -cut in graphs $G_1 = G/T_{\lambda_1}$ and $G_2 = G/S_{\lambda_2}$. The recursive call on G_1 is made with $(\lambda_{\min}, \lambda_{\max})$ set to $(\lambda_{\min}, \lambda_1)$, whereas the recursive call on G_2 uses $(\lambda_{\min}, \lambda_{\max}) := (\lambda_2, \lambda_{\max})$. Note that indeed we have $G_{1,\lambda_1} = G_{1,\lambda'}$ for $\lambda' > \lambda_1$ as required since the t -side of the minimal min s, t -cut in G_{1,λ_1} contains only t . Similarly, $G_{2,\lambda_2} = G_{2,\lambda'}$ for all $\lambda' < \lambda_2$.

Let $(\Lambda_1, \tau_1) = (\{\lambda_{1,1}, \dots, \lambda_{1,a}\}, \tau_1)$ and $(\Lambda_2, \tau_2) = (\{\lambda_{2,1}, \dots, \lambda_{2,b}\}, \tau_2)$ be the returned ε -approximate parametric min s, t -cuts of G_1 and G_2 respectively. We return (Λ, τ) as the ε -approximate parametric min- s, t -cut of G , where

$$\Lambda = \begin{cases} \Lambda_1 \cup \Lambda_2 & \text{if } |S_{\lambda_1}| = |S_{\lambda_2}| = n/2, \\ \Lambda_1 \cup \{\lambda_2\} \cup \Lambda_2 & \text{otherwise.} \end{cases} \quad \tau(v) = \begin{cases} \tau_1(v) & \text{if } v \in S_{\lambda_1}, \\ \tau_2(v) & \text{if } v \in T_{\lambda_2}, \\ \lambda_2 & \text{otherwise.} \end{cases}$$

Let us now prove the correctness of this algorithm. We proceed by induction on n . For $n \leq 2$

this is trivial, so suppose $n > 3$ and that recursive calls are made. Clearly, $\lambda_{1,a} \leq \lambda_1 < \lambda_2 < \lambda_{2,1}$.

Item (3) of Definition 9.3.12 follows easily by induction and the definition of λ_1, λ_2 . Let $\Lambda = \{\lambda'_1, \dots, \lambda'_k\}$. That $S(\lambda'_i) = \{v \in V : \tau(v) \leq \lambda'_i\}$ is a minimal min s, t -cut of $G_{\lambda'_i}$ for all $\lambda'_i \in \Lambda$ (i.e., item (2) of Definition 9.3.12) follows directly by Lemma 9.3.6 the definitions of λ_1, λ_2 .

Now consider item (1) of Definition 9.3.12. For some $j < k$ we have $\lambda'_j = \lambda_{1,a}$. For all $i = 0, \dots, k-1, i \neq j$, item (1), i.e., that $S_{\lambda'_i}$ is a minimal min s, t -cut for all $\lambda' \in [\lambda'_i, \lambda'_{i+1})$, follows directly inductively.

If $|S_{\lambda_1}| = |S_{\lambda_2}| = n/2$, then $S_{\lambda_{1,a}} = S_{\lambda_2}$. By induction it follows that $S_{\lambda_{1,a}}$ is a minimal min s, t -cut of $G_{\lambda'}$ for all $\lambda' \in [\lambda_2, \lambda_{2,1} - \varepsilon)$, and thus also for all $\lambda' \in [\lambda_{1,a}, \lambda_{2,1} - \varepsilon) = [\lambda'_j, \lambda'_{j+1} - \varepsilon)$.

If, on the other hand, $S_{\lambda_1} \subsetneq S_{\lambda_2}$, then $\lambda'_{j+1} = \lambda_2$. Since $S_{\lambda_{1,a}} = S_{\lambda_1}$, $S_{\lambda_{1,a}}$ is indeed a minimal min s, t -cut for all $\lambda' \in [\lambda_{1,a}, \lambda_2 - \varepsilon) = [\lambda'_j, \lambda'_{j+1} - \varepsilon)$ as $\lambda_2 - \varepsilon = \lambda_1$.

Note that the input graph of each of the recursive calls has at most $n/2 + 1$ vertices. Moreover, by merging the parallel edges (and summing their costs) after the contraction we can guarantee that $|E(G_1)| + |E(G_2)| \leq |E(G)| + n/2$. Indeed, observe that the only edges of G_2 that can also appear in G_1 are those incident to s in G_2 , and there are at most $|T_{\lambda_2}| \leq n/2$ of them.

There is one important technical detail here: even though the individual functions $c_\lambda(uv)$ (for the edges uv of the original input graph G) can be evaluated in constant time, after summing k of such functions in the process this cost can be as much as $\Theta(k)$. We now argue that this cannot happen in our case due to preprocessing G_1 and G_2 , and the evaluation cost is $O(1)$ for all edges in all recursive calls. More concretely, one can show that each edge capacity function in a recursive call can be expressed as the sum of at most one original capacity function $c_\lambda(xy)$ and a real number. Indeed, suppose this is the case for some call with input G . Then, each edge uv of G_1 either (1) is contained in G and has not resulted from merging some parallel edges after contraction, (2) has $u \notin \{s, t\}$ and $v = t$ and resulted from merging edges uz_1, \dots, uz_l such that $z_1, \dots, z_l \in T_{\lambda_1}$. The former case is trivial. In the latter case, for at least $l-1$ of these z_i we have $z_i \neq t$, so $c_\lambda(uz_i)$ is a constant function. For at most one z_j is of the form $c_\lambda(xy) + \Delta$ for some original capacity function $c_\lambda(xy)$ and $\Delta \in \mathbb{R}$. We conclude that the capacity function of uv in G_1 is of the same form and equals $c_\lambda(xy) + \Delta'$, where $\Delta' = \Delta + \sum_{i \neq j} c_\lambda(uz_i) \in \mathbb{R}$. The proof for G_2 is analogous.

Now let us analyze the running time of the algorithm. One can easily inductively prove that:

- Each graph at the i -th level of the recursion tree has less than $n/2^i + 2$ vertices; hence, there are no more than $\log_2 n + 1$ levels in the tree.
- The sum n_i of numbers of vertices through all the graphs at the i -th level is less than $2^i(n/2^i + 2) \leq n + 2^{i+1} \leq 3n$.
- Since, the sum m_i of numbers of edges in graphs at level $i > 0$ satisfies $m_i \leq m_{i-1} + n_{i-1}/2 \leq m_{i-1} + 3n/2$, we have $m_i \leq m + 3in/2 = O(m + n \log n)$.

By the above, and since the function $T_{\max\text{flow}}$ is convex, we conclude the total time cost at the i -th level is $O(T_{\max\text{flow}}(n, m + n \log n) \log(R/\varepsilon))$. Recall that there are $O(\log n)$ levels and therefore the total time is $O(T_{\max\text{flow}}(n, m + n \log n) \log(R/\varepsilon) \log n)$.

Exact Parametric Min s, t -Cut

In this section we show how Theorem 5.3.2 implies new bounds on computing *exact* parametric min s, t -cuts in a few interesting settings.

Integer Polynomial Costs. Suppose all the parametric costs $c_\lambda(uv)$ are of the form $c_\lambda(uv) = Q_{uv}(\lambda)$, where each Q_{uv} is a (possibly different) constant-degree polynomial with integer coefficients bounded in the absolute value by an integer $U > 0$ and take nonnegative values on \mathbb{D} . Recall Q_{uv} can have a positive degree only if $u = s$ or $v = t$. Moreover, if $u = s$, then Q_{uv} is increasing, whereas when $v = t$, then Q_{uv} is decreasing.

Observe that the parametric capacity $c_\lambda(S)$ of any S , $s \in S \subseteq V \setminus \{t\}$ is a constant-degree polynomial with integer coefficients bounded by nU in absolute value. The same applies to a difference polynomial $c_\lambda(S) - c_\lambda(S')$ for any two such sets S, S' .

It is known that for a constant-degree polynomials Q with integer coefficients bounded by W :

- The roots of Q are of absolute value $O(\text{poly}(W))$ (e.g., [168]).
- Any two distinct roots of Q are at least $\Omega(1/\text{poly}(W))$ apart. [119]

This means that by setting $\lambda_{\min} = -R/2$ and $\lambda_{\max} = R/2$ (or slightly less aggressively, if e.g., $R/2 \notin \mathbb{D}$) for a sufficiently large even integer $R = O(\text{poly } nU)$ such that $R/2$ exceeds the maximum possible absolute value of a root of any polynomial of the form $c_\lambda(S) - c_\lambda(S')$, we will indeed have $G_{\lambda_{\min}} = G_{\lambda'}$ for all $\lambda' < \lambda_{\min}$ and $G_{\lambda_{\max}} = G_{\lambda'}$ for all $\lambda' > \lambda_{\max}$.

Moreover, assume we compute an ε -approximate parametric min s, t -cut (Λ, τ) , where $\Lambda = \{\lambda_1, \dots, \lambda_k\}$. Suppose for some i there exists λ' , $\max(\lambda_i, \lambda_{i+1} - 2\varepsilon) \leq \lambda' < \lambda_{i+1}$, such that the minimal s, t -cut $S_{\lambda'}$ in $G_{\lambda'}$ satisfies $S_{\lambda_i} \subsetneq S_{\lambda'} \subsetneq S_{\lambda_{i+1}}$. Let $\lambda_i^* = \max(\lambda_i, \lambda_{i+1} - 2\varepsilon)$. Note that $S_{\lambda_i} = S_{\lambda_i^*}$ by Definition 9.3.12. Since $S_{\lambda'}$ is minimal, $c_\lambda(S_{\lambda_i})(\lambda_i^*) - c_\lambda(S_{\lambda'})(\lambda_i^*) \leq 0$ and $c_\lambda(S_{\lambda_i})(\lambda') - c_\lambda(S_{\lambda'})(\lambda') > 0$. So the polynomial $c_\lambda(S_{\lambda_i}) - c_\lambda(S_{\lambda'})$ is non-zero and has a root in the interval $[\lambda_i^*, \lambda']$. Similarly one can prove that the polynomial $c_\lambda(S_{\lambda'}) - c_\lambda(S_{\lambda_{i+1}})$ is non-zero and has a root in the interval $[\lambda', \lambda_{i+1}]$. We conclude that the product polynomial $[c_\lambda(S_{\lambda_i}) - c_\lambda(S_{\lambda'})] \cdot [c_\lambda(S_{\lambda'}) - c_\lambda(S_{\lambda_{i+1}})]$ is non-zero, has constant degree, has integer coefficients of order $O(\text{poly } nU)$, and has two distinct roots in the interval $[\lambda_i^*, \lambda_{i+1}]$, i.e., less than 2ε apart. Therefore, if we set ε so that $1/\varepsilon$ is a sufficiently large integer but still polynomial in nU , the assumption $S_{\lambda_i} \subsetneq S_{\lambda'} \subsetneq S_{\lambda_{i+1}}$ leads to a contradiction. As a result, for all such λ' , $S_{\lambda'}$ equals either $S_{\lambda'}$ or $S_{\lambda''}$.

In other words, computing an ε -approximate parametric min s, t -cut (Λ, τ) , where $\Lambda = \{\lambda_1, \dots, \lambda_k\}$, gives as the structure of all possible minimal min s, t -cuts S_λ in the following sense. Suppose $\Lambda^* = \{\lambda_1^*, \dots, \lambda_l^*\}$ is an *exact* parametric min s, t -cut. Then $k = l$ and $S(\lambda_i) = \{v \in V : \tau(v) \leq \lambda_i\} = S_{\lambda_i^*}$ for all $i = 1, \dots, k$. To compute λ_i^* , it is hence enough to find the *unique* $\lambda_i^* \in (\lambda_{i-1}, \lambda_i]$ such that $c_\lambda(S(\lambda_{i-1}))(\lambda_i^*) = c_\lambda(S(\lambda_i))(\lambda_i^*)$ which boils down to solving a polynomial equation of constant degree. It is well known that such equations can be solved exactly in constant time for degrees at most 4.

Observe that if we run our ε -approximate parametric min s, t -cut algorithm with $\lambda_{\min}, \lambda_{\max}, \varepsilon$ set as described above, maximum flow is always invoked on some minor H of G_λ for λ that is an integer multiple of ε . Since, $1/\varepsilon$ is an integer, by multiplying edge costs in H by $1/\varepsilon$, we only need a maximum flow algorithm that can handle integer edge capacities of order $O(\text{poly}(nU))$. By plugging in the best-known algorithms for computing max flow with integral capacities, we obtain the following.

Theorem 9.3.13. *Let G be a graph whose parameterized capacities are constant-degree polynomials with integer coefficients in $[-U, U]$. The structure of parametric min s, t -cut on G can be computed in:*

- $O(m \cdot \min(m^{1/2}, n^{2/3}) \cdot \text{polylog}\{n, U\})$ time using a combinatorial algorithm [75],

- $O((m + n^{1.5}) \cdot \text{polylog}\{n, U\})$ time using the algorithm of [25],
- $O(m^{1.497} \text{polylog}\{n, U\})$ time using the algorithm of [73].

The cut function can be found exactly in additional $O(n)$ time if the degrees of capacity polynomials are at most 4.

Discrete Domains. Let us now consider the case when \mathbb{D} is discrete and has ℓ elements. Suppose all parametric costs are arbitrary functions meeting the requirements of the parametric min s, t -cut problem. Then, we can make the ε -approximate algorithm exact by employing the following simple modifications. We start with $\lambda_{\min} = \min \mathbb{D}$ and $\lambda_{\max} = \max \mathbb{D}$. In the binary-search like step, we always choose the middle element of $\mathbb{D} \cap [\lambda_1, \lambda_2]$ as the next pivot. This way, all the max-flow computations are performed on minors of G_λ for $\lambda \in \mathbb{D}$.

Theorem 9.3.14. *Let G be a graph with arbitrary parameterized capacities $\mathbb{D} \rightarrow \mathbb{R}$ for a discrete domain $\mathbb{D} \subseteq \mathbb{R}$, where $\ell = |\mathbb{D}|$. Let $T_{\max\text{flow}}(n', m')$ be defined as in Theorem 5.3.2. Then exact parametric min s, t -cut on G can be computed in $O(T_{\max\text{flow}}(n, m \log n) \cdot \log \ell \cdot \log n)$ time.*

Planar Graphs. By Remark 9.3.7, all the max-flow computations in the algorithm of Theorem 5.3.2 are performed on minors of G . As a result, if the input graph G is planar, we can use state-of-the-art planar max s, t -flow algorithms to obtain better bounds on the parametric min- s, t -cut algorithms on planar graphs. Since maximum s, t -flow for planar graphs can be computed in $O(n \log n)$ time even for real capacities [22, 58], planar parametric min s, t -cut can be solved exactly:

- in $O(n \text{polylog}\{n, U\})$ time when parameterized capacities are constant degree polynomials with integer coefficients in $[-U, U]$,
- in $O(n \log^3 n \log \ell)$ time for discrete domains $\mathbb{D} \subseteq \mathbb{R}$ of size ℓ .

What may be surprising, our reduction is powerful enough to obtain an interesting subquadratic *strongly polynomial* exact algorithm computing parametric min s, t -cut in a planar graph with capacity functions that are arbitrary polynomials of degree no more than 4 and real coefficients.

We now sketch this algorithm. It is based on the *parametric search* technique [123] (see also [3]). Suppose we want to solve some decision problem $\mathcal{P}(\alpha)$, where $\alpha \in \mathbb{R}$, such that if $\mathcal{P}(\alpha_0)$ is a yes instance, then $\mathcal{P}(\alpha')$ for all $\alpha' < \alpha_0$ is also a yes instance. We wish to find the maximum α^* for which $\mathcal{P}(\alpha^*)$ is a yes instance. An example problem $\mathcal{P}(\alpha)$ could be “does an s, t -flow of value α exist in G ?”. Then, α^* clearly equals the maximum flow in G .

Suppose we have an efficient strongly polynomial algorithm solving the decision problem. Then, in practice one could find α^* via binary search given some initial interval containing α^* ; however, in general this would not lead to an exact algorithm for real values of α^* . Parametric search is a technique for converting a strongly polynomial *parallel* decision algorithm into a *sequential or parallel* strongly polynomial *optimization* algorithm as explained above. The only requirement to keep in mind is that the decision algorithm is governed by comparisons, each of which amounts to testing the sign of some low-degree (say, no more than 4) polynomial in α . Specifically, suppose we have a parallel decision algorithm \mathcal{A} that uses $W_{\mathcal{A}}$ work and $D_{\mathcal{A}}$ depth, and also a (possibly the same) another parallel decision algorithm \mathcal{B} with $W_{\mathcal{B}}$ work and $D_{\mathcal{B}}$ depth. Suppose for simplicity all these quantities are polynomial in n . Then, parametric search yields a strongly-polynomial optimization algorithm computing α^* in $\tilde{O}(D_{\mathcal{A}} \cdot W_{\mathcal{B}} + W_{\mathcal{A}})$ work and $\tilde{O}(D_{\mathcal{A}} \cdot D_{\mathcal{B}})$ depth.

Now back to planar graphs. We will use parametric search in a nested way. First of all, we will need a decent parallel max flow algorithm for planar graphs. It is well-known (e.g., [58]) that

the decision variant of the max s, t -flow problem on planar graphs is reducible to negative cycle detection in the dual graph (which is also planar). There exists a parallel negative cycle detection algorithm on planar graphs with $\tilde{O}(n + n^{3/2}/d^3)$ work and $\tilde{O}(d)$ depth for any $d \geq 1$ [93]. Hence, by using that algorithm as both \mathcal{A} (for $d = D^{3/7}$) and \mathcal{B} (for $d = D^{4/7}$), where D is a parameter, we have $W_{\mathcal{A}} = \tilde{O}(n + n^{3/2}/D^{9/7})$, $D_{\mathcal{A}} = \tilde{O}(D^{3/7})$, $W_{\mathcal{B}} = \tilde{O}(n + n^{3/2}/D^{12/7})$, $D_{\mathcal{B}} = \tilde{O}(D^{4/7})$. So parametric search yields a strongly-polynomial parallel max-flow algorithm for planar graphs with work $\tilde{O}(n + n^{3/2}/D^{9/7})$ and depth $\tilde{O}(D)$ for any $D \geq 1$.

Given a parallel algorithm for max flow in planar graphs, we can use parametric search (instead of binary search) once again when computing the pair λ_1, λ_2 in our recursive algorithm. More specifically, we would like λ_1 to be the largest such that $|S_{\lambda_1}| \leq n/2$, whereas λ_2 to be the smallest such that $|S_{\lambda_2}| \geq n/2$. It is easy to see that λ_1 and λ_2 are precisely neighboring (or the same) breakpoints of the cut function, i.e., belong to Λ from Definition 9.3.11. To actually compute λ_1, λ_2 , we use parametric search with \mathcal{A} set to the obtained parallel max-flow algorithm², and \mathcal{B} to the best known algorithm that computes a minimum min s, t -cut in a planar graph, i.e., a combination of the max-flow algorithm of [22, 58], and linear time graph search. So, in the outer parametric search instance we have $W_{\mathcal{A}} = \tilde{O}(n + n^{3/2}/D^{9/7})$, $D_{\mathcal{A}} = \tilde{O}(D)$, and $W_{\mathcal{B}} = \tilde{O}(n)$. Therefore, the obtained algorithm runs in $\tilde{O}(n^{3/2}/D^{9/7} + Dn)$ sequential time. By setting $D = n^{7/32}$, we obtain $\tilde{O}(n^{1+7/32}) = \tilde{O}(n^{1.21875})$ time.

We stress that all the algorithms [22, 58, 93] used above proceed by only adding and comparing edge weights. Adding polynomials cannot increase their degrees, so indeed when these algorithms are run “generically” for some λ , the control flow depends only on signs of some small degree polynomials.

Theorem 9.3.15. *Let G be a planar graph whose parameterized capacities are all polynomials of degree at most 4 with real coefficients. There exists a strongly polynomial algorithm computing parametric min s, t -cut in G exactly in $\tilde{O}(n^{1+7/32})$ time.*

9.3.4 Removing Assumptions on F_i

In this section we argue why the assumptions on the functions F_i introduced in Section 5.2 are valid without loss of generality. More precisely, we assumed that for all i , $F_i(\emptyset) = F_i(V_i) = 0$ and $F_i(S) \geq 0$ for all S . Here we show that a simple preprocessing step can enforce all of these conditions.

Without changing the original problem we can shift each F_i such that it evaluates to 0 on \emptyset , by defining $\bar{F}_i(S) = F_i(S) - F_i(\emptyset)$. This only changes F by a constant term without affecting the sets that minimize the parametric problem (5.1).

For each i , we use Lemma 9.3.16 to find a point $w_i \in B(\bar{F}_i)$. Using this point, we define $\overline{\bar{F}}_i(S) = \bar{F}_i(S) - w_i(S)$. Since by definition $w_i(V_i) = \bar{F}_i(V_i)$, we have that $\overline{\bar{F}}_i(V_i) = 0$. Also, we have $\overline{\bar{F}}_i(\emptyset) = \bar{F}_i(\emptyset) = 0$. Finally, since $w_i(S) \leq \bar{F}_i(S)$ for all S , we also have $\overline{\bar{F}}_i(S) \geq 0$.

²Actually, it is computing max-flow followed by a graph search to determine the minimal min s, t -cut. However, this latter step does not involve any comparisons on capacities, so its depth can be ignored.

Now we can equivalently rewrite the parametric problem

$$\begin{aligned}
F_\alpha(A) &= F(A) + \sum_{j \in A} \psi'_j(\alpha) \\
&= \bar{F}(A) + \sum_{j \in A} \psi'_j(\alpha) + \left(\sum_{i=1}^m F_i(\emptyset) \right) \\
&= \bar{\bar{F}}(A) + \sum_{j \in A} \left(\psi'_j(\alpha) + \sum_{i=1}^r w_i(j) \right) + \left(\sum_{i=1}^m F_i(\emptyset) \right).
\end{aligned}$$

Now we can solve the problem on $\bar{\bar{F}} = \sum_{i=1}^r \bar{\bar{F}}_i$ with the parametric penalties $\bar{\bar{\psi}}'_j(\alpha) = \psi'_j(\alpha) + \sum_{i=1}^r w_i(j)$, which maintain the validity of Assumption 5.2.2.

