

DISS. ETH NO. 31123

Fine-Grained Complexity and Algorithms for Structured Linear Equations and Linear Programs

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES
(Dr. sc. ETH Zurich)

presented by

Ming Ding

born on 16.02.1995

accepted on the recommendation of

Prof. Dr. Rasmus Kyng, examiner
Prof. Dr. David Steurer, co-examiner
Prof. Dr. Peng Zhang, co-examiner

2025

Abstract

Linear Equations (LEs) and Linear Programs (LPs) are fundamental tools in scientific research and practical applications, with widespread use in optimization, engineering, and data analysis. Many real-world problems exhibit specific structural properties, enabling structured problems to be solved significantly faster than general-purpose solvers for LEs and LPs. This thesis explores several structured LEs and LPs from both complexity and algorithmic perspectives.

For LPs, we investigate the computational complexity of the *two-commodity flow* (2CF) problem, a natural generalization of the well-studied single-commodity flow (i.e., maximum flow) problem. Both 2CF and maximum flow problems can be formulated as LPs, each defining a distinct family of structured LPs. Given the success of efficient algorithms for maximum flow and the structural similarities between the two problems, it was widely conjectured that 2CF could also be solved more efficiently than general LPs. Contrary to this belief, we prove that solving 2CF — either exactly or approximately to high accuracy — is as computationally hard as solving general LPs. We prove this by giving a nearly-linear time reduction that encodes any LP as a 2CF problem on a sparse directed graph with only a polylogarithmic blow-up in problem size. Furthermore, careful error analysis guarantees that any approximate 2CF solution can be mapped back into an approximate LP solution with only a polynomial increase in error.

For LEs, we focus on systems of linear equations whose coefficient matrices are *combinatorial Laplacians*, a class of generalized graph Laplacians that arise in higher-dimensional problems on *simplicial complexes*. Combinatorial Laplacians play a crucial role in homology (a central tool in topology), and have various applications in data analysis and physical modeling problems. While nearly-linear time solvers are known for graph Laplacians, efficient solvers for combinatorial Laplacians are only known for restricted classes of simplicial complexes. By constructing a nearly-linear time reduction from general LEs to combinatorial Laplacians of 2-complexes, we show that solving linear equations in combinatorial Laplacians is computationally as hard as solving general LEs. Notably, our reduction preserves problem sparsity up to polylogarithmic factors.

Finally on the algorithmic side, despite our hardness results for combinatorial Laplacians, we demonstrate that imposing geometric structure on simplicial complexes can lead to substantial computational improvements. In particular, we develop sub-quadratic time algorithms for approximately solving LEs in *1-Laplacians* (i.e., 1-dimensional combinatorial Laplacians) on well-shaped simplicial complexes up to high accuracy.

Zusammenfassung

Lineare Gleichungen (LEs) und Lineare Programme (LPs) sind grundlegende Werkzeuge in der wissenschaftlichen Forschung und in praktischen Anwendungen. Sie finden breite Anwendung in den Bereichen Optimierung, Ingenieurwesen und Datenanalyse. Viele reale Probleme weisen spezifische strukturelle Eigenschaften auf, wodurch sich strukturierte Probleme deutlich schneller lösen lassen als mit allgemeinen Lösungsverfahren für LEs und LPs. Diese Arbeit untersucht verschiedene strukturierte LEs und LPs sowohl aus komplexitätstheoretischer als auch aus algorithmischer Perspektive.

Für LPs analysieren wir die rechnerische Komplexität des *Zwei-Warenfluss-Problems* (2CF), einer natürlichen Verallgemeinerung des gut erforschten Ein-Warenfluss-Problems (d.h. Maximalflussproblem). Sowohl das 2CF- als auch das Maximum-Flow-Problem lassen sich als LPs formulieren, wobei jedes eine eigene Familie von strukturierten LPs definiert. Aufgrund des Erfolgs effizienter Algorithmen für das Maximalflussproblem und der strukturellen Ähnlichkeit beider Probleme wurde häufig vermutet, dass auch 2CF effizienter lösbar sein könnte als allgemeine LPs. Im Gegensatz zu dieser Annahme zeigen wir, dass die exakte oder hochgenaue approximative Lösung des 2CF-Problems genauso schwierig ist wie die Lösung allgemeiner LPs. Dies beweisen wir durch eine Reduktion in nahezu linearer Zeit, die jedes LP auf ein 2CF-Problem auf einem dünnbesetzten gerichteten Graphen abbildet, wobei die Problemgrösse lediglich um einen polylogarithmischen Faktor wächst. Eine sorgfältige Fehleranalyse stellt zudem sicher, dass jede approximative 2CF-Lösung in eine approximative LP-Lösung umgewandelt werden kann, wobei sich der Fehler nur polynomial vergrössert.

Bei LEs konzentrieren wir uns auf lineare Gleichungssysteme, deren Koeffizientenmatrizen *kombinatorische Laplace-Operatoren* sind, eine Klasse verallgemeinerter Graph-Laplace-Operatoren, welche bei höherdimensionalen Problemen auf *Simplizialkomplexen* auftreten. Kombinatorische Laplace-Operatoren spielen eine zentrale Rolle in der Homologie (einem wesentlichen Werkzeug der Topologie) und besitzen vielfältige Anwendungen in der Datenanalyse sowie in der physikalischen Modellierung. Während für Graph-Laplace-Operatoren nahezu lineare Lösungsverfahren bekannt sind, existieren für kombinatorische Laplace-Operatoren effiziente Algorithmen bislang nur für eingeschränkte Klassen von Simplizialkomplexen. Mittels einer nahezu linearen Reduktion von allgemeinen linearen Gleichungssystemen auf kombinatorische Laplace-Gleichungssysteme auf 2-Komplexen zeigen wir, dass das Lösen linearer Gleichungen mit kombinatorischen Laplace-Operatoren ebenso komplex ist wie das Lösen allgemeiner linearer Gleichungssysteme. Bemerkenswerterweise erhält unsere Reduktion dabei die Problem-Sparsität bis auf polylogarithmische Faktoren.

Abschliessend demonstrieren wir auf algorithmischer Ebene, dass trotz unserer Härteresultate für kombinatorische Laplace-Operatoren eine geometrische Strukturierung der Simplizialkomplexe erhebliche algorithmische Verbesserungen ermöglicht. Insbesondere entwickeln wir subquadratische Algorithmen zur hochgenauen Approximation lin-

earer Gleichungssysteme mit *1-Laplace-Operatoren* (also eindimensionalen kombinatorischen Laplace-Operatoren) auf geometrisch wohldefinierten Simplizialkomplexe.

Acknowledgements

I would like to express my deepest gratitude to my Ph.D. supervisor, Rasmus Kyng. I still vividly remember our enjoyable first interview, where we discussed the graph sparsification paper—it was my first exposure to the world of theoretical computer science. I was so excited to join your research group and start this journey.

I am especially grateful for your remarkable generosity with your time, as well as your patience in coaching me and solving problems together. I have always been amazed by your exceptional problem-solving skills—how you develop deep intuition and break down complex problems into tractable pieces. Your optimism and persistence, even when things seemed impossible, have been truly inspiring. Your guidance has not only shaped the research in this thesis but has also deeply influenced the way I approach challenges in life. Beyond that, I also sincerely appreciate the freedom you gave me to explore my own research interests and career aspirations.

I would also like to extend my sincere thanks to my co-examiners, David Steurer and Peng Zhang, for kindly agreeing to join my doctoral examination committee and for taking the time and effort to support my Ph.D. defense and graduation. A special thanks to Peng Zhang, who has also been my mentor and co-author. Peng, I am truly grateful for your hands-on tutoring in numerical linear algebra, and for the extra time you spent with me after our project meetings with Rasmus—patiently answering my questions and offering encouragement when I needed it most.

I am also lucky to have had the chance to work with amazing collaborators: Shunhua Jiang, Omri Weinstein, Maximilian Probst Gutenberg, and Deeksha Adil. I really enjoyed our discussions, tackling tough research questions together, and sharing both the excitement and struggles along the way.

A big thanks to my colleagues at ETH Zurich. I thank all the members of Rasmus’s research group. In particular, I want to thank my office mates—Federico Soldà, Simon Meierhans, and Aurelio Sulser—for the fun conversations, lunchtime discussions, and other group activities. I also want to thank Claudia Günthart and Andrea Salow for all their help with administrative matters.

Outside of research, I owe thanks to my friends, too many to name, who have made life outside of work fun and fulfilling. I feel incredibly lucky to have met my partner, Alexandre Binniger, through the summer retreat, one of the most memorable events organized by VMI at ETH Zurich. Alexandre, thank you for your support, encouragement, and companionship throughout this journey. You have brought immense happiness into my life. I am also grateful to your family for their warmth and kindness.

Finally, I want to thank my parents and brother, who, even from far away in China, have always given me unconditional love and support throughout my life.

Contents

1	Introduction	1
1.1	Solving Linear Equations	3
1.2	Solving Linear Programs	5
1.3	Fine-Grained Complexity Analysis	7
1.4	Our Results	9
1.5	Organization of Thesis	13
2	Preliminaries	15
2.1	Vectors and Matrices	15
2.2	Simplicial Homology	17
2.3	Approximately Solving Linear Equations	19
3	Hardness Results for Two-Commodity Flow	23
3.1	Prior Works	23
3.2	Our Contributions	23
3.3	Problem Definitions	24
3.3.1	Discussion of the Problem Assumptions	24
3.3.2	Problems in Algebra Space	26
3.3.3	Problems in Flow Space	27
3.4	Main Results	30
3.5	Algebra Space	31
3.5.1	Overview	31
3.5.2	LP(A) to LEN(A): Reducing Inequalities to Equalities	32
3.5.3	LEN(A) to 2-LEN(A): Reducing Integer Coefficients to $\{0, \pm 1, \pm 2\}$	33
3.5.4	2-LEN(A) to 1-LEN(A): Reducing Coefficients from ± 2 to ± 1	35
3.5.5	1-LEN(A) to FHF(A): Encoding Equations as Flows	37
3.6	Flow Space	39
3.6.1	Overview	39
3.6.2	FHF(A) to FPHF(A)	42
3.6.3	FPHF(A) to SFF(A): Dropping Homologous Flow Constraints	44
3.6.4	SFF(A) to 2CFF(A): Dropping Selective Flow Constraints	48
3.6.5	2CFF(A) to 2CFR(A): Dropping Fixed Flow Constraints	51
3.6.6	2CFR(A) to 2CF(A)	55
3.7	A Unified Framework for LP Transformations	57
3.7.1	Characterizing LP Transformations	57
3.7.2	Unified Correctness Analysis	58
3.7.3	Unified Error Analysis	59
3.7.4	Examples	60

3.7.5	Summary	65
4	Hardness Results for Combinatorial Laplacians	67
4.1	Prior Works	67
4.2	Our Contributions	67
4.3	Notations and Preliminaries	68
4.3.1	Matrix Classes	68
4.3.2	Reduction Between Linear Equations	69
4.4	Main Results and Reduction Outline	69
4.4.1	Overview of Our Proof	70
4.5	Reducing \mathcal{DA} to \mathcal{B}_2 in Feasible Case	72
4.5.1	Reduction Algorithm	72
4.5.2	Additional Notations	76
4.5.3	Algorithm Runtime and Problem Size	76
4.5.4	Relation Between Exact Solutions	77
4.5.5	Relation Between Approximate Solutions	78
4.5.6	Bounding the Condition Number of the New Matrix	80
4.6	Reducing \mathcal{DA} to \mathcal{B}_2 in General Case	83
4.6.1	Warm-Up: Reweighting Infeasible Equations to Preserve Solutions	83
4.6.2	Reduction Algorithm	84
4.6.3	Relation Between Exact Solutions	85
4.6.4	Relation Between Approximate Solutions	87
4.6.5	Bounding the Condition Number of the New Matrix	89
5	Hardness Results for More Structured Problems	91
5.1	Reducing General LE to Difference-Average LE	91
5.2	Reducing Difference-Average LE to 1-or-3 LE	93
5.2.1	Reducing Average Equations	95
5.2.2	Reducing Difference Equations	95
5.3	Reducing General LP to (Simplified) 1-or-3 LP	96
5.3.1	Reducing General LP to Scaled LP	97
5.3.2	Reducing Scaled LP to Difference-Average LP	98
5.3.3	Reducing Difference-Average LP to 1-or-3 LP	98
5.3.4	Reducing 1-or-3 LP to Simplified 1-or-3 LP	99
6	1-Laplacian Solver for Well-Shaped Simplicial Complexes	101
6.1	Prior Works	101
6.2	Motivations and Applications	102
6.3	Notations and Preliminaries	103
6.4	Main Results	104
6.5	Algorithm Overview	105
6.5.1	Down-Laplacian	105
6.5.2	Solver for Up-Laplacian	105
6.5.3	Projection onto the Image of Up-Laplacian	108
6.6	Solver for Down-Laplacian	109
6.7	Solver for Up-Laplacian	110
6.7.1	Solver for $\mathbf{L}_1^{\text{up}}(F, F)$	110
6.7.2	Solver for the Schur Complement	113
6.8	Projection onto the Image of Up-Laplacian	117

6.8.1	Preconditioning the Schur Complement	119
6.8.2	Proof of Lemma 6.5.9	122
6.9	Proof of the Main Theorem	123
6.10	A Union of Pure 3-Complexes	125
Bibliography		127
A Missing Proofs of Linear Algebra Facts		137
B Missing Proofs of Chapter 4		141
B.1	Reducing 2-Complex Boundary LE to Combinatorial Laplacian LE . . .	141
B.2	Connections With Interior Point Methods	142
C Missing Proofs of Chapter 6		145

Chapter 1

Introduction

Linear equations (LEs) and linear programs (LPs) are fundamental tools in both theory and practice. They arise in a wide range of fields, including computer science, engineering, economics, operations research, and the natural sciences. Moreover, efficient methods for solving LEs and LPs form the backbone of many core algorithmic paradigms, such as interior point methods, graph algorithms, and machine learning techniques. Given the importance of linear equations and linear programs, understanding their computational complexity is a central focus in theoretical computer science, numerical linear algebra, and optimization. Advances in this area not only deepen our theoretical understanding but also lead to faster and more scalable algorithms, driving progress in practical applications across science and engineering.

Given a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, a vector $\mathbf{b} \in \mathbb{R}^m$, a system of linear equations is defined as solving a vector $\mathbf{x} \in \mathbb{R}^n$ such that

$$\mathbf{Ax} = \mathbf{b}.$$

In practice, we allow small error $\epsilon > 0$, and the goal is to compute an approximate vector $\mathbf{x} \in \mathbb{R}^n$ such that

$$\|\mathbf{Ax} - \mathbf{b}\|_2 \leq \epsilon \|\mathbf{b}\|_2.$$

A linear program is an optimization problem with a linear objective function and linear constraints. Given a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, a vector $\mathbf{b} \in \mathbb{R}^m$, and a vector $\mathbf{c} \in \mathbb{R}^n$, a canonical form of LP is as follows:

$$\max_{\mathbf{x} \in \mathbb{R}^n} \{\mathbf{c}^\top \mathbf{x} : \mathbf{Ax} \leq \mathbf{b}, \mathbf{x} \geq \mathbf{0}\}.$$

Both LE and LP can be considered as special cases of a more general problem – *linear regression* problem. Given a (tall) matrix \mathbf{A} and a vector \mathbf{b} , it seeks a vector \mathbf{x} to minimize $\|\mathbf{Ax} - \mathbf{b}\|_p$ for some ℓ_p -norm. When $p = 1$ or ∞ , the regression problem is equivalent to solving linear programs; when $p = 2$, the regression problem is equivalent to solving linear equations.

LE and LP are also related via Interior Point Methods (IPM), a family of algorithms that solve linear programs in provably polynomial time. IPM reduces solving LP to solving a sequence of $\tilde{O}(\sqrt{\text{rank}(\mathbf{A})}L)$ systems of linear equations, where L is the bit complexity of the input [LS14]. Thus, any improvement in solving linear equations would imply a direct improvement in solving linear programs. Moreover, both linear programs and linear equations are frequently used as subroutines for other optimization problems.

Although significant progress has been made in designing fast solvers, the runtime of algorithms for solving linear equations and linear programs get stuck at the current matrix multiplicative time: the fastest solver for general system of n linear equations over $O(n)$ variables runs in time $O(n^\omega)$, where $\omega \approx 2.3728596\dots$ [AW21]; the fastest solver for general linear programs with n constraints over $O(n)$ variables runs in time $\tilde{O}(n^\omega)$ ¹ [CLS21; Jia+21]. However, these runtimes can be impractically slow for modern large-scale datasets.

In contrast to general LE and LP, many problems in practice possess additional structures that enable practically and provably faster solvers. Notable examples include: fast solvers for linear equations in graph Laplacians, which run in nearly-linear time [ST04; ST08; DS07; KS16; Kyn+16]; and fast algorithms for the maximum flow problem, a special structured LP problem, which can be solved in sub-quadratic and even almost-linear time [Mad13; Mad16; LS20; GLP21; Che+22]. These breakthroughs lead to a compelling open question:

To what extent can these efficient algorithms be generalized to broader classes of structured linear equations and linear programs?

There are two complementary perspectives to approach this question: the *complexity perspective* and the *algorithmic perspective*. From the complexity perspective, our goal is to establish computational lower bounds for structured problems. Unlike classical complexity theory, which focuses on distinguishing “P” from “NP”, our analysis falls under *fine-grained complexity*, as both LE and LP are already known to be in “P”. Specifically, we aim to show that certain structured LEs and LPs problems cannot be solved faster than a given time complexity, unless widely accepted complexity assumptions for LE and LP fail. Such lower bounds help identify the fundamental limits of the structured problems and how much room there is for further progress. From the algorithmic perspective, the focus is on designing faster algorithms by leveraging the problem’s specific structure. This helps narrow the gap between theoretical complexity barriers and performance of practical algorithms. Together, these two perspectives provide insights into the limitations and potentials of efficient solvers for structured problems.

In this thesis, we examine certain structured linear equations and linear programs from both the complexity perspective and the algorithmic perspective. From the complexity perspective, we investigate natural generalizations of problems with known fast algorithms. The two-commodity flow problem (2CF) generalizes the classical maximum flow problem.² Similarly, combinatorial Laplacians generalize graph Laplacians to higher-dimensional graphs – simplicial complexes. While 2CF and combinatorial Laplacians share structural similarities with their classical counterparts, no efficient algorithms have been discovered for these problems, apart from some specific restricted cases. In this thesis, we explain that the absence of efficient algorithms is not a coincidence. More specifically, we establish hardness results for 2CF, showing that approximately solving 2CF is equally hard as approximately solving general LP to high accuracy. And we prove that approximately solving linear equations in combinatorial Laplacians is as computationally hard as solving general linear equations to high accuracy. These results are achieved by constructing efficient reduction algorithms from a general LP or LE into these structured problems under study.

¹We use $\tilde{O}(\cdot)$ to hide a polylogarithmic factor of dimensions, condition number, and accuracy.

²The maximum flow problem can be viewed as a single-commodity flow problem.

Analogous to NP-completeness, we introduce new complexity classes, called *sparse-linear-equation (or sparse-linear-program) completeness*. The term “completeness” reflects the idea that if an algorithm with runtime $\tilde{O}((\text{number of non-zeros})^c)$ is known for the structured LEs (or LPs), then an algorithm with runtime $\tilde{O}((\text{number of non-zeros})^c)$ exists for general LEs (or LPs). In this terminology, the above results can be equivalently stated as 2CF and combinatorial Laplacians being sparse-LP and sparse-LE complete, respectively. Additionally, we introduce a simplified structured problem called *1-or-3 LE (or 1-or-3 LP)*, which is likewise sparse-LE (or sparse-LP) complete. We find that it plays a role similar to 3-SAT in NP-completeness. Our findings coincide with those of Průša and Werner (SODA’2017) [PW17], who further encoded 1-or-3 LP into a number of LP relaxations of NP-hard problems, showing that these LP relaxations are also sparse-LP complete.

Moreover, we investigate combinatorial Laplacians from an algorithmic perspective as well. While our hardness results indicate that the general structure of combinatorial Laplacians does not inherently provide computational advantages, we show that introducing mild *geometric assumptions* on the underlying simplicial complexes can lead to significant algorithmic improvements. Specifically, we develop a sub-quadratic time solver for linear equations in *1-Laplacians* (i.e., 1-dimensional combinatorial Laplacians) when the underlying simplicial complexes are *well-shaped* tetrahedral meshes.

1.1 Solving Linear Equations

Given linear equations $\mathbf{A}\mathbf{x} = \mathbf{b}$ where $\mathbf{A} \in \mathbb{R}^{n \times n}$, one can solve it by *Gaussian Elimination*, which runs in time $O(n^3)$. If applying the fast matrix multiplication by Strassen [Str69], one can speed up solving the linear equations in runtime $\tilde{O}(n^\omega)$ by inverting \mathbf{A} and then multiplying $\mathbf{A}^{-1}\mathbf{b}$, and it is by far the best known algorithm for solving general dense linear equations in dimensions $n \times n$ [AW21].

When the coefficient matrix \mathbf{A} is sparse, i.e., \mathbf{A} has $\tilde{O}(n)$ non-zero entries, the direct matrix inversion approach is not desirable as \mathbf{A}^{-1} cannot preserve sparsity of \mathbf{A} . Instead, iterative methods are more desirable to be used for solving sparse problems, where only matrix-vector multiplication and vector-vector operations are conducted in each iteration. For sparse linear equations with N non-zero coefficients and condition number κ , the best known approximate algorithms run in time $\tilde{O}(\min\{N^{2.27159}, N\kappa\})$, where the first runtime is from [PV21; Nie22] and the second is by conjugate gradient [HS+52]. In particular, if the coefficient matrix is symmetric positive semi-definite, the runtime for conjugate gradient is $\tilde{O}(N\sqrt{\kappa})$.

Graph Laplacians and Nearly-Linear Time Solvers. An $n \times n$ Laplacian matrix can be viewed as an undirected graph over n vertices and m edges. Each vertex of the graph corresponds to a single row and a single column of the matrix. Each edge corresponds to an n -dimensional vector whose non-zeros are only at the position of its two endpoints. Formally, a graph Laplacian is defined as

$$\mathbf{L} = \mathbf{D} - \mathbf{A},$$

where $\mathbf{D} \in \mathbb{R}^{n \times n}$ is a diagonal matrix with vertex degrees on the diagonal entries; and $\mathbf{A} \in \mathbb{R}^{n \times n}$ is a *adjacency matrix*, defined such that $\mathbf{A}(u, v) = 1$ if (u, v) is an edge and

0 otherwise. Alternatively, graph Laplacians can also be defined using the vertex-edge incidence matrix $\mathbf{B} \in \mathbb{R}^{n \times m}$, where

$$\mathbf{B}(v, e) = \begin{cases} 1 & \text{if } e = (u, v) \\ -1 & \text{if } e = (v, u) \\ 0 & \text{otherwise} \end{cases}. \quad (1.1)$$

In this case, the graph Laplacian is given by

$$\mathbf{L} = \mathbf{B}\mathbf{B}^\top.$$

It is well-known that linear equations in graph Laplacians can be approximately solved in nearly-linear time in the number of non-zeros of the graph Laplacians [ST14; KMP10; KMP11; Kel+13; LS13; PS14; Coh+14b; KS16; Kyn+16; JS21].

Linear equations in graph Laplacians arise from using IPM for combinatorial problems on graphs. As a result, these fast Laplacian solvers have led to significant developments in algorithm design for graph problems such as maximum flow [Mad13; Mad16; Che+22], minimum cost flow and lossy flow [LS14; DS08], and graph sparsification [SS08]. Collectively, using fast Laplacian solvers to speed up graph algorithms was formulated as the ‘‘Laplacian Paradigm’’ [Ten10]. Beyond graph Laplacians, nearly-linear time solvers were extended to symmetric M-matrices [DS08], Block Diagonally Dominant matrices and Connection Laplacians [Kyn+16], Directed Laplacians [Coh+17b].

Combinatorial Laplacians. Combinatorial Laplacians generalize graph Laplacians to higher dimensional simplicial complexes – a collection of 0-simplexes (vertices), 1-simplexes (edges), 2-simplexes (triangles), and their higher dimensional counterparts. Given a d -dimensional simplicial complex \mathcal{K} , for each $0 \leq i \leq d$, ∂_i is a linear map that maps every i -simplex to a signed sum of its boundary $(i - 1)$ -faces. We define the i -Laplacian \mathbf{L}_i to be

$$\mathbf{L}_i = \partial_i^\top \partial_i + \partial_{i+1} \partial_{i+1}^\top. \quad (1.2)$$

In particular, ∂_1 is the vertex-edge incidence matrix defined in Equation (1.1), and \mathbf{L}_0 is the graph Laplacian (following the convention, we define $\partial_0 = \mathbf{0}$).

Combinatorial Laplacians play an important role in both pure mathematics and applied areas. These matrices originate in the study of discrete Hodge decomposition [Eck44]: The kernel of \mathbf{L}_i is isomorphic to the i th homology space of \mathcal{K} . The properties of combinatorial Laplacians have been studied in many subsequent works [Fri98; DW02; DKM09; DKM15; MN21]. A central problem in homology theory is evaluating the *Betti number*³ of the i th homology space, which equals the rank of \mathbf{L}_i . In the case of homology over the reals, computing the rank of \mathbf{L}_i can be reduced to solving a polylogarithmic number of linear equations in \mathbf{L}_i [BV21]. In applied areas, the computation of Betti numbers over the reals is a key step in numerous problems in applied topology, computational topology, and topological data analysis [Zom05; Ghr08a; Car09; EH10; Cha+16]. In addition, combinatorial Laplacians have applications in statistical ranking [Jia+11; Xu+12], graphics and image processing [Ma+11a; Ton+03a], electromagnetism

³Informally, the i th Betti number is the number of i -dimensional holes on a topological surface. For example, the zeroth, first, and second Betti numbers represent the numbers of connected components, one-dimensional ‘‘circular’’ holes, and two-dimensional ‘‘voids’’ or ‘‘cavities,’’ respectively.

and fluids mechanics [DKT08], data representations [CMZ18], cryo-electron microscopy [YL17a], biology [Sch+20]. We refer to the readers to [Lim20] for an accessible survey.

Inspired by the success of graph Laplacians, [Coh+14a] initiated the study of fast approximate solvers for 1-Laplacian linear equations. They designed a nearly-linear time approximate solver for simplicial complexes *with zero Betti numbers and known collapsing sequences*.⁴ Later, [Bla+22] and [BN22] generalized this algorithm to subcomplexes of such a complex *with bounded first Betti numbers*.⁵ One concrete example studied in these papers is convex simplicial complexes that piecewise linearly triangulate a convex ball in \mathbb{R}^3 , for which a collapsing sequence exists and can be computed in linear time [Chi67; Chi80].

Lately, an algorithm about spectral sparsification for dense simplicial complexes was proposed by [Sav+25]. They proposed a probabilistic sampling method that reduces the number of $(k+1)$ -dimensional simplexes from $m_{k+1} \approx O(m_k^{\frac{k+1}{k}})$ (the complete-graph case) down to $\tilde{O}(m_k)$, in time $O(\delta^{-3} m_k^{1+\frac{4}{1+k}})$, where δ controls the approximation error. A direct application of this method is solving (dense) linear systems in the up-Laplacian $\partial_{k+1} \partial_{k+1}^\top \in \mathbb{R}^{m_k \times m_k}$. Specifically, by firstly sparsifying the $(k+1)$ -simplexes via [Sav+25]’s approach and then using the state-of-the-art sparse linear solver from [Nie22], the resulting runtime is $\tilde{O}((k^2 m_k)^{2.27})$. This combined approach beats the direct pseudo-inverse approach $O(m_{k+1}^\omega) = O(m_k^{(1+\frac{1}{k})\omega})$ for small k .

Geometric Structures and Nested Dissection. Another widely studied class of structured linear equations consists of those whose non-zero entries exhibit geometrically embedded structures. Given an $n \times n$ matrix \mathbf{A} , its non-zero structures can be viewed as a graph over n vertices, and the (i, j) th entry or the (j, i) th entry of \mathbf{A} is non-zero if there is an edge between vertex i and vertex j .

Asymptotically faster solvers exist if the geometric structures can be solved efficiently by *Nested Dissection* [Geo73; LRT79; MT90; AY10]. Specifically, Nested Dissection is a method to compute an ordering for Gaussian Elimination such that the number of *fill-ins* is small. Fill-ins refer to the number of new non-zeros introduced in the matrix when eliminating a variable. Using this idea, one can solve linear equations in a 2D planar in time $O(n^{1.5})$ [LRT79]. Additionally, faster solvers are also known for restricted classes of total-variation matrices [KMT11], stiffness matrices from elliptic finite element systems [BHV08a], and 2D and 3D truss stiffness matrices [DS07; Kyn+18].

1.2 Solving Linear Programs

The development of modern LP algorithms began with the introduction of the *Simplex Method* by George Dantzig in 1947. While this method has proven to be highly efficient in practice, its worst-case performance classifies it as an exponential-time algorithm in theory. In the 1970s, the *Ellipsoid Method* was introduced for solving certain non-linear optimization problems. In 1979, Khachiyan [Kha80a] demonstrated how this method could be adapted to solve LPs, marking a significant theoretical milestone as the first

⁴A collapsing sequence gives a good ordering of Gaussian Elimination so that it runs in linear time. However, deciding whether a general simplicial complex has a collapsing sequence is NP hard [Tan16].

⁵The solver has cubic dependence on the first Betti number.

polynomial-time algorithm for linear programming. However, the Ellipsoid Method is not competitive with the Simplex Method in practice.

The *Interior Point Methods* (IPMs) represent another major advancement in LP algorithms. Unlike the Simplex Method, which traverses along the boundary of the feasible polytope from one extreme point to another, or the Ellipsoid Method, which encircles the polytope, IPMs navigate through the interior of the feasible solution space. In 1984, Karmarkar [Kar84] proposed a polynomial-time IPM algorithm, further advancing the field of LP. In practice, IPMs often rival the Simplex Method and, in particular, outperform it on large-scale problems, making them a critical tool in modern optimization.

The fastest known solvers for general linear programs are based on Interior Point Methods, and in particular the central path methods [Ren88]. Recently, there has been significant progress on solvers for general linear programs, but the running time required to solve a linear program with n variables and $\tilde{O}(n)$ constraints (assuming polynomially bounded entries and polytope radius) remains stuck at the $\tilde{O}(n^{\max\{\omega, 2+1/18\}})$ [Jia+21].

Maximum Flow/Single-Commodity Flow Problem. Single-commodity flow problems have been an area of tremendous success for the development of graph algorithms, starting with an era of algorithms influenced by early results on maximum flow and minimum cut [FF56] and later the development of powerful combinatorial algorithms for maximum flow with polynomially bounded edge capacities [Din70; ET75; GR98]. Later, a breakthrough nearly-linear time algorithm for electrical flows by Spielman and Teng [ST04] lead to the *Laplacian paradigm*. A long line of work explored direct improvements and simplifications of this result [KMP10; KMP11; Kel+13; PS14; KS16; JS21]. The Laplacian paradigm also motivated research into undirected maximum flow [Chr+11; LRS13; Kel+14; She13], which subsequently led to faster algorithms for directed maximum flow and minimum cost flow problems [Mad13; Mad16; LS20; KLS20; van+21; GLP21] building on powerful tools using mixed- ℓ_2, ℓ_p -norm minimizing flows [Kyn+19] and inverse-maintenance ideas [Che+20]. A recent breakthrough is the development of an almost-linear time algorithm for maximum flow and minimum-cost flow by Chen et al. [Che+22], which currently represents the fastest approach. This algorithm computes exact maximum flows and minimum-cost flows on directed graphs with $|E|$ edges and polynomially bounded integral demands, costs, and capacities in $|E|^{1+o(1)}$ time.

To compare these advances with general-purpose LP solvers, consider the running times for state-of-the-art single-commodity maximum flow algorithms on a graph with $|V|$ vertices and $|E|$ edges. For sparse graphs where $|E| = \tilde{O}(|V|)$, formulating the maximum flow problem as a linear program and solving it with general LP solvers results in a running time of $\tilde{O}(|V|^{2.372\dots})$. In contrast, specialized maximum flow solvers achieve a much faster $\tilde{O}(|V|^{1+o(1)})$ runtime. On dense graphs with $|E| = \Theta(|V|^2)$, the gap narrows but remains significant: maximum flow solvers achieve a runtime of $\tilde{O}(|V|^2)$, compared to $\tilde{O}(|V|^{2.372\dots})$ for general LP solvers. Therefore, the structure of maxflow can be used to develop special-purpose solvers that outperform general LP solvers.

Multi-Commodity Flow Problem. Multi-commodity flow problems are a well-studied area in algorithm design and have been the subject of numerous surveys [Ken78; AMO93; OMV00; BKV09; Wan18], in part because a large number of problems can be expressed as variants of multi-commodity flow. Significant advancements have been made in solving the undirected multi-commodity flow problem, particularly in the low-accuracy

regime. Leighton et al. [Lei+95] showed that undirected capacitated k -commodity flow in a graph with $|E|$ edges and $|V|$ vertices can be approximately solved in $\tilde{O}(k|E||V|)$ time, completely routing all demands with $(1 + \epsilon)$ times the optimal congestion, albeit with a poor dependence on the error parameter ϵ . This beats the solve-times for linear programming in sparse graphs for small k , even with today’s LP solvers that run in current matrix multiplication time, albeit with much worse error. This result inspired a series of follow-up works focusing on improving low-accuracy algorithms [GK07; Fle00; Mad10]. Later breakthroughs in achieving almost- and nearly-linear time algorithms for undirected single-commodity maximum flow also led to faster algorithms for undirected k -commodity flow [Kel+14; She13; Pen16], culminating in Sherman’s introduction of area-convexity to build a $\tilde{O}(k|E|\epsilon^{-1})$ time algorithm for approximate undirected k -commodity flow [She17].

In contrast, directed 2-commodity flow problems were seemingly harder, despite the discovery of non-trivial algorithms for some special cases [Eva76; Eva78]. Alon Itai [Ita78] proved a polynomial-time reduction from linear programming to directed 2-commodity flow, before a polynomial-time algorithm for linear programming was known. For decades, the only major progress on solving directed multi-commodity flow to high accuracy came from improvements to general linear program solvers [Kha80b; Kar84; Ren88; Vai89].

Recently, a surprising advancement was made by van den Brand and Zhang [DZ23a], who developed a high-accuracy algorithm for the k -commodity flow problem with a runtime of $\tilde{O}(k^{2.5}\sqrt{|E||V|^{\omega-1/2}})$. For comparison, solving the problem using the state-of-the-art general LP solver [CLS21; Jia+21] takes runtime is $\tilde{O}((k|E|)^\omega)$. While this algorithm is slightly less efficient than general LP solvers in sparse graphs (i.e., $|E| = \tilde{O}(|V|)$), it outperforms them on dense graphs (i.e., $|E| = \Theta(|V|^2)$). This marks the first improvement to high accuracy multi-commodity flow algorithms that does not just stem from improvements to general linear program solvers.

1.3 Fine-Grained Complexity Analysis

In classical complexity theory, we typically classify decision problems into complexity classes such as “P” or “NP”, but do not often measure how difficult a polynomial-time problem might be. Fine-grained complexity theory addresses this gap by quantifying the precise difficulty of problems in P, going beyond the standard “P vs NP” framework. It explains conditional time-complexity barriers for polynomial-time solvable problems, relying on “core” conjectures (like the Strong Exponential-Time Hypothesis [IP01], the Orthogonal Vectors Conjecture [Wil05], and assumptions about matrix multiplication). By designing so-called *fine-grained reductions* from problem A to problem B , one can show that a seemingly faster algorithm for problem B would imply a surprising speedup for problem A . If problem A is believed to require a certain runtime, then problem B also inherits this hardness barrier. Over the last decade, fine-grained complexity has blossomed, yielding breakthroughs in explaining why certain polynomial-time problems remain “as slow as” the best-known methods. It has also led to the discovery of many meaningful relationships between problems, and equivalent classes. For a thorough treatment of the foundational ideas and recent developments in fine-grained complexity and algorithms, we refer interested readers to [Wil18].

Fine-Grained Reductions. The basic setup of fine-grained lower bounds is similar to classic NP-harness reductions: A fine-grained reduction from problem A to problem B

is an algorithm that given an instance I of size n for problem A , computes in time $t(n)$ an equivalent instance J of size $s(n)$ for problem B . Thus, if there exists an algorithm that solves problem B in time $T(n)$, then by this reduction there is an algorithm solving problem A in time $t(n) + T(s(n))$. In particular, if $t(n) + T(s(n))$ is faster than the hypothesized optimal time complexity of problem A , then problem B cannot be solved in time $T(n)$ assuming the hypothesis for A . This provides a conditional lower bound of problem B . For instance, under the rank-finding conjecture (i.e., no sub- n^ω method can compute the rank of an $\Theta(n) \times \Theta(n)$ matrix), [BV21] gave conditional lower bounds for general (dense) system of linear equations that even to get constant-factor or $(1 - O(1/n^c))$ -factor approximation, we likely need $O(n^\omega)$ time.

Sometimes the chain of reductions goes in both directions, showing that problems A and B are, in effect, “fine-grained equivalent”. In this case, if either problem’s time complexity were improved, it would yield an improvement for the other. For example, [WW18] introduced “subcubic reducibility”, which requires $t(n) = O(n^c)$ for some constant $c < 3$. This work showed that many important problems on graphs and matrices solvable in $O(n^3)$ time are equivalent under subcubic reductions (e.g., Negative-Triangle Detection, Matrix-Product Verification). A single-direction reduction can also suffice if problem B is a special case of A . Itai [Ita78] devised a polynomial-time reduction from a linear program to a 2-commodity flow problem. Since 2-commodity flow is a special case of linear programming, Itai concluded that solving 2-commodity flow and solving linear programs are polynomial-time equivalent. In a similar vein, Kyng and Zhang [KZ17] studied several structured linear systems, including two-commodity Laplacians, 2D truss stiffness matrices, and total variation matrices. By constructing reduction algorithms from general linear equations to these structured linear equations, it is concluded that if we can quickly solve either of these structured linear equations, then we can quickly solve linear equations in any matrix.

Error Analysis in Reductions. Many fine-grained reductions focus on exact versions of problems. For instance, Itai [Ita78] established that exactly solving linear programming is polynomial-time equivalent to exactly solving a sequence of related problems (e.g., linear equalities under nonnegative constraints, homologous flow, two-commodity flow). Dobkin and Reiss [DR80] later identified additional LP-complete problems in computational geometry, such as hyperplane intersection and extreme-point detection. Trevisan and Xhafa [TX98] proved that exactly solving packing LPs is P-complete. And [PW17] showed that solving LP relaxations of many NP-hard combinatorial optimization problems exactly is as hard as solving general LPs exactly.

In practice, however, approximate algorithms are more widely used. In fine-grained complexity analysis, they typically use the same reduction frameworks as in the exact setting, but must also account for *error analysis*: tracking the error propagation when mapping solutions to the reduced problem (problem B) back to that of the original problem (problem A). For instance, Kyng and Zhang [KZ17] not only provided reduction algorithms from general linear equations to certain structured systems, but also conducted error analysis to show that errors remain polynomially bounded when mapping solutions back to the original system. Moreover, they explicitly bounded the condition number of the new matrix with respect to that of the original one.

Approximate algorithms usually have runtimes dependent on error parameters. Error analysis in fine-grained complexity analysis can sometimes help to set a separation between low-accuracy and high-accuracy algorithms. A low-accuracy algorithm’s run-

time is polynomial in $1/\epsilon$, whereas high-accuracy algorithms depend polylogarithmically on $1/\epsilon$. Musco et al. [Mus+19] developed a low-accuracy algorithm for approximating several spectral sum problems, breaking the $O(n^\omega)$ barrier. On the hardness side, they showed that achieving milder error dependencies would imply breakthroughs on the triangle detection algorithms for general graphs running in faster than the state-of-the-art matrix multiplication time, indicating that fast high-accuracy spectrum approximation algorithms are unlikely without major algorithmic breakthroughs. With a similar spirit, Kyng, Wang, and Zhang [KWZ20] studied whether the error dependence in packing LP solvers can be improved, beyond the known $\tilde{O}(N/\epsilon)$ bound (where N is the number of non-zero entries). By creating a fine-grained reduction from solving dense or sparse systems of linear equations to such packing LPs, the paper shows that faster ϵ -dependence in packing LP solvers would contradict known hardness assumptions for systems of linear equations.

1.4 Our Results

Hardness Results for Two-Commodity Flow. We explore the hardness of *2-commodity maximum throughput flow*, which for brevity we refer to as the 2-commodity flow problem or 2CF. Given a directed graph with edge capacities and two source-sink pairs, this problem requires us to maximize the sum of the flows routed between the two source-sink pairs, while satisfying capacity constraints and flow conservation at the remaining nodes. We relate the difficulty of 2CF to that of LP by developing a highly efficient reduction from the former to the latter. The result presented below is joint work with Kyng and Zhang [DKZ22].

Theorem 1.4.1 (Hardness for 2CF (Informal)). *Consider any polynomially bounded linear program $\max\{\mathbf{c}^\top \mathbf{x} : \mathbf{Ax} \leq \mathbf{b}, \mathbf{x} \geq \mathbf{0}\}$ with integer coefficients and N non-zero entries. In nearly-linear time, this linear program can be transformed into a 2-commodity flow problem that is feasible if and only if the original linear program is feasible. The resulting 2-commodity flow instance has $\tilde{O}(N)$ edges and polynomially bounded integral edge capacities. Furthermore, any solution to the 2-commodity flow instance with at most ϵ additive error on each constraint and value at most ϵ from the optimum can be converted into a solution to the original linear program with additive error $\tilde{O}(\text{poly}(N)\epsilon)$ on each constraint and similarly value within $\tilde{O}(\text{poly}(N)\epsilon)$ of the optimum.*

This implies that, for any constant $c > 1$, if any 2-commodity flow instance with polynomially bounded integer capacities can be solved with ϵ additive error in time $\tilde{O}(|E|^c \cdot \text{poly} \log(1/\epsilon))$, then any polynomially bounded linear program can be solved with ϵ additive error in time $\tilde{O}(N^c \cdot \text{poly} \log(1/\epsilon))$.

We say a problem is *sparse-linear-program complete* if any polynomially bounded linear program with N non-zero entries can be reduced, in $\tilde{O}(N)$ time, to an instance of the problem such that: (1) The size of the resulting instance increases by at most a polylogarithmic factor in N ; (2) Any solution to the reduced problem with additive error ϵ can be efficiently mapped into a solution of the original linear program with additive error $\text{poly}(N)\epsilon$. With this terminology, the hardness result for 2CF can be concisely stated as: 2CF is sparse-linear-program complete.

The formal definitions of LP and 2CF are provided later in Chapter 3 (in Definitions 3.3.1 and 3.3.7). It is important to note that, under our definitions, any 2CF problem

is already an LP, and so no reduction in the other direction is necessary. Additionally, our definitions of approximate solutions for LPs (Definition 3.3.4) and 2CF problems (Definition 3.3.16) ensure compatibility: solving a 2CF problem as an LP and obtaining an approximate solution guarantees that the 2CF problem is also approximately solved according to our definition of approximate solutions for 2CF.

Our proof follows the outline of Itai’s polynomial-time reduction of a linear program to a 2-commodity flow problem [Ita78], which showed that *exactly* solving 2-commodity flow and *exactly* solving linear programming are polynomial-time equivalent. We improve Itai’s reduction to nearly preserve the problem representation size in each step. In addition, we establish an error bound for approximately solving each intermediate problem in the reduction, and show that the accumulated error is polynomially bounded. We remark that our hardness results only rule out possibilities of fast multi-commodity flow algorithms for *sparse directed* graphs in the *high accuracy* regime. Outside of this setting, efficient multi-commodity flow solvers may still exist [She17; Mad10; DZ23a]. We also remark that our reduction does not run in *strongly polynomial* time, and it remains an open question whether 2-commodity flow and linear programming are equivalent in the strongly polynomial time regime.

Hardness Results for Combinatorial Laplacians. We study linear equations in combinatorial Laplacians of k -dimensional simplicial complexes (k -complexes), a natural generalization of graph Laplacians. It is known that nearly-linear time solvers exist for graph Laplacians. However, nearly-linear time solvers for combinatorial Laplacians are only known for restricted classes of complexes. Our findings demonstrate that solving linear equations in combinatorial Laplacians of 2-complexes is as computationally hard as solving general linear equations. The following results are from the paper [Din+22] coauthored with Kyng, Probst Gutenberg, and Zhang.

Recall that approximately solving linear equation $\mathbf{A}\mathbf{x} = \mathbf{b}$ is to find $\tilde{\mathbf{x}}$ such that $\|\mathbf{A}\tilde{\mathbf{x}} - \mathbf{b}\|_2 \leq \epsilon \|\mathbf{b}\|_2$ for some ϵ . To formalize our result, we introduce the notion of *sparse-linear-equation complete*. Consider a family of matrices \mathcal{B} . We say that \mathcal{B} is sparse-linear-equation complete if, for any instance $(\mathbf{A}, \mathbf{b}, \epsilon)$ we can construct an instance $(\mathbf{B} \in \mathcal{B}, \mathbf{c}, \delta)$ in $\tilde{O}(\text{nnz}(\mathbf{A}))$ time such that: (1) $\text{nnz}(\mathbf{B}) = \tilde{O}(\text{nnz}(\mathbf{A}))$; (2) $\delta = \epsilon / \text{poly}(\text{nnz}(\mathbf{A}))$, and any approximate solution \mathbf{y} satisfying $\|\mathbf{B}\mathbf{y} - \mathbf{c}\|_2 \leq \delta \|\mathbf{c}\|_2$ can be mapped back to an approximate solution $\tilde{\mathbf{x}}$ to $\mathbf{A}\mathbf{x} = \mathbf{b}$ with error at most ϵ .

In this terminology, we succinctly state the main theorems derived in this study.

Theorem 1.4.2 (Hardness for Combinatorial Laplacians (Informal)). *Linear equations in combinatorial Laplacians of 2-complexes are sparse-linear-equation complete.*

The result is established through a three-step reduction process:

1. **Reduction to Difference-Average Equations:** In the first step, we show that a simple class of linear equations, which we term *difference-average equations*, is sparse-linear-equation complete. These equations take one of two forms:

$$\mathbf{x}(i) - \mathbf{x}(j) = b \quad \text{or} \quad \mathbf{x}(i) + \mathbf{x}(j) = 2\mathbf{x}(k).$$

This reduction was implicitly demonstrated in [KZ17] as an intermediate step, and leads to the following result:

Theorem 1.4.3 (Hardness for Difference-Average Matrices (Informal)). *Linear equations in the difference-average matrices are sparse-linear-equation complete.*

For completeness, we provide an explicit and simplified proof in Section 5.1.

2. **Reduction to Boundary Operators of 2-Complexes:** Next, we reduce the problem of solving difference-average equations to solving linear equations in the boundary operator ∂_2 of a 2-complex. This leads to the following result:

Theorem 1.4.4 (Hardness for Boundary Operators (Informal)). *Linear equations in the boundary operators ∂_2 of 2-complexes are sparse-linear-equation complete.*

Solving $\partial_2 \mathbf{f} = \mathbf{d}$ can be interpreted as computing a flow \mathbf{f} in the triangle space of a 2-complex subject to pre-specified edge demands \mathbf{d} . Our reduction is inspired by a reduction in [MN21] that proves NP-hardness of computing maximum *integral* flows in 2-complexes via a reduction from graph 3-coloring problem. An important aspect of our contribution is that we carefully control the number of non-zeros of the boundary operator matrix that we construct, and we bound the condition number of this matrix and how error propagates from an approximate solution to the boundary operator problem back to the original difference-average equations. Details are provided in Section 4.5 and 4.6.

3. **Reduction to Combinatorial Laplacians of 2-Complexes:** In the final step, we reduce the problem of solving in a boundary operator ∂_2 to solving in the corresponding combinatorial Laplacian \mathbf{L}_1 . Consequently, the combinatorial Laplacian problem is at least as hard as the boundary operator problem. This is captured in the following lemma:

Lemma 1.4.5 (Reduction from Boundary Operators to Combinatorial Laplacians (Informal)). *Suppose we can solve linear equations in combinatorial Laplacians of 2-complexes to high accuracy in nearly-linear time. Then, we can solve linear equations in boundary operators ∂_2 of 2-complexes to high accuracy in nearly-linear time.*

The proof of this lemma relies on standard arguments in linear algebra. A formal statement of the theorem, along with a detailed proof, is provided in Appendix B.1.

Again, we remark that our hardness results rule out the possibility of fast combinatorial Laplacian solvers for sparse simplicial complexes in the high-accuracy regime. However, efficient solvers may still exist outside this setting [Sav+25].

More Hardness Results. Motivated by the hardness results for two-commodity flow and combinatorial Laplacians, we aim to broaden the classes of sparse-linear-program complete and sparse-linear-equation complete problems, in a way analogous to the well-established theory of NP-completeness. However, we later found that our ideas align with those of Průša and Werner [PW17], who proved that LP relaxations of some NP-hard combinatorial optimization problems are sparse-LP complete. While the results presented here remain unpublished, we include them for completeness—particularly since our derivation techniques differ from those used in [PW17].

Our approach begins with *difference-average* equations, which we reduce to a simpler form that we call *1-or-3 equations*. In this form, each row has either one or three non-zero entries, taking one of two shapes:

$$\mathbf{x}(i) + \mathbf{x}(j) = \mathbf{x}(q) \quad \text{or} \quad \mathbf{x}(i) = 1.$$

We extend this reduction from linear equalities (LE) to linear programs (LP), obtaining (*simplified*) *1-or-3 linear programs* of the form

$$\{\mathbf{A}\mathbf{x} = \mathbf{1}, \quad \mathbf{x} \geq \mathbf{0}\},$$

where each equality constraint is either

$$\mathbf{x}(i) + \mathbf{x}(j) + \mathbf{x}(q) = 1 \quad \text{or} \quad \mathbf{x}(i) = 1.$$

From this simplified LP form, one can construct relaxations of various NP-hard combinatorial optimization problems—such as Set Cover, Set Packing, Maximum Satisfiability, and Maximum Independent Set. For additional details on these constructions, we refer readers to [PW17]. These constructions highlight how fundamental the 1-or-3 form is to understanding the complexity of other structured linear equations and linear programs.

Efficient Algorithm for 1-Laplacian Solvers. Previously, nearly-linear time approximate solvers for 1-Laplacian were developed for simplicial complexes with *known collapsing sequences* and *bounded Betti numbers*, such as those triangulating a three-ball in \mathbb{R}^3 [Coh+14a; Bla+22; BN22]. Additionally, Nested Dissection provides quadratic-time exact solvers for more general systems where the non-zero structures correspond to well-shaped simplicial complexes embedded in \mathbb{R}^3 . We generalize the specialized solvers for 1-Laplacians to simplicial complexes with additional geometric structures but *without collapsing sequences and bounded Betti numbers*. Furthermore, we improve the runtime of Nested Dissection for these broader settings. The results below are from the paper [DZ23b], coauthored with Zhang.

We focus on simplicial complexes possessing the following two geometric structures:

1. Each simplex has a bounded *aspect ratio*,⁶ referred to as being “*stable*”.
2. The complex can be partitioned into disjoint, balanced regions with well-shaped *interiors* and *boundaries*, a property we call “*r-hollowing*”.

Our first result applies to *pure* 3-complex⁷ with a given *r-hollowing*.

Theorem 1.4.6 (1-Laplacian Solver for a Pure 3-Complex (Informal)). *Let \mathcal{K} be a pure 3-complex embedded in \mathbb{R}^3 and composed of n stable simplexes. Given an r -hollowing for \mathcal{K} , for any $\epsilon > 0$, we can approximately solve a system in the 1-Laplacian of \mathcal{K} within error ϵ in time:*

$$O(nr + n^{4/3}r^{5/18}\log(n/\epsilon) + n^2r^{-2/3}).$$

⁶The aspect ratio of a geometric shape S is the radius of the smallest ball containing S divided by the radius of the largest ball contained in S .

⁷A simplicial complex is *pure* if every maximal simplex (i.e., a simplex that is not a proper subset of any other simplex in the complex) has the same dimension. For example, a pure 3-complex is a tetrahedron mesh that consists of tetrahedra and their subsimplexes.

In particular, the runtime is minimized (up to constant factors) when $r = \Theta(n^{3/5})$, resulting in:

$$O(n^{8/5} \log(n/\epsilon)).$$

In Section 9 of [DZ23b], sufficient conditions are established for 3-complexes that enable the computation of an r -hollowing in nearly-linear time. We omit this part from the thesis.

We remark that the runtime in Theorem 1.4.6 does not depend on the Betti numbers of \mathcal{K} and does not require collapsing sequences. When $r = o(n)$ and $r = \omega(1)$, the runtime is $o(n^2)$, asymptotically faster than Nested Dissection [MT90]. The solver in [BN22] for a 1-Laplacian system on \mathcal{K} , as stated in Theorem 1.4.6, has a runtime of $\tilde{O}(\beta^3 m)$, where m is the number of simplexes in $\mathcal{X} \supset \mathcal{K}$ with a known collapsing sequence and β is the first Betti number of \mathcal{K} . In the worst-case scenario, m can be as large as $\Omega(n^2)$. However, [BN22] does not require the simplexes to be stable or \mathcal{K} to have a known r -hollowing.

Next, we examine unions of pure 3-complexes glued together by identifying subsets of simplexes on the boundary components (called *exterior simplexes*). Each 3-complex chunk has a $\Theta(n_i^{3/5})$ -hollowing, where n_i is the number of simplexes in the i th chunk. We remark that such a union of 3-complexes, called \mathcal{U} , *may not be embeddable in \mathbb{R}^3* . As a result, previously established methods, including those from [Coh+14a; Bla+22; BN22] and Nested Dissection, cannot be directly applied in this scenario. Building on our algorithm from Theorem 1.4.6, we develop an efficient solver for \mathcal{U} with a runtime that scales sub-quadratically with the size of \mathcal{U} and polynomially with the number of chunks and the number of shared exterior simplexes by more than one chunk.

Theorem 1.4.7 (1-Laplacian Solver for a Union of Pure 3-Complexes (Informal)). *Let \mathcal{U} be a union of pure 3-complexes that are glued together by identifying certain subsets of their exterior simplexes. Assume that each 3-complex chunk is embedded in \mathbb{R}^3 , contains n_i stable simplexes, and has a known $\Theta(n_i^{3/5})$ -hollowing. For any $\epsilon > 0$, we can solve a system in the 1-Laplacian of \mathcal{U} within error ϵ in time*

$$\tilde{O}(n^{8/5} + n^{3/10}k^2 + k^3),$$

where n is the number of simplexes in \mathcal{U} and k is the number of exterior simplexes shared by more than one complex chunk.

When $h = \tilde{O}(1)$ and $k = \tilde{O}(n^{1/2})$, the runtime of the solver in Theorem 1.4.7 matches that of Theorem 1.4.6. Moreover, when $h = o(n^{2/5})$, $k = o(n^{3/5})$, the runtime is $o(n^2)$, asymptotically faster than Nested Dissection.

1.5 Organization of Thesis

In Chapter 2, we introduce the notations and foundational concepts from linear algebra and topology that will be used throughout the thesis. In Chapter 3, we establish the hardness of solving two-commodity flow problems, proving Theorem 1.4.1. Chapter 4 focuses on the hardness of solving linear equations in combinatorial Laplacians of simplicial complexes and proves Theorem 1.4.2. Chapter 5 presents additional hardness results, including reduction algorithms for difference-average equations, 1-or-3 equations,

and 1-or-3 linear programs. In Chapter 6, we present an efficient algorithm for solving 1-Laplacian systems, proving Theorems 1.4.6 and 1.4.7. In Appendix A, B, and C, we include missing proofs of Chapter 2, 4, and 6, respectively.

Chapter 2

Preliminaries

2.1 Vectors and Matrices

Indexing. Given a vector $\mathbf{x} \in \mathbb{R}^n$, for $1 \leq i \leq n$, we let $\mathbf{x}(i)$ be the i th entry of \mathbf{x} ; for $1 \leq i < j \leq n$, let $\mathbf{x}(i:j)$ be $[\mathbf{x}(i) \ \mathbf{x}(i+1) \ \dots \ \mathbf{x}(j)]^\top$. We use $\mathbf{1}_n, \mathbf{0}_n$ to denote an n -dimensional all-one vector and all-zero vector, respectively.

Given a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, for $1 \leq i \leq m, 1 \leq j \leq n$, we let $\mathbf{A}(i, j)$ be the (i, j) th entry of \mathbf{A} ; for $S_1 \subseteq \{1, \dots, m\}, S_2 \subseteq \{1, \dots, n\}$, let $\mathbf{A}(S_1, S_2)$ be the submatrix with row indices in S_1 and column indices in S_2 . Furthermore, we let $\mathbf{A}(S_1, :) = \mathbf{A}(S_1, \{1, \dots, n\})$ and $\mathbf{A}(:, S_2) = \mathbf{A}(\{1, \dots, m\}, S_2)$. In particular, we use $\mathbf{A}(i)$ to denote the i th row of \mathbf{A} . In addition, we use $\text{nnz}(\mathbf{A})$ to denote the number of non-zero entries of \mathbf{A} . Without loss of generality, we assume that $\text{nnz}(\mathbf{A}) \geq \max\{m, n\}$.

Subspaces. The image of \mathbf{A} is the linear span of the columns of \mathbf{A} , denoted by $\text{Im}(\mathbf{A})$, and the kernel of \mathbf{A} to be $\{\mathbf{x} \in \mathbb{R}^n : \mathbf{A}\mathbf{x} = \mathbf{0}\}$, denoted by $\text{Ker}(\mathbf{A})$. A fundamental theorem of Linear Algebra states

$$\mathbb{R}^n = \text{Im}(\mathbf{A}) \oplus \text{Ker}(\mathbf{A}^\top).$$

Fact 2.1.1. For any matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\text{Im}(\mathbf{A}) = \text{Im}(\mathbf{A}\mathbf{A}^\top)$.

All the facts in this section are well-known. For completeness, we include their proofs in Appendix A.

Pseudo-Inverse and Projection Matrix. The pseudo-inverse of \mathbf{A} is defined to be a matrix \mathbf{A}^\dagger that satisfies all the following four criteria: (1) $\mathbf{A}\mathbf{A}^\dagger\mathbf{A} = \mathbf{A}$, (2) $\mathbf{A}^\dagger\mathbf{A}\mathbf{A}^\dagger = \mathbf{A}^\dagger$, (3) $(\mathbf{A}\mathbf{A}^\dagger)^\top = \mathbf{A}\mathbf{A}^\dagger$, (4) $(\mathbf{A}^\dagger\mathbf{A})^\top = \mathbf{A}^\dagger\mathbf{A}$. The orthogonal projection matrix onto $\text{Im}(\mathbf{A})$ is

$$\Pi_{\text{Im}(\mathbf{A})} = \mathbf{A}(\mathbf{A}^\top\mathbf{A})^\dagger\mathbf{A}^\top. \quad (2.1)$$

Eigenvalues and Condition Numbers. Given a square matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, let $\lambda_{\max}(\mathbf{A})$ be the maximum eigenvalue of \mathbf{A} and $\lambda_{\min}(\mathbf{A})$ the minimum *non-zero* eigenvalue of \mathbf{A} . The condition number of \mathbf{A} is defined as the ratio of these two values:

$$\kappa(\mathbf{A}) = \frac{\lambda_{\max}(\mathbf{A})}{\lambda_{\min}(\mathbf{A})}.$$

A symmetric matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ is called *positive semi-definite (PSD)* if and only if $\mathbf{x}^\top \mathbf{A} \mathbf{x} \geq 0$ for all $\mathbf{x} \in \mathbb{R}^n$. Equivalently, \mathbf{A} is PSD if and only if all its eigenvalues of \mathbf{A} are non-negative.

For a non-square matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, the singular values of \mathbf{A} are defined as the square roots of the eigenvalues of $\mathbf{A}^\top \mathbf{A}$, which is always symmetric and PSD. Let $\sigma_{\max}(\mathbf{A})$ denote the maximum singular value of \mathbf{A} and $\sigma_{\min}(\mathbf{A})$ the minimum *non-zero* singular value of \mathbf{A} . If \mathbf{A} is a symmetric PSD matrix, then its eigenvalues and singular values coincide.

For two symmetric matrices $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times n}$, we say $\mathbf{A} \succcurlyeq \mathbf{B}$ if $\mathbf{A} - \mathbf{B}$ is PSD. The *condition number of \mathbf{A} relative to \mathbf{B}* is defined as:

$$\kappa(\mathbf{A}, \mathbf{B}) \stackrel{\text{def}}{=} \min \left\{ \frac{\alpha}{\beta} : \beta \cdot \mathbf{\Pi}_{\text{Im}(\mathbf{A})} \mathbf{B} \mathbf{\Pi}_{\text{Im}(\mathbf{A})} \preccurlyeq \mathbf{A} \preccurlyeq \alpha \cdot \mathbf{B} \right\},$$

where $\mathbf{\Pi}_{\text{Im}(\mathbf{A})}$ is the projection matrix as defined in Equation (2.1).

Fact 2.1.2. Let $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times n}$ be symmetric matrices such that $\mathbf{A} \preccurlyeq \mathbf{B}$. Then, for any $\mathbf{V} \in \mathbb{R}^{m \times n}$,

$$\mathbf{V} \mathbf{A} \mathbf{V}^\top \preccurlyeq \mathbf{V} \mathbf{B} \mathbf{V}^\top.$$

Norms. Given a vector $\mathbf{x} \in \mathbb{R}^n$, let¹

$$\|\mathbf{x}\|_1 = \sum_{i \in [n]} |\mathbf{x}(i)|, \quad \|\mathbf{x}\|_\infty = \max_{i \in [n]} |\mathbf{x}(i)|, \quad \|\mathbf{x}\|_2 = \sqrt{\sum_{i=1}^n \mathbf{x}(i)^2},$$

where $\|\mathbf{x}\|_2$ is called the *Euclidean norm*. When $\mathbf{M} \in \mathbb{R}^{n \times n}$ is a PSD, we also define

$$\|\mathbf{x}\|_{\mathbf{M}} = \sqrt{\mathbf{x}^\top \mathbf{M} \mathbf{x}}.$$

Next, for a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, we define its operator norm induced by a vector p -norm as

$$\|\mathbf{A}\|_p \stackrel{\text{def}}{=} \max_{\mathbf{x} \neq \mathbf{0}} \frac{\|\mathbf{A} \mathbf{x}\|_p}{\|\mathbf{x}\|_p}.$$

In particular, for $p = 1, \infty, 2$, we have

$$\|\mathbf{A}\|_1 = \max_{j \in [n]} \sum_{i \in [m]} |\mathbf{A}(i, j)|, \quad \|\mathbf{A}\|_\infty = \max_{i \in [m]} \sum_{j \in [n]} |\mathbf{A}(i, j)|, \quad \|\mathbf{A}\|_2 = \sqrt{\lambda_{\max}(\mathbf{A}^\top \mathbf{A})} = \sigma_{\max}(\mathbf{A}).$$

Moreover, we use $\|\mathbf{A}\|_{\max}$ to denote the maximum absolute entry of \mathbf{A} , i.e.,

$$\|\mathbf{A}\|_{\max} = \max_{i, j} |\mathbf{A}(i, j)|.$$

We define a function X that takes an arbitrary number of matrices $\mathbf{A}_1, \dots, \mathbf{A}_{k_1}$, vectors $\mathbf{b}_1, \dots, \mathbf{b}_{k_2}$, and scalars K_1, \dots, K_{k_3} as arguments, and returns the maximum of the entries of all the arguments, i.e.,

$$\begin{aligned} X(\mathbf{A}_1, \dots, \mathbf{A}_{k_1}, \mathbf{b}_1, \dots, \mathbf{b}_{k_2}, K_1, \dots, K_{k_3}) \\ = \max \{ \|\mathbf{A}_1\|_{\max}, \dots, \|\mathbf{A}_{k_1}\|_{\max}, \|\mathbf{b}_1\|_\infty, \dots, \|\mathbf{b}_{k_2}\|_\infty, |K_1|, \dots, |K_{k_3}| \}. \end{aligned}$$

¹We sometimes use $\|\mathbf{x}\|_{\max}$ for $\|\mathbf{x}\|_\infty$ interchangeably.

2.2 Simplicial Homology

Simplicial Complexes. A k -dimensional simplex (or k -simplex) $\sigma = \text{conv}\{v_0, \dots, v_k\}$ is the convex hull of $k + 1$ affinely independent points v_0, \dots, v_k . For example, 0, 1, 2-simplexes are vertices, edges, and triangles, respectively. A *face* of σ is the convex hull of a non-empty subset of $\{v_0, v_1, \dots, v_k\}$. An *orientation* of σ is given by an ordering Π of its vertices, written as $\sigma = [v_{\Pi(0)}, \dots, v_{\Pi(k)}]$, such that two orderings define the same orientation if and only if they differ by an even permutation. If Π is even, then $[v_{\Pi(0)}, \dots, v_{\Pi(k)}] = [v_0, \dots, v_k]$; if Π is odd, then $[v_{\Pi(0)}, \dots, v_{\Pi(k)}] = -[v_0, \dots, v_k]$.

A *simplicial complex* \mathcal{K} is a finite collection of simplexes such that: (1) for every $\sigma \in \mathcal{K}$ if $\tau \subset \sigma$ then $\tau \in \mathcal{K}$; and (2) for every $\sigma_1, \sigma_2 \in \mathcal{K}$, $\sigma_1 \cap \sigma_2$ is either empty or a face of both σ_1, σ_2 . The *dimensions* of \mathcal{K} is the maximum dimension of any simplex in \mathcal{K} . We refer to a simplicial complex in d dimensions as a d -complex. For $1 \leq k \leq d$, the k -skeleton of a d -complex \mathcal{K} is the subcomplex consisting of all the simplexes of \mathcal{K} of dimensions at most k . In particular, the 1-skeleton of \mathcal{K} is a graph. A simplicial complex is *pure* if every maximal simplex (i.e., one not properly contained in any larger simplex in \mathcal{K}) has the same dimension. For example, a tetrahedron (together with all its faces) forms a 3-dimensional pure simplicial complex: the only maximal simplex in it is the tetrahedron itself, which is of dimension 3.

Embedding and Triangulation. A *piecewise linear embedding* of a 3-complex in \mathbb{R}^3 maps a 0-simplex to a point, a 1-simplex to a line segment, a 2-simplex to a triangle, and a 3-simplex to a tetrahedron. In addition, the interiors of the images of simplexes are disjoint and the boundary of each simplex is mapped to the appropriate simplexes. Such an embedding of a simplicial complex \mathcal{K} defines an *underlying topological space* \mathbb{K} – the union of the images of all the simplexes of \mathcal{K} . We say \mathcal{K} is *convex* if \mathbb{K} is convex. We say \mathcal{K} *triangulates* a topological space \mathbb{X} if \mathbb{K} is homeomorphic to \mathbb{X} . A simplex σ of \mathcal{K} is a *exterior* simplex if σ is contained in the boundary of \mathbb{K} . A connected component of exterior simplexes is called a *boundary component* of \mathcal{K} .

More concretely, if \mathbb{K} is a 2-dimensional manifold, it can be triangulated by 2-complexes, where every edge in the 2-complex is contained in exactly one triangle (boundary edge) or two triangles (interior edge). An *oriented triangulation* of a 2-dimensional manifold is a triangulation together with an orientation for each triangle such that any two neighboring triangles induce opposite signs on their shared interior edge. Refer to Figure 2.1 for an example of (oriented) triangulation: the topological space is a disk; boundary edges are $[v_1, v_2], [v_2, v_3], [v_1, v_3]$; interior edges are $[v_1, v_4], [v_2, v_4], [v_3, v_4]$; the orientation for each triangle is clockwise.

Boundary Operators. Given an oriented d -dimensional simplicial complex \mathcal{K} , for each $0 \leq i \leq d$, let \mathcal{C}_i denote i -chain, which is a formal sum of the oriented i -simplexes in \mathcal{K} with the coefficients over \mathbb{R} . We can define a sequence of boundary operators:

$$\mathcal{C}_d(\mathcal{K}) \xrightarrow{\partial_d} \mathcal{C}_{d-1}(\mathcal{K}) \xrightarrow{\partial_{d-1}} \dots \xrightarrow{\partial_2} \mathcal{C}_1(\mathcal{K}) \xrightarrow{\partial_1} \mathcal{C}_0(\mathcal{K}),$$

where each linear map $\partial_k : \mathcal{C}_k(\mathcal{K}) \rightarrow \mathcal{C}_{k-1}(\mathcal{K})$ is a *boundary operator* that maps every k -simplex to a signed sum of its boundary $(k - 1)$ -faces. Specifically, for an oriented

k -simplex $\sigma = [v_0, v_1, \dots, v_k]$,

$$\partial_k(\sigma) = \sum_{i=0}^k (-1)^i [v_0, \dots, \hat{v}_i, \dots, v_k],$$

where $[v_0, \dots, \hat{v}_i, \dots, v_k]$ is the oriented $(k-1)$ -simplex obtained by removing v_i from σ , and $(-1)^i$ is its *induced orientation*. The operator ∂_k can be written as a matrix in $|\mathcal{C}_{k-1}| \times |\mathcal{C}_k|$ dimensions, where $|\mathcal{C}_{k-1}|, |\mathcal{C}_k|$ is the number of $(k-1)$ -simplexes and k -simplexes in \mathcal{K} , respectively. The (i, j) th entry of ∂_k is ± 1 if the i th $(k-1)$ -simplex is a face of the j th k -simplex where the sign is determined by the orientations, and 0 otherwise. See Figure 2.1 and Equation (2.2) for an example.




Figure 2.1: An example of boundary operator and oriented triangulation. We set a clockwise orientation for 2-simplexes, and set the orientation for 1-simplexes as the order of increasing vertex indices.

$$\partial_2 = \begin{matrix} & \begin{matrix} [v_1, v_4, v_2] & [v_2, v_4, v_3] & [v_1, v_3, v_4] \end{matrix} \\ \begin{matrix} [v_1, v_2] \\ [v_2, v_3] \\ [v_1, v_3] \\ [v_1, v_4] \\ [v_2, v_4] \\ [v_3, v_4] \end{matrix} & \begin{bmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \\ 0 & -1 & 1 \end{bmatrix} \end{matrix}. \quad (2.2)$$

An important property of the boundary operator is that applying the boundary operator twice results in the zero operator, i.e.,

$$\partial_k \partial_{k+1} = \mathbf{0}. \quad (2.3)$$

This implies

$$\text{Im}(\partial_{k+1}) \subseteq \text{Ker}(\partial_k).$$

Thus, we can define the *quotient space* $H_k = \text{Ker}(\partial_k) \setminus \text{Im}(\partial_{k+1})$, referred to as the k th *homology space* of \mathcal{K} . The dimension of H_k is the k th *Betti number* of \mathcal{K} , denoted as β_k . In particular, if $\beta_k = 0$, then

$$\text{Im}(\partial_k^\top) \oplus \text{Im}(\partial_{k+1}) = \mathbb{R}^{|\mathcal{C}_k|}.$$

Betti numbers play an important role in understanding the homology spaces. Intuitively, the k th Betti number refers to the number of k -dimensional “holes” on a topological surface. For example, β_0 represents the number of connected components; β_1 represents the number of “circular” holes; β_2 represents the number of “voids” or “cavities”.

Combinatorial Laplacians. Combinatorial Laplacians arise from the discrete Hodge decomposition.

Theorem 2.2.1 (Hodge Decomposition [Lim20]). *Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{B} \in \mathbb{R}^{n \times p}$ be matrices satisfying $\mathbf{AB} = \mathbf{0}$. Then, there is an orthogonal direct sum decomposition*

$$\mathbb{R}^n = \text{Im}(\mathbf{A}^\top) \oplus \text{Ker}(\mathbf{A}^\top \mathbf{A} + \mathbf{B}\mathbf{B}^\top) \oplus \text{Im}(\mathbf{B}).$$

By Equation (2.3), it is valid to set $\mathbf{A} = \partial_k$ and $\mathbf{B} = \partial_{k+1}$. The matrix we get in the middle term is the *combinatorial Laplacian*:

$$\mathbf{L}_k \stackrel{\text{def}}{=} \partial_k^\top \partial_k + \partial_{k+1} \partial_{k+1}^\top.$$

In particular, $\mathbf{L}_0 = \partial_1 \partial_1^\top$ is the graph Laplacian. The k th homology space $H_k(\mathcal{K})$ is isomorphic to $\text{Ker}(\mathbf{L}_k)$, and thus the k th Betti number of \mathcal{K} equals the dimension of $\text{Ker}(\mathbf{L}_k)$.

2.3 Approximately Solving Linear Equations

We define approximately solving linear equations in a general form, following [KZ17]. For more details, we refer the readers to Section 2.1 of [KZ17].

Definition 2.3.1 (Linear Equation Problem (LE)). Given a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, a vector $\mathbf{b} \in \mathbb{R}^m$, we refer to the linear equation problem for the tuple (\mathbf{A}, \mathbf{b}) , denoted by $\text{LE}(\mathbf{A}, \mathbf{b})$, as the problem of finding an $\mathbf{x} \in \mathbb{R}^n$ such that

$$\mathbf{x} \in \arg \min_{\mathbf{x} \in \mathbb{R}^n} \|\mathbf{Ax} - \mathbf{b}\|_2.$$

The exact solution to LE has a closed-form solution, as shown in the following fact.

Fact 2.3.2. *Let $\mathbf{x}^* \in \arg \min_{\mathbf{x} \in \mathbb{R}^n} \|\mathbf{Ax} - \mathbf{b}\|_2$. Then,*

$$\mathbf{Ax}^* = \mathbf{A}(\mathbf{A}^\top \mathbf{A})^\dagger \mathbf{A}^\top \mathbf{b} = \Pi_{\text{Im}(\mathbf{A})} \mathbf{b}$$

and

$$\|\mathbf{Ax}^* - \mathbf{b}\|_2^2 = \|(\mathbf{I} - \Pi_{\text{Im}(\mathbf{A})}) \mathbf{b}\|_2^2,$$

where $\Pi_{\text{Im}(\mathbf{A})}$ is the projection matrix as defined in Equation (2.1).

By the above fact, solving $\text{LE}(\mathbf{A}, \mathbf{b})$ is equivalent to finding an \mathbf{x} such that $\mathbf{Ax} = \Pi_{\text{Im}(\mathbf{A})} \mathbf{b}$. This equation is known as the *normal equation*, and it is always feasible. If $\mathbf{b} \in \text{Im}(\mathbf{A})$, then $\Pi_{\text{Im}(\mathbf{A})} \mathbf{b} = \mathbf{b}$.

In practice, we are more interested in *approximately* solving linear equations, since numerical errors are unavoidable in data collection and computation, and approximate solvers often run faster.

Definition 2.3.3 (Linear Equation Approximation Problem (LEA)). Given a matrix $\mathbf{A} \in \mathbb{R}^{d \times n}$, vectors $\mathbf{b} \in \mathbb{R}^d$, and an error parameter $\epsilon \in (0, 1]$, we refer to linear equation approximate problem for the tuple $(\mathbf{A}, \mathbf{b}, \epsilon)$, denoted by $\text{LEA}(\mathbf{A}, \mathbf{b}, \epsilon)$, as the problem of finding an $\mathbf{x} \in \mathbb{R}^n$ such that

$$\|\mathbf{Ax} - \Pi_{\text{Im}(\mathbf{A})} \mathbf{b}\|_2 \leq \epsilon \|\Pi_{\text{Im}(\mathbf{A})} \mathbf{b}\|_2.$$

The following fact demonstrates that the approximate error in Definition 2.3.3 is equivalent to several error notions that are commonly used in solving linear equations.

Fact 2.3.4.

$$\|\mathbf{A}\mathbf{x} - \Pi_{\text{Im}(\mathbf{A})}\mathbf{b}\|_2 = \|\mathbf{A}^\top \mathbf{A}\mathbf{x} - \mathbf{A}^\top \mathbf{b}\|_{(\mathbf{A}^\top \mathbf{A})^\dagger} = \|\mathbf{x} - \mathbf{x}^*\|_{\mathbf{A}^\top \mathbf{A}}.$$

After introducing the approximate solution problem in Definition 2.3.3, it is important to understand how the approximate solution relates to the exact solution. The following fact establishes a relationship between the error of the approximate solution and the exact solution error.

Fact 2.3.5. *Let \mathbf{x} be a solution to LEA $(\mathbf{A}, \mathbf{b}, \epsilon)$ and \mathbf{x}^* be the exact solution. Then,*

$$\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 \leq \|\mathbf{A}\mathbf{x}^* - \mathbf{b}\|_2^2 + \epsilon^2 \|\Pi_{\text{Im}(\mathbf{A})}\mathbf{b}\|_2^2.$$

Linear Equations in PSD Matrices. Consider the eigen-decomposition of a PSD rank- k matrix $\mathbf{A} = \sum_{1 \leq i \leq k} \lambda_i \mathbf{u}_i \mathbf{u}_i^\top$, where $\lambda_i \geq 0$ are eigenvalues and \mathbf{u}_i are the corresponding orthogonal eigenvectors. Let $\mathbf{A}^\dagger = \sum_{1 \leq i \leq k} \lambda_i^{-1} \mathbf{u}_i \mathbf{u}_i^\top$ be the pseudo-inverse of \mathbf{A} , and we have

$$\mathbf{A}\mathbf{A}^\dagger = \mathbf{A}^\dagger \mathbf{A} = \sum_{1 \leq i \leq k} \mathbf{u}_i \mathbf{u}_i^\top = \Pi_{\text{Im}(\mathbf{A})}.$$

Furthermore, let $\mathbf{A}^{1/2} = \sum_{1 \leq i \leq k} \lambda_i^{1/2} \mathbf{u}_i \mathbf{u}_i^\top$ be the square root \mathbf{A} , and we have $\mathbf{A}^{1/2} \mathbf{A}^{1/2} = \mathbf{A}$. Similarly, let $\mathbf{A}^{\dagger/2} = \sum_{1 \leq i \leq k} \lambda_i^{-1/2} \mathbf{u}_i \mathbf{u}_i^\top$ be square root of \mathbf{A}^\dagger . By Fact 2.1.1, we have $\text{Im}(\mathbf{A}^{1/2}) = \text{Im}(\mathbf{A})$, thus

$$\mathbf{A}^{1/2} \mathbf{A}^\dagger \mathbf{A}^{1/2} = \mathbf{A}^{1/2} \mathbf{A}^{\dagger/2} = \mathbf{A}^{\dagger/2} \mathbf{A}^{1/2} = \Pi_{\text{Im}(\mathbf{A})}. \quad (2.4)$$

It is easy to see that $\mathbf{A}^{1/2}$, \mathbf{A}^\dagger , $\mathbf{A}^{\dagger/2}$ are also PSD.

The following fact provides two useful results about the perturbation of PSD matrices, which will help analyze approximation errors and stability properties in such systems of linear equations.

Fact 2.3.6. *Let $\mathbf{A}, \mathbf{Z} \in \mathbb{R}^{n \times n}$ be two symmetric PSD matrices.*

1. *If $(1 - \epsilon)\mathbf{A}^\dagger \preceq \mathbf{Z} \preceq (1 + \epsilon)\mathbf{A}^\dagger$, then for any $\mathbf{b} \in \text{Im}(\mathbf{A})$,*

$$\|\mathbf{A}\mathbf{Z}\mathbf{b} - \mathbf{b}\|_2 \leq \epsilon \sqrt{\kappa(\mathbf{A})} \|\mathbf{b}\|_2.$$

2. *If $\|\mathbf{A}\mathbf{Z}\mathbf{b} - \mathbf{b}\|_2 \leq \epsilon \|\mathbf{b}\|_2$ for any $\mathbf{b} \in \text{Im}(\mathbf{A})$, then*

$$(1 - \epsilon)\mathbf{A}^\dagger \preceq \Pi_{\text{Im}(\mathbf{A})} \mathbf{Z} \Pi_{\text{Im}(\mathbf{A})} \preceq (1 + \epsilon)\mathbf{A}^\dagger.$$

Schur Complement. Let $\mathbf{A} \in \mathbb{R}^{n \times n}$, and let $F \cup C$ be a partition of $\{1, \dots, n\}$. We write \mathbf{A} as a block matrix:

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}(F, F) & \mathbf{A}(F, C) \\ \mathbf{A}(C, F) & \mathbf{A}(C, C) \end{bmatrix}. \quad (2.5)$$

We define the (*generalized*) *Schur complement of \mathbf{A} onto C* to be

$$\text{Sc}(\mathbf{A})_C = \mathbf{A}(C, C) - \mathbf{A}(C, F)\mathbf{A}(F, F)^\dagger \mathbf{A}(F, C).$$

The Schur complement appears in performing a block Gaussian Elimination on matrix \mathbf{A} to eliminate the indices in F .

Fact 2.3.7. For any fixed vector \mathbf{y} ,

$$\min_{\mathbf{x}} [\mathbf{x}^\top \quad \mathbf{y}^\top] \mathbf{A} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} = \mathbf{y}^\top \text{Sc}(\mathbf{A})_C \mathbf{y}.$$

Using Fact 2.3.7, we can obtain the following fact that Schur complement preserves PSD.

Fact 2.3.8. Let \mathbf{A} be a PSD matrix defined in Equation (2.5). Then, $\text{Sc}(\mathbf{A})_C$ is also PSD.

The following fact provides a useful matrix identity that factorizes the matrix \mathbf{A} in terms of its submatrices and Schur complement.

Fact 2.3.9. Let \mathbf{A} be a PSD matrix defined in Equation (2.5). Then,

$$\mathbf{A} = \begin{bmatrix} \mathbf{I} & \\ \mathbf{A}(C, F)\mathbf{A}(F, F)^\dagger & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{A}(F, F) & \\ & \text{Sc}(\mathbf{A})_C \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{A}(F, F)^\dagger \mathbf{A}(F, C) \\ & \mathbf{I} \end{bmatrix}.$$

In the special case where $\mathbf{A} = \mathbf{B}\mathbf{B}^\top$, the Schur complement $\text{Sc}(\mathbf{A})_C$ can be expressed directly in terms of \mathbf{B} .

Fact 2.3.10. Let \mathbf{A} be a PSD matrix defined in Equation (2.5). Let $\mathbf{A} = \mathbf{B}\mathbf{B}^\top$, and we decompose $\mathbf{B} = \begin{bmatrix} \mathbf{B}_F \\ \mathbf{B}_C \end{bmatrix}$ accordingly. Then,

$$\text{Sc}(\mathbf{A})_C = \mathbf{B}_C \mathbf{\Pi}_{\text{Ker}(\mathbf{B}_F)} \mathbf{B}_C^\top,$$

where $\mathbf{\Pi}_{\text{Ker}(\mathbf{B}_F)}$ is the projection onto the kernel of \mathbf{B}_F : $\mathbf{\Pi}_{\text{Ker}(\mathbf{B}_F)} = \mathbf{I} - \mathbf{B}_F^\top (\mathbf{B}_F \mathbf{B}_F^\top)^\dagger \mathbf{B}_F$.

Chapter 3

Hardness Results for Two-Commodity Flow

This chapter is based on [DKZ22] and focuses on proving Theorem 1.4.1. Here in Section 3.5 and 3.6, we improve the proof structure from the published version [DKZ22], making it modular and streamlined. Furthermore, we introduce a unified linear programming (LP) transformation framework in Section 3.7. This framework provides a unified characterization and a systematic analysis of various reduction steps.

3.1 Prior Works

Previous work by Kyng and Zhang [KZ17] had shown that fast algorithms for multi-commodity flow were unlikely to arise from combining interior point methods with special-purpose linear equation solvers. Concretely, they showed that the linear equations that arise in interior point methods for multi-commodity flow are as hard to solve as arbitrary linear equations. This ruled out algorithms following the pattern of the known fast algorithms for high-accuracy single-commodity flow problems. However, it left open the broader question if some other families of algorithms could succeed. We now show that, in general, a separation between multi-commodity flow and linear programming is not possible.

Our paper follows the outline of Itai’s polynomial-time reduction of a linear program to a 2-commodity flow problem. In addition, it is also inspired by recent works on hardness for structured linear equations [KZ17] and packing/covering LPs [KWZ20], which focused on obtaining nearly-linear time reductions in related settings. These works in turn were motivated by the last decade’s substantial progress on fine-grained complexity for a range of polynomial time solvable problems, e.g. see [WW18]. Also notable is the result by Musco et al. [Mus+19] on hardness for matrix spectrum approximation.

3.2 Our Contributions

The many successes in developing high-accuracy algorithms for single-commodity flow problems highlight an important open question: Can multi-commodity flow be solved to high accuracy faster than general linear programs? We rule out this possibility by proving that any linear program (assuming it is polynomially bounded and has integer entries) can be encoded as a multi-commodity flow problem in nearly-linear time. This implies

that any improvement in the running time of (high-accuracy) algorithms for sparse multi-commodity flow problems would directly translate to a faster algorithm for solving sparse linear programs to high accuracy, with only a polylogarithmic increase in running time.

Theorem 1.4.1 demonstrates several key improvements over Itai’s original polynomial time reduction from LP to 2CF. Firstly, while Itai produced a 2CF with the number of edges on the order of $\Theta(N^2 \log^2 X)$, where N is the number of non-zeros and X is the magnitude of the largest entry in the original LP, we show that an improved gadget can reduce this to $O(N \log X)$. Thus, in the case of polynomially bounded linear programs, where $\log X = O(\log N)$, we get only a polylogarithmic multiplicative increase in the number of non-zero entries from N to $\tilde{O}(N)$, whereas Itai had an increase in the number of non-zeros by a factor $\tilde{O}(N)$, i.e., from N to $\tilde{O}(N^2)$.

Secondly, Itai used very large graph edge capacities that require $O((N \log X)^{1.01})$ many bits *per edge*, letting the capacities grow exponentially given an LP with polynomially bounded entries. We show that when the feasible polytope radius R is bounded, we can ensure capacities remain a polynomial function of the initial parameters N, R , and X . In the important case of polynomially bounded linear programs, this means the capacities stay polynomially bounded.

Thirdly, while Itai only analyzed the chain of reductions under the case with exact solutions, we generalize the analysis to the case with approximate solutions by establishing an error analysis along the chain of reductions. We show that the error only grows polynomially during the reduction.

Finally, we introduce a framework that unifies all steps in the reduction chain. Rather than analyzing each reduction individually, this framework offers a unified characterization of problem reductions and enables a systematic approach to analyzing the correctness and error propagation of these reductions. It serves not only as a powerful tool for validating existing reductions but also as a valuable guide for designing and verifying new reduction steps within the LP realm.

3.3 Problem Definitions

This section introduces the intermediate problems that form our nine-step chain of reduction algorithms. We group these problems into two categories for clarity: the *Algebra Space*, which includes linear programming and linear equation problems, and the *Flow Space*, which involves flow routing in graphs. We also define the error notions used to evaluate approximate solutions in both domains. Before delving into the formal problem definitions, we begin by discussing the key assumptions involved.

3.3.1 Discussion of the Problem Assumptions

Polynomially-Bounded LP. Current research on fast algorithms for solving linear programs generally relies on assuming bounds on: (1) the size of the program entries, and (2) the norm of all feasible solutions. Generally, the algorithm running time depends logarithmically on these quantities, and hence to make these factors negligible, entry size and feasible solution norms are assumed to be polynomially bounded, for example in [CLS21; Jia+21]. We will refer to a linear program satisfying these assumptions as *polynomially bounded*. More precisely, we say a linear program with N non-zero coefficients is *polynomially bounded* if it has coefficients in the range $[-X, X]$ and $\|\mathbf{x}\|_1 \leq R$

for all feasible \mathbf{x} (i.e., the polytope of feasible solutions has radius of R in ℓ_1 norm), and $X, R \leq O(N^c)$ for some constant c . In fact, if there exists a feasible solution \mathbf{x} satisfying $\|\mathbf{x}\|_1 \leq R$, then we can add a constraint $\|\mathbf{x}\|_1 \leq R$ to the LP (which can be rewritten as linear inequality constraints) so that in the new LP, all feasible solutions have ℓ_1 norm at most R . This only increases the number of non-zeros in the LP by at most a constant factor.

Interior Point Methods and Reductions with Fixed Point Arithmetic. Modern fast interior point methods for linear programming, such as [CLS21], are analyzed in the RealRAM model. In order to implement these algorithms using *fixed point arithmetic* with polylogarithmic bit complexity per number, instead of RealRAM, additional assumptions are required. For example, this class of algorithms relies on computing matrix inverses, and these must be approximately representable using polylogarithmic bit complexity per entry. This is not possible, if the inverses have exponentially large entries, which may occur even in polynomially bounded linear programs. For example, consider a linear program feasibility problem $\{\mathbf{Ax} \leq \mathbf{b}, \mathbf{x} \geq \mathbf{0}\}$, with constraint matrix $\mathbf{A} \in \mathbb{R}^{2n \times 2n}$ given by

$$\mathbf{A}(i, j) = \begin{cases} 1 & \text{if } i < n \text{ and } i = j \\ -2 & \text{if } i < n \text{ and } i + 1 = j \\ 2 & \text{if } i \geq n \text{ and } i = j \\ -1 & \text{if } i \geq n \text{ and } i + 1 = j \\ 0 & \text{o.w.} \end{cases}.$$

Such a linear program is polynomially bounded for many choices of \mathbf{b} , e.g., $\mathbf{b} = \mathbf{e}_{2n}$. Unfortunately, for the vector $\mathbf{x} \in \mathbb{R}^{2n}$ given by

$$\mathbf{x}(i) = \begin{cases} 2^{-i-1} & \text{if } i \leq n \\ 0 & \text{o.w.} \end{cases},$$

we have $\mathbf{Ax} = 2^{-n} \mathbf{e}_n$, and from this one can see that \mathbf{A}^{-1} must have entries of size at least $\Omega(2^n/n)$. This will cause algorithms such as [CLS21] to perform intermediate calculations with n bits per number, increasing the running time by a factor roughly n .

Modern interior point methods can be translated to fixed precision arithmetic with various different assumptions leading to different per entry bit complexity (see [CLS21] for a discussion of one standard sufficient condition). Furthermore, if the problem has polynomially bounded condition number (when appropriately defined), then we expect that polylogarithmic bit complexity per entry should suffice, at least for highly accurate approximate solutions, by relying on fast stable numerical linear algebra [DDH07], although we are not aware of a complete analysis of this translation.

If a linear program with integer entries is solved to sufficiently small additive error, the approximate solution can be converted into an exact solution, e.g. see [Ren88; LS14; CLS21] for a discussion of the necessary precision and for a further discussion of numerical stability properties of interior point methods. Polynomially bounded linear programs with integer coefficients may still require exponentially small additive error for this rounding to succeed.

We give a reduction from general linear programming to 2-commodity flow, and like [CLS21], we assume the program is polynomially bounded to carry out the reduction. We also use an additional assumption, namely that the linear program is written using

integral entries.¹ We do not make additional assumptions about polynomially bounded condition number of the problem. This means we can apply our reduction to programs such as the one above, despite [CLS21] not obtaining a reasonable running time on such programs using fixed point arithmetic.

Our analysis of our reduction uses the RealRAM model like [CLS21; Jia+21] and other modern interior point method analysis, however, it should be straightforward to translate our reduction and error analysis to fixed point arithmetic with polylogarithmically many bits, because all our mappings are simple linear transformations, and we never need to compute or apply a matrix inverse.

Rounding Linear Programs to Have Integer Entries. It is possible to give some fairly general and natural sufficient conditions for when a polynomially bounded linear program can be rounded to have integral entries, one example of this is having a polynomially bounded *Renegar's condition number*. Renegar introduced this condition number for linear programs in [Ren95]. For a given linear program, suppose that perturbing the entries of the program by at most δ each does not change the feasibility of the linear program, and let δ^* be the largest such δ . Let U denote the maximum absolute value of entries in the linear program. Then $\kappa = U/\delta^*$ is Renegar's condition number for the linear program.

Suppose we are given a polynomially bounded linear program $\max\{\mathbf{c}^\top \mathbf{x} : \mathbf{A}\mathbf{x} \leq \mathbf{b}, \mathbf{x} \geq \mathbf{0}\}$ (also referred to as $(\mathbf{A}, \mathbf{b}, \mathbf{c})$), with polytope radius at most R , and Renegar's condition number κ is also bounded by a polynomial. We wish to compute a vector $\mathbf{x} \geq \mathbf{0}$ with an ϵ additive error on each constraint and in the optimal value. We can reduce this problem, for instance, $(\mathbf{A}, \mathbf{b}, \mathbf{c})$ to a polynomially bounded linear program instance with integral input numbers. Specifically, we round the entries of \mathbf{A} down to $\tilde{\mathbf{A}}$ and those of \mathbf{b}, \mathbf{c} up to $\tilde{\mathbf{b}}, \tilde{\mathbf{c}}$ all by at most $\min\{\frac{\epsilon}{3R}, \frac{U}{\kappa R}\}$ such that each entry of $\tilde{\mathbf{A}}, \tilde{\mathbf{b}}, \tilde{\mathbf{c}}$ only needs a logarithmic number of bits. Suppose $\tilde{\mathbf{x}}$ is a solution to $(\tilde{\mathbf{A}}, \tilde{\mathbf{b}}, \tilde{\mathbf{c}})$ with $\frac{\epsilon}{3}$ additive error on each constraint and the optimal value. Then,

$$\begin{aligned} \mathbf{A}\tilde{\mathbf{x}} &= \tilde{\mathbf{A}}\tilde{\mathbf{x}} + (\mathbf{A} - \tilde{\mathbf{A}})\tilde{\mathbf{x}} \leq \mathbf{b} + \epsilon \mathbf{1}, \\ \mathbf{c}^\top \tilde{\mathbf{x}} &\geq \tilde{\mathbf{c}}^\top \tilde{\mathbf{x}}^* - \frac{2\epsilon}{3}. \end{aligned}$$

Here, $\mathbf{1}$ is the all-one vector, $\tilde{\mathbf{x}}^*$ is an optimal solution to $(\tilde{\mathbf{A}}, \tilde{\mathbf{b}}, \tilde{\mathbf{c}})$. In addition, the optimal value of $(\tilde{\mathbf{A}}, \tilde{\mathbf{b}}, \tilde{\mathbf{c}})$ is greater than or equal to that of $(\mathbf{A}, \mathbf{b}, \mathbf{c})$. So, $\tilde{\mathbf{x}}$ is a solution to $(\mathbf{A}, \mathbf{b}, \mathbf{c})$ with ϵ additive error as desired. Since each entry of $\tilde{\mathbf{A}}, \tilde{\mathbf{b}}, \tilde{\mathbf{c}}$ has a logarithmic number of bits, we can scale all of them to polynomially bounded integers without changing the feasible set and the optimal solutions. Thus we see that if we restrict ourselves to polynomially linear programs with polynomially bounded Renegar's condition number, and we wish to solve the program with small additive error, we can assume, without loss of generality, that the program has integer coefficients.

3.3.2 Problems in Algebra Space

For the convenience of our reduction, we define linear programming as a “decision” problem. We can solve the optimization problem $\max\{\mathbf{c}^\top \mathbf{x} : \mathbf{A}\mathbf{x} \leq \mathbf{b}, \mathbf{x} \geq \mathbf{0}\}$ by binary

¹W.l.o.g., by scaling, this is the same as assuming the program is written with polynomially bounded fixed precision numbers of the form k/D where k is an integer, and D is an integral denominator shared across all entries, and both k and D are polynomially bounded.

searching its optimal value via the decision problem.

Definition 3.3.1 (Linear Programming (LP)). Given a matrix $\mathbf{A} \in \mathbb{Z}^{m \times n}$, vectors $\mathbf{b} \in \mathbb{Z}^m$ and $\mathbf{c} \in \mathbb{Z}^n$, an integer K , and $R \geq \max\{1, \max\{\|\mathbf{x}\|_1 : \mathbf{A}\mathbf{x} \leq \mathbf{b}, \mathbf{x} \geq \mathbf{0}\}\}$, we refer to the LP problem for $(\mathbf{A}, \mathbf{b}, \mathbf{c}, K, R)$ as the problem of finding a vector $\mathbf{x} \in \mathbb{R}_{\geq 0}^n$ satisfying

$$\mathbf{A}\mathbf{x} \leq \mathbf{b} \text{ and } \mathbf{c}^\top \mathbf{x} \geq K$$

if such an \mathbf{x} exists and returning “infeasible” otherwise.

We will reduce linear programs to linear equations with non-negative variables (LEN), and then to linear equations with non-negative variables and small integer coefficients (k -LEN).

Definition 3.3.2 (Linear Equations with Non-negative Variables (LEN)). Given $\mathbf{A} \in \mathbb{Z}^{m \times n}$, $\mathbf{b} \in \mathbb{Z}^m$, and $R \geq \max\{1, \max\{\|\mathbf{x}\|_1 : \mathbf{A}\mathbf{x} = \mathbf{b}, \mathbf{x} \geq \mathbf{0}\}\}$, we refer to the LEN problem for $(\mathbf{A}, \mathbf{b}, R)$ as the problem of finding a vector $\mathbf{x} \in \mathbb{R}_{\geq 0}^n$ satisfying $\mathbf{A}\mathbf{x} = \mathbf{b}$ if such an \mathbf{x} exists and returning “infeasible” otherwise.

Definition 3.3.3 (k -LEN (k -LEN)). The k -LEN problem is an LEN problem $(\mathbf{A}, \mathbf{b}, R)$ where the entries of \mathbf{A} are integers in $[-k, k]$ for some given $k \in \mathbb{Z}_{>0}$.

We employ the following additive error notion. We append a letter “A” to each problem name to denote its approximation version, e.g., LP Approximate Problem is abbreviated to LPA.

Definition 3.3.4 (Error Notions in the Algebra Space). We always require $\mathbf{x} \geq \mathbf{0}$. In addition,

- The inequality constraint $\mathbf{A}\mathbf{x} \leq \mathbf{b}$ is relaxed to $\mathbf{A}\mathbf{x} - \mathbf{b} \leq \epsilon \mathbf{1}$, where $\mathbf{1}$ is the all-1 vector;
- The equality constraint $\mathbf{A}\mathbf{x} = \mathbf{b}$ is relaxed to $\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_\infty \leq \epsilon$.

3.3.3 Problems in Flow Space

A *flow network* is a directed graph $G = (V, E)$, where V is the set of vertices and $E \subseteq V \times V$ is the set of edges, together with a vector of edge capacities $\mathbf{u} \in \mathbb{Z}_{>0}^{|E|}$ that upper bounds the amount of flow passing each edge. A *2-commodity flow network* is a flow network together with two source-sink pairs $s_i, t_i \in V$ for each commodity $i \in \{1, 2\}$.

Definition 3.3.5 (Capacity Constraints). Given a 2-commodity flow network, we say the flows $\mathbf{f}_1, \mathbf{f}_2 \geq \mathbf{0}$ satisfy *capacity constraints* if

$$\mathbf{f}_1(e) + \mathbf{f}_2(e) \leq \mathbf{u}(e), \quad \forall e \in E.$$

Definition 3.3.6 (Flow Conservation Constraints). Given a 2-commodity flow network, we say the flows $\mathbf{f}_1, \mathbf{f}_2 \geq \mathbf{0}$ satisfy *flow conservation constraints* if

$$\sum_{u:(u,v) \in E} \mathbf{f}_i(u, v) = \sum_{w:(v,w) \in E} \mathbf{f}_i(v, w), \quad \forall i \in \{1, 2\}, v \in V \setminus \{s_i, t_i\}.$$

Given a 2-commodity flow network $(G = (V, E), \mathbf{u}, s_1, t_1, s_2, t_2)$, a *feasible 2-commodity flow* is a pair of flows $\mathbf{f}_1, \mathbf{f}_2 \in \mathbb{R}_{\geq 0}^{|E|}$ that satisfies both the capacity and flow conservation constraints.

In particular, in a *1-commodity flow network*, capacity constraints refer to $\mathbf{0} \leq \mathbf{f}(e) \leq \mathbf{u}(e)$, $\forall e \in E$; and flow conservation constraints become $\sum_{u:(u,v) \in E} \mathbf{f}(u, v) = \sum_{w:(v,w) \in E} \mathbf{f}(v, w)$, $v \in V \setminus \{s, t\}$. A *feasible 1-commodity flow* is a flow $\mathbf{f} \in \mathbb{R}_{\geq 0}^{|E|}$ that satisfies both the capacity and flow conservation constraints.

Similar to the definition of LP, we define 2-commodity flow problem as a decision problem. The corresponding optimization problem can be solved by reducing it to a sequence of decision problems.

Definition 3.3.7 (2-Commodity Flow Problem (2CF)). Given a 2-commodity flow network $(G, \mathbf{u}, s_1, t_1, s_2, t_2)$ together with $R \geq 0$, we refer to the 2CF problem for $(G, \mathbf{u}, s_1, t_1, s_2, t_2, R)$ as the problem of finding a feasible 2-commodity flow $\mathbf{f}_1, \mathbf{f}_2 \geq \mathbf{0}$ satisfying

$$F_1 + F_2 \geq R$$

if such flows exist and returning “infeasible” otherwise.

A closely related flow problem to 2CF is 2CFR, which imposes an additional requirement that a specific amount of flow must be routed for each commodity.

Definition 3.3.8 (2-Commodity Flow with Required Flow Amount (2CFR)). Given a 2-commodity flow network $(G, \mathbf{u}, s_1, t_1, s_2, t_2)$ together with $R_1, R_2 \geq 0$, we refer to the 2CFR for $(G, \mathbf{u}, s_1, t_1, s_2, t_2, R_1, R_2)$ as the problem of finding a feasible 2-commodity flow $\mathbf{f}_1, \mathbf{f}_2 \geq \mathbf{0}$ satisfying

$$F_1 \geq R_1, \quad F_2 \geq R_2$$

if such flows exist and returning “infeasible” otherwise.

Reducing LP to 2CF needs to go through a sequence of variants of flow problems with additional constraints beyond capacity and flow conservation constraints.

Definition 3.3.9 (Homologous Flow Constraints). Given a collection of disjoint subsets of edges $\mathcal{H} = \{H_1, \dots, H_h\}$ in a 1-commodity flow network, we say the flow $\mathbf{f} \geq \mathbf{0}$ satisfies *homologous flow constraints on \mathcal{H}* if

$$\mathbf{f}(e_1) = \mathbf{f}(e_2), \quad \forall e_1, e_2 \in H_k, \quad \forall H_k \in \mathcal{H}.$$

If $|H_k| = 2$ for all $H_k \in \mathcal{H}$, we say \mathbf{f} satisfy *pairwise homologous flow constraints*.

Definition 3.3.10 (Fixed Flow Constraints). Given a set $F \subseteq E$ in a 2-commodity flow network, we say the flows $\mathbf{f}_1, \mathbf{f}_2 \geq \mathbf{0}$ satisfy *fixed flow constraints on F* if

$$\mathbf{f}_1(e) + \mathbf{f}_2(e) = \mathbf{u}(e), \quad \forall e \in F.$$

Similarly, given a set $F \subseteq E$ in a 1-commodity flow network, we say the flow $\mathbf{f} \geq \mathbf{0}$ satisfies *fixed flow constraints on F* if

$$\mathbf{f}(e) = \mathbf{u}(e), \quad \forall e \in F.$$

Definition 3.3.11 (Selective Flow Constraints). Given two disjoint subsets of edges $S_1, S_2 \subseteq E$ in a 2-commodity flow network, we say the flows $\mathbf{f}_1, \mathbf{f}_2 \geq \mathbf{0}$ satisfy *selective flow constraints on $S_1 \cup S_2$* if

$$\mathbf{f}_i(e) > 0, \quad \mathbf{f}_{\bar{i}} = 0, \quad \forall e \in S_i, \quad i \in \{1, 2\}, \quad \bar{i} \neq i.$$

Intermediate problems are defined with some combinations of the above additional flow constraints.

Definition 3.3.12 (Fixed Homologous Flow Problem (FHF)). Given a flow network with a single source-sink pair (G, \mathbf{u}, s, t) together with a collection of disjoint subsets of edges $\mathcal{H} = \{H_1, \dots, H_h\}$ and a subset of edges $F \subseteq E$ such that F is disjoint from all the sets in \mathcal{H} , we refer to the FHF problem for $(G, F, \mathcal{H}, \mathbf{u}, s, t)$ as the problem of finding a feasible flow $\mathbf{f} \geq \mathbf{0}$ such that \mathbf{f} satisfies the homologous constraints on \mathcal{H} , and satisfies fixed flow constraints on F , if such flows exist, and returning “infeasible” otherwise.

Definition 3.3.13 (Fixed Pair Homologous Flow Problem (FPHF)). An FPHF is an FHF problem $(G, F, \mathcal{H}, \mathbf{u}, s, t)$ where every set in \mathcal{H} has size 2.

Definition 3.3.14 (Selective Fixed Flow Problem (SFF)). Given a 2-commodity network $(G, \mathbf{u}, s_1, t_1, s_2, t_2)$ together with three edge sets $F, S_1, S_2 \subseteq E$, we refer to the SFF problem for $(G, F, S_1, S_2, \mathbf{u}, s_1, t_1, s_2, t_2)$ as the problem of finding a feasible 2-commodity flow $\mathbf{f}_1, \mathbf{f}_2 \geq \mathbf{0}$ such that $\mathbf{f}_1, \mathbf{f}_2$ satisfy the selective flow constraints on S_1, S_2 , and satisfy the fixed flow constraints on F , if such flows exist, and returning “infeasible” otherwise.

Definition 3.3.15 (2-Commodity Fixed Flow Problem (2CFF)). Given a 2-commodity flow network $(G, \mathbf{u}, s_1, t_1, s_2, t_2)$ together with a subset of edges $F \subseteq E$, we refer to the 2CFF problem for the tuple $(G, F, \mathbf{u}, s_1, t_1, s_2, t_2)$ as the problem of finding a feasible 2-commodity flow $\mathbf{f}_1, \mathbf{f}_2 \geq \mathbf{0}$ which also satisfies the fixed flow constraints on F if such flows exist and returning “infeasible” otherwise.

Now, we define error notions for the above flow problems.

Definition 3.3.16 (Error Notions in the Flow Space). We always require flows $\mathbf{f}_1, \mathbf{f}_2 \geq \mathbf{0}$.² In addition,

1. *Error in congestion.*

- Relax the capacity constraints: $\mathbf{f}_1(e) + \mathbf{f}_2(e) \leq \mathbf{u}(e) + \epsilon, \quad \forall e \in E;$
- Relax the fixed flow constraints: $|\mathbf{f}_1(e) + \mathbf{f}_2(e) - \mathbf{u}(e)| \leq \epsilon, \quad \forall e \in F.$

2. *Error in demand.*

- Relax the flow conservation constraints for vertices other than the source and sink:

$$\left| \sum_{u:(u,v) \in E} \mathbf{f}_i(u, v) - \sum_{w:(v,w) \in E} \mathbf{f}_i(v, w) \right| \leq \epsilon, \quad \forall v \in V \setminus \{s_i, t_i\}, \quad i \in \{1, 2\}.$$

²Let $\mathbf{f}_2 = \mathbf{0}$ in 1-commodity flow problems.

- Relax the amount of flow requirements for source and sink nodes:

$$\sum_{w:(s_i,w) \in E} \mathbf{f}_i(s_i, w) \geq F_i - \epsilon, \quad \sum_{u:(u,t_i) \in E} \mathbf{f}_i(u, t_i) \geq F_i - \epsilon, \quad i \in \{1, 2\}.$$

3. *Error in selection.*

Relax the selective flow constraints: $\mathbf{f}_{\bar{i}}(e) \leq \epsilon, \quad \forall e \in S_i, \bar{i} \neq i.$

4. *Error in (pairwise) homology.*

Relax the (pairwise) homologous flow constraints: $|\mathbf{f}(e_1) - \mathbf{f}(e_2)| \leq \epsilon, \quad \forall e_1, e_2 \in H_k, H_k \in \mathcal{H}.$

We remark that the error notions in the flow space are essentially equivalent to those in the algebra space, as defined in Definition 3.3.4. For instance, relaxing capacity constraints corresponds to relaxing inequality constraints, while relaxing flow conservation constraints corresponds to relaxing equality constraints.

3.4 Main Results

Before formally stating our main results, we first formalize the approximate versions of LP and 2CF by applying the error notions defined earlier in Definition 3.3.4 and 3.3.16.

Definition 3.4.1 (LP Approximate Problem (LPA)). An LPA instance is given by an LP instance $(\mathbf{A}, \mathbf{b}, \mathbf{c}, K, R)$ and an error parameter $\epsilon \in [0, 1]$, which we collect in a tuple $(\mathbf{A}, \mathbf{b}, \mathbf{c}, K, R, \epsilon)$. We say an algorithm solves the LPA problem, if, given any LPA instance, it returns a vector $\mathbf{x} \geq \mathbf{0}$ such that

$$\begin{aligned} \mathbf{c}^\top \mathbf{x} &\geq K - \epsilon, \\ \mathbf{A}\mathbf{x} &\leq \mathbf{b} + \epsilon \mathbf{1}, \end{aligned}$$

where $\mathbf{1}$ is the all-1 vector, or it correctly declares that the associated LP instance is infeasible.

Definition 3.4.2 (2CF Approximate Problem (2CFA)). A 2CFA instance is given by a 2CF instance $(G, \mathbf{u}, s_1, t_1, s_2, t_2, R)$ and an error parameter $\epsilon \in [0, 1]$, which we collect in a tuple $(G, \mathbf{u}, s_1, t_1, s_2, t_2, R, \epsilon)$. We say an algorithm solves the 2CFA problem, if, given any 2CFA instance, it returns a pair of flows $\mathbf{f}_1, \mathbf{f}_2 \geq 0$ that satisfies

$$\mathbf{f}_1(e) + \mathbf{f}_2(e) \leq \mathbf{u}(e) + \epsilon, \quad \forall e \in E, \quad (\text{error in congestion})$$

$$\left| \sum_{u:(u,v) \in E} \mathbf{f}_i(u, v) - \sum_{w:(v,w) \in E} \mathbf{f}_i(v, w) \right| \leq \epsilon, \quad \forall v \in V \setminus \{s_i, t_i\}, i \in \{1, 2\}, \quad (\text{error in demand})$$

$$\left| \sum_{w:(s_i,w) \in E} \mathbf{f}_i(s_i, w) - F_i \right| \leq \epsilon, \quad \left| \sum_{u:(u,t_i) \in E} \mathbf{f}_i(u, t_i) - F_i \right| \leq \epsilon, \quad i \in \{1, 2\}, \quad (\text{error in demand})$$

where $F_1 + F_2 = R$;³ or it correctly declares that the associated 2CF instance is infeasible.

³If we encode 2CF as an LP instance, and approximately solve the LP with at most ϵ additive error. Then, the approximate solution also agrees with the error notions of 2CF, except that we get $F_1 + F_2 \geq R - \epsilon$ instead of $F_1 + F_2 \geq R$. This inconsistency can be eliminated by setting $\epsilon' = 2\epsilon$, and slightly adjusting F_1, F_2 to F'_1, F'_2 such that $F'_1 + F'_2 \geq R$. This way, we obtain an approximate solution to 2CF with at most ϵ' additive error.

With the approximate versions of LP and 2CF defined, we now present the main result formally, which provides a fine-grained complexity on the problem size and error tolerance in the reduction from LPA to 2CFA.

Theorem 3.4.3 (Main Theorem). *Given an LPA instance $(\mathbf{A}, \mathbf{b}, \mathbf{c}, K, R, \epsilon^{lp})$ where $\mathbf{A} \in \mathbb{Z}^{m \times n}$, $\mathbf{b} \in \mathbb{Z}^m$, $\mathbf{c} \in \mathbb{Z}^n$, $K \in \mathbb{Z}$, $\epsilon^{lp} \geq 0$ and \mathbf{A} has $\text{nnz}(\mathbf{A})$ non-zero entries, we can reduce it to a 2CFA instance $(G = (V, E), \mathbf{u}, s_1, t_1, s_2, t_2, R^{2cf}, \epsilon^{2cf})$ in time $O(\text{nnz}(\mathbf{A}) \log X)$ where $X = X(\mathbf{A}, \mathbf{b}, \mathbf{c}, K)$, such that*

$$\begin{aligned} |V|, |E| &= O(\text{nnz}(\mathbf{A}) \log X), \\ \|\mathbf{u}\|_{\max}, R^{2cf} &= O(\text{nnz}^3(\mathbf{A}) R X \log^2 X), \\ \epsilon^{2cf} &= \Omega\left(\frac{1}{\text{nnz}^6(\mathbf{A}) R X \log^6 X}\right) \epsilon^{lp}. \end{aligned}$$

If the LP instance $(\mathbf{A}, \mathbf{b}, \mathbf{c}, K, R)$ has a solution, then the 2CF instance $(G^{2cf}, \mathbf{u}^{2cf}, s_1, t_1, s_2, t_2, R^{2cf})$ has a solution. Furthermore, if \mathbf{f}^{2cf} is a solution to the 2CFA (2CF) instance, then in time $O(\text{nnz}(\mathbf{A}) \log X)$, we can compute a solution \mathbf{x} to the LPA (LP, respectively) instance, where the exact case holds when $\epsilon^{2cf} = \epsilon^{lp} = 0$.

Our main theorem immediately implies the following corollary.

Corollary 3.4.4. *If we can solve any 2CFA instance $(G = (V, E), \mathbf{u}, s_1, t_1, s_2, t_2, R^{2cf}, \epsilon)$ in time $O(|E|^c \text{poly log}(\frac{\|\mathbf{u}\|_1}{\epsilon}))$ for some small constant $c \geq 1$, then we can solve any LPA instance $(\mathbf{A}, \mathbf{b}, \mathbf{c}, K, R, \epsilon)$ in time $O(\text{nnz}^c(\mathbf{A}) \text{poly log}(\frac{\text{nnz}(\mathbf{A}) R X(\mathbf{A}, \mathbf{b}, \mathbf{c}, K)}{\epsilon}))$.*

3.5 Algebra Space

3.5.1 Overview

In this section, we describe the general process of reduction and solution mapping in the algebra space. Consider the task of reducing an instance A of the form $\{\mathbf{A}^A \mathbf{x}^A = \mathbf{b}^A, \mathbf{x}^A \geq \mathbf{0}\}$ to another instance B , defined as $\{\mathbf{A}^B \mathbf{x}^B = \mathbf{b}^B, \mathbf{x}^B \geq \mathbf{0}\}$. We analyze the reductions for both the exact and approximate cases.

Exact Case. The reduction algorithms are designed such that A has a feasible solution if and only if the reduced instance B also has a feasible solution. It implies that instance A is infeasible if and only if B is infeasible.

Reduction:

Given \mathbf{A}^A and \mathbf{b}^A , while the reduction algorithm to construct \mathbf{A}^B and \mathbf{b}^B is problem-specific, the corresponding \mathbf{x}^B is constructed by introducing auxiliary variables \mathbf{x}^{aux} such that:

$$\mathbf{x}^B = \begin{bmatrix} \mathbf{x}^A \\ \mathbf{x}^{\text{aux}} \end{bmatrix}.$$

We keep track of the increase in the problem size and problem magnitude for the reduction algorithm.

Solution mapping:

The solution mapping algorithm is the same for all reduction steps, with a straightforward runtime analysis: If a solver returns a feasible solution \mathbf{x}^B for instance B , then we extract the appropriate subset of entries \mathbf{x}^A from \mathbf{x}^B in time $O(\dim(\mathbf{x}^A))$ and return \mathbf{x}^A as the solution to the instance A ; if the solver returns “infeasible” for instance B , then we return “infeasible” for instance A as well.

Approximate Case. For approximate problems, both the reduction and solution mapping algorithms are applied in the same way. However, there is a key difference from the exact case: the approximate problem does not require to provide a certificate of “infeasibility”, but never incorrectly asserts “infeasibility”.

In addition to tracking the increase in problem size and magnitude, we perform an error analysis to monitor error propagation. Specifically, let $(\mathbf{x}^B, \epsilon^B)$ be an approximate solution to instance B with error ϵ^B , and let $(\mathbf{x}^A, \epsilon^A)$ be an approximate solution to instance A with error ϵ^A , obtained by applying the solution mapping algorithm. Our goal is to analyze how ϵ^A increases compared to ϵ^B . We remark that when $\epsilon^A, \epsilon^B = 0$, the statement for the approximate case is reduced to the exact case.

In the remainder of this section, we describe each reduction step in algebra space in detail. For each step, we first describe the reduction algorithm that is used to construct $\mathbf{A}^B, \mathbf{b}^B$ from $\mathbf{A}^A, \mathbf{b}^A$. We then present a lemma for the exact case, proving the correctness of the reduction algorithm while tracking the problem size and magnitude. And finally we extend the lemma for the approximate case, analyzing how errors propagate through the solution mapping algorithm.

3.5.2 LP(A) to LEN(A): Reducing Inequalities to Equalities

Given an LP instance $(\mathbf{A}, \mathbf{b}, \mathbf{c}, K, R)$ where $R \geq \max\{1, \max\{\|\mathbf{x}\|_1 : \mathbf{A}\mathbf{x} \leq \mathbf{b}, \mathbf{x} \geq \mathbf{0}\}\}$, we want to compute a vector $\mathbf{x} \geq \mathbf{0}$ s.t.

$$\mathbf{c}^\top \mathbf{x} \geq K, \quad \mathbf{A}\mathbf{x} \leq \mathbf{b}.$$

To reduce it to an LEN instance, we introduce \mathbf{x}^{aux} as slack variables $\begin{bmatrix} \mathbf{s} \\ \alpha \end{bmatrix} \geq \mathbf{0}$, and construct

$$\tilde{\mathbf{A}} = \begin{bmatrix} \mathbf{c}^\top & \mathbf{0} & -1 \\ \mathbf{A} & \mathbf{I} & \mathbf{0} \end{bmatrix}, \quad \tilde{\mathbf{x}} = \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \\ \alpha \end{bmatrix}, \quad \tilde{\mathbf{b}} = \begin{bmatrix} K \\ \mathbf{b} \end{bmatrix}. \quad (3.1)$$

This defines an LEN instance $(\tilde{\mathbf{A}}, \tilde{\mathbf{b}}, \tilde{R})$ where $\tilde{R} = \max\{1, \max\{\|\tilde{\mathbf{x}}\|_1 : \tilde{\mathbf{A}}\tilde{\mathbf{x}} = \tilde{\mathbf{b}}, \tilde{\mathbf{x}} \geq \mathbf{0}\}\}$.

Lemma 3.5.1 (LP to LEN). *Given an LP instance $(\mathbf{A}, \mathbf{b}, \mathbf{c}, K, R)$ with N non-zeros and magnitude X , we can construct, in $O(N)$ time, an LEN instance $(\tilde{\mathbf{A}}, \tilde{\mathbf{b}}, \tilde{R})$ with \tilde{N} non-zeros and magnitude \tilde{X} such that*

$$\tilde{N} = O(N), \quad \tilde{R} = O(NRX), \quad \tilde{X} = O(X).$$

If the LP instance has a solution, then the LEN instance has a solution.

Proof. Based on the reduction construction shown in Eq. (3.1), it is obvious for the linear reduction time and $\tilde{N} = O(N)$, $\tilde{X} = O(X)$. Regarding the radius of polytope $\tilde{R} = \max\{1, \max\{\|\tilde{\mathbf{x}}\|_1 : \tilde{\mathbf{A}}\tilde{\mathbf{x}} = \tilde{\mathbf{b}}, \tilde{\mathbf{x}} \geq \mathbf{0}\}\}$, we have

$$\begin{aligned} \|\tilde{\mathbf{x}}\|_1 &\leq \|\mathbf{x}\|_1 + \|\mathbf{x}^{\text{aux}}\|_1 && \text{by triangle inequality} \\ &\leq \|\mathbf{x}\|_1 + \|\mathbf{Ax} - \mathbf{b}\|_1 + \|\mathbf{c}^\top \mathbf{x} - K\|_1 \\ &\leq R + N(\|\mathbf{A}\|_{\max} \|\mathbf{x}\|_1 + \|\mathbf{b}\|_{\max} + \|\mathbf{c}\|_{\max} \|\mathbf{x}\|_1 + |K|) \\ &\leq O(NRX). \end{aligned}$$

If \mathbf{x} is a solution to the LP instance, we can derive a solution $\tilde{\mathbf{x}} = (\mathbf{x}^\top, \mathbf{s}^\top, \alpha)^\top$ to the LEN instance by setting

$$\mathbf{s} = \mathbf{b} - \mathbf{Ax}, \quad \alpha = \mathbf{c}^\top \mathbf{x} - K.$$

□

Now we extend Lemma 3.5.1 for the exact case to the approximate case.

Lemma 3.5.2 (LPA to LENA). *Given an LPA instance $(\mathbf{A}, \mathbf{b}, \mathbf{c}, K, R, \epsilon^{\text{lp}})$ with N non-zeros and magnitude X , we can reduce it to an LENA instance $(\tilde{\mathbf{A}}, \tilde{\mathbf{b}}, \tilde{R}, \epsilon^{\text{le}})$ with \tilde{N} non-zeros and magnitude \tilde{X} , by letting*

$$\epsilon^{\text{le}} = \epsilon^{\text{lp}},$$

and using Lemma 3.5.1 to construct an LEN instance $(\tilde{\mathbf{A}}, \tilde{\mathbf{b}}, \tilde{R})$ from the LP instance $(\mathbf{A}, \mathbf{b}, \mathbf{c}, K, R)$. If $\tilde{\mathbf{x}}$ is a solution to the LENA (or LEN) instance, then in time $O(N)$, we can compute a solution \mathbf{x} to the LPA (or LP respectively) instance, where the exact case holds when $\epsilon^{\text{le}} = \epsilon^{\text{lp}} = 0$.

Proof. If $\tilde{\mathbf{x}} = [\mathbf{x}^\top, \mathbf{s}^\top, \alpha]^\top$ is an approximate solution to the LENA instance with error ϵ^{le} , according to error notions in algebra space (Definition 3.3.4),

$$\begin{aligned} |\mathbf{c}^\top \mathbf{x} - \alpha - K| &\leq \epsilon^{\text{le}}, \\ |\mathbf{Ax} + \mathbf{s} - \mathbf{b}| &\leq \epsilon^{\text{le}} \mathbf{1}. \end{aligned}$$

Taking one direction of the absolute value inequalities,

$$\begin{aligned} \mathbf{c}^\top \mathbf{x} &\geq \alpha + K - \epsilon^{\text{le}} \geq K - \epsilon^{\text{le}}, \\ \mathbf{Ax} &\leq -\mathbf{s} + \mathbf{b} + \epsilon^{\text{le}} \mathbf{1} \leq \mathbf{b} + \epsilon^{\text{le}} \mathbf{1}, \end{aligned}$$

As we set in the reduction that $\epsilon^{\text{le}} = \epsilon^{\text{lp}}$, \mathbf{x} is a solution to the LPA instance $(\mathbf{A}, \mathbf{b}, \mathbf{c}, K, R, \epsilon^{\text{lp}})$. □

3.5.3 LEN(A) to 2-LEN(A): Reducing Integer Coefficients to $\{0, \pm 1, \pm 2\}$

In this section, we show the reduction from an LEN instance $(\tilde{\mathbf{A}}, \tilde{\mathbf{b}}, \tilde{R})$ to a 2-LEN instance $(\bar{\mathbf{A}}, \bar{\mathbf{b}}, \bar{R})$, where the coefficients of $\bar{\mathbf{A}}$ are in $\{0, \pm 1, \pm 2\}$. The reduction is based on *bitwise decomposition*. We illustrate the reduction algorithm with a concrete example.

$$5x_1 + 3x_2 - 7x_3 = -1$$

$$\Downarrow$$

$$(2^0 + 2^2)\mathbf{x}_1 + (2^0 + 2^1)\mathbf{x}_2 - (2^0 + 2^1 + 2^2)\mathbf{x}_3 = -2^0$$

$$\Downarrow$$

$$(\mathbf{x}_1 + \mathbf{x}_2 - \mathbf{x}_3)2^0 + (\mathbf{x}_2 - \mathbf{x}_3)2^1 + (\mathbf{x}_1 - \mathbf{x}_3)2^2 = -1 \cdot 2^0$$

It can be decomposed to 3 linear equations, together with carry terms $(\mathbf{c}_i - \mathbf{d}_i)$, where $\mathbf{c}_i, \mathbf{d}_i \geq 0$:

$$\begin{aligned} \mathbf{x}_1 + \mathbf{x}_2 - \mathbf{x}_3 - 2(\mathbf{c}_0 - \mathbf{d}_0) &= -1, \\ \mathbf{x}_2 - \mathbf{x}_3 + (\mathbf{c}_0 - \mathbf{d}_0) - 2(\mathbf{c}_1 - \mathbf{d}_1) &= 0, \\ \mathbf{x}_1 - \mathbf{x}_3 + (\mathbf{c}_1 - \mathbf{d}_1) &= 0. \end{aligned} \tag{3.2}$$

Moreover, since $\mathbf{c}_i, \mathbf{d}_i$ can be arbitrarily large, we also impose an upper bound for carry variables $\mathbf{c}_i, \mathbf{d}_i \leq \bar{X}$, where \bar{X} will be specified shortly.

Finally, we introduce slack variables $\mathbf{s} \geq \mathbf{0}$ to reduce inequalities to equalities:

$$\begin{bmatrix} \mathbf{c} \\ \mathbf{d} \end{bmatrix} + \mathbf{s} = \bar{X}\mathbf{1}.$$

We repeat the above procedure for all equations in LEN. Notice that in this step, \mathbf{x}^{aux} is composed of $\mathbf{c}, \mathbf{d}, \mathbf{s}$.

Lemma 3.5.3 (LEN to 2-LEN). *Given an LEN instance $(\tilde{\mathbf{A}}, \tilde{\mathbf{b}}, \tilde{R})$ with \tilde{N} non-zeros and magnitude \tilde{X} , we can construct, in $O(\tilde{N} \log \tilde{X})$ time, a 2-LEN instance $(\bar{\mathbf{A}}, \bar{\mathbf{b}}, \bar{R})$ with \bar{N} non-zeros and magnitude \bar{X} such that*

$$\bar{N} = O(\tilde{N} \log \tilde{X}), \quad \bar{R} = O(\tilde{N} \tilde{R} \log \tilde{X}), \quad \bar{X} = O(\tilde{R}).$$

If the LEN instance has a solution, then the 2-LEN instance has a solution.

Proof. By construction, for each equation in LEN, it is decomposed into at most $O(\log \tilde{X})$ equations so that the number of non-zero entries is increased by a factor $O(\log \tilde{X})$. Therefore, the reduction can be conducted in $O(\tilde{N} \log \tilde{X})$ time and $\bar{N} = O(\tilde{N} \log \tilde{X})$.