To compute a point in the base polytope of a submodular function we use the following folklore lemma, which shows that the running time of our initialization procedure is $O(\sum_{i=1}^r |V_i| \cdot \text{EO}_i)$:

Lemma 9.3.16 ([70]). *Let $F : 2^V \rightarrow \mathbb{Z}$ be a submodular set function, with $F(\emptyset) = 0$, and let $B(F)$ be its base polytope. Given any $x \in \mathbb{R}^{|V|}$, one can compute*

$$\arg \max_{w \in B(F)} \langle x, w \rangle$$

using $O(|V|)$ calls to an evaluation oracle for F . Furthermore w is integral.

9.3.5 Deferred Proofs

Proof of Lemma 5.4.3

We define the primal and dual optima of this problem, which will be useful for the proof.

Definition 9.3.17 (Graph subproblem minimizers). *Let \tilde{x}^* be the minimizer of*

$$\min_x g(x) + \phi(x), \tag{9.52}$$

and \tilde{w}^* be the minimizer of

$$\min_{w \in B(G)} \phi^*(-w). \tag{9.53}$$

The main tool that we will use for this proof will be the following two structural statements, which can be extracted from Propositions 4.2 and 8.3 in [15].

Lemma 9.3.18 ([15]). *Consider any submodular function $F : 2^V \rightarrow \mathbb{R}$.*

1. *Fix some $x \in \mathbb{R}^n$. For any $w \in \mathbb{R}^n$, w is an optimizer of $\max_{w \in B(F)} \langle w, x \rangle$ if and only if there exists a permutation π of $[n]$ such that $x_{\pi_1} \geq x_{\pi_2} \geq \dots \geq x_{\pi_n}$ and for all $u \in V$ we have*

$$w_u = \begin{cases} F(\{\pi_1\}) & \text{if } u = 1 \\ F(\{\pi_1, \pi_2, \dots, \pi_u\}) - F(\{\pi_1, \pi_2, \dots, \pi_{u-1}\}) & \text{if } u \geq 2 \end{cases}.$$

2. Given a function ϕ that satisfies the conditions in Definition 5.2.2, the optimal solution \tilde{x}^* to the problem

$$\min_x f(x) + \phi(x),$$

where f is the Lovász extension of F , is given by

$$\tilde{x}_u^* = -\inf(\{\lambda \in \mathbb{R} : u \in S(\lambda)\})$$

for all $u \in V$, where

$$S(\lambda) = \operatorname{argmin}_{S \subseteq V} F(S) + \sum_{u \in S} \phi'_u(-\lambda).$$

Additionally, we present two simple lemmas which will be useful in the proof. The first one upper bounds the ℓ_1 diameter of a base polytope, and the second one upper bounds the ℓ_1 norm of the gradient of a function in the base polytope.

Lemma 9.3.8. *For any submodular function $F : 2^V \rightarrow \mathbb{R}_{\geq 0}$ and $F(S) \leq F_{\max}$ for all $S \subseteq V$, we have that*

$$\max_{w \in B(F)} \|w\|_1 \leq 2nF_{\max}.$$

Proof. By definition of $B(F)$, for all $u \in V$ we have $w_u \leq F(\{u\}) \leq F_{\max}$, so $\sum_{u \in V: w_u \geq 0} w_u \leq nF_{\max}$.

Also, $\sum_{u \in V} w_u = F(V)$, so we conclude that

$$\begin{aligned} \|w\|_1 &= \sum_{u \in V: w_u \geq 0} w_u - \sum_{u \in V: w_u < 0} w_u \\ &= 2 \sum_{u \in V: w_u \geq 0} w_u - F(V) \\ &\leq 2nF_{\max} - F(V) \\ &\leq 2nF_{\max} \end{aligned}$$

□

Lemma 9.3.9. *For any submodular function $F : 2^V \rightarrow \mathbb{R}_{\geq 0}$, $F(S) \leq F_{\max}$ for all $S \subseteq V$, and function $\psi : \mathbb{R}^n \rightarrow \mathbb{R}$ satisfying the conditions of Definition 5.2.2 we have that*

$$\max_{w \in B(F)} \|\nabla \psi^*(-w)\|_1 \leq \frac{2nF_{\max}}{\sigma} + \|\nabla \psi^*(0)\|_1.$$

Proof.

$$\begin{aligned} \|\nabla \psi^*(-w)\|_1 &\leq \|\nabla \psi^*(-w) - \nabla \psi^*(0)\|_1 + \|\nabla \psi^*(0)\|_1 \\ &\leq \frac{1}{\sigma} \|w\|_1 + \|\nabla \psi^*(0)\|_1 \\ &\leq \frac{2nF_{\max}}{\sigma} + \|\nabla \psi^*(0)\|_1, \end{aligned}$$

where we used the triangle inequality, the $\frac{1}{\sigma}$ -smoothness of the ψ_u^* 's, and Lemma 9.3.8. \square

We are now ready to proceed with the proof.

Proof of Lemma 5.4.3. We let $\Lambda = \{\lambda_1, \dots, \lambda_k\}$, where $\lambda_1 < \dots < \lambda_k$, and define $S(\lambda)$ to be a minimal set in

$$\operatorname{argmin}_{S \subseteq V} G(S) + \sum_{u \in S} \phi'_u(-\lambda)$$

for all $\lambda \in \mathbb{R}$. Note that this can be equivalently written as

$$\begin{aligned} & \operatorname{argmin}_{S \subseteq V} G(S) + \sum_{u \in S} \max\{0, \phi'_u(-\lambda)\} - \sum_{u \in S} \max\{0, -\phi'_u(-\lambda)\} \\ &= \operatorname{argmin}_{S \subseteq V} G(S) + \sum_{u \in S} \max\{0, \phi'_u(-\lambda)\} + \sum_{u \in V \setminus S} \max\{0, -\phi'_u(-\lambda)\} - \sum_{u \in V} \max\{0, -\phi'_u(-\lambda)\} \\ &= \operatorname{argmin}_{S \subseteq V} G(S) + \sum_{u \in S} \max\{0, \phi'_u(-\lambda)\} + \sum_{u \in V \setminus S} \max\{0, -\phi'_u(-\lambda)\} \\ &= \operatorname{argmin}_{S \subseteq V} c_\lambda^+(S \cup \{s\}), \end{aligned}$$

by the definition of the parametric capacities c_λ , where $c_\lambda^+(S \cup \{s\}) = \sum_{\substack{u \in S \cup \{s\} \\ v \in V \setminus (S \cup \{s\})}} c_\lambda(u, v)$. Addition-

ally, we denote $\varepsilon = \frac{1}{3L}$ for convenience. By the second item of Lemma 9.3.18, we know that the minimizer of $\min_x f(x) + \phi(x)$ is defined as

$$\tilde{x}_u^* = -\inf\{\lambda \in \mathbb{R} : u \in S(\lambda)\}.$$

For all $u \in V$, let $i_u = \operatorname{argmin}_{i \in [k]} \{\lambda_i \mid u \in S(\lambda_i)\}$ and $\tilde{x}_u = -\lambda_{i_u}$. We will first prove that $\|\tilde{x} - \tilde{x}^*\|_\infty \leq \varepsilon$.

Now, by definition we have that $\tilde{x}_u^* \geq \tilde{x}_u$. Additionally, setting $\lambda_0 = -\infty$ for convenience, we have $u \notin S(\lambda_{i_u-1})$, and $u \in S(-\tilde{x}_u^*)$, so $S(\lambda_{i_u-1}) \subset S(-\tilde{x}_u^*)$. By the first item of Definition 9.3.12, this implies that

$$-\tilde{x}_u^* \geq \lambda_{i_u} - \varepsilon = -\tilde{x}_u - \varepsilon \Leftrightarrow \tilde{x}_u^* \leq \tilde{x}_u + \varepsilon.$$

Therefore, we have concluded that $|\tilde{x}_u - \tilde{x}_u^*| \leq \varepsilon$ for all $u \in V$, i.e. $\|\tilde{x} - \tilde{x}^*\|_\infty \leq \varepsilon$.

We define a dual solution $\hat{w} = -\nabla\phi(\tilde{x})$. We will show that \tilde{w}^* can be retrieved by rounding \hat{w} . Using the fact that ϕ_u 's are L -smooth and the optimality condition $\tilde{w}^* = -\nabla\phi(\tilde{x}^*)$ from Lemma 9.3.10, we get that

$$\|\hat{w} - \tilde{w}^*\|_\infty = \|\nabla\phi(\tilde{x}) - \nabla\phi(\tilde{x}^*)\|_\infty \leq L\|\tilde{x} - \tilde{x}^*\|_\infty \leq L\varepsilon = 1/3.$$

On the other hand, by optimality of \tilde{w}^* , it is a maximizer of $\max_{w \in B(G)} \langle w, \tilde{x}^* \rangle$. By the first item of Lemma 9.3.18, there exists a permutation π_1, \dots, π_n of V such that $\tilde{w}_u^* = G(\{\pi_1, \dots, \pi_u\}) - G(\{\pi_1, \dots, \pi_{u-1}\})$ for $u \in V$. As G takes integral values, we have $\tilde{w}_u^* \in \Delta \cdot \mathbb{Z}$ for all $u \in V$, and since $|\hat{w}_u - \tilde{w}_u^*| < 1/2$, we can exactly recover \tilde{w}^* by rounding each entry of \hat{w} to the closest integer.

Our next goal is to compute a G -decomposition of \tilde{w}^* , which we will do by computing an exact primal solution and then again applying the first item of Lemma 9.3.18. Given \tilde{w}^* , we can easily

recover the primal optimum $\tilde{x}^* = \nabla\phi(-\tilde{w}^*)$. In order to recover a decomposition $\tilde{w}^* = \sum_{i=1}^r \tilde{w}^{*i}$, we use the well-known fact [54] that

$$\max_{w \in B(G)} \langle w, x \rangle = \max_{w^i \in B(G_i)} \sum_{i=1}^r \langle w^i, x \rangle,$$

so for any $i \in [r]$, \tilde{w}^{*i} necessarily maximizes

$$\max_{w^i \in B(G_i)} \langle w^i, \tilde{x}^* \rangle.$$

Therefore, by the first item of Lemma 9.3.18, \tilde{w}^{*i} can be recovered by sorting the entries of \tilde{x}^* in decreasing order, such that $\tilde{x}_{\pi_1}^* \geq \tilde{x}_{\pi_2}^* \geq \dots \geq \tilde{x}_{\pi_n}^*$ for some permutation π of V , and then setting

$$\tilde{w}_u^{*i} = G_i(\{\pi_1, \dots, \pi_u\}) - G_i(\{\pi_1, \dots, \pi_{u-1}\}) \quad (9.54)$$

for all $u \in V$. Note that \tilde{w}^{*i} 's are in \mathbb{Z}^n .

The runtime is dominated by the computation of the decomposition in (9.54), which involves computing prefix cuts for each G_i and by Lemma 9.3.10 takes time $O\left(\sum_{i=1}^r |V_i|^2\right)$. Therefore, the total runtime is $O\left(n + \sum_{i=1}^r |V_i|^2\right)$. \square

Lemma 9.3.10 (Computing all prefix cut values). *Given a graph $G(V, E, c)$ with $V = \{1, 2, \dots, n\}$, we can compute the values $c^+([u])$ for all $u \in [n]$ in time $O(n^2)$.*

Proof. We note that $c^+(\emptyset) = 0$ and for any $u \geq 1$ we have

$$c^+([u]) = c^+([u-1]) + \sum_{v=u+1}^n c_{uv} - \sum_{v=1}^{u-1} c_{vu}. \quad (9.55)$$

Therefore $c^+([u])$ can be computed in $O(n)$ given $c^+([u-1])$. As we apply (9.55) n times, the total runtime is $O(n^2)$. \square

Proof of Lemma 5.4.4

We first prove the following lemma, which helps us bound the range of parameters for parametric min s, t -cut.

Lemma 9.3.11. *Consider a graph $G(V \cup \{s, t\}, E, c \geq 0)$ and a function $\phi(x) = \sum_{u \in V} \phi_u(x_u)$ that satisfies Assumption 5.2.2. Additionally, let $G(S) = c^+(S \cup \{s\})$ for all $S \subseteq V$ be the cut function associated with the graph. For any $\lambda \in \mathbb{R}$, we set $S(\lambda)$ to be the smallest set that minimizes*

$$\min_{S \subseteq V} G(S) + \sum_{u \in S} \phi'_u(-\lambda).$$

Let $\rho = \max_{u \in V} |\phi'_u(0)|$ and $G_{\max} = \max_{S \subseteq V} G(S)$. Then, $S(\lambda_{\min}) = \emptyset$ and $S(\lambda_{\max}) = V$, where $\lambda_{\min} = -2\frac{\rho + G_{\max}}{\sigma}$ and $\lambda_{\max} = 2\frac{\rho + G_{\max}}{\sigma}$.

Proof. We first note that by the σ -strong convexity of the ϕ_u 's, and since $-\lambda_{\min} > 0 > -\lambda_{\max}$, we have that

$$\phi'_u(-\lambda_{\min}) \geq \phi'_u(0) + \sigma|\lambda_{\min}|.$$

and

$$\phi'_u(-\lambda_{\max}) \leq \phi'_u(0) - \sigma|\lambda_{\max}|.$$

Therefore for any $\emptyset \neq S \subseteq V$ we have

$$\begin{aligned} G(S) + \sum_{u \in S} \phi'_u(-\lambda_{\min}) &\geq G(S) + \sum_{u \in S} (\phi'_u(0) + \sigma|\lambda_{\min}|) \\ &\geq G(S) - |S|\rho + |S|\sigma|\lambda_{\min}| \\ &= G(S) - |S|\rho + 2|S|\sigma \frac{\rho + G_{\max}}{\sigma} \\ &> G(S) + G_{\max} \\ &\geq G(\emptyset), \end{aligned}$$

where we used the fact that $G(S) \geq 0$ and $G(\emptyset) \leq G_{\max}$, so $S(\lambda_{\min}) = \emptyset$. Similarly, for any $S \subset V$ we have

$$\begin{aligned} G(S) + \sum_{u \in S} \phi'_u(-\lambda_{\max}) &= G(S) + \sum_{u \in V} \phi'_u(-\lambda_{\max}) - \sum_{u \in V \setminus S} \phi'_u(-\lambda_{\max}) \\ &\geq G(S) + \sum_{u \in V} \phi'_u(-\lambda_{\max}) - \sum_{u \in V \setminus S} (\phi'_u(0) - \sigma|\lambda_{\max}|) \\ &\geq G(S) + \sum_{u \in V} \phi'_u(-\lambda_{\max}) + |V \setminus S|(\sigma|\lambda_{\max}| - \rho) \\ &= G(S) + \sum_{u \in V} \phi'_u(-\lambda_{\max}) + |V \setminus S| \left(2\sigma \frac{\rho + G_{\max}}{\sigma} - \rho \right) \\ &> G(S) + \sum_{u \in V} \phi'_u(-\lambda_{\max}) + G_{\max} \\ &\geq G(V) + \sum_{u \in V} \phi'_u(-\lambda_{\max}), \end{aligned}$$

where we used the fact that $G(S) \geq 0$ and $G(V) \leq G_{\max}$, so $S(\lambda_{\max}) = V$. □

We are now ready for the proof.

Proof of Lemma 5.4.4. We first shift the polytope $B(F)$ so that w is translated to 0. Specifically, for all $S \subseteq V$, we let $\widehat{F}(S) = F(S) - w(S)$ and $\widehat{F}_i(S) = F_i(S) - w^i(S)$ for all $i \in [r]$. As we are just subtracting a linear function, \widehat{F} and the \widehat{F}_i 's are still submodular functions, and $B(\widehat{F}) = B(F) - w$, $B(\widehat{F}_i) = B(F_i) - w^i$ for all $i \in [r]$. Note that $w^i \in B(F_i)$ implies that the \widehat{F}_i 's (and thus also \widehat{F}) are non-negative, since

$$\widehat{F}_i(S) = F_i(S) - w^i(S) \geq 0,$$

and additionally $\widehat{F}_i(\emptyset) = F_i(\emptyset) = 0$ and $\widehat{F}_i(V_i) = F_i(V_i) - w^i(V_i) = 0$ for all $i \in [r]$ (also implying $\widehat{F}(\emptyset) = \widehat{F}(V) = 0$).

We run the algorithm from Lemma 9.3.5 on the \widehat{F}_i 's to obtain directed graphs $G_i(V, E, c^i \geq 0)$ whose $(V_i$ -restricted) cut functions $G_i(S) = c^{i+}(S)$ α -approximate $\widehat{F}_i(S)$, where $\alpha = \max_{i \in [r]} \{|V_i|^2/4 +$

$|V_i|$ }. More specifically,

$$\frac{1}{\alpha} \widehat{F}_i(S) \leq G_i(S) \leq \widehat{F}_i(S) \text{ for all } S \subseteq V_i, \quad G_i(V_i) = \widehat{F}_i(V_i) \quad (9.56)$$

and

$$\frac{1}{\alpha} (B(F_i) - w^i) = \frac{1}{\alpha} B(\widehat{F}_i) \subseteq B(G_i) \subseteq B(\widehat{F}_i) = B(F_i) - w^i, \quad (9.57)$$

We also define the graph $G(V, E, c \geq 0)$, where $c = \sum_{i=1}^r c^i$ and has cut function $G(S) = \sum_{i=1}^r G_i(S)$ for all $S \subseteq V$. Then,

$$\frac{1}{\alpha} \widehat{F}(S) \leq G(S) \leq \widehat{F}(S) \text{ for all } S \subseteq V, \quad G(V) = \widehat{F}(V) \quad (9.58)$$

and

$$\frac{1}{\alpha} (B(F) - w) \subseteq B(G) \subseteq B(F) - w. \quad (9.59)$$

We absorb the linear term that we subtracted from F into the parametric function. Concretely, we define ϕ as $\phi(x) = \psi(x) + \langle w, x \rangle$ for all $x \in \mathbb{R}^n$. It is easy to see that $\phi(x)$ is coordinate-wise separable, as $\phi(x) = \sum_{u \in V} \phi_u(x_u)$ where $\phi_u(x_u) = \psi_u(x_u) + w_u x_u$ for all $u \in V$ and that it satisfies Assumption 5.2.2, since $\phi_u''(x_u) = \psi_u''(x_u)$ and

$$|\phi_u'(0)| = |\psi_u'(0) + w_u| \leq |\psi_u'(0)| + F(\{u\}) \leq |\psi_u'(0)| + F_{\max} = n^{O(1)}.$$

Additionally, the Fenchel dual of ϕ_u is a shifted version of ψ_u , i.e.

$$\phi_u^*(z) = \max_{y \in \mathbb{R}} zy - \phi(y) = \max_{y \in \mathbb{R}} zy - \psi(y) - w_u y = \max_{y \in \mathbb{R}} (z - w_u)y - \psi(y) = \psi_u^*(z - w_u).$$

We will now run the algorithm from Lemma 5.4.3 on graphs G_i and parametric function ϕ to obtain a dual solution vector \tilde{w} . We note that, as F_i 's and w^i 's take integer values, G_i 's take integer values too. This algorithm takes a $\frac{1}{3L}$ -approximate parametric min s, t -cut as input, which we first compute using Theorem 5.3.2, with range of parameters $[\lambda_{\min}, \lambda_{\max}]$ given by Lemma 9.3.11. We have

$$\begin{aligned} \lambda_{\max} - \lambda_{\min} &= 4 \frac{\max_{u \in V} |\phi_u'(0)| + \max_{S \subseteq V} G(S)}{\sigma} \\ &\leq 4 \frac{\max_{u \in V} |\phi_u'(0)| + \widehat{F}_{\max}}{\sigma} \\ &\leq 4 \frac{\max_{u \in V} |\phi_u'(0)| + F_{\max} + \|w\|_1}{\sigma} \\ &\leq 4 \frac{\max_{u \in V} |\phi_u'(0)| + (2n + 1)F_{\max}}{\sigma} \\ &= n^{O(1)}, \end{aligned}$$

where we used Lemma 9.3.8 and the fact that the quantities $|\phi_u'(0)|, F_{\max}, \frac{1}{\sigma}$ are bounded by $n^{O(1)}$ (Assumption 5.2.2).

Based on Theorem 5.3.2, the time to obtain the $\frac{1}{3L}$ -approximate parametric min s, t -cut will be

$$O\left(T_{\max\text{flow}}(n, |E'| \log n) \log \frac{\lambda_{\max} - \lambda_{\min}}{1/3L} \log n\right) = \tilde{O}\left(T_{\max\text{flow}}\left(n, n + \sum_{i=1}^r |V_i|^2\right)\right).$$

So, by applying Lemma 5.4.3, we obtained a dual solution $\tilde{w} = \sum_{i=1}^r \tilde{w}^i$ for which $\tilde{w}^i \in B(G_i)$ and

$$\tilde{w} = \underset{\tilde{w} \in B(G)}{\operatorname{argmin}} \phi^*(-\tilde{w}) = \underset{\tilde{w} \in B(G)}{\operatorname{argmin}} \psi^*(-w - \tilde{w}). \quad (9.60)$$

For all $u \in V$ and $i \in [r]$, we set $w_u^i = w_u^i + \tilde{w}_u^i$ and $w'_u = \sum_{i=1}^r w_u^i$. Note that these quantities are still integral. We now prove the two parts of the lemma statement (feasibility and optimality) separately.