The key to bounding the radius of polytope \bar{R} and magnitude \bar{X} is to analyze the magnitude of carry terms. For an arbitrary equation $\tilde{\mathbf{A}}(q)\tilde{\mathbf{x}} = \tilde{\mathbf{b}}(q)$, denote the introduced carry terms of the decomposed equations as $\mathbf{c}_q, \mathbf{d}_q \in \mathbb{R}_{\geq 0}^{N_q}$, where $N_q \leq O(\log \tilde{X})$ denotes the number of decomposed equations corresponding to original equation. Starting from the lowest bit, we have

$$2|\mathbf{c}_q(0) - \mathbf{d}_q(0)| \leq 1 + \|\tilde{\mathbf{x}}\|_1 = O(\tilde{R}).$$

Then for the second lowest bit,

$$2|\mathbf{c}_q(1) - \mathbf{d}_q(1)| \leq 1 + \|\tilde{\mathbf{x}}\|_1 + |\mathbf{c}_q(0) - \mathbf{d}_q(0)| = \left(1 + \frac{1}{2}\right) O(\tilde{R}).$$

Repeating the same until the second-highest bit,

$$2|\mathbf{c}_q(N_q - 1) - \mathbf{d}_q(N_q - 1)| \leq \left(1 + \frac{1}{2} + \cdots + \frac{1}{2^{N_q-1}}\right) O(\tilde{R}) = O(\tilde{R}).$$

Therefore, it suffices to bound \mathbf{c}, \mathbf{d} by $O(\tilde{R})$, thus the same for the slack variables $\mathbf{s} \leq O(\tilde{R})$.

Now, we can bound the magnitude of 2-LENA by $\bar{X} = O(\tilde{R})$ as $\|\bar{\mathbf{A}}\|_{\max} = 2$ and $\|\bar{\mathbf{b}}\|_{\max} \leq O(\tilde{R})$. To bound the radius of the polytope, we have

$$\|\bar{\mathbf{x}}\|_1 \leq \|\tilde{\mathbf{x}}\|_1 + \|\mathbf{c}\|_1 + \|\mathbf{d}\|_1 + \|\mathbf{s}\|_1 \leq \tilde{R} + O(\tilde{N} \log \tilde{X}) \times \|\mathbf{c}\|_{\max} \leq O(\tilde{N} \tilde{R} \log \tilde{X}).$$

Finally, we describe how to construct a solution to 2-LEN given a solution to LENA. If $\tilde{\mathbf{x}}$ is a feasible solution to the LENA instance such that $\tilde{\mathbf{A}}\tilde{\mathbf{x}} = \tilde{\mathbf{b}}$, we are able to derive a solution $\bar{\mathbf{x}} = (\tilde{\mathbf{x}}^\top, \mathbf{c}^\top, \mathbf{d}^\top, \mathbf{s}^\top)^\top$ to the 2-LEN instance. More concretely, we plug in $\tilde{\mathbf{x}}$ into 2-LEN, and for each group of equations, we start from the lowest bit and compute the smallest nonnegative $\mathbf{c}_0, \mathbf{d}_0$. Using forward substitution, we can compute all \mathbf{c}, \mathbf{d} . In the end, given \mathbf{c} and \mathbf{d} , \mathbf{s} can be computed as well. \square

Lemma 3.5.4 (LENA to 2-LENA). *Given an LENA instance $(\tilde{\mathbf{A}}, \tilde{\mathbf{b}}, \tilde{R}, \epsilon^{le})$ with \tilde{N} non-zeros and magnitude \tilde{X} , we can reduce it to an 2-LENA instance $(\bar{\mathbf{A}}, \bar{\mathbf{b}}, \bar{R}, \epsilon^{2le})$ with \bar{N} non-zeros and magnitude \bar{X} by letting*

$$\epsilon^{2le} = \Omega\left(\frac{1}{\tilde{X}}\right) \epsilon^{le},$$

and using Lemma 3.5.3 to construct a 2-LEN instance $(\bar{\mathbf{A}}, \bar{\mathbf{b}}, \bar{R})$ from the LENA instance $(\tilde{\mathbf{A}}, \tilde{\mathbf{b}}, \tilde{R})$. If $\bar{\mathbf{x}}$ is a solution to the 2-LENA (or 2-LEN) instance, then in time $O(\tilde{n})$, we can compute a solution $\tilde{\mathbf{x}}$ to the LENA (or LEN respectively) instance, where the exact case holds when $\epsilon^{2le} = \epsilon^{le} = 0$.

Proof. Suppose $\bar{\mathbf{x}} = [\tilde{\mathbf{x}}^\top, \mathbf{c}^\top, \mathbf{d}^\top, \mathbf{s}^\top]^\top$ is an approximate solution to the 2-LENA instance. We consider an arbitrary equation $\tilde{\mathbf{A}}(q)\tilde{\mathbf{x}} = \tilde{\mathbf{b}}(q)$ in LENA. By Definition 3.3.4, for the corresponding equations in 2-LENA, for any i th bit,

$$|\bar{\mathbf{A}}_q(i)\bar{\mathbf{x}} - \bar{\mathbf{b}}_q(i)| \leq \epsilon^{2le}.$$

The original equation $\tilde{\mathbf{A}}(q)\tilde{\mathbf{x}} = \tilde{\mathbf{b}}(q)$ in LENA can be recovered by combining all the corresponding bit equations in 2-LENA as follows:

$$\tilde{\mathbf{A}}(q)\tilde{\mathbf{x}} = \sum_{i=0}^{N_q} 2^i \times \bar{\mathbf{A}}_q(i)^\top \bar{\mathbf{x}}, \quad \tilde{\mathbf{b}}(q) = \sum_{i=0}^{N_q} 2^i \times \bar{\mathbf{b}}_q(i).$$

Hence,

$$\left| \tilde{\mathbf{A}}(q)\tilde{\mathbf{x}} - \tilde{\mathbf{b}}(q) \right| = \sum_{i=0}^{N_q} 2^i \times |\bar{\mathbf{A}}_q(i)\bar{\mathbf{x}} - \bar{\mathbf{b}}_q(i)| \leq O(\tilde{X}) \cdot \epsilon^{2le}.$$

\square

3.5.4 2-LEN(A) to 1-LEN(A): Reducing Coefficients from ± 2 to ± 1

The reduction from a 2-LEN instance $(\bar{\mathbf{A}}, \bar{\mathbf{b}}, \bar{R})$ to a 1-LEN instance $(\hat{\mathbf{A}}, \hat{\mathbf{b}}, \hat{R})$ is relatively straightforward: For each variable $\bar{\mathbf{x}}(j)$ that has a ± 2 coefficient, we introduce a new variable $\bar{\mathbf{x}}'(j)$, replace every $\pm 2\bar{\mathbf{x}}(j)$ with $\pm(\bar{\mathbf{x}}(j) + \bar{\mathbf{x}}'(j))$, and add an additional equation $\bar{\mathbf{x}}(j) - \bar{\mathbf{x}}'(j) = 0$.

Lemma 3.5.5 (2-LEN to 1-LEN). *Given a 2-LEN instance $(\bar{\mathbf{A}}, \bar{\mathbf{b}}, \bar{R})$ with \bar{N} non-zeros and magnitude \bar{X} , we can construct, in $O(\bar{N})$ time, a 1-LEN instance $(\hat{\mathbf{A}}, \hat{\mathbf{b}}, \hat{R})$ with \hat{N} non-zeros and magnitude \hat{X} such that*

$$\hat{N} = O(\bar{N}), \quad \hat{R} = O(\bar{R}), \quad \hat{X} = O(\bar{X}).$$

If the 2-LEN instance has a solution, then the 1-LEN instance has a solution.

Proof. By construction, we only make changes to each ± 2 coefficient, which will generate four non-zero entries, hence $\hat{N} = O(\bar{N})$ and both \hat{R} and \hat{X} will remain the same.

If $\bar{\mathbf{x}}$ is a solution to the 2-LEN instance such that $\bar{\mathbf{A}}\bar{\mathbf{x}} = \bar{\mathbf{b}}$, to derive a solution $\hat{\mathbf{x}} = (\bar{\mathbf{x}}^\top, (\bar{\mathbf{x}}')^\top)^\top$ to the 1-LEN instance, we split each $\bar{\mathbf{x}}(j)$ having a ± 2 coefficient in $\bar{\mathbf{A}}$ into $\bar{\mathbf{x}}'(j), \bar{\mathbf{x}}(j)$ such that $\bar{\mathbf{x}}'(j) = \bar{\mathbf{x}}(j)$. We can check that $\hat{\mathbf{x}}$ is a solution to the 1-LEN instance. \square

Lemma 3.5.6 (2-LENA to 1-LENA). *Given a 2-LENA instance $(\bar{\mathbf{A}}, \bar{\mathbf{b}}, \bar{R})$ with \bar{N} non-zeros and magnitude \bar{X} , we can reduce it to a 1-LENA instance $(\hat{\mathbf{A}}, \hat{\mathbf{b}}, \hat{R})$ with \hat{N} non-zeros and magnitude \hat{X} by letting*

$$\epsilon^{1le} = \Omega\left(\frac{1}{\bar{N}}\right) \epsilon^{2le},$$

and using Lemma 3.5.5 to construct a 1-LEN instance $(\hat{\mathbf{A}}, \hat{\mathbf{b}}, \hat{R})$ from the 2-LEN instance $(\bar{\mathbf{A}}, \bar{\mathbf{b}}, \bar{R})$. If $\hat{\mathbf{x}}$ is a solution to the 1-LENA (1-LEN) instance, then in time $O(\bar{n})$, we can compute a solution $\bar{\mathbf{x}}$ to the 2-LENA (2-LEN, respectively) instance, where the exact case holds when $\epsilon^{1le} = \epsilon^{2le} = 0$.

Proof. Suppose $\hat{\mathbf{x}}$ is an approximate solution to the 1-LENA instance. Then for any pair of $(\bar{\mathbf{x}}(j), \bar{\mathbf{x}}'(j))$, we have

$$|\bar{\mathbf{x}}(j) - \bar{\mathbf{x}}'(j)| \leq \epsilon^{1le},$$

and for any equations in 1-LENA, we have

$$\left| \hat{\mathbf{A}}(q)\hat{\mathbf{x}} - \hat{\mathbf{b}}(q) \right| \leq \epsilon^{1le}.$$

Therefore, for any equation in 2-LENA, assume S_q is the set of indices with coefficients ± 2 , then

$$\begin{aligned} \bar{\mathbf{A}}(q)\bar{\mathbf{x}} &= \sum_{k \notin S_q} \bar{\mathbf{A}}(q, k)\bar{\mathbf{x}}(k) + \sum_{k \in S_q} (\pm 2)\bar{\mathbf{x}}(k) \\ &= \sum_{k \notin S_q} \bar{\mathbf{A}}(q, k)\bar{\mathbf{x}}(k) + \sum_{k \in S_q} \pm ((\bar{\mathbf{x}}(k) + \bar{\mathbf{x}}'(k)) + (\bar{\mathbf{x}}(k) - \bar{\mathbf{x}}'(k))) \\ &= \hat{\mathbf{A}}(q)\hat{\mathbf{x}} + \sum_{k \in S_q} \pm (\bar{\mathbf{x}}(k) - \bar{\mathbf{x}}'(k)). \end{aligned}$$

Therefore,

$$\begin{aligned} |\bar{\mathbf{A}}(q)\bar{\mathbf{x}} - \bar{\mathbf{b}}(q)| &= \left| \hat{\mathbf{A}}(q)\hat{\mathbf{x}} - \hat{\mathbf{b}}(q) + \sum_{k \in S_q} \pm (\bar{\mathbf{x}}(k) - \bar{\mathbf{x}}'(k)) \right| \\ &\leq \left| \hat{\mathbf{A}}(q)\hat{\mathbf{x}} - \hat{\mathbf{b}}(q) \right| + \sum_{k \in S_q} |(\bar{\mathbf{x}}(k) - \bar{\mathbf{x}}'(k))| \\ &\leq (1 + |S_q|)\epsilon^{1le} \leq O(\bar{N})\epsilon^{1le}. \end{aligned}$$

\square

3.5.5 1-LEN(A) to FHF(A): Encoding Equations as Flows

This section describes the reduction from a 1-LEN instance $(\hat{\mathbf{A}}, \hat{\mathbf{b}}, \hat{R})$ to an FHF instance $(G^h, F^h, \mathbf{u}^h, \mathcal{H}^h, s, t)$. Recall the form of 1-LEN $\hat{\mathbf{A}}\hat{\mathbf{x}} = \hat{\mathbf{b}}$, where $\hat{\mathbf{A}} \in \mathbb{R}^{\hat{m} \times \hat{n}}$ and entries of $\hat{\mathbf{A}}$ are in $\{0, \pm 1\}$. For an arbitrary equation i in the 1-LEN instance, $\hat{\mathbf{A}}(i)\hat{\mathbf{x}} = \hat{\mathbf{b}}(i)$, $i \in [\hat{m}]$, let $J_i^+ = \{j | \hat{\mathbf{A}}(i, j) = 1\}$ and $J_i^- = \{j | \hat{\mathbf{A}}(i, j) = -1\}$ denote the set of indices of variables with coefficients being 1 and -1 in equation i , respectively. Then, each equation can be rewritten as a difference of the sum of variables with coefficient 1 and -1:

$$\sum_{j \in J_i^+} \hat{x}(j) - \sum_{j \in J_i^-} \hat{x}(j) = \hat{b}(i), \quad i \in [\hat{m}].$$




Figure 3.1: The reduction from 1-LEN to FHF.

Representing in this form, we can encode 1-LEN with a graph using the gadget as shown in Figure 3.1 – a graph that is composed of a number of *homologous* edges and *fixed flow* edges. More specifically, the fixed homologous flow network consists of a source s , a sink t , and \hat{m} sections representing the \hat{m} linear equations in 1-LEN. Inside each section i , there are 2 vertices $\{J_i^-, J_i^+\}$ and a number of edges:

- For the incoming edges of $\{J_i^-, J_i^+\}$,
 - if $\hat{\mathbf{A}}(i, j) = 1$, then s is connected to J_i^+ by edge \hat{x}_i^j with capacity \hat{R} ;
 - if $\hat{\mathbf{A}}(i, j) = -1$, then s is connected to J_i^- by edge \hat{x}_i^j with capacity \hat{R} ;
 - if $\hat{\mathbf{A}}(i, j) = 0$, no edge is needed.

Note that those incoming edges that correspond to the same variable are forced to route the same amount of flow, i.e., $(\hat{x}_1^j, \dots, \hat{x}_m^j), j \in [\hat{n}]$ constitute a *homologous* edge set that corresponds to the variable $\hat{x}(j)$.

- For the outgoing edges of $\{J_i^-, J_i^+\}$,

- J_i^+ is connected to t by a *fixed flow* edge \hat{b}_i that routes $\hat{\mathbf{b}}(i)$ units of flow;
- J_i^+ and J_i^- are connected to t by a pair of *homologous* edges e_i^+, e_i^- with capacity \hat{R} .

Lemma 3.5.7 (1-LEN to FHF). *Given a 1-LEN instance $(\hat{\mathbf{A}}, \hat{\mathbf{b}}, \hat{R})$ with \hat{N} non-zeros and magnitude \hat{X} , we can construct, in time $O(\hat{N})$, an FHF instance $(G^h, F^h, \mathbf{u}^h, \mathcal{H}^h = (H_1, \dots, H_h), s, t)$ such that*

$$|E^h| = O(\hat{N}), \quad \|\mathbf{u}^h\|_{\max} = \max \{ \hat{R}, \hat{X} \}.$$

If the 1-LEN instance has a solution, then the FHF instance has a solution.

Proof. According to Figure 3.1, the constructed graph G^h preserves the sparsity of $\hat{\mathbf{A}}$ and the number of edges $|E^h|$ can be bounded by $O(\hat{N})$. Moreover, the maximum edge capacity is bounded by

$$\|\mathbf{u}^h\|_{\max} = \max \left\{ \hat{R}, \|\hat{\mathbf{b}}\|_{\max} \right\} \leq \max \{ \hat{R}, \hat{X} \}.$$

If $\hat{\mathbf{x}}$ is a solution to the 1-LEN instance, we can derive a feasible flow \mathbf{f}^h for the FHF instance by setting

$$\mathbf{f}^h(\hat{x}_1^j) = \dots = \mathbf{f}^h(\hat{x}_m^j) = \hat{\mathbf{x}}(j) \leq \hat{R}, \quad \forall j \in [\hat{n}],$$

$$\mathbf{f}^h(\hat{b}_i) = \hat{\mathbf{b}}(i) \quad \text{and} \quad \mathbf{f}^h(e_i^+) = \mathbf{f}^h(e_i^-) = \sum_{j \in J_i^-} \hat{\mathbf{x}}(j) = \sum_{j \in J_i^+} \hat{\mathbf{x}}(j) - \hat{\mathbf{b}}(i), \quad \forall i \in [\hat{m}].$$

It is a feasible flow since both capacity constraints and flow conservation constraints are satisfied. \square

In the approximate case, we first formally define FHFA by applying the corresponding error notions (Definition 3.3.16) to FHF.

Definition 3.5.8 (FHF Approximate Problem (FHFA)). An FHFA instance is given by an FHF instance $(G, F, \mathbf{u}, \mathcal{H}, s, t)$ as in Definition 3.3.12, and an error parameter $\epsilon \in [0, 1]$, which we collect in a tuple $(G, F, \mathcal{H}, \mathbf{u}, s, t, \epsilon)$. We say an algorithm solves the FHFA problem, if, given any FHFA instance, it returns a flow $\mathbf{f} \geq \mathbf{0}$ that satisfies

$$\mathbf{u}(e) - \epsilon \leq \mathbf{f}(e) \leq \mathbf{u}(e) + \epsilon, \quad \forall e \in F, \quad (\text{error in congestion})$$

$$0 \leq \mathbf{f}(e) \leq \mathbf{u}(e) + \epsilon, \quad \forall e \in E \setminus F, \quad (\text{error in congestion})$$

$$\left| \sum_{u:(u,v) \in E} \mathbf{f}(u,v) - \sum_{w:(v,w) \in E} \mathbf{f}(v,w) \right| \leq \epsilon, \quad \forall v \in V \setminus \{s, t\}, \quad (\text{error in demand})$$

$$|\mathbf{f}(v,w) - \mathbf{f}(v',w')| \leq \epsilon, \quad \forall (v,w), (v',w') \in H_i, \quad \forall H_i \in \mathcal{H}, \quad (\text{error in homology})$$

or it correctly declares that the associated FHF instance is infeasible.

Lemma 3.5.9 (1-LENA to FHFA). *Given a 1-LENA instance $(\hat{\mathbf{A}}, \hat{\mathbf{b}}, \hat{R}, \epsilon^{1le})$ with \hat{N} non-zeros and magnitude \hat{X} , we can reduce it to an FHFA instance $(G^h, F^h, \mathcal{H}^h, \mathbf{u}^h, s, t, \epsilon^h)$ by letting*

$$\epsilon^h = \Omega \left(\frac{1}{\hat{N}} \right) \epsilon^{1le},$$

and using Lemma 3.5.7 to construct an FHF instance $(G^h, F^h, \mathcal{H}^h, \mathbf{u}^h, s, t)$ from the 1-LEN instance $(\hat{\mathbf{A}}, \hat{\mathbf{b}}, \hat{R})$. If \mathbf{f}^h is a solution to the FHFA (FHF) instance, then in time $O(\hat{n})$, we can compute a solution $\hat{\mathbf{x}}$ to the 1-LENA (1-LEN, respectively) instance, where the exact case holds when $\epsilon^h = \epsilon^{1e} = 0$.

Proof. Let \mathbf{f}^h be a feasible solution to the FHFA instance, so that error in demand, error in congestion, and error in homology can all be bounded by ϵ^h . Let $\hat{\mathbf{x}}_i(j) = \mathbf{f}^h(\hat{x}_i^j)$ denote the amount of flow routed through in the i th section of G^h . For an arbitrary equation $\hat{\mathbf{A}}(i)\hat{\mathbf{x}} - \hat{\mathbf{b}}(i)$ in 1-LENA, its error can be bounded as follows:

$$\begin{aligned} |\hat{\mathbf{A}}(i)\hat{\mathbf{x}} - \hat{\mathbf{b}}(i)| &= |\hat{\mathbf{A}}(i)\hat{\mathbf{x}} - \hat{\mathbf{A}}(i)\hat{\mathbf{x}}_i + \hat{\mathbf{A}}(i)\hat{\mathbf{x}}_i - \hat{\mathbf{b}}(i)| \\ &\leq |\hat{\mathbf{A}}(i)\hat{\mathbf{x}} - \hat{\mathbf{A}}(i)\hat{\mathbf{x}}_i| + |\hat{\mathbf{A}}(i)\hat{\mathbf{x}}_i - \hat{\mathbf{b}}(i)|. \end{aligned}$$

To bound the first term, we have

$$|\hat{\mathbf{A}}(i)\hat{\mathbf{x}} - \hat{\mathbf{A}}(i)\hat{\mathbf{x}}_i| \leq \|\hat{\mathbf{A}}(i)\|_1 \|\hat{\mathbf{x}} - \hat{\mathbf{x}}_i\|_{\max} \leq \|\hat{\mathbf{A}}(i)\|_1 \epsilon^h,$$

where the last step follows from the error in homology in G^h . To bound the second term, we have

$$\begin{aligned} &|\hat{\mathbf{A}}(i)\hat{\mathbf{x}}_i - \hat{\mathbf{b}}(i)| \\ &= |\mathbf{f}^h(s, J_i^+) - \mathbf{f}^h(s, J_i^-) - \hat{\mathbf{b}}(i)| \quad \text{by construction} \\ &\leq |\mathbf{f}^h(s, J_i^+) - (\mathbf{f}^h(e_i^+) + \mathbf{f}^h(\hat{b}_i))| + |\mathbf{f}^h(e_i^+) + \mathbf{f}^h(\hat{b}_i) - \mathbf{f}^h(s, J_i^-) - \hat{\mathbf{b}}(i)| \\ &\leq \epsilon^h + |\mathbf{f}^h(e_i^+) + \mathbf{f}^h(\hat{b}_i) - \mathbf{f}^h(s, J_i^-) - \hat{\mathbf{b}}(i)| \quad \text{by error in demand of vertex } J_i^+ \\ &\leq \epsilon^h + |\mathbf{f}^h(e_i^+) - \mathbf{f}^h(e_i^-)| + |\mathbf{f}^h(e_i^-) - \mathbf{f}^h(s, J_i^-)| + |\mathbf{f}^h(\hat{b}_i) - \hat{\mathbf{b}}(i)| \\ &\leq 4\epsilon^h, \end{aligned}$$

where the last step follows from error in homology between edge e_i^+, e_i^- , error in demand of vertex J_i^- , and error in congestion of fixed flow edge \hat{b}_i , respectively.

Combining the bounds of the two terms, we obtain the error of an arbitrary equation in 1-LENA as

$$|\hat{\mathbf{A}}(i)\hat{\mathbf{x}} - \hat{\mathbf{b}}(i)| \leq \left(\|\hat{\mathbf{A}}(i)\|_1 + 4 \right) \epsilon^h \leq O(\hat{N})\epsilon^h.$$

□

3.6 Flow Space

3.6.1 Overview

In the flow space, the starting problem is FHF (Definition 3.3.12), whose additional edge constraints include *homologous flow constraints* (Definition 3.3.9) and *fixed flow constraints* (Definition 3.3.10). The target problem is 2CF (Definition 3.3.7), which has no additional edge constraints. Therefore, we are tasked to drop additional edge constraints step by step in the flow space. Note that new additional edge constraints may emerge as we drop ones, for example, *selective flow constraints* (Definition 3.3.11).

Exact Case. To reduce instance A to instance B , we focus on dropping specific edge constraints. Let the graph of instance A be $G^A = (E^A, V^A)$ and let $S^A \subseteq E^A$ denote the set of edges in A with the constraints to be dropped. The reduced graph $G^B = (E^B, V^B)$ replaces each edge $e \in S^A$ with a constant-sized gadget.⁴ These gadgets simplify or remove the associated constraints.

The general reduction and solution mapping framework is shown in Figure 3.2.

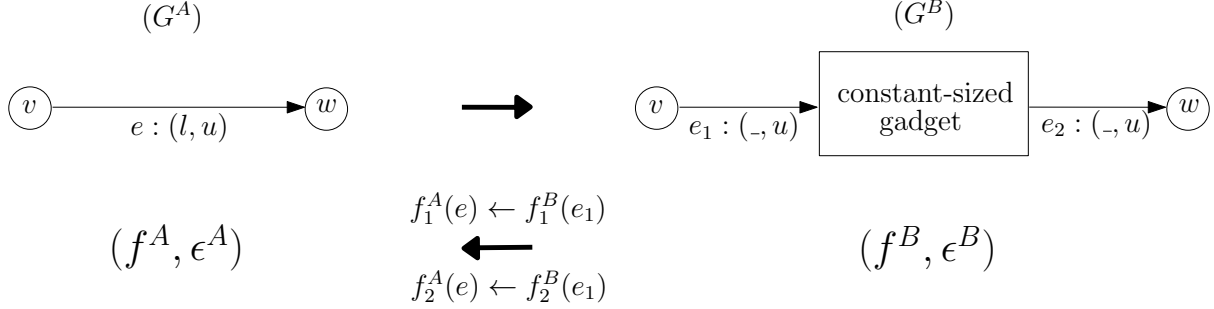


Figure 3.2: Reduction and solution mapping algorithms in the flow space.

Reduction:

For each special edge $e = (v, w) \in S^A$ with lower and upper capacities (l, u) ,⁵ the reduction replaces e with two edges e_1 and e_2 , along with a constant sized gadget between e_1 and e_2 . The upper capacity of e_1 and e_2 remains u ; and the lower capacity is either l or reduced to 0, depending on the reduction step.

Solution mapping:

If a solver returns f^B for instance B , the flow for instance A is constructed as follows:⁶

$$f_i^A(e) = f_i^B(e_1), \quad \text{for } e \in S^A, \quad i \in \{1, 2\}.$$

For regular edges in $e \in E^A \setminus S^A$, the flow is directly mapped:

$$f_i^A(e) = f_i^B(e), \quad \text{for } e \in E^A \setminus S^A, \quad i \in \{1, 2\}.$$

If the solver returns “infeasible” for instance B , then we return “infeasible” for instance A as well.

Based on the general reduction and solution mapping framework, the following results apply to every step within the flow space.

Lemma 3.6.1 (Problem Size and Runtime for Reduction Algorithm). *In the flow space, when reducing instance A to instance B , the number of edges in B satisfies $|E^B| = O(|E^A|)$. Moreover, the reduction algorithm takes runtime $|E^B| = O(|E^A|)$.*

Proof. There are at most $|E^A|$ edges in G^A that are replaced by gadgets of constant size. The total size of G^B is thus $O(|E^A|)$. The runtime is linear to the number of edges in G^B , thus $O(|E^B|)$. \square

⁴Only a few steps may utilize gadgets that deviate slightly from the general pattern. If this occurs, we will provide a note in the relevant section.

⁵The lower edge capacity is set to $l = 0$ unless the edge is a fixed-flow edge, in which case $l = u$.

⁶For 1-commodity flow problem, let $f_2 = \mathbf{0}$.

Lemma 3.6.2 (Runtime for Solution Mapping Algorithm). *In the flow space, mapping a solution to instance B to a solution to instance A takes $O(|E^A|)$ runtime.*

Proof. The runtime is linear to the number of edges in G^A , as solution mapping involves directly transferring or reconstructing flow values for each edge. Thus, the runtime is $O(|E^A|)$. \square

Approximate Case. We apply the same reduction and solution mapping algorithm for approximate problems. Hence, Lemma 3.6.1, 3.6.2 also hold for the approximate case. Additional analyses focus on the error propagation during solution mapping. Specifically, if $(\mathbf{f}^B, \epsilon^B)$ is an approximate solution to instance B with error ϵ^B , then the mapped solution $(\mathbf{f}^A, \epsilon^A)$ to instance A has an error ϵ^A that can be upper bounded. In particular, when $\epsilon^A, \epsilon^B = 0$, the statement for the approximate case is reduced to the exact case.

Different types of errors are analyzed separately for problems in the flow space. As the capacity constraints and the flow conservation constraints are required for all problems in the flow space, the corresponding error in congestion and error in demand are involved in the error analysis for all steps. The following error bounds are applicable at every step within the flow space.

Lemma 3.6.3 (Error in Congestion). *Let $\epsilon_l^A, \epsilon_u^A$ and $\epsilon_l^B, \epsilon_u^B$ be the errors in congestion w.r.t. lower and upper capacity for \mathbf{f}^A and \mathbf{f}^B , respectively. If applying the solution mapping algorithm, then for edges in G^A with $l = 0$,*

$$\epsilon_l^A = 0,$$

and for all edges in G^A ,

$$\epsilon_u^A \leq O(\epsilon_u^B).$$

Proof. With the non-negativity constraints satisfied by \mathbf{f}^A for edges in G^A with $l = 0$, ϵ_l^A can be trivially bounded, regardless of ϵ_l^B .

For the upper edge capacity, the error bounds are straightforward if e is a regular edge, as the flow is copied directly from G^A and G^B . However, if e is a special edge, the upper-capacity error in G^A remains bounded by ϵ^B , as e_1 and e have the same upper edge capacity u according to Figure 3.2. \square

Remark that for edges in G^A with $l = u$, lower-capacity error in congestion requires case-by-case analysis.

The following helper lemma is frequently used to analyze error in demand across various reduction steps.

Lemma 3.6.4 (Error in Demand). *Let $\epsilon_{d,i}^A$ and $\epsilon_{d,i}^B$ be the error in demand of \mathbf{f}^A and \mathbf{f}^B for commodity i , $i \in \{1, 2\}$, respectively. If applying the solution mapping algorithm, we have*

$$\epsilon_{d,i}^A \leq \epsilon_{d,i}^B + \max_{v \in V^A \setminus \{s_i, t_i\}} \sum_{e=(u,v) \in S^A} |\mathbf{f}_i^B(e_1) - \mathbf{f}_i^B(e_2)|.$$

Proof. By Definition 3.3.16, error in demand for vertices other than the source/sink nodes is defined as

$$\left| \sum_{u:(u,v) \in E} \mathbf{f}_i(u,v) - \sum_{w:(v,w) \in E} \mathbf{f}_i(v,w) \right| \leq \epsilon, \quad \forall v \in V \setminus \{s_i, t_i\}, \quad i \in \{1, 2\}.$$

As illustrated in Figure 3.3, consider an arbitrary vertex $v \in V^B \setminus \{s_i, t_i\}$ with a number of incoming and outgoing edges. Some of these edges are derived from the reduction of special edges in G^A , as indicated by dashed boxes in the figure.

When mapping \mathbf{f}^B to \mathbf{f}^A , the incoming/outgoing flows of v along regular edges remain unchanged, while flows along special edges may vary due to the gap between e_1 and e_2 .




Figure 3.3: An illustration of the incoming and outgoing flows of vertex $v \in V^B \setminus \{s_i, t_i\}$ when get reduced to G^B .

For flow of commodity $i \in \{1, 2\}$, we have

$$\begin{aligned}
\epsilon_{d,i}^A &= \max_{v \in V^A \setminus \{s_i, t_i\}} \left| \sum_{(u,v) \in EA} \mathbf{f}_i^A(u, v) - \sum_{(v,w) \in EA} \mathbf{f}_i^A(v, w) \right| \\
&\stackrel{(1)}{=} \max_{v \in V^A \setminus \{s_i, t_i\}} \left| \left(\sum_{(u,v) \in SA} \mathbf{f}_i^A(u, v) + \sum_{(u,v) \in EA \setminus SA} \mathbf{f}_i^A(u, v) \right) - \left(\sum_{(v,w) \in SA} \mathbf{f}_i^A(v, w) + \sum_{(v,w) \in EA \setminus SA} \mathbf{f}_i^A(v, w) \right) \right| \\
&\stackrel{(2)}{=} \max_{v \in V^A \setminus \{s_i, t_i\}} \left| \left(\sum_{e=(u,v) \in SA} \mathbf{f}_i^B(e_1) + \sum_{(u,v) \in EA \setminus SA} \mathbf{f}_i^B(u, v) \right) - \left(\sum_{e=(v,w) \in SA} \mathbf{f}_i^B(e_1) + \sum_{(v,w) \in EA \setminus SA} \mathbf{f}_i^B(v, w) \right) \right| \\
&\stackrel{(3)}{\leq} \epsilon_{d,i}^B + \max_{v \in V^A \setminus \{s_i, t_i\}} \left| \left(\sum_{e=(u,v) \in SA} \mathbf{f}_i^B(e_1) + \sum_{(u,v) \in EA \setminus SA} \mathbf{f}_i^B(u, v) \right) - \left(\sum_{e=(u,v) \in SA} \mathbf{f}_i^B(e_2) + \sum_{(u,v) \in EA \setminus SA} \mathbf{f}_i^B(u, v) \right) \right| \\
&= \epsilon_{d,i}^B + \max_{v \in V^A \setminus \{s_i, t_i\}} \left| \sum_{e=(u,v) \in SA} \mathbf{f}_i^B(e_1) - \sum_{e=(u,v) \in SA} \mathbf{f}_i^B(e_2) \right| \\
&= \epsilon_{d,i}^B + \max_{v \in V^A \setminus \{s_i, t_i\}} \left| \sum_{e=(u,v) \in SA} (\mathbf{f}_i^B(e_1) - \mathbf{f}_i^B(e_2)) \right| \\
&\leq \epsilon_{d,i}^B + \max_{v \in V^A \setminus \{s_i, t_i\}} \sum_{e=(u,v) \in SA} |\mathbf{f}_i^B(e_1) - \mathbf{f}_i^B(e_2)|.
\end{aligned}$$

For step (1), we separate special edges from regular edges that are incident to vertex v . For step (2), we apply the solution mapping algorithm from \mathbf{f}^B back to \mathbf{f}^A . For step (3), we can replace the sum of outgoing flows of v by the sum of its incoming flows with an error in demand of instance B being introduced. \square

3.6.2 FHF(A) to FPHF(A)

In this section, we show the reduction from an FHF instance $(G^h, H^h, \mathcal{H}^h, \mathbf{u}^h, s, t)$ to an FPHF instance $(G^p, F^p, \mathcal{H}^p, \mathbf{u}^p, s, t)$. This step does not drop any additional edge constraints, but simplifies the edge constraints for future dropping steps. More specifically, in G^h , it is possible for more than two edges to be constrained to route the same amount of flow. The goal is to reduce G^h to G^p such that each edge is homologous to at most one other edge.

Suppose that $(v_1, w_1), \dots, (v_k, w_k) \in E^h$ belong to a set of homologous edges of size k in G^h . As shown in Figure 3.4,⁷ we replace $(v_i, w_i), i \in \{2, \dots, k-1\}$ by two edges (v_i, z_i) and (z_i, w_i) such that z_i is a new vertex incident only to these two edges, and edge capacities of the two new edges are the same as that of the original edge (v_i, w_i) . Then, we can construct $k-1$ pairs of homologous edges: (v_1, w_1) and $(v_2, z_2); (z_2, w_2)$ and $(v_3, z_3); \dots; (z_i, w_i)$ and $(v_{i+1}, z_{i+1}); \dots; (z_{k-1}, w_{k-1})$ and (v_k, w_k) . In addition, no reduction is performed on non-homologous edges in G^h , and we trivially copy these edges to G^p . Compared with Figure 3.2, the constant-sized gadget in this step corresponds to the inserted vertex z_i , as indicated by the dashed boxes.

Lemma 3.6.5 (FHF to FPHF). *Given an FHF instance $(G^h, F^h, \mathcal{H}^h, \mathbf{u}^h, s, t)$, we can construct, in time $O(|E^p|)$, an FPHF instance $(G^p, F^p, \mathcal{H}^p, \mathbf{u}^p, s, t)$ such that*

$$|E^p| = O(|E^h|), \quad \|\mathbf{u}^p\|_{\max} = \|\mathbf{u}^h\|_{\max}.$$

If the FHF instance has a solution, then the FPHF instance has a solution.

Proof. The runtime of reduction and the bound for $|E^p|$ follow from Lemma 3.6.1. Moreover, as the reduction of this step only inserts new vertices without modifying edge capacities, we have $\|\mathbf{u}^p\|_{\max} = \|\mathbf{u}^h\|_{\max}$.

If \mathbf{f}^h is a feasible flow to the FHF instance, it is easy to derive a solution \mathbf{f}^p to the FPHF instance. Concretely, for any homologous edge $e \in E^h$ that is split into two edges $e_1, e_2 \in E^p$, we set $\mathbf{f}^p(e_1) = \mathbf{f}^p(e_2) = \mathbf{f}^h(e)$. Since the vertex between e_1 and e_2 is only incident to these two edges, then the conservation of flows is satisfied on the inserted vertices. Moreover, since the edge capacities of e_1 and e_2 are the same as that of e , and they route the same amount of flows, thus the capacity constraint is also satisfied on the split edges. In addition, the flow conservation and capacity constraints are trivially satisfied for the rest vertices and edges. Therefore, \mathbf{f}^p is a feasible flow to the FPHF instance. \square

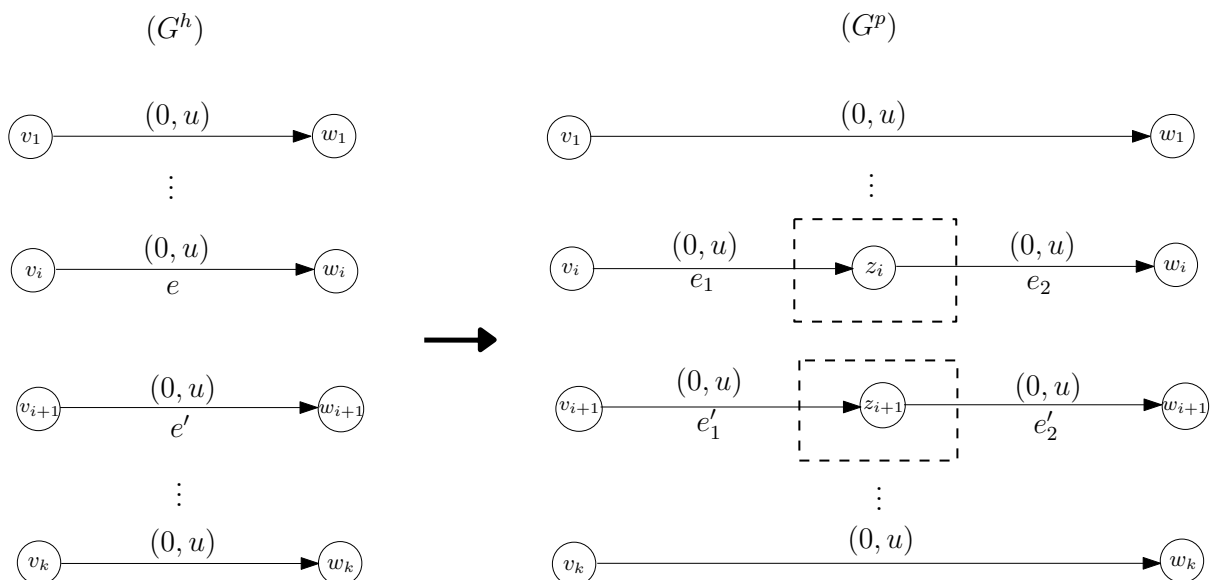


Figure 3.4: The reduction from FHF to FPHF.

⁷We use $(0, u)$ for an non-fixed flow edge of capacity u .

In the approximate case, we first formally define FPHFA by applying the corresponding error notions (Definition 3.3.16) to FPHF.

Definition 3.6.6 (FPHF Approximate Problem (FPHFA)). An FPHFA instance is given by an FPHF instance $(G, F, \mathcal{H}, \mathbf{u}, s, t)$ as in Definition 3.3.13, and an error parameter $\epsilon \in [0, 1]$, which we collect in a tuple $(G, F, \mathcal{H}, \mathbf{u}, s, t, \epsilon)$. We say an algorithm solves the FPHFA problem, if, given any FPHFA instance, it returns a flow $\mathbf{f} \geq \mathbf{0}$ that satisfies

$$\mathbf{u}(e) - \epsilon \leq \mathbf{f}(e) \leq \mathbf{u}(e) + \epsilon, \quad \forall e \in F, \quad (\text{error in congestion})$$

$$0 \leq \mathbf{f}(e) \leq \mathbf{u}(e) + \epsilon, \quad \forall e \in E \setminus F, \quad (\text{error in congestion})$$

$$\left| \sum_{u:(u,v) \in E} \mathbf{f}(u,v) - \sum_{w:(v,w) \in E} \mathbf{f}(v,w) \right| \leq \epsilon, \quad \forall v \in V \setminus \{s, t\}, \quad (\text{error in demand})$$

$$|\mathbf{f}(v,w) - \mathbf{f}(y,z)| \leq \epsilon, \quad \forall (v,w), (y,z) \in H_i, \quad \forall H_i \in \mathcal{H}, \quad (\text{error in pairwise homology})$$

or it correctly declares that the associated FPHF instance is infeasible.

Lemma 3.6.7 (FHFA to FPHFA). *Given an FHFA instance $(G^h, F^h, \mathcal{H}^h, \mathbf{u}^h, s, t, \epsilon^h)$, we can reduce it to an FPHFA instance $(G^p, F^p, \mathcal{H}^p, \mathbf{u}^p, s, t, \epsilon^p)$ by letting*

$$\epsilon^p = \Omega\left(\frac{1}{|E^h|}\right) \epsilon^h,$$

and using Lemma 3.6.5 to construct an FPHF instance $(G^p, F^p, \mathcal{H}^p, \mathbf{u}^p, s, t)$ from the FHF instance $(G^h, H^h, \mathcal{H}^h, \mathbf{u}^h, s, t)$. If \mathbf{f}^p is a solution to the FPHFA (FPHF) instance, then in time $O(|E^h|)$, we can compute a solution \mathbf{f}^h to the FHFA (FHF, respectively) instance, where the exact case holds when $\epsilon^p = \epsilon^h = 0$.

Proof. To bound error in congestion, we have $\epsilon_u^h \leq O(\epsilon^p)$ according to Lemma 3.6.3. For ϵ_l^h , we have $\epsilon_l^h = 0$ if $e \in E^p \setminus F^p$. If $e \in F^p$, since we just copy them without making reductions for fixed flow edges in this step, we also have $\epsilon_l^h \leq O(\epsilon^p)$.

To bound error in demand, by Lemma 3.6.4 and error in demand of auxiliary vertex z , we have

$$\epsilon_d^h \leq \epsilon^p + \max_{v \in V^h \setminus \{s, t\}} \sum_{e=(u,v) \in H^h} |\mathbf{f}^p(e_1) - \mathbf{f}^p(e_2)| \leq O(|E^h|) \epsilon^p.$$

Finally, we bound error in homology in G^h . It is observed that error in homology in G^h for a homologous edge set of size k gets accumulated by $(k-1)$ times of error in pair homology in G^p , and we have $k \leq |E_h|$. Thus, $\epsilon_h^h \leq O(|E^h|) \epsilon^p$. Putting it all together, we have

$$\epsilon^h = \max\{\epsilon_l^h, \epsilon_u^h, \epsilon_d^h, \epsilon_h^h\} \leq O(|E^h|) \epsilon^p.$$

□

3.6.3 FPHF(A) to SFF(A): Dropping Homologous Flow Constraints

We show the reduction from an FPHF instance $(G^p, F^p, \mathcal{H}^p, \mathbf{u}^p, s, t)$ to an SFF instance $(G^s, F^s, S_1, S_2, \mathbf{u}^s, s_1, t_1, s_2, t_2)$. Assume that $\{e, \hat{e}\} \in \mathcal{H}^p$ is an arbitrary pair of homologous edges in G^p . As shown in Figure 3.5, we map $\{e, \hat{e}\}$ in G^p to a constant-sized

Moreover, we set

$$\mathbf{f}_2^s(e_3) = \mathbf{f}_2^s(e_4) = \mathbf{f}_2^s(e_5) = \mathbf{f}_2^s(\hat{e}_4) = \mathbf{f}_2^s(\hat{e}_5) = u - \mathbf{f}^p(e) \leq u,$$

$$\mathbf{f}_1^s(e_3) = \mathbf{f}_1^s(e_4) = \mathbf{f}_1^s(e_5) = \mathbf{f}_1^s(\hat{e}_4) = \mathbf{f}_1^s(\hat{e}_5) = 0.$$

Then it is obvious that \mathbf{f}^s satisfies the selective constraint and the capacity constraint on edges e_3, e_5, \hat{e}_5 , and satisfies the flow conservation constraint of commodity 2 on vertices vw, vw', yz, yz' .

It remains to verify if \mathbf{f}^s satisfies the fixed flow constraint on edges e_4, \hat{e}_4 . According to the above constructions, we have (we abuse the notation to also let $\mathbf{f}^s = \mathbf{f}_1^s + \mathbf{f}_2^s$)

$$\mathbf{f}^s(e_4) = \mathbf{f}_1^s(e_4) + \mathbf{f}_2^s(e_4) = \mathbf{f}^p(e) + (u - \mathbf{f}^p(e)) = u,$$

$$\mathbf{f}^s(\hat{e}_4) = \mathbf{f}_1^s(\hat{e}_4) + \mathbf{f}_2^s(\hat{e}_4) = \mathbf{f}^p(\hat{e}) + (u - \mathbf{f}^p(\hat{e})) = u.$$

- For any non-homologous edge $e' \in E^p$, we also have $e' \in E^s$ since no reduction is made on this edge, and we copy it trivially to G^s . We set

$$\mathbf{f}_1^s(e') = \mathbf{f}^p(e'), \quad \mathbf{f}_2^s(e') = 0.$$

Since \mathbf{f}^p is a feasible flow in G^p , it is easy to check that \mathbf{f}^s also satisfies the selective constraint for commodity 1 and the capacity constraint on non-homologous edges, as well as conservation of flows on vertices incident to non-homologous edges.

To conclude, \mathbf{f}^s is a feasible flow to the SFF instance. \square

Definition 3.6.9 (SFF Approximate Problem (SFFA)). An SFFA instance is given by an SFF instance $(G, F, S_1, S_2, \mathbf{u}, s_1, t_1, s_2, t_2)$ as in Definition 3.3.14, and an error parameter $\epsilon \in [0, 1]$, which we collect in a tuple $(G, F, S_1, S_2, \mathbf{u}, s_1, t_1, s_2, t_2, \epsilon)$. We say an algorithm solves the SFFA problem, if, given any SFFA instance, it returns a pair of flows $\mathbf{f}_1, \mathbf{f}_2 \geq \mathbf{0}$ that satisfies

$$\mathbf{u}(e) - \epsilon \leq \mathbf{f}_1(e) + \mathbf{f}_2(e) \leq \mathbf{u}(e) + \epsilon, \quad \forall e \in F, \quad (\text{error in congestion})$$

$$0 \leq \mathbf{f}_1(e) + \mathbf{f}_2(e) \leq \mathbf{u}(e) + \epsilon, \quad \forall e \in E \setminus F, \quad (\text{error in congestion})$$

$$\left| \sum_{u:(u,v) \in E} \mathbf{f}_i(u, v) - \sum_{w:(v,w) \in E} \mathbf{f}_i(v, w) \right| \leq \epsilon, \quad \forall v \in V \setminus \{s_i, t_i\}, \quad i \in \{1, 2\},$$

(error in demand)

$$\mathbf{f}_{\bar{i}}(e) \leq \epsilon, \quad \forall e \in S_i, \quad \bar{i} = \{1, 2\} \setminus i, \quad (\text{error in selection})$$

or it correctly declares that the associated SFF instance is infeasible.

Lemma 3.6.10 (FPHFA to SFFA). Given an FPHFA instance $(G^p, F^p, \mathcal{H}^p, \mathbf{u}^p, s, t, \epsilon^p)$, we can reduce it to an SFFA instance $(G^s, F^s, S_1, S_2, \mathbf{u}^s, s_1, t_1, s_2, t_2, \epsilon^s)$ by letting

$$\epsilon^s = \Omega\left(\frac{1}{|E^p|}\right) \epsilon^p,$$

and using Lemma 3.6.8 to construct an SFF instance $(G^s, F^s, S_1, S_2, \mathbf{u}^s, s_1, t_1, s_2, t_2)$ from the FPHF instance $(G^p, F^p, \mathcal{H}^p, \mathbf{u}^p, s, t)$. If \mathbf{f}^s is a solution to the SFFA (SFF) instance, then in time $O(|E^p|)$, we can compute a solution \mathbf{f}^p to the FPHFA (FPHF, respectively) instance, where the exact case holds when $\epsilon^s = \epsilon^p = 0$.

Proof. To bound error in congestion, we have $\epsilon_u^p \leq O(\epsilon^s)$ according to Lemma 3.6.3. For ϵ_l^p , we have $\epsilon_l^p = 0$ if $e \in E^p \setminus F^p$, whereas if $e \in F^p$, we have

$$\begin{aligned} \mathbf{f}^p(e) &= \mathbf{f}_1^s(e) && \text{by solution mapping for non-homologous edges} \\ &\geq \mathbf{f}^s(e) - \epsilon^s && \text{by error in selection} \\ &\geq \mathbf{u}^s(e) - 2\epsilon^s && \text{by error in congestion} \\ &= \mathbf{u}^p(e) - 2\epsilon^s. \end{aligned}$$

Hence, $\epsilon_l^p \leq 2\epsilon^s$.

To bound error in demand, by Lemma 3.6.4 and solution mapping algorithm, we have

$$\epsilon_d^p \leq \epsilon^s + \max_{w \in V^p \setminus \{s, t\}} \sum_{e=(v, w) \in H^p} |\mathbf{f}_1^s(e_1) - \mathbf{f}_1^s(e_2)|, \quad (3.3)$$

so the key is to bound $|\mathbf{f}_1^s(e_1) - \mathbf{f}_1^s(e_2)|$. Again, we can decompose the difference into several error terms.

$$\begin{aligned} |\mathbf{f}_1^s(e_1) - \mathbf{f}_1^s(e_2)| &\leq |\mathbf{f}_1^s(e_1) + \mathbf{f}_1^s(e_3) - \mathbf{f}_1^s(e_4)| + |\mathbf{f}_1^s(e_4) - \mathbf{f}_1^s(e_3) - \mathbf{f}_1^s(e_2)| \\ &\leq \epsilon^s + |\mathbf{f}_1^s(e_4) - \mathbf{f}_1^s(e_3) - \mathbf{f}_1^s(e_2)| && \text{by error in demand of vertex } vw \\ &\leq \epsilon^s + |\mathbf{f}_1^s(e_4) - \mathbf{f}_1^s(e_2) - \mathbf{f}_1^s(e_5)| + |\mathbf{f}_1^s(e_5) - \mathbf{f}_1^s(e_3)| \\ &\leq 2\epsilon^s + |\mathbf{f}_1^s(e_5) - \mathbf{f}_1^s(e_3)| && \text{by error in demand of vertex } vw' \\ &\leq 2\epsilon^s + \max\{|\mathbf{f}_1^s(e_5)|, |\mathbf{f}_1^s(e_3)|\} \\ &\leq 3\epsilon^s, && \text{by error in type of edge } e_3 \text{ and } e_5 \end{aligned}$$

Since $|H^p| = O(|E^p|)$, plugging the above back to Eq. (3.3), we have $\epsilon_d^p \leq O(|E^p|)\epsilon^s$.

Finally, we bound error in pairwise homology in G^p . For any pair of homologous edges,

$$\begin{aligned} |\mathbf{f}_1^s(e_1) - \mathbf{f}_1^s(\hat{e}_1)| &= |(\mathbf{f}^s(e_1) - \mathbf{f}_2^s(e_1)) - (\mathbf{f}^s(\hat{e}_1) - \mathbf{f}_2^s(\hat{e}_1))| \\ &\leq |\mathbf{f}^s(e_1) - \mathbf{f}^s(\hat{e}_1)| + |\mathbf{f}_2^s(e_1) - \mathbf{f}_2^s(\hat{e}_1)| \\ &\leq \epsilon^s + |\mathbf{f}^s(e_1) - \mathbf{f}^s(\hat{e}_1)| && \text{by error in type of edge } e_1 \text{ and } \hat{e}_1 \\ &\leq \epsilon^s + |(\mathbf{f}^s(e_1) + \mathbf{f}^s(e_3)) - (\mathbf{f}^s(\hat{e}_1) + \mathbf{f}^s(\hat{e}_3))| + |\mathbf{f}^s(e_3) - \mathbf{f}^s(\hat{e}_3)|. \end{aligned}$$

To analyze $|(\mathbf{f}^s(e_1) + \mathbf{f}^s(e_3)) - (\mathbf{f}^s(\hat{e}_1) + \mathbf{f}^s(\hat{e}_3))|$, we decompose as follows:

$$\begin{aligned} &|(\mathbf{f}^s(e_1) + \mathbf{f}^s(e_3)) - (\mathbf{f}^s(\hat{e}_1) + \mathbf{f}^s(\hat{e}_3))| \\ &\leq |\mathbf{f}^s(e_1) + \mathbf{f}^s(e_3) - \mathbf{f}^s(e_4)| + |\mathbf{f}^s(e_4) - (\mathbf{f}^s(\hat{e}_1) + \mathbf{f}^s(\hat{e}_3))| \\ &\leq 2\epsilon^s + |\mathbf{f}^s(e_4) - (\mathbf{f}^s(\hat{e}_1) + \mathbf{f}^s(\hat{e}_3))| \\ &\quad \text{by error in demand of vertex } vw \text{ for two commodities} \\ &\leq 2\epsilon^s + |\mathbf{f}^s(\hat{e}_4) - (\mathbf{f}^s(\hat{e}_1) + \mathbf{f}^s(\hat{e}_3))| + |\mathbf{f}^s(e_4) - \mathbf{f}^s(\hat{e}_4)| \\ &\leq 4\epsilon^s + |\mathbf{f}^s(e_4) - \mathbf{f}^s(\hat{e}_4)| && \text{by error in demand of vertex } yz \text{ for two commodities} \\ &\leq 6\epsilon^s, && \text{by error in congestion on } e_4 \text{ and } \hat{e}_4 \end{aligned}$$

For $|\mathbf{f}^s(e_3) - \mathbf{f}^s(\hat{e}_3)|$, similarly, we have

$$\begin{aligned} |\mathbf{f}^s(e_3) - \mathbf{f}^s(\hat{e}_3)| &= |(\mathbf{f}_1^s(e_3) + \mathbf{f}_2^s(e_3)) - (\mathbf{f}_1^s(\hat{e}_3) + \mathbf{f}_2^s(\hat{e}_3))| \\ &\leq |\mathbf{f}_1^s(e_3) - \mathbf{f}_1^s(\hat{e}_3)| + |\mathbf{f}_2^s(e_3) - \mathbf{f}_2^s(\hat{e}_3)| \\ &\leq \epsilon^s + |\mathbf{f}_2^s(e_3) - \mathbf{f}_2^s(\hat{e}_3)| && \text{by error in type in } G^s \text{ and } e_3, \hat{e}_3 \\ &\leq 4\epsilon^s, && \text{by symmetry from the bound of } |\mathbf{f}_1^s(e_1) - \mathbf{f}_1^s(e_2)| \end{aligned}$$

Therefore, we can bound the error in pair homology by

$$\epsilon_h^p \leq 11\epsilon^s.$$

Putting it altogether, we have

$$\epsilon^p = \max\{\epsilon_l^p, \epsilon_u^p, \epsilon_d^p, \epsilon_h^p\} \leq O(|E^p|) \epsilon^s.$$

□

3.6.4 SFF(A) to 2CFF(A): Dropping Selective Flow Constraints

We show the reduction from an SFF instance $(G^s, F^s, S_1, S_2, \mathbf{u}^s, s_1, t_1, s_2, t_2)$ to a 2CFF instance $(G^f, F^f, \mathbf{u}^f, s_1, t_1, s_2, t_2)$. Assume that $e \in S_i$ is an arbitrary selective edge for commodity i in G^p . As shown in Figure 3.6, we map e in G^s to a gadget consisting of edges $\{e_1, e_2, e_3, e_4, e_5\}$ in G^f . Note that a selective edge e can be either a fixed flow edge or a non-fixed flow edge. In Figure 3.6, $l = u$ if e is a fixed flow edge and $l = 0$ if e is a non-fixed flow edge. Moreover, no reduction is performed on non-selective edges in G^s , and we trivially copy these edges to G^f .

The key idea to remove the selective requirement is utilizing edge directions and the source-sink pair (s_i, t_i) to simulate a selective edge e for commodity i . More specifically, in the gadget, the flow of commodity i routes through three directed paths: (1) $e_1 \rightarrow e_4$, (2) $e_5 \rightarrow e_3 \rightarrow e_4$, (3) $e_5 \rightarrow e_2$. The selective requirement is realized because e_4 is the only outgoing edge of xy and only the flow of commodity i is allowed in e_4 (since its tail is t_i), thus in e_1 . Similarly, e_5 is the only incoming edge of xy' and only the flow of commodity i is allowed in e_5 , thus in e_2 . In addition, to ensure that e_1 and e_2 route the same amount of flow, flows in e_4 and e_5 must be equal by the conservation of flows. Therefore, we impose the fixed flow constraint on e_4 and e_5 by setting the fixed flow to be u . We remove e_3 if e is a fixed flow edge (in which case e_3 has capacity 0), and set the capacity of e_3 to be u otherwise (in which case $u - l = u - 0 = u$).




Figure 3.6: The reduction from SFF to 2CFF. $l = u$ if e is a fixed flow edge, and $l = 0$ if e is a non-fixed flow edge.

Lemma 3.6.11 (SFF to 2CFF). *Given an SFF instance $(G^s, F^s, S_1, S_2, \mathbf{u}^s, s_1, t_1, s_2, t_2)$, we can construct, in time $O(|E^s|)$, a 2CFF instance $(G^f, F^f, \mathbf{u}^f, s_1, t_1, s_2, t_2)$ such that*

$$|E^f| = O(|E^s|), \quad \|\mathbf{u}^f\|_{\max} = \|\mathbf{u}^s\|_{\max}.$$

If the SFF instance has a solution, then the 2CFF instance has a solution.

Proof. If \mathbf{f}^s is a solution to the SFF instance, it is easy to derive a solution \mathbf{f}^f to the 2CFF instance as follows:

- For any selective edge $e \in S_i$, in its corresponding gadget in G^f , we set

$$l \leq \mathbf{f}_i^f(e_1) = \mathbf{f}_i^f(e_3) = \mathbf{f}_i^s(e) \leq u,$$

$$\mathbf{f}_i^f(e_4) = \mathbf{f}_i^f(e_5) = u,$$

$$0 \leq \mathbf{f}_i^f(e_2) = u - \mathbf{f}_i^s(e) \leq u - l,$$

where $l = 0$ if e is a non-fixed flow edge, or $l = u = \mathbf{u}^s(e)$ if e is a fixed flow edge. And we set

$$\mathbf{f}_{\bar{i}}^f(e_1) = \mathbf{f}_{\bar{i}}^f(e_2) = \mathbf{f}_{\bar{i}}^f(e_3) = \mathbf{f}_{\bar{i}}^f(e_4) = \mathbf{f}_{\bar{i}}^f(e_5) = 0, \quad \bar{i} = \{1, 2\} \setminus i.$$

It is obvious that \mathbf{f}^f satisfies the capacity constraint on edges (e_1, \dots, e_5) , and the flow conservation constraint on vertices xy, xy' .

- For any non-selective edge $e' \in E^s \setminus (S_1 \cup S_2)$, we also have $e' \in E^f$ since no reduction is made on this edge, and we copy it trivially to G^f . We set

$$\mathbf{f}_i^f(e') = \mathbf{f}_i^s(e'), \quad i \in \{1, 2\}.$$

Since \mathbf{f}^s is a feasible flow in G^s , it is easy to check that \mathbf{f}^f also satisfies the capacity constraint on non-selective edges, as well as the flow conservation constraint on vertices incident to non-selective edges. Moreover, if e' is a fixed flow edge, the fixed flow edge constraint is also satisfied.

To conclude, \mathbf{f}^f is a feasible flow to the 2CFF instance. □

Definition 3.6.12 (2CFF Approximate Problem (2CFFA)). A 2CFFA instance is given by a 2CFF instance $(G, F, \mathbf{u}, s_1, t_1, s_2, t_2)$ as in Definition 3.3.15, and an error parameter $\epsilon \in [0, 1]$, which we collect in a tuple $(G, F, \mathbf{u}, s_1, t_1, s_2, t_2, \epsilon)$. We say an algorithm solves the 2CFFA problem, if, given any 2CFFA instance, it returns a pair of flows $\mathbf{f}_1, \mathbf{f}_2 \geq \mathbf{0}$ that satisfies

$$\mathbf{u}(e) - \epsilon \leq \mathbf{f}_1(e) + \mathbf{f}_2(e) \leq \mathbf{u}(e) + \epsilon, \quad \forall e \in F, \quad (\text{error in congestion})$$

$$0 \leq \mathbf{f}_1(e) + \mathbf{f}_2(e) \leq \mathbf{u}(e) + \epsilon, \quad \forall e \in E \setminus F, \quad (\text{error in congestion})$$

$$\left| \sum_{u:(u,v) \in E} \mathbf{f}_i(u, v) - \sum_{w:(v,w) \in E} \mathbf{f}_i(v, w) \right| \leq \epsilon, \quad \forall v \in V \setminus \{s_i, t_i\}, \quad i \in \{1, 2\},$$

(error in demand)

or it correctly declares that the associated 2CFF instance is infeasible.

Lemma 3.6.13 (SFFA to 2CFFA). *Given an SFFA instance $(G^s, F^s, S_1, S_2, \mathbf{u}^s, s_1, t_1, s_2, t_2, \epsilon^s)$, we can reduce it to a 2CFFA instance $(G^f, F^f, \mathbf{u}^f, s_1, t_1, s_2, t_2, \epsilon^f)$ by letting*

$$\epsilon^f = \Omega\left(\frac{1}{|E^p|}\right) \epsilon^s,$$

and using Lemma 3.6.11 to construct a 2CFF instance $(G^f, F^f, \mathbf{u}^f, s_1, t_1, s_2, t_2)$ from the SFF instance $(G^s, F^s, S_1, S_2, \mathbf{u}^s, s_1, t_1, s_2, t_2)$. If \mathbf{f}^f is a solution to the 2CFFA (2CFF) instance, then in time $O(|E^s|)$, we can compute a solution \mathbf{f}^s to the SFFA (SFF, respectively) instance, where the exact case holds when $\epsilon^f = \epsilon^s = 0$.

Proof. To bound error in congestion, we have $\epsilon_u^s \leq O(\epsilon^f)$ according to Lemma 3.6.3. For ϵ_l^s , we have $\epsilon_l^s = 0$ if $e \in E^s \setminus F^s$. If $e \in F^s$, since we just copy them without making reductions for fixed flow edges in this step, we also have $\epsilon_l^s \leq O(\epsilon^f)$.

To bound error in demand, we analyze commodity 1 and commodity 2 separately. Let S_i be the set of selective edges for commodity i , $i \in \{1, 2\}$, and let $S = S_1 \cup S_2$. Also, let $\epsilon_{d,i}^s$ be the error in demand for commodity i . By Lemma 3.6.4 and solution mapping algorithm, we have

$$\begin{aligned} \epsilon_{d,i}^s &\leq \epsilon^f + \max_{w \in V^p \setminus \{s_i, t_i\}} \sum_{e=(v,w) \in S} |\mathbf{f}_i^s(e_1) - \mathbf{f}_i^s(e_2)| \\ &= \epsilon^f + \max_{w \in V^p \setminus \{s_i, t_i\}} \left(\sum_{e=(v,w) \in S_i} |\mathbf{f}_i^f(e_1) - \mathbf{f}_i^f(e_2)| + \sum_{e=(v,w) \in S_{\bar{i}}} |\mathbf{f}_i^f(e_1) - \mathbf{f}_i^f(e_2)| \right). \end{aligned} \quad (3.4)$$

The problem is reduced to bounding $|\mathbf{f}_i^f(e_1) - \mathbf{f}_i^f(e_2)|$.

$$\begin{aligned} &|\mathbf{f}_i^f(e_1) - \mathbf{f}_i^f(e_2)| \\ &\leq |\mathbf{f}_i^f(e_1) + \mathbf{f}_i^f(e_3) - \mathbf{f}_i^f(e_4)| + |\mathbf{f}_i^f(e_4) - \mathbf{f}_i^f(e_3) - \mathbf{f}_i^f(e_2)| \\ &\leq \epsilon^f + |\mathbf{f}_i^f(e_4) - \mathbf{f}_i^f(e_3) - \mathbf{f}_i^f(e_2)| \quad \text{by error in demand of vertex } xy \\ &\leq \epsilon^f + |\mathbf{f}_i^f(e_4) - \mathbf{f}_i^f(e_5)| + |\mathbf{f}_i^f(e_5) - \mathbf{f}_i^f(e_3) - \mathbf{f}_i^f(e_2)| \\ &\leq 2\epsilon^f + |\mathbf{f}_i^f(e_4) - \mathbf{f}_i^f(e_5)|, \quad \text{by error in demand of vertex } xy' \end{aligned}$$

If $e \in S_{\bar{i}}$, then, by error in demand of vertex $t_{\bar{i}}$ and $s_{\bar{i}}$, we have

$$|\mathbf{f}_i^f(e_4) - \mathbf{f}_i^f(e_5)| \leq 2\epsilon^f.$$

However, if $e \in S_i$, we have

$$\begin{aligned} |\mathbf{f}_i^f(e_4) - \mathbf{f}_i^f(e_5)| &= |(\mathbf{f}^f(e_4) - \mathbf{f}_i^f(e_4)) - (\mathbf{f}^f(e_5) - \mathbf{f}_i^f(e_5))| \\ &\leq |\mathbf{f}^f(e_4) - \mathbf{f}^f(e_5)| + |\mathbf{f}_i^f(e_4) - \mathbf{f}_i^f(e_5)| \\ &\leq 2\epsilon^f + |\mathbf{f}_i^f(e_4) - \mathbf{f}_i^f(e_5)| \quad \text{by error in congestion of edge } e_4 \text{ and } e_5 \\ &\leq 4\epsilon^f, \quad \text{by error in demand of vertex } t_i \text{ and } s_i \end{aligned}$$

Hence, since $|S| = O(|E^s|)$, plugging the above back to Eq. (3.4), we have

$$\epsilon_d^s = \max\{\epsilon_{d,1}^s, \epsilon_{d,2}^s\} \leq O(|E^s|)\epsilon^f.$$

Finally, we bound the error in selection in G^s . Suppose $e \in S_i$, we have

$$\begin{aligned} f_i^s(e) &= f_i^f(e_1) \\ &\leq f_i^f(e_1) + f_i^f(e_3) && \text{by non-negativity constraint} \\ &\leq f_i^f(e_4) + \left| f_i^f(e_1) + f_i^f(e_3) - f_i^f(e_4) \right| \\ &\leq \epsilon^f + \left| f_i^f(e_1) + f_i^f(e_3) - f_i^f(e_4) \right|, && \text{by error in demand of vertex } t_i \\ &\leq 2\epsilon^f, && \text{by error in demand of vertex } xy \end{aligned}$$

Therefore, we can bound the error in type by

$$\epsilon_s^s \leq 2\epsilon^f.$$

Putting it all together, we have

$$\epsilon^s = \max\{\epsilon_l^s, \epsilon_u^s, \epsilon_d^s, \epsilon_s^s\} \leq O(|E^s|)\epsilon^f.$$

□

3.6.5 2CFF(A) to 2CFR(A): Dropping Fixed Flow Constraints

In this section, we show the reduction from a 2CFF instance $(G^f, F^f, \mathbf{u}^f, s_1, t_1, s_2, t_2)$ to a 2CFR instance $(G^r, \mathbf{u}^r, \bar{s}_1, \bar{t}_1, \bar{s}_2, \bar{t}_2, R_1, R_2)$. First of all, we add two new sources \bar{s}_1, \bar{s}_2 and two new sinks \bar{t}_1, \bar{t}_2 . Then, for each edge $e \in E^f$, we map it to a gadget consisting edges $\{e_1, e_2, e_3, e_4, e_5, e_6, e_7\}$ in G^r , as shown in the upper part of Figure 3.7. Additionally, there is another gadget with 5 edges in G^r that connects the original sink t_i and source s_i , as shown in the lower part of Figure 3.7. Capacity of these edges is the sum capacities of all edges in G^f , i.e., $M^f = \sum_{e \in E^f} \mathbf{u}^f(e)$. Additionally, we set $R_1 = R_2 = 2M^f$, indicating that at least $2M^f$ unit of flow should be routed from \bar{s}_i to \bar{t}_i , $i \in \{1, 2\}$.

The key idea to remove the fixed flow constraint is utilizing edge directions and the requirements that $2M^f$ units of the flow of commodity i to be routed from the new source \bar{s}_i to the new sink \bar{t}_i , $i \in \{1, 2\}$. It is noticed that all edges that are incident to the new sources and sinks should be saturated to fulfill the requirements. Therefore, for a fixed flow edge e in the first gadget, the incoming flow of vertex xy' and the outgoing flow of vertex xy must be $2u$. Since the capacity of e_3 is $2u - u = u$, then the flows of e_1, e_2, e_3 are forced to be u . As such, the fixed flow constraint can be simulated. Note that instead of simply copying non-fixed flow edges to G^r , we also need to map non-fixed flow edges to the designed gadget in this step. This guarantees that the requirement on flow values can be satisfied if the 2CFF instance is feasible.

Lemma 3.6.14 (2CFF to 2CFR). *Given a 2CFF instance $(G^f, F^f, \mathbf{u}^f, s_1, t_1, s_2, t_2)$, we can construct, in time $O(|E^f|)$, a 2CFR instance $(G^r, \mathbf{u}^r, \bar{s}_1, \bar{t}_1, \bar{s}_2, \bar{t}_2, R_1, R_2)$ such that*

$$|E^r| = O(|E^f|), \quad \|\mathbf{u}^r\|_{\max} = M^f,$$

where $M^f = \sum_{e \in E^f} \mathbf{u}^f(e)$. *If the 2CFF instance has a solution, then the 2CFR instance has a solution.*

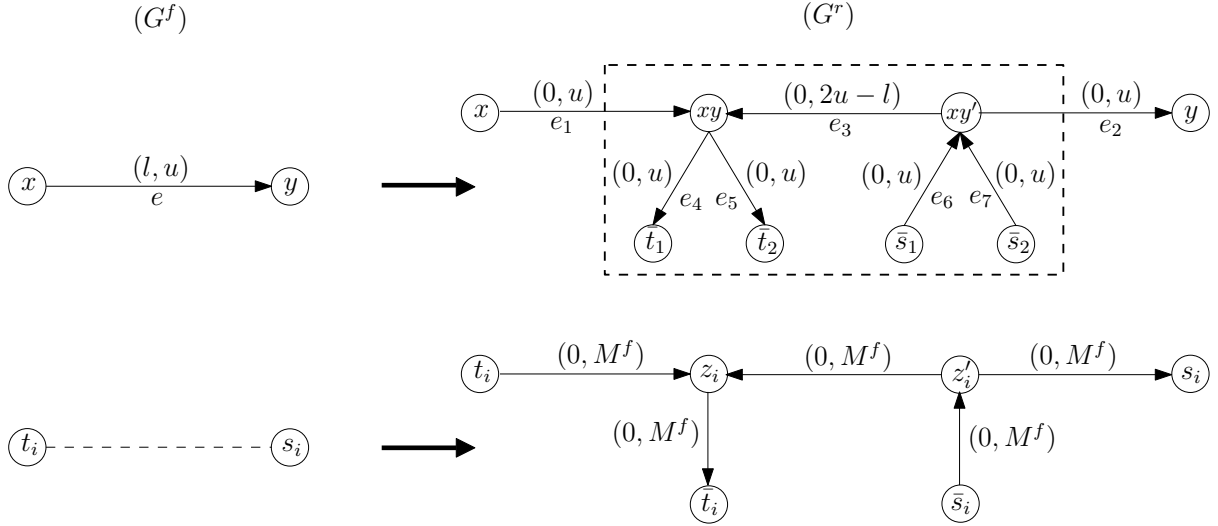


Figure 3.7: The reduction from 2CFF to 2CFR. $l = u$ if e is a fixed flow edge, and $l = 0$ if e is a non-fixed flow edge.

Proof. Suppose \mathbf{f}^f is a feasible solution to the 2CFF instance, then we can derive a feasible solution \mathbf{f}^r to the 2CFR instance. Concretely, we define a feasible flow \mathbf{f}^r as follows. For any edge $e \in E^f$, we set:

$$\begin{aligned} \mathbf{f}_i^r(e_1) &= \mathbf{f}_i^r(e_3) = \mathbf{f}_i^f(e), & i \in \{1, 2\}, \\ \mathbf{f}_i^r(e_2) &= u - \mathbf{f}_i^f(e), & i \in \{1, 2\}, \\ \mathbf{f}_1^r(e_4) &= \mathbf{f}_1^r(e_6) = u, & \mathbf{f}_2^r(e_5) = \mathbf{f}_2^r(e_7) = u, \\ \mathbf{f}_2^r(e_4) &= \mathbf{f}_2^r(e_6) = \mathbf{f}_1^r(e_5) = \mathbf{f}_1^r(e_7) = 0, \end{aligned}$$

where $u = \mathbf{u}^f(e)$. And we set

$$\begin{aligned} \mathbf{f}_i^r(t_i, z_i) &= \mathbf{f}_i^r(z'_i, s_i) = \sum_{w: (s_i, w) \in E^f} \mathbf{f}_i^f(s_i, w) \leq M^f, & i \in \{1, 2\}, \\ \mathbf{f}_i^r(z'_i, z_i) &= M^f - \sum_{w: (s_i, w) \in E^f} \mathbf{f}_i^f(s_i, w), & i \in \{1, 2\}, \\ \mathbf{f}_i^r(\bar{s}_i, z'_i) &= \mathbf{f}_i^r(z_i, \bar{t}_i) = M^f, & i \in \{1, 2\}, \\ \mathbf{f}_i^r(t_i, z_i) &= \mathbf{f}_i^r(z'_i, s_i) = \mathbf{f}_i^r(z'_i, z_i) = \mathbf{f}_i^r(\bar{s}_i, z'_i) = \mathbf{f}_i^r(z_i, \bar{t}_i) = 0, & \bar{i} \in \{1, 2\} \setminus i. \end{aligned}$$

To prove that \mathbf{f}^r is a solution to the 2CFR instance, it is obvious that capacity constraint and flow conservation constraint are satisfied by construction. For the satisfaction of the requirement constraint $R_1 = R_2 = 2M^f$, we have, for commodity 1,

$$\begin{aligned} R_1 &= \sum_{w: (\bar{s}_1, w) \in E^r} \mathbf{f}_1^r(\bar{s}_1, w) = \mathbf{f}^r(\bar{s}_1, z'_1) + \sum_{e \in E^f} \mathbf{f}_1^r(e_6) = M^f + \sum_{e \in E^f} \mathbf{u}^f(e) = 2M^f, \\ R_1 &= \sum_{u: (u, \bar{t}_1) \in E^r} \mathbf{f}_1^r(u, \bar{t}_1) = \mathbf{f}^r(z_1, \bar{t}_1) + \sum_{e \in E^f} \mathbf{f}_1^r(e_4) = M^f + \sum_{e \in E^f} \mathbf{u}^f(e) = 2M^f. \end{aligned}$$

We can prove that the requirement $R_2 = 2M^f$ is satisfied similarly. Therefore, \mathbf{f}^r is a feasible flow to the 2CFR instance. \square

Definition 3.6.15 (2CFR Approximate Problem (2CFRA)). A 2CFRA instance is given by a 2CFR instance $(G, \mathbf{u}, s_1, t_1, s_2, t_2, R_1, R_2)$ as in Definition 3.3.8, and an error parameter $\epsilon \in [0, 1]$, which we collect in a tuple $(G, \mathbf{u}, s_1, t_1, s_2, t_2, R_1, R_2, \epsilon)$. We say an algorithm solves the 2CFRA problem, if, given any 2CFRA instance, it returns a pair of flows $\mathbf{f}_1, \mathbf{f}_2 \geq \mathbf{0}$ that satisfies

$$\begin{aligned} \mathbf{f}_1(e) + \mathbf{f}_2(e) &\leq \mathbf{u}(e) + \epsilon, \quad \forall e \in E, & (\text{error in congestion}) \\ \left| \sum_{u:(u,v) \in E} \mathbf{f}_i(u,v) - \sum_{w:(v,w) \in E} \mathbf{f}_i(v,w) \right| &\leq \epsilon, \quad \forall v \in V \setminus \{s_i, t_i\}, i \in \{1, 2\}, & (\text{error in demand}) \\ \left| \sum_{w:(s_i,w) \in E} \mathbf{f}_i(s_i, w) - R_i \right| &\leq \epsilon, \quad \left| \sum_{u:(u,t_i) \in E} \mathbf{f}_i(u, t_i) - R_i \right| \leq \epsilon, \quad i \in \{1, 2\}, & (\text{error in demand}) \end{aligned}$$

or it correctly declares that the associated 2CFR instance is infeasible.

Lemma 3.6.16 (2CFFA to 2CFRA). *Given a 2CFFA instance $(G^f, F^f, \mathbf{u}^f, s_1, t_1, s_2, t_2, \epsilon^f)$, we can reduce it to a 2CFRA instance $(G^r, \mathbf{u}^r, \bar{s}_1, \bar{t}_1, \bar{s}_2, \bar{t}_2, R_1, R_2, \epsilon^r)$ by letting*

$$\epsilon^r = \Omega\left(\frac{1}{|E^f|}\right) \epsilon^f.$$

and using Lemma 3.6.14 to construct a 2CFR instance $(G^r, \mathbf{u}^r, \bar{s}_1, \bar{t}_1, \bar{s}_2, \bar{t}_2, R_1, R_2)$ from the SFF instance $(G^f, F^f, \mathbf{u}^f, s_1, t_1, s_2, t_2)$. If \mathbf{f}^r is a solution to the 2CFRA (2CFR) instance, then in time $O(|E^f|)$, we can compute a solution \mathbf{f}^f to the 2CFFA (2CFF, respectively) instance, where the exact case holds when $\epsilon^r = \epsilon^f = 0$.

Proof. To bound error in congestion, we have $\epsilon_u^f \leq \epsilon^r$ by Lemma 3.6.3. Regarding to ϵ_l^f , we have $\epsilon_l^f = 0$ if $e \in E^f \setminus F^f$; whereas for an arbitrary fixed flow edge $e \in F^f$, by solution mapping, we have

$$\begin{aligned} \mathbf{f}^f(e) &= \mathbf{f}^r(e_1) = \mathbf{f}_1^r(e_1) + \mathbf{f}_2^r(e_1) \\ &\geq \sum_{i \in \{1, 2\}} (\mathbf{f}_i^r(e_4) + \mathbf{f}_i^r(e_5) - \mathbf{f}_i^r(e_3) - \epsilon^r) \\ &\geq \mathbf{f}_1^r(e_4) + \mathbf{f}_2^r(e_5) - \mathbf{f}^r(e_3) - 2\epsilon^r, \end{aligned} \tag{3.5}$$

where the last step follows from the non-negativity of $\mathbf{f}_1^r(e_5)$ and $\mathbf{f}_2^r(e_4)$.

Now, to bound $\mathbf{f}_1^r(e_4)$ and $\mathbf{f}_2^r(e_5)$, we have to make use of the requirement constraint of \bar{t}_1, \bar{t}_2 . By definition,

$$\begin{aligned} \left| \sum_{e \in E^f} \mathbf{f}_1^r(e_4) + \mathbf{f}^r(z_1, \bar{t}_1) - 2M^f \right| &\leq \epsilon^r, \\ \left| \sum_{e \in E^f} \mathbf{f}_2^r(e_5) + \mathbf{f}^r(z_2, \bar{t}_2) - 2M^f \right| &\leq \epsilon^r. \end{aligned}$$

Rearranging and taking the lower bound direction, we have

$$\begin{aligned} \mathbf{f}_1^r(e_4) &\geq 2M^f - \epsilon^r - \sum_{\hat{e} \in E^f \setminus e} \mathbf{f}_1^r(e_4) - \mathbf{f}^r(z_1, \bar{t}_1) \\ &\geq \max \{0, 2M^f - \epsilon^r - (M^f - \mathbf{u}^f(e) + (|E^f| - 1)\epsilon^r) - (M^f + \epsilon^r)\} \\ &= \max \{0, \mathbf{u}^f(e) - (|E^f| + 1)\epsilon^r\}. \end{aligned}$$

Similarly, we have $\mathbf{f}_2^r(e_5) \geq \max \{0, \mathbf{u}^f(e) - (|E^f| + 1)\epsilon^r\}$.

Plugging back to Eq. (3.5),

$$\begin{aligned} \mathbf{f}^f(e) &\geq \max \{0, 2(\mathbf{u}^f(e) - (|E^f| + 1)\epsilon^r) - \mathbf{f}^r(e_3) - 2\epsilon^r\} \\ &\quad \text{by lower bound of } \mathbf{f}_1^r(e_4) \text{ and } \mathbf{f}_2^r(e_5) \\ &\geq \max \{0, 2(\mathbf{u}^f(\hat{e}) - (|E^f| + 1)\epsilon^r) - (\mathbf{u}^f(e) + \epsilon^r) - 2\epsilon^r\} \\ &\quad \text{by upper bound of } \mathbf{f}^r(e_3) \\ &= \max \{0, \mathbf{u}^f(e) - (2|E^f| + 5)\epsilon^r\}. \end{aligned}$$

Hence, we can bound

$$\epsilon_l^f \leq O(|E^f|) \epsilon^r.$$

Next, we analyze error in demand. By Lemma 3.6.4 and solution mapping algorithm, we have for $i \in \{1, 2\}$,

$$\epsilon_{d,i}^f \leq \epsilon^r + \max_{v \in V^f \setminus \{s_i, t_i\}} \sum_{e=(u,v) \in E^f} |\mathbf{f}_i^r(e_1) - \mathbf{f}_i^r(e_2)|. \quad (3.6)$$

The problem is reduced to bounding $|\mathbf{f}_i^r(e_1) - \mathbf{f}_i^r(e_2)|$.

$$\begin{aligned} &|\mathbf{f}_i^r(e_1) - \mathbf{f}_i^r(e_2)| \\ &\leq |\mathbf{f}_i^r(e_1) + \mathbf{f}_i^r(e_3) - \mathbf{f}_i^r(e_4) - \mathbf{f}_i^r(e_5)| + |\mathbf{f}_i^r(e_4) + \mathbf{f}_i^r(e_5) - \mathbf{f}_i^r(e_3) - \mathbf{f}_i^r(e_2)| \\ &\leq \epsilon^r + |\mathbf{f}_i^r(e_4) + \mathbf{f}_i^r(e_5) - \mathbf{f}_i^r(e_3) - \mathbf{f}_i^r(e_2)| \quad \text{by error in demand of vertex } xy \\ &\leq \epsilon^r + |\mathbf{f}_i^r(e_4) + \mathbf{f}_i^r(e_5) - \mathbf{f}_i^r(e_6) - \mathbf{f}_i^r(e_7)| + |\mathbf{f}_i^r(e_6) + \mathbf{f}_i^r(e_7) - \mathbf{f}_i^r(e_3) - \mathbf{f}_i^r(e_2)| \\ &\leq 2\epsilon^r + |\mathbf{f}_i^r(e_4) + \mathbf{f}_i^r(e_5) - \mathbf{f}_i^r(e_6) - \mathbf{f}_i^r(e_7)|, \quad \text{by error in demand of vertex } xy' \end{aligned}$$

By error in demand of \bar{t}_i and \bar{s}_i , we have

$$\mathbf{f}_1^r(e_5), \mathbf{f}_1^r(e_7) \leq \epsilon^r; \quad \mathbf{f}_2^r(e_4), \mathbf{f}_2^r(e_6) \leq \epsilon^r.$$

Hence,

$$\begin{aligned} |\mathbf{f}_1^r(e_1) - \mathbf{f}_1^r(e_2)| &\leq 4\epsilon^r + |\mathbf{f}_1^r(e_4) - \mathbf{f}_1^r(e_6)|, \\ |\mathbf{f}_2^r(e_1) - \mathbf{f}_2^r(e_2)| &\leq 4\epsilon^r + |\mathbf{f}_2^r(e_5) - \mathbf{f}_2^r(e_7)|. \end{aligned}$$

Claim 3.6.17.

$$\sum_{e \in E^f} |\mathbf{f}_1^r(e_4) - \mathbf{f}_1^r(e_6)| \leq 6|E^f|\epsilon^r, \quad \sum_{e \in E^f} |\mathbf{f}_2^r(e_5) - \mathbf{f}_2^r(e_7)| \leq 6|E^f|\epsilon^r.$$

Combining the above bounds and Claim 3.6.17 and plugging back to Eq. (3.6), we have

$$\epsilon_{d,1}^f, \epsilon_{d,2}^f \leq O(|E^f|) \epsilon^r.$$

Putting it altogether, we have

$$\epsilon^f = \max\{\epsilon_u^f, \epsilon_l^f, \epsilon_{d,1}^f, \epsilon_{d,2}^f\} \leq O(|E^f|) \epsilon^r.$$

□

Proof of Claim 3.6.17. We divide all edges in E^f into two groups:

$$E_I = \{e \in E^f \text{ s.t. } \mathbf{f}_1^r(e_4) \leq \mathbf{f}_1^r(e_6)\},$$

$$E_{II} = \{e \in E^f \text{ s.t. } \mathbf{f}_1^r(e_4) > \mathbf{f}_1^r(e_6)\}.$$

We denote

$$T_I = \sum_{e \in E_I} \mathbf{f}_1^r(e_4), \quad S_I = \sum_{e \in E_I} \mathbf{f}_1^r(e_6);$$

$$T_{II} = \sum_{e \in E_{II}} \mathbf{f}_1^r(e_4), \quad S_{II} = \sum_{e \in E_{II}} \mathbf{f}_1^r(e_6).$$

Then,

$$\sum_{e \in E^f} |\mathbf{f}_1^r(e_4) - \mathbf{f}_1^r(e_6)| = (S_I - T_I) + (T_{II} - S_{II}).$$

On one hand, by error in congestion, we have

$$S_I + T_{II} \leq \sum_{e \in E^f} \mathbf{u}^f(e) + |E^f| \epsilon^r = M^f + |E^f| \epsilon^r. \quad (3.7)$$

On the other hand, applying error in demand on vertex t_i and s_i , we have

$$T_I + T_{II} \geq (2M^f - \epsilon^r) - \mathbf{f}_1^r(z_1, \bar{t}_1) \geq 2M^f - \epsilon^r - M^f - \epsilon^r = M^f - 2\epsilon^r,$$

$$S_I + S_{II} \geq (2M^f - \epsilon^r) - \mathbf{f}_1^r(\bar{s}_1, z'_1) \geq 2M^f - \epsilon^r - M^f - \epsilon^r = M^f - 2\epsilon^r.$$

Thus,

$$S_I + T_{II} = S_I + (T_I + T_{II}) - T_I \geq M^f - 2\epsilon^r + (S_I - T_I),$$

$$S_I + T_{II} = (S_I + S_{II}) + T_{II} - S_{II} \geq M^f - 2\epsilon^r + (T_{II} - S_{II}).$$

Together with Eq. (3.7), we obtain

$$S_I - T_I \leq S_I + T_{II} - M^f + 2\epsilon^r \leq M^f + |E^f| \epsilon^r - M^f + 2\epsilon^r = (|E^f| + 2)\epsilon^r,$$

$$T_{II} - S_{II} \leq S_I + T_{II} - M^f + 2\epsilon^r \leq M^f + |E^f| \epsilon^r - M^f + 2\epsilon^r = (|E^f| + 2)\epsilon^r.$$

Hence, we have

$$\sum_{e \in E^f} |\mathbf{f}_1^r(e_4) - \mathbf{f}_1^r(e_6)| = (S_I - T_I) + (T_{II} - S_{II}) \leq 2(|E^f| + 2)\epsilon^r \leq 6|E^f| \epsilon^r,$$

which finishes the proof. \square

3.6.6 2CFR(A) to 2CF(A)

In this section, we show the reduction from a 2CFR instance $(G^r, \mathbf{u}^r, \bar{s}_1, \bar{t}_1, \bar{s}_2, \bar{t}_2, R_1, R_2)$ to a 2CF instance $(G^{2cf}, \mathbf{u}^{2cf}, \bar{s}_1, \bar{t}_1, \bar{s}_2, \bar{t}_2, R^{2cf})$. To drop the required flow constraint, we add to G^{2cf} with two new sources \bar{s}_1, \bar{s}_2 and two new edges $(\bar{s}_1, \bar{s}_1), (\bar{s}_2, \bar{s}_2)$ with capacity R_1, R_2 , respectively. We set $R^{2cf} = R_1 + R_2$.

For solution mapping, if a 2CF solver returns \mathbf{f}^{2cf} for the 2CF instance $(G^{2cf}, \mathbf{u}^{2cf}, \bar{s}_1, \bar{t}_1, \bar{s}_2, \bar{t}_2, R^{2cf})$, then we return \mathbf{f}^r for the 2CFR instance $(G^r, \mathbf{u}^r, \bar{s}_1, \bar{t}_1, \bar{s}_2, \bar{t}_2, R_1, R_2)$ by setting $\mathbf{f}_i^r(e) = \mathbf{f}_i^{2cf}(e), \forall e \in E^r, i \in \{1, 2\}$. If the 2CF solver returns “infeasible” for the 2CF instance, then we return “infeasible” for the 2CFR instance.

Note that the required flow constraint is not an edge constraint, so conclusions from Section 3.6.1 cannot be directly applied. Instead, we analyze this step directly.

Lemma 3.6.18 (2CFR to 2CF). *Given a 2CFR instance $(G^r, \mathbf{u}^r, \bar{s}_1, \bar{t}_1, \bar{s}_2, \bar{t}_2, R_1, R_2)$, we can construct, in time $O(|E^r|)$, a 2CF instance $(G^{2cf}, \mathbf{u}^{2cf}, \bar{s}_1, \bar{t}_1, \bar{s}_2, \bar{t}_2, R^{2cf})$ such that*

$$|E^{2cf}| = O(|E^r|), \quad \|\mathbf{u}^{2cf}\|_{\max} = \max\{R_1, R_2\}, \quad R^{2cf} = R_1 + R_2.$$

If the 2CFR instance has a solution, then the 2CF instance has a solution.

Proof. The runtime for reduction and solution mapping is trivial, as we only add two vertices and two edges. And by construction, $\|\mathbf{u}^{2cf}\|_{\max} = \max\{R_1, R_2\}$.

If there exists a solution \mathbf{f}^r to the 2CFR instance such that $F_1^r \geq R_1, F_2^r \geq R_2$, then there also exists a solution \mathbf{f}^{2cf} to the 2CF instance because $F_1^{2cf} = R_1, F_2^{2cf} = R_2$, and thus $F_1^{2cf} + F_2^{2cf} = R^{2cf}$. \square

Lemma 3.6.19 (2CFRA to 2CFA). *Given a 2CFRA instance $(G^r, \mathbf{u}^r, \bar{s}_1, \bar{t}_1, \bar{s}_2, \bar{t}_2, R_1, R_2, \epsilon^r)$, we can reduce it to a 2CFA instance $(G^{2cf}, \mathbf{u}^{2cf}, \bar{s}_1, \bar{t}_1, \bar{s}_2, \bar{t}_2, R^{2cf}, \epsilon^{2cf})$ by letting*

$$\epsilon^{2cf} = \Omega(\epsilon^r),$$

and using Lemma 3.6.18 to construct a 2CF instance $(G^{2cf}, \mathbf{u}^{2cf}, \bar{s}_1, \bar{t}_1, \bar{s}_2, \bar{t}_2, R^{2cf})$ from the 2CFR instance $(G^r, \mathbf{u}^r, \bar{s}_1, \bar{t}_1, \bar{s}_2, \bar{t}_2, R_1, R_2)$. If \mathbf{f}^{2cf} is a solution to the 2CFA (2CF) instance, then in time $O(|E^r|)$, we can compute a solution \mathbf{f}^r to the 2CFRA (2CFR, respectively) instance, where the exact case holds when $\epsilon^{2cf} = \epsilon^r = 0$.

Proof. Error in congestion does not increase since flows and edge capacities are unchanged for those edges other than $(\bar{s}_1, \bar{s}_1), (\bar{s}_2, \bar{s}_2)$. Hence, $\epsilon^r \leq O(\epsilon^{2cf})$.

Then for error in demand, notice that only the incoming flow of vertex \bar{s}_1, \bar{s}_2 changes by solution mapping. Therefore, error in demand does not increase for vertices other than \bar{s}_1, \bar{s}_2 , thus we only need to bound the error in demand of \bar{s}_1, \bar{s}_2 . We consider commodity $i, i \in \{1, 2\}$. We first bound the error in demand of vertex \bar{s}_i .

$$\begin{aligned} \left| \sum_{w: (\bar{s}_i, w) \in E} \mathbf{f}_i^{2cf}(\bar{s}_i, w) - R_i \right| &\leq \left| \sum_{w: (\bar{s}_i, w) \in E} \mathbf{f}_i^{2cf}(\bar{s}_i, w) - \mathbf{f}_i^{2cf}(\bar{s}_i, \bar{s}_i) \right| + \left| \mathbf{f}_i^{2cf}(\bar{s}_i, \bar{s}_i) - R_i \right| \\ &\leq \epsilon^{2cf} + \left| \mathbf{f}_i^{2cf}(\bar{s}_i, \bar{s}_i) - R_i \right|, \text{ by error in demand of vertex } \bar{s}_i \end{aligned}$$

For $\left| \mathbf{f}_i^{2cf}(\bar{s}_i, \bar{s}_i) - R_i \right|$, by error in congestion, we have

$$\mathbf{f}_i^{2cf}(\bar{s}_i, \bar{s}_i) - R_i \leq \epsilon^{2cf}.$$

In the other direction, we have

$$\begin{aligned} \mathbf{f}_i^{2cf}(\bar{s}_i, \bar{s}_i) - R_i &\geq F_i^{2cf} - \epsilon^{2cf} - R_i && \text{by error in demand of vertex } \bar{s}_i \\ &= R^{2cf} - F_i^{2cf} - R_i - \epsilon^{2cf} && \text{by construction} \\ &= R_i - F_i^{2cf} - \epsilon^{2cf} && \text{by } R_1 + R_2 = R^{2cf} \\ &\geq \mathbf{f}_i^{2cf}(\bar{s}_i, \bar{s}_i) - F_i^{2cf} - 2\epsilon^{2cf} && \text{by error in congestion of edge } (\bar{s}_i, \bar{s}_i) \\ &\geq -3\epsilon^{2cf}, && \text{by error in demand of vertex } \bar{s}_i \end{aligned}$$

Therefore, we have $\left| \mathbf{f}_i^{2cf}(\bar{s}_i, \bar{s}_i) - R_i \right| \leq 3\epsilon^{2cf}$, and thus

$$\left| \sum_{w: (\bar{s}_i, w) \in E} \mathbf{f}_i^{2cf}(\bar{s}_i, w) - R_i \right| \leq 4\epsilon^{2cf}.$$

To bound the error in demand of vertex \bar{s}_i , we have $\sum_{w: (\bar{s}_i, w) \in E} \mathbf{f}_i^r(\bar{s}_i, w) \leq 2\epsilon^{2cf}$ because error in demand is accumulated twice over two vertices \bar{s}_i, \bar{s}_i .

To conclude,

$$\epsilon^r \leq O(\epsilon^{2cf}).$$

□

3.7 A Unified Framework for LP Transformations

In Section 3.5 and 3.6, we described a series of reduction steps in both the linear algebra space and the flow space. While we have identified certain reduction patterns in Section 3.5.1 and 3.6.1, each step uses its own gadget and error propagation arguments. We aim to provide a unified characterization of problem reduction and how the reductions can be analyzed systematically.

Note that every intermediate problem we consider (both in the algebra space and in the flow space) can be represented as a feasibility LP of the form:

$$\{\mathbf{x} : \mathbf{A}\mathbf{x} \leq \mathbf{b}\},$$

where \mathbf{A} and \mathbf{b} define linear inequalities enforcing the constraints of that problem. Note that a linear equality constraint $\mathbf{A}\mathbf{x} = \mathbf{b}$ is equivalent to the pair of inequalities $\mathbf{A}\mathbf{x} \leq \mathbf{b}$ and $-\mathbf{A}\mathbf{x} \leq -\mathbf{b}$. Similarly, non-negativity constraints $\mathbf{x} \geq \mathbf{0}$ can be expressed as $-\mathbf{I}\mathbf{x} \leq \mathbf{0}$.

This section introduces a framework that enables us to view each reduction step as an LP-to-LP transformation characterized by linear mappings on both variables and constraints. This viewpoint simplifies the correctness and error-propagation analyses, showing that each step follows the same underlying principles.

3.7.1 Characterizing LP Transformations

Suppose we have an original problem A in LP form:

$$\mathcal{F}^A = \{\mathbf{x}^A \in \mathbb{R}^{n^A} : \mathbf{A}^A \mathbf{x}^A \leq \mathbf{b}^A\}, \text{ where } \mathbf{A}^A \in \mathbb{R}^{m^A \times n^A}, \mathbf{b}^A \in \mathbb{R}^{m^A}.$$

We want to reduce it to a new problem B , also representable as an LP:

$$\mathcal{F}^B = \{\mathbf{x}^B \in \mathbb{R}^{n^B} : \mathbf{A}^B \mathbf{x}^B \leq \mathbf{b}^B\}, \text{ where } \mathbf{A}^B \in \mathbb{R}^{m^B \times n^B}, \mathbf{b}^B \in \mathbb{R}^{m^B}.$$

By construction, typically $n^A \leq n^B$ and $m^A \leq m^B$, as we add auxiliary variables and constraints in the new problem.

LP transformation can be characterized with two linear operators: *variable operator* \mathbf{T} and *constraint operator* \mathbf{P} .

Variable Operator \mathbf{T} . Each reduction step embeds the original variable vector \mathbf{x}^A into the larger variable vector \mathbf{x}^B . Concretely, we define $\mathbf{T} \in \mathbb{R}^{n^A \times n^B}$ such that

$$\mathbf{x}^A = \mathbf{T}\mathbf{x}^B. \quad (3.8)$$

Typically, by design, each row of \mathbf{T} typically has a single 1 and the rest zeros, so that each $\mathbf{x}^A(j)$ is copied from the appropriate position in \mathbf{x}^B . Equivalently, we can view \mathbf{x}^B as

$$\mathbf{x}^B = \begin{bmatrix} \mathbf{x}^A \\ \mathbf{x}^{\text{aux}} \end{bmatrix}.$$

Hence, applying \mathbf{T} just discards the auxiliary variables when extracting \mathbf{x}^A from \mathbf{x}^B .

Constraint Operator \mathbf{P} . We regard the original system $\mathbf{A}^A \mathbf{x}^A \leq \mathbf{b}^A$ as being derived from \mathbf{A}^B (and \mathbf{b}^B) by applying a linear operator $\mathbf{P} \in \mathbb{R}_{\geq 0}^{m^A \times m^B}$. Specifically, we require

$$\mathbf{A}^A \mathbf{T} = \mathbf{P} \mathbf{A}^B, \quad \mathbf{b}^A = \mathbf{P} \mathbf{b}^B. \quad (3.9)$$

That is, each constraint of problem A is obtained by taking a *non-negative* linear combination of constraints in problem B . The non-negativity of \mathbf{P} is essential when working with inequalities, as it guarantees that feasibility is preserved: if $\mathbf{A}^B \mathbf{x}^B \leq \mathbf{b}^B$, then $\mathbf{P}(\mathbf{A}^B \mathbf{x}^B) \leq \mathbf{P} \mathbf{b}^B$, ensuring that $\mathbf{A}^A \mathbf{x}^A \leq \mathbf{b}^A$.

We remark that for transformation among equality constraints, the sign of \mathbf{P} is irrelevant. In this case, if $\mathbf{A}^B \mathbf{x}^B = \mathbf{b}^B$, then any linear operator \mathbf{P} can preserve feasibility $\mathbf{P}(\mathbf{A}^B \mathbf{x}^B) = \mathbf{P} \mathbf{b}^B$.

3.7.2 Unified Correctness Analysis

Recall that the correctness analysis of a reduction algorithm involves verifying that problem A has a feasible solution if and only if the reduced problem B also has a feasible solution. In other words, there are two main points to check:

1. Forward feasibility: From any feasible \mathbf{x}^A for problem A , we can construct a feasible \mathbf{x}^B for problem B . Usually, this forward-feasibility construction is direct or “gadget-based”⁸.
2. Backward feasibility: Given a feasible \mathbf{x}^B for problem B , we must show $\mathbf{x}^A = \mathbf{T}\mathbf{x}^B$ is feasible for problem A .

Lemma 3.7.1 (Correctness of Reduction Algorithm). *Let \mathbf{T} and \mathbf{P} satisfy the relations in Equation (3.8) and (3.9). Additionally, assume there is a forward-feasibility mapping $\mathbf{L} : \mathbb{R}^{n^A} \rightarrow \mathbb{R}^{n^B}$ such that $\mathbf{L}(\mathcal{F}^A) \subseteq \mathcal{F}^B$. Then problem A has a feasible solution if and only if the reduced problem B also has a feasible solution.*

Proof. If \mathbf{x}^A is a feasible solution to A , by $\mathbf{L}(\mathcal{F}^A) \subseteq \mathcal{F}^B$, then a feasible solution to B can be constructed by applying \mathbf{L} to \mathbf{x}^A as $\mathbf{x}^B = \mathbf{L}(\mathbf{x}^A)$.

⁸In each step of Sections 3.5 and 3.6, the solution construction approach is detailed in the proofs of the lemmas for the exact cases.

It remains to show that if \mathbf{x}^B is a feasible solution to B , then $\mathbf{T}\mathbf{x}^B$ is a feasible solution to A , i.e., $\mathbf{T}(\mathcal{F}^B) \subseteq \mathcal{F}^A$. For any feasible solution $\mathbf{x}^B \in \mathcal{F}^B$, we have

$$\begin{aligned}
\mathbf{A}^B \mathbf{x}^B &\leq \mathbf{b}^B \\
&\Downarrow \\
\mathbf{P} \mathbf{A}^B \mathbf{x}^B &\leq \mathbf{P} \mathbf{b}^B && \text{by non-negativity of } \mathbf{P} \\
&\Downarrow \\
\mathbf{A}^A \mathbf{T} \mathbf{x}^B &\leq \mathbf{b}^A && \text{by Eq. (3.9)} \\
&\Downarrow \\
\mathbf{A}^A \mathbf{x}^A &\leq \mathbf{b}^A && \text{by Eq. (3.8)}
\end{aligned}$$

□

3.7.3 Unified Error Analysis

We extend the above analysis for feasible solutions to solutions with additive approximations. We say \mathbf{x} is an ϵ -approximate solution if⁹

$$\mathbf{A}\mathbf{x} - \mathbf{b} \leq \epsilon \mathbf{1}.$$

Given \mathbf{x}^B an ϵ^B -approximate solution to problem B , the goal of the error analysis is to understand how the approximation error of $\mathbf{x}^A = \mathbf{T}\mathbf{x}^B$ depends on ϵ^B .

Lemma 3.7.2 (Error Analysis). *Let \mathbf{x}^B be an ϵ^B -approximate solution to problem B , then $\mathbf{x}^A = \mathbf{T}\mathbf{x}^B$ is an ϵ^A -approximate solution to problem A with*

$$\epsilon^A \leq \|\mathbf{P}\|_\infty \epsilon^B.$$

Proof. We have

$$\begin{aligned}
\mathbf{A}^B \mathbf{x}^B - \mathbf{b}^B &\leq \epsilon^B \mathbf{1}_{n^B} \\
&\Downarrow \\
\mathbf{P}(\mathbf{A}^B \mathbf{x}^B - \mathbf{b}^B) &\leq \epsilon^B \mathbf{P} \mathbf{1}_{n^B} && \text{by non-negativity of } \mathbf{P} \\
&\Downarrow \\
\mathbf{A}^A \mathbf{T} \mathbf{x}^B - \mathbf{b}^A &\leq \epsilon^B \|\mathbf{P}\|_\infty \mathbf{1}_{n^A} && \text{by Eq. (3.9)} \\
&\Downarrow \\
\mathbf{A}^A \mathbf{x}^A - \mathbf{b}^A &\leq \epsilon^B \|\mathbf{P}\|_\infty \mathbf{1}_{n^A} && \text{by Eq. (3.8)}
\end{aligned}$$

□

Therefore, bounding $\|\mathbf{P}\|_\infty$ by a modest polynomial factor at each step ensures the overall additive error does not grow too large throughout the entire chain of reductions. Lemma 3.7.2 unifies the more specialized error analyses found in each step of Sections 3.5 and 3.6.

⁹If we assume the non-negativity constraints are always satisfied, which is the case in our proofs in Section 3.5 and 3.6, the right-hand side vector can be further tightened.

3.7.4 Examples

As an illustration, we revisit several steps in the reduction chain and analyze them using this unified framework.

LP(A) to LEN(A)

We revisit the step that transform an LP instance $\{\mathbf{c}^\top \mathbf{x} \geq K, \mathbf{A}\mathbf{x} \leq \mathbf{b}, \mathbf{x} \geq \mathbf{0}\}$ to an LEN instance $\{\tilde{\mathbf{A}}\tilde{\mathbf{x}} = \tilde{\mathbf{b}}, \tilde{\mathbf{x}} \geq \mathbf{0}\}$. To be compatible with the form used in the framework, we rewrite the constraints in LP instance (problem A) as

$$\begin{bmatrix} -\mathbf{c}^\top \\ \mathbf{A} \\ -\mathbf{I} \end{bmatrix} \mathbf{x} \leq \begin{bmatrix} -K \\ \mathbf{b} \\ \mathbf{0} \end{bmatrix} \Rightarrow \mathbf{A}^A = \begin{bmatrix} -\mathbf{c}^\top \\ \mathbf{A} \\ -\mathbf{I} \end{bmatrix}, \quad \mathbf{b}^A = \begin{bmatrix} -K \\ \mathbf{b} \\ \mathbf{0} \end{bmatrix}.$$

The LEN instance (problem B) is defined by appending slack variables $\mathbf{s} \geq \mathbf{0}$ and $\alpha \geq 0$. We have

$$\tilde{\mathbf{A}} = \begin{bmatrix} -\mathbf{c}^\top & \mathbf{0} & 1 \\ \mathbf{A} & \mathbf{I} & \mathbf{0} \end{bmatrix}, \quad \tilde{\mathbf{b}} = \begin{bmatrix} -K \\ \mathbf{b} \end{bmatrix}.$$

Constructing $\mathbf{A}^B, \mathbf{b}^B$ from $\tilde{\mathbf{A}}, \tilde{\mathbf{b}}$ is straightforward:

$$\mathbf{A}^B = \begin{bmatrix} \tilde{\mathbf{A}} \\ -\tilde{\mathbf{A}} \\ -\mathbf{I} \end{bmatrix}, \quad \mathbf{b}^B = \begin{bmatrix} \tilde{\mathbf{b}} \\ -\tilde{\mathbf{b}} \\ \mathbf{0} \end{bmatrix}.$$

The operator \mathbf{P} performs row additions on \mathbf{A}^B to match rows in $\mathbf{A}^A \mathbf{T}$. Concretely, each row of $\mathbf{A}^A \mathbf{T}$ is constructed by adding the corresponding row in \mathbf{A}^B to a specific row of $-\mathbf{I}$, ensuring that any extra non-zero entry in the auxiliary variables is canceled out.

The forward-feasibility mapping \mathbf{L} is described in the proof of Lemma 3.5.1. Thus, by Lemma 3.7.1, the correctness of the reduction can be proven. Moreover, since $\|\mathbf{P}\|_\infty$ is a constant here, the approximate error ϵ^A grows only by a constant factor from ϵ^B , consistent with Lemma 3.5.2.

LEN(A) to 2-LEN(A)

We use a concrete example to illustrate the construction of the constraint operator \mathbf{P} for this step. Starting with a linear constraint $5\mathbf{x}_1 + 3\mathbf{x}_2 - 7\mathbf{x}_3 = -1$. In matrix form, this can be rewritten as

$$\begin{bmatrix} 5 & 3 & -7 \end{bmatrix} \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \end{bmatrix} = [-1].$$

The reduction algorithm of bitwise decomposition (see Eq. (3.2)) leads to the following reduced problem:

$$\left[\begin{array}{ccc|ccc|} 1 & 1 & -1 & -2 & 2 & 0 & 0 & \\ 0 & 1 & -1 & 1 & -1 & -2 & 2 & \\ 1 & 0 & -1 & 0 & 0 & 1 & -1 & \\ & & & & \mathbf{I}_4 & & & \end{array} \right] \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \\ \mathbf{c}_0 \\ \mathbf{d}_0 \\ \mathbf{c}_1 \\ \mathbf{d}_1 \\ \mathbf{s} \end{bmatrix} = \begin{bmatrix} -1 \\ 0 \\ 0 \\ \bar{X}\mathbf{1}_4 \end{bmatrix}.$$

The first three rows represent the bitwise decomposition constraints, corresponding to bit weights $2^0, 2^1$ and 2^2 , respectively; and the last block row with identity matrices \mathbf{I} corresponds to slack variables \mathbf{s} added to carry terms \mathbf{c}, \mathbf{d} .

Notice that

$$[2^0 \quad 2^1 \quad 2^2 \quad \mathbf{0}_4^\top] \left[\begin{array}{ccc|ccc|} 1 & 1 & -1 & -2 & 2 & 0 & 0 \\ 0 & 1 & -1 & 1 & -1 & -2 & 2 \\ 1 & 0 & -1 & 0 & 0 & 1 & -1 \\ & & & & \mathbf{I}_4 & & \end{array} \right] = [5 \quad 3 \quad -7 \quad \mathbf{0}_4^\top \quad \mathbf{0}_4^\top],$$

and

$$[2^0 \quad 2^1 \quad 2^2 \quad \mathbf{0}_4^\top] \begin{bmatrix} -1 \\ 0 \\ 0 \\ \bar{X}\mathbf{1}_4 \end{bmatrix} = [-1].$$

They correspond to $\mathbf{P}\mathbf{A}^B = \mathbf{A}^A\mathbf{T}$ and $\mathbf{P}\mathbf{b}^B = \mathbf{b}^A$, respectively. Hence, we can construct the constraint operator \mathbf{P} (up to column rearrangements) as

$$\mathbf{P} = \begin{bmatrix} 2^0 & 2^1 & 2^2 & \mathbf{0}_4^\top & \\ & & & & \mathbf{I}_3 \end{bmatrix},$$

where the first row corresponds to the reverse of the bitwise decomposition process, summing up contributions weighted by 2^i ; the last block \mathbf{I}_3 corresponds to non-negativity constraints for \mathbf{x} .

We have

$$\|\mathbf{P}\|_\infty = O(\bar{X}),$$

where \bar{X} is the magnitude of LEN. Again, it matches the bound stated in Lemma 3.5.4.

2-LEN(A) to 1-LEN(A)

The reduction from a 2-LEN instance $(\bar{\mathbf{A}}, \bar{\mathbf{b}}, \bar{R})$ to a 1-LEN instance $(\hat{\mathbf{A}}, \hat{\mathbf{b}}, \hat{R})$ is relatively straightforward. We illustrate with a concrete example from the previous step: $\mathbf{x}_1 + \mathbf{x}_2 - \mathbf{x}_3 - 2(\mathbf{c}_0 - \mathbf{d}_0) = -1$. In matrix form,

$$\begin{bmatrix} 1 & 1 & -1 & -2 & -2 \end{bmatrix} \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \\ \mathbf{c}_0 \\ \mathbf{d}_0 \end{bmatrix} = [-1].$$

Introducing auxiliary variables \mathbf{c}'_0 and \mathbf{d}'_0 , we reduce it to

$$\begin{aligned} \mathbf{x}_1 + \mathbf{x}_2 - \mathbf{x}_3 - \mathbf{c}_0 - \mathbf{c}'_0 + \mathbf{d}_0 + \mathbf{d}'_0 &= -1 \\ \mathbf{c}_0 - \mathbf{c}'_0 &= 0 \\ \mathbf{d}_0 - \mathbf{d}'_0 &= 0, \end{aligned}$$

or in matrix form,

$$\begin{bmatrix} 1 & 1 & -1 & -1 & -1 & 1 & 1 \\ 0 & 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & -1 \end{bmatrix} \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \\ \mathbf{c}_0 \\ \mathbf{c}'_0 \\ \mathbf{d}_0 \\ \mathbf{d}'_0 \end{bmatrix} = \begin{bmatrix} -1 \\ 0 \\ 0 \end{bmatrix}.$$

Notice that a linear combination of the reduced problem recovers the original problem. Specifically,

$$\begin{bmatrix} 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 1 & -1 & -1 & -1 & 1 & 1 \\ 0 & 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & -1 \end{bmatrix} = \begin{bmatrix} 1 & 1 & -1 & -2 & -2 & 0 & 0 \end{bmatrix}.$$

The constraint operator \mathbf{P} can thus be constructed accordingly. In the general case, each coefficient of ± 2 in the 2-LEN instance leads to the introduction of two auxiliary variables and two additional equations in the 1-LEN instance. Consequently, each row of \mathbf{P} includes a number of non-zero entries proportional to the number of variables with coefficients ± 2 . Since the number of such variables is trivially bounded by the number of non-zeros \bar{N} in 2-LEN, we obtain a matching bound as Lemma 3.5.6:

$$\|\mathbf{P}\|_\infty \leq O(\bar{N}).$$

1-LEN(A) to FHF(A)

Recall the form of 1-LEN $\{\hat{\mathbf{A}}\hat{\mathbf{x}} = \hat{\mathbf{b}}, \hat{\mathbf{x}} \geq \mathbf{0}\}$, where $\hat{\mathbf{A}} \in \mathbb{R}^{\hat{m} \times \hat{n}}$ and entries of $\hat{\mathbf{A}}$ are in $\{0, \pm 1\}$. For each linear equation in 1-LEN, saying

$$\sum_{j \in J^+} \hat{\mathbf{x}}(j) - \sum_{j \in J^-} \hat{\mathbf{x}}(j) = b,$$

or in matrix form,

$$\begin{bmatrix} \mathbf{1}_{|J^+|}^\top & -\mathbf{1}_{|J^-|}^\top \end{bmatrix} \begin{bmatrix} \hat{\mathbf{x}}(J^+) \\ \hat{\mathbf{x}}(J^-) \end{bmatrix} = [b], \quad \mathbf{x} \geq \mathbf{0}.$$

It is encoded by several flow constraints as follows:

$$\begin{aligned} \sum_{j \in J^+} \hat{\mathbf{x}}(j) &= \mathbf{f}_1 + \mathbf{f}_2, & (\text{flow conservation}) \\ \sum_{j \in J^-} \hat{\mathbf{x}}(j) &= \mathbf{f}_3, & (\text{flow conservation}) \\ \mathbf{f}_1 &= b, & (\text{fixed flow}) \\ \mathbf{f}_2 - \mathbf{f}_3 &= 0, & (\text{homologous flow}) \\ \begin{bmatrix} \hat{\mathbf{x}} \\ \mathbf{f} \end{bmatrix} &\leq \mathbf{u}. & (\text{capacity}) \end{aligned}$$

In matrix form, the flow constraints are

$$\begin{bmatrix} \mathbf{1}_{|J^+|}^\top & & -1 & -1 & 0 \\ & -\mathbf{1}_{|J^-|}^\top & 0 & 0 & 1 \\ & & 1 & 0 & 0 \\ & & 0 & 1 & -1 \end{bmatrix} \begin{bmatrix} \hat{\mathbf{x}}^{(J^+)} \\ \hat{\mathbf{x}}^{(J^-)} \\ \mathbf{f}_1 \\ \mathbf{f}_2 \\ \mathbf{f}_3 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ b \\ 0 \end{bmatrix}, \quad \mathbf{0} \leq \begin{bmatrix} \hat{\mathbf{x}} \\ \mathbf{f} \end{bmatrix} \leq \mathbf{u}.$$

Observe that by taking a linear combination of the flow constraints, we can recover the original 1-LEN equation. Specifically,

$$\begin{bmatrix} 1 & 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} \mathbf{1}_{|J^+|}^\top & & -1 & -1 & 0 \\ & -\mathbf{1}_{|J^-|}^\top & 0 & 0 & 1 \\ & & 1 & 0 & 0 \\ & & 0 & 1 & -1 \end{bmatrix} = \begin{bmatrix} \mathbf{1}_{|J^+|}^\top & -\mathbf{1}_{|J^-|}^\top & 0 & 0 & 0 \end{bmatrix}.$$

For a single equation, we can derive the above constraint operator \mathbf{P} and $\|\mathbf{P}\|_\infty = \Theta(1)$. When handling multiple equations, more homologous flow constraints are introduced—each linking multiple flows corresponding to the same variable across different sections of the network. This increases the complexity of the constraint operator \mathbf{P} . However, the number of non-zero entries in each row of \mathbf{P} remains bounded by $O(\hat{N})$, where \hat{N} is the number of non-zeros in the matrix $\hat{\mathbf{A}}$. Thus, in the general case,

$$\|\mathbf{P}\|_\infty \leq O(\hat{N}),$$

which matches the bound established in Lemma 3.5.9.

Dropping Homologous Flow Constraints

We show that every homologous flow constraint can be encoded as a linear combination of a constant number of flow conservation constraints, selective flow constraints, and fixed flow constraints.

Refer to Figure 3.5. Our target is to encode $\mathbf{f}^p(e) = \mathbf{f}^p(\hat{e})$, or before solution mapping, $\mathbf{f}_1^s(e_1) = \mathbf{f}_1^s(\hat{e}_1)$.

We start by flow conservation of vertex vw, vw', yz, yz' :

$$\begin{aligned} \mathbf{f}_1^s(e_1) + \mathbf{f}_1^s(e_3) &= \mathbf{f}_1^s(e_4) = \mathbf{f}_1^s(e_2) + \mathbf{f}_1^s(e_5), \\ \mathbf{f}_2^s(e_1) + \mathbf{f}_2^s(e_3) &= \mathbf{f}_2^s(e_4) = \mathbf{f}_2^s(e_2) + \mathbf{f}_2^s(e_5), \\ \mathbf{f}_1^s(\hat{e}_1) + \mathbf{f}_1^s(e_5) &= \mathbf{f}_1^s(\hat{e}_4) = \mathbf{f}_1^s(\hat{e}_2) + \mathbf{f}_1^s(e_5), \\ \mathbf{f}_2^s(\hat{e}_1) + \mathbf{f}_2^s(e_5) &= \mathbf{f}_2^s(\hat{e}_4) = \mathbf{f}_2^s(\hat{e}_2) + \mathbf{f}_2^s(e_5). \end{aligned} \tag{flow conservation}$$

Then plugging in the selective constraints on edge $e_1, e_2, \hat{e}_1, \hat{e}_2$ (i.e., $\mathbf{f}_2^s(e_1) = \mathbf{f}_2^s(e_2) = \mathbf{f}_2^s(\hat{e}_1) = \mathbf{f}_2^s(\hat{e}_2) = 0$) and on edge e_3, e_5, \hat{e}_5 (i.e., $\mathbf{f}_1^s(e_3) = \mathbf{f}_1^s(e_5) = \mathbf{f}_1^s(\hat{e}_5) = 0$), we have

$$\begin{aligned} \mathbf{f}_1^s(e_1) &= \mathbf{f}_1^s(e_4) = \mathbf{f}_1^s(e_2), \\ \mathbf{f}_1^s(\hat{e}_1) &= \mathbf{f}_1^s(\hat{e}_4) = \mathbf{f}_1^s(\hat{e}_2), \\ \mathbf{f}_2^s(e_3) &= \mathbf{f}_2^s(e_4) = \mathbf{f}_2^s(e_5) = \mathbf{f}_2^s(\hat{e}_4) = \mathbf{f}_2^s(\hat{e}_5). \end{aligned} \tag{flow conservation + selective}$$

Finally, the fixed flow constraints on edge e_4 and \hat{e}_4 enforce $\mathbf{f}_1^s(e_4) + \mathbf{f}_2^s(e_4) = u$ and $\mathbf{f}_1^s(\hat{e}_4) + \mathbf{f}_2^s(\hat{e}_4) = u$. Combining with the above immediately yields

$$\mathbf{f}_1^s(e_1) = \mathbf{f}_1^s(\hat{e}_1), \tag{flow conservation + selective + fixed}$$

which encodes the homologous flow constraint. Crucially, each original homologous-flow constraint is encoded by $O(1)$ constraints in the reduced problem, so the corresponding row sum of \mathbf{P} is bounded by a constant.

Dropping Selective Flow Constraints

Consider a selective edge e for commodity 1 (Figure 3.6). The target is to encode $\mathbf{f}_2^s(e) = 0$, or before solution mapping, $\mathbf{f}_2^f(e_1) = 0$.

By flow conservation of vertex xy and t_1 , because t_1 has no outgoing edges for commodity 2, we have

$$\mathbf{f}_2^f(e_1) + \mathbf{f}_2^f(e_3) = \mathbf{f}_2^f(e_4) = 0. \quad (\text{flow conservation})$$

Since every flow is non-negative, it implies

$$\mathbf{f}_2^f(e_1) = 0, \quad (\text{flow conservation} + \text{non-negativity})$$

which encodes the selective flow constraint.

Dropping Fixed Flow Constraints

Figure 3.7 shows the gadget that replaces an edge with the fixed capacity constraint

$$\mathbf{f}_1^f(e) + \mathbf{f}_2^f(e) = u.$$

Before the solution mapping, this means enforcing

$$\mathbf{f}_1^r(e_1) + \mathbf{f}_2^r(e_1) = u.$$

By flow conservation of vertex xy and no outgoing edges for commodity 2 and 1 for \bar{t}_1 and \bar{t}_2 , respectively, we have

$$\begin{aligned} \mathbf{f}_1^r(e_1) + \mathbf{f}_1^r(e_3) &= \mathbf{f}_1^r(e_4), \\ \mathbf{f}_2^r(e_1) + \mathbf{f}_2^r(e_3) &= \mathbf{f}_2^r(e_5). \end{aligned} \quad (\text{flow conservation})$$

Moreover, because of the required amount of flow, all incoming edges of \bar{t}_1, \bar{t}_2 saturate. That is, \bar{t}_i receive exactly u units of commodity i :

$$\begin{aligned} \mathbf{f}_1^r(e_1) + \mathbf{f}_1^r(e_3) &= \mathbf{f}_1^r(e_4) = u, \\ \mathbf{f}_2^r(e_1) + \mathbf{f}_2^r(e_3) &= \mathbf{f}_2^r(e_5) = u. \end{aligned} \quad (\text{flow conservation} + \text{requirement} + \text{capacity})$$

Then combined with capacity constraint of edge e_1 and e_3 (i.e., $\mathbf{f}_1^r(e_1) + \mathbf{f}_2^r(e_1) \leq u$, $\mathbf{f}_1^r(e_3) + \mathbf{f}_2^r(e_3) \leq u$), this gives

$$\mathbf{f}_1^r(e_1) + \mathbf{f}_2^r(e_1) = u, \quad (\text{flow conservation} + \text{requirement} + \text{capacity})$$

which encodes the fixed flow constraint.

The examples above demonstrate that every special constraint can be encoded as a linear combination of a constant number of other constraints. We omit the detailed encoding of other ordinary constraints associated with the reduced edges, such as flow

conservation and capacity constraints. Due to the constant size of the gadgets, the transformed graph G^B contains $O(|E^A|)$ edges, thus $O(|E^A|)$ constraints. Therefore, every original constraint can be recovered as the sum of at most $O(|E^A|)$ constraints in G^B , i.e.,

$$\|\mathbf{P}\|_\infty \leq O(|E^A|),$$

yielding the matching error bounds achieved in these reduction steps:

$$\epsilon^A \leq O(|E^A|)\epsilon^B.$$

3.7.5 Summary

Under this framework, we can treat each reduction step as a composition of two linear mappings, one on the variables (\mathbf{T}) and one on the constraints (\mathbf{P}), plus the forward-feasibility mapping \mathbf{L} . Once we construct $(\mathbf{T}, \mathbf{P}, \mathbf{L})$ for a reduction step, its correctness follows immediately from Lemma 3.7.1.

In addition, this framework provides a systematic approach to analyzing error propagation. Rather than redoing an error analysis from scratch for each reduction gadget, it suffices to bound the ℓ_∞ norm of the constraint operator $\|\mathbf{P}\|_\infty$ at each step. This immediately yields how ϵ^B translates into ϵ^A .

Intuitively, the role of $\|\mathbf{P}\|_\infty$ in bounding error growth corresponds to how we break down an error term into several intermediate terms in our actual proofs:

$$|\mathbf{A}^A(i)\mathbf{x}^A - \mathbf{b}^A(i)| \leq \sum_{j \in J} |\mathbf{A}^B(j)\mathbf{x}^B - \mathbf{b}^B(j)|,$$

where J is the set of indices of constraints in B that are related to the i th constraint in A . Combining these terms is equivalent to forming a linear combination of constraints in problem B —an operation precisely captured by the constraint operator \mathbf{P} .

Although constructing \mathbf{P} still requires step-specific analysis for individual reduction gadgets, the framework serves as a powerful tool both for validating existing reductions and for guiding the design and verification of new reduction steps within the chain of LP reductions.

Chapter 4

Hardness Results for Combinatorial Laplacians

This chapter is based on [Din+22]. We focus on proving Theorem 1.4.4. Specifically, we describe the reduction algorithm and analyze it for approximate solvers in Section 4.5. Then in Section 4.6, we slightly modify the reduction for approximate solvers to handle the general case, when the right-hand side vector may not be in the image of the coefficient matrix. Theorem 1.4.2 can be derived from Theorem 1.4.4 and standard arguments in linear algebra, with the proof provided in Appendix B.1.

4.1 Prior Works

Kyng and Zhang [KZ17] initiated the study of hardness results for solving structured linear equations. They showed that 2-commodity Laplacians, 2-dimensional truss stiffness matrices, and 2-total-variation matrices are all sparse-linear-equation complete. [KWZ20] considered a larger family of hardness assumptions based on linear equations. Through defining a parameterized family of hypotheses for runtime of solving sparse linear equations, they prove the hardness of approximately solving packing and covering linear programs. For example, if one can solve a packing linear program up to ϵ accuracy in time $\tilde{O}(\text{number of non-zeros} \times \epsilon^{-0.165})$, then one can solve a system of linear equations in time asymptotically faster than $\tilde{O}(\text{number of non-zeros} \times \text{condition number of matrix})$, which is the runtime of conjugate gradient.

4.2 Our Contributions

We prove that solving linear equations in boundary operators and combinatorial Laplacians of 2-complexes is as hard to solve as general linear equations. As discussed in Section 1.1, there are extensive applications of combinatorial Laplacians in both pure math and applied areas. In addition, the problem of solving linear equations in $\partial_2 \partial_2^\top$ also arises when using Interior Point Methods to solve a generalized max-flow problem in higher-dimensional simplicial complexes as defined in [MN21]. We sketch how this inverse problem arises when using an Interior Point Method in Appendix B.2. By a similar argument as Lemma 1.4.5, we can show that if we can solve linear equations in $\partial_2 \partial_2^\top$ to high accuracy in nearly-linear time, then we can solve linear equations in ∂_2 to high accuracy in nearly-linear time.