Feasibility. For any $i \in [r]$ and $S \subseteq V_i$, we have that

$$w'^i(S) = w^i(S) + \tilde{w}^i(S) \leq w^i(S) + G_i(S) \leq w^i(S) + \widehat{F}_i(S) = w^i(S) + F_i(S) - w^i(S) = F_i(S),$$

where the second inequality follows from the fact that $\tilde{w}^{*i} \in G_i$. Similarly, we have

$$w'^i(V_i) = w^i(V_i) + \tilde{w}^i(V_i) = w^i(V_i) + G_i(V_i) = w^i(V_i) + \widehat{F}_i(V_i) = F_i(V_i).$$

So we conclude that $w'^i \in B(F_i)$ for all $i \in [r]$.

Optimality. Let's set $h(z) := \psi^*(-z)$ for all $z \in \mathbb{R}^n$ for notational convenience, so that our goal is to prove that

$$h(w') - h(w^*) \leq \left(1 - \frac{1}{\alpha}\right) (h(w) - h(w^*)).$$

We will prove a slightly different statement where w' is replaced by a solution on the path from w to w^* , which is enough because w' is optimal in $w + B(G)$. Concretely, we set $\bar{w} = \frac{1}{\alpha}(w^* - w)$ and instead will prove

$$h(w + \bar{w}) - h(w^*) \leq \left(1 - \frac{1}{\alpha}\right) (h(w) - h(w^*)).$$

Now, h is a convex function, so applying convexity twice we have

$$\begin{aligned} h(w) &\geq h(w + \bar{w}) + \langle \nabla h(w + \bar{w}), -\bar{w} \rangle \\ &= h(w + \bar{w}) - \frac{1}{\alpha} \langle \nabla h(w + \bar{w}), w^* - w \rangle \end{aligned} \quad (9.61)$$

and

$$\begin{aligned} h(w^*) &\geq h(w + \bar{w}) + \langle \nabla h(w + \bar{w}), w^* - w - \bar{w} \rangle \\ &= h(w + \bar{w}) + \frac{\alpha - 1}{\alpha} \langle \nabla h(w + \bar{w}), w^* - w \rangle. \end{aligned} \quad (9.62)$$

We divide (9.62) by $\alpha - 1$ and then sum it with (9.61), getting

$$h(w) + \frac{1}{\alpha - 1} h(w^*) \geq h(w + \bar{w}) + \frac{1}{\alpha - 1} h(w + \bar{w}).$$

Equivalently,

$$\frac{\alpha}{\alpha - 1} (h(w + \bar{w}) - h(w^*)) \leq h(w) - h(w^*).$$

So by rearranging,

$$h(w + \bar{w}) - h(w^*) \leq \left(1 - \frac{1}{\alpha}\right) (h(w) - h(w^*)). \quad (9.63)$$

Thus we can equivalently write that

$$\psi^*(-w - \bar{w}) - \psi^*(-w^*) \leq \left(1 - \frac{1}{\alpha}\right) (\psi^*(-w) - \psi^*(-w^*)). \quad (9.64)$$

Now, since by (9.59) we have $\frac{1}{\alpha}(B(F) - w) \subseteq B(G)$ and $w^* \in B(F)$, we have $\bar{w} = \frac{1}{\alpha}(w^* - w) \in B(G)$. Combining the fact that \tilde{w} is a minimizer of $\min_{\tilde{w}^* \in B(G)} \psi^*(-w - \tilde{w}^*)$ with (9.60) and the fact that $\bar{w} \in B(G)$, we have

$$\psi^*(-w') = \psi^*(-w - \tilde{w}) \leq \psi^*(-w - \bar{w}).$$

Combining this with (9.64), we obtain the desired claim:

$$\psi^*(-w') - \psi^*(-w^*) \leq \left(1 - \frac{1}{\alpha}\right) (\psi^*(-w) - \psi^*(-w^*)). \quad (9.65)$$

The running time to compute graphs G_i is $O\left(\sum_{i=1}^r |V_i|^2 \mathcal{O}_i\right)$ and the time to run the algorithm from Lemma 5.4.3 is $\tilde{O}\left(n + \sum_{i=1}^r |V_i|^2\right)$, so the total running time is

$$\tilde{O}\left(\sum_{i=1}^r |V_i|^2 \mathcal{O}_i + T_{\max\text{flow}}\left(n, n + \sum_{i=1}^r |V_i|^2\right)\right).$$

□

Algorithm 24 Finding all minimum cuts

- 1: **function** FINDMINCUTS($G(V, E, c), \phi, \varepsilon$)
- 2: $V' = V \cup \{s, t\}$, $E' = E \cup \bigcup_{u \in V} \{(u, t)\} \cup \bigcup_{u \in V} \{(s, u)\}$
- 3: Define parametric capacities

$$c_\lambda(u, v) = \begin{cases} \max\{0, \phi'_u(-\lambda)\} & \text{if } u \in V, v = t \\ \max\{0, -\phi'_u(-\lambda)\} & \text{if } u = s, v \in V \\ c_{uv} & \text{otherwise} \end{cases}$$

- 4: Set $(\Lambda, \tau) = \text{APXPARAMETRICMINCUT}(G'(V', E'), c_\lambda, \lambda_{\min} = -n^{O(1)}, \lambda_{\max} = n^{O(1)})$
 - 5: Set $\tilde{w}_u = -\phi'_u(-\tau(u))$ for all $u \in V$ **return** \tilde{w}
-

9.4 Appendix for Chapter 6

9.4.1 Deferred Proofs

Proof of Lemma 6.3.8

Proof. Φ^t is a quadratic restricted on R^t

$$\begin{aligned}
 & \Phi(y) - \Phi(x) - \nabla\Phi(x)^T(y - x) \\
 &= \frac{\rho_2^+}{2} \left(\|y_{R^t}\|_2^2 - \|x_{R^t}\|_2^2 - 2x_{R^t}^T(y_{R^t} - x_{R^t}) \right) \\
 &= \frac{\rho_2^+}{2} \|y_{R^t} - x_{R^t}\|_2^2 \\
 &\in \left[0, \frac{\rho_2^+}{2} \|y - x\|_2^2 \right]
 \end{aligned}$$

and so for any x, y with $|\text{supp}(y - x)| \leq s + s^*$ (resp. $|\text{supp}(y - x)| \leq 1$) we have

$$\begin{aligned}
 & g(y) - g(x) - \nabla g(x)^T(y - x) \\
 &= f(y) - f(x) - \nabla f(x)^T(y - x) + \Phi(y) - \Phi(x) - \nabla\Phi(x)^T(y - x) \\
 &\geq \frac{\rho_{s+s^*}^-}{2} \|y - x\|_2^2 \text{ (resp. } \leq \rho_2^+ \|y - x\|_2^2 \text{)}.
 \end{aligned}$$

□

Proof of Lemma 6.3.3

Proof. By definition, and setting $\tau = \frac{1}{s}$, in each Type 1 iteration we have

$$\begin{aligned}
 & g(x^t) - g(x^{t+1}) \geq \tau (g(x^t) - \text{opt}) \\
 & \Rightarrow g(x^{t+1}) - \text{opt} \leq (1 - \tau)(g(x^t) - \text{opt})
 \end{aligned}$$

and in each Type 2 iteration we have

$$g(x^{t+1}) - \text{opt} \leq g(x^t) - \text{opt}$$

(since g can only decrease when unregularizing), therefore

$$\begin{aligned}
 f(x^T) - \text{opt} &\leq g(x^T) - \text{opt} \\
 &\leq (1 - \tau)^{T_1} (g(x^0) - \text{opt}) \\
 &\leq e^{-\tau T_1} (g(x^0) - \text{opt}) \\
 &\leq \varepsilon,
 \end{aligned}$$

where we used the fact that $T_1 = \frac{1}{\tau} \log \frac{g(x^0) - \text{opt}}{\varepsilon}$.

□

```

import numpy as np

def IHT(n, s):
    x = np.zeros(n)
    for _ in range(T):
        x_new = x - eta * grad(x)
        x_new[np.argsort(np.abs(x_new))[:s]] = 0
        x = x_new
    return x

def RegIHT(n, s):
    x, w = np.zeros(n), np.ones(n)
    for _ in range(T):
        x_new = (1 - 0.5 * w) * x - 0.5 * eta * grad(x)
        x_new[np.argsort(np.abs(x_new))[:s]] = 0
        reg = np.sum(w * x**2)
        if reg != 0:
            w = w * (1 - c * w * x**2 / reg)
            w[w <= round_th] = 0
        x = x_new
    return x

```

Figure 9-1: Our python implementations of IHT, RegIHT, where `grad` is the gradient function, `n` is the total number of features, `s` is the desired sparsity level, `eta` is the step size, `c` is the weight step size, and `round_th` is the weight rounding threshold, which we set to 0.5. Note that `grad(x) = np.dot(A.T, np.dot(A, x) - b)` for linear regression and `grad(x) = np.dot(A.T, expit(np.dot(A,x)) - b)` for logistic regression, where `expit` is the sigmoid function.

Proof of Lemma 6.4.1

Proof. We have

$$\begin{aligned}
 \left(\sqrt{f(x^t) - f(\tilde{x}^t)} + \sqrt{f(x^*) - f(\tilde{x}^t)} \right)^2 &\geq \frac{\rho^-}{2} (\|x^t - \tilde{x}^t\|_2 + \|x^* - \tilde{x}^t\|_2)^2 \\
 &\geq \frac{\rho^-}{2} \|x^t - x^*\|_2^2 \\
 &\geq \frac{\rho^-}{2} \left(\|x_{S^* \setminus S^t}^*\|_2^2 + \|x_{S^t \setminus S^*}^t\|_2^2 \right),
 \end{aligned}$$

where the first inequality follows by applying strong convexity to lower bound $f(x^t) - f(\tilde{x}^t)$ and $f(x^*) - f(\tilde{x}^t)$ combined with the fact that by definition of \tilde{x}^t , $\nabla_{S^t \cup S^*} f(\tilde{x}^t) = \vec{0}$, and the second is a triangle inequality. \square

9.5 Appendix for Chapter 7

9.5.1 Python implementation

In Figure 9-1 we have python implementations of the IHT and regularized IHT algorithms that we use for our experiments. As can be seen, both implementations are pretty short.

9.5.2 Proof of Theorem 7.1.1

Theorem 7.1.1 (Regularized IHT). *Let $f \in \mathbb{R}^n \rightarrow \mathbb{R}$ be a convex function that is β -smooth and α -strongly convex, with condition number $\kappa = \beta/\alpha$, and \mathbf{x}^* be an (unknown) s -sparse solution with support S^* . Then, running Algorithm 16 with $\eta = (2\beta)^{-1}$ and $c = s'/(4T)$ for*

$$T = O\left(\kappa \log \frac{f(\mathbf{x}^0) + (\beta/2) \|\mathbf{x}^0\|_2^2 - f(\mathbf{x}^*)}{\varepsilon}\right)$$

iterations starting from an arbitrary $s' = O(s\kappa)$ -sparse solution \mathbf{x}^0 , the algorithm returns an s' -sparse solution \mathbf{x}^T such that $f(\mathbf{x}^T) \leq f(\mathbf{x}^) + \varepsilon$. Furthermore, each iteration requires $O(1)$ evaluations of f , ∇f , and $O(n)$ additional time.*

Proof. We repeatedly apply Lemma 7.3.1 for

$$T = 64(\kappa + 1) \log \frac{f(\mathbf{x}^0) + (\beta/2) \|\mathbf{x}^0\|_2^2 - f(\mathbf{x}^*)}{\varepsilon}$$

iterations. We define the regularized function

$$g^t(\mathbf{x}) := f(\mathbf{x}) + (\beta/2) \|\mathbf{x}\|_{\mathbf{w}^t, 2}^2,$$

where \mathbf{w}^t are the weights before iteration $t \in [0, T-1]$. Specifically, for each t we apply Lemma 7.3.1 on the current solution \mathbf{x}^t and obtain the solution \mathbf{x}^{t+1} .

Before moving forward, we give an intuitive summary of the proof and the role of Lemma 7.3.1. As long as IHT makes “sufficient” progress on the regularized function g^t , this is satisfactory for the original function f as well, because $f(\mathbf{x}) \leq g^t(\mathbf{x})$ for all \mathbf{x} . This is the case of the first bullet of Lemma 7.3.1. If it stops making sufficient progress, this means we are at an (approximate) sparse optimum for g^t , although it is not necessarily a good sparse solution for f , which is the objective we are aiming to minimize. This is where the second bullet of Lemma 7.3.1 comes in, which gives necessary conditions for the above non-progress phenomenon (in other words, a partial characterization of the local minima encountered when running IHT on a regularized function). Specifically, the following condition is central to our approach:

$$\|\mathbf{x}_{S^*}^t\|_{(\mathbf{w}^t)_2}^2 \geq (4\kappa + 6)^{-1} \|\mathbf{x}^t\|_{\mathbf{w}^t, 2}^2.$$

We use this condition in the second part of the proof (after Case 2) to motivate a weight update from \mathbf{w}^t to \mathbf{w}^{t+1} , and show that, exactly because of this condition, a lot of the weight decrease is concentrated inside the optimal support S^* . As the total weight decrease in S^* is bounded by s , this gives a bound on the total number of iterations with insufficient decrease of g^t . If not for this condition, we would not be able to bound the number of such iterations and would have potentially remained forever stuck at a local minimum.

Now we are ready to move to the technical proof. In order to make sure that $g^{t+1}(\mathbf{x}^{t+1}) \leq g^{t+1}(\mathbf{x}^t)$, we revert to the previous solution if the one returned by Lemma 7.3.1 has a larger value of g^{t+1} . This is exactly what the conditional in Algorithm 16 is for. The property that $g^{t+1}(\mathbf{x}^{t+1}) \leq g^{t+1}(\mathbf{x}^t)$ is only used in the very last part of the proof.

Let us assume that $g^t(\mathbf{x}^t) > f(\mathbf{x}^*)$ at all times, as otherwise the statement holds by the fact that $g^t(\mathbf{x}^t)$ is non-increasing as a function of t and upper bounds $f(\mathbf{x}^t)$ for all t . We have $g^{t+1}(\mathbf{x}^{t+1}) \leq g^t(\mathbf{x}^t)$ by the fact that $\mathbf{w}^{t+1} \leq \mathbf{w}^t$ and $g^t(\mathbf{x}^{t+1}) \leq g^t(\mathbf{x}^t)$ by the guarantees of Lemma 7.3.1.

If the first bullet of Lemma 7.3.1 holds, we have that the value of g decreases considerably on iteration t , i.e.

$$\begin{aligned} g^{t+1}(\mathbf{x}^{t+1}) &\leq g^t(\mathbf{x}^{t+1}) \\ &\leq g^t(\mathbf{x}^t) - (16\kappa)^{-1}(g^t(\mathbf{x}^t) - f(\mathbf{x}^*)). \end{aligned}$$

Let us call these iterations *progress iterations*, and the other ones (where the second bullet of Lemma 7.3.1 holds) *weight iterations*. Now, since $g^t(\mathbf{x}^t)$ is non-increasing as a function of t , after $16\kappa \log \frac{g^0(\mathbf{x}^0) - f(\mathbf{x}^*)}{\varepsilon}$ progress iterations we will have

$$f(\mathbf{x}^t) \leq g^t(\mathbf{x}^t) \leq f(\mathbf{x}^*) + \varepsilon,$$

and so we will be done. From now on let us assume this is not the case, so there are at least

$$T - 16\kappa \log \frac{g^0(\mathbf{x}^0) - f(\mathbf{x}^*)}{\varepsilon} \geq 3T/4$$

weight iterations.

We remind that in each weight iteration, we have

$$\|\mathbf{x}_{S^*}^t\|_{(\mathbf{w}^t)^2, 2}^2 \geq (4\kappa + 6)^{-1} \|\mathbf{x}^t\|_{\mathbf{w}^t, 2}^2 \quad (9.66)$$

$$(\beta/2) \|\mathbf{x}_{S^*}^t\|_{(\mathbf{w}^t)^2, 2}^2 \geq (8\kappa + 8)^{-1} (g^t(\mathbf{x}^t) - f(\mathbf{x}^*)). \quad (9.67)$$

In words, (9.66) roughly implies that at least an $\Omega(1/\kappa)$ fraction of the mass of $(\mathbf{w}^t \mathbf{x}^t)^2$ lies inside S^* . Therefore, if we decrease \mathbf{w}^t by a quantity proportional to $(\mathbf{w}^t \mathbf{x}^t)^2$, the total sum of weights will decrease at most $O(\kappa)$ times faster than the sum of weights inside S^* . As the latter quantity can only decrease by s overall, the total decrease of weights will be $O(s\kappa)$.

Concretely, after each iteration we update the regularization weights as follows:

$$\mathbf{w}^{t+1} = \left(\mathbf{w}^t - c \cdot (\mathbf{w}^t \mathbf{x}^t)^2 / \|\mathbf{x}^t\|_{\mathbf{w}^t, 2}^2 \right)_{\geq 1/2},$$

for some $c > 0$ to be determined later. First of all, note that the weights are non-increasing. Now, if not for the thresholding operation, it is easy to see that the total weight decrease is at most c . The thresholding operation can only double this weight decrease to $2c$. Concretely, for all t we define a vector $\bar{\mathbf{w}}^t$ such that

$$\bar{w}_i^t = \begin{cases} w_i^t & \text{if } w_i^t \geq 1/2 \\ 1/2 & \text{if } w_i^t = 0. \end{cases}$$

Clearly, $\frac{1}{2}(\mathbf{1} - \mathbf{w}^t) \leq \mathbf{1} - \bar{\mathbf{w}}^t \leq \mathbf{1} - \mathbf{w}^t$. Now, we have

$$\begin{aligned} &\|\bar{\mathbf{w}}^t\|_1 - \|\bar{\mathbf{w}}^{t+1}\|_1 \\ &\leq c \|\mathbf{x}^t\|_{(\mathbf{w}^t)^2, 2}^2 / \|\mathbf{x}^t\|_{\mathbf{w}^t, 2}^2 \\ &\leq c, \end{aligned}$$

and, summing up for all t we get

$$\|\mathbf{1} - \bar{\mathbf{w}}^T\|_1 = \|\bar{\mathbf{w}}^0\|_1 - \|\bar{\mathbf{w}}^T\|_1 \leq cT.$$

Therefore,

$$\|\mathbf{1} - \mathbf{w}^T\|_1 \leq 2\|\mathbf{1} - \bar{\mathbf{w}}^T\|_1 \leq 2cT,$$

and so $\|\mathbf{w}^T\|_1 \geq n - 2cT$.

Therefore, the condition $\|\mathbf{w}^t\|_1 \geq n - s'/2$ of Lemma 7.3.1 is satisfied for all t as long as $c \leq s'/(4T)$. In order to bound the number of iterations, we distinguish two cases for the sum of weights inside S^* .

Case 1: The sum of weights inside S^* decreases by $\geq 4s/T$.

This case cannot happen more than $T/4$ times since the sum of weights inside S^* can only decrease by s in total. Therefore, case 2 below happens at least $T/2$ times.

Case 2: The sum of weights inside S^* decreases by $< 4s/T$.