Our reduction is inspired by a reduction in [MN21] that proves NP-hardness of computing maximum *integral* flows in 2-complexes via a reduction from the graph 3-coloring problem. However, the correctness of their reduction heavily relies on the fact that the flow values in the 2-complex are 0-1 integers, which does not apply in our setting. In addition, it is unclear how to encode linear equations as a graph coloring problem even if fractional colors are allowed.

We employ some basic building blocks used in [MN21], including punctured spheres and tubes. However, we need to carefully arrange and orient the triangles in the 2-complex to encode both the positive and negative coefficients in difference-average equations, and we need to express the averaging relations not covered by the previous work.

An important aspect of our contribution is that we carefully control the number of non-zeros of the boundary operator matrix that we construct, and we bound the condition number of this matrix, and how the error propagates from an approximate solution to the boundary operator problem back to the original difference-average equations. In order to do so, we develop explicit triangulation algorithms that specify the precise number of triangles needed to triangulate each building block and allow a detailed error and condition number analysis.

We remark that our constructed 2-complex does not admit an embedding into a sphere in 3 dimensions. Recent work [Bla+22] has shown that simplicial complexes with a known embedding into \mathbb{R}^3 have non-trivial linear equation solvers, but the full extent to which embeddability can lead to better solvers remains an open question.

We analyze our construction in the RealRAM model. However, it can be transferred to the fixed point arithmetic model with $(\log N)^{O(1)}$ bits per number, where N is the size of the problem instance.

4.3 Notations and Preliminaries

4.3.1 Matrix Classes

We are interested in linear equations whose coefficient matrices belong to the following matrix classes.

1. \mathcal{G} refers to the class of *General Matrices* that have integer entries and do not have all-0 rows and all-0 columns. We refer to linear equations whose coefficient matrix is in \mathcal{G} as *general linear equations*.
2. \mathcal{DA} refers to the class of *Difference-Average Matrices* whose rows fall into two categories:

- (a) A *difference row* which has exactly two non-zero entries 1 and -1 ;
- (b) An *average row* which has exactly three non-zero entries 1, 1, and -2 .

Multiplying a difference row vector to a column vector \mathbf{x} gives $\mathbf{x}(i) - \mathbf{x}(j)$; multiplying an average row vector to \mathbf{x} gives $\mathbf{x}(i) + \mathbf{x}(j) - 2\mathbf{x}(k)$. We refer to linear equations whose coefficient matrix is in \mathcal{DA} as *difference-average linear equations*.

3. \mathcal{B}_2 refers to the class of *Boundary Operator Matrices* ∂_2 in 2-complexes. We refer to linear equations whose coefficient matrix is in \mathcal{B}_2 as *2-complex boundary linear equations*.

Our definition of “general matrices” specifies that the matrix must have integer entries. However, when the input matrix is invertible, using a simple rounding argument, we can convert any linear equation into a linear equation with integer entries $\tilde{O}(1)$ bits per entry. We caution the reader that this relies on our definition of $\tilde{O}(\cdot)$ as hiding polylogarithmic factors in the input condition number. In general, the condition number can be exponentially large – however, our results are mainly of interest when the condition number is quasipolynomially bounded.

4.3.2 Reduction Between Linear Equations

We will again follow the definition of efficient reductions in [KZ17].

We say LEA over matrix class \mathcal{M}_1 is *nearly-linear time reducible* to LEA over matrix class \mathcal{M}_2 , denoted by $\mathcal{M}_1 \leq_{nlt} \mathcal{M}_2$, if the following holds:

1. There is an algorithm that maps an arbitrary instance LEA $(\mathbf{M}_1, \mathbf{c}_1, \epsilon_1)$ where $\mathbf{M}_1 \in \mathcal{M}_1$ to an instance LEA $(\mathbf{M}_2, \mathbf{c}_2, \epsilon_2)$ where $\mathbf{M}_2 \in \mathcal{M}_2$ such that there is another algorithm that can map a solution to LEA $(\mathbf{M}_2, \mathbf{c}_2, \epsilon_2)$ to a solution to LEA $(\mathbf{M}_1, \mathbf{c}_1, \epsilon_1)$.
2. Both the two algorithms run in time $\tilde{O}(\text{nnz}(\mathbf{M}_1))$.
3. In addition, we can guarantee $\text{nnz}(\mathbf{M}_2) = \tilde{O}(\text{nnz}(\mathbf{M}_1))$, and

$$\epsilon_2^{-1}, \kappa(\mathbf{M}_2), U(\mathbf{M}_2, \mathbf{b}_2) = \text{poly}(\text{nnz}(\mathbf{M}_1), \epsilon_1^{-1}, \kappa(\mathbf{M}_1), U(\mathbf{M}_1, \mathbf{b}_1)).$$

We do not require a nearly-linear time reduction to preserve the number of variables or constraints (dimensions) of a system of linear equations. The dimensions of the new linear equation instance that we construct can be much larger than that of the original instance. On the other hand, a reduction that *only* preserves dimensions may construct a dense linear equation instance even if the original instance is sparse. A nearly-linear time reduction that preserves *both* the number of non-zeros and dimensions would be stronger than what we achieve.

Fact 4.3.1. *If $\mathcal{M}_1 \leq_{nlt} \mathcal{M}_2$ and $\mathcal{M}_2 \leq_{nlt} \mathcal{M}_3$, then $\mathcal{M}_1 \leq_{nlt} \mathcal{M}_3$.*

Definition 4.3.2 (Sparse-Linear-Equation Complete (SLE-complete)). We say LEA over a matrix class \mathcal{M} is *sparse-linear-equation complete* if $\mathcal{G} \leq_{nlt} \mathcal{M}$.

Fact 4.3.3. *Suppose LEA over \mathcal{M} is SLE-complete. If one can solve all instances LEA $(\mathbf{A}, \mathbf{b}, \epsilon)$ with $\mathbf{A} \in \mathcal{M}$ in time $\tilde{O}(\text{nnz}(\mathbf{A})^c)$ where $c \geq 1$, then one can solve all instances LEA $(\mathbf{A}', \mathbf{b}', \epsilon')$ with $\mathbf{A}' \in \mathcal{G}$ in time $\tilde{O}(\text{nnz}(\mathbf{A}')^c)$.*

Under the above definitions, [KZ17] implicitly shows the following result.

Theorem 4.3.4 (Implicitly Stated in [KZ17]). *LEA over \mathcal{DA} is SLE-complete.*

4.4 Main Results and Reduction Outline

Our main result is stated in the following theorem.

Theorem 4.4.1. *LEA over \mathcal{B}_2 is SLE-complete.*

Although our main theorem focuses on linear equation approximate problems, we construct nearly-linear time reductions for both linear equation problem LE and its approximate counterpart LEA. We first reduce LE instances (\mathbf{A}, \mathbf{b}) (and LEA instances $(\mathbf{A}, \mathbf{b}, \epsilon)$) over difference-average matrices to those over 2-complex boundary operator matrices, *under the assumption* $\mathbf{b} \in \text{Im}(\mathbf{A})$ (stated in Theorem 4.4.2 and 4.4.3). In this case, the constructed 2-complexes have unit edge weights. We then provide a slightly modified nearly-linear time reduction for LEA $(\mathbf{A}, \mathbf{b}, \epsilon)$ over difference-average matrices to LEA over 2-complex boundary operator matrices *without assuming* $\mathbf{b} \in \text{Im}(\mathbf{A})$ (stated in Theorem 4.4.4). In this case, we introduce polynomially bounded edge weights for the constructed 2-complexes.

Theorem 4.4.2. *Given a linear equation instance LE (\mathbf{A}, \mathbf{b}) where $\mathbf{A} \in \mathcal{DA}$ and $\mathbf{b} \in \text{Im}(\mathbf{A})$, we can reduce it to an instance LE (∂_2, γ) where $\partial_2 \in \mathcal{B}_2$, in time $O(\text{nnz}(\mathbf{A}))$, such that a solution to LE (∂_2, γ) can be mapped to a solution to LE (\mathbf{A}, \mathbf{b}) in time $O(\text{nnz}(\mathbf{A}))$.*

Theorem 4.4.3. *Given a linear equation instance LEA $(\mathbf{A}, \mathbf{b}, \epsilon^{DA})$ where $\mathbf{A} \in \mathcal{DA}$ and $\mathbf{b} \in \text{Im}(\mathbf{A})$, we can reduce it to an instance LEA $(\partial_2, \gamma, \epsilon^{B_2})$ where $\partial_2 \in \mathcal{B}_2$ and $\epsilon^{B_2} \leq \frac{\epsilon^{DA}}{42 \text{nnz}(\mathbf{A})}$, in time $O(\text{nnz}(\mathbf{A}))$, such that a solution to LEA $(\partial_2, \gamma, \epsilon^{B_2})$ can be mapped to a solution to LEA $(\mathbf{A}, \mathbf{b}, \epsilon^{DA})$ in time $O(\text{nnz}(\mathbf{A}))$.*

Theorem 4.4.4. *Given an instance LEA $(\mathbf{A}, \mathbf{b}, \epsilon^{DA})$ where $\mathbf{A} \in \mathcal{DA}$, we can reduce it to an instance LEA $(\mathbf{W}^{1/2}\partial_2, \mathbf{W}^{1/2}\gamma, \epsilon^{B_2})$ where $\partial_2 \in \mathcal{B}_2$ and \mathbf{W} is a diagonal matrix with non-negative diagonals, in time $O(\text{nnz}(\mathbf{A}))$. Let s, ϵ, K, U denote $\text{nnz}(\mathbf{A}), \epsilon^{DA}, \kappa(\mathbf{A}), U(\mathbf{A}, \mathbf{b})$, respectively. Then, we can guarantee that*

$$\begin{aligned} \text{nnz}(\partial_2) &= O(s), \quad U(\mathbf{W}^{1/2}\partial_2, \mathbf{W}^{1/2}\gamma) = O(sU\epsilon^{-1}), \\ \epsilon^{B_2} &= \Omega(\epsilon U^{-1}s^{-1}), \quad \kappa(\mathbf{W}^{1/2}\partial_2) = O(s^{15/2}K^2\epsilon^{-2}) \end{aligned}$$

and a solution to LEA $(\mathbf{W}^{1/2}\partial_2, \mathbf{W}^{1/2}\gamma, \epsilon^{B_2})$ can be mapped to a solution to LEA $(\mathbf{A}, \mathbf{b}, \epsilon^{DA})$ in time $O(\text{nnz}(\mathbf{A}))$.

We will prove Theorem 4.4.2 and Theorem 4.4.3 in Section 4.5, and we will prove Theorem 4.4.4 in Section 4.6.

4.4.1 Overview of Our Proof

Multiplying a 2-complex boundary operator $\partial_2 \in \mathbb{R}^{m \times t}$ to a vector $\mathbf{f} \in \mathbb{R}^t$ can be interpreted as transforming flows in the triangle space to demands in the edge space. Given $\mathbf{d} \in \mathbb{R}^d$, solving $\partial_2 \mathbf{f} = \mathbf{d}$ can be interpreted as finding flows in the triangle space subject to edge demands in \mathbf{d} . We will encode difference-average linear equations as a 2-complex flow network.

Encoding a Single Equation. We observe a simple fact: If we glue two triangles Δ_1, Δ_2 with the same orientation, then the net flow $\partial_2 \mathbf{f}$ on the shared interior edge is $\mathbf{f}(\Delta_1) - \mathbf{f}(\Delta_2)$ (see Figure 4.1 (a)); if we glue two triangles Δ_1, Δ_2 with opposite orientations, then the net flow $\partial_2 \mathbf{f}$ on the shared interior edge is $\mathbf{f}(\Delta_1) + \mathbf{f}(\Delta_2)$ (see Figure 4.1 (b)). Given an equation $\mathbf{a}^\top \mathbf{x} = b$ with the non-zero coefficients being ± 1 , we can

encode it by gluing more triangles as above and setting the demand of the shared interior edge to be b . To handle the coefficient -2 in an average equation, say $\mathbf{x}(i) + \mathbf{x}(j) - 2\mathbf{x}(k)$, we implicitly interpret it as $\mathbf{x}(i) + \mathbf{x}(j) - \mathbf{x}(k_1) - \mathbf{x}(k_2)$ together with an additional difference equation $\mathbf{x}(k_1) = \mathbf{x}(k_2)$ (see Figure 4.1 (c)).




Figure 4.1: An illustration for encoding a single equation and encoding a variable.

Encoding a Variable. We use a sphere to encode a variable involved in many equations. We can obtain an oriented triangulation of the sphere and set all the edge demand to be 0 so that all the triangles on the sphere must have an equal flow value (see Figure 4.1 (d)).

Putting It All Together. For each variable $\mathbf{x}(i)$ and the sphere for $\mathbf{x}(i)$, we create a “hole” for each equation that involves $\mathbf{x}(i)$, and then attach a tube. We can have an oriented triangulation of the tubes so that the triangles on the tubes have equal value as the triangles on the sphere. We then connect these tubes properly to encode each given difference and average equation.

Discussion.

- *Why encode difference-average equations rather than directly encoding general equations with integer coefficients?*

We can generalize the above encoding method to encode a general equation $\mathbf{g}^\top \mathbf{y} = c$ with arbitrary integer coefficients into a 2-complex with roughly $\|\mathbf{g}\|_1$ tubes. However, the encoding size required to express a general system of linear equations $\mathbf{G}\mathbf{y} = \mathbf{c}$ this way can be as large as $\Omega(\text{nnz}(\mathbf{G}) \|\mathbf{G}\|_{\max})$. This dependence on $\|\mathbf{G}\|_{\max}$ is prohibitive, and makes for a fairly weak result.

On the other hand, we can first reduce the general linear equations $\mathbf{G}\mathbf{y} = \mathbf{c}$ into difference-average linear equations $\mathbf{A}\mathbf{x} = \mathbf{b}$, where $\|\mathbf{A}\|_{\max} = 2$ and $\text{nnz}(\mathbf{A}) = O(\text{nnz}(\mathbf{G}) \log \|\mathbf{G}\|_{\max})$ (by Lemma 5.1.1). Then we can encode $\mathbf{A}\mathbf{x} = \mathbf{b}$ into a 2-complex. The encoding size required to express the difference-average linear equations as a 2-complex is thus $O(\text{nnz}(\mathbf{A}))$ (by Lemma 4.5.2). Thus, the overall encoding size required to express the original linear equation $\mathbf{G}\mathbf{y} = \mathbf{c}$ is now $\tilde{O}(\text{nnz}(\mathbf{G}))$, exponentially improving the dependence on $\|\mathbf{G}\|_{\max}$.

Therefore, the two-step reduction is a nearly-linear time reduction while the one-step reduction is not.

- *Why encode into a 2-complex rather than a 1-complex?*

We do not expect that general linear equations with integer coefficients can be efficiently encoded using a 1-complex. This would immediately imply a nearly-linear

time solver for general linear equations, as fast solvers for 1-complex operators exist (using Laplacian linear equation solvers).

4.5 Reducing \mathcal{DA} to \mathcal{B}_2 in Feasible Case

In this section, we show that the algorithm REDUCEDATOB_2 and the algorithm $\text{MAPSOLNB}_2\text{TO}\mathcal{DA}$ reduce instances $\text{LEA } (\mathbf{A}, \mathbf{b}, \epsilon)$ over \mathcal{DA} to instances $\text{LEA } (\partial_2, \gamma, \epsilon')$ over \mathcal{B}_2 , under the assumption that \mathbf{b} is in the image of \mathbf{A} . Specifically, we show that by a proper choice of ϵ , a solution to $\text{LEA } (\mathbf{A}, \mathbf{b}, \epsilon)$ can be converted to a solution to $\text{LEA } (\partial_2, \gamma, \epsilon')$ in Section 4.5.5; we upper bound the condition number of ∂_2 in Section 4.5.6.

Recall that an instance $\text{LE } (\mathbf{A}, \mathbf{b})$ over \mathcal{DA} only consists of two types of linear equations:

1. *Difference equation:* $\mathbf{x}(i) - \mathbf{x}(j) = \mathbf{b}(q)$;
2. *Average equation:* $\mathbf{x}(i) + \mathbf{x}(j) - 2\mathbf{x}(k) = 0$.

Suppose $\text{LE } (\mathbf{A}, \mathbf{b})$ has d_1 difference equations and d_2 average equations. Without loss of generality, we reorder all the equations so that the first d_1 equations are difference equations and the rest are average equations.

4.5.1 Reduction Algorithm

Given an instance $\text{LE } (\mathbf{A}, \mathbf{b})$ where \mathbf{A} is a $d \times n$ matrix in \mathcal{DA} , the following algorithm REDUCEDATOB_2 constructs a 2-complex and a system of linear equations in its boundary operator.

Algorithm ReduceDATOB_2

Input: an instance $\text{LE } (\mathbf{A}, \mathbf{b})$ where $\mathbf{A} \in \mathcal{DA}$ is a $d \times n$ matrix and $\mathbf{b} \in \mathbb{R}^d$.

Output: $(\partial_2, \gamma, \Delta^c)$ where $\partial_2 \in \mathcal{B}_2$ is an $m \times t$ matrix, $\gamma \in \mathbb{R}^m$, and Δ^c is a set of n triangles.

1. For each $i \in [n]$ and variable $\mathbf{x}(i)$ in $\text{LE } (\mathbf{A}, \mathbf{b})$, we construct a sphere \mathcal{S}_i .
2. For each $q \in [d_1]$, which corresponds to a difference equation $\mathbf{x}(i) - \mathbf{x}(j) = \mathbf{b}(q)$, we add a *loop* α_q with a net flow demand $\mathbf{b}(q)$. Then,
 - (a) we add a boundary component $\beta_{q,i}$ on \mathcal{S}_i , and a boundary component $\beta_{q,j}$ on \mathcal{S}_j ;
 - (b) we construct a tube $\mathcal{T}_{q,i}$ with boundary components $\{-\beta_{q,i}, \alpha_q\}$, and a tube $\mathcal{T}_{q,j}$ with boundary components $\{-\beta_{q,j}, -\alpha_q\}$.

See Figure 4.2 for an illustration.¹

3. For each $q \in \{d_1 + 1, \dots, d\}$, which corresponds to an average equation $\mathbf{x}(i) + \mathbf{x}(j) - 2\mathbf{x}(k) = \mathbf{b}(q) = 0$, we add a loop α_q with zero net flow demand. Then,

¹Note that since the loop α_q has demand $\mathbf{b}(q)$, our construction is different from identifying the boundary component α_q of $\mathcal{T}_{q,i}$ and the boundary component $-\alpha_q$ of $\mathcal{T}_{q,j}$.




Figure 4.2: The construction for a difference equation $\mathbf{x}(i) - \mathbf{x}(j) = \mathbf{b}(q)$.

- (a) we add a boundary component $\beta_{q,i}$ on \mathcal{S}_i , a boundary component $\beta_{q,j}$ on \mathcal{S}_j , and two boundary components $\beta_{q,k,1}, \beta_{q,k,2}$ on \mathcal{S}_k ;
- (b) we construct a tube $\mathcal{T}_{q,i}$ with boundary components $\{-\beta_{q,i}, \alpha_q\}$, a tube $\mathcal{T}_{q,j}$ with boundary components $\{-\beta_{q,j}, \alpha_q\}$, and two tubes $\mathcal{T}_{q,k,1}, \mathcal{T}_{q,k,2}$ with boundary components $\{-\beta_{q,k,1}, -\alpha_q\}$ and $\{-\beta_{q,k,2}, -\alpha_q\}$, respectively.

See Figure 4.3 for an illustration.²




Figure 4.3: The construction for an average equation $\mathbf{x}(i) + \mathbf{x}(j) - 2\mathbf{x}(k) = 0$.

4. For each $i \in [n]$, the punctured sphere \mathcal{S}_i and the tubes connected to \mathcal{S}_i form a continuous topological space. We construct an oriented triangulation for this space

²As four tubes are connected to a single loop, to avoid the intersection of tubes before attaching the loop, a higher-dimensional space is required.

such that the induced orientation of each edge on a loop α_q is consistent with the orientation of α_q . We will describe this oriented triangulation subroutine in Section 4.5.1. Let \mathcal{K} be the oriented 2-complex. Let ∂_2 be the boundary operator of \mathcal{K} .

5. Each edge on a loop α_q has net demand $\mathbf{b}(q)$; each other edge has net demand 0. Let γ be the vector of the net flow demands.
6. On each triangulated sphere \mathcal{S}_i , we choose an arbitrary triangle $\Delta_i \in \mathcal{S}_i$ as the *central triangle*. Let Δ^c be the set of all the central triangles.
7. We return $(\partial_2, \gamma, \Delta^c)$.

The following algorithm $\text{MAPSOLN}\mathcal{B}_2\text{TO}\mathcal{DA}$ maps a solution \mathbf{f} to LE (∂_2, γ) to a solution \mathbf{x} to LE (\mathbf{A}, \mathbf{b}) .

Algorithm MapSoln \mathcal{B}_2 to \mathcal{DA}

Input: a tuple $(\mathbf{A}, \mathbf{b}, \mathbf{f}, \Delta^c)$, where $\mathbf{A} \in \mathcal{DA}$ is a $d \times n$ matrix, $\mathbf{b} \in \mathbb{R}^d$, $\mathbf{f} \in \mathbb{R}^t$, and Δ^c is the set of n central triangles.

Output: a vector $\mathbf{x} \in \mathbb{R}^n$.

1. If $\mathbf{A}^\top \mathbf{b} = \mathbf{0}$, we return $\mathbf{x} = \mathbf{0}$.
2. Otherwise, we set $\mathbf{x}(i) = \mathbf{f}(\Delta_i)$, where $\Delta_i \in \Delta^c$ is the central triangle on sphere \mathcal{S}_i .

Oriented Triangulation

We provide a concrete triangulation subroutine here for the benefit of analyzing our reduction algorithm.

Oriented Triangulation for Punctured Spheres. By our construction, each sphere \mathcal{S}_i has $b_i = \sum_{q=1}^d |\mathbf{A}(q, i)|$ boundary components. We will create \tilde{t}_i triangles and \tilde{m}_i edges on \mathcal{S}_i , based on b_i .

1. If $b_i = 1$ (see Figure 4.4 (a)), the punctured sphere is topologically equivalent to a disk. In this case, \mathcal{S}_i can be triangulated using a single triangle $[v_{(1)}^1, v_{(1)}^2, v_{(1)}^3]$, thus $\tilde{t}_i = 1, \tilde{m}_i = 3$.
2. If $b_i = 2$ (see Figure 4.4 (b)), the punctured sphere is topologically equivalent to an annulus. We subdivide the triangle $[v_{(1)}^1, v_{(1)}^2, v_{(1)}^3]$ obtained in the previous case by adding 6 interior edges between vertices of the inner and the outer boundaries: $[v_{(1)}^1, v_{(2)}^1], [v_{(1)}^1, v_{(2)}^2], [v_{(1)}^2, v_{(2)}^1], [v_{(1)}^2, v_{(2)}^2], [v_{(1)}^3, v_{(2)}^2], [v_{(1)}^3, v_{(2)}^3]$, thus $\tilde{t}_i = 6, \tilde{m}_i = 12$.
3. Generally, if $b_i = k$ (see Figure 4.4 (c)), we subdivide the rightmost triangle $[v_{(1)}^1, v_{(1)}^2, v_{(k-1)}^1]$ obtained in the case of $b_i = k - 1$ with the same method. By induction, we have

$$\tilde{t}_i = 5b_i - 4, \quad \tilde{m}_i = 9b_i - 6, \quad \text{for } b_i \geq 1. \quad (4.1)$$




Figure 4.4: Oriented triangulation of punctured spheres. The light area represents the “holes” defined by boundary components.

The orientation for triangles on the same sphere should be identical. Without loss of generality, we orient all triangles clockwise. Note that with this triangulation method, all boundary components are composed of 3 edges.

Oriented Triangulation for Tubes. A tube is defined by two boundary components. By our construction, for every tube connected to \mathcal{S}_i , one of the two boundary components is always $-\beta_{q,i,*}$ ³ and the other one is $\pm\alpha_q$, whose orientation depends on the sign of the entry $\mathbf{A}(q, i)$. Without loss of generality, we orient anti-clockwise for all α_q , thus clockwise for all $-\alpha_q$. Therefore, there are two possibilities of boundary component combinations.

1. If $\mathbf{A}(q, i) > 0$ (see Figure 4.5 (a)), then the two boundary components have opposite orientations: $-\beta_{q,i,*} = [v_{q,i,*}^1, v_{q,i,*}^3, v_{q,i,*}^2]$ and $\alpha_q = [v_q^1, v_q^2, v_q^3]$. We triangulate by matching $v_{q,i,*}^1$ to v_q^1 , $v_{q,i,*}^2$ to v_q^2 , and $v_{q,i,*}^3$ to v_q^3 .
2. If $\mathbf{A}(q, i) < 0$ (see Figure 4.5 (b)), then the two boundary components have identical orientations: $-\beta_{q,i,*} = [v_{q,i,*}^1, v_{q,i,*}^3, v_{q,i,*}^2]$ and $-\alpha_q = [v_q^1, v_q^3, v_q^2]$. We triangulate by matching $v_{q,i,*}^1$ to v_q^1 , $v_{q,i,*}^3$ to v_q^3 , and $v_{q,i,*}^2$ to v_q^2 .




Figure 4.5: Oriented triangulation of tubes with opposite or identical boundary orientations.

In either case, only 6 triangles and 12 edges are required for an oriented triangulation of any tube $\mathcal{T}_{q,i,*}$. Again, we orient all triangles clockwise.

³We introduce a third element $* \in \{1, 2\}$ in the subscript of $\beta_{q,k,*}$, which is activated only when $\mathbf{A}(q, k) = -2$.

4.5.2 Additional Notations

To facilitate our analysis, we introduce several notions and corresponding notations about the constructed 2-complex. These notions and notations will be used in the rest of the paper.

For each $i \in [n]$, let \mathcal{K}_i be the union of the triangulated \mathcal{S}_i and the triangulated tubes that are connected to \mathcal{S}_i , which we refer to as the i th *complex group*; let t_i be the number of triangles in \mathcal{K}_i ; let m_i be the number of the interior edges in \mathcal{K}_i .

We refer to an edge on a loop α_q as a *boundary edge* and an edge not on any loop as an *interior edge*. According to our triangulation, each loop α_q has three boundary edges, denoted by $\alpha_q^1 = [v_q^1, v_q^2]$, $\alpha_q^2 = [v_q^2, v_q^3]$, $\alpha_q^3 = [v_q^3, v_q^1]$. A triangle containing a boundary edge is called a *boundary triangle*. For each boundary edge α_q^r , where $q \in [d_1]$ and $r \in [3]$, corresponding to equation $\mathbf{x}(i) - \mathbf{x}(j) = \mathbf{b}(q)$, we denote the boundary triangles by $\Delta_{q,i,1}^r, \Delta_{q,j,1}^r$ where $\Delta_{q,i,1}^r \in \mathcal{T}_{q,i}$, $\Delta_{q,j,1}^r \in \mathcal{T}_{q,j}$; for each boundary edge α_q^r , where $q \in [d_1+1 : d_2]$ and $r \in [3]$, corresponding to equation $\mathbf{x}(i) + \mathbf{x}(j) = 2\mathbf{x}(k)$, we denote the boundary triangles by $\Delta_{q,i,1}^r, \Delta_{q,j,1}^r, \Delta_{q,k,1}^r, \Delta_{q,k,2}^r$ where $\Delta_{q,i,1}^r \in \mathcal{T}_{q,i}$, $\Delta_{q,j,1}^r \in \mathcal{T}_{q,j}$, $\Delta_{q,k,1}^r, \Delta_{q,k,2}^r \in \mathcal{T}_{q,k}$.

Given any two triangles $\Delta, \Delta' \in \mathcal{K}$, a *triangle path* from Δ to Δ' is an ordered collection of triangles $\mathcal{P} = [\Delta^{(0)} = \Delta, \Delta^{(1)}, \dots, \Delta^{(l)} = \Delta']$ such that every pair of neighboring triangles shares an edge. The length of \mathcal{P} is l . A triangle path can also be defined by an ordered collection of edges $\mathcal{P} = [e^{(1)}, \dots, e^{(l)}]$, where $e^{(i)}$ denotes the edge shared by $\Delta^{(i-1)}$ and $\Delta^{(i)}$, for $i \in [l]$.

4.5.3 Algorithm Runtime and Problem Size

In this section, we firstly show that the reduction algorithm $\text{REDUCE}\mathcal{DATo}\mathcal{B}_2$ and the solution mapping algorithm $\text{MAPSOLN}\mathcal{B}_2\text{To}\mathcal{DA}$ both run in linear time, and $\text{REDUCE}\mathcal{DATo}\mathcal{B}_2$ constructs a 2-complex whose size is linear in the number of non-zeros in the input linear equations.

Lemma 4.5.1 (Runtime). *Given a difference-average instance $\text{LE}(\mathbf{A}, \mathbf{b})$ where $\mathbf{A} \in \mathbb{R}^{d \times n}$, Algorithm $\text{REDUCE}\mathcal{DATo}\mathcal{B}_2(\mathbf{A}, \mathbf{b})$ returns $(\partial_2, \gamma, \Delta^c)$ in time $O(\text{nnz}(\mathbf{A}))$. Given a solution \mathbf{f} to $\text{LE}(\partial_2, \gamma)$, Algorithm $\text{MAPSOLN}\mathcal{B}_2\text{To}\mathcal{DA}(\mathbf{A}, \mathbf{b}, \mathbf{f}, \Delta^c)$ returns \mathbf{x} in time $O(n)$.*

Proof. For reduction, $\text{REDUCE}\mathcal{DATo}\mathcal{B}_2(\mathbf{A}, \mathbf{b})$ calls the tube triangulation subroutine for $\|\mathbf{A}\|_1$ times, and the punctured sphere triangulation subroutine for n times. The tube triangulation subroutine runs in time $O(1)$ since there are a constant number of triangles in a tube; and the punctured sphere triangulation subroutine runs in time $O(\|\mathbf{A}(:, j)\|_1)$ for the j th call, $j \in [n]$. Putting all together, the total runtime of $\text{REDUCE}\mathcal{DATo}\mathcal{B}_2(\mathbf{A}, \mathbf{b})$ is $O\left(\|\mathbf{A}\|_1 + \sum_{j \in [n]} \|\mathbf{A}(:, j)\|_1\right) \leq O(\text{nnz}(\mathbf{A}))$, where we use the fact $\|\mathbf{A}\|_{\max} = 2$.

For solution mapping, the runtime of the algorithm $\text{MAPSOLN}\mathcal{B}_2\text{To}\mathcal{DA}$ is obvious. \square

Lemma 4.5.2 (Size of ∂_2). *Given a difference-average instance $\text{LE}(\mathbf{A}, \mathbf{b})$, let $(\partial_2, \gamma, \Delta^c)$ be returned by $\text{REDUCE}\mathcal{DATo}\mathcal{B}_2(\mathbf{A}, \mathbf{b})$. Suppose $\partial_2 \in \mathbb{R}^{m \times t}$. Then,*

- $t \leq 22 \text{nnz}(\mathbf{A})$;
- $m \leq 33 \text{nnz}(\mathbf{A})$;
- $\text{nnz}(\partial_2) \leq 66 \text{nnz}(\mathbf{A})$.

Proof. We first compute the total number of triangles in the constructed 2-complex \mathcal{K} . For sphere \mathcal{S}_j , we have $\tilde{t}_j = 5b_j - 4$ triangles by (4.1), where $b_j = \sum_{i \in [d]} |\mathbf{A}(i, j)|$. Therefore, the number of triangles of all spheres is

$$\sum_{j=1}^n \tilde{t}_j = \sum_{j=1}^n \left(5 \sum_{i \in [d]} |\mathbf{A}(i, j)| - 4 \right) = 5 \|\mathbf{A}\|_1 - 4n.$$

Moreover, each boundary component on spheres corresponds to a tube, and each tube has 6 triangles. Hence, the number of triangles of all tubes is $6 \|\mathbf{A}\|_1$. Putting spheres and tubes together, we get

$$t = 11 \|\mathbf{A}\|_1 - 4n \leq 22 \text{nnz}(\mathbf{A}),$$

where the last inequality is because entries of \mathbf{A} are bounded by 2.

Next, we compute the total number of edges in \mathcal{K} . By construction, each triangle has 3 incident edges and each edge is shared by a constant number of triangles (2 for interior edges, and 4 for boundary edges). Thus, we have

$$m \leq 1.5t \leq 33 \text{nnz}(\mathbf{A}).$$

Since each column of ∂_2 has exactly 3 non-zero entries, we have

$$\text{nnz}(\partial_2) = 3t \leq 66 \text{nnz}(\mathbf{A}).$$

□

4.5.4 Relation Between Exact Solutions

In this section, we show that the algorithm REDUCEDATOB_2 and the algorithm $\text{MAPSOLNB}_2\text{TO}\mathcal{DA}$ reduce instances LE (\mathbf{A}, \mathbf{b}) over \mathcal{DA} to instances LE over \mathcal{B}_2 by exploring the relationship between exact solutions.

Given LE (\mathbf{A}, \mathbf{b}) where \mathbf{A} is a $d \times n$ matrix in \mathcal{DA} , let $(\partial_2, \gamma, \Delta^c)$ be returned by $\text{REDUCEDATOB}_2(\mathbf{A}, \mathbf{b})$, and let \mathcal{K} be the 2-complex constructed in $\text{REDUCEDATOB}_2(\mathbf{A}, \mathbf{b})$ and the boundary operator of \mathcal{K} is ∂_2 . Let \mathbf{f} be a solution to LE (\mathbf{A}, \mathbf{b}) .

To simplify the analysis, we reorder the columns and rows of ∂_2 . The columns $[1 : t_1]$ of ∂_2 correspond to the triangles in \mathcal{K}_1 , the columns $[t_1 + 1 : t_2]$ correspond to the triangles in \mathcal{K}_2 , and so on. Then, \mathbf{f} can be written as

$$\mathbf{f} = \begin{bmatrix} \mathbf{f}_1 \\ \vdots \\ \mathbf{f}_n \end{bmatrix}, \quad \text{where } \mathbf{f}_i \in \mathbb{R}^{t_i}, \forall i \in [n].$$

The rows of ∂_2 and the entries of γ are:

$$\partial_2 = \begin{bmatrix} & \mathbf{B}_1 & & \\ & \vdots & & \\ & \mathbf{B}_d & & \\ \mathbf{M}_1 & & & \\ & \ddots & & \\ & & \mathbf{M}_n & \end{bmatrix}, \quad \gamma = \begin{bmatrix} \mathbf{b}(1)\mathbf{1}_3 \\ \vdots \\ \mathbf{b}(d)\mathbf{1}_3 \\ \mathbf{0}_{m_1} \\ \vdots \\ \mathbf{0}_{m_n} \end{bmatrix}, \quad (4.2)$$

Here, each submatrix $\mathbf{B}_q \in \{0, \pm 1\}^{3 \times t}$ corresponds to the three boundary edges $\{\alpha_q^1, \alpha_q^2, \alpha_q^3\}$; each submatrix $\mathbf{M}_i \in \{0, \pm 1\}^{m_i \times t_i}$ corresponds to all the interior edges in \mathcal{K}_i . Interior edges in \mathcal{K}_i and those in \mathcal{K}_j do not share endpoints if $i \neq j$. Let $\mathbf{M} = \text{diag}(\mathbf{M}_1, \mathbf{M}_2, \dots, \mathbf{M}_n)$.

Claim 4.5.3. *For each $i \in [n]$, $\mathbf{f}_i = \alpha \mathbf{1}_{t_i}$ for some $\alpha \in \mathbb{R}$.*

Proof. For each $i \in [n]$, we have $\mathbf{M}_i \mathbf{f}_i = \mathbf{0}$. This means that for any two triangles Δ, Δ' in \mathcal{K}_i sharing an interior edge, we have $\mathbf{f}_i(\Delta) = \mathbf{f}_i(\Delta')$. By our construction of \mathcal{K}_i , for any two triangles Δ, Δ' in \mathcal{K}_i , there exists a triangle path connecting Δ and Δ' . The values of \mathbf{f}_i at the triangles in this triangle path are equal; in particular, $\mathbf{f}_i(\Delta) = \mathbf{f}_i(\Delta')$. Thus, the values of \mathbf{f}_i at all the triangles in \mathcal{K}_i are equal, that is, $\mathbf{f}_i = \alpha \mathbf{1}_{t_i}$ for some $\alpha \in \mathbb{R}$. \square

Lemma 4.5.4 (Exact Solvers in Feasible Case). *Given a difference-average instance LE (\mathbf{A}, \mathbf{b}) where $\mathbf{b} \in \text{Im}(\mathbf{A})$, let $(\partial_2, \gamma, \Delta^c)$ be returned by $\text{REDUCEDATOB}_2(\mathbf{A}, \mathbf{b})$, and let \mathbf{f} be a solution to LE (∂_2, γ) . Then, $\mathbf{x} \leftarrow \text{MAPSOLNB}_2\text{TO}\mathcal{DA}(\mathbf{A}, \mathbf{b}, \mathbf{f}, \Delta^c)$ is a solution to LE (\mathbf{A}, \mathbf{b}) .*

Proof. By Claim 4.5.3, we can write \mathbf{f} as

$$\mathbf{f} = \begin{bmatrix} \alpha_1 \mathbf{1}_{t_1} \\ \alpha_2 \mathbf{1}_{t_2} \\ \vdots \\ \alpha_n \mathbf{1}_{t_n} \end{bmatrix}, \quad \text{where } \alpha_1, \dots, \alpha_n \in \mathbb{R}.$$

According to Algorithm $\text{MAPSOLNB}_2\text{TO}\mathcal{DA}(\mathbf{A}, \mathbf{b}, \mathbf{f}, \Delta^c)$, for each $i \in [n]$, $\mathbf{x}(i) = \alpha_i$. Our goal is to show $\mathbf{Ax} = \mathbf{b}$.

- For each difference equation in LE (\mathbf{A}, \mathbf{b}) , say $\mathbf{x}(i) - \mathbf{x}(j) = \mathbf{b}(q)$, we look at the equations in LE (∂_2, γ) related to \mathbf{B}_q :

$$\mathbf{f}_i(\Delta_{q,i}^r) - \mathbf{f}_j(\Delta_{q,j}^r) = \mathbf{b}(q), \quad \forall r \in \{1, 2, 3\}.$$

Thus $\mathbf{x}(i) - \mathbf{x}(j) = \mathbf{b}(q)$ holds.

- For each average equation in LE (\mathbf{A}, \mathbf{b}) , say $\mathbf{x}(i) + \mathbf{x}(j) - 2\mathbf{x}(k) = \mathbf{b}(q) = 0$, we look at the equations in LE (∂_2, γ) related to \mathbf{B}_q :

$$\mathbf{f}_i(\Delta_{q,i}^r) + \mathbf{f}_j(\Delta_{q,j}^r) - \mathbf{f}_k(\Delta_{q,k,1}^r) - \mathbf{f}_k(\Delta_{q,k,2}^r) = 0, \quad \forall r \in \{1, 2, 3\}.$$

Thus $\mathbf{x}(i) + \mathbf{x}(j) - 2\mathbf{x}(k) = 0$ holds. \square

4.5.5 Relation Between Approximate Solutions

Lemma 4.5.5 (Approximate Solvers in Feasible Case). *Given a difference-average instance LEA $(\mathbf{A}, \mathbf{b}, \epsilon^{DA})$ where $\mathbf{b} \in \text{Im}(\mathbf{A})$, let $(\partial_2, \gamma, \Delta^c)$ be returned by $\text{REDUCEDATOB}_2(\mathbf{A}, \mathbf{b})$. Suppose \mathbf{f} is a solution to LEA $(\partial_2, \gamma, \epsilon^{B_2})$ where*

$$\epsilon^{B_2} \leq \frac{\epsilon^{DA}}{42 \text{nnz}(\mathbf{A})},$$

and \mathbf{x} is returned by $\text{MAPSOLNB}_2\text{TO}\mathcal{DA}(\mathbf{A}, \mathbf{b}, \mathbf{f}, \Delta^c)$. Then, \mathbf{x} is a solution to LEA $(\mathbf{A}, \mathbf{b}, \epsilon^{DA})$.

Proof. Since $\mathbf{b} \in \text{Im}(\mathbf{A})$, we have $\|\partial_2 \mathbf{f} - \gamma\|_\infty \leq \|\partial_2 \mathbf{f} - \gamma\|_2 \leq \epsilon^{B_2} \|\gamma\|_2$.

We claim

$$\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_\infty \leq 24 \text{nnz}(\mathbf{A})^{1/2} \cdot \epsilon^{B_2} \|\gamma\|_2. \quad (4.3)$$

Then,

$$\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2 \leq 24 \text{nnz}(\mathbf{A}) \cdot \epsilon^{B_2} \|\gamma\|_2.$$

By $\text{REDUCEDATOB}_2(\mathbf{A}, \mathbf{b})$, $\|\gamma\|_2 \leq \sqrt{3} \|\mathbf{b}\|_2$. Thus,

$$\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2 \leq 42 \text{nnz}(\mathbf{A})^{1/2} \cdot \epsilon^{B_2} \|\mathbf{b}\|_2 \leq \epsilon^{DA} \|\mathbf{b}\|_2,$$

that is, \mathbf{x} is a solution to LEA $(\mathbf{A}, \mathbf{b}, \epsilon^{DA})$.

To prove Eq. (4.3), consider an arbitrary equation in LEA $(\mathbf{A}, \mathbf{b}, \epsilon^{DA})$, say $a_i \mathbf{x}(i) + a_j \mathbf{x}(j) + a_k \mathbf{x}(k) = \mathbf{b}(q)$ where $a_i, a_j, a_k \in \{-2, -1, 0, 1\}$. According to $\text{MAPSOLNB}_2\text{TO}\mathcal{DA}(\mathbf{A}, \mathbf{b}, \mathbf{f}, \Delta^c)$, for each $l \in \{i, j, k\}$, $\mathbf{x}(l) = \mathbf{f}(\Delta_l)$ where $\Delta_l \in \mathcal{K}_l$ is the l th central triangle in Δ^c . Then,

$$a_i \mathbf{x}(i) + a_j \mathbf{x}(j) + a_k \mathbf{x}(k) = a_i \mathbf{f}(\Delta_i) + a_j \mathbf{f}(\Delta_j) + a_k \mathbf{f}(\Delta_k).$$

Note that the equation in LEA $(\partial_2, \gamma, \epsilon^{B_2})$ related to the boundary edge α_q^1 , shared by triangles $\Delta_{q,i,1}^1, \Delta_{q,j,1}^1$ and $\Delta_{q,k,1}^1, \Delta_{q,k,2}^1$ (if the equation is an average equation), satisfies

$$\left| a_i \mathbf{f}(\Delta_{q,i,1}^1) + a_j \mathbf{f}(\Delta_{q,j,1}^1) + \frac{1}{2} a_k (\mathbf{f}(\Delta_{q,k,1}^1) + \mathbf{f}(\Delta_{q,k,2}^1)) - \mathbf{b}(q) \right| \leq \epsilon^{B_2} \|\gamma\|_2.$$

For each $\Delta \in \{\Delta_{q,i,1}^1, \Delta_{q,j,1}^1, \Delta_{q,k,1}^1, \Delta_{q,k,2}^1\}$, we will replace $\mathbf{f}(\Delta)$ with its corresponding central triangle Δ^c . We can find a triangle path connecting Δ and Δ^c , say $\mathcal{P} = [\Delta^{(0)} = \Delta, \dots, \Delta^{(l_q)} = \Delta^c]$, such that two adjacent triangles $\Delta^{(l)}, \Delta^{(l+1)}$ share an interior edge $e^{(l)}$. Then,

$$\begin{aligned} |\mathbf{f}(\Delta) - \mathbf{f}(\Delta^c)| &= \left| \sum_{l=1}^{l_q} \mathbf{f}(\Delta^{(l-1)}) - \mathbf{f}(\Delta^{(l)}) \right| \leq \sum_{l=1}^{l_q} |\mathbf{f}(\Delta^{(l-1)}) - \mathbf{f}(\Delta^{(l)})| = \sum_{l=1}^{l_q} |[\partial_2 \mathbf{f}](e^{(l)})| \\ &= \sum_{l=1}^{l_q} |[\partial_2 \mathbf{f} - \gamma](e^{(l)})| \quad \text{since } \gamma(e^{(l)}) = 0 \text{ for interior edges} \\ &= \|[\partial_2 \mathbf{f} - \gamma](e^{(1)} : e^{(l_q)})\|_1 \\ &\quad \text{where the subvector corresponds to } [e^{(1)}, \dots, e^{(l_q)}] \\ &\leq \sqrt{t_i} \|[\partial_2 \mathbf{f} - \gamma](e^{(1)} : e^{(l_q)})\|_2 \quad \text{since } l_q \leq t_i \\ &\leq \sqrt{t_i} \|\partial_2 \mathbf{f} - \gamma\|_2 \\ &\leq \sqrt{t_i} \cdot \epsilon^{B_2} \|\gamma\|_2. \end{aligned}$$

Thus,

$$\begin{aligned}
& |a_i \mathbf{x}(i) + a_j \mathbf{x}(j) + a_k \mathbf{x}(k) - \mathbf{b}(q)| \\
&= |a_i \mathbf{f}(\Delta_i) + a_j \mathbf{f}(\Delta_j) + a_k \mathbf{f}(\Delta_k) - \mathbf{b}(q)| \\
&\leq \underbrace{\left| a_i \mathbf{f}(\Delta_{q,i,1}^1) + a_j \mathbf{f}(\Delta_{q,j,1}^1) + \frac{1}{2} a_k (\mathbf{f}(\Delta_{q,k,1}^1) + \mathbf{f}(\Delta_{q,k,2}^1)) - \mathbf{b}(q) \right|}_{\leq \epsilon^{B_2} \|\gamma\|_2} \\
&\quad + \underbrace{|\mathbf{f}(\Delta_{q,i,1}^1) - \mathbf{f}(\Delta_i)|}_{\leq \sqrt{t_i} \epsilon^{B_2} \|\gamma\|_2} + \underbrace{|\mathbf{f}(\Delta_{q,j,1}^1) - \mathbf{f}(\Delta_j)|}_{\leq \sqrt{t_j} \epsilon^{B_2} \|\gamma\|_2} + \underbrace{|\mathbf{f}(\Delta_{q,k,1}^1) - \mathbf{f}(\Delta_k)|}_{\leq \sqrt{t_k} \epsilon^{B_2} \|\gamma\|_2} + \underbrace{|\mathbf{f}(\Delta_{q,k,2}^1) - \mathbf{f}(\Delta_k)|}_{\leq \sqrt{t_k} \epsilon^{B_2} \|\gamma\|_2} \\
&\leq 5\sqrt{t} \cdot \epsilon^{B_2} \|\gamma\|_2 \\
&\leq 24 \text{nnz}(\mathbf{A})^{1/2} \cdot \epsilon^{B_2} \|\gamma\|_2, \tag*{by Lemma 4.5.2}
\end{aligned}$$

That is, $\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_\infty \leq 24 \text{nnz}(\mathbf{A})^{1/2} \cdot \epsilon^{B_2} \|\gamma\|_2$. \square

4.5.6 Bounding the Condition Number of the New Matrix

In this section, we show that the condition number of ∂_2 is upper bounded by a polynomial of $\text{nnz}(\mathbf{A})$, $\kappa(\mathbf{A})$.

Lemma 4.5.6 (Condition Number of ∂_2). *Given a difference-average instance $\text{LEA}(\mathbf{A}, \mathbf{b}, \epsilon^{DA})$ where $\mathbf{b} \in \text{Im}(\mathbf{A})$, let $(\partial_2, \gamma, \Delta^c)$ be returned by $\text{REDUCEDATOB}_2(\mathbf{A}, \mathbf{b})$. Then,*

$$\kappa(\partial_2) \leq 10^9 \text{nnz}(\mathbf{A})^{9/2} \kappa(\mathbf{A})^2.$$

Note that

$$\kappa^2(\partial_2) = \kappa(\partial_2^\top \partial_2) = \frac{\lambda_{\max}(\partial_2^\top \partial_2)}{\lambda_{\min}(\partial_2^\top \partial_2)}.$$

We will upper bound $\lambda_{\max}(\partial_2^\top \partial_2)$ and lower bound $\lambda_{\min}(\partial_2^\top \partial_2)$. Our proof will heavily rely on the Courant-Fischer theorem.

Theorem 4.5.7 (The Courant-Fischer Theorem). *Let \mathbf{M} be a symmetric matrix in $\mathbb{R}^{n \times n}$ where $\lambda_{\max}, \lambda_{\min}$ are its maximum and minimum non-zero eigenvalue, respectively. Then*

$$\lambda_{\max} = \max_{\mathbf{x} \neq \mathbf{0}} \frac{\mathbf{x}^\top \mathbf{M} \mathbf{x}}{\mathbf{x}^\top \mathbf{x}}, \quad \lambda_{\min} = \min_{\mathbf{x} \perp \text{Ker}(\mathbf{M}), \mathbf{x} \neq \mathbf{0}} \frac{\mathbf{x}^\top \mathbf{M} \mathbf{x}}{\mathbf{x}^\top \mathbf{x}}.$$

The Maximum Eigenvalue

Lemma 4.5.8 (Maximum Eigenvalue). $\lambda_{\max}(\partial_2^\top \partial_2) \leq 12$.

Proof. By the Courant-Fischer theorem,

$$\lambda_{\max}(\partial_2^\top \partial_2) = \max_{\mathbf{f}: \|\mathbf{f}\|_2=1} \mathbf{f}^\top \partial_2^\top \partial_2 \mathbf{f}.$$

Then for any \mathbf{f} with $\|\mathbf{f}\|_2 = 1$,

$$\mathbf{f}^\top \partial_2^\top \partial_2 \mathbf{f} = \sum_{i=1}^m \left(\sum_{\Delta: e_i \in \Delta} \partial_2(e_i, \Delta) \mathbf{f}(\Delta) \right)^2 \leq 4 \sum_{i=1}^m \sum_{\Delta: e_i \in \Delta} \mathbf{f}^2(\Delta) = 12 \|\mathbf{f}\|_2^2 = 12.$$

where the inequality is by the Cauchy-Schwarz inequality. \square

The Minimum Non-Zero Eigenvalue

We start with proving a relation between the null space of \mathbf{A} and that of ∂_2 .

Lemma 4.5.9. *Let*

$$\mathbf{H} = \begin{bmatrix} \mathbf{1}_{t_1} & & \\ & \ddots & \\ & & \mathbf{1}_{t_n} \end{bmatrix} \in \mathbb{R}^{t \times n}.$$

Then, \mathbf{H} is a bijection from $\text{Ker}(\mathbf{A})$ to $\text{Ker}(\partial_2)$.

Proof. Let $\mathbf{x} \in \text{Ker}(\mathbf{A})$. By our construction of ∂_2 , we have $\mathbf{H}\mathbf{x} \in \text{Ker}(\partial_2)$. For any $\mathbf{f} \in \text{Ker}(\partial_2)$, by the proof of Lemma 4.5.4, $\mathbf{f} = \mathbf{H}\mathbf{x}$ for some $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{A}\mathbf{x} = \mathbf{0}$. \square

Lemma 4.5.10 (Minimum Non-Zero Eigenvalue).

$$\lambda_{\min}(\partial_2^\top \partial_2) \geq \frac{\min\{\lambda_{\min}(\mathbf{A}^\top \mathbf{A})^2, 1\}}{10^{16}d^7}. \quad (4.4)$$

Proof. By the Courant-Fischer theorem,

$$\lambda_{\min}(\partial_2^\top \partial_2) = \min_{\substack{\mathbf{f} \in \mathbb{R}^t: \mathbf{f} \perp \text{Ker}(\partial_2) \\ \|\mathbf{f}\|_2=1}} \mathbf{f}^\top \partial_2^\top \partial_2 \mathbf{f}.$$

Let

$$C = \frac{\min\{\lambda_{\min}(\mathbf{A}^\top \mathbf{A}), 1\}}{10^6 d^{2.5}}.$$

We will exhaust all the vectors in $\{\mathbf{f} : \mathbf{f} \perp \text{Ker}(\partial_2), \|\mathbf{f}\|_2 = 1\}$ by the following two cases.

Case 1. Suppose there exists $i \in [n]$ such that \mathcal{K}_i contains two triangles Δ, Δ' satisfying $|\mathbf{f}(\Delta) - \mathbf{f}(\Delta')| \geq C$. Consider a triangle path in \mathcal{K}_i connecting Δ and Δ' , say $[\Delta^{(0)} = \Delta, \Delta^{(1)}, \dots, \Delta^{(l)} = \Delta']$ where $l \leq 22d$. There must exist $i^* \in [l]$ such that

$$|\mathbf{f}(\Delta^{(i^*-1)}) - \mathbf{f}(\Delta^{(i^*)})| \geq \frac{C}{l}.$$

Note that

$$\mathbf{f}^\top \partial_2^\top \partial_2 \mathbf{f} \geq (\mathbf{f}(\Delta^{(i^*-1)}) - \mathbf{f}(\Delta^{(i^*)}))^2 \geq \left(\frac{C}{l}\right)^2 \geq \frac{\min\{\lambda_{\min}(\mathbf{A}^\top \mathbf{A})^2, 1\}}{10^{16}d^7}.$$

Case 2. Suppose for every $i \in [n]$ and every two $\Delta, \Delta' \in \mathcal{K}_i$, we have $|\mathbf{f}(\Delta) - \mathbf{f}(\Delta')| < C$. We write \mathbf{f} as

$$\mathbf{f} = \tilde{\mathbf{f}} + \boldsymbol{\epsilon} = \begin{bmatrix} \alpha_1 \mathbf{1}_{t_1} \\ \alpha_2 \mathbf{1}_{t_2} \\ \vdots \\ \alpha_n \mathbf{1}_{t_n} \end{bmatrix} + \boldsymbol{\epsilon},$$

where α_i is the value of \mathbf{f} at the central triangle of \mathcal{K}_i . Then,

$$\begin{aligned} \|\boldsymbol{\epsilon}\|_2 &< \sqrt{t}C \leq \sqrt{22d}C = o(1), \\ \|\tilde{\mathbf{f}}\|_2 &= \|\mathbf{f} - \boldsymbol{\epsilon}\|_2 \in \left(\frac{1}{2}, 2\right). \end{aligned}$$

We lower bound the quadratic value:

$$\begin{aligned} \mathbf{f}^\top \partial_2^\top \partial_2 \mathbf{f} &= \tilde{\mathbf{f}}^\top \partial_2^\top \partial_2 \tilde{\mathbf{f}} + 2\tilde{\mathbf{f}}^\top \partial_2^\top \partial_2 \boldsymbol{\epsilon} + \boldsymbol{\epsilon}^\top \partial_2^\top \partial_2 \boldsymbol{\epsilon} \\ &\geq \tilde{\mathbf{f}}^\top \partial_2^\top \partial_2 \tilde{\mathbf{f}} - 2 \left\| \partial_2^\top \partial_2 \tilde{\mathbf{f}} \right\|_2 \|\boldsymbol{\epsilon}\|_2 \\ &\geq \tilde{\mathbf{f}}^\top \partial_2^\top \partial_2 \tilde{\mathbf{f}} - 4 \left\| \partial_2^\top \partial_2 \right\|_2 \|\boldsymbol{\epsilon}\|_2 \\ &\geq \tilde{\mathbf{f}}^\top \partial_2^\top \partial_2 \tilde{\mathbf{f}} - 48\sqrt{22d}C, \end{aligned} \quad \text{by Lemma 4.5.8}$$

Note that

$$\tilde{\mathbf{f}}^\top \partial_2^\top \partial_2 \tilde{\mathbf{f}} = 3 \|\mathbf{A}\boldsymbol{\alpha}\|_2^2,$$

where $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_n)^\top$. Write $\boldsymbol{\alpha} = \boldsymbol{\alpha}_\perp + \boldsymbol{\alpha}_0$, where $\boldsymbol{\alpha}_\perp$ is orthogonal to the null space of \mathbf{A} and $\boldsymbol{\alpha}_0$ is in the null space of \mathbf{A} . Then,

$$\tilde{\mathbf{f}}^\top \partial_2^\top \partial_2 \tilde{\mathbf{f}} \geq 3\lambda_{\min}(\mathbf{A}^\top \mathbf{A}) \|\boldsymbol{\alpha}_\perp\|_2^2. \quad (4.5)$$

It remains to lower bound $\|\boldsymbol{\alpha}_\perp\|_2$. We can write $\tilde{\mathbf{f}} = \mathbf{H}\boldsymbol{\alpha}_\perp + \mathbf{H}\boldsymbol{\alpha}_0$. By Lemma 4.5.9, $\mathbf{H}\boldsymbol{\alpha}_0 \in \text{Ker}(\partial_2)$ and $\mathbf{H}\boldsymbol{\alpha}_\perp \perp \text{Ker}(\partial_2)$. On the other hand, $\tilde{\mathbf{f}} = \mathbf{f} - \boldsymbol{\epsilon}$ and $\mathbf{f} \perp \text{Ker}(\partial_2)$. We know

$$\|\mathbf{H}\boldsymbol{\alpha}_\perp\|_2 \geq \|\mathbf{f}\|_2 - \|\boldsymbol{\epsilon}\|_2,$$

and thus

$$\|\boldsymbol{\alpha}_\perp\|_2 \geq \frac{1}{\sqrt{t}} \|\mathbf{H}\boldsymbol{\alpha}_\perp\|_2 \geq \frac{1}{\sqrt{t}} - C > \frac{1}{\sqrt{22d}}.$$

Together with Eq. (4.5),

$$\tilde{\mathbf{f}}^\top \partial_2^\top \partial_2 \tilde{\mathbf{f}} \geq \frac{3\lambda_{\min}(\mathbf{A}^\top \mathbf{A})}{22d}.$$

This completes the proof. □

Proof of Lemma 4.5.6. The proof follows by combining Lemma 4.5.8 and 4.5.10. □

4.6 Reducing \mathcal{DA} to \mathcal{B}_2 in General Case

In this section, we show how to reduce LEA $(\mathbf{A}, \mathbf{b}, \epsilon^{DA})$ to instances LEA over \mathcal{B}_2 without requiring the assumption that \mathbf{b} is in the image of \mathbf{A} as we did in earlier sections. In this more general case, LEA $(\mathbf{A}, \mathbf{b}, \epsilon^{DA})$ aims to compute an *approximate* solution to $\arg \min_{\mathbf{x}} \|\mathbf{Ax} - \mathbf{b}\|_2 = \arg \min_{\mathbf{x}} \|\mathbf{Ax} - \Pi_{\mathbf{A}} \mathbf{b}\|_2$.

We remark that our reduction for this general case does not work for LE (\mathbf{A}, \mathbf{b}) , which computes an *exact* solution to $\min_{\mathbf{x}} \|\mathbf{Ax} - \mathbf{b}\|_2$. Approximate linear equation solvers, however, can be more interesting in practice, since numerical errors occur unavoidably during data collection and computation, and approximate solutions may be computed much faster than exact solutions.

4.6.1 Warm-Up: Reweighting Infeasible Equations to Preserve Solutions

There is a crucial difference between the reductions we use for LE (\mathbf{A}, \mathbf{b}) and LEA $(\mathbf{A}, \mathbf{b}, \epsilon^{DA})$ when $\mathbf{b} \in \text{Im}(\mathbf{A})$ and in the general case when we may have $\mathbf{b} \notin \text{Im}(\mathbf{A})$.

To understand this, consider the following feasible system of just two linear equations in two variables x and y .

$$\begin{aligned} x - y &= 1, \\ -x + y &= -1. \end{aligned}$$

A feasible solution is $x = 1$ and $y = 0$. Now, suppose we add another linear equation that is satisfied by all solutions to the previous equations, for example, we can simply repeat the first constraint $x - y = 1$. It remains true that the existing solutions are feasible.

Now, in contrast, consider an infeasible system of two linear equations in variables x and y .

$$\begin{aligned} x - y &= 1, \\ -x + y &= 0. \end{aligned}$$

This linear equation is *not* feasible. In particular, we can consider the associated minimization problem $\arg \min_{\mathbf{x}} \|\mathbf{Ax} - \mathbf{b}\|_2$ with $\mathbf{A} = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$ and $\mathbf{b} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$, for which one minimizing solution is $\mathbf{x} = \begin{bmatrix} 1/2 \\ 0 \end{bmatrix}$. Notice that if we add a row which simply repeats the first constraint, i.e. $x - y = 1$, then the resulting minimization problem has

$$\mathbf{A} = \begin{bmatrix} 1 & -1 \\ -1 & 1 \\ 1 & -1 \end{bmatrix} \text{ and } \mathbf{b} = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix},$$

and now $\mathbf{x} = \begin{bmatrix} 1/2 \\ 0 \end{bmatrix}$ is no longer a minimizing solution to $\arg \min_{\mathbf{x}} \|\mathbf{Ax} - \mathbf{b}\|_2$. In particular, $\mathbf{x} = \begin{bmatrix} 1/2 \\ 0 \end{bmatrix}$ achieves a value of $\sqrt{3}/2$, while $\mathbf{x}' = \begin{bmatrix} 2/3 \\ 0 \end{bmatrix}$ achieves the smaller value $\sqrt{6}/3$.

However, if we reweigh the first and last row of our system of inequalities by a factor $1/\sqrt{2}$, so that

$$\mathbf{A} = \begin{bmatrix} 1/\sqrt{2} & -1/\sqrt{2} \\ -1 & 1 \\ 1/\sqrt{2} & -1/\sqrt{2} \end{bmatrix} \quad \text{and} \quad \mathbf{b} = \begin{bmatrix} 1/\sqrt{2} \\ 0 \\ 1/\sqrt{2} \end{bmatrix},$$

then we in fact maintain that the original solution $\mathbf{x} = \begin{bmatrix} 1/2 \\ 0 \end{bmatrix}$ stays optimal. Thus, when we modify an infeasible system of linear equations, we have to be very careful about the weight we assign to different constraints if we are to (approximately) preserve the correspondence between the optimal solutions of the original and the final problems. In the following section, we describe a reweighting scheme which can be combined with our existing reductions to ensure that optimal solutions are (approximately) preserved, even when the original problem is infeasible.

4.6.2 Reduction Algorithm

Our reduction is almost the same as the algorithm REDUCEDATOB_2 , except that we assign edge weights for the constructed 2-complex.

Suppose we are given an instance LEA $(\mathbf{A}, \mathbf{b}, \epsilon^{DA})$ where \mathbf{A} is a $d \times n$ matrix in \mathcal{DA} . Let $(\partial_2, \gamma, \Delta^c) \leftarrow \text{REDUCEDATOB}_2(\mathbf{A}, \mathbf{b})$. Let \mathcal{K} be the 2-complex whose boundary operator is ∂_2 . We compute the edge weights of \mathcal{K} as follows.

1. For each boundary triangle $\Delta_{q,i,*}^1$ where $q \in [d], i \in [n]$ and $* \in \{1, 2\}$, we find a minimal triangle path $\mathcal{P}_{q,i,*}^1$ in \mathcal{K}_i from the central triangle Δ_i to $\Delta_{q,i,*}^1$. Let $E_{q,i,*}^1$ be the set of edges shared by neighboring triangles in $\mathcal{P}_{q,i,*}^1$.⁴
2. For each interior edge e and equation q , let $k_{q,e}$ be the number of triangle paths indexed by equation q that contain e . For each equation q , let l_q be the sum of the lengths of all the triangle paths indexed by equation q . Then, we set the weight for edge e to be

$$w_e = \begin{cases} 1, & \text{if } e \text{ is a boundary edge} \\ \alpha \sum_{q \in [d]} k_{q,e} l_q, & \text{if } e \text{ is an interior edge} \end{cases}, \quad (4.6)$$

where

$$\alpha = \frac{2}{(\epsilon^{DA})^2}.$$

Let \mathbf{W} be a diagonal matrix whose diagonals are the edge weights. We return $(\mathbf{W}, \partial_2, \gamma, \Delta^c)$.

Note that an edge in \mathcal{K} may have weight 0. If we want to make all the edge weights positive, we can impose polynomially small edge weights for these 0-weight edges, which only affects the error propagation and condition number up to polynomial factors.

We refer to the above algorithm as REDUCEREGDATOB_2 . We use the algorithm $\text{MAPSOLNB}_2\text{TO}\mathcal{DA}$ to map a solution to LEA $(\mathbf{W}^{1/2}\partial_2, \mathbf{W}^{1/2}\gamma, \epsilon^{B_2})$ back to a solution to LEA $(\mathbf{A}, \mathbf{b}, \epsilon^{DA})$, where we choose

$$\epsilon^{B_2} \leq \frac{\epsilon^{DA}}{\sqrt{3 \left(1 + \frac{1}{\alpha} \|\mathbf{b}\|_2^2 \text{nnz}(\mathbf{A}) \|\mathbf{A}\|_{\max}^2\right)}}.$$

⁴Note that any two neighboring triangles in $\mathcal{P}_{q,i,*}^1$ share exactly one edge. So, the length of $\mathcal{P}_{q,i,*}^1$ equals $|E_{q,i,*}^1|$.

Computing the Edge Weights in Linear Time. Since the weight of an interior edge e only depends on \mathcal{K}_i that contains edge e , we will compute the edge weights for the edges in each \mathcal{K}_i separately.

For each $i \in [n]$, consider a graph G_i whose vertices are the triangles in \mathcal{K}_i and the two vertices are adjacent if and only if the two corresponding triangles share an edge in \mathcal{K}_i . We run the breadth-first-search (BFS) to construct a shortest-path tree T_i of G_i rooted at Δ_i . For each boundary triangle $\Delta_{q,i,*}^1$, we choose the triangle path $\mathcal{P}_{q,i,*}^1$ to be the triangle path induced by T_i from the root Δ_i to the node corresponding to $\Delta_{q,i,*}^1$, whose length is the height of the node $\Delta_{q,i,*}^1$. Since we are only interested in these triangle paths, for simplicity, we remove a node from T_i if it is not a boundary triangle $\Delta_{q,i,*}^1$ for some q and none of its descendant is a boundary triangle $\Delta_{q,i,*}^1$ for some q . After this operation, every leaf of T_i is a boundary triangle $\Delta_{q,i,*}^1$ for some q .

Our goal is to count $\sum_{q \in [d]} k_{q,e} l_q$ for each edge in the tree T_i . We observe that an edge e appears in a triangle path from the root Δ_i to a boundary triangle $\Delta_{q,i,*}^1$ if and only if $\Delta_{q,i,*}^1$ is a descendant of an end-node of e with the higher height. First, we BFS traverse the tree T_i to store the height of each boundary triangle node. Then, we traverse the tree T_i to store l_q at each boundary triangle node $\Delta_{q,i,*}^1$ for each $q \in [d]$. Next, we traverse the tree T_i from the leaf nodes to the root to count $\sum_{q \in [d]} k_{q,e} l_q$ for each edge e . The total runtime is linear in the number of triangles in \mathcal{K}_j .

4.6.3 Relation Between Exact Solutions

In this section, we show that by reweighting all the edges in \mathcal{K} , an exact solution to $(\mathbf{W}^{1/2} \partial_2, \mathbf{W}^{1/2} \gamma)$ is close to an exact solution to (\mathbf{A}, \mathbf{b}) , stated in Claim 4.6.1. Claim 4.6.1 plays a key role in analyzing approximate solutions and condition numbers.

Claim 4.6.1. *For any $\mathbf{f} \in \mathbb{R}^t$ and $\mathbf{x} \leftarrow \text{MAPSOLNB}_2\text{TO}\mathcal{DA}(\mathbf{A}, \mathbf{b}, \mathbf{f}, \Delta^c)$,*

$$\frac{\alpha}{\alpha + 1} \|\mathbf{Ax} - \mathbf{b}\|_2^2 \leq \left\| \mathbf{W}^{1/2} \partial_2 \mathbf{f} - \mathbf{W}^{1/2} \gamma \right\|_2^2.$$

Note that if \mathbf{f} satisfies $\partial_2 \mathbf{f} = \gamma$, then $\|\mathbf{Ax} - \mathbf{b}\|_2^2 = \|\partial_2 \mathbf{f} - \gamma\|_2^2$. Claim 4.6.1 states that by our choice of weights in \mathbf{W} , we can generalize this relation to more general \mathbf{f} .

Proof. For the convenience of analysis, we construct an auxiliary boundary matrix $\hat{\partial}_2$. For each interior edge e , let k_e be the number of all the triangle paths containing e , and we split the row $\partial_2(e)$ into k_e copies. For each copy related to equation q , we assign its weight to be $\alpha \cdot l_q$. Let $\hat{\mathbf{W}}$ be the auxiliary weight matrix, and let $\hat{\gamma}$ be the auxiliary demand vector. We can check

$$\left\| \mathbf{W}^{1/2} \partial_2 \mathbf{f} - \mathbf{W}^{1/2} \gamma \right\|_2^2 = \left\| \hat{\mathbf{W}}^{1/2} \hat{\partial}_2 \mathbf{f} - \hat{\mathbf{W}}^{1/2} \hat{\gamma} \right\|_2^2. \quad (4.7)$$

We reorder rows of the matrices $\hat{\mathbf{W}}, \hat{\partial}_2$ and the vector $\hat{\gamma}$ in the following way. For each $q \in [d]$, let E_q be a (multi)set that is the union of the shared edges in the triangle paths indexed by q ⁵. Then, we reorder rows of $\hat{\mathbf{W}}, \hat{\partial}_2, \hat{\gamma}$ by grouping those corresponding

⁵If the q th equation is a difference equation, then every edge appears in E_q at most once; if the q th equation is an average equation, then some edges may appear twice.

to the edges in $E_q \cup \{\alpha_q^1\}$:

$$\hat{\partial}_2 = \begin{bmatrix} \mathbf{G}_1 \\ \vdots \\ \mathbf{G}_d \end{bmatrix}, \quad \text{where } \mathbf{G}_q = \begin{bmatrix} \mathbf{B}_q \\ \mathbf{M}_q \end{bmatrix}. \quad (4.8)$$

where \mathbf{B}_q corresponds to the boundary edge α_q^1 and \mathbf{M}_q is the submatrix corresponding to all the interior edges in E_q . Correspondingly, we write

$$\hat{\mathbf{W}} = \begin{bmatrix} \hat{\mathbf{W}}_1 & & \\ & \ddots & \\ & & \hat{\mathbf{W}}_d \end{bmatrix} \quad \text{and} \quad \hat{\gamma} = \begin{bmatrix} \hat{\gamma}_1 \\ \vdots \\ \hat{\gamma}_d \end{bmatrix}, \quad \text{where } \hat{\gamma}_q = \begin{bmatrix} \mathbf{b}(q) \\ \mathbf{0}_{|E_q|} \end{bmatrix}. \quad (4.9)$$

Let

$$\hat{w}_{q,e} \stackrel{\text{def}}{=} \hat{\mathbf{W}}_q(e, e) = \begin{cases} 1, & \text{if } e \text{ is a boundary edge} \\ \alpha \cdot l_q, & \text{if } e \text{ is an interior edge} \end{cases}.$$

For each $q \in [d]$, we define

$$\epsilon_q = \mathbf{A}(q)\mathbf{x} - \mathbf{b}(q), \quad \xi_q = \mathbf{B}_q\mathbf{f} - \mathbf{b}(q), \quad \delta_e = \mathbf{M}_q(e)\mathbf{f}.$$

We can check

$$\epsilon_q = \mathbf{1}^\top (\mathbf{G}_q\mathbf{f} - \hat{\gamma}_q) = \xi_q + \sum_{e \in E_q} \delta_e.$$

By the Cauchy-Schwarz inequality,

$$\epsilon_q^2 = \left(\xi_q + \sum_{e \in E_q} \delta_e \right)^2 \leq \left(1 + \sum_{e \in E_q} \frac{1}{\hat{w}_{q,e}} \right) \left(\xi_q^2 + \sum_{e \in E_q} \hat{w}_{q,e} \delta_e^2 \right) = \left(1 + \frac{1}{\alpha} \right) \left(\xi_q^2 + \sum_{e \in E_q} \hat{w}_{q,e} \delta_e^2 \right).$$

That is,

$$\|\mathbf{A}(q)\mathbf{x} - \mathbf{b}(q)\|_2^2 \leq \left(1 + \frac{1}{\alpha} \right) \left\| \hat{\mathbf{W}}_q^{1/2} \mathbf{G}_q \mathbf{f} - \hat{\mathbf{W}}_q^{1/2} \hat{\gamma}_q \right\|_2^2.$$

Summing over all d equations, we get

$$\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 \leq \left(1 + \frac{1}{\alpha} \right) \left\| \hat{\mathbf{W}}^{1/2} \hat{\partial}_2 \mathbf{f} - \hat{\mathbf{W}}^{1/2} \hat{\gamma} \right\|_2^2.$$

By Eq. (4.7), we get the inequality in the statement. \square

Claim 4.6.1 implies the following relation between the exact solutions to (\mathbf{A}, \mathbf{b}) and those to $(\mathbf{W}^{1/2} \partial_2, \mathbf{W}^{1/2} \gamma)$.

Lemma 4.6.2. *Given any $d \times n$ matrix $\mathbf{A} \in \mathcal{DA}$ and vector $\mathbf{b} \in \mathbb{R}^d$, let $(\mathbf{W}, \partial_2, \gamma, \Delta^c) \leftarrow \text{REDUCEREGDATO}\mathcal{B}_2(\mathbf{A}, \mathbf{b}, \epsilon^{DA})$. Then,*

$$\frac{\alpha}{\alpha + 1} \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 \leq \min_{\mathbf{f}} \left\| \mathbf{W}^{1/2} \partial_2 \mathbf{f} - \mathbf{W}^{1/2} \gamma \right\|_2^2 \leq \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2.$$

Remark that Lemma 4.6.2 can also be stated as

$$\frac{\alpha}{\alpha+1} \|(\mathbf{I} - \Pi_{\mathbf{A}})\mathbf{b}\|_2^2 \leq \left\| (\mathbf{I} - \Pi_{\mathbf{W}^{1/2}\partial_2}) \mathbf{W}^{1/2}\gamma \right\|_2^2 \leq \|(\mathbf{I} - \Pi_{\mathbf{A}})\mathbf{b}\|_2^2.$$

We do not prove equalities. But as $\alpha \rightarrow \infty$, the leftmost hand side and the rightmost hand side are equal.

Proof. Let $\mathbf{f}^* \in \arg \min_{\mathbf{f}} \left\| \mathbf{W}^{1/2}\partial_2\mathbf{f} - \mathbf{W}^{1/2}\gamma \right\|_2$ and $\mathbf{x}^* \in \arg \min_{\mathbf{x}} \|\mathbf{Ax} - \mathbf{b}\|_2$. Let $\mathbf{x} \leftarrow \text{MAPSOLNB}_2\text{TO}\mathcal{DA}(\mathbf{A}, \mathbf{b}, \mathbf{W}^{1/2}\partial_2, \mathbf{f}^*)$. By Claim 4.6.1

$$\frac{\alpha}{\alpha+1} \|\mathbf{Ax}^* - \mathbf{b}\|_2^2 \leq \frac{\alpha}{\alpha+1} \|\mathbf{Ax} - \mathbf{b}\|_2^2 \leq \left\| \mathbf{W}^{1/2}\partial_2\mathbf{f}^* - \mathbf{W}^{1/2}\gamma \right\|_2^2,$$

which is the first inequality in the lemma statement. Let

$$\mathbf{f} = \begin{bmatrix} \mathbf{x}^*(1)\mathbf{1}_{t_1} \\ \vdots \\ \mathbf{x}^*(n)\mathbf{1}_{t_n} \end{bmatrix}.$$

Then,

$$\left\| \mathbf{W}^{1/2}\partial_2\mathbf{f}^* - \mathbf{W}^{1/2}\gamma \right\|_2^2 \leq \left\| \mathbf{W}^{1/2}\partial_2\mathbf{f} - \mathbf{W}^{1/2}\gamma \right\|_2^2 = \|\mathbf{Ax}^* - \mathbf{b}\|_2^2,$$

which is the second inequality in the lemma statement. \square

4.6.4 Relation Between Approximate Solutions

Linear equation problem LE (\mathbf{A}, \mathbf{b}) aims to find a vector \mathbf{x} such that $\mathbf{Ax} = \Pi_{\mathbf{A}}\mathbf{b}$. In our setting, both \mathbf{A} and \mathbf{b} have integer entries. We will need the following claim to lower bound $\mathbf{Ax} = \Pi_{\mathbf{A}}\mathbf{b}$.

Claim 4.6.3. *Let $\mathbf{A} \in \mathbb{R}^{d \times n}$ and $\mathbf{b} \in \mathbb{R}^d$ such that $\|\mathbf{A}^\top \mathbf{b}\|_2 \geq 1$. Then,*

$$\|\Pi_{\mathbf{A}}\mathbf{b}\|_2^2 \geq \frac{1}{\lambda_{\max}(\mathbf{A}^\top \mathbf{A})}.$$

Proof. Note that

$$\|\Pi_{\mathbf{A}}\mathbf{b}\|_2^2 = \mathbf{b}^\top \mathbf{A}(\mathbf{A}^\top \mathbf{A})^\dagger \mathbf{A}^\top \mathbf{b} \geq \lambda_{\min}((\mathbf{A}^\top \mathbf{A})^\dagger) \|\mathbf{A}^\top \mathbf{b}\|_2^2 \geq \frac{1}{\lambda_{\max}(\mathbf{A}^\top \mathbf{A})}.$$

\square

Claim 4.6.3 and Lemma 4.6.2 enable us to relate $\|\Pi_{\mathbf{A}}\mathbf{b}\|_2$ with $\left\| \Pi_{\mathbf{W}^{1/2}\partial_2} \mathbf{W}^{1/2}\gamma \right\|_2$.

Claim 4.6.4.

$$\left\| \Pi_{\mathbf{W}^{1/2}\partial_2} \mathbf{W}^{1/2}\gamma \right\|_2^2 \leq \left(1 + \frac{1}{\alpha} \lambda_{\max}(\mathbf{A}^\top \mathbf{A}) \|\mathbf{b}\|_2^2 \right) \|\Pi_{\mathbf{A}}\mathbf{b}\|_2^2.$$

Proof. We apply Lemma 4.6.2.

$$\begin{aligned}
\|\Pi_A \mathbf{b}\|_2^2 &= \|\mathbf{b}\|_2^2 - \|(\mathbf{I} - \Pi_A) \mathbf{b}\|_2^2 \\
&\geq \left\| \mathbf{W}^{1/2} \gamma \right\|_2^2 - \frac{\alpha + 1}{\alpha} \left\| (\mathbf{I} - \Pi_{\mathbf{W}^{1/2} \partial_2}) \mathbf{W}^{1/2} \gamma \right\|_2^2 \\
&= \left\| \Pi_{\mathbf{W}^{1/2} \partial_2} \mathbf{W}^{1/2} \gamma \right\|_2^2 - \frac{1}{\alpha} \left\| (\mathbf{I} - \Pi_{\mathbf{W}^{1/2} \partial_2}) \mathbf{W}^{1/2} \gamma \right\|_2^2 \\
&\geq \left\| \Pi_{\mathbf{W}^{1/2} \partial_2} \mathbf{W}^{1/2} \gamma \right\|_2^2 - \frac{1}{\alpha} \|(\mathbf{I} - \Pi_A) \mathbf{b}\|_2^2.
\end{aligned}$$

Since \mathbf{A}, \mathbf{b} have integer entries, by Claim 4.6.3,

$$\|(\mathbf{I} - \Pi_A) \mathbf{b}\|_2^2 \leq \|\mathbf{b}\|_2^2 \leq \|\mathbf{b}\|_2^2 \cdot \lambda_{\max}(\mathbf{A}^\top \mathbf{A}) \|\Pi_A \mathbf{b}\|_2^2.$$

Thus,

$$\left\| \Pi_{\mathbf{W}^{1/2} \partial_2} \mathbf{W}^{1/2} \gamma \right\|_2^2 \leq \left(1 + \frac{1}{\alpha} \lambda_{\max}(\mathbf{A}^\top \mathbf{A}) \|\mathbf{b}\|_2^2 \right) \|\Pi_A \mathbf{b}\|_2^2.$$

□

Now, we apply Claim 4.6.1 and Lemma 4.6.4 to prove a relation between approximate solutions.

Lemma 4.6.5 (Approximate Solvers in General Case). *Given a difference-average instance LEA $(\mathbf{A}, \mathbf{b}, \epsilon^{DA})$, let $(\mathbf{W}, \partial_2, \gamma, \Delta^c) \leftarrow \text{REDUCEREGDATO}\mathcal{B}_2(\mathbf{A}, \mathbf{b}, \alpha)$. Suppose \mathbf{f} is a solution to LEA $(\mathbf{W}^{1/2} \partial_2, \mathbf{W}^{1/2} \gamma, \epsilon^{B_2})$, where $\epsilon^{B_2} \leq \frac{\epsilon^{DA}}{10}$. Let $\mathbf{x} \leftarrow \text{MAPSOLN}\mathcal{B}_2\text{TO}\mathcal{DA}(\mathbf{A}, \mathbf{b}, \mathbf{f}, \Delta^c)$. Then, \mathbf{x} is a solution to LEA $(\mathbf{A}, \mathbf{b}, \epsilon^{DA})$.*

Proof. By Claim 4.6.1, we have

$$\|\mathbf{Ax} - \mathbf{b}\|_2^2 \leq \frac{\alpha + 1}{\alpha} \left\| \mathbf{W}^{1/2} \partial_2 \mathbf{f} - \mathbf{W}^{1/2} \gamma \right\|_2^2. \quad (4.10)$$

Also note that

$$\begin{aligned}
\|\mathbf{Ax} - \mathbf{b}\|_2^2 &= \|\mathbf{Ax} - \Pi_A \mathbf{b}\|_2^2 + \|(\mathbf{I} - \Pi_A) \mathbf{b}\|_2^2, \\
\left\| \mathbf{W}^{1/2} \partial_2 \mathbf{f} - \mathbf{W}^{1/2} \gamma \right\|_2^2 &= \left\| \mathbf{W}^{1/2} \partial_2 \mathbf{f} - \Pi_{\mathbf{W}^{1/2} \partial_2} \mathbf{W}^{1/2} \gamma \right\|_2^2 + \left\| (\mathbf{I} - \Pi_{\mathbf{W}^{1/2} \partial_2}) \mathbf{W}^{1/2} \gamma \right\|_2^2.
\end{aligned}$$

Plugging these into Eq. (4.10) and apply Lemma 4.6.2,

$$\begin{aligned}
\|\mathbf{Ax} - \Pi_A \mathbf{b}\|_2^2 &\leq \frac{\alpha + 1}{\alpha} \left(\left\| \mathbf{W}^{1/2} \partial_2 \mathbf{f} - \Pi_{\mathbf{W}^{1/2} \partial_2} \mathbf{W}^{1/2} \gamma \right\|_2^2 + \left\| (\mathbf{I} - \Pi_{\mathbf{W}^{1/2} \partial_2}) \mathbf{W}^{1/2} \gamma \right\|_2^2 \right) \\
&\quad - \|(\mathbf{I} - \Pi_A) \mathbf{b}\|_2^2 \\
&\leq \frac{\alpha + 1}{\alpha} \left\| \mathbf{W}^{1/2} \partial_2 \mathbf{f} - \Pi_{\mathbf{W}^{1/2} \partial_2} \mathbf{W}^{1/2} \gamma \right\|_2^2 + \frac{1}{\alpha} \|(\mathbf{I} - \Pi_A) \mathbf{b}\|_2^2.
\end{aligned}$$

By the fact that \mathbf{f} is a solution to LEA $(\mathbf{W}^{1/2} \partial_2, \mathbf{W}^{1/2} \gamma, \epsilon^{B_2})$ and by Claim 4.6.4,

$$\left\| \mathbf{W}^{1/2} \partial_2 \mathbf{f} - \Pi_{\mathbf{W}^{1/2} \partial_2} \mathbf{W}^{1/2} \gamma \right\|_2^2 \leq (\epsilon^{B_2})^2 \left\| \Pi_{\mathbf{W}^{1/2} \partial_2} \mathbf{W}^{1/2} \gamma \right\|_2^2 \leq \frac{(\epsilon^{DA})^2}{3} \|\Pi_A \mathbf{b}\|_2^2.$$

In addition,

$$\frac{1}{\alpha} \|(\mathbf{I} - \Pi_A) \mathbf{b}\|_2^2 \leq \frac{1}{\alpha} \|\mathbf{b}\|_2^2 \leq \frac{(\epsilon^{DA})^2}{2} \|\Pi_A \mathbf{b}\|_2^2.$$

Thus,

$$\|\mathbf{Ax} - \Pi_A \mathbf{b}\|_2^2 \leq (\epsilon^{DA})^2 \|\Pi_A \mathbf{b}\|_2^2,$$

that is, \mathbf{x} is a solution to LEA $(\mathbf{A}, \mathbf{b}, \epsilon^{DA})$. □

4.6.5 Bounding the Condition Number of the New Matrix

We will upper bound the maximum eigenvalue of $\partial_2^\top \mathbf{W} \partial_2$ and lower bound its minimum non-zero eigenvalue. The proofs are similar to those in Section 4.5.6, which bound the eigenvalues of $\partial_2^\top \partial_2$.

Lemma 4.6.6. $\kappa(\partial_2^\top \mathbf{W} \partial_2) = O((\epsilon^{DA})^{-2} \text{nnz}(\mathbf{A})^{15/2} \kappa(\mathbf{A}^\top \mathbf{A}))$.

Proof. The proof follows the same proof line in Section 4.5.6 for $\mathbf{W} = \mathbf{I}$. Here, we lose a factor w_{\max} when we upper bound $\lambda_{\max}(\partial_2^\top \mathbf{W} \partial_2)$, and we lose a factor $\frac{w_{\max}}{w_{\min}}$ when we lower bound $\lambda_{\min}(\partial_2^\top \mathbf{W} \partial_2)$, where w_{\max} is the maximum diagonal in \mathbf{W} and w_{\min} is the minimum non-zero diagonal in \mathbf{W} . By our setting, $w_{\max} = O(\alpha \text{nnz}(\mathbf{A})^2)$ and $w_{\min} = \alpha$, where $\alpha = 2(\epsilon^{DA})^{-2}$. \square

Chapter 5

Hardness Results for More Structured Problems

This chapter extends the reduction from Chapter 4 to derive additional hardness results.

We begin by presenting the reduction algorithm from general linear equations to difference-average equations for completeness, and prove Theorem 1.4.3: LEA over \mathcal{DA} is sparse-linear-equation complete. Next, using matrices in \mathcal{DA} as our starting point, we introduce an even simpler form called *One-Or-Three* matrices, in which each row has either one or three non-zero entries (all of which are ± 1). We denote this matrix class by \mathcal{OT} . Building on this simplification, we extend the reduction of linear equations to the reduction of linear programs, leading to what we call *1-or-3* linear programs. We believe these programs serve as a strong foundation for further reductions to more structured LP problems. We remark that this result is unpublished.

For simplicity, throughout this chapter we focus on the reduction algorithms. Analyses such as error propagation and condition number are omitted.

5.1 Reducing General LE to Difference-Average LE

In this section, we prove Theorem 4.3.4: LEA over \mathcal{DA} is sparse-linear-equation complete.

We show a nearly-linear time reduction from linear equation problems over matrices in \mathcal{G} to linear equation problems over matrices in \mathcal{DA} . This reduction is implicit in Section 4 of [KZ17], as an intermediate step to reduce linear equation problems over matrices in \mathcal{G} to matrices in a slight generalization of Laplacians. We explicitly separate the reduction step and simplify the proofs in [KZ17], which might be of independent interest.

Recall that a matrix in \mathcal{G} has polynomially bounded integer entries and polynomially bounded condition numbers, and they do not have all-0 rows or all-0 columns; a matrix $\mathbf{A} \in \mathcal{DA}$ only has two types of rows such that if we multiply \mathbf{A} to a vector \mathbf{x} , then the entries of \mathbf{Ax} are in the form of either $\mathbf{x}(i) - \mathbf{x}(j)$ or $\mathbf{x}(i) + \mathbf{x}(j) - 2\mathbf{x}(k)$.

The reduction in [KZ17] has three steps:

1. Reduce linear equation problems over matrices in \mathcal{G} to matrices in \mathcal{G}_z , a subset of \mathcal{G} containing matrices with row sum 0. Given an instance (\mathbf{G}, \mathbf{b}) where $\mathbf{G} \in \mathcal{G}$, we construct a new instance $(\mathbf{G}', \mathbf{b})$ where $\mathbf{G}' = [\mathbf{G} \quad -\mathbf{G}\mathbf{1}] \in \mathcal{G}_z$.
2. Reduce linear equation problems over matrices in \mathcal{G}_z to matrices in $\mathcal{G}_{z,2}$, a subset of \mathcal{G}_z containing matrices such that the sum of positive entries in each row is a

power of 2. Given an instance $(\mathbf{G}', \mathbf{b})$ where $\mathbf{G}' \in \mathcal{G}_z$, we construct a new instance $(\mathbf{G}'', \mathbf{b}'')$ where $\mathbf{G}'' = \begin{bmatrix} \mathbf{G}' & \mathbf{g} & -\mathbf{g} \\ \mathbf{0} & 1 & -1 \end{bmatrix} \in \mathcal{G}_{z,2}$ and $\mathbf{b}'' = \begin{bmatrix} \mathbf{b} \\ 0 \end{bmatrix}$.

3. Reduce linear equation problems over matrices in $\mathcal{G}_{z,2}$ to matrices in \mathcal{DA} .

The first two steps are proved in Section 7 of [KZ17], by standard tricks. We will focus on the third step.

In the rest of this section, to be consistent with the notations used in [KZ17], we use subscripts to denote entries of a matrix or a vector. Let \mathbf{A}_i denote the i th row of a matrix \mathbf{A} , and \mathbf{A}_{ij} denote the (i, j) th entry of \mathbf{A} . Let \mathbf{x}_i denote the i th entry of a vector \mathbf{x} , and $\mathbf{x}_{i:j}$ denote the subvector of entries $\mathbf{x}_i, \mathbf{x}_{i+1}, \dots, \mathbf{x}_j$. Moreover, we let $\|\mathbf{A}\|_{\max}$ be the maximum magnitude of the entries of \mathbf{A} .

Given an instance of linear equation problems over $\mathcal{G}_{z,2}$, our construction of an instance over \mathcal{DA} does not depend on the error parameter. Therefore, we will describe our construction without the error parameters; then we will explain how to set the error parameter for the construction instance over \mathcal{DA} .

Let $(\mathbf{A}, \mathbf{c}^A)$ be an instance over $\mathcal{G}_{z,2}$. The idea is to write each equation in $(\mathbf{A}, \mathbf{c}^A)$ as a set of difference equations and average equations, via bitwise decomposition. Same as [KZ17], we first explain the idea by an example:

$$3\mathbf{x}_1 + 5\mathbf{x}_2 - \mathbf{x}_3 - 7\mathbf{x}_4 = 1. \quad (5.1)$$

We will manipulate the variables with positive coefficients and the variables with negative coefficients separately. Let us take the positive coefficients as an example. We pair all the variables with the odd positive coefficients and replace each pair with a new variable via an average equation. In this example, we pair $\mathbf{x}_1, \mathbf{x}_2$; we then create a new variable \mathbf{x}_{12} and a new average equation

$$\mathbf{x}_1 + \mathbf{x}_2 - 2\mathbf{x}_{12} = 0.$$

Plugging this into Eq. (5.1), we get

$$2\mathbf{x}_1 + 4\mathbf{x}_2 + 2\mathbf{x}_{12} - \mathbf{x}_3 - 7\mathbf{x}_4 = 1.$$

We pull out a factor 2:

$$2(\mathbf{x}_1 + 2\mathbf{x}_2 + \mathbf{x}_{12}) - \mathbf{x}_3 - 7\mathbf{x}_4 = 1.$$

We repeat the pair-and-replace process to pair $\mathbf{x}_1, \mathbf{x}_{12}$ with a new variables $\mathbf{x}_{1,12}$ and a new average equation $\mathbf{x}_1 + \mathbf{x}_{12} - 2\mathbf{x}_{1,12} = 0$:

$$2(2\mathbf{x}_{1,12} + 2\mathbf{x}_2) - \mathbf{x}_3 - 7\mathbf{x}_4 = 1.$$

Pull out a factor 2:

$$4(\mathbf{x}_{1,12} + \mathbf{x}_2) - \mathbf{x}_3 - 7\mathbf{x}_4 = 1.$$

Repeat pair-and-replace with $\mathbf{x}_{1,12} + \mathbf{x}_2 - 2\mathbf{x}_{1,12,2} = 0$:

$$8\mathbf{x}_{1,12,2} - \mathbf{x}_3 - 7\mathbf{x}_4 = 1.$$

Similarly, we can use a sequence of average equations for the variables with odd coefficients, and the above equation becomes:

$$8\mathbf{x}_{1,12,2} - 8\mathbf{x}_{34} = 1.$$

where \mathbf{x}_{34} is a new variable. The final equation is a difference equation.

The above reduction relies on the fact that the sum of all the coefficients is 0 and the sum of all the positive coefficients is a power of 2.

Pseudo-code of the reduction from linear equation problems over $\mathcal{G}_{z,2}$ to \mathcal{DA} is shown in Algorithm 1.

Lemma 5.1.1. *Let $\mathbf{A} \in \mathcal{G}_{z,2}$, and let \mathbf{B} be returned by running Algorithm REDUCE $\mathcal{G}_{z,2}$ TO \mathcal{DA} on \mathbf{A} . Then,*

- $\text{nnz}(\mathbf{B}) = O(\text{nnz}(\mathbf{A}) \log \|\mathbf{A}\|_{\max})$;
- *both the dimensions of \mathbf{B} are $O(\text{nnz}(\mathbf{A}) \log \|\mathbf{A}\|_{\max})$.*

Proof. The second statement is implied by the first one.

To prove the first statement, we focus on Algorithm REDUCE $\mathcal{G}_{z,2}$ TO \mathcal{DA} for a single equation in \mathbf{A}_i . The number of the while-iterations from line 16 to 27 is at most $\log \|\mathbf{A}\|_{\max}$. We observe that (1) in each iteration, the number of newly created variables and equations is at most the value of $\text{nnz}(\mathcal{A}_i)$ at the beginning of the iteration; (2) once an auxiliary variable is added to \mathcal{A}_i , it must be replaced with a new auxiliary variable in the next iteration (if exists). Therefore, $\text{nnz}(\mathcal{A}_i) = O(\text{nnz}(\mathbf{A}_i))$ for every iteration. So, $\text{nnz}(\mathbf{B}) = O(\text{nnz}(\mathbf{A}) \log \|\mathbf{A}\|_{\max})$. \square

Knowing that the first two steps of reduction only increase the problem size by a constant factor, we obtain the following result directly.

Corollary 5.1.2. *Let $\mathbf{A} \in \mathcal{G}$ and we reduce a polynomially bounded general LE (\mathbf{A}, \mathbf{b}) to a difference-average LE (\mathbf{B}, \mathbf{c}) , where $\mathbf{B} \in \mathcal{DA}$. Then,*

$$\text{nnz}(\mathbf{B}) = O(\text{nnz}(\mathbf{A}) \log \|\mathbf{A}\|_{\max}).$$

5.2 Reducing Difference-Average LE to 1-or-3 LE

Recall that a matrix in \mathcal{DA} only has two types of rows such that for any q th equation in $\mathbf{A}\mathbf{x} = \mathbf{b}$, it takes one of the following forms:

1. Average equation: $\mathbf{x}(q_1) + \mathbf{x}(q_2) = 2\mathbf{x}(q_3)$;
2. Difference equation: $\mathbf{x}(q_1) - \mathbf{x}(q_2) = b_q$.

In this section, we present a reduction algorithm to reduce \mathcal{DA} to \mathcal{OT} , a class of matrices with a further simplified structure. Specifically, a matrix $\mathbf{A} \in \mathcal{OT}$ only has two types of rows such that for any q th equation in $\mathbf{A}\mathbf{x} = \mathbf{b}$, it has one of the following forms:

¹Note that given two singleton multi-sets each containing a single equation, e.g. $\{\mathbf{a}_1^\top \mathbf{x} = c_1\}$ and $\{\mathbf{a}_2^\top \mathbf{x} = c_2\}$ where $\mathbf{a}_1, \mathbf{a}_2$ are vectors, we define $\{\mathbf{a}_1^\top \mathbf{x} = c_1\} + \{\mathbf{a}_2^\top \mathbf{x} = c_2\} = \{\mathbf{a}_1^\top \mathbf{x} + \mathbf{a}_2^\top \mathbf{x} = c_1 + c_2\}$ and we define $\{\mathbf{a}_1^\top \mathbf{x} = c_1\} \cup \{\mathbf{a}_2^\top \mathbf{x} = c_2\} = \{\mathbf{a}_1^\top \mathbf{x} = c_1, \mathbf{a}_2^\top \mathbf{x} = c_2\}$.

Algorithm 1: REDUCE $\mathcal{G}_{z,2}$ TO \mathcal{DA}

Input: $(\mathbf{A}, \mathbf{c}^A)$ where $\mathbf{A} \in \mathcal{G}_{z,2}$ is an $m \times n$ matrix, $\mathbf{c}^A \in \mathbb{R}^n$, and $\alpha > 0$.
Output: $(\mathbf{B}, \mathbf{c}^B)$ where $\mathbf{B} \in \mathcal{DA}$ is an $m \times n$ matrix, and $\mathbf{c}^B \in \mathbb{R}^n$.

- 1 $\mathcal{X} \leftarrow \{\mathbf{u}_1, \dots, \mathbf{u}_n\}$; // \mathcal{DA} variables and index of new variables
- 2 Let \mathbf{x} be the vector of variables corresponding to the set of variables \mathcal{X}
- 3 $t \leftarrow n + 1$
- 4 $\mathcal{A} \leftarrow \emptyset, \mathcal{B} \leftarrow \emptyset$; // Multisets of main and \mathcal{DA} auxiliary equations respectively
- 5 **for** each equation $1 \leq i \leq m$ in \mathbf{A} **do**
- 6 **if** \mathbf{A}_i is already a difference equation or an average equation **then**
- 7 $\mathcal{A} \leftarrow \mathcal{A} \cup \{\mathbf{A}_{ij_+} \mathbf{u}_{j_+} - \mathbf{A}_{ij_-} \mathbf{u}_{j_-} = \mathbf{c}_i\}$
- 8 $w_i \leftarrow \frac{\alpha}{\alpha+1}$
- 9 **end**
- 10 **else**
- 11 Let $\mathcal{I}^{+1} \leftarrow \{j : \mathbf{A}_{ij} > 0\}, \mathcal{I}^{-1} \leftarrow \{j : \mathbf{A}_{ij} < 0\}$
- 12 $\mathcal{A}_i \leftarrow \{\mathbf{A}_i \mathbf{u} = \mathbf{c}_i\}$
- 13 $\mathcal{B}_i \leftarrow \emptyset$
- 14 **for** $s = -1, +1$ **do**
- 15 $r \leftarrow 0$
- 16 **while** \mathbf{A}_i is neither a difference equation nor an average equation **do**
- 17 For each j , let $\hat{\mathbf{A}}_{ij}$ be the coefficient of \mathbf{u}_j in \mathbf{A}_i .
- 18 $\mathcal{I}_{odd}^s \leftarrow \{j \in \mathcal{I}^s : \lfloor |\hat{\mathbf{A}}_{ij}|/2^r \rfloor \text{ is odd}\}$
- 19 Pair the indices of \mathcal{I}_{odd}^s into disjoint pairs $(j_1, l_1), (j_2, l_2), \dots$
- 20 **for** each pair of indices (j_k, l_k) **do**
- 21 $\mathcal{A}_i \leftarrow \mathcal{A}_i + s \cdot 2^r \{(2\mathbf{u}_t - (\mathbf{u}_{j_k} + \mathbf{u}_{l_k})) = 0\}$ ¹
- 22 $\mathcal{B}_i \leftarrow \mathcal{B}_i \cup s \cdot 2^r \{(2\mathbf{u}_t - (\mathbf{u}_{j_k} + \mathbf{u}_{l_k})) = 0\}$
- 23 $\mathcal{X} \leftarrow \mathcal{X} \cup \{\mathbf{u}_t\}$, update \mathbf{x} accordingly
- 24 $t \leftarrow t + 1$
- 25 $r \leftarrow r + 1$
- 26 **end**
- 27 **end**
- 28 **end**
- 29 $w_i \leftarrow \alpha |\mathcal{B}_i|$
- 30 $\mathcal{B} \leftarrow \mathcal{B} \cup w_i^{1/2} \cdot \mathcal{B}_i$
- 31 $\mathcal{A} \leftarrow \mathcal{A} \cup \mathcal{A}_i$.
- 32 **end**
- 33 **end**
- 34 **return** \mathbf{B}, \mathbf{c} s.t. $\mathbf{B}\mathbf{x} = \mathbf{c}$ corresponds to the equations in $\mathcal{A} \cup \mathcal{B}$, on the variable set \mathcal{X} .

1. Three-variable equation: $\mathbf{x}(q_1) + \mathbf{x}(q_2) = \mathbf{x}(q_3)$;

2. Fixed-value equation: $\mathbf{x}(q) = 1$.

We independently reduce each average or difference equation into the \mathcal{OT} form. The reduction process is presented step by step.

5.2.1 Reducing Average Equations

For each average equation $\mathbf{x}(q_1) + \mathbf{x}(q_2) = 2\mathbf{x}(q_3)$, we first introduce auxiliary variables \mathbf{s}^{avg} and replace it with a constant number of equations:

$$\begin{aligned} \mathbf{x}(q_1) + \mathbf{x}(q_2) &= \mathbf{s}^{\text{avg}}(q_1), \\ \mathbf{x}(q_3) + \mathbf{s}^{\text{avg}}(q_2) &= \mathbf{s}^{\text{avg}}(q_1), \\ \mathbf{x}(q_3) + \mathbf{s}^{\text{avg}}(q_3) &= \mathbf{s}^{\text{avg}}(q_2), \\ \mathbf{s}^{\text{avg}}(q_3) &= 0. \end{aligned} \tag{Average reduction}$$

Compared to the \mathcal{OT} form, it remains to reduce the last equation $\mathbf{s}^{\text{avg}}(q_3) = 0$ because its right-hand-side (RHS) is not 1. For this, we reduce it to

$$\begin{aligned} \mathbf{s}^{\text{avg}}(q_3) + \mathbf{s}^{\text{avg}}(q_4) &= \mathbf{s}^{\text{avg}}(q_5), \\ \mathbf{s}^{\text{avg}}(q_4) &= 1, \\ \mathbf{s}^{\text{avg}}(q_5) &= 1. \end{aligned} \tag{RHS reduction from 0 to 1}$$

Now, the reduced equations are in the form of \mathcal{OT} . And we can check that this reduction preserves equivalence with the original average equation.

5.2.2 Reducing Difference Equations

For each difference equation $\mathbf{x}(q_1) - \mathbf{x}(q_2) = b_q$, the major challenge lies in the reduction of the RHS b_q to 1. Without loss of generality, we can assume $b_q \geq 0$ because we can switch $\mathbf{x}(q_1), \mathbf{x}(q_2)$ otherwise. Moreover, we assume that entries of \mathbf{b} are integers and can be represented with polylogarithmic bits.

We present the reduction in several steps.

1. First, we reduce the difference equation to the following two equations:

$$\begin{aligned} \mathbf{x}(q_1) + \mathbf{s}^{\text{diff}}(q) &= \mathbf{x}(q_2), \\ \mathbf{s}^{\text{diff}}(q) &= b_q. \end{aligned}$$

2. We then conduct bitwise decomposition of b_q :

$$b_q = \sum_{i=0}^K r_q^{(i)} 2^i, \quad \text{where } r_q^{(i)} \in \{0, 1\} \text{ and } K = \log b_q.$$

Using more auxiliary variables, we break down $\mathbf{s}^{\text{diff}}(q) = b_q$ into K equations:

$$\begin{aligned} \mathbf{s}^{\text{diff}}(q_0) &= r_q^{(0)} \in \{0, 1\}, \\ \mathbf{s}^{\text{diff}}(q_i) &= \mathbf{s}^{\text{diff}}(q_{i-1}) + \mathbf{y}_q(i), \quad \forall i \in \{1, \dots, K\}, \\ \mathbf{y}_q(i) &= r_q^{(i)} 2^i, \quad \forall i \in \{1, \dots, K\}. \end{aligned}$$

3. In this step, we reduce $\mathbf{y}_q(i) = r_q^{(i)} 2^i$. If $r_q^{(i)} = 0$, then we skip. In the case of $r_q^{(i)} = 1$, we reduce it to a series of i equations:

$$\begin{aligned} \mathbf{y}_q(i_0) &= 1, \\ \mathbf{y}_q(i_j) &= 2\mathbf{y}_q(i_{j-1}), \quad \forall j \in \{1, \dots, i\}. \end{aligned}$$

4. For each equation $\mathbf{y}_q(i_j) = 2\mathbf{y}_q(i_{j-1})$, we reduce it to

$$\begin{aligned}\mathbf{y}_q(i_{j-1,1}) + \mathbf{y}_q(i_{j-1,2}) &= \mathbf{y}_q(i_j), \\ \mathbf{y}_q(i_{j-1,1}) - \mathbf{y}_q(i_{j-1,2}) &= 0.\end{aligned}$$

Notice that the second equation is a difference equation with RHS being 0, applying the reduction in step 1 gives

$$\begin{aligned}\mathbf{y}_q(i_{j-1,1}) + \mathbf{y}_q(i_{j-1,3}) &= \mathbf{y}_q(i_{j-1,2}), \\ \mathbf{y}_q(i_{j-1,3}) &= 0.\end{aligned}$$

5. Finally, for all equation with only one variable and 0 RHS, we apply the reduction in (RHS reduction from 0 to 1).

Now, all the equations are reduced to the two types of equations in \mathcal{OT} : $\{\mathbf{x}(q_1) + \mathbf{x}(q_2) = \mathbf{x}(q_3), \mathbf{x}(q) = 1\}$.

Lemma 5.2.1. *Let $\mathbf{A} \in \mathcal{DA}$ and we reduce a difference-average LE (\mathbf{A}, \mathbf{b}) to a 1-or-3 LE (\mathbf{B}, \mathbf{c}) , where $\mathbf{B} \in \mathcal{OT}$ by applying the above reduction algorithm. Then,*

$$\text{nnz}(\mathbf{B}) = O(\text{nnz}(\mathbf{A}) \log^2 \|\mathbf{b}\|_{\max}).$$

Proof. We first have that each average equation is reduced to a constant number of equations. For each difference equation with RHS b_q , step 1, 4, and 5 only increase the number of equations by a constant factor, while step 2 and 3 increase the reduced number of equations by a factor of $O(\log b_q)$, respectively. Thus, each difference equation is reduced to $O(\log^2 b_q)$ equations.

To conclude, $\text{nnz}(\mathbf{B}) = O(\text{nnz}(\mathbf{A}) \log^2 \|\mathbf{b}\|_{\max})$. \square

The following corollary is obtained by combining Corollary 5.1.2 and Lemma 5.2.1.

Corollary 5.2.2. *Let $\mathbf{A} \in \mathcal{G}$ and we reduce a polynomially bounded general LE (\mathbf{A}, \mathbf{b}) to a 1-or-3 LE (\mathbf{B}, \mathbf{c}) , where $\mathbf{B} \in \mathcal{OT}$ and entries of \mathbf{c} are in $\{0, 1\}$. Then,*

$$\text{nnz}(\mathbf{B}) = O(\text{nnz}(\mathbf{A}) \log \|\mathbf{A}\|_{\max} \log^2 \|\mathbf{b}\|_{\max}).$$

5.3 Reducing General LP to (Simplified) 1-or-3 LP

In this section, we present a reduction algorithm from general linear programs to (simplified) 1-or-3 linear programs.

We start with an arbitrary polynomially-bounded LP

$$\{\mathbf{Ax} = \mathbf{b}, \quad \mathbf{x} \geq \mathbf{0}\}, \quad (\text{General LP})$$

where entries of \mathbf{A}, \mathbf{b} are polynomially bounded integers and $\|\mathbf{x}\|_1$ is also polynomially bounded. Our goal is to reduce it to 1-or-3 LP, which is in the form of

$$\{\mathbf{Ax} = \mathbf{b}, \quad \mathbf{0} \leq \mathbf{x} \leq \mathbf{1}\}, \quad \text{where } \mathbf{A} \in \mathcal{OT}. \quad (1\text{-or-3 LP})$$

Compared to the reduction algorithms for general LE described in Section 5.1 and 5.2, reducing general linear programs (LPs) requires extra caution to preserve non-negativity

constraints. Moreover, to reduce to 1-or-3 LP, we have to reduce and maintain the scale of \mathbf{x} to be in the range of $[0, 1]$.

Interestingly, the upper bound of variables $\mathbf{x} \leq \mathbf{1}$ in a 1-or-3 LP can be relaxed, leading to a further simplified problem:

$$\{\mathbf{A}\mathbf{x} = \mathbf{1}, \quad \mathbf{x} \geq \mathbf{0}\}, \quad (\text{Simplified 1-or-3 LP})$$

The reduction algorithm for transforming a 1-or-3 LP into its simplified version is detailed in Section 5.3.4. The reduction algorithm from (General LP) to (1-or-3 LP) is presented in Section 5.3.1, 5.3.2 and 5.3.3, following a chain of reductions through several intermediate problems:

$$\text{General LP} \rightarrow \text{Scaled LP} \rightarrow \text{Difference-Average LP} \rightarrow \text{1-or-3 LP},$$

where a scaled LP takes the form

$$\{\mathbf{A}\mathbf{x} = \mathbf{b}, \mathbf{x} \geq \mathbf{0}, \|\mathbf{x}\|_{\max} \leq c\}, \quad \text{where } c \leq \frac{1}{2} \text{ is a constant}, \quad (\text{Scaled LP})$$

and a difference-average LP is defined as

$$\{\mathbf{A}\mathbf{x} = \mathbf{b}, \mathbf{x} \geq \mathbf{0}, \|\mathbf{x}\|_{\max} \leq \frac{1}{2}\}, \quad \text{where } \mathbf{A} \in \mathcal{DA}. \quad (\text{Difference-Average LP})$$

5.3.1 Reducing General LP to Scaled LP

Since (General LP) is polynomially bounded, let

$$B = \|\mathbf{x}\|_{\max} \leq \|\mathbf{x}\|_1 = \text{poly}(n).$$

And set k and C be integers such that

$$k = \lceil \log_2 \|\mathbf{b}\|_{\max} \rceil, \quad C = 2^{\max\{k, 1 + \lceil \log_2 B \rceil\}}.$$

Then we set

$$\tilde{\mathbf{A}} = \frac{C\mathbf{A}}{2^k}, \quad \tilde{\mathbf{b}} = \frac{\mathbf{b}}{2^k}, \quad \tilde{\mathbf{x}} = \frac{\mathbf{x}}{C},$$

and obtain a scaled LP

$$\left\{ \tilde{\mathbf{A}}\tilde{\mathbf{x}} = \tilde{\mathbf{b}}, \quad \tilde{\mathbf{x}} \geq \mathbf{0}, \quad \|\tilde{\mathbf{x}}\|_{\max} \leq \frac{B}{C} \right\}, \quad (5.2)$$

where we can notice that

- $\tilde{\mathbf{A}} = \frac{C\mathbf{A}}{2^k}$ is a matrix with integer entries since $\frac{C}{2^k}$ is an integer;
- $\|\tilde{\mathbf{b}}\|_{\max} \leq 1$;
- $\|\tilde{\mathbf{x}}\|_{\max} \leq \frac{B}{C} \leq \frac{1}{2}$.

5.3.2 Reducing Scaled LP to Difference-Average LP

We apply the reduction algorithm described in Section 5.1. Recall the three steps to reduce a general LE to \mathcal{DA} , and we closely keep track of the non-negativity and the infinity norm of variables.

1. To reduce $\tilde{\mathbf{A}}$ to a matrix with row sum 0, we set

$$\tilde{\mathbf{A}}' = \begin{bmatrix} \tilde{\mathbf{A}} & -\tilde{\mathbf{A}}\mathbf{1} \end{bmatrix}, \quad \tilde{\mathbf{b}}' = \tilde{\mathbf{b}}, \quad \tilde{\mathbf{x}}' = \begin{bmatrix} \tilde{\mathbf{x}} \\ x_1^{\text{aux}} \end{bmatrix},$$

and we add an additional equality constraint

$$x_1^{\text{aux}} = 0.$$

Then we can check that $\tilde{\mathbf{x}}' \geq \mathbf{0}$ and $\|\tilde{\mathbf{x}}'\|_{\max} \leq 1/2$.

2. To reduce $\tilde{\mathbf{A}}'$ to a matrix such that the sum of positive entries in each row is a power of 2, we set

$$\tilde{\mathbf{A}}'' = \begin{bmatrix} \tilde{\mathbf{A}}' & \mathbf{g} & -\mathbf{g} \\ \mathbf{0} & 1 & -1 \end{bmatrix}, \quad \tilde{\mathbf{b}}'' = \begin{bmatrix} \tilde{\mathbf{b}}' \\ 0 \end{bmatrix}, \quad \tilde{\mathbf{x}}'' = \begin{bmatrix} \tilde{\mathbf{x}}' \\ x_2^{\text{aux}} \\ x_3^{\text{aux}} \end{bmatrix},$$

and we add additional equality constraints

$$x_2^{\text{aux}} = 0, \quad x_3^{\text{aux}} = 0.$$

Again, we can check that $\tilde{\mathbf{x}}'' \geq \mathbf{0}$ and $\|\tilde{\mathbf{x}}''\|_{\max} \leq 1/2$.

3. To reduce $\tilde{\mathbf{A}}''$ to a difference-average matrix, notice that each time we introduce a new variable, its value is the average of two old variables. Hence, new variables are also non-negative, and the infinity norm of the auxiliary variables will not exceed that of the original variables. The resulting problem is a difference-average LP, denoted as

$$\left\{ \bar{\mathbf{A}}\bar{\mathbf{x}} = \bar{\mathbf{b}}, \quad \bar{\mathbf{x}} \geq \mathbf{0}, \quad \|\bar{\mathbf{x}}\|_{\max} \leq \frac{1}{2} \right\}, \quad \text{where } \bar{\mathbf{A}} \in \mathcal{DA}. \quad (5.3)$$

In particular, $\bar{\mathbf{b}} = \begin{bmatrix} \tilde{\mathbf{b}}'' \\ \mathbf{0} \end{bmatrix}$, thus $\|\bar{\mathbf{b}}\|_{\max} \leq 1$.

5.3.3 Reducing Difference-Average LP to 1-or-3 LP

In this step, we extend the reduction algorithm from \mathcal{DA} to \mathcal{OT} , as described in Section 5.2. To reduce average equations as shown in ([Average reduction](#)), we have

$$\mathbf{s}^{\text{avg}} \geq \mathbf{0}, \quad \|\mathbf{s}^{\text{avg}}\|_{\max} \leq 2\|\bar{\mathbf{x}}\|_{\max} \leq 1.$$

The reduction for difference equations needs to be slightly adjusted. The reason is that in Section 5.2, we assume that entries of the RHS vector are polynomially bounded

integers, while here we have the RHS $\|\bar{\mathbf{b}}\|_{\max} \leq 1$. Therefore, for each difference equation $\bar{\mathbf{x}}(q_1) - \bar{\mathbf{x}}(q_2) = \bar{b}_q$, conducting bitwise decomposition of \bar{b}_q gives

$$\bar{b}_q = \sum_{i=0}^k r_q^{(i)} 2^{-i}, \quad \text{where } r_q^{(i)} \in \{0, 1\} \text{ and } k = \lceil \log \|\mathbf{b}\|_{\max} \rceil.$$

And to reduce $\mathbf{y}_q(i) = r_q^{(i)} 2^{-i}$ in the case of $r_q^{(i)} = 1$ and for $i \in \{1, \dots, k\}$, we reduce it to the following series of i equations:

$$\begin{aligned} \mathbf{y}_q(i_0) &= 1, \\ \mathbf{y}_q(i_j) &= \mathbf{y}_q(i_{j-1})/2, \quad \forall j \in \{1, \dots, i\}. \end{aligned}$$

And for each equations $\mathbf{y}_q(i_j) = \mathbf{y}_q(i_{j-1})/2$, we reduce it to

$$\begin{aligned} \mathbf{y}_q(i_{j-1,1}) + \mathbf{y}_q(i_{j-1,2}) &= 2\mathbf{y}_q(i_j), \\ \mathbf{y}_q(i_{j-1,2}) &= 0. \end{aligned}$$

Notice that the first equation is an average equation, which can be further reduced by applying (**Average reduction**). As the last step, we reduce one-variable equations with 0 right-hand side by applying (**RHS reduction from 0 to 1**). Again, we have the newly introduced variables to satisfy

$$\mathbf{s}^{\text{diff}}, \mathbf{y} \geq \mathbf{0} \quad \text{and} \quad \|\mathbf{s}^{\text{diff}}\|_{\max}, \|\mathbf{y}\|_{\max} \leq 1.$$

Putting it altogether, we reduce the difference-average LP (5.3) to a 1-or-3 LP in the form of (**1-or-3 LP**).

Lemma 5.3.1. *Let $\mathbf{A} \in \mathcal{G}$ be the coefficient matrix in (**General LP**) and let $\mathbf{B} \in \mathcal{OT}$ be the coefficient matrix in (**1-or-3 LP**). Then,*

$$\text{nnz}(\mathbf{B}) = O(\text{nnz}(\mathbf{A}) \log \|\mathbf{A}\|_{\max} \log^2 \|\mathbf{b}\|_{\max}).$$

Proof. It follows directly from Corollary 5.2.2. □

5.3.4 Reducing 1-or-3 LP to Simplified 1-or-3 LP

Actually, (**1-or-3 LP**) can be further simplified by dropping the upper bound of variables $\mathbf{x} \leq \mathbf{1}$, and achieving the following form:

$$\{\mathbf{A}\mathbf{x} = \mathbf{1}, \quad \mathbf{x} \geq \mathbf{0}\},$$

where equality constraints have one of the following forms:

1. $\mathbf{x}(q_1) + \mathbf{x}(q_2) + \mathbf{x}(q_3) = 1$;
2. $\mathbf{x}(q) = 1$.

For this, because of $\mathbf{x} \leq \mathbf{1}$, we can replace each equation in the form of $\mathbf{x}(q_1) + \mathbf{x}(q_2) = \mathbf{x}(q_3)$ with the following equations:

$$\begin{aligned} \mathbf{x}(q_1) + \mathbf{x}(q_2) + \mathbf{s}(q_1) &= 1, \\ \mathbf{x}(q_3) + \mathbf{s}(q_1) + \mathbf{s}(q_2) &= 1, \\ \mathbf{s}(q_1) &\geq 0, \\ \mathbf{s}(q_2) &= 0. \end{aligned} \tag{Dropping upper bound}$$

And the equation $\mathbf{s}(q_2) = 0$ can be further reduced to

$$\begin{aligned} \mathbf{s}(q_2) + \mathbf{s}(q_3) + \mathbf{s}(q_4) &= 1, \\ \mathbf{s}(q_3) &= 1, \\ \mathbf{s}(q_4) &\geq 0. \end{aligned} \tag{Another RHS reduction from 0 to 1}$$

The simplified 1-or-3 LP has a much more direct and intuitive structure, making it easier to relate to 3-SAT. Like 3-SAT, it serves as a promising starting point for proving the hardness of other problems, and for ultimately establishing a *sparse-linear-program complete* problem class, similar to the NP-complete problem class. This perspective aligns with the findings of [PW17], which also reduced general LPs to the simplified 1-or-3 LP but used a different reduction algorithm. Moreover, [PW17] extended this work by reducing the simplified 1-or-3 LP to LP relaxations of several combinatorial optimization problems, such as Set Cover, Set Packing, Maximum Satisfiability, and Maximum Independent Set. These results concluded that the LP relaxations of these combinatorial problems are sparse-linear-program complete.

Chapter 6

1-Laplacian Solver for Well-Shaped Simplicial Complexes

This chapter is based on [DZ23b]. We will prove Theorem 1.4.6 and 1.4.7. In the published version of [DZ23b], we further developed additional technical components that are not included in this thesis but may be of independent interest (e.g., sufficient conditions for 3-complexes that enable efficient r -hollowing computation; a nearly-linear time hollowing algorithm; a faster solver for up-Laplacian).

6.1 Prior Works

Cohen, Fasy, Miller, Nayyeri, Peng, and Walkington [Coh+14a] initiated the study of fast approximate solvers for 1-Laplacian linear equations. They designed a nearly-linear time approximate solver for simplicial complexes *with zero Betti numbers and known collapsing sequences*. Later, Black, Maxwell, Nayyeri, and Winkelman [Bla+22], and Black and Nayyeri [BN22] generalized this algorithm to subcomplexes of such a complex *with bounded first Betti numbers*. One concrete example studied in these papers is convex simplicial complexes that piecewise linearly triangulate a convex ball in \mathbb{R}^3 , for which a collapsing sequence exists and can be computed in linear time [Chi67; Chi80]. However, deciding whether a simplicial complex has a collapsing sequence is NP-hard in general [Tan16]; computing the Betti numbers is as hard as computing the ranks of general $\{0, 1\}$ matrices [EP14]. In addition, 1-Laplacian systems for general simplicial complexes embedded in \mathbb{R}^4 are as hard to solve as general sparse linear equations [Din+22], for which the best-known algorithms need super-quadratic time [PV21; Nie22].

In addition to the specialized solvers for 1-Laplacian systems mentioned above, *Nested Dissection* can solve 1-Laplacian systems in quadratic time for simplicial complexes in \mathbb{R}^3 with additional geometric structures [Geo73; LRT79; MT90] such as bounded aspect ratios of individual tetrahedra. Furthermore, iterative methods such as *Preconditioned Conjugate Gradient* approximately solve 1-Laplacian systems in time $\tilde{O}(n\sqrt{\kappa})$,¹ where n is the number of simplexes and κ is the condition number of the coefficient matrix.

Inspired by solvers that leverage both geometric structures and spectral properties, we develop efficient 1-Laplacian approximate solvers for well-shaped simplicial complexes embedded in \mathbb{R}^3 *without known collapsing sequences and with arbitrary Betti num-*

¹We use $\tilde{O}(\cdot)$ to hide polylog factors on the number of simplexes, the condition number, and the inverse of the error parameter.

bers. Our solver adapts the Incomplete Nested Dissection algorithm, proposed by Kyng, Peng, Schwieterman, and Zhang [Kyn+18] for solving linear equations in well-shaped 3-dimensional truss stiffness matrices. These matrices represent another generalization of graph Laplacians; however, they differ quite from the 1-Laplacians studied in this paper. A primary distinction is that the kernel of a truss stiffness matrix has an explicit and well-understood form, while computing a 1-Laplacian’s kernel is as hard as that for a general matrix.

6.2 Motivations and Applications

In the past decades, combinatorial Laplacians have played a crucial role in the development of computational topology and topological data analysis in various domains, such as statistics [Jia+11; ODO13], graphics and imaging [Ma+11b; Ton+03b], brain networks [Lee+19], deep learning [Bro+17], signal processing [BS20], and cryo-electron microscope [YL17b]. We recommend readers consult accessible surveys [Ghr08b; Car09; EH10; Lim20] for more information.

Combinatorial Laplacians have their roots in the study of discrete Hodge decomposition [Eck44], which states that the kernel of the i -Laplacian \mathbf{L}_i is isomorphic to the i th homology group of the simplicial complex. Among the many applications of combinatorial Laplacians, a central problem is determining the Betti numbers – the ranks of the homology groups – which are important topological invariants. Additionally, discrete Hodge decomposition allows for the extraction of meaningful information from data by decomposing them into three mutually orthogonal components: gradient (in the image of ∂_i^\top), curl (in the image of ∂_{i+1}), and harmonic (in the kernel of \mathbf{L}_i) components. For instance, the three components of edge flows in a graph capture the global trends, local circulations, and “noise”.

The computation of both Betti numbers and discrete Hodge decomposition of higher-order flows can be achieved by solving systems of linear equations in combinatorial Laplacians [Fri96; Lim20]. The rank of a matrix \mathbf{L}_i can be determined by solving a logarithmic number of linear equation systems in \mathbf{L}_i [BV21]. The discrete Hodge decomposition can be calculated by solving least square problems involving boundary operators or combinatorial Laplacians, which in turn reduces to solving linear equations in these matrices.

Furthermore, an important question in numerical linear algebra concerns whether the nearly-linear time solvers for graph Laplacian linear equations can be generalized to larger classes of linear equations. Researchers have achieved success with elliptic finite element systems [BHV08b], connection Laplacians [Kyn+16], directed Laplacians [Coh+17a; Coh+18], well-shaped truss stiffness matrices [DS07; ST08; Kyn+18]. It would be intriguing to determine what structures of linear equations facilitate faster solvers. Another theoretically compelling reason for developing efficient solvers for 1-Laplacians stems from the “equivalence” of time complexity between solving 1-Laplacian systems and general sparse systems of linear equations [Din+22]. If one can solve all 1-Laplacian systems in time $\tilde{O}((\text{number of simplexes})^c)$ where $c \geq 1$ is a constant, then one can solve all general systems of linear equations in time $\tilde{O}((\text{number of non-zeros})^c)$.

6.3 Notations and Preliminaries

Aspect Ratio. The *aspect ratio* of a set $S \subset \mathbb{R}^3$ is the radius of the smallest ball containing S divided by the radius of the largest ball contained in S . The aspect ratio of S is always greater than or equal to 1. We say a simplex σ is *stable* if it has $O(1)$ aspect ratio and $\Theta(1)$ weight. Miller and Thurston proved the following lemma. As a corollary, the numbers of the vertices, the edges, the triangles, and the tetrahedra of a 3-complex \mathcal{K} that is composed of stable tetrahedra are all equal up to a constant factor.

Lemma 6.3.1 (Lemma 4.1 of [MT90]). *Let \mathcal{K} be a 3-complex in \mathbb{R}^3 in which each tetrahedron has $O(1)$ aspect ratio. Then, each vertex of \mathcal{K} is contained in at most $O(1)$ tetrahedra.*

r -Hollowings. Let \mathcal{K} be a pure 3-complex with n simplexes. A set of triangles $\hat{\Delta}_1, \dots, \hat{\Delta}_k$ form a *triangle path* of length $k - 1$ if for any $1 \leq i \leq k - 1$, $\hat{\Delta}_i$ and $\hat{\Delta}_{i+1}$ share an edge. The *triangle distance* between two triangles Δ_1 and Δ_2 is the shortest triangle path length between Δ_1 and Δ_2 . The *triangle diameter* of \mathcal{K} is the longest triangle distance between any two triangles. A *spherical shell* is $\{\mathbf{x} \in \mathbb{R}^3 : R_1 \leq \|\mathbf{x}\|_2 \leq R_2\}$ where $R_1 < R_2$. If \mathcal{K} triangulates a spherical shell, we define the *shell width* to be the shortest triangle distance between any two triangles where one is on the outer sphere and one is on the inner sphere.