Note that the decrease in the sum of weights in S^* is exactly equal to

$$\sum_{i \in S^*} \begin{cases} c \cdot (w_i^t x_i^t)^2 / \|\mathbf{x}^t\|_{\mathbf{w}^t, 2}^2 & \text{if this is } \leq w_i^t - 1/2 \\ w_i^t & \text{otherwise.} \end{cases}$$

Let T^* be the set of indices $i \in S^*$ for which the second case is true, i.e.

$$c \cdot (w_i^t x_i^t)^2 / \|\mathbf{x}^t\|_{\mathbf{w}^t, 2}^2 > w_i^t - 1/2.$$

The total weight decrease from elements in $S^* \setminus T^*$ is then

$$\begin{aligned} & \sum_{i \in S^* \setminus T^*} c \cdot (w_i^t x_i^t)^2 / \|\mathbf{x}^t\|_{\mathbf{w}^t, 2}^2 \\ &= c \left\| \mathbf{x}_{S^* \setminus T^*}^t \right\|_{(\mathbf{w}^t)^2, 2}^2 / \|\mathbf{x}^t\|_{\mathbf{w}^t, 2}^2 \\ &\geq \frac{c}{4\kappa + 6} \left\| \mathbf{x}_{S^* \setminus T^*}^t \right\|_{(\mathbf{w}^t)^2, 2}^2 / \left\| \mathbf{x}_{S^*}^t \right\|_{(\mathbf{w}^t)^2, 2}^2, \end{aligned}$$

where we used (9.66). As we have assumed that this decrease is less than $4s/T$, we have that

$$\begin{aligned} \left\| \mathbf{x}_{T^*}^t \right\|_{(\mathbf{w}^t)^2, 2}^2 &= \left\| \mathbf{x}_{S^*}^t \right\|_{(\mathbf{w}^t)^2, 2}^2 - \left\| \mathbf{x}_{S^* \setminus T^*}^t \right\|_{(\mathbf{w}^t)^2, 2}^2 \\ &\geq \left(1 - \frac{4s(4\kappa + 6)}{cT} \right) \left\| \mathbf{x}_{S^*}^t \right\|_{(\mathbf{w}^t)^2, 2}^2 \\ &\geq (1/2) \left\| \mathbf{x}_{S^*}^t \right\|_{(\mathbf{w}^t)^2, 2}^2, \end{aligned} \tag{9.68}$$

as long as $c \geq 8s(4\kappa + 6)/T$. We can pick such a c as long as

$$8s(4\kappa + 6)/T \leq c \leq s'/(4T) \Leftrightarrow s' \geq 32(4\kappa + 6)s.$$

Now, to deal with the fact that the sum weights in T^* might not decrease sufficiently, note that all

the weights in T^* are being set to 0, i.e. $w_i^{t+1} = 0$ for all $i \in T^*$. Together with (9.68) and (9.67) this means that we can make significant progress in function value. To see this, note that

$$\begin{aligned}
& g^{t+1}(\mathbf{x}^{t+1}) \\
& \leq g^{t+1}(\mathbf{x}^t) \\
& \leq g^t(\mathbf{x}^t) - (\beta/2) \|\mathbf{x}_{T^*}^t\|_{\mathbf{w}^t, 2}^2 \\
& \leq g^t(\mathbf{x}^t) - (\beta/2) \|\mathbf{x}_{T^*}^t\|_{(\mathbf{w}^t)^2, 2}^2 \\
& \leq g^t(\mathbf{x}^t) - (\beta/4) \|\mathbf{x}_{S^*}^t\|_{(\mathbf{w}^t)^2, 2}^2 \\
& \leq g^t(\mathbf{x}^t) - (16\kappa + 16)^{-1} (g^t(\mathbf{x}^t) - f(\mathbf{x}^*)) ,
\end{aligned}$$

which can happen at most

$$16(\kappa + 1) \log \frac{g^0(\mathbf{x}^0) - f(\mathbf{x}^*)}{\varepsilon} \leq T/4$$

times. □

9.5.3 Proof of Lemma 7.3.1

Lemma 7.3.1 (Regularized IHT step progress). *Let $f \in \mathbb{R}^n \rightarrow \mathbb{R}$ be a convex function that is β -smooth and α -strongly convex, $\kappa = \beta/\alpha$ be its condition number, and \mathbf{x}^* be any s -sparse solution.*

Given any s' -sparse solution $\mathbf{x} \in \mathbb{R}^n$ where

$$s' \geq (128\kappa + 2)s$$

and a weight vector $\mathbf{w} \in (\{0\} \cup [1/2, 1])^n$ such that $\|\mathbf{w}\|_1 \geq n - s'/2$, we make the following update:

$$\mathbf{x}' = H_{s'} \left((\mathbf{1} - 0.5\mathbf{w})\mathbf{x} - (2\beta)^{-1} \nabla f(\mathbf{x}) \right) .$$

Then, at least one of the following two conditions holds:

- *Updating \mathbf{x} makes regularized progress:*

$$g(\mathbf{x}') \leq g(\mathbf{x}) - (16\kappa)^{-1} (g(\mathbf{x}) - f(\mathbf{x}^*)) ,$$

where

$$g(\mathbf{x}) := f(\mathbf{x}) + (\beta/2) \|\mathbf{x}\|_{\mathbf{w}, 2}^2$$

is the ℓ_2 -regularized version of f with weights given by \mathbf{w} . Note: The regularized progress statement is true as long as \mathbf{x} is suboptimal, i.e. $g(\mathbf{x}) > f(\mathbf{x}^)$. Otherwise, we just have $g(\mathbf{x}') \leq g(\mathbf{x})$.*

- *\mathbf{x} is significantly correlated to the optimal support $S^* := \text{supp}(\mathbf{x}^*)$:*

$$\|\mathbf{x}_{S^*}\|_{\mathbf{w}^2, 2}^2 \geq (4\kappa + 6)^{-1} \|\mathbf{x}\|_{\mathbf{w}, 2}^2 ,$$

and the regularization term restricted to S^ is non-negligible:*

$$(\beta/2) \|\mathbf{x}_{S^*}\|_{\mathbf{w}^2, 2}^2 \geq (8\kappa + 8)^{-1} (g(\mathbf{x}) - f(\mathbf{x}^*)) .$$

Proof. By using the fact that f is β -smooth, and so g is 2β -smooth due to $\mathbf{w} \leq \mathbf{1}$, for any \mathbf{x}' we obtain

$$g(\mathbf{x}') - g(\mathbf{x}) \leq \langle \nabla g(\mathbf{x}), \mathbf{x}' - \mathbf{x} \rangle + \beta \|\mathbf{x}' - \mathbf{x}\|_2^2. \quad (9.69)$$

We let S be the support of \mathbf{x} and S' the support of \mathbf{x}' , i.e. the set of s' indices of the largest magnitude entries of the vector. Since $\nabla g(\mathbf{x}) = \nabla f(\mathbf{x}) + \beta \mathbf{w} \mathbf{x}$, we have

$$\bar{\mathbf{x}} = \mathbf{x} - \eta \nabla g(\mathbf{x}) = (\mathbf{1} - 0.5\mathbf{w})\mathbf{x} - \eta \nabla f(\mathbf{x}),$$

where $\eta = (2\beta)^{-1}$. We let $A = S' \setminus S$ be the newly inserted entries and $B = S \setminus S'$ be the entries that were just removed from the support. Note that

$$\begin{aligned} \mathbf{x}' &= [\mathbf{x} - \eta \nabla g(\mathbf{x})]_{S'} \\ &= \mathbf{x} - \eta \nabla_{S'} g(\mathbf{x}) - \mathbf{x}_B \\ &= \mathbf{x} - \eta \nabla_{S' \cup B} g(\mathbf{x}) - \bar{\mathbf{x}}_B. \end{aligned}$$

Using (9.69), we have

$$\begin{aligned} g(\mathbf{x}') - g(\mathbf{x}) &\leq \langle \nabla g(\mathbf{x}), -\eta \nabla_{S' \cup B} g(\mathbf{x}) - \bar{\mathbf{x}}_B \rangle + \beta \|-\eta \nabla_{S' \cup B} g(\mathbf{x}) - \bar{\mathbf{x}}_B\|_2^2 \\ &= -(4\beta)^{-1} \|\nabla_{S' \cup B} g(\mathbf{x})\|_2^2 + \beta \|\bar{\mathbf{x}}_B\|_2^2 \\ &= -\beta \|\eta \nabla_{S \cup A} g(\mathbf{x})\|_2^2 + \beta \|\bar{\mathbf{x}}_B\|_2^2 \\ &\leq -\beta \|\eta \nabla_{S \cup A'} g(\mathbf{x})\|_2^2 + \beta \|\bar{\mathbf{x}}_{B'}\|_2^2, \end{aligned} \quad (9.70)$$

for any two sets $A' \in [n] \setminus S$ and $B' \subseteq S$ with $|A'| = |B'|$. The latter inequality follows because of the following lemma about IHT:

Lemma 9.5.1. *Suppose that we run one step of IHT on vector \mathbf{x} supported on S for some function g , and let the updated solution vector be $\mathbf{x}' = \bar{\mathbf{x}}_{(S \cup A) \setminus B}$, where $\bar{\mathbf{x}} = \mathbf{x} - \eta \nabla g(\mathbf{x})$. Then, for any $A' \subseteq [n] \setminus S$ and $B' \subseteq S$ with $|A'| = |B'|$, we have*

$$-\|\eta \nabla_{A'} g(\mathbf{x})\|_2^2 + \|\bar{\mathbf{x}}_B\|_2^2 \leq -\|\eta \nabla_{A'} g(\mathbf{x})\|_2^2 + \|\bar{\mathbf{x}}_{B'}\|_2^2. \quad (9.71)$$

Proof. If we denote $|A| = |B| = t$ and $|A'| = |B'| = t'$, then note that by definition of IHT, A are the t largest entries in

$$|\bar{\mathbf{x}}_{[n] \setminus S}| = \eta |\nabla_{[n] \setminus S} g(\mathbf{x})|,$$

and B are the t smallest entries in $|\bar{\mathbf{x}}_S|$. Similarly, we can assume that A' are the t' largest entries in $\eta |\nabla_{[n] \setminus S} g(\mathbf{x})|$ and B' are the t' smallest entries in $|\bar{\mathbf{x}}_S|$, since this way the right hand side of (9.71) takes its minimum value. If $t' = t$, we are done. We consider two cases:

1. $t' > t$: In this case we have $A' \supseteq A$, $B' \supseteq B$, so

$$\begin{aligned}
& - \|\eta \nabla_{A'} g(\mathbf{x})\|_2^2 + \|\bar{\mathbf{x}}_{B'}\|_2^2 + \|\eta \nabla_A g(\mathbf{x})\|_2^2 - \|\bar{\mathbf{x}}_B\|_2^2 \\
& = - \|\eta \nabla_{A' \setminus A} g(\mathbf{x})\|_2^2 + \|\bar{\mathbf{x}}_{B' \setminus B}\|_2^2 \\
& = - \|\bar{\mathbf{x}}_{A' \setminus A}\|_2^2 + \|\bar{\mathbf{x}}_{B' \setminus B}\|_2^2 \\
& \geq (t' - t) \left(- \max_{i \in A' \setminus A} (\bar{x}_i)^2 + \min_{j \in B' \setminus B} (\bar{x}_j)^2 \right) \\
& \geq 0,
\end{aligned}$$

where the last inequality follows since, by definition of the IHT step, $|\bar{x}_i| \leq |\bar{x}_j|$ for any $i \in A' \setminus A$ and $j \in B' \setminus B$. Otherwise, i would have taken j 's place in S' .

2. $t' < t$: In this case we have $A' \subseteq A$, $B' \subseteq B$. Similarly to the previous case,

$$\begin{aligned}
& - \|\eta \nabla_{A'} g(\mathbf{x})\|_2^2 + \|\bar{\mathbf{x}}_{B'}\|_2^2 + \|\eta \nabla_A g(\mathbf{x})\|_2^2 - \|\bar{\mathbf{x}}_B\|_2^2 \\
& = \|\eta \nabla_{A \setminus A'} g(\mathbf{x})\|_2^2 - \|\bar{\mathbf{x}}_{B \setminus B'}\|_2^2 \\
& = \|\bar{\mathbf{x}}_{A \setminus A'}\|_2^2 - \|\bar{\mathbf{x}}_{B \setminus B'}\|_2^2 \\
& \geq (t - t') \left(\min_{i \in A \setminus A'} (\bar{x}_i)^2 - \max_{j \in B \setminus B'} (\bar{x}_j)^2 \right) \\
& \geq 0,
\end{aligned}$$

where the last inequality follows since, by definition of the IHT step, $|\bar{x}_i| \geq |\bar{x}_j|$ for any $i \in A \setminus A'$ and $j \in B \setminus B'$. Otherwise i wouldn't have taken j 's place in S' .

□

Now, let us assume that the first bullet in the lemma statement is false, i.e.

$$g(\mathbf{x}') - g(\mathbf{x}) > -(16\kappa)^{-1}(g(\mathbf{x}) - f(\mathbf{x}^*)).$$

Setting $A' = B' = \emptyset$ in (9.70), we get that

$$g(\mathbf{x}') - g(\mathbf{x}) \leq -\beta \|\eta \nabla_S g(\mathbf{x})\|_2^2,$$

so we conclude that

$$\|\nabla_S g(\mathbf{x})\|_2^2 < \frac{1}{16\kappa\beta\eta^2} (g(\mathbf{x}) - f(\mathbf{x}^*)) = \frac{\alpha}{4} (g(\mathbf{x}) - f(\mathbf{x}^*)). \quad (9.72)$$

Now, we again use (9.70) but we set A' to be the s entries from $[n] \setminus S$ on which $\nabla g(\mathbf{x})$ has the largest magnitude, and B' to be the s entries from S on which $\bar{\mathbf{x}}$ has the smallest magnitude. Also,

let R be an arbitrary subset of $S \setminus S^*$ with size $r \geq 2s$. We then have

$$\begin{aligned}
& g(\mathbf{x}') - g(\mathbf{x}) \\
& \leq -(4\beta)^{-1} \|\nabla_{S \cup A'} g(\mathbf{x})\|_2^2 + \beta \|\bar{\mathbf{x}}_{B'}\|_2^2 \\
& \leq -(4\beta)^{-1} \|\nabla_{S \cup S^*} g(\mathbf{x})\|_2^2 + \beta \|\bar{\mathbf{x}}_{B'}\|_2^2 \\
& \leq -(4\beta)^{-1} \|\nabla_{S \cup S^*} g(\mathbf{x})\|_2^2 + \frac{\beta s}{|R \setminus S^*|} \|\bar{\mathbf{x}}_{R \setminus S^*}\|_2^2 \\
& \leq -(4\beta)^{-1} \|\nabla_{S \cup S^*} g(\mathbf{x})\|_2^2 + \frac{\beta s}{r-s} \|\bar{\mathbf{x}}_{R \setminus S^*}\|_2^2,
\end{aligned} \tag{9.73}$$

where we used the fact that

$$\|\nabla_{A'} g(\mathbf{x})\|_2^2 \geq \|\nabla_{S^* \setminus S} g(\mathbf{x})\|_2^2$$

by definition of A' (and since $|S^* \setminus S| \leq s$), and the fact that, by definition of B' (and since $|R \setminus S^*| \geq 2s - s = |B'|$),

$$\frac{1}{|B'|} \|\bar{\mathbf{x}}_{B'}\|_2^2 \leq \frac{1}{|R \setminus S^*|} \|\bar{\mathbf{x}}_{R \setminus S^*}\|_2^2.$$

In fact, we will let $R = \{i \in S \mid w_i > 0\}$ be the set of elements that are being regularized. To lower bound the size r of this set, note that by the guarantee of the lemma statement,

$$\begin{aligned}
n - s'/2 & \leq \|\mathbf{w}\|_1 \\
& \leq n - |\{i \in S \mid w_i = 0\}| \\
& = n - (s' - r),
\end{aligned}$$

so $r \geq s'/2$. We conclude that $r \geq 2s$ since $s' \geq 4s$.

Now, because of the fact that f is α -strongly convex, we have

$$\begin{aligned}
& f(\mathbf{x}^*) - f(\mathbf{x}) \\
& \geq \langle \nabla f(\mathbf{x}), \mathbf{x}^* - \mathbf{x} \rangle + (\alpha/2) \|\mathbf{x}^* - \mathbf{x}\|_2^2 \\
& = \langle \nabla g(\mathbf{x}), \mathbf{x}^* - \mathbf{x} \rangle - \beta \langle \mathbf{w}\mathbf{x}, \mathbf{x}^* - \mathbf{x} \rangle + (\alpha/2) \|\mathbf{x}^* - \mathbf{x}\|_2^2 \\
& \geq -\alpha^{-1} \|\nabla_{S \cup S^*} g(\mathbf{x})\|_2^2 - \beta \langle \mathbf{w}\mathbf{x}, \mathbf{x}^* - \mathbf{x} \rangle + (\alpha/4) \|\mathbf{x}^* - \mathbf{x}\|_2^2,
\end{aligned} \tag{9.74}$$

where we used the inequality

$$\langle \mathbf{a}, \mathbf{b} \rangle + (\alpha/4) \|\mathbf{b}\|_2^2 \geq -\alpha^{-1} \|\mathbf{a}\|_2^2.$$

By re-arranging and plugging (9.74) into (9.73), we get

$$\begin{aligned}
& g(\mathbf{x}') - g(\mathbf{x}) \\
& \leq -(4\kappa)^{-1} \left(f(\mathbf{x}) - f(\mathbf{x}^*) - \beta \langle \mathbf{w}\mathbf{x}, \mathbf{x}^* - \mathbf{x} \rangle + (\alpha/4) \|\mathbf{x}^* - \mathbf{x}\|_2^2 \right) + \frac{\beta s}{r-s} \|\bar{\mathbf{x}}_{R \setminus S^*}\|_2^2 \\
& = -(4\kappa)^{-1} \left(g(\mathbf{x}) - f(\mathbf{x}^*) - \beta \langle \mathbf{w}\mathbf{x}, \mathbf{x}^* - \mathbf{x} \rangle - (\beta/2) \|\mathbf{x}\|_{\mathbf{w},2}^2 + (\alpha/4) \|\mathbf{x}^* - \mathbf{x}\|_2^2 - \frac{4\kappa\beta s}{r-s} \|\bar{\mathbf{x}}_{R \setminus S^*}\|_2^2 \right).
\end{aligned} \tag{9.75}$$

Now, note that by definition of $\bar{\mathbf{x}}$ we have

$$\begin{aligned} & \|\bar{\mathbf{x}}_{R \setminus S^*}\|_2^2 \\ & \leq 2 \|\mathbf{x}_{R \setminus S^*}\|_2^2 + 2(2\beta)^{-2} \|\nabla_{R \setminus S^*} g(\mathbf{x})\|_2^2 \end{aligned}$$

and, since $w_i \geq 1/2$ for each $i \in R$,

$$\|\mathbf{x}_{R \setminus S^*}\|_2^2 \leq 2 \|\mathbf{x}_{R \setminus S^*}\|_{\mathbf{w},2}^2 \leq 2 \|\mathbf{x}_{S \setminus S^*}\|_{\mathbf{w},2}^2.$$

Therefore,

$$\begin{aligned} & \frac{4\kappa\beta s}{r-s} \|\bar{\mathbf{x}}_{R \setminus S^*}\|_2^2 \\ & \leq \frac{16\kappa\beta s}{r-s} \|\mathbf{x}_{S \setminus S^*}\|_{\mathbf{w},2}^2 + \frac{2\kappa s}{\beta(r-s)} \|\nabla_{R \setminus S^*} g(\mathbf{x})\|_2^2 \\ & \leq \frac{16\kappa\beta s}{r-s} \|\mathbf{x}_{S \setminus S^*}\|_{\mathbf{w},2}^2 + \frac{s}{2(r-s)} (g(\mathbf{x}) - f(\mathbf{x}^*)), \end{aligned}$$

where the last inequality follows from (9.72) since $R \setminus S^* \subseteq S$. Plugging this back into (9.75), we get

$$\begin{aligned} & g(\mathbf{x}') - g(\mathbf{x}) \\ & \leq -(4\kappa)^{-1} \left(\left(1 - \frac{s}{2(r-s)} \right) (g(\mathbf{x}) - f(\mathbf{x}^*)) \right. \\ & \quad \left. - \beta \langle \mathbf{w}\mathbf{x}, \mathbf{x}^* - \mathbf{x} \rangle - (\beta/2) \|\mathbf{x}\|_{\mathbf{w},2}^2 + (\alpha/4) \|\mathbf{x}^* - \mathbf{x}\|_2^2 - \frac{16\kappa\beta s}{r-s} \|\mathbf{x}_{S \setminus S^*}\|_{\mathbf{w},2}^2 \right) \\ & = -(4\kappa)^{-1} \left(\left(1 - \frac{s}{2(r-s)} \right) (g(\mathbf{x}) - f(\mathbf{x}^*)) \right. \\ & \quad \left. - \underbrace{\beta \langle \mathbf{w}\mathbf{x}_{S \cap S^*}, \mathbf{x}^* - \mathbf{x} \rangle + (\alpha/4) \|\mathbf{x}^* - \mathbf{x}\|_2^2}_{\geq -(\kappa\beta) \|\mathbf{x}_{S \cap S^*}\|_{\mathbf{w},2}^2} + \beta \|\mathbf{x}_{S \setminus S^*}\|_{\mathbf{w},2}^2 - (\beta/2) \|\mathbf{x}\|_{\mathbf{w},2}^2 - \underbrace{\frac{16\kappa\beta s}{r-s} \|\mathbf{x}_{S \setminus S^*}\|_{\mathbf{w},2}^2}_{\geq -(\beta/4) \|\mathbf{x}_{S \setminus S^*}\|_{\mathbf{w},2}^2} \right) \\ & \leq -(4\kappa)^{-1} \left(\left(1 - \frac{s}{2(r-s)} \right) (g(\mathbf{x}) - f(\mathbf{x}^*)) - (\kappa\beta) \|\mathbf{x}_{S \cap S^*}\|_{\mathbf{w},2}^2 - (\beta/2) \|\mathbf{x}_{S \cap S^*}\|_{\mathbf{w},2}^2 + (\beta/4) \|\mathbf{x}_{S \setminus S^*}\|_{\mathbf{w},2}^2 \right) \\ & \leq -(4\kappa)^{-1} \left(0.5 (g(\mathbf{x}) - f(\mathbf{x}^*)) - (\kappa + 1)\beta \|\mathbf{x}_{S \cap S^*}\|_{\mathbf{w},2}^2 + (\beta/4) \|\mathbf{x}_{S \setminus S^*}\|_{\mathbf{w},2}^2 \right), \end{aligned} \tag{9.76}$$

where we used the fact that

$$\frac{s}{2(r-s)} \leq 1/2,$$

which holds as long as $s' \geq 4s$, and

$$\frac{16\kappa\beta s}{r-s} \leq \beta/4,$$

which holds as long as $r \geq s'/2$ and $s' \geq (128\kappa + 2)s$. In the last inequality we also used the property $\mathbf{w}/2 \leq \mathbf{w}^2$, which is by definition of \mathbf{w} .

Now, note that, because we have assumed that the first bullet of the statement doesn't hold, it

has to be the case that

$$(1/4)(g(\mathbf{x}) - f(\mathbf{x}^*)) - (\kappa + 1)\beta \|\mathbf{x}_{S \cap S^*}\|_{\mathbf{w},2}^2 + (\beta/4) \|\mathbf{x}_{S \setminus S^*}\|_{\mathbf{w},2}^2 \leq 0.$$

This immediately implies that

$$\begin{aligned} (\kappa + 1)\beta \|\mathbf{x}_{S \cap S^*}\|_{\mathbf{w},2}^2 &\geq (\beta/4) \|\mathbf{x}_{S \setminus S^*}\|_{\mathbf{w},2}^2 \\ \Rightarrow (4\kappa + 4) \|\mathbf{x}_{S \cap S^*}\|_{\mathbf{w},2}^2 &\geq \|\mathbf{x}_{S \setminus S^*}\|_{\mathbf{w},2}^2 \\ \Rightarrow (4\kappa + 6) \|\mathbf{x}_{S \cap S^*}\|_{\mathbf{w},2}^2 &\geq \|\mathbf{x}\|_{\mathbf{w},2}^2, \end{aligned}$$

so

$$\|\mathbf{x}_{S \cap S^*}\|_{\mathbf{w},2}^2 \geq (4\kappa + 6)^{-1} \|\mathbf{x}\|_{\mathbf{w},2}^2.$$

Similarly we also have

$$\begin{aligned} (\kappa + 1)\beta \|\mathbf{x}_{S \cap S^*}\|_{\mathbf{w},2}^2 &\geq (1/4)(g(\mathbf{x}) - f(\mathbf{x}^*)) \\ \Rightarrow (\beta/2) \|\mathbf{x}_{S \cap S^*}\|_{\mathbf{w},2}^2 &\geq (8\kappa + 8)^{-1}(g(\mathbf{x}) - f(\mathbf{x}^*)). \end{aligned}$$

Therefore the second bullet of the statement is true, and we are done. □

9.5.4 Low Rank Minimization

Preliminaries

We will use the following simple lemma about Frobenius products between low rank projections and symmetric PSD matrices. We remind the reader that $H_r(\mathbf{A})$ is the matrix consisting of the top r components from the singular value decomposition of \mathbf{A} .