Definition 6.3.2 (r -Hollowing). Let \mathcal{K} be a 3-complex with n simplexes, and let $r = o(n)$ be a positive number. We divide \mathcal{K} into $O(n/r)$ *regions* each of $O(r)$ simplexes, $O(r^{2/3})$ *boundary simplexes*, and $O(r^{2/3})$ *exterior simplexes*. Non-boundary simplexes are called *interior simplexes*.² Interior simplexes from different regions do not share any subsimplexes. The boundary of each region triangulates a spherical shell in \mathbb{R}^3 and has triangle diameter $O(r^{1/3})$ and shell width at least 5. The union of all boundary simplexes of each region is referred to as an *r -hollowing* of \mathcal{K} .

Figure 6.1 illustrates an example of r -hollowing by showing a cross-section of a 3-complex. The left figure presents a cross-section of a 3-complex \mathcal{K} with two holes inside, depicted as two empty discs. The outlines in black represent the “exterior simplexes” of \mathcal{K} . On the right-hand side of Figure 6.1, the gray area represents an r -hollowing of \mathcal{K} . It divides \mathcal{K} into 4 balanced regions $\{(1), (2), (3), (4)\}$, each indicated by the area inside the red, blue, green, or orange outline, respectively. Region (1) includes the smaller hole, while the larger hole “intersects” regions (3) and (4). The unshaded area inside the squares corresponds to “interior simplexes,” which encompasses the exterior simplexes of the smaller hole and the left half of the larger hole. In contrast, the gray area represents “boundary simplexes,” including the exterior simplexes of the right half of the larger hole. The shell width is indicated by arrows.

In the full version, [DZ23b] examines sufficient conditions for 3-complexes that enable us to compute an r -hollowing in nearly-linear time. Here, we focus on the case where \mathcal{K} is already equipped with a known r -hollowing.

²We would like to emphasize that “exterior simplex” is defined for any 3-complex, while “boundary simplex” and “interior simplex” are defined for r -hollowing. Although boundary and interior simplexes are mutually exclusive, an exterior simplex can be either a boundary or an interior simplex for a region of an r -hollowing.




Figure 6.1: A cross-sectional illustration of an r -hollowing.

Algorithms for Solving Linear Equations. Our algorithm combines two of the most important tools for solving linear systems: Nested Dissection and preconditioning.

Theorem 6.3.3 (Nested Dissection [MT90]). *Let $\mathbf{A} \in \mathbb{R}^{n \times n}$ be a symmetric matrix defined on a simple tetrahedral mesh. A Cholesky factorization $\mathbf{A} = \mathbf{P}\mathbf{L}\mathbf{L}^\top\mathbf{P}^\top$ can be computed in time $O(n^2)$, in which \mathbf{P} is a permutation matrix and \mathbf{L} is a lower triangular matrix with $O(n^{4/3})$ non-zero entries. As a result, a linear system in \mathbf{A} can be solved in time $O(n^2)$ by Gaussian Elimination.*

Theorem 6.3.4 (Preconditioned Conjugate Gradient [Axe85]). *Let $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times n}$ be two symmetric PSD matrices, and let $\mathbf{b} \in \mathbb{R}^n$. Each iteration of Preconditioned Conjugate Gradient multiplies one vector with \mathbf{A} , solves one system of linear equations in \mathbf{B} , and performs a constant number of vector operations. For any $\epsilon > 0$, the algorithm outputs an \mathbf{x} satisfying $\|\mathbf{A}\mathbf{x} - \Pi_{\text{Im}(\mathbf{A})}\mathbf{b}\|_2 \leq \epsilon \|\Pi_{\text{Im}(\mathbf{A})}\mathbf{b}\|_2$ in $O(\sqrt{\kappa} \log(\kappa/\epsilon))$ such iterations, where $\Pi_{\text{Im}(\mathbf{A})}$ is the orthogonal projection matrix onto the image of \mathbf{A} and $\kappa = \kappa(\mathbf{A}, \mathbf{B})$.*

6.4 Main Results

We formally state our main results as follows.

Theorem 6.4.1 (A Pure 3-Complex). *Let \mathcal{K} be a pure 3-complex embedded in \mathbb{R}^3 consisting of n stable simplexes and with a known r -hollowing. Let \mathbf{L}_1 be the 1-Laplacian operator of \mathcal{K} , and let Π_1 be the orthogonal projection matrix onto the image of \mathbf{L}_1 . For any vector \mathbf{b} and $\epsilon > 0$, we can find a solution $\tilde{\mathbf{x}}$ such that $\|\mathbf{L}_1\tilde{\mathbf{x}} - \Pi_1\mathbf{b}\|_2 \leq \epsilon \|\Pi_1\mathbf{b}\|_2$ in time*

$$O(nr + n^{4/3}r^{5/18} \log(n/\epsilon) + n^2r^{-2/3}).$$

We will overview our algorithm for Theorem 6.4.1 in Section 6.5 and provide detailed proofs in Section 6.6, 6.7, 6.8, and 6.9.

Theorem 6.4.2 (A Union of Pure 3-Complexes). *Let \mathcal{U} be a union of pure 3-complexes glued together by identifying certain subsets of their exterior simplexes. Assume that each 3-complex chunk is embedded in \mathbb{R}^3 , contains n_i stable simplexes, and has a known $\Theta(n_i^{3/5})$ -hollowing. Let \mathbf{L}_1 be the 1-Laplacian operator of \mathcal{U} , and let Π_1 be the orthogonal projection matrix onto the image of \mathbf{L}_1 . For any vector \mathbf{b} and $\epsilon > 0$, we can find a solution $\tilde{\mathbf{x}}$ such that $\|\mathbf{L}_1\tilde{\mathbf{x}} - \Pi_1\mathbf{b}\|_2 \leq \epsilon \|\Pi_1\mathbf{b}\|_2$ in time*

$$\tilde{O}(n^{8/5} + n^{3/10}k^2 + k^3),$$

where n is the number of simplexes in \mathcal{U} and k is the number of exterior simplexes shared by more than one chunk.

We will defer our proof of Theorem 6.4.2 to Section 6.10.

6.5 Algorithm Overview

Cohen, Fasy, Miller, Nayyeri, Peng, and Walkington [Coh+14a] observed that

$$\mathbf{L}_1^\dagger = (\mathbf{L}_1^{\text{down}})^\dagger + (\mathbf{L}_1^{\text{up}})^\dagger,$$

where $\mathbf{L}_1^{\text{down}} = \partial_1^\top \mathbf{W}_0 \partial_1$ is the down-Laplacian and $\mathbf{L}_1^{\text{up}} = \partial_2 \mathbf{W}_2 \partial_2^\top$ is the up-Laplacian.

The orthogonal projection matrices onto $\text{Im}(\partial_1^\top)$ and $\text{Im}(\partial_2)$ are:

$$\Pi_1^{\text{down}} \stackrel{\text{def}}{=} \partial_1^\top (\partial_1 \partial_1^\top)^\dagger \partial_1, \quad \Pi_1^{\text{up}} \stackrel{\text{def}}{=} \partial_2 (\partial_2^\top \partial_2)^\dagger \partial_2^\top.$$

Lemma 6.5.1 (Lemma 4.1 of [Coh+14a]). *Let \mathbf{b} be a vector. Consider the systems of linear equations: $\mathbf{L}_1 \mathbf{x} = \Pi_1 \mathbf{b}$, $\mathbf{L}_1^{\text{up}} \mathbf{x}^{\text{up}} = \Pi_1^{\text{up}} \mathbf{b}$, $\mathbf{L}_1^{\text{down}} \mathbf{x}^{\text{down}} = \Pi_1^{\text{down}} \mathbf{b}$. Then, $\mathbf{x} = \Pi_1^{\text{up}} \mathbf{x}^{\text{up}} + \Pi_1^{\text{down}} \mathbf{x}^{\text{down}}$.*

Lemma 6.5.1 implies that four operators are needed to approximate \mathbf{L}_1^\dagger : (1) an approximate projection operator $\tilde{\Pi}_1^{\text{down}} \approx \Pi_1^{\text{down}}$, (2) an approximate projection operator $\tilde{\Pi}_1^{\text{up}} \approx \Pi_1^{\text{up}}$, (3) a down-Laplacian solver $\mathbf{Z}_1^{\text{down}}$ such that $\mathbf{L}_1^{\text{down}} \mathbf{Z}_1^{\text{down}} \mathbf{b} \approx \mathbf{b}$ for any $\mathbf{b} \in \text{Im}(\mathbf{L}_1^{\text{up}})$, and (4) an up-Laplacian solver \mathbf{Z}_1^{up} such that $\mathbf{L}_1^{\text{up}} \mathbf{Z}_1^{\text{up}} \mathbf{b} \approx \mathbf{b}$ for any $\mathbf{b} \in \text{Im}(\mathbf{L}_1^{\text{up}})$.

6.5.1 Down-Laplacian

We will apply the same approximate orthogonal projection $\tilde{\Pi}_1^{\text{down}}$ given in [Coh+14a], which does not depend on Betti numbers. Our solver for the first down-Laplacian is a slight modification of the one in [Coh+14a] to incorporate simplex weights. We state the two lemmas but defer their proofs to Section 6.6.

Lemma 6.5.2 (Down-Projection Operator, Lemma 3.2 of [Coh+14a]). *Let \mathcal{K} be a 3-complex with n simplexes. For any $\epsilon > 0$, there exists a linear operator $\tilde{\Pi}_1^{\text{down}}$, computable in nearly-linear time, such that*

$$(1 - \epsilon) \Pi_1^{\text{down}} \preceq \tilde{\Pi}_1^{\text{down}}(\epsilon) \preceq \Pi_1^{\text{down}}.$$

Lemma 6.5.3 (Down-Laplacian Solver). *Let \mathcal{K} be a weighted simplicial complex, and let $\mathbf{b} \in \text{Im}(\mathbf{L}_1^{\text{down}})$. There exists an operator $\mathbf{Z}_1^{\text{down}}$ such that $\mathbf{L}_1^{\text{down}} \mathbf{Z}_1^{\text{down}} \mathbf{b} = \mathbf{b}$. In addition, we can compute $\mathbf{Z}_1^{\text{down}} \mathbf{b}$ in linear time.*

6.5.2 Solver for Up-Laplacian

One of our primary technical contributions is the development of an efficient solver for the up-Laplacian system, stated in Lemma 6.5.4. We will describe the key idea behind our solver in this section.

Lemma 6.5.4 (Up-Laplacian Solver). *Let \mathcal{K} be a pure 3-complex embedded in \mathbb{R}^3 and composed of n stable simplexes. Suppose we are given an r -hollowing for \mathcal{K} where $r = o(n)$. Then for any $\epsilon > 0$, there exists an operator \mathbf{Z}_1^{up} such that*

$$\forall \mathbf{b} \in \text{Im}(\mathbf{L}_1^{up}), \quad \|\mathbf{L}_1^{up} \mathbf{Z}_1^{up} \mathbf{b} - \mathbf{b}\|_2 \leq \epsilon \|\mathbf{b}\|_2.$$

In addition, $\mathbf{Z}_1^{up} \mathbf{b}$ can be computed in time $O(nr + n^{4/3} r^{5/18} \log(n/\epsilon) + n^2 r^{-2/3})$.

We remark that Lemma 6.5.4 can be improved to $\tilde{O}(n^{3/2})$ by using a slightly different r -hollowing (proved in Section 10 of [DZ23b]), which might be of independent interest. Since the bottleneck of our solver for 1-Laplacians is from the projection for up 1-Laplacians, we use the same r -hollowing here.

The given r -hollowing suggests a partition of the edges in \mathcal{K} into $F \cup C$. We will explain the concrete partition shortly. We have the following matrix identity:

$$\mathbf{L}_1^{up} = \begin{bmatrix} \mathbf{I} & \\ \mathbf{L}_1^{up}(C, F) \mathbf{L}_1^{up}(F, F)^\dagger & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{L}_1^{up}(F, F) & \\ & \text{Sc}(\mathbf{L}_1^{up})_C \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{L}_1^{up}(F, F)^\dagger \mathbf{L}_1^{up}(F, C) \\ & \mathbf{I} \end{bmatrix},$$

where

$$\text{Sc}(\mathbf{L}_1^{up})_C = \mathbf{L}_1^{up}(C, C) - \mathbf{L}_1^{up}(C, F) \mathbf{L}_1^{up}(F, F)^\dagger \mathbf{L}_1^{up}(F, C).$$

The following Lemma 6.5.5 reduces (approximately) solving a system in \mathbf{L}_1^{up} to (approximately) solving two systems in $\mathbf{L}_1^{up}(F, F)$ and one system in $\text{Sc}(\mathbf{L}_1^{up})_C$, which is proved in Appendix C. It is worth noting that Lemma 6.5.5 holds if we replace \mathbf{L}_1^{up} with an arbitrary symmetric PSD matrix, and we will apply it or its variants for different PSD matrices in our solvers. To avoid introducing additional notations, we state the lemma below in terms of \mathbf{L}_1^{up} .

Lemma 6.5.5. *Suppose we have two operators:*

1. $\text{UPLAPFSOLVER}(\cdot)$ such that given any $\mathbf{b} \in \text{Im}(\mathbf{L}_1^{up}(F, F))$, $\text{UPLAPFSOLVER}(\mathbf{b})$ returns a vector \mathbf{x} satisfying $\mathbf{L}_1^{up}(F, F) \mathbf{x} = \mathbf{b}$;
2. $\text{SCHURSOLVER}(\cdot, \cdot)$ such that for any $\mathbf{h} \in \text{Im}(\text{Sc}(\mathbf{L}_1^{up})_C)$ and $\delta > 0$, $\text{SCHURSOLVER}(\mathbf{h}, \delta)$ returns $\tilde{\mathbf{x}}$ satisfying $\|\text{Sc}(\mathbf{L}_1^{up})_C \tilde{\mathbf{x}} - \mathbf{h}\|_2 \leq \delta \|\mathbf{h}\|_2$.

Given any $\mathbf{b} = \begin{bmatrix} \mathbf{b}_F \\ \mathbf{b}_C \end{bmatrix} \in \text{Im}(\mathbf{L}_1^{up})$ and any $\epsilon > 0$, let

$$\begin{aligned} \mathbf{h} &= \mathbf{b}_C - \mathbf{L}_1^{up}(F, F) \cdot \text{UPLAPFSOLVER}(\mathbf{b}_F), \\ \tilde{\mathbf{x}}_C &= \text{SCHURSOLVER}(\mathbf{h}, \delta), \\ \tilde{\mathbf{x}}_F &= \text{UPLAPFSOLVER}(\mathbf{b}_F - \mathbf{L}_1^{up}(F, C) \tilde{\mathbf{x}}_C), \end{aligned} \tag{6.1}$$

where $\delta \leq \frac{\epsilon}{1 + \|\mathbf{L}_1^{up}(F, F) \mathbf{L}_1^{up}(F, F)^\dagger\|_2}$. Then,

$$\|\mathbf{L}_1^{up} \tilde{\mathbf{x}} - \mathbf{b}\|_2 \leq \epsilon \|\mathbf{b}\|_2,$$

where $\tilde{\mathbf{x}} = \begin{bmatrix} \tilde{\mathbf{x}}_F \\ \tilde{\mathbf{x}}_C \end{bmatrix}$. Let $m_F = |F|$ and $m_C = |C|$, and let UPLAPFSOLVER have runtime $t_1(m_F)$ and SCHURSOLVER have runtime $t_2(m_C)$. Then, we can compute $\tilde{\mathbf{x}}$ in time $O(t_1(m_F) + t_2(m_C) + m_F + m_C)$.

The proof of Lemma 6.5.5 can be found in Appendix C.

Partitioning the Edges

As suggested by Lemma 6.5.5, we want to partition the edges of \mathcal{K} into $F \cup C$ so that both systems in $\mathbf{L}_1^{\text{up}}(F, F)$ and the Schur complement $\text{Sc}(\mathbf{L}_1^{\text{up}})_C$ can be efficiently solved. The given $O(n^{3/5})$ -hollowing divides \mathcal{K} into “disjoint” and “balanced” regions with small boundary. Let F be the set of the interior edges of the regions and C be the set of the boundary edges.

We first show that the interiors of different regions are “disjoint” in the sense that $\mathbf{L}_1^{\text{up}}(F, F)$ is a block diagonal matrix where each diagonal block corresponds to the interior of a region. We can write \mathbf{L}_1^{up} as the sum of rank-1 matrices that each corresponds to a triangle in \mathcal{K} :

$$\mathbf{L}_1^{\text{up}} = \partial_2 \mathbf{W}_2 \partial_2^\top = \sum_{\sigma: \text{triangle in } \mathcal{K}} \mathbf{W}_2(\sigma, \sigma) \cdot \partial_2(:, \sigma) \partial_2(:, \sigma)^\top.$$

For any two edges e_1, e_2 , $\mathbf{L}_1^{\text{up}}(e_1, e_2) = 0$ if and only if no triangle in \mathcal{K} contains both e_1, e_2 . By our definition of r -hollowing in Definition 6.3.2, for different regions R_1, R_2 of \mathcal{K} w.r.t. an r -hollowing, no triangle contains both an edge from R_1 and an edge from R_2 .

In addition, the following lemma shows that the boundaries of the regions well approximate the Schur complement onto the boundaries. The proof is in Section 6.7.2.

Lemma 6.5.6 (Spectral Bounds for r -Hollowing). *Let \mathcal{K} be a pure 3-complex embedded in \mathbb{R}^3 composed of stable simplexes. Let \mathcal{T} be an r -hollowing of \mathcal{K} , and let C be the edges of \mathcal{T} . Then,*

$$\mathbf{L}_{1,\mathcal{T}}^{\text{up}} \preceq \text{Sc}(\mathbf{L}_1^{\text{up}})_C \preceq O(r) \mathbf{L}_{1,\mathcal{T}}^{\text{up}}.$$

Proof of Lemma 6.5.4 for Up-Laplacian Solver

Algorithm 2 sketches a pseudo-code for our up-Laplacian solver.

Algorithm 2: UPLAPSOLVER($\mathcal{K}, \mathcal{T}, \mathbf{b}, \epsilon$)

Data: A pure 3-complex \mathcal{K} of n stable simplexes with up-Laplacian \mathbf{L}_1^{up} , an r -hollowing \mathcal{T} , a vector $\mathbf{b} \in \text{Im}(\mathbf{L}_1^{\text{up}})$, an error parameter $\epsilon > 0$

Result: An approximate solution $\tilde{\mathbf{x}}$ such that $\|\mathbf{L}_1^{\text{up}} \tilde{\mathbf{x}} - \mathbf{b}\|_2 \leq \epsilon \|\mathbf{b}\|_2$

- 1 $F \leftarrow$ the interior edges of regions of \mathcal{K} w.r.t. \mathcal{T} ; $C \leftarrow$ the boundary edges of regions.
 - 2 UPLAPFSOLVER(\cdot) \leftarrow a solver by Nested Dissection that satisfies the requirement in Lemma 6.5.5.
 - 3 SCHURSOLVER(\cdot, \cdot) \leftarrow a solver by Preconditioned Conjugate Gradient with the preconditioner being the up-Laplacian of \mathcal{T} that satisfies the requirement in Lemma 6.5.5.
 - 4 $\tilde{\mathbf{x}} \leftarrow$ computed by Equation (6.1)
 - 5 **return** solution $\tilde{\mathbf{x}}$
-

By Lemma 6.5.5, the $\tilde{\mathbf{x}}$ returned by Algorithm 2 satisfies $\|\mathbf{L}_1^{\text{up}} \tilde{\mathbf{x}} - \mathbf{b}\|_2 \leq \epsilon \|\mathbf{b}\|_2$. To bound the runtime of Algorithm 2, we need the following lemmas for lines 2 and 3.

Lemma 6.5.7 (Solver for the “F” Part). *Let \mathcal{K} be a pure 3-complex embedded in \mathbb{R}^3 and composed of n stable simplexes. Let \mathcal{T} be an r -hollowing of \mathcal{K} , and let F be the set of interior edges in each region of \mathcal{K} w.r.t. \mathcal{T} . Then, with a pre-processing time $O(nr)$, there exists a solver $\text{UPLAPFSOLVER}(\cdot)$ such that given any $\mathbf{b}_F \in \text{im}(\mathbf{L}_1^{\text{up}}(F, F))$, $\text{UPLAPFSOLVER}(\mathbf{b}_F)$ returns an \mathbf{x}_F such that $\mathbf{L}_1^{\text{up}}(F, F)\mathbf{x}_F = \mathbf{b}_F$ in time $O(nr^{1/3})$.*

By our choice of F , the matrix $\mathbf{L}_1^{\text{up}}(F, F)$ can be written as a block diagonal matrix where each block corresponds to a region of \mathcal{K} w.r.t. the r -hollowing \mathcal{T} . Since each region is a 3-complex in which every tetrahedron has an aspect ratio $O(1)$, we can construct the solver UPLAPFSOLVER by Nested Dissection [MT90]. However, since each row or column of $\mathbf{L}_1^{\text{up}}(F, F)$ corresponds to an edge in \mathcal{K} , we need to turn the good *vertex separators* in [MT90] into good *edge separators* for regions of \mathcal{K} . We prove Lemma 6.5.7 in Section 6.7.1.

Lemma 6.5.8 (Solver for the Schur Complement). *Let \mathcal{K} be a pure 3-complex embedded in \mathbb{R}^3 and composed of n stable simplexes. Let \mathcal{T} be an r -hollowing of \mathcal{K} , and let C be the set of boundary edges of each region of \mathcal{K} w.r.t. \mathcal{T} . Then, with a pre-processing time $O(nr + n^2r^{-2/3})$ there exists a solver $\text{SCHURSOLVER}(\cdot, \cdot)$ such that for any $\mathbf{h} \in \text{Im}(\text{Sc}(\mathbf{L}_1^{\text{up}})_C)$ and $\delta > 0$, $\text{SCHURSOLVER}(\mathbf{h}, \delta)$ returns an $\tilde{\mathbf{x}}_C$ such that $\|\text{Sc}(\mathbf{L}_1^{\text{up}})_C \tilde{\mathbf{x}}_C - \mathbf{h}\|_2 \leq \delta \|\mathbf{h}\|_2$ in time $\tilde{O}(nr^{5/6} + n^{4/3}r^{5/18})$.*

Our solver SCHURSOLVER is based on the Preconditioned Conjugate Gradient (PCG) with the preconditioner $\mathbf{L}_{1,\mathcal{T}}^{\text{up}}$, the up-Laplacian operator of \mathcal{T} . By Theorem 6.3.4 and Lemma 6.5.6, the number of PCG iterations is $\tilde{O}(\sqrt{r})$. In each PCG iteration, we solve the system in $\mathbf{L}_{1,\mathcal{T}}^{\text{up}}$ via Nested Dissection. We prove Lemma 6.5.8 in Section 6.7.2.

Given the above lemmas, we prove Lemma 6.5.4.

Proof of Lemma 6.5.4. The correctness of Algorithm 2 is by Lemma 6.5.5. By Lemma 6.5.7 and 6.5.8, the total runtime of the algorithm is

$$\tilde{O}(nr^{5/6} + n^{4/3}r^{5/18} + nr + n^2r^{-2/3}).$$

□

6.5.3 Projection onto the Image of Up-Laplacian

As the first Betti number of \mathcal{K} can be arbitrary, the approximate projection operators for the up 1-Laplacian provided in [Coh+14a; Bla+22; BN22] are not applicable here. Our approximate projection operator follows a similar approach to our up-Laplacian solver, which is based on an incomplete Nested Dissection for *triangles*, instead of edges.

Lemma 6.5.9 (Up-Projection Operator). *Let \mathcal{K} be a pure 3-complex embedded in \mathbb{R}^3 and composed of n stable simplexes. Suppose we are given an r -hollowing for \mathcal{K} . Then, for any $\epsilon > 0$, there exists an operator $\tilde{\Pi}_1^{\text{up}}$ such that for any \mathbf{b} , $\tilde{\Pi}_1^{\text{up}}\mathbf{b}$ is in the image of \mathbf{L}_1^{up} , and*

$$\|\tilde{\Pi}_1^{\text{up}}\mathbf{b} - \Pi_1^{\text{up}}\mathbf{b}\|_2 \leq \epsilon \|\Pi_1^{\text{up}}\mathbf{b}\|_2.$$

In addition, $\tilde{\Pi}_1^{\text{up}}\mathbf{b}$ can be computed in time $O(nr + n^{4/3}r^{5/18} \log(n/\epsilon) + n^2r^{-2/3})$.

A detailed proof of Lemma 6.5.9 is deferred to Section 6.8. The subsequent Lemma offers a helpful formula of Π_1^{up} , the orthogonal projection matrix onto the image of \mathbf{L}_1^{up} .

Lemma 6.5.10. *Let \mathcal{K} be a simplicial complex with boundary operator ∂_2 . For any partition $F \cup C$ of the 2-simplexes of \mathcal{K} , the orthogonal projection Π_1^{up} for \mathcal{K} can be decomposed as*

$$\Pi_1^{up} = \Pi_{\text{Im}(\partial_2(:,F))} + \Pi_{\text{Ker}(\partial_2^\top(F,:))} \partial_2(:,C) (Sc(\mathbf{L}_2^{down})_C)^\dagger \partial_2^\top(C,:) \Pi_{\text{Ker}(\partial_2^\top(F,:))},$$

where \mathbf{L}_2^{down} is the down 2-Laplacian.

Once more, an r -hollowing offers a natural partition of the triangles within \mathcal{K} . We assign all the interior triangles to F and all the boundary triangles to C . As such, Nested Dissection can be utilized to compute $\Pi_{\text{Im}(\partial_2(:,F))}$ and $\Pi_{\text{Ker}(\partial_2^\top(F,:))}$, and PCG to solve the system in the Schur complement. The primary technical challenge arises when bounding the relative condition number of the Schur complement and the preconditioner, which requires a different approach.

6.6 Solver for Down-Laplacian

This section shows how to solve $\mathbf{L}_1^{down} \mathbf{x}^{down} = \mathbf{b}$ for any $\mathbf{b} \in \text{Im}(\mathbf{L}_1^{down})$ in linear time and proves Lemma 6.5.3. Recall that $\mathbf{L}_1^{down} = \partial_1^\top \mathbf{W}_0 \partial_1$. Our down-Laplacian solver works for *any* simplicial complexes and returns a solution *without error*. Our approach is a slight modification of the down-Laplacian solver in [Coh+14a] to incorporate the vertex weights \mathbf{W}_0 . Specifically, we compute \mathbf{x}^{down} by three steps: solve $\partial_1^\top \mathbf{W}_0^{1/2} \mathbf{y} = \mathbf{b}$, project \mathbf{y} onto $\text{Im}(\mathbf{W}_0^{1/2} \partial_1)$ and get a new vector \mathbf{y}' , then solve $\mathbf{W}_0^{1/2} \partial_1 \mathbf{x}^{down} = \mathbf{y}'$. The first and the last steps can be solved by the approach in Lemma 4.2 of [Coh+14a] (stated below), and the second step can be explicitly solved by utilizing that ∂_1 is a vertex-edge incidence matrix of an oriented graph.

Lemma 6.6.1 (Restatement of Lemma 4.2 of [Coh+14a]). *Given any $\mathbf{b}_1 \in \text{Im}(\partial_1^\top)$ (respectively, $\mathbf{b}_0 \in \text{Im}(\partial_1)$), there is a linear operator $\partial_1^{+\top}$ such that $\partial_1^\top \partial_1^{+\top} \mathbf{b}_1 = \mathbf{b}_1$ (respectively, $\partial_1^+ = (\partial_1^{+\top})^\top$ such that $\partial_1 \partial_1^+ \mathbf{b}_0 = \mathbf{b}_0$). In addition, we can compute $\partial_1^{+\top} \mathbf{b}_1$ (respectively, $\partial_1^+ \mathbf{b}_0$) in linear time.*

We remark that the operators ∂_1^+ and $\partial_1^{+\top}$ in Lemma 6.6.1 are not necessary to be the pseudo-inverses of ∂_1 and ∂_1^\top .

Claim 6.6.2 explicitly characterizes the image of $\mathbf{W}_0^{1/2} \partial_1$.

Claim 6.6.2. *Let \mathcal{K} be a simplicial complex whose 1-skeleton is connected, and let v be the number of vertices in \mathcal{K} . Let $\mathbf{W}_0 = \text{diag}(w_1, \dots, w_v)$ where $w_1, \dots, w_v > 0$. Then,*

$$\text{Ker}(\partial_1^\top \mathbf{W}_0^{1/2}) = \text{span}\{\mathbf{u}\}, \text{ where } \mathbf{u} = \left(\frac{1}{\sqrt{w_1}}, \dots, \frac{1}{\sqrt{w_v}} \right)^\top.$$

Proof. Let $\dim(\mathcal{V})$ be the dimension of a space \mathcal{V} . Since all the diagonals of \mathbf{W}_0 are positive,

$$\dim(\text{Ker}(\partial_1^\top \mathbf{W}_0^{1/2})) = \dim(\text{Ker}(\partial_1^\top)) = v - \dim(\text{Im}(\partial_1)) = 1,$$

where the last equality holds since the 1-skeleton of \mathcal{K} is connected. We can check that

$$\partial_1^\top \mathbf{W}_0^{1/2} \mathbf{u} = \partial_1^\top \mathbf{1} = \mathbf{0}.$$

Here, $\mathbf{1}$ is the all-one vector, and the second equality holds since ∂_1 is the vertex-edge incidence matrix of an oriented (weakly) connected graph. Thus, the statement holds. \square

We construct our down-Laplacian solver by combining Lemma 6.6.1 and Claim 6.6.2. We remark that this down-Laplacian solver applies to any simplicial complex.

Proof of Lemma 6.5.3. Without loss of generality, we assume the 1-skeleton of \mathcal{K} is connected. Otherwise, we can write $\mathbf{L}_1^{\text{down}}$ as a block diagonal matrix, each corresponding to a connected component of the 1-skeleton, and we reduce solving a system in $\mathbf{L}_1^{\text{down}}$ into solving several smaller down-Laplacian systems.

Let $\partial_1^+, \partial_1^{+\top}$ be the linear operators in Lemma 6.6.1 for \mathcal{K} , and let \mathbf{u} be the vector in Claim 6.6.2. We define

$$\mathbf{Z}_1^{\text{down}} \stackrel{\text{def}}{=} \partial_1^+ \mathbf{W}_0^{-1/2} \left(\mathbf{I} - \frac{\mathbf{u}\mathbf{u}^\top}{\|\mathbf{u}\|_2^2} \right) \mathbf{W}_0^{-1/2} \partial_1^{+\top}.$$

We can compute $\mathbf{Z}_1^{\text{down}} \mathbf{b}$ in linear time. In addition, we define

$$\mathbf{y} = \partial_1^{+\top} \mathbf{b}, \mathbf{z} = \mathbf{W}_0^{-1/2} \mathbf{y}, \mathbf{z}_1 = \left(\mathbf{I} - \frac{\mathbf{u}\mathbf{u}^\top}{\|\mathbf{u}\|_2^2} \right) \mathbf{z}, \text{ and } \mathbf{x} = \partial_1^+ \mathbf{W}_0^{-1/2} \mathbf{z}_1.$$

Then, $\mathbf{z}_1 \in \text{Im}(\mathbf{W}_0^{1/2} \partial_1)$ and $\mathbf{z} - \mathbf{z}_1 \in \text{Ker}(\partial_1^\top \mathbf{W}_0^{1/2})$.

$$\begin{aligned} \mathbf{L}_1^{\text{down}} \mathbf{Z}_1^{\text{down}} \mathbf{b} &= \partial_1^\top \mathbf{W}_0 \partial_1 \mathbf{x} = \partial_1^\top \mathbf{W}_0 \partial_1 \partial_1^+ \mathbf{W}_0^{-1/2} \mathbf{z}_1 \\ &= \partial_1^\top \mathbf{W}_0 \mathbf{W}_0^{-1/2} \mathbf{z}_1 && \text{since } \mathbf{W}_0^{-1/2} \mathbf{z}_1 \in \text{Im}(\partial_1) \\ &= \partial_1^\top \mathbf{W}_0^{1/2} \mathbf{z} && \text{since } \mathbf{z} - \mathbf{z}_1 \in \text{Ker}(\partial_1^\top \mathbf{W}_0^{1/2}) \\ &= \partial_1^\top \mathbf{y} = \mathbf{b}. \end{aligned}$$

□

6.7 Solver for Up-Laplacian

In this section, we will prove the two key lemmas, Lemma 6.5.7 and Lemma 6.5.8, for building the first up-Laplacian solver.

6.7.1 Solver for $\mathbf{L}_1^{\text{up}}(F, F)$

In this section, we construct an efficient solver $\text{UPLAPFSOLVER}(\mathbf{b})$ that returns an \mathbf{x} such that $\mathbf{L}_1^{\text{up}}(F, F) \mathbf{x} = \mathbf{b}$ for any $\mathbf{b} \in \text{Im}(\mathbf{L}_1^{\text{up}}(F, F))$ and prove Lemma 6.5.7.

We first show that the interiors of different regions are “disjoint” in the sense that $\mathbf{L}_1^{\text{up}}(F, F)$ can be written as a block diagonal matrix where each diagonal block corresponds to the interior of a region after proper row and column permutation. By definition, we can write \mathbf{L}_1^{up} as the sum of rank-1 matrices that each corresponds to a triangle in \mathcal{K} :

$$\mathbf{L}_1^{\text{up}} = \partial_2 \mathbf{W}_2 \partial_2^\top = \sum_{\sigma: \text{triangle in } \mathcal{K}} \mathbf{W}_2(\sigma, \sigma) \cdot \partial_2(:, \sigma) \partial_2(:, \sigma)^\top.$$

For any two edges e_1, e_2 , $\mathbf{L}_1^{\text{up}}(e_1, e_2) = 0$ if and only if no triangle in \mathcal{K} contains both e_1, e_2 . We say such e_1 and e_2 are Δ -disjoint. The following claim shows that interior edges from different regions are Δ -disjoint.

Claim 6.7.1. *Let \mathcal{K} be a stable 3-complex and \mathcal{T} be an r -hollowing of \mathcal{K} . Let R_1, R_2 be two different regions of \mathcal{K} w.r.t. \mathcal{T} , and let e_1 be an interior edge in R_1 and e_2 an interior edge in R_2 . Then, e_1 and e_2 are Δ -disjoint.*

Proof. Assume by contradiction there exists a triangle $\sigma \in \mathcal{K}$ that contains both e_1, e_2 . Let e_3 be the third edge in σ . Since e_1, e_2 are interior edges of different regions R_1, R_2 , e_3 must cross the boundary of R_1, R_2 . This contradicts the fact that regions can only intersect on their boundaries. Thus, e_1, e_2 must be Δ -disjoint. \square

By the above Claim 6.7.1, computing an \mathbf{x} such that $\mathbf{L}_1^{\text{up}}(F, F)\mathbf{x} = \mathbf{b}$ reduces to computing a solution for each diagonal block submatrix per region. For this reason, Lemma 6.5.7 is a corollary of the following Lemma 6.7.2.

Lemma 6.7.2. *Let \mathcal{X} be a 3-complex with $O(r)$ simplexes embedded in \mathbb{R}^3 such that (1) each tetrahedron of \mathcal{X} has $O(1)$ aspect ratio, and (2) \mathcal{X} has $O(r^{2/3})$ exterior simplexes. Let F be the set of interior edges of \mathcal{X} , and let $\mathbf{L}_{1,\mathcal{X}}^{\text{up}}$ be the up-Laplacian of \mathcal{X} and $\mathbf{M} \stackrel{\text{def}}{=} \mathbf{L}_{1,\mathcal{X}}^{\text{up}}(F, F)$. Then, there is a permutation matrix \mathbf{P} and a lower triangular matrix \mathbf{L} with $O(r^{4/3})$ non-zeros such that*

$$\mathbf{M} = \mathbf{P}\mathbf{L}\mathbf{L}^\top\mathbf{P}^\top. \quad (6.2)$$

In addition, we can find such \mathbf{P} and \mathbf{L} in time $O(r^2)$. Given the above factorization, for any $\mathbf{b} \in \text{Im}(\mathbf{M})$, we can compute an \mathbf{x} such that $\mathbf{M}\mathbf{x} = \mathbf{b}$ in $O(r^{4/3})$ time.

The factorization in Equation (6.2) is called *Cholesky factorization*. We will utilize the geometric structures of \mathcal{X} to find a sparse Cholesky factorization efficiently. Miller and Thurston [MT90] studied vertex separators of the 1-skeleton of a 3-complex \mathcal{X} in which each tetrahedron has $O(1)$ aspect ratio. A subset of vertices S of a graph over v vertices δ -separates if the remaining vertices can be partitioned into two sets A, B such that there are no edges between A and B , and $|A|, |B| \leq \delta v$. The set S is an $f(v)$ -separator if there exists a constant $\delta < 1$ such that S δ -separates and $|S| \leq f(v)$. These separators can be incorporated with Nested Dissection to efficiently compute a sparse Cholesky factorization of a matrix in which the non-zero structure encodes the 1-skeleton of \mathcal{X} .

Theorem 6.7.3 (Vertex Separator for a 3-Complex, Theorem 1.5 of [MT90] and [Mil+98]). *Let \mathcal{X} be a 3-complex in \mathbb{R}^3 in which each tetrahedron has $O(1)$ aspect ratio, and suppose \mathcal{X} has t tetrahedra and \bar{v} exterior vertices. Then, the 1-skeleton of \mathcal{X} has a $O(t^{2/3} + \bar{v})$ -separator that $4/5$ -separates \mathcal{X} . In addition, such a separator can be found in linear time.*

We need a slightly modified version of Theorem 6.7.3 to apply to our matrix \mathbf{M} .

Corollary 6.7.4 (Edge Separator for a 3-Complex). *Let \mathcal{X} be a 3-complex that satisfies the requirements in Theorem 6.7.3. In addition, the 1-skeleton of \mathcal{X} is connected. Let m be the number of edges in \mathcal{X} . Then, there exists an algorithm that removes $O(t^{2/3} + \bar{v})$ edges in linear time so that the remaining edges can be partitioned into two Δ -disjoint sets, each of size at most cm for some constant $c < 1$.*

Proof. Let $v = O(r)$ be the number of vertices in \mathcal{X} . By Lemma 6.3.1, $v = \Theta(m)$. Let S be the set of $O(t^{2/3} + \bar{v})$ vertices that $4/5$ -separates \mathcal{X} (by Theorem 6.7.3), and let A, B

be two disjoint sets of the remaining vertices after removing S such that $|A|, |B| \leq 4v/5$. Let E be the set of edges in \mathcal{X} incident to some vertex in S . Since each vertex has $O(1)$ degree by Lemma 6.3.1, $|E| = O(|S|) = O(t^{2/3} + \bar{v})$. Now, we remove edges in E from the edges of \mathcal{X} . Let E_A be the set of the remaining edges incident to a vertex in A , and let E_B be the set of the remaining edges incident to a vertex in B . Since A, B are disjoint, E_A and E_B are Δ -disjoint. We then show that $|E_A|, |E_B| \leq cm$ for some constant $c < 1$. Let E'_A be the set of edges in \mathcal{X} that are incident to some vertex in A . Since the 1-skeleton of \mathcal{X} is a connected graph, $|E'_A| \geq |A| \geq (1 - o(1))\frac{v}{5}$. Besides, $|E'_A| = |E_A| + |E''_A|$, where E''_A contains edges with one endpoint in A and the other in S . Since $|E''_A| = O(t^{2/3} + \bar{v})$, we know $|E_A| \geq (1 - o(1))\frac{v}{5} > c'm$ for some constant c' . Since E_A and E_B are disjoint, we know $|E_B| \leq m - |E_A| \leq (1 - c')m$. By symmetry, we have $|E_A| \leq (1 - c')m$. \square

We need the following theorem about Nested Dissection from [LRT79].

Theorem 6.7.5 (Nested Dissection, Theorem 6 of [LRT79]). *Let \mathcal{G} be any class of graphs closed under subgraph on which a v^α -separator exists for $\alpha > 1/2$. Let $\mathbf{A} \in \mathbb{R}^{v \times v}$ be symmetric and positive definite (that is, all eigenvalues of \mathbf{A} are positive). Let G_A be a graph over vertices $\{1, \dots, v\}$ where vertices i, j are connected if and only if $\mathbf{A}(i, j) \neq 0$. If $G_A \in \mathcal{G}$, then we can find a permutation matrix \mathbf{P} and a lower triangular matrix \mathbf{L} with $O(v^{2\alpha})$ non-zeros in $O(v^{3\alpha})$ time such that $\mathbf{A} = \mathbf{P}\mathbf{L}\mathbf{L}^\top\mathbf{P}^\top$.*

Combining Corollary 6.7.4 and Theorem 6.7.5, we prove Lemma 6.7.2.

Proof of Lemma 6.7.2. Our approach for finding a sparse Cholesky factorization is the same as Section 4 of [MT90]. Specifically, we find a family of separators for \mathcal{X} by recursively applying Corollary 6.7.4 to the remaining sets of edges A, B . Since the separator size grows as a function of the exterior vertices, at the top level of the recursion, we include all the exterior edges $O(r^{2/3})$ in the root separator, which only increases the root separator by a constant factor. In each remaining recursion step, suppose we want to separate \mathcal{Y} , a sub-complex of \mathcal{X} ; we let $\bar{\mathcal{Y}}$ be the complex consisting of \mathcal{Y} and all the simplexes in \mathcal{X} that contain a vertex in \mathcal{Y} . Then we apply a slightly modified version of Corollary 6.7.4 (obtained by a slightly modified Theorem 6.7.3) to $\bar{\mathcal{Y}}$ in which we only separate \mathcal{Y} . Here, the exterior-vertex term $O(\bar{v})$ in the size of the separator can be ignored since all the exterior edges of $\bar{\mathcal{Y}}$ have already been included in upper-level separators; we also use the fact that the number of exterior vertices and the number of exterior edges are equal up to a constant factor. This separator family provides an elimination ordering for the edges of \mathcal{K} , which is the permutation matrix \mathbf{P} in Equation (6.2), and the ordering uniquely determines the matrix \mathbf{L} . By Theorem 6.7.5, \mathbf{L} has $O(r^{4/3})$ non-zeros, and \mathbf{P}, \mathbf{L} can be found in time $O(r^2)$.

One issue left is that \mathbf{M} in Lemma 6.7.2 is positive semidefinite but *not* positive definite. During the process of numeric factorization, the first row and column of some Schur complements are all-zero. We simply ignore these zeros and proceed (Ref: Chapter 4.2.8 of [GV96]). This produces a Cholesky factorization of $\mathbf{P}^\top\mathbf{M}\mathbf{P} = \mathbf{L}\mathbf{L}^\top$ such that only k columns of \mathbf{L} are non-zero, where k is the rank of \mathbf{M} . We permute the rows and the columns of \mathbf{L} by multiplying permutation matrices $\mathbf{P}_1, \mathbf{P}_2$ so that

$$\mathbf{P}_1\mathbf{L}\mathbf{P}_2 = \begin{bmatrix} \mathbf{T} & \mathbf{0} \end{bmatrix} \stackrel{\text{def}}{=} \begin{bmatrix} \mathbf{H}_1 & \mathbf{0} \\ \mathbf{H}_2 & \mathbf{0} \end{bmatrix},$$

where \mathbf{H}_1 is a $k \times k$ non-singular lower-triangular matrix. Then, solving $\mathbf{M}\mathbf{x} = \mathbf{b}$ is equivalent to solving

$$\mathbf{T}\mathbf{T}^\top \mathbf{z} = \mathbf{P}_1 \mathbf{P}^\top \mathbf{b} \stackrel{\text{def}}{=} \mathbf{f}, \quad \mathbf{P}_1 \mathbf{P}^\top \mathbf{x} = \mathbf{z}.$$

We solve the first system by solving $\mathbf{T}\mathbf{y} = \mathbf{f}$ and $\mathbf{T}^\top \mathbf{z} = \mathbf{y}$. We let \mathbf{y} satisfy $\mathbf{H}_1 \mathbf{y} = \mathbf{f}[1 : k]$. Since \mathbf{H}_1 has full rank and $O(r^{4/3})$ non-zeros, such a \mathbf{y} exists and can be found in $O(r^{4/3})$ time. Since $\mathbf{y} \in \text{Im}(\mathbf{T})$, we know $\mathbf{T}\mathbf{y} = \mathbf{f}$. Then we let $\mathbf{z}[k+1 : v] = 0$, where v is the number of vertices in \mathcal{X} , and we let $\mathbf{H}_1^\top \mathbf{z}[1 : k] = \mathbf{y}$. Again, such \mathbf{z} exists and can be found in $O(r^{4/3})$ time. Finally, we compute $\mathbf{x} = \mathbf{P}\mathbf{P}_1^\top \mathbf{z}$ in linear time. \square

Proof of Lemma 6.5.7. We apply Lemma 6.7.2 to each diagonal block matrix of $\mathbf{L}_1^{\text{up}}(F, F)$ to compute a sparse Cholesky factorization as in Equation (6.2) in time

$$O\left(\frac{n}{r} \cdot r^2\right) = O(nr).$$

Given this Cholesky factorization, we can solve a system in $\mathbf{L}_1^{\text{up}}(F, F)$ in time

$$O\left(\frac{n}{r} \cdot r^{4/3}\right) = O(nr^{1/3}).$$

\square

6.7.2 Solver for the Schur Complement

In this section, we establish a fast approximate solver for $\text{Sc}(\mathbf{L}_1^{\text{up}})_C$ and prove Lemma 6.5.8. Recall that C contains all the edges in \mathcal{T} (an r -hollowing of \mathcal{K}), and $\text{Sc}(\mathbf{L}_1^{\text{up}})_C$ is the Schur complement of the up-Laplacian operator of \mathcal{K} onto C . The idea is to run the Preconditioned Conjugate Gradient (PCG) for systems in $\text{Sc}(\mathbf{L}_1^{\text{up}})_C$ with preconditioner $\mathbf{L}_{1,\mathcal{T}}^{\text{up}} = \partial_{2,\mathcal{T}} \partial_{2,\mathcal{T}}^\top$, which is the first up-Laplacian operator of \mathcal{T} . By Theorem 6.3.4, the number of PCG iterations is $O(\sqrt{\kappa} \log(\kappa/\epsilon))$ where $\kappa = \kappa(\text{Sc}(\mathbf{L}_1^{\text{up}})_C, \mathbf{L}_{1,\mathcal{T}}^{\text{up}})$ is the relative condition number and ϵ is the error parameter; in each PCG iteration, we need to solve a system in $\mathbf{L}_{1,\mathcal{T}}^{\text{up}}$, multiply $\text{Sc}(\mathbf{L}_1^{\text{up}})_C$ with $O(1)$ vectors, and implement $O(1)$ vector operations. In Section 6.7.2, we upper bound the relative condition number. In Section 6.7.2, we prove Lemma 6.5.8 about the solver for Schur complement.

We will need the following observation.

Claim 6.7.6. *Let \mathcal{X} be a simplicial complex. Changing the orientations of the triangles in \mathcal{X} does not change its first up-Laplacian operator.*

Proof. Let $\partial_{2,\mathcal{X}}$ be the second boundary operator of \mathcal{X} , $\mathbf{W}_{2,\mathcal{X}}$ the diagonal matrix for the triangle weights, and $\mathbf{L}_{1,\mathcal{X}}^{\text{up}}$ the first up-Laplacian. Changing the orientations of the triangles in \mathcal{X} corresponds to multiplying a ± 1 diagonal matrix \mathbf{X} to the right of $\partial_{2,\mathcal{X}}$. Observe

$$\partial_{2,\mathcal{X}}(\mathbf{X} \mathbf{W}_{2,\mathcal{X}} \mathbf{X}) \partial_{2,\mathcal{X}}^\top = \partial_{2,\mathcal{X}} \mathbf{W}_{2,\mathcal{X}} \partial_{2,\mathcal{X}}^\top = \mathbf{L}_{1,\mathcal{X}}^{\text{up}}.$$

Thus the statement holds. \square

Preconditioning the Schur Complement

We will upper bound the relative condition number of $\text{Sc}(\mathbf{L}_1^{\text{up}})_C$ and $\mathbf{L}_{1,\mathcal{T}}^{\text{up}}$ and prove will Lemma 6.5.6.

We decompose the r -hollowing \mathcal{T} into two parts. Recall that the boundary of each region triangulates a spherical shell in \mathbb{R}^3 . Let \mathcal{B}_1 be a 2-complex of the union of all the inner spheres of the boundaries, and let \mathcal{B}_2 be a 2-complex of the rest of the boundaries. Let C_1 be the set of edges of \mathcal{B}_1 . Let $\widehat{\mathbf{L}}_1^{\text{up}}$ be the first up-Laplacian of the union of \mathcal{B}_1 and all the interior simplexes in all the regions. Then,

$$\mathbf{L}_{1,\mathcal{T}}^{\text{up}} = \mathbf{L}_{1,\mathcal{B}_1}^{\text{up}} + \mathbf{L}_{1,\mathcal{B}_2}^{\text{up}} \text{ and } \text{Sc}(\mathbf{L}_1^{\text{up}})_C = \text{Sc}(\widehat{\mathbf{L}}_1^{\text{up}})_{C_1} + \mathbf{L}_{1,\mathcal{B}_2}^{\text{up}}.$$

Claim 6.7.7. *If $\mathbf{L}_{1,\mathcal{B}_1}^{\text{up}} \preceq \text{Sc}(\widehat{\mathbf{L}}_1^{\text{up}})_{C_1} \preceq \alpha \mathbf{L}_{1,\mathcal{B}_1}^{\text{up}}$, where $\alpha > 1$, then*

$$\mathbf{L}_{1,\mathcal{T}}^{\text{up}} \preceq \text{Sc}(\mathbf{L}_1^{\text{up}})_C \preceq \alpha \mathbf{L}_{1,\mathcal{T}}^{\text{up}}.$$

Proof. We have

$$\begin{aligned} \text{Sc}(\mathbf{L}_1^{\text{up}})_C &\succcurlyeq \mathbf{L}_{1,\mathcal{B}_1}^{\text{up}} + \mathbf{L}_{1,\mathcal{B}_2}^{\text{up}} = \mathbf{L}_{1,\mathcal{T}}^{\text{up}}, \\ \text{Sc}(\mathbf{L}_1^{\text{up}})_C &\preceq \alpha \mathbf{L}_{1,\mathcal{B}_1}^{\text{up}} + \mathbf{L}_{1,\mathcal{B}_2}^{\text{up}} \preceq \alpha \mathbf{L}_{1,\mathcal{T}}^{\text{up}}. \end{aligned}$$

□

To prove Lemma 6.5.6, it suffices to show the following lemma.

Lemma 6.7.8.

$$\mathbf{L}_{1,\mathcal{B}_1}^{\text{up}} \preceq \text{Sc}(\widehat{\mathbf{L}}_1^{\text{up}})_{C_1} \preceq O(r) \mathbf{L}_{1,\mathcal{B}_1}^{\text{up}}.$$

The first inequality of Lemma 6.7.8 follows immediately from a well-known fact about Schur complements, stated in Fact 2.3.7. For completeness, we restate it below.

Fact 6.7.9. *Let \mathbf{A} be a symmetric and PSD matrix:*

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}(F, F) & \mathbf{A}(F, C) \\ \mathbf{A}(C, F) & \mathbf{A}(C, C) \end{bmatrix}.$$

Let $\text{Sc}(\mathbf{A})_C = \mathbf{A}(C, C) - \mathbf{A}(C, F)\mathbf{A}(F, F)^{\dagger}\mathbf{A}(F, C)$ be the Schur complement of \mathbf{A} onto C . Then, for any $\mathbf{x} \in \mathbb{R}^{|C|}$,

$$\min_{\mathbf{y} \in \mathbb{R}^{|F|}} \begin{bmatrix} \mathbf{y}^{\top} & \mathbf{x}^{\top} \end{bmatrix} \mathbf{A} \begin{bmatrix} \mathbf{y} \\ \mathbf{x} \end{bmatrix} = \mathbf{x}^{\top} \text{Sc}(\mathbf{A})_C \mathbf{x}.$$

As a corollary, $\text{Sc}(\mathbf{A})_C$ is symmetric and PSD.

In the rest of the section, we prove the second inequality in Lemma 6.7.8. Since both $\widehat{\mathbf{L}}_1^{\text{up}}$ and $\mathbf{L}_{1,\mathcal{B}_1}^{\text{up}}$ can be written as a block diagonal matrix where each block corresponds to a region w.r.t. \mathcal{T} (after proper row and column permutation), it suffices to show the inequality in Lemma 6.7.8 holds for each region, restated in the following lemma.

Lemma 6.7.10. *Consider an r -hollowing region. Let \mathcal{B} be the 2-complex of the inner sphere of the boundary, and let B be the set of edges in \mathcal{B} . Let \mathcal{X} be a 2-complex of the union of \mathcal{B} and the interior simplexes. Then,*

$$\text{Sc}(\mathbf{L}_{1,\mathcal{X}}^{\text{up}})_B \preceq O(r) \mathbf{L}_{1,\mathcal{B}}^{\text{up}}.$$

We first show that the images of $\text{Sc}(\mathbf{L}_{1,\mathcal{X}}^{\text{up}})_C$ and $\mathbf{L}_{1,\mathcal{B}}^{\text{up}}$ are equal.

Claim 6.7.11.

$$\text{Im}(\text{Sc}(\mathbf{L}_{1,\mathcal{X}}^{\text{up}})_B) = \text{Im}(\mathbf{L}_{1,\mathcal{B}}^{\text{up}}).$$

Proof. In the proof, we drop the subscript \mathcal{X} to simplify our notations by letting $\mathbf{L}_1^{\text{up}} = \mathbf{L}_{1,\mathcal{X}}^{\text{up}}$ and $\partial_2 = \partial_{2,\mathcal{X}}$. We let A be the set of edges in \mathcal{X} but not in B , and let V_A be the set of vertices in \mathcal{X} but not in \mathcal{B} and V_B the set of vertices in \mathcal{B} .

Since both the two matrices $\text{Sc}(\mathbf{L}_1^{\text{up}})_B$, $\mathbf{L}_{1,\mathcal{B}}^{\text{up}}$ are symmetric and PSD, the statement in the claim is equivalent to $\text{Ker}(\text{Sc}(\mathbf{L}_1^{\text{up}})_B) = \text{Ker}(\mathbf{L}_{1,\mathcal{B}}^{\text{up}})$. By Fact 6.7.9,

$$\mathbf{L}_{1,\mathcal{B}}^{\text{up}} \preceq \text{Sc}(\mathbf{L}_1^{\text{up}})_B \implies \text{Ker}(\text{Sc}(\mathbf{L}_1^{\text{up}})_B) \subseteq \text{Ker}(\mathbf{L}_{1,\mathcal{B}}^{\text{up}}).$$

It remains to show $\text{Ker}(\text{Sc}(\mathbf{L}_1^{\text{up}})_B) \supseteq \text{Ker}(\mathbf{L}_{1,\mathcal{B}}^{\text{up}})$.

Let \mathbf{x}_B be an arbitrary vector in $\text{Ker}(\mathbf{L}_{1,\mathcal{B}}^{\text{up}})$. We want to show $\mathbf{x}_B \in \text{Ker}(\text{Sc}(\mathbf{L}_1^{\text{up}})_B)$, that is, $\mathbf{x}_B^\top \text{Sc}(\mathbf{L}_1^{\text{up}})_B \mathbf{x}_B = 0$. By Fact 6.7.9, it suffices to show that there exists an $\mathbf{x} = \begin{bmatrix} \mathbf{x}_A \\ \mathbf{x}_B \end{bmatrix}$ such that $\mathbf{x}^\top \mathbf{L}_1^{\text{up}} \mathbf{x} = 0$. This is equivalent to $\mathbf{x} \in \text{Ker}(\partial_2^\top)$. Suppose we add tetrahedra, triangles, and necessary edges to \mathcal{X} and get a new simplicial complex \mathcal{X}' so that \mathcal{B} is the boundary of \mathcal{X}' and \mathcal{X}' triangulates a 3-ball. Since operator ∂_2^\top maps the vector space of edges to the vector space of triangles and no new edges in \mathcal{X}' can appear in a triangle in \mathcal{X} , it suffices to show there exists an $\mathbf{x}' = \begin{bmatrix} \mathbf{x}'_A \\ \mathbf{x}_B \end{bmatrix}$ such that $\mathbf{x}' \in \text{Ker}(\partial_{2,\mathcal{X}'}^\top)$ (restricting \mathbf{x}'_A to \mathbf{x}_A gives $\mathbf{x} \in \text{Ker}(\partial_{2,\mathcal{X}}^\top)$). Since the first Betti number of \mathcal{X}' is zero,

$$\mathbf{x} \in \text{Ker}(\partial_{2,\mathcal{X}'}^\top) \iff \mathbf{x} \perp \text{Im}(\partial_{2,\mathcal{X}'}^\top) \iff \mathbf{x} \in \text{Im}(\partial_{1,\mathcal{X}'}^\top). \quad (6.3)$$

Let A' be the set of interior edges in \mathcal{X}' , and $V_{A'}$ the set of interior vertices in \mathcal{X}' . We can write

$$\partial_{1,\mathcal{X}'}^\top = \begin{bmatrix} \partial_{1,\mathcal{X}'}^\top(A', V_{A'}) & \partial_{1,\mathcal{X}'}^\top(A', V_B) \\ \mathbf{0} & \partial_{1,\mathcal{X}'}^\top(B, V_B) \end{bmatrix},$$

where $\partial_{1,\mathcal{X}'}^\top(V_B, B) = \partial_{1,\mathcal{B}}^\top$. Since $\mathbf{x}_B \in \text{Ker}(\partial_{2,\mathcal{B}}^\top)$ and the first Betti number of \mathcal{B} is 0, by an argument similar to Equation (6.3), we have $\mathbf{x}_B \in \text{Im}(\partial_{1,\mathcal{B}}^\top)$, that is, $\mathbf{x}_B = \partial_{1,\mathcal{B}}^\top \mathbf{y}_B$ for some \mathbf{y}_B . Setting $\mathbf{x}'_A = \partial_{1,\mathcal{X}'}^\top(A', V_B) \mathbf{y}_B$, we have $\mathbf{x} \in \text{Im}(\partial_{1,\mathcal{X}'}^\top)$. \square

Claim 6.7.12.

$$\lambda_{\max}(\text{Sc}(\mathbf{L}_{1,\mathcal{X}}^{\text{up}})_B) = O(w_{\max}),$$

where w_{\max} is the maximum triangle weight in \mathcal{X} .

Proof. We drop the subscript \mathcal{X} in the proof to simplify our notations. The Courant-Fischer Minimax Theorem (see Theorem 4.5.7, or Theorem 8.1.2 of [GV96]) states that for any symmetric matrix \mathbf{A} ,

$$\lambda_{\max}(\mathbf{A}) = \max_{\mathbf{x}: \|\mathbf{x}\|_2=1} \mathbf{x}^\top \mathbf{A} \mathbf{x}.$$

Apply this theorem and Fact 6.7.9,

$$\lambda_{\max}(\text{Sc}(\mathbf{L}_1^{\text{up}})_B) = \max_{\mathbf{x}: \|\mathbf{x}\|_2=1} \mathbf{x}^\top \text{Sc}(\mathbf{L}_1^{\text{up}})_B \mathbf{x} \leq \max_{\mathbf{x}: \|\mathbf{x}\|_2=1} \begin{bmatrix} \mathbf{0}^\top & \mathbf{x}^\top \end{bmatrix} \mathbf{L}_1^{\text{up}} \begin{bmatrix} \mathbf{0} \\ \mathbf{x} \end{bmatrix} \leq \lambda_{\max}(\mathbf{L}_1^{\text{up}}).$$

Below, we bound $\lambda_{\max}(\mathbf{L}_1^{\text{up}})$:

$$\begin{aligned}\lambda_{\max}(\mathbf{L}_1^{\text{up}}) &= \max_{\mathbf{x}: \|\mathbf{x}\|_2=1} \mathbf{x}^\top \mathbf{L}_1^{\text{up}} \mathbf{x} = \max_{\mathbf{x}: \|\mathbf{x}\|_2=1} \sum_{\sigma: \text{triangle in } \mathcal{X}} \mathbf{W}_2(\sigma, \sigma) (\partial_2(:, \sigma)^\top \mathbf{x})^2 \\ &\leq 3w_{\max} \cdot \max_{\mathbf{x}: \|\mathbf{x}\|_2=1} \sum_{\substack{\sigma=[i,j,k]: \\ \text{triangle in } \mathcal{X}}} (\mathbf{x}(i)^2 + \mathbf{x}(j)^2 + \mathbf{x}(k)^2) = O(w_{\max}).\end{aligned}$$

The last inequality holds since each edge appears in at most $O(1)$ triangles by Lemma 6.3.1. Thus, $\lambda_{\max}(\text{Sc}(\mathbf{L}_1^{\text{up}})_B) = O(w_{\max})$. \square

One more piece we need is a lower bound for $\lambda_{\min}(\mathbf{L}_{1,\mathcal{B}}^{\text{up}})$, which will be established via eigenvalues of graph Laplacian matrices.

Theorem 6.7.13 (Section 4.2 of [Moh91]). *Let G be an unweighted graph over n vertices with diameter D , the length of the longest path in G . Let \mathbf{L}_G be the graph Laplacian matrix of G . Then,*

$$\lambda_{\min}(\mathbf{L}_G) \geq \frac{4}{nD}.$$

Claim 6.7.14.

$$\lambda_{\min}(\mathbf{L}_{1,\mathcal{B}}^{\text{up}}) = \Omega(w_{\min} \cdot r^{-1}),$$

where w_{\min} is the minimum triangle weight in \mathcal{B} .

Proof. We drop the subscript \mathcal{B} in the proof to simplify our notations. Again we decompose the matrix as a sum of rank-1 matrices for each triangle in \mathcal{B}' :

$$\mathbf{L}_1^{\text{up}} = \sum_{\sigma: \text{triangle in } \mathcal{B}} \mathbf{W}_2(\sigma, \sigma) \partial_2(:, \sigma) \partial_2(:, \sigma)^\top \succcurlyeq w_{\min} \cdot \partial_2 \partial_2^\top.$$

So,

$$\lambda_{\min}(\mathbf{L}_1^{\text{up}}) \geq w_{\min} \cdot \lambda_{\min}(\partial_2 \partial_2^\top) = w_{\min} \cdot \lambda_{\min}(\partial_2^\top \partial_2).$$

Since \mathcal{B} triangulates a two-sphere, each edge of \mathcal{B} appears in exactly two triangles. Since changing the orientations of the triangles in \mathcal{B} does not change \mathbf{L}_1^{up} (by Claim 6.7.6), we assume all the triangles in \mathcal{B} are oriented clockwise. Then, each column of ∂_2^\top has exactly one entry with value 1 and one entry -1 and all others 0. That is, $\partial_2^\top \partial_2$ is the Laplacian of the dual graph of \mathcal{B} : the vertices are the triangles in \mathcal{B} , and two vertices are adjacent if and only if the corresponding two triangles share a common edge. The dual graph has $O(r^{2/3})$ vertices and diameter $O(r^{1/3})$. By Theorem 6.7.13, $\lambda_{\min}(\partial_2^\top \partial_2) = \Omega(w_{\min} \cdot r^{-1})$. \square

Combining all the claims above, we prove Lemma 6.7.10.