Lemma 9.5.2. *For any two symmetric PSD matrices $\mathbf{\Pi}, \mathbf{A} \in \mathbb{R}^{n \times n}$, where $\text{rank}(\mathbf{\Pi}) \leq r$ and $\|\mathbf{\Pi}\|_2 \leq 1$, we have that*

$$|\langle \mathbf{\Pi}, \mathbf{A} \rangle| \leq \text{Tr}[H_r(\mathbf{A})].$$

Proof. We will use the following inequality for singular values

$$\sum_{i=1}^k \sigma_i(AB) \leq \sum_{i=1}^k \sigma_i(A)\sigma_i(B)$$

for $k = 1, \dots, n$, $A, B \in \mathbb{R}^{n \times n}$ and $\sigma_1(A) \geq \dots \geq \sigma_n(A)$ are singular values of matrix A (see page

Algorithm 25 Regularized Local Search

\mathbf{A}^0 : initial rank- r' solution

$\mathbf{W}^0 = \mathbf{Y}^0 = \mathbf{I}$: initial regularization weights

η : step size, T : #iterations

c : weight step size

for $t = 0, \dots, T - 1$ **do**

$$\Phi(\mathbf{A}) := (\beta/4) (\langle \mathbf{W}^t, \mathbf{A}\mathbf{A}^\top \rangle + \langle \mathbf{Y}^t, \mathbf{A}^\top \mathbf{A} \rangle)$$

$$g(\mathbf{A}) := f(\mathbf{A}) + \Phi(\mathbf{A})$$

$$\bar{\mathbf{A}} = H_{s'-1}(\mathbf{A}^t) - 0.5H_1(\eta\nabla g(\mathbf{A}^t))$$

$$\mathbf{P} = (\mathbf{W}^t)^{1/2} \mathbf{A}^t (\mathbf{A}^t)^\top (\mathbf{W}^t)^{1/2}$$

$$\mathbf{Q} = (\mathbf{Y}^t)^{1/2} (\mathbf{A}^t)^\top \mathbf{A}^t (\mathbf{Y}^t)^{1/2}$$

$$\Delta = g(\mathbf{A}^t) - f(\mathbf{A}^*)$$

if $g(\mathbf{A}^t) - g(\bar{\mathbf{A}}) \geq (r')^{-1}\Delta$ **then**

Let $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ be the SVD of $\bar{\mathbf{A}}$

$$\mathbf{A}^{t+1} = \operatorname{argmin}_{\mathbf{A}} g(\mathbf{A})$$

$$\mathbf{W}^{t+1}, \mathbf{Y}^{t+1} = \mathbf{W}^t, \mathbf{Y}^t$$

else if $\max\{\operatorname{Tr}[H_r(\mathbf{P})], \operatorname{Tr}[H_r(\mathbf{Q})]\} \geq (0.4/\beta)\Delta$ **then**

$$\mathbf{A}^{t+1} = \mathbf{A}^t$$

$$\mathbf{W}^{t+1} = (\mathbf{W}^t)^{1/2} (\mathbf{I} - r^{-1} \mathbf{\Pi}_{\operatorname{im}(\mathbf{P})}) (\mathbf{W}^t)^{1/2}$$

$$\mathbf{Y}^{t+1} = (\mathbf{Y}^t)^{1/2} (\mathbf{I} - r^{-1} \mathbf{\Pi}_{\operatorname{im}(\mathbf{Q})}) (\mathbf{Y}^t)^{1/2}$$

else

$$\mathbf{A}^{t+1} = \mathbf{A}^t$$

$$\mathbf{W}^{t+1} = \mathbf{W}^t - \frac{\mathbf{W}^t \mathbf{A}^t (\mathbf{A}^t)^\top \mathbf{W}^t}{\langle \mathbf{W}^t, \mathbf{A}^t (\mathbf{A}^t)^\top \rangle}$$

$$\mathbf{Y}^{t+1} = \mathbf{Y}^t - \frac{\mathbf{Y}^t (\mathbf{A}^t)^\top \mathbf{A}^t \mathbf{Y}^t}{\langle \mathbf{Y}^t, (\mathbf{A}^t)^\top \mathbf{A}^t \rangle}$$

177 in [82]). Then

$$\begin{aligned}
|\langle \Pi, \mathbf{A} \rangle| &= \text{Tr} \left[\Pi \mathbf{A}^\top \right] \\
&= \text{Tr} \left[\Pi \mathbf{A} \right] \\
&= \sum_{i=1}^n \sigma_i(\Pi \mathbf{A}) \\
&\leq \sum_{i=1}^n \sigma_i(\Pi) \sigma_i(\mathbf{A}) \\
&= \sum_{i=1}^r \sigma_i(\Pi) \sigma_i(\mathbf{A}) \\
&\leq \sum_{i=1}^r \sigma_i(\mathbf{A}) \\
&= \text{Tr} \left[H_r(\mathbf{A}) \right].
\end{aligned}$$

□

Analysis

This section is devoted to proving Theorem 7.1.2, which analyzes an algorithm for low rank optimization that uses adaptive regularization.

Theorem 7.1.2 (Adaptive Regularization for Low Rank Optimization). *Let $f \in \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ be a convex function with condition number κ and consider the low rank minimization problem*

$$\min_{\text{rank}(\mathbf{A}) \leq r} f(\mathbf{A}). \quad (9.77)$$

For any error parameter $\varepsilon > 0$, there exists a polynomial time algorithm that returns a matrix \mathbf{A} with $\text{rank}(\mathbf{A}) \leq O\left(r\left(\kappa + \log \frac{f(\mathbf{O}) - f(\mathbf{A}^)}{\varepsilon}\right)\right)$ and $f(\mathbf{A}) \leq f(\mathbf{A}^*) + \varepsilon$, where \mathbf{A}^* is any rank- r matrix.*

Proof of Theorem 7.1.2. Let the smoothness and strong convexity parameters of f be β, α . We repeatedly apply Lemma 9.5.3 $T \geq O\left(r\kappa \log \frac{f(\mathbf{A}^0) + (\beta/2)\|\mathbf{A}^0\|_F^2 - f(\mathbf{A}^*)}{\varepsilon}\right)$ times starting from solution $\mathbf{A}^0 = \mathbf{O}$ and weight matrices $\mathbf{W}^0 = \mathbf{I}$, $\mathbf{Y}^0 = \mathbf{I}$. Thus, we obtain solutions $\mathbf{A}^0, \dots, \mathbf{A}^T$, and weights $\mathbf{W}^0, \mathbf{W}^1, \dots, \mathbf{W}^T$ and $\mathbf{Y}^0, \mathbf{Y}^1, \dots, \mathbf{Y}^T$. We let

$$\begin{aligned}
g^t(\mathbf{A}) &= f(\mathbf{A}) + (\beta/4) \left(\langle \mathbf{W}^t, \mathbf{A}^t (\mathbf{A}^t)^\top \rangle + \langle \mathbf{Y}^t, (\mathbf{A}^t)^\top \mathbf{A}^t \rangle \right)
\end{aligned}$$

be the regularized function at iteration t .

We denote by T_i the total number of iterations for which item $i \in \{1, 2, 3\}$ from the statement of Lemma 9.5.3 holds.

Consider the T_2 iterations for which item 2 from the statement of Lemma 9.5.3 holds. Without loss of generality, \mathbf{W} is updated at least $T_2/2$ times. Letting $\mathbf{A}^* = \mathbf{U}^* \Sigma^* \mathbf{V}^{*\top}$ be the singular value

decomposition of \mathbf{A}^* , for each such iteration we have

$$\begin{aligned} & \text{Tr} [\mathbf{H}_{\text{im}(\mathbf{U}^*)} \mathbf{W}^{t+1} \mathbf{H}_{\text{im}(\mathbf{U}^*)}] \\ & \leq \text{Tr} [\mathbf{H}_{\text{im}(\mathbf{U}^*)} \mathbf{W}^t \mathbf{H}_{\text{im}(\mathbf{U}^*)}] - (10\kappa)^{-1}, \end{aligned}$$

and for all other types of iterations we have $\mathbf{W}^{t+1} \preceq \mathbf{W}^t$. Therefore,

$$\begin{aligned} & \text{Tr} [\mathbf{H}_{\text{im}(\mathbf{U}^*)} \mathbf{W}^T \mathbf{H}_{\text{im}(\mathbf{U}^*)}] \\ & \leq \text{Tr} [\mathbf{H}_{\text{im}(\mathbf{U}^*)} \mathbf{W}^0 \mathbf{H}_{\text{im}(\mathbf{U}^*)}] - \frac{T_2}{2} (10\kappa)^{-1}. \end{aligned}$$

However, note that by the guarantee of Lemma 9.5.3 that $\mathbf{W}^T \succeq \mathbf{O}$, we have

$$\text{Tr} [\mathbf{H}_{\text{im}(\mathbf{U}^*)} \mathbf{W}^T \mathbf{H}_{\text{im}(\mathbf{U}^*)}] \geq 0,$$

and because $\mathbf{W}^0 = \mathbf{I}$ we also know that

$$\text{Tr} [\mathbf{H}_{\text{im}(\mathbf{U}^*)} \mathbf{W}^0 \mathbf{H}_{\text{im}(\mathbf{U}^*)}] = \text{Tr} [\mathbf{H}_{\text{im}(\mathbf{U}^*)}] \leq r.$$

This implies that $T_2 \leq 20\kappa r$.

Now, if $T_1 \geq 16r\kappa \log \frac{g^0(\mathbf{A}^0) - f(\mathbf{A}^*)}{\varepsilon}$, and since $g^t(\mathbf{A}^t)$ is non-increasing for all t , we have

$$\begin{aligned} & g^T(\mathbf{A}^T) - f(\mathbf{A}^*) \\ & \leq (1 - (16r\kappa)^{-1})^{T_1} (g^0(\mathbf{A}^0) - f(\mathbf{A}^*)) \\ & \leq \varepsilon, \end{aligned}$$

so $T_1 \leq 16r\kappa \log \frac{g^0(\mathbf{A}^0) - f(\mathbf{A}^*)}{\varepsilon}$.

Similarly, if $T_3 \geq 10r \log \frac{g^0(\mathbf{A}^0) - f(\mathbf{A}^*)}{\varepsilon}$ we have

$$\begin{aligned} & g^T(\mathbf{A}^T) - f(\mathbf{A}^*) \\ & \leq (1 - (10r)^{-1})^{T_3} (g^0(\mathbf{A}^0) - f(\mathbf{A}^*)) \\ & \leq \varepsilon, \end{aligned}$$

so $T_3 \geq 10r \log \frac{g^0(\mathbf{A}^0) - f(\mathbf{A}^*)}{\varepsilon}$.

Overall, we have that the total number of iterations is

$$T = \sum T_i \leq 36r(\kappa + 1) \log \frac{g^0(\mathbf{A}^0) - f(\mathbf{A}^*)}{\varepsilon}.$$

The only thing left is to ensure that the conditions

$$\begin{aligned} \text{Tr} [\mathbf{I} - \mathbf{W}^t] & \leq r'/2 \\ \text{Tr} [\mathbf{I} - \mathbf{Y}^t] & \leq r'/2 \end{aligned}$$

of Lemma 9.5.3 are satisfied for all t . By the guarantees of Lemma 9.5.3, if one of items 2, 3 holds, then

$$\text{Tr} [\mathbf{I} - \mathbf{W}^{t+1}] \leq \text{Tr} [\mathbf{I} - \mathbf{W}^t] + 1,$$

and if item 1 holds, then

$$\mathrm{Tr} [\mathbf{I} - \mathbf{W}^{t+1}] = \mathrm{Tr} [\mathbf{I} - \mathbf{W}^t] .$$

As $\mathrm{Tr} [\mathbf{I} - \mathbf{W}^0] = 0$, we have

$$\begin{aligned} & \mathrm{Tr} [\mathbf{I} - \mathbf{W}^T] \\ & \leq T_2 + T_3 \\ & \leq 20\kappa r + 10r \log \frac{g^0(\mathbf{A}^0) - f(\mathbf{A}^*)}{\varepsilon} \\ & \leq r'/2, \end{aligned}$$

where the last inequality holds as long as

$$r' \geq 20r \left(2\kappa + \log \frac{g^0(\mathbf{A}^0) - f(\mathbf{A}^*)}{\varepsilon} \right) .$$

□

Lemma 9.5.3 (Low rank minimization step analysis). *Let $f : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ be a β -smooth and α -strongly convex function with condition number $\kappa = \beta/\alpha$, and $\mathbf{W} \in \mathbb{R}^{m \times m}$, $\mathbf{Y} \in \mathbb{R}^{n \times n}$ be symmetric positive semi-definite weight matrices with spectral norm bounded by 1 and such that $\mathrm{Tr} [\mathbf{I} - \mathbf{W}] \leq r'/2$ and $\mathrm{Tr} [\mathbf{I} - \mathbf{Y}] \leq r'/2$ for fixed parameter $r' \geq 256r$. We define the regularized function*

$$g(\mathbf{A}) := f(\mathbf{A}) + \underbrace{(\beta/4) \left(\langle \mathbf{W}, \mathbf{A}\mathbf{A}^\top \rangle + \langle \mathbf{Y}, \mathbf{A}^\top \mathbf{A} \rangle \right)}_{\Phi(\mathbf{A})} .$$

Now, consider a rank- r' matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ with singular value decomposition

$$\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{V}^\top = \sum_{j \in S} \lambda_j \mathbf{u}_j \mathbf{v}_j^\top$$

and with the property that

$$\mathbf{\Pi}_{\mathrm{im}(\mathbf{U})} \cdot \nabla g(\mathbf{A}) \cdot \mathbf{\Pi}_{\mathrm{im}(\mathbf{V})} = \mathbf{O} .$$

For any rank- r solution \mathbf{A}^* where $r' \geq 256r$, there is a procedure that updates \mathbf{A} , \mathbf{W} , \mathbf{Y} , and for which exactly one of the following scenarios holds:

1. \mathbf{A} is updated to a rank- r' matrix \mathbf{A}' , and \mathbf{W} , \mathbf{Y} are not updated. We have sufficient progress in the regularized function:

$$g(\mathbf{A}') \leq g(\mathbf{A}) - (16\kappa r)^{-1} (g(\mathbf{A}) - f(\mathbf{A}^*)) .$$

2. Exactly one of \mathbf{W} or \mathbf{Y} is updated (wlog \mathbf{W}) to a symmetric PSD $\mathbf{W}' \preceq \mathbf{W}$, and \mathbf{A} is not updated. We have

$$\mathrm{Tr} [\mathbf{I} - \mathbf{W}'] \leq \mathrm{Tr} [\mathbf{I} - \mathbf{W}] + 1$$

and

$$\begin{aligned} & \text{Tr} [\mathbf{H}_{\text{im}(\mathbf{U}^*)} \mathbf{W}' \mathbf{H}_{\text{im}(\mathbf{U}^*)}] \\ & \leq \text{Tr} [\mathbf{H}_{\text{im}(\mathbf{U}^*)} \mathbf{W} \mathbf{H}_{\text{im}(\mathbf{U}^*)}] - (10\kappa)^{-1}. \end{aligned}$$

Respectively, for \mathbf{Y} :

$$\begin{aligned} & \text{Tr} [\mathbf{H}_{\text{im}(\mathbf{V}^*)} \mathbf{Y}' \mathbf{H}_{\text{im}(\mathbf{V}^*)}] \\ & \leq \text{Tr} [\mathbf{H}_{\text{im}(\mathbf{V}^*)} \mathbf{Y} \mathbf{H}_{\text{im}(\mathbf{V}^*)}] - (10\kappa)^{-1}. \end{aligned}$$

3. Exactly one of \mathbf{W} or \mathbf{Y} is updated (wlog \mathbf{W}) to a symmetric PSD $\mathbf{W}' \preceq \mathbf{W}$, and \mathbf{A} is not updated. We have sufficient progress in the regularized function, where g' is the regularized function with the new weights:

$$g'(\mathbf{A}) \leq g(\mathbf{A}) - (10r)^{-1} (g(\mathbf{A}) - f(\mathbf{A}^*)).$$

Additionally,

$$\text{Tr} [\mathbf{I} - \mathbf{W}'] \leq \text{Tr}[\mathbf{I} - \mathbf{W}] + 1$$

Proof. We attempt to make the update $\mathbf{A} \rightarrow \mathbf{A}'$ as defined in Lemma 9.5.4. If it makes enough progress, i.e.

$$g(\mathbf{A}') \leq g(\mathbf{A}) - (16\kappa r)^{-1} (g(\mathbf{A}) - f(\mathbf{A}^*)),$$

we are done. Otherwise, one of the items 2-5 in the statement of Lemma 9.5.4 must hold. Let us take them one by one.

Item 2:

$$\langle \mathbf{H}_{\text{im}(\mathbf{U}^*)}, \mathbf{W} \mathbf{A} \mathbf{A}^\top \mathbf{W} \rangle \geq (10\kappa)^{-1} \langle \mathbf{W}, \mathbf{A} \mathbf{A}^\top \rangle.$$

We update \mathbf{W} as

$$\mathbf{W}' = \mathbf{W} - c \cdot \mathbf{W} \mathbf{A} \mathbf{A}^\top \mathbf{W},$$

where $c = \langle \mathbf{W}, \mathbf{A} \mathbf{A}^\top \rangle^{-1}$. Note that this update preserves symmetry, and

$$\mathbf{O} \preceq \mathbf{W}' \preceq \mathbf{W}.$$

This is because

$$c \mathbf{W}^{1/2} \mathbf{A} \mathbf{A}^\top \mathbf{W}^{1/2} \preceq c \langle \mathbf{W}, \mathbf{A} \mathbf{A}^\top \rangle \cdot \mathbf{I} \preceq \mathbf{I},$$

so

$$\mathbf{W}' = \mathbf{W}^{1/2} \left(\mathbf{I} - c \mathbf{W}^{1/2} \mathbf{A} \mathbf{A}^\top \mathbf{W}^{1/2} \right) \mathbf{W}^{1/2} \succeq \mathbf{O}$$

and

$$\mathbf{W}' = \mathbf{W} - c \mathbf{W} \mathbf{A} \mathbf{A}^\top \mathbf{W} \preceq \mathbf{W}.$$

Now, note that

$$\begin{aligned}\mathrm{Tr} [\mathbf{I} - \mathbf{W}'] &= \mathrm{Tr} [\mathbf{I} - \mathbf{W}] + c \langle \mathbf{W}^2, \mathbf{A} \mathbf{A}^\top \rangle \\ &\leq \mathrm{Tr} [\mathbf{I} - \mathbf{W}] + c \langle \mathbf{W}, \mathbf{A} \mathbf{A}^\top \rangle \\ &= \mathrm{Tr} [\mathbf{I} - \mathbf{W}] + 1,\end{aligned}$$

where we used the fact that $\mathbf{W}^2 \preceq \mathbf{W}$, and (letting $\mathbf{\Pi}^* = \mathbf{\Pi}_{\mathrm{im}(\mathbf{U}^*)}$ for convenience),

$$\begin{aligned}\mathrm{Tr} [\mathbf{\Pi}^* \mathbf{W}' \mathbf{\Pi}^*] &= \mathrm{Tr} [\mathbf{\Pi}^* \mathbf{W} \mathbf{\Pi}^*] - c \langle \mathbf{\Pi}^*, \mathbf{W} \mathbf{A} \mathbf{A}^\top \mathbf{W} \rangle \\ &\leq \mathrm{Tr} [\mathbf{\Pi}^* \mathbf{W} \mathbf{\Pi}^*] - c / (10\kappa) \langle \mathbf{W}, \mathbf{A} \mathbf{A}^\top \rangle \\ &= \mathrm{Tr} [\mathbf{\Pi}^* \mathbf{W} \mathbf{\Pi}^*] - (10\kappa)^{-1},\end{aligned}\tag{9.78}$$

Item 3:

$$\langle \mathbf{\Pi}_{\mathrm{im}(\mathbf{V}^*)}, \mathbf{Y} \mathbf{A}^\top \mathbf{A} \mathbf{Y} \rangle \geq (10\kappa)^{-1} \langle \mathbf{Y}, \mathbf{A}^\top \mathbf{A} \rangle.$$

This is entirely analogous to the previous case.

Item 4:

$$(\beta/4) \mathrm{Tr} \left[H_r \left(\mathbf{A}^\top \mathbf{W} \mathbf{A} \right) \right] \geq 10^{-1} (g(\mathbf{A}) - f(\mathbf{A}^*)). \tag{9.79}$$

After considering the eigendecomposition

$$\mathbf{W}^{1/2} \mathbf{A} \mathbf{A}^\top \mathbf{W}^{1/2} = \sum_{i \in [r']} \bar{\lambda}_i \bar{\mathbf{v}}_i \bar{\mathbf{v}}_i^\top$$

with $\bar{\lambda}_1 \geq \bar{\lambda}_2 \geq \dots \geq \bar{\lambda}_{r'} \geq 0$, (9.79) can be re-phrased as

$$(\beta/4) \sum_{i \in [r]} \bar{\lambda}_i > (1/10) (g(\mathbf{A}) - f(\mathbf{A}^*)).$$

We update \mathbf{W} as

$$\mathbf{W}' = \mathbf{W}^{1/2} \left(\mathbf{I} - r^{-1} \sum_{i \in [r]} \bar{\mathbf{v}}_i \bar{\mathbf{v}}_i^\top \right) \mathbf{W}^{1/2}$$

and let g' be the new regularized objective. First of all, note that this operation preserves symmetry, and that $\mathbf{O} \preceq \mathbf{W}' \preceq \mathbf{I}$, since $\sum_{i \in [r]} \bar{\mathbf{v}}_i \bar{\mathbf{v}}_i^\top \preceq \mathbf{I}$. Additionally,

$$\begin{aligned}\mathrm{Tr} [\mathbf{I} - \mathbf{W}'] &= \mathrm{Tr} [\mathbf{I} - \mathbf{W}] + r^{-1} \sum_{i \in [r]} \bar{\mathbf{v}}_i^\top \mathbf{W} \bar{\mathbf{v}}_i \\ &\leq \mathrm{Tr} [\mathbf{W}] + 1\end{aligned}$$

and

$$\begin{aligned}
& g'(\mathbf{A}) - g(\mathbf{A}) \\
&= (\beta/4)\langle \mathbf{W}', \mathbf{A}\mathbf{A}^\top \rangle - (\beta/4)\langle \mathbf{W}, \mathbf{A}\mathbf{A}^\top \rangle \\
&= -(\beta/(4r)) \left\langle \mathbf{W}^{1/2} \left(\sum_{i \in [r]} \bar{\mathbf{v}}_i \bar{\mathbf{v}}_i^\top \right) \mathbf{W}^{1/2}, \mathbf{A}\mathbf{A}^\top \right\rangle \\
&= -(\beta/(4r)) \sum_{i \in [r]} \bar{\lambda}_i \\
&\leq -(10r)^{-1}(g(\mathbf{A}) - f(\mathbf{A}^*)),
\end{aligned}$$

Item 5:

$$(\beta/4) \operatorname{Tr} \left[H_r \left(\mathbf{A} \mathbf{Y} \mathbf{A}^\top \right) \right] \geq 10^{-1} (g(\mathbf{A}) - f(\mathbf{A}^*)).$$

This is entirely analogous to the previous case. \square

Lemma 9.5.4. *Let $f : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ be a β -smooth and α -strongly convex function with condition number $\kappa = \beta/\alpha$, and $\mathbf{W} \in \mathbb{R}^{m \times m}$, $\mathbf{Y} \in \mathbb{R}^{n \times n}$ be symmetric positive semi-definite weight matrices with spectral norm bounded by 1 and such that $\operatorname{Tr}[\mathbf{I} - \mathbf{W}]$, $\operatorname{Tr}[\mathbf{I} - \mathbf{Y}] \leq r'/2$ for some parameter $r' \geq 0$. We define the regularized function*

$$g(\mathbf{A}) := f(\mathbf{A}) + \underbrace{(\beta/4) \left(\langle \mathbf{W}, \mathbf{A}\mathbf{A}^\top \rangle + \langle \mathbf{Y}, \mathbf{A}^\top \mathbf{A} \rangle \right)}_{\Phi(\mathbf{A})}.$$

Now, consider a rank- r' matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ with singular value decomposition

$$\mathbf{A} = \mathbf{U} \mathbf{\Lambda} \mathbf{V}^\top = \sum_{j \in S} \lambda_j \mathbf{u}_j \mathbf{v}_j^\top$$

and with the property that

$$\mathbf{\Pi}_{\operatorname{im}(\mathbf{U})} \cdot \nabla g(\mathbf{A}) \cdot \mathbf{\Pi}_{\operatorname{im}(\mathbf{V})} = \mathbf{O}.$$

We define an updated solution

$$\mathbf{A}' = \mathbf{A} - \eta \cdot H_1(\nabla g(\mathbf{A})) - \lambda_j \mathbf{u}_j \mathbf{v}_j^\top,$$

where $\eta = (2\beta)^{-1}$, $H_1(\cdot)$ returns the top singular component, and $j \in S$ is picked to minimize λ_j .

Then, for any rank- r solution \mathbf{A}^* , where $r' \geq 256r$, and its singular value decomposition $\mathbf{A}^* = \mathbf{U}^* \mathbf{\Lambda}^* \mathbf{V}^{*\top}$, at least one of the following conditions holds:

1. We have sufficient progress in the regularized function:

$$g(\mathbf{A}') \leq g(\mathbf{A}) - (16\kappa r)^{-1} (g(\mathbf{A}) - f(\mathbf{A}^*)).$$

2. $\mathbf{W} \mathbf{A} \mathbf{A}^\top \mathbf{W}$ is significantly correlated to \mathbf{U}^* :

$$\langle \mathbf{\Pi}_{\operatorname{im}(\mathbf{U}^*)}, \mathbf{W} \mathbf{A} \mathbf{A}^\top \mathbf{W} \rangle \geq (10\kappa)^{-1} \langle \mathbf{W}, \mathbf{A} \mathbf{A}^\top \rangle.$$

3. $\mathbf{Y} \mathbf{A}^\top \mathbf{A} \mathbf{Y}$ is significantly correlated to \mathbf{V}^* :

$$\langle \mathbf{H}_{\text{im}(\mathbf{V}^*)}, \mathbf{Y} \mathbf{A}^\top \mathbf{A} \mathbf{Y} \rangle \geq (10\kappa)^{-1} \langle \mathbf{Y}, \mathbf{A}^\top \mathbf{A} \rangle .$$

4. The spectrum of $\mathbf{A}^\top \mathbf{W} \mathbf{A}$ is highly concentrated and responsible for a constant fraction of the error:

$$(\beta/4) \text{Tr} \left[H_r \left(\mathbf{A}^\top \mathbf{W} \mathbf{A} \right) \right] \geq 10^{-1} (g(\mathbf{A}) - f(\mathbf{A}^*)) .$$

and

5. The spectrum of $\mathbf{A} \mathbf{Y} \mathbf{A}^\top$ is highly concentrated and responsible for a constant fraction of the error:

$$(\beta/4) \text{Tr} \left[H_r \left(\mathbf{A} \mathbf{Y} \mathbf{A}^\top \right) \right] \geq 10^{-1} (g(\mathbf{A}) - f(\mathbf{A}^*)) .$$

Proof. Note that g is a 2β -smooth function. This follows because

$$\nabla g(\mathbf{A}) = \nabla f(\mathbf{A}) + (\beta/2) (\mathbf{W} \mathbf{A} + \mathbf{A} \mathbf{Y}) ,$$

and so for any two matrices \mathbf{A}, \mathbf{A}' ,

$$\begin{aligned} & \|\nabla g(\mathbf{A}') - \nabla g(\mathbf{A})\|_F \\ & \leq \|\nabla f(\mathbf{A}') - \nabla f(\mathbf{A})\|_F + (\beta/2) \|\mathbf{W}(\mathbf{A}' - \mathbf{A})\|_F + (\beta/2) \|(\mathbf{A}' - \mathbf{A}) \mathbf{Y}\|_F \\ & \leq 2\beta \|\mathbf{A}' - \mathbf{A}\|_F , \end{aligned}$$

which is known to imply 2β -smoothness of g . Here we used the triangle inequality and the fact that $\mathbf{W}, \mathbf{Y} \preceq \mathbf{I}$. Therefore, we have

$$\begin{aligned} & g(\mathbf{A}') - g(\mathbf{A}) \\ & \leq \langle \nabla g(\mathbf{A}), \mathbf{A}' - \mathbf{A} \rangle + \|\nabla g(\mathbf{A}') - \nabla g(\mathbf{A})\|_F \|\mathbf{A}' - \mathbf{A}\|_F \\ & \leq \langle \nabla g(\mathbf{A}), \mathbf{A}' - \mathbf{A} \rangle + \beta \|\mathbf{A}' - \mathbf{A}\|_F^2 \\ & \leq -\eta \|\nabla g(\mathbf{A})\|_2^2 + 2\beta\eta^2 \|\nabla g(\mathbf{A})\|_2^2 + 2\beta\lambda_j^2 \\ & = -(8\beta)^{-1} \|\nabla g(\mathbf{A})\|_2^2 + 2\beta\lambda_j^2 , \end{aligned} \tag{9.80}$$

where in the second inequality we used the facts that

$$\begin{aligned} & \langle \nabla g(\mathbf{A}), -\lambda_j \mathbf{u}_j \mathbf{v}_j^\top \rangle \\ & = \langle \mathbf{H}_{\text{im}(\mathbf{U})} \nabla g(\mathbf{A}) \mathbf{H}_{\text{im}(\mathbf{V})}, -\lambda_j \mathbf{u}_j \mathbf{v}_j^\top \rangle \\ & = 0 \end{aligned}$$

and that, for any two matrices \mathbf{B}, \mathbf{C} ,

$$\|\mathbf{B} + \mathbf{C}\|_F^2 \leq 2\|\mathbf{B}\|_F^2 + 2\|\mathbf{C}\|_F^2 .$$

The last equality follows by our choice of η . In order to lower bound $\|\nabla g(\mathbf{A})\|_2^2$, we use the strong

convexity of f as follows:

$$\begin{aligned}
& f(\mathbf{A}^*) - f(\mathbf{A}) \\
& \geq \langle \nabla f(\mathbf{A}), \mathbf{A}^* - \mathbf{A} \rangle + (\alpha/2) \|\mathbf{A}^* - \mathbf{A}\|_F^2 \\
& = \langle \nabla g(\mathbf{A}), \mathbf{A}^* - \mathbf{A} \rangle - \langle \nabla \Phi(\mathbf{A}), \mathbf{A}^* - \mathbf{A} \rangle + (\alpha/2) \|\mathbf{A}^* - \mathbf{A}\|_F^2 \\
& = \underbrace{\langle \nabla g(\mathbf{A}), \mathbf{A}^* - \mathbf{A} \rangle + (\alpha/4) \|\mathbf{A}^* - \mathbf{A}\|_F^2}_P - \langle \nabla \Phi(\mathbf{A}), \mathbf{A}^* - \mathbf{A} \rangle + (\alpha/4) \|\mathbf{A}^* - \mathbf{A}\|_F^2 .
\end{aligned} \tag{9.81}$$

Bounding P . We let $\mathbf{\Pi}_{\text{im}(\mathbf{U})}$, $\mathbf{\Pi}_{\text{im}(\mathbf{V})}$ be the orthogonal projections onto the images of \mathbf{U} and \mathbf{V} respectively, so we can write

$$\begin{aligned}
& \mathbf{A}^* - \mathbf{A} \\
& = \mathbf{\Pi}_{\text{im}(\mathbf{U})} (\mathbf{A}^* - \mathbf{A}) \mathbf{\Pi}_{\text{im}(\mathbf{V})} + (\mathbf{I} - \mathbf{\Pi}_{\text{im}(\mathbf{U})}) (\mathbf{A}^* - \mathbf{A}) \mathbf{\Pi}_{\text{im}(\mathbf{V})} + (\mathbf{A}^* - \mathbf{A}) (\mathbf{I} - \mathbf{\Pi}_{\text{im}(\mathbf{V})}) \\
& = \mathbf{\Pi}_{\text{im}(\mathbf{U})} (\mathbf{A}^* - \mathbf{A}) \mathbf{\Pi}_{\text{im}(\mathbf{V})} + (\mathbf{I} - \mathbf{\Pi}_{\text{im}(\mathbf{U})}) \mathbf{A}^* \mathbf{\Pi}_{\text{im}(\mathbf{V})} + \mathbf{A}^* (\mathbf{I} - \mathbf{\Pi}_{\text{im}(\mathbf{V})}) .
\end{aligned}$$

Now, note that

$$\begin{aligned}
& \langle \nabla g(\mathbf{A}), \mathbf{A}^* - \mathbf{A} \rangle \\
& = \langle \nabla g(\mathbf{A}), (\mathbf{I} - \mathbf{\Pi}_{\text{im}(\mathbf{U})}) \mathbf{A}^* \mathbf{\Pi}_{\text{im}(\mathbf{V})} \rangle + \langle \nabla g(\mathbf{A}), \mathbf{A}^* (\mathbf{I} - \mathbf{\Pi}_{\text{im}(\mathbf{V})}) \rangle ,
\end{aligned}$$

where we used the fact that

$$\begin{aligned}
& \langle \nabla g(\mathbf{A}), \mathbf{\Pi}_{\text{im}(\mathbf{U})} (\mathbf{A}^* - \mathbf{A}) \mathbf{\Pi}_{\text{im}(\mathbf{V})} \rangle \\
& = \langle \mathbf{\Pi}_{\text{im}(\mathbf{U})} \nabla g(\mathbf{A}) \mathbf{\Pi}_{\text{im}(\mathbf{V})}, \mathbf{A}^* - \mathbf{A} \rangle \\
& = 0 ,
\end{aligned}$$

and

$$\begin{aligned}
& (\alpha/4) \|\mathbf{A}^* - \mathbf{A}\|_F^2 \\
& \geq (\alpha/4) \|(\mathbf{I} - \mathbf{\Pi}_{\text{im}(\mathbf{U})}) \mathbf{A}^* \mathbf{\Pi}_{\text{im}(\mathbf{V})}\|_F^2 + (\alpha/4) \|\mathbf{A}^* (\mathbf{I} - \mathbf{\Pi}_{\text{im}(\mathbf{V})})\|_F^2 .
\end{aligned}$$

Additionally, note that for any rank- r matrix \mathbf{B} , we have

$$\begin{aligned}
& \langle \nabla g(\mathbf{A}), \mathbf{B} \rangle + (\alpha/4) \|\mathbf{B}\|_F^2 \\
& \geq -\alpha^{-1} \|H_r(\nabla g(\mathbf{A}))\|_F^2 \\
& \geq -\alpha^{-1} r \|\nabla g(\mathbf{A})\|_2^2 ,
\end{aligned}$$

a proof of which can be found e.g. in Lemma A.6 of [13]. Applying this inequality with

$$\mathbf{B} = (\mathbf{I} - \mathbf{\Pi}_{\text{im}(\mathbf{U})}) \mathbf{A}^* \mathbf{\Pi}_{\text{im}(\mathbf{V})}$$

and

$$\mathbf{B} = \mathbf{A}^* (\mathbf{I} - \mathbf{\Pi}_{\text{im}(\mathbf{V})})$$

and summing them up, we obtain

$$\begin{aligned} P &= \langle \nabla g(\mathbf{A}), \mathbf{A}^* - \mathbf{A} \rangle + (\alpha/4) \|\mathbf{A}^* - \mathbf{A}\|_F^2 \\ &\geq -2\alpha^{-1}r \|\nabla g(\mathbf{A})\|_2^2. \end{aligned}$$

Plugging this into (9.81) and re-arranging, we get

$$\begin{aligned} &\|\nabla g(\mathbf{A})\|_2^2 \\ &\geq \alpha/(2r) \left(f(\mathbf{A}) - f(\mathbf{A}^*) - \langle \nabla \Phi(\mathbf{A}), \mathbf{A}^* - \mathbf{A} \rangle + (\alpha/4) \|\mathbf{A}^* - \mathbf{A}\|_F^2 \right) \\ &= \alpha/(2r) \left(g(\mathbf{A}) - f(\mathbf{A}^*) - \underbrace{\Phi(\mathbf{A}) - \langle \nabla \Phi(\mathbf{A}), \mathbf{A}^* - \mathbf{A} \rangle + (\alpha/4) \|\mathbf{A}^* - \mathbf{A}\|_F^2}_Q \right). \end{aligned} \tag{9.82}$$

Bounding Q . We know that

$$-\langle \nabla \Phi(\mathbf{A}), \mathbf{A}^* - \mathbf{A} \rangle = -(\beta/2) \langle \mathbf{W}\mathbf{A} + \mathbf{A}\mathbf{Y}, \mathbf{A}^* - \mathbf{A} \rangle.$$

If we let

$$\mathbf{A}^* = \mathbf{U}^* \mathbf{\Lambda}^* \mathbf{V}^{*\top}$$

be the SVD of \mathbf{A}^* and $\mathbf{\Pi}_{\text{im}(\mathbf{U}^*)}$, $\mathbf{\Pi}_{\text{im}(\mathbf{V}^*)}$ be the orthogonal projections onto the images of \mathbf{U}^* and \mathbf{V}^* respectively, then we have

$$\begin{aligned} &-(\beta/2) \langle \mathbf{W}\mathbf{A}, \mathbf{A}^* - \mathbf{A} \rangle \\ &= -(\beta/2) \langle \mathbf{W}\mathbf{A}, \mathbf{\Pi}_{\text{im}(\mathbf{U}^*)}(\mathbf{A}^* - \mathbf{A})\mathbf{\Pi}_{\text{im}(\mathbf{V}^*)} \rangle + (\beta/2) \langle \mathbf{W}, \mathbf{A}\mathbf{A}^\top \rangle - (\beta/2) \langle \mathbf{W}\mathbf{A}, \mathbf{\Pi}_{\text{im}(\mathbf{U}^*)}\mathbf{A}\mathbf{\Pi}_{\text{im}(\mathbf{V}^*)} \rangle. \end{aligned}$$

Looking at the first term of this, we have

$$\begin{aligned} &-(\beta/2) \langle \mathbf{W}\mathbf{A}, \mathbf{\Pi}_{\text{im}(\mathbf{U}^*)}(\mathbf{A}^* - \mathbf{A})\mathbf{\Pi}_{\text{im}(\mathbf{V}^*)} \rangle + (\alpha/8) \|\mathbf{A}^* - \mathbf{A}\|_F^2 \\ &= -(\beta/2) \langle \mathbf{\Pi}_{\text{im}(\mathbf{U}^*)}\mathbf{W}\mathbf{A}\mathbf{\Pi}_{\text{im}(\mathbf{V}^*)}, \mathbf{A}^* - \mathbf{A} \rangle + (\alpha/8) \|\mathbf{A}^* - \mathbf{A}\|_F^2 \\ &\geq -\beta^2/(2\alpha) \|\mathbf{\Pi}_{\text{im}(\mathbf{U}^*)}\mathbf{W}\mathbf{A}\mathbf{\Pi}_{\text{im}(\mathbf{V}^*)}\|_F^2. \end{aligned}$$

Similarly for the terms containing \mathbf{Y} , we get

$$\begin{aligned} &-(\beta/2) \langle \mathbf{A}\mathbf{Y}, \mathbf{A}^* - \mathbf{A} \rangle \\ &= -(\beta/2) \langle \mathbf{A}\mathbf{Y}, \mathbf{\Pi}_{\text{im}(\mathbf{U}^*)}(\mathbf{A}^* - \mathbf{A})\mathbf{\Pi}_{\text{im}(\mathbf{V}^*)} \rangle + (\beta/2) \langle \mathbf{Y}, \mathbf{A}^\top \mathbf{A} \rangle - (\beta/2) \langle \mathbf{A}\mathbf{Y}, \mathbf{\Pi}_{\text{im}(\mathbf{U}^*)}\mathbf{A}\mathbf{\Pi}_{\text{im}(\mathbf{V}^*)} \rangle. \end{aligned}$$

and

$$\begin{aligned} &-(\beta/2) \langle \mathbf{A}\mathbf{Y}, \mathbf{\Pi}_{\text{im}(\mathbf{U}^*)}(\mathbf{A}^* - \mathbf{A})\mathbf{\Pi}_{\text{im}(\mathbf{V}^*)} \rangle + (\alpha/8) \|\mathbf{A}^* - \mathbf{A}\|_F^2 \\ &\geq -\beta^2/(2\alpha) \|\mathbf{\Pi}_{\text{im}(\mathbf{U}^*)}\mathbf{A}\mathbf{Y}\mathbf{\Pi}_{\text{im}(\mathbf{V}^*)}\|_F^2. \end{aligned}$$

In summary, we have

$$\begin{aligned}
Q &= -\langle \nabla \Phi(\mathbf{A}), \mathbf{A}^* - \mathbf{A} \rangle + (\alpha/4) \|\mathbf{A}^* - \mathbf{A}\|_F^2 \\
&\geq -\beta^2/(2\alpha) \|\mathbf{\Pi}_{\text{im}(\mathbf{U}^*)} \mathbf{W} \mathbf{A} \mathbf{\Pi}_{\text{im}(\mathbf{V}^*)}\|_F^2 - \beta^2/(2\alpha) \|\mathbf{\Pi}_{\text{im}(\mathbf{U}^*)} \mathbf{A} \mathbf{Y} \mathbf{\Pi}_{\text{im}(\mathbf{V}^*)}\|_F^2 \\
&\quad + (\beta/2) \langle \mathbf{W}, \mathbf{A} \mathbf{A}^\top \rangle + (\beta/2) \langle \mathbf{Y}, \mathbf{A}^\top \mathbf{A} \rangle \\
&\quad - (\beta/2) \langle \mathbf{W} \mathbf{A}, \mathbf{\Pi}_{\text{im}(\mathbf{U}^*)} \mathbf{A} \mathbf{\Pi}_{\text{im}(\mathbf{V}^*)} \rangle - (\beta/2) \langle \mathbf{A} \mathbf{Y}, \mathbf{\Pi}_{\text{im}(\mathbf{U}^*)} \mathbf{A} \mathbf{\Pi}_{\text{im}(\mathbf{V}^*)} \rangle.
\end{aligned} \tag{9.83}$$

Now, let us assume that all items 2-5 from the lemma statement are false. For the first term of (9.83), we have

$$\begin{aligned}
& -\beta^2/(2\alpha) \|\mathbf{\Pi}_{\text{im}(\mathbf{U}^*)} \mathbf{W} \mathbf{A} \mathbf{\Pi}_{\text{im}(\mathbf{V}^*)}\|_F^2 \\
& \geq -\beta^2/(2\alpha) \|\mathbf{\Pi}_{\text{im}(\mathbf{U}^*)} \mathbf{W} \mathbf{A}\|_F^2 \\
& = -\beta^2/(2\alpha) \langle \mathbf{\Pi}_{\text{im}(\mathbf{U}^*)}, \mathbf{W} \mathbf{A} \mathbf{A}^\top \mathbf{W} \rangle \\
& \geq -(\beta/20) \langle \mathbf{W}, \mathbf{A} \mathbf{A}^\top \rangle,
\end{aligned}$$

where we used item 2 from the lemma statement, and similarly for the second term of (9.83),

$$\begin{aligned}
& -\beta^2/(2\alpha) \|\mathbf{\Pi}_{\text{im}(\mathbf{U}^*)} \mathbf{A} \mathbf{Y} \mathbf{\Pi}_{\text{im}(\mathbf{V}^*)}\|_F^2 \\
& \geq -(\beta/20) \langle \mathbf{Y}, \mathbf{A}^\top \mathbf{A} \rangle.
\end{aligned}$$

Now we look at the second to last term of (9.83), i.e.

$$\begin{aligned}
& -(\beta/2) \langle \mathbf{W} \mathbf{A}, \mathbf{\Pi}_{\text{im}(\mathbf{U}^*)} \mathbf{A} \mathbf{\Pi}_{\text{im}(\mathbf{V}^*)} \rangle \\
& = -(\beta/2) \langle \mathbf{W} \mathbf{A} \mathbf{\Pi}_{\text{im}(\mathbf{V}^*)} \mathbf{A}^\top, \mathbf{\Pi}_{\text{im}(\mathbf{U}^*)} \rangle.
\end{aligned}$$

Now, we use the matrix Holder inequality

$$\begin{aligned}
& -(\beta/2) \langle \mathbf{W} \mathbf{A} \mathbf{\Pi}_{\text{im}(\mathbf{V}^*)} \mathbf{A}^\top, \mathbf{\Pi}_{\text{im}(\mathbf{U}^*)} \rangle \\
& \geq -(\beta/2) \left\| \mathbf{W} \mathbf{A} \mathbf{\Pi}_{\text{im}(\mathbf{V}^*)} \mathbf{A}^\top \right\|_* \|\mathbf{\Pi}_{\text{im}(\mathbf{U}^*)}\|_2 \\
& \geq -(\beta/2) \left\| \mathbf{W} \mathbf{A} \mathbf{\Pi}_{\text{im}(\mathbf{V}^*)} \mathbf{A}^\top \right\|_*,
\end{aligned}$$

which can be proved by applying von Neumann's trace inequality and then the classical Holder inequality. Now, note that the matrix $\mathbf{W} \mathbf{A} \mathbf{\Pi}_{\text{im}(\mathbf{V}^*)} \mathbf{A}^\top$ is similar to $\mathbf{W}^{1/2} \mathbf{A} \mathbf{\Pi}_{\text{im}(\mathbf{V}^*)} \mathbf{A}^\top \mathbf{W}^{1/2}$, and so they have the same eigenvalues. Furthermore, the latter is a symmetric PSD matrix, and so the former has real positive eigenvalues as well. This means that its singular values are the same as its eigenvalues, and as a result the nuclear norm is equal to the trace, i.e.

$$\begin{aligned}
& -(\beta/2) \left\| \mathbf{W} \mathbf{A} \mathbf{\Pi}_{\text{im}(\mathbf{V}^*)} \mathbf{A}^\top \right\|_* \\
& = -(\beta/2) \text{Tr} \left(\mathbf{W} \mathbf{A} \mathbf{\Pi}_{\text{im}(\mathbf{V}^*)} \mathbf{A}^\top \right) \\
& = -(\beta/2) \langle \mathbf{\Pi}_{\text{im}(\mathbf{V}^*)}, \mathbf{A}^\top \mathbf{W} \mathbf{A} \rangle \\
& \geq -(\beta/2) \text{Tr} \left[H_r \left(\mathbf{A}^\top \mathbf{W} \mathbf{A} \right) \right] \\
& \geq -(1/5) (g(\mathbf{A}) - f(\mathbf{A}^*)).
\end{aligned}$$

where we also used Lemma 9.5.2 and item 4 from the lemma statement. So we derived that

$$\begin{aligned} & -(\beta/2)\langle \mathbf{W}\mathbf{A}, \mathbf{\Pi}_{\text{im}(\mathbf{U}^*)}\mathbf{A}\mathbf{\Pi}_{\text{im}(\mathbf{V}^*)}\rangle \\ & \geq -(1/5)(g(\mathbf{A}) - f(\mathbf{A}^*)) , \end{aligned}$$

and similarly for the last term of (9.83),

$$\begin{aligned} & -(\beta/2)\langle \mathbf{A}\mathbf{Y}, \mathbf{\Pi}_{\text{im}(\mathbf{U}^*)}\mathbf{A}\mathbf{\Pi}_{\text{im}(\mathbf{V}^*)}\rangle \\ & \geq -(1/5)(g(\mathbf{A}) - f(\mathbf{A}^*)) . \end{aligned}$$

Plugging the four inequalities that we derived back into (9.83), we get

$$\begin{aligned} Q & \geq (\beta/2 - \beta/20)\langle \mathbf{W}, \mathbf{A}\mathbf{A}^\top \rangle + (\beta/2 - \beta/20)\langle \mathbf{Y}, \mathbf{A}^\top \mathbf{A} \rangle - (2/5)(g(\mathbf{A}) - f(\mathbf{A}^*)) \\ & = (9/5)\Phi(\mathbf{A}) - (2/5)(g(\mathbf{A}) - f(\mathbf{A}^*)) \\ & > (3/2)\Phi(\mathbf{A}) - (2/5)(g(\mathbf{A}) - f(\mathbf{A}^*)) . \end{aligned}$$

Finally, combining this with the smoothness inequality (9.80) and the lower bound on $\|\nabla g(\mathbf{A})\|_2^2$ (9.82), we derive

$$\begin{aligned} & g(\mathbf{A}') - g(\mathbf{A}) \\ & \leq -(16\kappa r)^{-1}\left(g(\mathbf{A}) - f(\mathbf{A}^*) + (1/2)\Phi(\mathbf{A})\right) + 2\beta\lambda_j^2 \\ & = -(16\kappa r)^{-1}\left(g(\mathbf{A}) - f(\mathbf{A}^*)\right) - (32\kappa r)^{-1}\Phi(\mathbf{A}) + 2\beta\lambda_j^2 . \end{aligned}$$

What remains is the bound the sum of the last two terms. We remind the reader that $\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{V}^\top$. Now, letting \mathbf{z} equal to the vectorized diagonal of $\mathbf{U}^\top \mathbf{W} \mathbf{U}$ and $\boldsymbol{\lambda}$ to the vectorized diagonal of $\mathbf{\Lambda}$, note that

$$\|\boldsymbol{\lambda}\|_z^2 = \langle \mathbf{\Lambda}^2, \mathbf{U}^\top \mathbf{W} \mathbf{U} \rangle = \langle \mathbf{W}, \mathbf{A}\mathbf{A}^\top \rangle ,$$

using which we derive

$$\begin{aligned} \lambda_j^2 & = \min_{j \in S} \lambda_j^2 \leq \frac{\|\boldsymbol{\lambda}\|_z^2}{\|\mathbf{z}\|_1} \\ & = \frac{\langle \mathbf{W}, \mathbf{A}\mathbf{A}^\top \rangle}{\text{Tr}[\mathbf{U}^\top \mathbf{W} \mathbf{U}]} \\ & = \frac{\langle \mathbf{W}, \mathbf{A}\mathbf{A}^\top \rangle}{\text{Tr}[\mathbf{U}^\top \mathbf{U}] - \text{Tr}[\mathbf{U}^\top (\mathbf{I} - \mathbf{W}) \mathbf{U}]} \\ & \leq \frac{\langle \mathbf{W}, \mathbf{A}\mathbf{A}^\top \rangle}{r' - \text{Tr}[\mathbf{I} - \mathbf{W}]} \\ & \leq \frac{\langle \mathbf{W}, \mathbf{A}\mathbf{A}^\top \rangle}{r'/2} \\ & \leq \frac{\langle \mathbf{W}, \mathbf{A}\mathbf{A}^\top \rangle}{128r\kappa} , \end{aligned}$$

where we used the fact that

$$\begin{aligned} & \text{Tr}[\mathbf{U}^\top(\mathbf{I} - \mathbf{W})\mathbf{U}] \\ &= \text{Tr}\left[(\mathbf{I} - \mathbf{W})^{1/2}\mathbf{U}\mathbf{U}^\top(\mathbf{I} - \mathbf{W})^{1/2}\right] \\ &\leq \text{Tr}[\mathbf{I} - \mathbf{W}], \end{aligned}$$

because the columns of \mathbf{U} are orthonormal. We also used the property that $\text{Tr}[\mathbf{I} - \mathbf{W}] \leq r'/2$ and the fact that $r' \geq 256r\kappa$ by the lemma statement.

Similarly, we derive that

$$\lambda_j^2 \leq \frac{\langle \mathbf{Y}, \mathbf{A}^\top \mathbf{A} \rangle}{128r\kappa},$$

and, adding these two inequalities, we have

$$2\beta\lambda_j^2 \leq (32r\kappa)^{-1}\Phi(\mathbf{A}),$$

finally concluding that

$$g(\mathbf{A}') - g(\mathbf{A}) \leq -(16\kappa r)^{-1}\left(g(\mathbf{A}) - f(\mathbf{A}^*)\right).$$

□

9.5.5 Lower Bounds

Lemma 9.5.5 (IHT lower bound). *Let $f(\mathbf{x}) := (1/2)\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2$. For any $\kappa, s \geq 1$, $s' \leq 0.6s\kappa^2$, there exists a (diagonal) matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ and a vector $\mathbf{b} \in \mathbb{R}^n$ where $n = s(\kappa^2 + \kappa + 1)$, f is 1-strongly convex and κ -smooth, as well as an s -sparse solution \mathbf{x}^* and an s' -sparse solution \mathbf{x} , such that*

$$f(\mathbf{x}) \geq f(\mathbf{x}^*) + 0.1s\kappa^2$$

but

$$\mathbf{x} = H_{s'}(\mathbf{x} - \beta^{-1}\nabla f(\mathbf{x})),$$

i.e. \mathbf{x} is a fixpoint for IHT.

Proof. We use the same example as in [14], Section 5.2: \mathbf{A} is diagonal with

$$\mathbf{A}_{ii} = \begin{cases} 1 & \text{if } i \in I_1 \\ \sqrt{\kappa} & \text{if } i \in I_2 \\ 1 & \text{if } i \in I_3, \end{cases}$$

where $I_1 = [s]$, $I_2 = [s + 1, s(\kappa + 1)]$, $I_3 = [s(\kappa + 1) + 1, s(\kappa^2 + \kappa + 1)]$, and \mathbf{b} is defined as

$$b_i = \begin{cases} \kappa\sqrt{1 - 4\delta} & \text{if } i \in I_1 \\ \sqrt{\kappa}\sqrt{1 - 2\delta} & \text{if } i \in I_2 \\ 1 & \text{if } i \in I_3, \end{cases}$$

for some sufficiently small $\delta > 0$ used for tie-breaking. We define

$$x_i^* = \begin{cases} \kappa\sqrt{1-4\delta} & \text{if } i \in I_1 \\ 0 & \text{otherwise} \end{cases}$$

and, for some arbitrary s' -sized $S \subseteq I_3$

$$x_i = \begin{cases} 0 & \text{if } i \in I_1 \cup I_2 \cup I_3 \setminus S \\ 1 & \text{otherwise.} \end{cases}$$

Note that $f(\mathbf{x}) - f(\mathbf{x}^*) = 0.5s\kappa^2(1-4\delta) - 0.5s' \geq 0.1s\kappa^2$. Furthermore, the gradient is equal to

$$\begin{aligned} \nabla f(\mathbf{x}) &= \mathbf{A}^\top (\mathbf{A}\mathbf{x} - \mathbf{b}) \\ &= \begin{cases} -\kappa\sqrt{1-4\delta} & \text{if } i \in I_1 \\ -\kappa\sqrt{1-2\delta} & \text{if } i \in I_2 \\ -1 & \text{if } i \in I_3 \setminus S \\ 0 & \text{if } i \in S, \end{cases} \end{aligned}$$

and since we have $\beta = \kappa$,

$$\mathbf{x} - \beta^{-1}\nabla f(\mathbf{x}) = \begin{cases} \sqrt{1-4\delta} & \text{if } i \in I_1 \\ \sqrt{1-2\delta} & \text{if } i \in I_2 \\ 1/\kappa & \text{if } i \in I_3 \setminus S \\ 1 & \text{if } i \in S, \end{cases}$$

implying that $H_{s'}(\mathbf{x} - \beta^{-1}\nabla f(\mathbf{x})) = \mathbf{x}$. □

9.6 Appendix for Chapter 8

Preliminaries and Notation

Given an positive integer k , we denote $[k] = \{1, 2, \dots, k\}$. Given a matrix A , we denote by $\|A\|_F$ its Frobenius norm, i.e. the ℓ_2 norm of the entries of A (or equivalently of the singular values of A). The following lemma is a simple corollary of the definition of the Frobenius norm:

Lemma 9.6.1. *Given two matrices $A \in \mathbb{R}^{m \times n}, B \in \mathbb{R}^{m \times n}$, we have $\|A + B\|_F^2 \leq 2(\|A\|_F^2 + \|B\|_F^2)$.*

Proof.

$$\|A + B\|_F^2 = \sum_{ij} (A + B)_{ij}^2 \leq 2 \sum_{ij} (A_{ij}^2 + B_{ij}^2) = 2(\|A\|_F^2 + \|B\|_F^2)$$

□

Definition 9.6.2 (Rank-restricted smoothness, strong convexity, condition number). *Given a convex function $R \in \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ and an integer parameter r , the rank-restricted smoothness of R at rank r is the minimum constant $\rho_r^+ \geq 0$ such that for any two matrices $A \in \mathbb{R}^{m \times n}, B \in \mathbb{R}^{m \times n}$ such that*

$\text{rank}(A - B) \leq r$, we have

$$R(B) \leq R(A) + \langle \nabla R(A), B - A \rangle + \frac{\rho_r^+}{2} \|B - A\|_F^2.$$

Similarly, the rank-restricted strong convexity of R at rank r is the maximum constant $\rho_r^- \geq 0$ such that for any two matrices $A \in \mathbb{R}^{m \times n}$, $B \in \mathbb{R}^{m \times n}$ such that $\text{rank}(A - B) \leq r$, we have

$$R(B) \geq R(A) + \langle \nabla R(A), B - A \rangle + \frac{\rho_r^-}{2} \|B - A\|_F^2.$$

Given that ρ_r^+, ρ_r^- exist and are nonzero, the rank-restricted condition number of R at rank r is then defined as

$$\kappa_r = \frac{\rho_r^+}{\rho_r^-}$$

Note that ρ_r^+ is increasing and ρ_r^- is decreasing in r . Therefore, even though our bounds are proven in terms of the constants $\frac{\rho_1^+}{\rho_r^+}$ and $\frac{\rho_2^-}{\rho_r^-}$, these quantities are always at most $\frac{\rho_r^+}{\rho_r^-} = \kappa_r$ as long as $r \geq 2$, and so they directly imply the same bounds in terms of the constant κ_r .

Definition 9.6.3 (Spectral norm). Given a matrix $A \in \mathbb{R}^{m \times n}$, we denote its spectral norm by $\|A\|_2$. The spectral norm is defined as

$$\|A\|_2 = \max_{x \in \mathbb{R}^n} \frac{\|Ax\|_2}{\|x\|_2},$$

Definition 9.6.4 (Singular value thresholding operator). Given a matrix $A \in \mathbb{R}^{m \times n}$ of rank k , a singular value decomposition $A = U\Sigma V^\top$ such that $\Sigma_{11} \geq \Sigma_{22} \geq \dots \geq \Sigma_{kk}$, and an integer $1 \leq r \leq k$, we define $H_r(A) = U\Sigma'V^\top$, here Σ' is a diagonal matrix with

$$\Sigma'_{ii} = \begin{cases} \Sigma_{ii} & \text{if } i \leq r \\ 0 & \text{otherwise} \end{cases}$$

In other words, $H_r(\cdot)$ is an operator that eliminates all but the top r highest singular values of a matrix.

Lemma 9.6.5 (Weyl's inequality). For any matrix A and integer $i \geq 1$, let $\sigma_i(A)$ be the i -th largest singular value of A or 0 if $i > \text{rank}(A)$. Then, for any two matrices A, B and integers $i \geq 1, j \geq 1$:

$$\sigma_{i+j-1}(A + B) \leq \sigma_i(A) + \sigma_j(B)$$

A proof of the previous fact can be found e.g. in [59].

Lemma 9.6.6 ($H_r(\cdot)$ optimization problem). Let $A \in \mathbb{R}^{m \times n}$ be a rank- k matrix and $r \in [k]$ be an integer parameter. Then $M = \frac{1}{\lambda} H_r(A)$ is an optimal solution to the following optimization problem:

$$\max_{\text{rank}(M) \leq r} \{ \langle A, M \rangle - \frac{\lambda}{2} \|M\|_F^2 \} \quad (9.84)$$

Proof. Let $U\Sigma V^\top = \sum_i \Sigma_{ii} U_i V_i^\top$ be a singular value decomposition of A . We note that (9.84) is equivalent to

$$\min_{\text{rank}(M) \leq r} \|A - \lambda M\|_F^2 := f(M) \quad (9.85)$$

Now, note that $f(\frac{1}{\lambda} H_r(A)) = \|A - H_r(A)\|_F^2 = \sum_{i=r+1}^k \Sigma_{ii}^2$. On the other hand, by applying Weyl's inequality (Lemma 9.6.5) for $j = r + 1$,

$$f(M) = \|A - \lambda M\|_F^2 = \sum_{i=1}^{k+r} \sigma_i^2(A - \lambda M) \geq \sum_{i=1}^{k+r} (\sigma_{i+r}(A) - \sigma_{r+1}(\lambda M))^2 = \sum_{i=r+1}^k \Sigma_{ii}^2,$$

where the last equality follows from the fact that $\text{rank}(A) = k$ and $\text{rank}(M) \leq r$. Therefore, $M = \frac{1}{\lambda} H_r(A)$ minimizes (9.85) and thus maximizes (9.84). \square

Proof of Theorem 8.2.1 (greedy)

We will start with the following simple lemma about the Frobenius norm of a sum of matrices with orthogonal columns or rows:

Lemma 9.6.7. *Let $U \in \mathbb{R}^{m \times r}$, $V \in \mathbb{R}^{n \times r}$, $X \in \mathbb{R}^{m \times r}$, $Y \in \mathbb{R}^{n \times r}$ be such that the columns of U are orthogonal to the columns of X or the columns of V are orthogonal to the columns of Y . Then $\|UV^\top + XY^\top\|_F^2 = \|UV^\top\|_F^2 + \|XY^\top\|_F^2$.*

Proof. If the columns of U are orthogonal to those of X , then $U^\top X = \mathbf{0}$ and if the columns of V are orthogonal to those of Y , then $Y^\top V = \mathbf{0}$. Therefore in any case $\langle UV^\top, XY^\top \rangle = \text{Tr}(VU^\top XY^\top) = \text{Tr}(U^\top XY^\top V) = \mathbf{0}$, implying

$$\|UV^\top + XY^\top\|_F^2 = \|UV^\top\|_F^2 + \|XY^\top\|_F^2 + 2\langle UV^\top, XY^\top \rangle = \|UV^\top\|_F^2 + \|XY^\top\|_F^2$$

\square

Additionally, we have the following lemma regarding the optimality conditions of (8.2):

Lemma 9.6.8. *Let $A = UXV^\top$ where $U \in \mathbb{R}^{m \times r}$, $X \in \mathbb{R}^{r \times r}$, and $V \in \mathbb{R}^{n \times r}$, such that X is the optimal solution to (8.2). Then for any $u \in \text{im}(U)$ and $v \in \text{im}(V)$ we have that $\langle \nabla R(A), uv^\top \rangle = 0$.*

Proof. By the optimality condition of 8.2, we have that

$$U^\top \nabla R(A) V = \mathbf{0}$$

Now, for any $u = Ux$ and $v = Vy$ we have

$$\langle \nabla R(A), uv^\top \rangle = u^\top \nabla R(A) v = x^\top U^\top \nabla R(A) V y = 0$$

\square

We are now ready for the proof of Theorem 8.2.1.

Proof. Let A_{t-1} be the current solution UV^\top before iteration $t-1 \geq 0$. Let $u \in \mathbb{R}^m$ and $v \in \mathbb{R}^m$ be left and right singular vectors of matrix $\nabla R(A)$, i.e. unit vectors maximizing $|\langle \nabla R(A), uv^\top \rangle|$. Let

$$\mathcal{B}_t = \{B \mid B = A_{t-1} + \eta uv^\top, \eta \in \mathbb{R}\}.$$

By smoothness we have

$$\begin{aligned} R(A_{t-1}) - R(A_t) &\geq \max_{B \in \mathcal{B}_t} \{R(A_{t-1}) - R(B)\} \\ &\geq \max_{B \in \mathcal{B}_t} \left\{ -\langle \nabla R(A_{t-1}), B - A_{t-1} \rangle - \frac{\rho_1^+}{2} \|B - A_{t-1}\|_F^2 \right\} \\ &\geq \max_{\eta} \left\{ \eta \langle \nabla R(A_{t-1}), uv^\top \rangle - \eta^2 \frac{\rho_1^+}{2} \right\} \\ &= \max_{\eta} \left\{ \eta \|\nabla R(A_{t-1})\|_2 - \eta^2 \frac{\rho_1^+}{2} \right\} \\ &= \frac{\|\nabla R(A_{t-1})\|_2^2}{2\rho_1^+} \end{aligned}$$

where $\|\cdot\|_2$ is the spectral norm (i.e. maximum magnitude of a singular value).

On the other hand, by strong convexity and noting that

$$\text{rank}(A^* - A_{t-1}) \leq \text{rank}(A^*) + \text{rank}(A_{t-1}) \leq r^* + r,$$

$$R(A^*) - R(A_{t-1}) \geq \langle \nabla R(A_{t-1}), A^* - A_{t-1} \rangle + \frac{\rho_{r+r^*}^-}{2} \|A^* - A_{t-1}\|_F^2. \quad (9.86)$$

Let $A_{t-1} = UV^\top$ and $A^* = U^*V^{*\top}$. We let $\Pi_{\text{im}(U)} = U(U^\top U)^+U^\top$ and $\Pi_{\text{im}(V)} = V(V^\top V)^+V^\top$ denote the orthogonal projections onto the images of U and V respectively. We now write

$$A^* = U^*V^{*\top} = (U^1 + U^2)(V^1 + V^2)^\top = U^1V^{1\top} + U^1V^{2\top} + U^2V^{*\top}$$

where $U^1 = \Pi_{\text{im}(U)}U^*$ is a matrix where every column of U^* is replaced by its projection on $\text{im}(U)$ and $U^2 = U^* - U^1$ and similarly $V^1 = \Pi_{\text{im}(V)}V^*$ is a matrix where every column of V^* is replaced by its projection on $\text{im}(V)$ and $V^2 = V^* - V^1$. By setting $U' = (-U \mid U^1)$ and $V' = (V \mid V^1)$ we can write

$$A^* - A_{t-1} = U'V'^\top + U^1V^{2\top} + U^2V^{*\top}$$

where $\text{im}(U') = \text{im}(U)$ and $\text{im}(V') = \text{im}(V)$. Also, note that

$$\text{rank}(U^1V^{2\top}) \leq \text{rank}(V^2) \leq \text{rank}(V^*) = \text{rank}(A^*) \leq r^*$$

and similarly $\text{rank}(U^2V^{*\top}) \leq r^*$. So now the right hand side of (9.86) can be reshaped as

$$\begin{aligned} &\langle \nabla R(A_{t-1}), A^* - A_{t-1} \rangle + \frac{\rho_{r+r^*}^-}{2} \|A^* - A_{t-1}\|_F^2 \\ &= \langle \nabla R(A_{t-1}), U'V'^\top + U^1V^{2\top} + U^2V^{*\top} \rangle + \frac{\rho_{r+r^*}^-}{2} \|U'V'^\top + U^1V^{2\top} + U^2V^{*\top}\|_F^2 \end{aligned}$$

Now, note that since by definition the columns of U' are in $\text{im}(U)$ and the columns of V' are in

$\text{im}(V)$, Lemma 9.6.8 implies that $\langle \nabla R(A_{t-1}), U'V'^\top \rangle = 0$. Therefore the above is equal to

$$\begin{aligned}
& \langle \nabla R(A_{t-1}), U^1V^{2\top} + U^2V^{*\top} \rangle + \frac{\bar{\rho}_{r+r^*}}{2} \|U'V'^\top + U^1V^{2\top} + U^2V^{*\top}\|_F^2 \\
& \geq \langle \nabla R(A_{t-1}), U^1V^{2\top} \rangle + \langle \nabla R(A_{t-1}), U^2V^{*\top} \rangle + \frac{\bar{\rho}_{r+r^*}}{2} \left(\|U^1V^{2\top}\|_F^2 + \|U^2V^{*\top}\|_F^2 \right) \\
& \geq 2 \min_{\text{rank}(M) \leq r^*} \left\{ \langle \nabla R(A_{t-1}), M \rangle + \frac{\bar{\rho}_{r+r^*}}{2} \|M\|_F^2 \right\} \\
& = -2 \frac{\|H_{r^*}(\nabla R(A_{t-1}))\|_F^2}{2\bar{\rho}_{r+r^*}} \\
& \geq -r^* \frac{\|\nabla R(A_{t-1})\|_2^2}{\bar{\rho}_{r+r^*}}
\end{aligned}$$

where the first equality follows by noticing that the columns of V' and V^1 are orthogonal to those of V^2 and the columns of U' and U^1 are orthogonal to those of U^2 , and applying Lemma 9.6.7. The last equality is a direct application of Lemma 9.6.6 and the last inequality states that the largest squared singular value is not smaller than the average of the top r^* squared singular values. Therefore we have concluded that

$$\|\nabla R(A_{t-1})\|_2^2 \geq \frac{\bar{\rho}_{r+r^*}}{r^*} (R(A_{t-1}) - R(A^*))$$

Plugging this back into the smoothness inequality, we get

$$R(A_{t-1}) - R(A_t) \geq \frac{1}{2r^*\kappa} (R(A_{t-1}) - R(A^*))$$

or equivalently

$$R(A_t) - R(A^*) \leq \left(1 - \frac{1}{2r^*\kappa}\right) (R(A_{t-1}) - R(A^*)).$$

Therefore after $L = 2r^*\kappa \log \frac{R(A_0) - R(A^*)}{\varepsilon}$ iterations we have

$$\begin{aligned}
R(A_T) - R(A^*) & \leq \left(1 - \frac{1}{2r^*\kappa}\right)^L (R(A_0) - R(A^*)) \\
& \leq e^{-\frac{L}{2r^*\kappa}} (R(A_0) - R(A^*)) \\
& \leq \varepsilon
\end{aligned}$$

Since $A_0 = \mathbf{0}$, the result follows. \square

Proof of Theorem 8.2.2 (local search)

Proof. Similarly to Section 9.6, we let A_{t-1} be the current solution before iteration $t - 1 \geq 0$. Let $u \in \mathbb{R}^m$ and $v \in \mathbb{R}^m$ be left and right singular vectors of matrix $\nabla R(A)$, i.e. unit vectors maximizing $|\langle \nabla R(A), uv^\top \rangle|$ and let

$$\mathcal{B}_t = \{B | B = A_{t-1} + \eta uv^\top - \sigma_{\min} xy^\top, \eta \in \mathbb{R}\},$$

where $\sigma_{\min}xy^\top = A_{t-1} - H_{r-1}(A_{t-1})$ is the rank-1 term corresponding to the minimum singular value of A_{t-1} . By smoothness we have

$$\begin{aligned}
& R(A_{t-1}) - R(A_t) \\
& \geq \max_{B \in \mathcal{B}_t} \{R(A_{t-1}) - R(B)\} \\
& \geq \max_{B \in \mathcal{B}_t} \left\{ -\langle \nabla R(A_{t-1}), B - A_{t-1} \rangle - \frac{\rho_2^+}{2} \|B - A_{t-1}\|_F^2 \right\} \\
& = \max_{\eta \in \mathbb{R}} \left\{ -\langle \nabla R(A_{t-1}), \eta uv^\top - \sigma_{\min}xy^\top \rangle - \frac{\rho_2^+}{2} \|\eta uv^\top - \sigma_{\min}xy^\top\|_F^2 \right\} \\
& \geq \max_{\eta \in \mathbb{R}} \left\{ -\langle \nabla R(A_{t-1}), \eta uv^\top \rangle - \eta^2 \rho_2^+ - \sigma_{\min}^2 \rho_2^+ \right\} \\
& = \max_{\eta \in \mathbb{R}} \left\{ \eta \|\nabla R(A_{t-1})\|_2 - \eta^2 \rho_2^+ - \sigma_{\min}^2 \rho_2^+ \right\} \\
& = \frac{\|\nabla R(A_{t-1})\|_2^2}{4\rho_2^+} - \sigma_{\min}^2 \rho_2^+,
\end{aligned}$$

where in the last inequality we used the fact that $\langle \nabla R(A_{t-1}), xy^\top \rangle = 0$ following from Lemma 9.6.8, as well as Lemma 9.6.1.

On the other hand, by strong convexity,

$$R(A^*) - R(A_{t-1}) \geq \langle \nabla R(A_{t-1}), A^* - A_{t-1} \rangle + \frac{\rho_{r+r^*}^-}{2} \|A^* - A_{t-1}\|_F^2.$$

Let $A_{t-1} = UV^\top$ and $A^* = U^*V^{*\top}$. We write

$$A^* = U^*V^{*\top} = (U^1 + U^2)(V^1 + V^2)^\top = U^1V^{1\top} + U^1V^{2\top} + U^2V^{*\top}$$

where U^1 is a matrix where every column of U^* is replaced by its projection on $\text{im}(U)$ and $U^2 = U^* - U^1$ and similarly V^1 is a matrix where every column of V^* is replaced by its projection on $\text{im}(V)$ and $V^2 = V^* - V^1$. By setting $U' = (-U \mid U^1)$ and $V' = (V \mid V^1)$ we can write

$$A^* - A_{t-1} = U'V'^\top + U^1V^{2\top} + U^2V^{*\top}$$

where $\text{im}(U') = \text{im}(U)$ and $\text{im}(V') = \text{im}(V)$. Also, note that

$$\text{rank}(U^1V^{2\top}) \leq \text{rank}(V^2) \leq \text{rank}(V^*) = \text{rank}(A^*) \leq r^*$$

and similarly $\text{rank}(U^2V^{*\top}) \leq r^*$. So we now have

$$\begin{aligned}
& \langle \nabla R(A_{t-1}), A^* - A_{t-1} \rangle + \frac{\rho_{r+r^*}^-}{2} \|A^* - A_{t-1}\|_F^2 \\
&= \langle \nabla R(A_{t-1}), U'V'^\top + U^1V^{2\top} + U^2V^{*\top} \rangle + \frac{\rho_{r+r^*}^-}{2} \|U'V'^\top + U^1V^{2\top} + U^2V^{*\top}\|_F^2 \\
&= \langle \nabla R(A_{t-1}), U^1V^{2\top} + U^2V^{*\top} \rangle + \frac{\rho_{r+r^*}^-}{2} \|U'V'^\top + U^1V^{2\top} + U^2V^{*\top}\|_F^2 \\
&= \langle \nabla R(A_{t-1}), U^1V^{2\top} + U^2V^{*\top} \rangle + \frac{\rho_{r+r^*}^-}{2} \left(\|U'V'^\top\|_F^2 + \|U^1V^{2\top}\|_F^2 + \|U^2V^{*\top}\|_F^2 \right) \\
&\geq \langle \nabla R(A_{t-1}), U^1V^{2\top} \rangle + \langle \nabla R(A_{t-1}), U^2V^{*\top} \rangle + \frac{\rho_{r+r^*}^-}{2} \left(\|U^1V^{2\top}\|_F^2 + \|U^2V^{*\top}\|_F^2 \right) \\
&+ \frac{\rho_{r+r^*}^-}{2} \|U'V'^\top\|_F^2 \\
&\geq 2 \min_{\text{rank}(M) \leq r^*} \left\{ \langle \nabla R(A_{t-1}), M \rangle + \frac{\rho_{r+r^*}^-}{2} \|M\|_F^2 \right\} + \frac{\rho_{r+r^*}^-}{2} \|U'V'^\top\|_F^2 \\
&= -2 \frac{\|H_{r^*}(\nabla R(A_{t-1}))\|_F^2}{2\rho_{r+r^*}^-} + \frac{\rho_{r+r^*}^-}{2} \|U'V'^\top\|_F^2 \\
&\geq -r^* \frac{\|\nabla R(A_{t-1})\|_2^2}{\rho_{r+r^*}^-} + \frac{\rho_{r+r^*}^-}{2} \|U'V'^\top\|_F^2
\end{aligned}$$

where the second equality follows from the fact that $\langle \nabla R(A_{t-1}), uv^\top \rangle = 0$ for any $u \in \text{im}(U), v \in \text{im}(V)$, the third equality from the fact that $\text{im}(U^2) \perp \text{im}(U') \cup \text{im}(U^1)$ and $\text{im}(V^2) \perp \text{im}(V')$ and by applying Lemma 9.6.7, and the last inequality from the fact that the largest squared singular value is not smaller than the average of the top r^* squared singular values. Now, note that since $\text{rank}(U^1V^{1\top}) \leq r^* < r = \text{rank}(UV^\top)$,

$$\begin{aligned}
\|U'V'^\top\|_F^2 &= \|U^1V^{1\top} - UV^\top\|_F^2 \\
&= \sum_{i=1}^r \sigma_i^2(U^1V^{1\top} - UV^\top) \\
&\geq \sum_{i=1}^r (\sigma_{i+r^*}(UV^\top) - \sigma_{r^*+1}(U^1V^{1\top}))^2 \\
&= \sum_{i=r^*+1}^r \sigma_i^2(UV^\top) \\
&\geq (r - r^*) \sigma_{\min}^2(UV^\top) \\
&= (r - r^*) \sigma_{\min}^2(A_{t-1}),
\end{aligned}$$

where we used the fact that $\text{rank}(U^1V^{1\top}) \leq r^*$ together with Lemma 9.6.5. Therefore we have concluded that

$$\|\nabla R(A_{t-1})\|_2^2 \geq \frac{\rho_{r+r^*}^-}{r^*} (R(A_{t-1}) - R(A^*)) + \frac{(\rho_{r+r^*}^-)^2 (r - r^*)}{2r^*} \sigma_{\min}^2$$

Plugging this back into the smoothness inequality and setting $\tilde{\kappa} = \frac{\rho_2^+}{\rho_{r+r^*}^-}$, we get

$$\begin{aligned} R(A_{t-1}) - R(A_t) &\geq \frac{1}{4r^*\tilde{\kappa}}(R(A_{t-1}) - R(A^*)) + \left(\frac{\rho_{r+r^*}^-(r-r^*)}{8r^*\tilde{\kappa}} - \rho_2^+ \right) \sigma_{\min}^2(A_{t-1}) \\ &\geq \frac{1}{4r^*\tilde{\kappa}}(R(A_{t-1}) - R(A^*)) \end{aligned}$$

as long as $r \geq r^*(1 + 8\tilde{\kappa}^2)$, or equivalently,

$$R(A_t) - R(A^*) \leq \left(1 - \frac{1}{4r^*\tilde{\kappa}} \right) (R(A_{t-1}) - R(A^*)).$$

Therefore after $L = 4r^*\tilde{\kappa} \log \frac{R(A_0) - R(A^*)}{\varepsilon}$ iterations we have

$$\begin{aligned} R(A_T) - R(A^*) &\leq \left(1 - \frac{1}{4r^*\tilde{\kappa}} \right)^L (R(A_0) - R(A^*)) \\ &\leq e^{-\frac{L}{4r^*\tilde{\kappa}}} (R(A_0) - R(A^*)) \\ &\leq \varepsilon \end{aligned}$$

Since $A_0 = \mathbf{0}$ and $\tilde{\kappa} \leq \kappa_{r+r^*}$, the result follows. \square

Tightness of the analysis

It is important to note that the κ_{r+r^*} factor that appears in the rank bounds of both Theorems 8.2.1 and 8.2.2 is inherent in these algorithms and not an artifact of our analysis. In particular, such lower bounds based on the restricted condition number have been previously shown for the problem of sparse linear regression. More specifically, [63] showed that there is a family of instances in which the analogues of Greedy and Local Search for sparse optimization require the sparsity to be $\Omega(s^*\kappa')$ for constant error $\varepsilon > 0$, where s^* is the optimal sparsity and κ' is the *sparsity*-restricted condition number. These instances can be easily adjusted to give a *rank* lower bound of $\Omega(r^*\kappa_{r+r^*})$ for constant error $\varepsilon > 0$, implying that the κ dependence in Theorem 8.2.1 is tight for Greedy. Furthermore, specifically for Local Search, [14] additionally showed that there is a family of instances in which the analogue of Local Search for sparse optimization requires a sparsity of $\Omega(s^*(\kappa')^2)$. Adapting these instances to the setting of rank-constrained convex optimization is less trivial, but we conjecture that it is possible, which would lead to a rank lower bound of $\Omega(r^*\kappa_{r+r^*}^2)$ for Local Search.

We present the following lemma, which essentially states that sparse optimization lower bounds for Orthogonal Matching Pursuit (OMP, [135]) (resp. Orthogonal Matching Pursuit with Replacement (OMPR, [87])) in which the optimal sparse solution is also a global optimum, immediately carry over (up to constants) to rank-constrained convex optimization lower bounds for Greedy (resp. Local Search).

Lemma 9.6.9. *Let $f \in \mathbb{R}^n \rightarrow \mathbb{R}$ and $x^* \in \mathbb{R}^n$ be an s^* -sparse vector that is also a global minimizer of f . Also, let f have restricted smoothness parameter β at sparsity level $s + s^*$ for some $s \geq s^*$ and restricted strong convexity parameter α at sparsity level $s + s^*$. Then we can define the rank-constrained problem, with $R : \mathbb{R}^{n \times n} \rightarrow \mathbb{R}$,*

$$\min_{\text{rank}(A) \leq s^*} R(A) := f(\text{diag}(A)) + \frac{\beta}{2} \|A - \text{diag}(A)\|_F^2, \quad (9.87)$$

where $\text{diag}(A)$ is a vector containing the diagonal of A . R has rank-restricted smoothness at rank $s + s^*$ at most 2β and rank-restricted strong convexity at rank $s + s^*$ at least α . Suppose that we run t iterations of OMP (resp. OMPR) starting from a solution x , to get solution x' , and similarly run t iterations of Greedy (resp. Local Search) starting from solution $A = \text{diag}(x)$ (where $\text{diag}(x)$ is a diagonal matrix with x on the diagonal) to get solution A' . Then A' is diagonal and $\text{diag}(A') = x'$. In other words, in this scenario OMP and Greedy (resp. OMPR and Local Search) are equivalent.

Proof. Note that for any solution \hat{A} of R we have $R(\hat{A}) \geq f(\text{diag}(\hat{A})) \geq f(x^*)$, with equality only if \hat{A} is diagonal. Furthermore, $\text{rank}(\text{diag}(x^*)) \leq s^*$, meaning that $\text{diag}(x^*)$ is an optimal solution of (9.87). Now, given any diagonal solution A of (9.87) such that $A = \text{diag}(x)$, we claim that one step of either Greedy or Local Search keeps it diagonal. This is because

$$\nabla R(A) = \text{diag}(\nabla f(x)) + \frac{\beta}{2}(A - \text{diag}(A)) = \text{diag}(\nabla f(x)).$$

Therefore the largest eigenvalue of $\nabla R(A)$ has corresponding eigenvector $\mathbf{1}_i$ for some i , which implies that the rank-1 component which will be added is a multiple of $\mathbf{1}_i \mathbf{1}_i^\top$. For the same reason the rank-1 component removed by Local Search will be a multiple of $\mathbf{1}_j \mathbf{1}_j^\top$ for some j . Therefore running Greedy (resp. Local Search) on such an instance is identical to running OMP (resp. OMPR) on the diagonal. \square

Together with the lower bound instances of [63] (in which the global minimum property is true), it immediately implies a rank lower bound of $\Omega(r^* \kappa_{r+r^*})$ for getting a solution with constant error for rank-constrained convex optimization. On the other hand, the lower bound instances of [14] give a *quadratic* lower bound in κ for OMPR. The above lemma cannot be directly applied since the sparse solutions are not global minima, but we conjecture that a similar proof will give a rank lower bound of $\Omega(r^* \kappa_{r+r^*}^2)$ for rank-constrained convex optimization with Local Search.

Addendum to Section 8.4


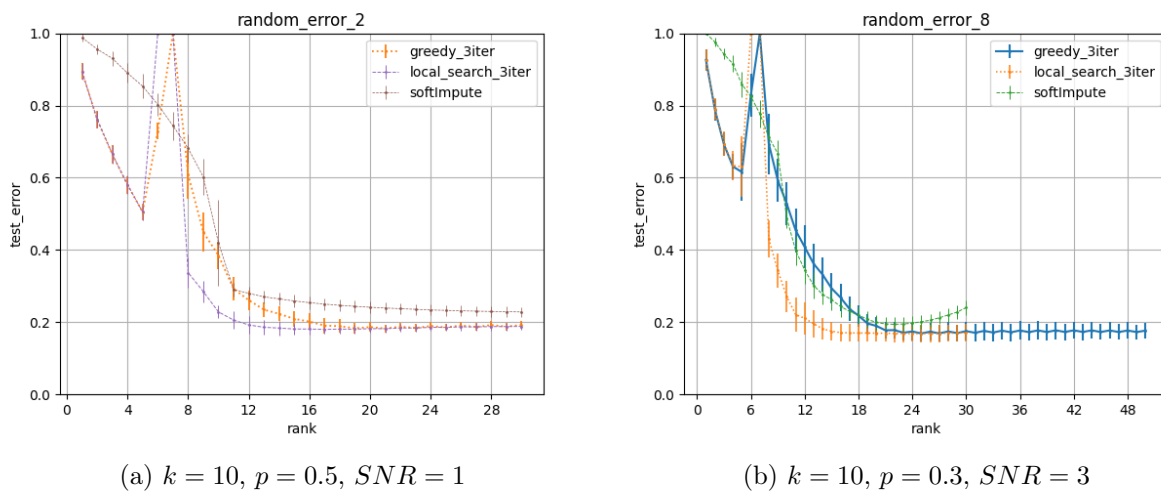



Figure 9-2: One of the splits of the Movielens 100K dataset. We can see that for small ranks the Fast Local Search solution is better and more stable, but for larger ranks it does not provide any improvement over the Fast Greedy algorithm.




(a) $k = 10, p = 0.5, SNR = 1$

(b) $k = 10, p = 0.3, SNR = 3$

Figure 9-3: Test error vs rank in the matrix completion problem of Section 8.4.2. Bands of ± 1 standard error are shown.



(a) Train error vs rank



(b) Test error vs rank

Figure 9-4: Performance of greedy with fully solving the inner optimization problem (left) and applying 3 iterations of the LSQR algorithm (right) in the matrix completion problem of Section 8.4.2. $k = 5$, $p = 0.2$, $SNR = 10$. Bands of ± 1 standard error are shown. This experiment shows why it is crucial to apply some kind of regularization to the Fast Greedy and Fast Local Search algorithms for machine learning applications.