Proof of Lemma 6.7.10. Let $\mathbf{\Pi}$ be the orthogonal projection matrix onto $\text{Im}(\mathbf{L}_{1,\mathcal{B}}^{\text{up}}) = \text{Im}(\text{Sc}(\mathbf{L}_{1,\mathcal{X}}^{\text{up}})_B)$ (by Claim 6.7.11). By Claim 6.7.12, $\text{Sc}(\mathbf{L}_{1,\mathcal{X}}^{\text{up}})_B \preccurlyeq O(w_{\max})\mathbf{\Pi}$. By Claim 6.7.14, $\mathbf{\Pi} \preccurlyeq O(\frac{r}{w_{\min}})\mathbf{L}_{1,\mathcal{B}}^{\text{up}}$. Combining all together, we have

$$\text{Sc}(\mathbf{L}_{1,\mathcal{X}}^{\text{up}})_B \preccurlyeq O(rU)\mathbf{L}_{1,\mathcal{B}}^{\text{up}},$$

where $U = \frac{w_{\max}}{w_{\min}} = O(1)$. \square

Proof of Lemma 6.5.8

By Theorem 6.3.4 and Lemma 6.5.6, the number of Preconditioned Conjugate Gradient (PCG) iterations is

$$\tilde{O}\left(\kappa\left(\text{Sc}(\mathbf{L}_1^{\text{up}})_C, \mathbf{L}_{1,\mathcal{T}}^{\text{up}}\right)^{1/2}\right) = \tilde{O}(r^{1/2}).$$

In each PCG iteration, we solve a system in $\mathbf{L}_{1,\mathcal{T}}^{\text{up}}$ and multiply $\text{Sc}(\mathbf{L}_1^{\text{up}})_C$ with $O(1)$ vectors. Recall

$$\text{Sc}(\mathbf{L}_1^{\text{up}})_C = \mathbf{L}_1^{\text{up}}(C, C) - \mathbf{L}_1^{\text{up}}(C, F)\mathbf{L}_1^{\text{up}}(F, F)^{\dagger}\mathbf{L}_1^{\text{up}}(F, C).$$

In our preprocessing, we compute a Cholesky factorization of $\mathbf{L}_1^{\text{up}}(F, F)$ in time $O(r^2 \cdot \frac{n}{r}) = O(nr)$; then, we can multiply $\text{Sc}(\mathbf{L}_1^{\text{up}})_C$ onto a vector in time $O(r^{4/3} \cdot \frac{n}{r}) = O(nr^{1/3})$. Similarly, we solve a system in $\mathbf{L}_{1,\mathcal{T}}^{\text{up}}$ by Nested Dissection. By our construction, \mathcal{T} has $O(r^{2/3} \cdot \frac{n}{r}) = O(nr^{-1/3})$ triangles. The Cholesky factorization runs in time $O(n^2 r^{-2/3})$, and each system solve runs in time $O(n^{4/3} r^{-4/9})$. So, the runtime per PCG iteration is

$$O\left(nr^{1/3} + n^{4/3}r^{-4/9}\right).$$

Therefore, the total time is

$$\tilde{O}\left(r^{1/2}\left(nr^{1/3} + n^{4/3}r^{-4/9}\right) + nr + n^2 r^{-2/3}\right) = \tilde{O}\left(nr + n^{4/3}r^{5/18} + n^2 r^{-2/3}\right).$$

This completes the proof of Lemma 6.5.8.

6.8 Projection onto the Image of Up-Laplacian

In this section, we will show how to approximately project a vector onto the image of the first up-Laplacian of a well-shaped 3-complex with a given r -hollowing and prove Lemma 6.5.9. The following lemma gives an explicit formula for the orthogonal projection onto the up-Laplacian.

Lemma 6.8.1. *Let \mathcal{K} be a simplicial complex with boundary operator ∂_2 . For any partition $F \cup C$ of the 2-simplexes of \mathcal{K} , the orthogonal projection Π_1^{up} for \mathcal{K} can be decomposed as*

$$\Pi_1^{\text{up}} = \Pi_{\text{Im}(\partial_2(\cdot, F))} + \Pi_{\text{Ker}(\partial_2^{\top}(F, \cdot))} \partial_2(\cdot, C) (\text{Sc}(\mathbf{L}_2^{\text{down}})_C)^{\dagger} \partial_2^{\top}(C, \cdot) \Pi_{\text{Ker}(\partial_2^{\top}(F, \cdot))}. \quad (6.4)$$

In addition, the second matrix on the right-hand side of the above equation is the orthogonal projection matrix onto the image of $\Pi_{\text{Ker}(\partial_2^{\top}(F, \cdot))} \partial_2(\cdot, C)$.

Proof. For any $\mathbf{b} \in \mathbb{R}^n$, we define

$$\begin{aligned} \mathbf{f}_F &= \partial_2^{\top}(F, \cdot) \mathbf{b}, \quad \mathbf{f}_C = \partial_2^{\top}(C, \cdot) \mathbf{b}, \\ \mathbf{h} &= \mathbf{f}_C - \mathbf{L}_2^{\text{down}}(C, F) (\mathbf{L}_2^{\text{down}}(F, F))^{\dagger} \mathbf{f}_F, \\ \mathbf{x}_C &= (\text{Sc}(\mathbf{L}_2^{\text{down}})_C)^{\dagger} \mathbf{h}, \\ \mathbf{x}_F &= (\mathbf{L}_2^{\text{down}}(F, F))^{\dagger} (\mathbf{f}_F - \mathbf{L}_2^{\text{down}}(F, C) \mathbf{x}_C). \end{aligned} \quad (6.5)$$

Let $\mathbf{x} = \begin{bmatrix} \mathbf{x}_C \\ \mathbf{x}_F \end{bmatrix}$. Applying Lemma 6.5.5 with $\delta = 0$, we have $\partial_2^{\top} \partial_2 \mathbf{x} = \partial_2^{\top} \mathbf{b}$. That is, \mathbf{x} can be written as $\mathbf{x} = \mathbf{x}_1 + \mathbf{x}_2$ where $\mathbf{x}_1 = (\partial_2^{\top} \partial_2)^{\dagger} \partial_2^{\top} \mathbf{b}$ and \mathbf{x}_2 is in $\text{Ker}(\partial_2)$. Then,

$$\partial_2 \mathbf{x} = \partial_2 \mathbf{x}_1 = \Pi_1^{\text{up}} \mathbf{b}.$$

By Equation (6.5),

$$\Pi_1^{\text{up}} \mathbf{b} = \partial_2 \mathbf{x} = \Pi_{\text{Im}(\partial_2(:, F))} \mathbf{b} + \Pi_{\text{Ker}(\partial_2^\top(F, :))} \partial_2(:, C) (\text{Sc}(\mathbf{L}_2^{\text{down}})_C)^\dagger \partial_2^\top(C, :) \Pi_{\text{Ker}(\partial_2^\top(F, :))} \mathbf{b}.$$

Thus, the equation in the statement holds. By Fact 2.3.10,

$$\text{Sc}(\mathbf{L}_2^{\text{down}})_C = \partial_2^\top(C, :) \Pi_{\text{Ker}(\partial_2^\top(F, :))} \partial_2(:, C).$$

The second matrix on the right-hand side of the equation in the statement is the orthogonal projection onto the image of $\Pi_{\text{Ker}(\partial_2^\top(F, :))} \partial_2(:, C)$. \square

To apply Π_1^{up} to a vector \mathbf{b} , we will need to (1) project \mathbf{b} onto $\text{Im}(\partial_2(:, F))$, (2) project \mathbf{b} onto $\text{Ker}(\partial_2^\top(F, :))$, and (3) solve a system in $\text{Sc}(\mathbf{L}_2^{\text{down}})_C$. We let F be the set of all the interior triangles of \mathcal{K} w.r.t. the given r -hollowing, and let C be the set of all the boundary triangles. Similar to our up-Laplacian solver, we will apply the Nested Dissection for the “ F ” part and the Preconditioned Conjugate Gradient for the Schur complement onto C .

With a slight modification of the edge separator for a 3-complex (Corollary 6.7.4), we obtain a triangle separator for a 3-complex as the corollary below.

Corollary 6.8.2 (Triangle Separator for a 3-Complex). *Let \mathcal{X} be a 3-complex that satisfies the requirements in Theorem 6.7.3 and has a connected 2-skeleton. Let p be the number of triangles in \mathcal{X} . Then, there exists an algorithm that removes $O(t^{2/3} + \bar{v})$ triangle in linear time. The remaining triangles can be divided into two sets, each containing at most cp triangles (where $c < 1$ is a constant), with no shared edges between the two sets.*

Proof. The proof is very similar to that of Corollary 6.7.4, by replacing edges with triangles. \square

Given a triangle separator algorithm for a 3-complex, we can obtain a solver for $\mathbf{L}_2^{\text{down}}(F, F) = \partial_2^\top(F, :) \partial_2(:, F)$ by Nested Dissection. Note $\mathbf{L}_2^{\text{down}}[\sigma_1, \sigma_2] \neq 0$ if and only if two triangles σ_1, σ_2 share a common edge.

Lemma 6.8.3. *Let \mathcal{X} be a 3-complex with $O(r)$ simplexes embedded in \mathbb{R}^3 such that (1) each tetrahedron of \mathcal{X} has $O(1)$ aspect ratio and (2) \mathcal{X} has $O(r^{2/3})$ exterior simplexes. Let F be the set of interior triangles of \mathcal{X} , and let $\mathbf{L}_{2, \mathcal{X}}^{\text{down}}$ be the second down-Laplacian of \mathcal{X} and $\mathbf{M} \stackrel{\text{def}}{=} \mathbf{L}_{2, \mathcal{X}}^{\text{down}}(F, F)$. Then, there is a permutation matrix \mathbf{P} and a lower triangular matrix \mathbf{L} with $O(r^{4/3})$ non-zeros such that $\mathbf{M} = \mathbf{P} \mathbf{L} \mathbf{L}^\top \mathbf{P}^\top$. In addition, we can find such \mathbf{P} and \mathbf{L} in time $O(r^2)$. Given the above factorization, for any $\mathbf{b} \in \text{Im}(\mathbf{M})$, we can compute an \mathbf{x} such that $\mathbf{M} \mathbf{x} = \mathbf{b}$ in $O(r^{4/3})$ time.*

Claim 6.8.4. *Given any 1-chain vector \mathbf{b} , we can compute $\Pi_{\text{Im}(\partial_2(:, F))} \mathbf{b}$ and $\Pi_{\text{Ker}(\partial_2^\top(F, :))} \mathbf{b}$ in time $O(nr)$.*

Proof. We have

$$\Pi_{\text{Im}(\partial_2(:, F))} \mathbf{b} = \partial_2(:, F) (\partial_2^\top(F, :) \partial_2(:, F))^\dagger \partial_2^\top(F, :) \mathbf{b}.$$

We can compute $\partial_2^\top(F, :) \mathbf{b}$ in time $O(n)$. To apply $(\partial_2^\top(F, :) \partial_2(:, F))^\dagger$, we compute a Cholesky factorization of $\partial_2^\top(F, :) \partial_2(:, F)$ in time $O(r^2 \cdot \frac{n}{r}) = O(nr)$ via Nested Dissection

(Lemma 6.8.3), where the Cholesky factorization has $O(r^{4/3} \cdot \frac{n}{r}) = O(nr^{1/3})$ non-zeros; given such a Cholesky factorization, we can apply $(\partial_2^\top(F, :) \partial_2(:, F))^\dagger$ to a vector in time $O(nr^{1/3})$. Then, multiplying $\partial_2(:, F)$ with a vector runs in linear time. So, the total runtime of computing $\Pi_{\text{Im}(\partial_2(:, F))} \mathbf{b}$ is $O(nr)$. Since $\Pi_{\text{Ker}(\partial_2^\top(F, :))} = \mathbf{I} - \Pi_{\text{Im}(\partial_2(:, F))}$, we can apply $\Pi_{\text{Ker}(\partial_2^\top(F, :))}$ to a vector in the same time up to a constant. \square

It remains to bound the runtime of (approximately) solving a system in the Schur complement $\text{Sc}(\mathbf{L}_2^{\text{down}})_C$. We run the Preconditioned Conjugate Gradient and precondition the Schur complement by the boundary itself $\partial_2^\top(C, :) \partial_2(:, C)$.

6.8.1 Preconditioning the Schur Complement

In this section, we prove the following lemma for the relative condition number of the preconditioner for the Schur complement.

Lemma 6.8.5.

$$\kappa(\text{Sc}(\mathbf{L}_2^{\text{down}})_C, \partial_2^\top(C, :) \partial_2(:, C)) = O(r).$$

We will bound the relative condition number for each region and then combine them to get Lemma 6.8.5. Recall that in each region of an r -hollowing, the boundary triangulates a spherical shell in \mathbb{R}^3 . We call the boundary triangles in a region containing an edge on the inner sphere of the boundary *the boundary layer*. Since each region boundary has its shell width of at least 5, boundary layers from different regions are disjoint. For the i th region, let \mathbf{S}_i be the Schur complement of the interior triangles and the triangles in the boundary layer of the region i onto its boundary layer. We can write the Schur complement

$$\text{Sc}(\mathbf{L}_2^{\text{down}})_C = \text{diag}(\mathbf{S}_1, \dots, \mathbf{S}_k) + \mathbf{A},$$

where \mathbf{A} is a PSD matrix. Similarly, let \mathbf{M}_i be the boundary layer of the region i . We write the boundary

$$\partial_2^\top(C, :) \partial_2(:, C) = \text{diag}(\mathbf{M}_1, \dots, \mathbf{M}_k) + \mathbf{A}.$$

By an argument similar to Claim 6.7.7, proving Lemma 6.8.5 reduces to proving the following lemma.

Lemma 6.8.6. *For each i ,*

$$\Omega(r^{-1}) \Pi_i \mathbf{M}_i \Pi_i \preceq \mathbf{S}_i \preceq \mathbf{M}_i.$$

In the rest of this subsection, we prove Lemma 6.8.6. We locally use ∂_2 for the boundary operator of the region i and drop the region index i in the rest of this subsection. We write

$$\partial_2^\top = \begin{bmatrix} \mathbf{B}_{\text{int}} \\ \mathbf{B}_{\text{bd}} \end{bmatrix},$$

where the rows of \mathbf{B}_{int} correspond to the interior triangles and the rows of \mathbf{B}_{bd} the boundary layer triangles. The Schur complement onto the boundary layer triangles, by Fact 2.3.10,

$$\mathbf{S} = \mathbf{B}_{\text{bd}} \Pi_{\text{Ker}(\mathbf{B}_{\text{int}})} \mathbf{B}_{\text{bd}}^\top.$$

Let $\mathbf{M} \stackrel{\text{def}}{=} \mathbf{B}_{\text{bd}} \mathbf{B}_{\text{bd}}^\top$, and let Π be the orthogonal projection onto the image of \mathbf{S} . Then, $\mathbf{S} \preceq \mathbf{M}$, and

$$\mathbf{S} \succcurlyeq \Pi \mathbf{M}^{1/2} \mathbf{M}^{\dagger/2} \mathbf{S} \mathbf{M}^{\dagger/2} \mathbf{M}^{1/2} \Pi \succcurlyeq \lambda_{\min}(\mathbf{M}^{\dagger/2} \mathbf{S} \mathbf{M}^{\dagger/2}) \cdot \Pi \mathbf{M} \Pi.$$

Claim 6.8.7.

$$\lambda_{\min} \left(\mathbf{M}^{\dagger/2} \mathbf{S} \mathbf{M}^{\dagger/2} \right) = \lambda_{\min} \left(\Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \Pi_{\text{Ker}(\mathbf{B}_{int})} \Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \right).$$

Proof. We take the singular value decomposition: $\mathbf{B}_{bd} = \mathbf{U} \mathbf{D} \mathbf{V}^\top$, where the columns of \mathbf{U} (resp., \mathbf{V}) form an orthonormal basis of $\text{Im}(\mathbf{B}_{bd})$ (resp., $\text{Im}(\mathbf{B}_{bd}^\top)$). Then,

$$\mathbf{M}^{\dagger/2} \mathbf{S} \mathbf{M}^{\dagger/2} = \mathbf{U} \mathbf{D}^{-1} \mathbf{U}^\top \mathbf{U} \mathbf{D} \mathbf{V}^\top \Pi_{\text{Ker}(\mathbf{B}_{int})} \mathbf{V} \mathbf{D} \mathbf{U}^\top \mathbf{U} \mathbf{D}^{-1} \mathbf{U}^\top = \mathbf{U} \mathbf{V}^\top \Pi_{\text{Ker}(\mathbf{B}_{int})} \mathbf{V} \mathbf{U}^\top.$$

Since \mathbf{U} 's columns are orthonormal,

$$\lambda_{\min}(\mathbf{M}^{\dagger/2} \mathbf{S} \mathbf{M}^{\dagger/2}) = \lambda_{\min}(\mathbf{V}^\top \Pi_{\text{Ker}(\mathbf{B}_{int})} \mathbf{V}).$$

We claim

$$\lambda_{\min}(\mathbf{V}^\top \Pi_{\text{Ker}(\mathbf{B}_{int})} \mathbf{V}) = \lambda_{\min}(\mathbf{V} \mathbf{V}^\top \Pi_{\text{Ker}(\mathbf{B}_{int})} \mathbf{V} \mathbf{V}^\top), \quad (6.6)$$

which implies the claim statement. Let λ be an eigenvalue of $\mathbf{V}^\top \Pi_{\text{Ker}(\mathbf{B}_{int})} \mathbf{V}$ with eigenvector \mathbf{u} . Then,

$$\begin{aligned} \mathbf{V}^\top \Pi_{\text{Ker}(\mathbf{B}_{int})} \mathbf{V} \mathbf{u} = \lambda \mathbf{u} &\implies \mathbf{V} \mathbf{V}^\top \Pi_{\text{Ker}(\mathbf{B}_{int})} \mathbf{V} \mathbf{V}^\top \mathbf{V} \mathbf{u} = \lambda \mathbf{V} \mathbf{u} \\ &\implies \Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \Pi_{\text{Ker}(\mathbf{B}_{int})} \Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \mathbf{V} \mathbf{u} = \lambda \mathbf{V} \mathbf{u}. \end{aligned}$$

That is, λ is an eigenvalue of $\Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \Pi_{\text{Ker}(\mathbf{B}_{int})} \Pi_{\text{Im}(\mathbf{B}_{bd}^\top)}$ with eigenvector $\mathbf{V} \mathbf{u}$. Let μ be an eigenvalue of $\Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \Pi_{\text{Ker}(\mathbf{B}_{int})} \Pi_{\text{Im}(\mathbf{B}_{bd}^\top)}$ with eigenvector \mathbf{w} . Then,

$$\Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \Pi_{\text{Ker}(\mathbf{B}_{int})} \Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \mathbf{w} = \mu \mathbf{w} \implies \mathbf{V}^\top \Pi_{\text{Ker}(\mathbf{B}_{int})} \mathbf{V} \mathbf{V}^\top \mathbf{w} = \mu \mathbf{V}^\top \mathbf{w}.$$

That is, μ is an eigenvalue of $\mathbf{V}^\top \Pi_{\text{Ker}(\mathbf{B}_{int})} \mathbf{V}$ with eigenvector $\mathbf{V}^\top \mathbf{w}$. So, Equation (6.6) holds. \square

Claim 6.8.8. *The image of $\Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \Pi_{\text{Ker}(\mathbf{B}_{int})} \Pi_{\text{Im}(\mathbf{B}_{bd}^\top)}$ is*

$$\mathcal{U} = \{\mathbf{u} \in \text{Im}(\mathbf{B}_{bd}^\top) : \mathbf{u} \perp \text{Im}(\mathbf{B}_{bd}) \cap \text{Im}(\mathbf{B}_{int}^\top)\}.$$

Proof. We first find an orthogonal basis of the kernel of $\Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \Pi_{\text{Ker}(\mathbf{B}_{int})} \Pi_{\text{Im}(\mathbf{B}_{bd}^\top)}$, which is $\mathcal{U}' \stackrel{\text{def}}{=} \{\mathbf{u} : \Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \Pi_{\text{Ker}(\mathbf{B}_{int})} \Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \mathbf{u} = \mathbf{0}\} = \{\mathbf{u} : \Pi_{\text{Ker}(\mathbf{B}_{int})} \Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \mathbf{u} = \mathbf{0}\}$. Clearly, $\text{Ker}(\mathbf{B}_{bd}) \subset \mathcal{U}'$. We consider $\mathbf{u} \in \text{Im}(\mathbf{B}_{bd}^\top)$. In this case, $\Pi_{\text{Ker}(\mathbf{B}_{int})} \Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \mathbf{u} = \Pi_{\text{Ker}(\mathbf{B}_{int})} \mathbf{u} = \mathbf{0}$ if and only if $\mathbf{u} \in \text{Im}(\mathbf{B}_{int}^\top)$. So, the claim statement holds. \square

Lemma 6.8.9.

$$\lambda_{\min} \left(\Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \Pi_{\text{Ker}(\mathbf{B}_{int})} \Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \right) = \Omega(r^{-1}).$$

Proof. By the Courant-Fischer min-max theorem and Claim 6.8.8,

$$\lambda_{\min} \left(\Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \Pi_{\text{Ker}(\mathbf{B}_{int})} \Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \right) = \min_{\mathbf{u} \in \mathcal{U}} \frac{\mathbf{u}^\top \Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \Pi_{\text{Ker}(\mathbf{B}_{int})} \Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \mathbf{u}}{\mathbf{u}^\top \mathbf{u}}, \quad (6.7)$$

where $\mathcal{U} = \{\mathbf{u} \in \text{Im}(\mathbf{B}_{bd}^\top) : \mathbf{u} \perp \text{Im}(\mathbf{B}_{bd}) \cap \text{Im}(\mathbf{B}_{int}^\top)\} \setminus \{\mathbf{0}\}$.

Recall that in each region, the boundary triangulates a spherical shell in \mathbb{R}^3 . We further decompose ∂_2 according to the inner sphere of the boundary:

$$\partial_2 = \begin{bmatrix} \mathbf{B}_{int}^\top & | & \mathbf{B}_{bd}^\top \end{bmatrix} = \begin{bmatrix} \mathbf{B}_{11} & \mathbf{B}_{12} & | & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_2 & | & \mathbf{B}_3 \\ \mathbf{0} & \mathbf{0} & | & \mathbf{B}_4 \end{bmatrix}. \quad (6.8)$$

Here, the blocks of the rows from top to bottom correspond to the interior edges, the boundary edges on the boundary inner sphere, and the other boundary edges, respectively; the blocks of the columns from left to right correspond to the interior triangles with only interior edges, the interior triangles with both interior and boundary edges, and the boundary triangles, respectively. The column in ∂_2 corresponding to a triangle whose edges are all on the boundary inner sphere can be written as a linear combination of the other boundary triangle columns. We remove these boundary inner sphere triangles so that every remaining triangle has at most one edge on the boundary inner sphere. Without loss of generality, we can orient and reorder the edges and the triangles so that

$$\mathbf{B}_2 = \text{diag}(\mathbf{1}^\top, \dots, \mathbf{1}^\top), \quad \mathbf{B}_3 = \text{diag}(\mathbf{1}^\top, \dots, \mathbf{1}^\top). \quad (6.9)$$

In addition, in each triangle with one edge on the boundary inner sphere, we let the other two edges point away from the inner sphere so that every column of \mathbf{B}_{12} and of \mathbf{B}_4 has exactly two non-zero entries with values 1 and -1 each.

We want to characterize $\text{Im}(\mathbf{B}_{int}^\top) \cap \text{Im}(\mathbf{B}_{bd}^\top)$. Let

$$\mathcal{V} \stackrel{\text{def}}{=} \left\{ \begin{bmatrix} \mathbf{0} \\ \mathbf{B}_3 \mathbf{x} \\ \mathbf{0} \end{bmatrix} : \mathbf{x} \in \text{Ker}(\mathbf{B}_4) \right\}.$$

By Equation (6.8), $\text{Im}(\mathbf{B}_{int}^\top) \cap \text{Im}(\mathbf{B}_{bd}^\top) \subseteq \mathcal{V}$. We will show $\mathcal{V} \subseteq \text{Im}(\mathbf{B}_{int}^\top) \cap \text{Im}(\mathbf{B}_{bd}^\top)$, that is, $\mathcal{V} \subseteq \text{Im}(\mathbf{B}_{int}^\top)$. Since the boundary layer triangulates a spherical shell in \mathbb{R}^3 , whose first Betti number is zero, \mathcal{V} is orthogonal to the image of ∂_1^\top of the boundary layer. Without loss of generality, we can also assume the 2-complex of the interior triangles touching the boundary inner sphere has the first Betti number being zero. Otherwise, we shift the boundary inner sphere towards the boundary outer sphere and include the boundary layer in the interior part. This can be done since, in the given r -hollowing, each region boundary triangulates a spherical shell with a “hop” shell width of at least 5.

Under this assumption, $\mathcal{V} \subseteq \text{Im} \left(\begin{bmatrix} \mathbf{B}_{12} \\ \mathbf{B}_2 \\ \mathbf{0} \end{bmatrix} \right) \subseteq \text{Im}(\mathbf{B}_{int}^\top)$. Thus, $\mathcal{V} = \text{Im}(\mathbf{B}_{int}^\top) \cap \text{Im}(\mathbf{B}_{bd}^\top)$.

Then,

$$\mathcal{U} = \left\{ \begin{bmatrix} \mathbf{0} \\ \mathbf{B}_3 \mathbf{y} \\ \mathbf{B}_4 \mathbf{y} \end{bmatrix} : \mathbf{B}_3 \mathbf{y} \perp \mathbf{B}_3 \mathbf{x}, \forall \mathbf{x} \in \text{Ker}(\mathbf{B}_4) \right\} \setminus \{\mathbf{0}\}.$$

By Equation (6.7) and (6.8),

$$\begin{aligned} \lambda_{\min} \left(\Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \Pi_{\text{Ker}(\mathbf{B}_{int})} \Pi_{\text{Im}(\mathbf{B}_{bd}^\top)} \right) &\geq \min_{\mathbf{y} \neq \mathbf{0}: \mathbf{B}_3^\top \mathbf{B}_3 \mathbf{y} \perp \text{Ker}(\mathbf{B}_4)} \frac{\|\mathbf{B}_4 \mathbf{y}\|^2}{\|\mathbf{B}_3 \mathbf{y}\|^2 + \|\mathbf{B}_4 \mathbf{y}\|^2} \\ &= \min_{\mathbf{y} \neq \mathbf{0}: \mathbf{B}_3^\top \mathbf{B}_3 \mathbf{y} \perp \text{Ker}(\mathbf{B}_4)} \frac{1}{\|\mathbf{B}_3 \mathbf{y}\|^2 / \|\mathbf{B}_4 \mathbf{y}\|^2 + 1}. \end{aligned}$$

It suffices to show $\frac{\|\mathbf{B}_3 \mathbf{y}\|^2}{\|\mathbf{B}_4 \mathbf{y}\|^2} = O(r)$.

Without loss of generality, we can assume each diagonal block of \mathbf{B}_3 in Equation (6.9) has equal dimensions c by duplicating columns in \mathbf{B}_{bd}^\top (which does not change its image).

$$\mathbf{B}_3^\top \mathbf{B}_3 = \text{diag}(\mathbf{J}, \dots, \mathbf{J}) \stackrel{\text{def}}{=} c\mathbf{\Pi}$$

is a multiple of projection matrix where \mathbf{J} is the all-one matrix in dimensions $c \times c$. Since $\mathbf{B}_3^\top \mathbf{B}_3 \mathbf{y} \perp \text{Ker}(\mathbf{B}_4)$, we know $\mathbf{B}_3^\top \mathbf{B}_3 \mathbf{y} = c\mathbf{\Pi} \mathbf{y} \in \text{Im}(\mathbf{B}_4^\top)$. Since $\mathbf{\Pi}$ is an orthogonal projection, we have $\|\mathbf{\Pi} \mathbf{y}\| \leq \|\mathbf{\Pi}_{\text{Im}(\mathbf{B}_4^\top)} \mathbf{y}\|$. Note $\mathbf{B}_4^\top \mathbf{B}_4$ can be treated as a graph Laplacian matrix. By the eigenvalue bound in Theorem 6.7.13,

$$\|\mathbf{B}_4 \mathbf{y}\|^2 \geq \Omega(r^{-1}) \|\mathbf{\Pi}_{\text{Im}(\mathbf{B}_4^\top)} \mathbf{y}\|^2 \geq \Omega(r^{-1}) \|\mathbf{\Pi} \mathbf{y}\|^2 = \Omega(r^{-1}) \cdot \frac{1}{c} \|\mathbf{B}_3 \mathbf{y}\|^2.$$

Therefore,

$$\frac{\|\mathbf{B}_3 \mathbf{y}\|^2}{\|\mathbf{B}_4 \mathbf{y}\|^2} = O(r). \quad \square$$

Combining all the lemmas and claims above, we prove Lemma 6.8.6.

6.8.2 Proof of Lemma 6.5.9

Given a vector \mathbf{b} , we approximate $\mathbf{\Pi}_1^{\text{up}} \mathbf{b}$ by the following steps: (1) compute $\mathbf{b}_1 = \mathbf{\Pi}_{\text{Im}(\partial_2(\cdot, F))} \mathbf{b}$; (2) compute $\mathbf{b}_2 = \partial_2^\top(C, \cdot) \mathbf{\Pi}_{\text{Ker}(\partial_2^\top(F, \cdot))} \mathbf{b}$; (3) approximately solve $\text{Sc}(\mathbf{L}_2^{\text{down}})_C \mathbf{b}_3 = \mathbf{b}_2$ via Preconditioned Conjugate Gradient and get an approximate solution $\widetilde{\mathbf{b}}_3$ up to error $\delta \leq \frac{\epsilon}{\|\mathbf{L}_1^{\text{up}}\|}$; (4) compute $\mathbf{b}_4 = \mathbf{\Pi}_{\text{Ker}(\partial_2^\top(F, \cdot))} \partial_2^\top(\cdot, C) \widetilde{\mathbf{b}}_3$; (5) compute $\mathbf{b}_5 = \mathbf{b}_1 + \mathbf{b}_4$. Let $\widetilde{\mathbf{\Pi}}_1^{\text{up}}$ be the above operator so that $\mathbf{b}_5 = \widetilde{\mathbf{\Pi}}_1^{\text{up}} \mathbf{b}$.

Claim 6.8.10. $\widetilde{\mathbf{\Pi}}_1^{\text{up}} \mathbf{b}$ is in the image of \mathbf{L}_1^{up} .

Proof. By Lemma 6.8.1, the image of \mathbf{L}_1^{up} can be decomposed as a direct sum of two orthogonal subspaces: $\text{Im}(\partial_2(\cdot, F))$ and $\text{Im}(\mathbf{\Pi}_{\text{Ker}(\partial_2^\top(F, \cdot))} \partial_2^\top(\cdot, C))$. Note $\widetilde{\mathbf{\Pi}}_1^{\text{up}} \mathbf{b} = \mathbf{b}_1 + \mathbf{b}_4$, where $\mathbf{b}_1 \in \text{Im}(\partial_2(\cdot, F))$ and $\mathbf{b}_4 \in \text{Im}(\mathbf{\Pi}_{\text{Ker}(\partial_2^\top(F, \cdot))} \partial_2^\top(\cdot, C))$. So, $\widetilde{\mathbf{\Pi}}_1^{\text{up}} \mathbf{b} \in \text{Im}(\mathbf{L}_1^{\text{up}})$. \square

We bound the error of $\widetilde{\mathbf{\Pi}}_1^{\text{up}} \mathbf{b}$:

$$\begin{aligned} \|\widetilde{\mathbf{\Pi}}_1^{\text{up}} \mathbf{b} - \mathbf{\Pi}_1^{\text{up}} \mathbf{b}\| &= \left\| \mathbf{\Pi}_{\text{Ker}(\partial_2^\top(F, \cdot))} \partial_2^\top(\cdot, C) ((\text{Sc}(\mathbf{L}_2^{\text{down}})_C)^\dagger - \mathbf{S}^\dagger) \partial_2^\top(C, \cdot) \mathbf{\Pi}_{\text{Ker}(\partial_2^\top(F, \cdot))} \mathbf{b} \right\| \\ &\leq \delta \left\| \partial_2^\top(C, \cdot) \mathbf{\Pi}_{\text{Ker}(\partial_2^\top(F, \cdot))} \mathbf{b} \right\| \left\| \partial_2^\top(C, \cdot) \mathbf{\Pi}_{\text{Ker}(\partial_2^\top(F, \cdot))} \right\| \\ &\leq \delta \lambda_{\max}(\mathbf{L}_1^{\text{up}}) \|\mathbf{\Pi}_1^{\text{up}} \mathbf{b}\| \\ &\leq \epsilon \|\mathbf{\Pi}_1^{\text{up}} \mathbf{b}\|, \end{aligned} \quad \text{by the setting of } \delta$$

Finally, we bound the runtime of our approximate projection algorithm. The proof is similar to that of Lemma 6.5.8. By Theorem 6.3.4 and Lemma 6.8.5, the number of Preconditioned Conjugate Gradient (PCG) iterations is

$$\widetilde{O}(\kappa(\text{Sc}(\mathbf{L}_2^{\text{down}})_C, \partial_2^\top(C, \cdot) \partial_2(\cdot, C))^{1/2}) = \widetilde{O}(r^{1/2}).$$

In each PCG iteration, we solve a system in $\partial_2^\top(C, :)\partial_2(:, C)$ and multiply $\text{Sc}(\mathbf{L}_2^{\text{down}})_C$ with $O(1)$ vectors. Recall

$$\text{Sc}(\mathbf{L}_2^{\text{down}})_C = \partial_2^\top(C, :)\mathbf{\Pi}_{\text{Ker}(\partial_2^\top(F, :))}\partial_2(:, C).$$

In our preprocessing, we compute a Cholesky factorization of $\partial_2^\top(F, :)\mathbf{\Pi}_{\text{Ker}(\partial_2^\top(F, :))}\partial_2(:, F)$ in time $O(nr)$ (by the proof of Claim 6.8.4); then, we can multiply $\text{Sc}(\mathbf{L}_2^{\text{down}})_C$ onto a vector in time $O(nr^{1/3})$. Similarly, we solve a system in $\partial_2^\top(C, :)\partial_2(:, C)$ by nested dissection. By our construction, C has $O(r^{2/3} \cdot \frac{n}{r}) = O(nr^{-1/3})$ triangles. The Cholesky factorization runs in time $O(n^2r^{-2/3})$, and each system solve runs in time $O(n^{4/3}r^{-4/9})$. So, the runtime per PCG iteration is

$$O(nr^{1/3} + n^{4/3}r^{-4/9}).$$

Therefore, the total time is

$$\tilde{O}(r^{1/2}(nr^{1/3} + n^{4/3}r^{-4/9}) + nr + n^2r^{-2/3}) = \tilde{O}(nr + n^{4/3}r^{5/18} + n^2r^{-2/3}).$$

We finish the proof of Lemma 6.5.9.

6.9 Proof of the Main Theorem

Given all the four operators in Lemma 6.5.2, 6.5.3, 6.5.4, and 6.5.9, we prove the following main theorem.

Theorem 6.9.1 (Restatement of Theorem 6.4.1). *Let \mathcal{K} be a pure 3-complex embedded in \mathbb{R}^3 consisting of n stable simplexes and with a known r -hollowing. Let \mathbf{L}_1 be the 1-Laplacian operator of \mathcal{K} , and let $\mathbf{\Pi}_1$ be the orthogonal projection matrix onto the image of \mathbf{L}_1 . For any vector \mathbf{b} and $\epsilon > 0$, we can find a solution $\tilde{\mathbf{x}}$ such that $\|\mathbf{L}_1\tilde{\mathbf{x}} - \mathbf{\Pi}_1\mathbf{b}\|_2 \leq \epsilon\|\mathbf{\Pi}_1\mathbf{b}\|_2$ in time $O(nr + n^{4/3}r^{5/18}\log(n/\epsilon) + n^2r^{-2/3})$.*

Proof. Let κ be the maximum of $\kappa(\mathbf{L}_1^{\text{down}})$ and $\kappa(\mathbf{L}_1^{\text{up}})$. Let $\delta > 0$ be a parameter to be determined later. Let $\tilde{\mathbf{\Pi}}_1^{\text{down}} = \tilde{\mathbf{\Pi}}_1^{\text{down}}(\delta)$, $\tilde{\mathbf{\Pi}}_1^{\text{up}} = \tilde{\mathbf{\Pi}}_1^{\text{up}}(\delta)$ be defined in Lemma 6.5.2 and 6.5.9, and let $\mathbf{Z}_1^{\text{down}}$ be the operator in Lemma 6.5.3 with no error and \mathbf{Z}_1^{up} in Lemma 6.5.4 with error δ . Let

$$\begin{aligned}\tilde{\mathbf{b}}^{\text{up}} &\stackrel{\text{def}}{=} \tilde{\mathbf{\Pi}}_1^{\text{up}}\mathbf{b}, \quad \tilde{\mathbf{b}}^{\text{down}} \stackrel{\text{def}}{=} \tilde{\mathbf{\Pi}}_1^{\text{down}}\mathbf{b}, \\ \tilde{\mathbf{x}}^{\text{up}} &\stackrel{\text{def}}{=} \mathbf{Z}_1^{\text{up}}\tilde{\mathbf{b}}^{\text{up}}, \quad \tilde{\mathbf{x}}^{\text{down}} \stackrel{\text{def}}{=} \mathbf{Z}_1^{\text{down}}\tilde{\mathbf{b}}^{\text{down}}, \\ \tilde{\mathbf{x}} &\stackrel{\text{def}}{=} \tilde{\mathbf{\Pi}}_1^{\text{up}}\tilde{\mathbf{x}}^{\text{up}} + \tilde{\mathbf{\Pi}}_1^{\text{down}}\tilde{\mathbf{x}}^{\text{down}}.\end{aligned}$$

Then,

$$\begin{aligned}\|\mathbf{L}_1\tilde{\mathbf{x}} - \mathbf{\Pi}_1\mathbf{b}\|_2 &\leq \left\| \mathbf{L}_1^{\text{up}}\tilde{\mathbf{\Pi}}_1^{\text{up}}\tilde{\mathbf{x}}^{\text{up}} - \tilde{\mathbf{b}}^{\text{up}} \right\|_2 + \left\| \mathbf{L}_1^{\text{down}}\tilde{\mathbf{\Pi}}_1^{\text{down}}\tilde{\mathbf{x}}^{\text{down}} - \tilde{\mathbf{b}}^{\text{down}} \right\|_2 + \left\| \tilde{\mathbf{b}}^{\text{up}} + \tilde{\mathbf{b}}^{\text{down}} - \mathbf{\Pi}_1\mathbf{b} \right\|_2.\end{aligned}$$

We will upper bound the three terms on the right-hand side separately.

- For the first term,

$$\begin{aligned} \left\| \mathbf{L}_1^{\text{up}} \tilde{\Pi}_1^{\text{up}} \tilde{\mathbf{x}}^{\text{up}} - \tilde{\mathbf{b}}^{\text{up}} \right\|_2 &\leq \left\| \mathbf{L}_1^{\text{up}} \tilde{\Pi}_1^{\text{up}} \tilde{\mathbf{x}}^{\text{up}} - \mathbf{L}_1^{\text{up}} \Pi_1^{\text{up}} \tilde{\mathbf{x}}^{\text{up}} \right\|_2 + \left\| \mathbf{L}_1^{\text{up}} \tilde{\mathbf{x}}^{\text{up}} - \tilde{\mathbf{b}}^{\text{up}} \right\|_2 \\ &\leq \left\| \mathbf{L}_1^{\text{up}} \tilde{\Pi}_1^{\text{up}} \tilde{\mathbf{x}}^{\text{up}} - \mathbf{L}_1^{\text{up}} \Pi_1^{\text{up}} \tilde{\mathbf{x}}^{\text{up}} \right\|_2 + \delta \left\| \tilde{\mathbf{b}}^{\text{up}} \right\|_2, \text{ by Lemma 6.5.4} \end{aligned}$$

By Lemma 6.5.9,

$$\begin{aligned} \left\| \mathbf{L}_1^{\text{up}} \tilde{\Pi}_1^{\text{up}} \tilde{\mathbf{x}}^{\text{up}} - \mathbf{L}_1^{\text{up}} \Pi_1^{\text{up}} \tilde{\mathbf{x}}^{\text{up}} \right\|_2 &\leq \left\| \mathbf{L}_1^{\text{up}} \right\|_2 \left\| (\tilde{\Pi}_1^{\text{up}} - \Pi_1^{\text{up}}) \Pi_1^{\text{up}} \tilde{\mathbf{x}}^{\text{up}} \right\|_2 \\ &\leq \delta \left\| \mathbf{L}_1^{\text{up}} \right\|_2 \left\| \Pi_1^{\text{up}} \tilde{\mathbf{x}}^{\text{up}} \right\|_2. \end{aligned}$$

Let $\mathbf{y} \stackrel{\text{def}}{=} (\mathbf{L}_1^{\text{up}})^\dagger \tilde{\mathbf{b}}^{\text{up}}$. By Lemma 6.5.4,

$$\left\| \Pi_1^{\text{up}} \tilde{\mathbf{x}}^{\text{up}} - \mathbf{y} \right\|_2 \leq \left\| (\mathbf{L}_1^{\text{up}})^\dagger \right\|_2 \left\| \mathbf{L}_1^{\text{up}} \Pi_1^{\text{up}} \tilde{\mathbf{x}}^{\text{up}} - \tilde{\mathbf{b}}^{\text{up}} \right\|_2 \leq \delta \left\| (\mathbf{L}_1^{\text{up}})^\dagger \right\|_2 \left\| \tilde{\mathbf{b}}^{\text{up}} \right\|_2.$$

By the triangle inequality,

$$\left\| \Pi_1^{\text{up}} \tilde{\mathbf{x}}^{\text{up}} \right\|_2 \leq \left\| \mathbf{y} \right\|_2 + \delta \left\| (\mathbf{L}_1^{\text{up}})^\dagger \right\|_2 \left\| \tilde{\mathbf{b}}^{\text{up}} \right\|_2 \leq (1 + \delta) \left\| (\mathbf{L}_1^{\text{up}})^\dagger \right\|_2 \left\| \tilde{\mathbf{b}}^{\text{up}} \right\|_2.$$

So,

$$\left\| \mathbf{L}_1^{\text{up}} \tilde{\Pi}_1^{\text{up}} \tilde{\mathbf{x}}^{\text{up}} - \mathbf{L}_1^{\text{up}} \Pi_1^{\text{up}} \tilde{\mathbf{x}}^{\text{up}} \right\|_2 \leq \delta(1 + \delta) \kappa \left\| \tilde{\mathbf{b}}^{\text{up}} \right\|_2,$$

and

$$\left\| \mathbf{L}_1^{\text{up}} \tilde{\Pi}_1^{\text{up}} \tilde{\mathbf{x}}^{\text{up}} - \tilde{\mathbf{b}}^{\text{up}} \right\|_2 \leq 3\delta\kappa \left\| \tilde{\mathbf{b}}^{\text{up}} \right\|_2.$$

- For the second term, the operator $\mathbf{Z}_1^{\text{down}}$ has no error, which means $\mathbf{L}_1^{\text{down}} \tilde{\mathbf{x}}^{\text{down}} = \tilde{\mathbf{b}}^{\text{down}}$. Then,

$$\begin{aligned} \left\| \mathbf{L}_1^{\text{down}} \tilde{\Pi}_1^{\text{down}} \tilde{\mathbf{x}}^{\text{down}} - \tilde{\mathbf{b}}^{\text{down}} \right\|_2 &= \left\| \mathbf{L}_1^{\text{down}} \tilde{\Pi}_1^{\text{down}} \tilde{\mathbf{x}}^{\text{down}} - \mathbf{L}_1^{\text{down}} \tilde{\mathbf{x}}^{\text{down}} \right\|_2 \\ &\leq \delta(1 + \delta) \kappa \left\| \tilde{\mathbf{b}}^{\text{down}} \right\|_2. \end{aligned}$$

- For the third term,

$$\begin{aligned} \left\| \tilde{\mathbf{b}}^{\text{up}} + \tilde{\mathbf{b}}^{\text{down}} - \Pi_1 \mathbf{b} \right\|_2^2 &= \left\| (\tilde{\Pi}^{\text{up}} - \Pi^{\text{up}}) \mathbf{b} \right\|_2^2 + \left\| (\tilde{\Pi}^{\text{down}} - \Pi^{\text{down}}) \mathbf{b} \right\|_2^2 \\ &\leq \delta^2 \left(\left\| \Pi^{\text{up}} \mathbf{b} \right\|_2^2 + \left\| \Pi^{\text{down}} \mathbf{b} \right\|_2^2 \right) \\ &\quad \text{by Lemma 6.5.2, 6.5.9, Fact 2.3.6} \\ &= \delta^2 \left\| \Pi_1 \mathbf{b} \right\|_2^2. \end{aligned}$$

Combining all the above inequalities,

$$\begin{aligned} \left\| \mathbf{L}_1 \tilde{\mathbf{x}} - \Pi_1 \mathbf{b} \right\|_2 &\leq 3\delta\kappa \left\| \tilde{\mathbf{b}}^{\text{up}} \right\|_2 + 2\delta\kappa \left\| \tilde{\mathbf{b}}^{\text{down}} \right\|_2 + \delta \left\| \Pi_1 \mathbf{b} \right\|_2 \\ &\leq 3\delta\kappa(1 + \delta) \left\| \Pi_1^{\text{up}} \mathbf{b} \right\|_2 + 2\delta\kappa(1 + \delta) \left\| \Pi_1^{\text{down}} \mathbf{b} \right\|_2 + \delta \left\| \Pi_1 \mathbf{b} \right\|_2 \\ &\leq 11\delta\kappa \left\| \Pi_1 \mathbf{b} \right\|_2. \end{aligned}$$

Choosing $\delta \leq \frac{\epsilon}{11\kappa}$, we have

$$\left\| \mathbf{L}_1 \tilde{\mathbf{x}} - \Pi_1 \mathbf{b} \right\|_2 \leq \epsilon \left\| \Pi_1 \mathbf{b} \right\|_2.$$

The runtime bound follows Lemma 6.5.2, 6.5.3, 6.5.4, and 6.5.9. \square

6.10 A Union of Pure 3-Complexes

In this section, we consider a union \mathcal{U} of h pure 3-complex chunks $\mathcal{K}_1, \dots, \mathcal{K}_h$ that are composed of stable simplexes, and these chunks are glued together by identifying certain subsets of their exterior simplexes. Let k be the number of simplexes shared by more than one chunk. We remark that \mathcal{U} may *not* be embeddable in \mathbb{R}^3 and the first and second Betti numbers of \mathcal{U} are *no longer zero*. This makes designing efficient solvers for 1-Laplacian systems of \mathcal{U} much harder. Edelsbrunner and Parsa [EP14] showed that computing the first Betti number of a simplicial complex linearly embedded in \mathbb{R}^4 with m simplexes is as hard as computing the rank of a 0-1 matrix with $\Theta(m)$ non-zeros; Ding, Kyng, Probst Gutenberg, and Zhang [Din+22] showed that (approximately) solving 1-Laplacian systems for simplicial complexes in \mathbb{R}^4 is as hard as (approximately) solving general sparse systems of linear equations.

We design a 1-Laplacian solver for \mathcal{U} whose runtime is comparable to that of the 1-Laplacian solver for a single pure 3-complex when both h and k are small and prove Theorem 6.4.2. The approximate down-projection operator for in Lemma 6.5.2 and the approximate solver for systems in the down-Laplacian in Lemma 6.5.3 hold for any simplicial complexes. Thus, it suffices to generalize the approximate up-projection operator in Lemma 6.5.9 and the approximate solver for systems in the up-Laplacian in Lemma 6.5.4 from a single chunk to a union of chunks.

Lemma 6.10.1 (Up-Laplacian Solver). *Let \mathcal{U} be a union of pure 3-complexes glued together by identifying certain subsets of their exterior simplexes. Each 3-complex chunk is embedded in \mathbb{R}^3 , comprises n_i stable simplexes, and has a known $\Theta(n_i^{3/5})$ -hollowing. For any $\mathbf{b} \in \text{Im}(\mathbf{L}_1^{\text{up}})$ and $\epsilon > 0$, we can compute an $\tilde{\mathbf{x}}$ such that $\|\mathbf{L}_1^{\text{up}} \tilde{\mathbf{x}} - \mathbf{b}\|_2 \leq \epsilon \|\mathbf{b}\|_2$ in time $\tilde{O}(n^{8/5} + n^{3/10}k^2 + k^3)$ where n is the number of simplexes in \mathcal{U} and k is the number of simplexes shared by more than one chunk.*

Lemma 6.10.2 (Up-Projection Operator). *Let \mathcal{U} be a 3-complex satisfying the conditions in Lemma 6.10.2. For any $\epsilon > 0$, there exists an operator $\tilde{\Pi}_1^{\text{up}}$ such that*

$$\forall \mathbf{b}, \quad \left\| \tilde{\Pi}_1^{\text{up}} \mathbf{b} - \Pi_1^{\text{up}} \mathbf{b} \right\|_2 \leq \epsilon \|\Pi_1^{\text{up}} \mathbf{b}\|_2.$$

In addition, $\tilde{\Pi}_1^{\text{up}} \mathbf{b}$ can be computed in time $\tilde{O}(n^{8/5} + n^{3/10}k^2 + k^3)$, where n is the number of exterior simplexes of \mathcal{U} and k is the number of simplexes shared by more than one chunk.

Our approaches align with those for a single chunk. We partition the simplexes of \mathcal{K} into $F \cup C$, then deal with the “ F ” part and the Schur complement separately. Given the definition of r -hollowing, the exterior simplexes of each chunk \mathcal{K}_i must belong to the boundary of some region. We let C be the union of the hollowing boundary of each chunk and let F be the remaining simplexes. Then, F is a union of disjoint subcomplexes, each embedded in \mathbb{R}^3 and can be handled by Nested Dissection. We precondition the Schur complement by the union of the hollowing boundaries. However, systems in this preconditioner cannot be approximately solved Nested Dissection directly since it may not allow an embedding in \mathbb{R}^3 . We will need a slightly more careful treatment.

Proof of Lemma 6.10.1. We let C be the union of the hollowing boundary edges in each chunk and let F be the union of the hollowing interior edges. Suppose \mathcal{U} has h pure 3-complex chunks. By our definition of r -hollowing, we can write $\mathbf{L}_1^{\text{up}}[F, F] =$

$\text{diag}(\mathbf{L}_1^{\text{up}}[F_1, F_1], \dots, \mathbf{L}_1^{\text{up}}[F_h, F_h])$ where F_i contains all the hollowing interior edges in the i th chunk. Let $r_i = n_i^{3/5}$ be the hollowing parameter for the i th chunk. By Lemma 6.5.7, with a pre-processing time

$$O\left(\sum_{i=1}^h n_i r_i\right) = O\left(\sum_{i=1}^h n_i^{8/5}\right) = O(n^{8/5}),$$

for any $\mathbf{b} \in \text{Im}(\mathbf{L}_1^{\text{up}}[F, F])$, we can find \mathbf{x} such that $\mathbf{L}_1^{\text{up}}[F, F]\mathbf{x} = \mathbf{b}$ in time

$$O\left(\sum_{i=1}^h n_i r_i^{1/3}\right) = O\left(\sum_{i=1}^h n_i^{6/5}\right) = O(n^{6/5}).$$

To solve the system in the Schur complement $\text{Sc}[\mathbf{L}_1^{\text{up}}]_C$, we precondition it by the union of the hollowing boundaries of each chunk, denoted by \mathcal{T}_U . By Claim 6.7.7 and Lemma 6.7.8, the relative condition number is $O(\max_i n_i^{3/5})$. Let $C_1 \subset C$ contain the edges shared by more than one chunk and $C_2 = C \setminus C_1$. Then, the submatrix $\mathbf{L}_{1, \mathcal{T}_U}^{\text{up}}[C_2, C_2]$ is a block diagonal matrix where each block corresponds to a chunk. We solve $\mathbf{L}_{1, \mathcal{T}_U}^{\text{up}}$ by Lemma 6.5.5: We solve a system in $\mathbf{L}_{1, \mathcal{T}_U}^{\text{up}}[C_2, C_2]$ by Nested Dissection and solve the Schur complement onto C_1 by directly inverting the Schur complement. With a pre-processing time

$$O\left(\sum_{i=1}^h \left(\frac{n_i}{r_i} \cdot r_i^{2/3}\right)^2 + k^3\right) = O\left(\sum_{i=1}^h n_i^{8/5} + k^3\right) = O(n^{8/5} + k^3),$$

we can solve a system in $\mathbf{L}_{1, \mathcal{T}_U}^{\text{up}}$ in time

$$O\left(\sum_{i=1}^h \left(\frac{n_i}{r_i} \cdot r_i^{2/3}\right)^{4/3} + k^2\right) = O\left(\sum_{i=1}^h n_i^{16/15} + k^2\right) = O(n^{16/15} + k^2).$$

Therefore, the total runtime of approximately solving a system in \mathbf{L}_1^{up} is $\tilde{O}(n^{8/5} + n^{3/10}k^2 + k^3)$. \square

We can prove Lemma 6.10.2 by a similar argument.

Bibliography

- [AMO93] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. “Network Flows”. In: (1993).
- [AW21] J. Alman and V. V. Williams. “A refined laser method and faster matrix multiplication”. In: *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2021, pp. 522–539.
- [Axe85] O. Axelsson. “A survey of preconditioned iterative methods for linear systems of algebraic equations”. In: *BIT Numerical Mathematics* 25.1 (1985), pp. 165–187.
- [AY10] N. Alon and R. Yuster. “Solving linear systems through nested dissection”. In: *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*. IEEE. 2010, pp. 225–234.
- [BHV08a] E. G. Boman, B. Hendrickson, and S. Vavasis. “Solving elliptic finite element systems in near-linear time with support preconditioners”. In: *SIAM Journal on Numerical Analysis* 46.6 (2008), pp. 3264–3284.
- [BHV08b] E. G. Boman, B. Hendrickson, and S. Vavasis. “Solving elliptic finite element systems in near-linear time with support preconditioners”. In: *SIAM Journal on Numerical Analysis* 46.6 (2008), pp. 3264–3284.
- [BKV09] C. Barnhart, N. Krishnan, and P. H. Vance. “Multicommodity Flow Problems”. en. In: *Encyclopedia of Optimization*. Ed. by C. A. Floudas and P. M. Pardalos. Boston, MA: Springer US, 2009, pp. 2354–2362. ISBN: 978-0-387-74759-0.
- [Bla+22] M. Black, W. Maxwell, A. Nayyeri, and E. Winkelman. “Computational Topology in a Collapsing Universe: Laplacians, Homology, Cohomology”. In: *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2022, pp. 226–251.
- [BN22] M. Black and A. Nayyeri. “Hodge Decomposition and General Laplacian Solvers for Embedded Simplicial Complexes”. In: *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik. 2022.
- [Bro+17] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst. “Geometric deep learning: going beyond euclidean data”. In: *IEEE Signal Processing Magazine* 34.4 (2017), pp. 18–42.
- [BS20] S. Barbarossa and S. Sardellitti. “Topological signal processing over simplicial complexes”. In: *IEEE Transactions on Signal Processing* 68 (2020), pp. 2992–3007.

- [BV21] M. Bafna and N. Vyas. “Optimal fine-grained hardness of approximation of linear equations”. In: *48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*. Ed. by N. Bansal, E. Merelli, and J. Worrell. Vol. 198. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021, 20:1–20:19. ISBN: 978-3-95977-195-5.
- [Car09] G. Carlsson. “Topology and Data”. In: *Bulletin of the American Mathematical Society* 46.2 (2009), pp. 255–308.
- [Cha+16] F. Chazal, V. de Silva, M. Glisse, and S. Oudot. *The Structure and Stability of Persistence Modules*. Springer, 2016. ISBN: 978-3-319-42545-0.
- [Che+20] L. Chen, G. Goranci, M. Henzinger, R. Peng, and T. Saranurak. “Fast Dynamic Cuts, Distances and Effective Resistances via Vertex Sparsifiers”. In: *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. 2020, pp. 1135–1146.
- [Che+22] L. Chen, R. Kyng, Y. P. Liu, R. Peng, M. P. Gutenberg, and S. Sachdeva. “Maximum flow and minimum-cost flow in almost-linear time”. In: *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2022, pp. 612–623.
- [Chi67] D. R. Chillingworth. “Collapsing three-dimensional convex polyhedra”. In: *Mathematical Proceedings of the Cambridge Philosophical Society*. Vol. 63. Cambridge University Press. 1967, pp. 353–357.
- [Chi80] D. R. Chillingworth. “Collapsing three-dimensional convex polyhedra: correction”. In: *Mathematical Proceedings of the Cambridge Philosophical Society*. Vol. 88. Cambridge University Press. 1980, pp. 307–310.
- [Chr+11] P. Christiano, J. A. Kelner, A. Madry, D. A. Spielman, and S.-H. Teng. “Electrical Flows, Laplacian Systems, and Faster Approximation of Maximum Flow in Undirected Graphs”. In: *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*. 2011, pp. 273–282.
- [CLS21] M. B. Cohen, Y. T. Lee, and Z. Song. “Solving Linear Programs in the Current Matrix Multiplication Time”. In: *Journal of the ACM (JACM)* 68.1 (2021), pp. 1–39.
- [CMZ18] C. K. Chui, H. Mhaskar, and X. Zhuang. “Representation of functions on big data associated with directed graphs”. In: *Applied and Computational Harmonic Analysis* 44.1 (2018), pp. 165–188.
- [Coh+14a] M. B. Cohen, B. T. Fasy, G. L. Miller, A. Nayyeri, R. Peng, and N. Walkington. “Solving 1-laplacians in nearly linear time: Collapsing and expanding a topological ball”. In: *Proceedings of the 2014 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2014, pp. 204–216.
- [Coh+14b] M. B. Cohen, R. Kyng, G. L. Miller, J. W. Pachocki, R. Peng, A. B. Rao, and S. C. 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*. ACM. 2014, pp. 343–352.

- [Coh+17a] M. B. Cohen, J. Kelner, J. Peebles, R. Peng, A. B. Rao, A. Sidford, and A. Vladu. “Almost-linear-time algorithms for markov chains and new spectral primitives for directed graphs”. In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. 2017, pp. 410–419.
- [Coh+17b] M. B. Cohen, J. Kelner, J. Peebles, R. Peng, A. B. Rao, A. Sidford, and A. Vladu. “Almost-linear-time Algorithms for Markov Chains and New Spectral Primitives for Directed Graphs”. In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2017. Montreal, Canada: ACM, 2017, pp. 410–419. ISBN: 978-1-4503-4528-6.
- [Coh+18] M. B. Cohen, J. Kelner, R. Kyng, J. Peebles, R. Peng, A. B. Rao, and A. Sidford. “Solving directed Laplacian systems in nearly-linear time through sparse LU factorizations”. In: *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2018, pp. 898–909.
- [DDH07] J. Demmel, I. Dumitriu, and O. Holtz. “Fast Linear Algebra Is Stable”. In: *Numerische Mathematik* 108.1 (2007), pp. 59–91.
- [Din+22] M. Ding, R. Kyng, M. P. Gutenberg, and P. Zhang. “Hardness results for laplacians of simplicial complexes via sparse-linear equation complete gadgets”. In: *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*. Ed. by M. Bojańczyk, E. Merelli, and D. P. Woodruff. Vol. 229. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022, 53:1–53:19. ISBN: 978-3-95977-235-8.
- [Din70] E. A. Dinic. “Algorithm for Solution of a Problem of Maximum Flow in Networks with Power Estimation”. In: *Soviet Math. Doklady*. Vol. 11. 1970, pp. 1277–1280.
- [DKM09] A. Duval, C. Klivans, and J. Martin. “Simplicial matrix-tree theorems”. In: *Transactions of the American Mathematical Society* 361.11 (2009), pp. 6073–6114.
- [DKM15] A. M. Duval, C. J. Klivans, and J. L. Martin. “Cuts and flows of cell complexes”. In: *Journal of Algebraic Combinatorics* 41.4 (2015), pp. 969–999.
- [DKT08] M. Desbrun, E. Kanso, and Y. Tong. “Discrete differential forms for computational modeling”. In: *Discrete differential geometry*. Springer, 2008, pp. 287–324.
- [DKZ22] M. Ding, R. Kyng, and P. Zhang. “Two-Commodity Flow Is Equivalent to Linear Programming Under Nearly-Linear Time Reductions”. In: *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*. Ed. by M. Bojańczyk, E. Merelli, and D. P. Woodruff. Vol. 229. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022, 54:1–54:19. ISBN: 978-3-95977-235-8.
- [DR80] D. P. Dobkin and S. P. Reiss. “The complexity of linear programming”. In: *Theoretical Computer Science* 11.1 (1980), pp. 1–18.
- [DS07] S. I. Daitch and D. A. Spielman. “Support-Graph Preconditioners for 2-Dimensional Trusses”. In: *CoRR* abs/cs/0703119 (2007). arXiv: [cs / 0703119](https://arxiv.org/abs/cs/0703119).

- [DS08] S. I. Daitch and D. A. Spielman. “Faster approximate lossy generalized flow via interior point algorithms”. In: *Proceedings of the 40th annual ACM symposium on Theory of computing*. STOC ’08. Available at <http://arxiv.org/abs/0803.0988>. Victoria, British Columbia, Canada: ACM, 2008, pp. 451–460. ISBN: 978-1-60558-047-0.
- [DW02] X. Dong and M. L. Wachs. “Combinatorial Laplacian of the matching complex”. In: *the electronic journal of combinatorics* (2002), R17–R17.
- [DZ23a] J. van Den Brand and D. J. Zhang. “Faster high accuracy multi-commodity flow from single-commodity techniques”. In: *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2023, pp. 493–502.
- [DZ23b] M. Ding and P. Zhang. “Efficient 1-Laplacian Solvers for Well-Shaped Simplicial Complexes: Beyond Betti Numbers and Collapsing Sequences”. In: *31st Annual European Symposium on Algorithms (ESA 2023)*. Ed. by I. L. Gørtz, M. Farach-Colton, S. J. Puglisi, and G. Herman. Vol. 274. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023, 41:1–41:19. ISBN: 978-3-95977-295-2.
- [Eck44] B. Eckmann. “Harmonische Funktionen Und Randwertaufgaben in Einem Komplex”. In: *Commentarii Mathematici Helvetici* 17.1 (1944), pp. 240–255. ISSN: 1420-8946.
- [EH10] H. Edelsbrunner and J. Harer. *Computational Topology: An Introduction*. American Mathematical Soc., 2010.
- [EP14] H. Edelsbrunner and S. Parsa. “On the computational complexity of Betti numbers: reductions from matrix rank”. In: *Proceedings of the twenty-fifth annual ACM-SIAM symposium on discrete algorithms*. SIAM. 2014, pp. 152–160.
- [ET75] S. Even and R. E. Tarjan. “Network Flow and Testing Graph Connectivity”. In: *SIAM journal on computing* 4.4 (1975), pp. 507–518.
- [Eva76] J. R. Evans. “A Combinatorial Equivalence between A Class of Multicommodity Flow Problems and the Capacitated Transportation Problem”. en. In: *Mathematical Programming* 10.1 (1976), pp. 401–404. ISSN: 1436-4646.
- [Eva78] J. R. Evans. “The Simplex Method for Integral Multicommodity Networks”. en. In: *Naval Research Logistics Quarterly* 25.1 (1978), pp. 31–37. ISSN: 00281441, 19319193.
- [FF56] L. R. Ford and D. R. Fulkerson. “Maximal Flow through a Network”. In: *Canadian journal of Mathematics* 8 (1956), pp. 399–404.
- [Fle00] L. K. Fleischer. “Approximating Fractional Multicommodity Flow Independent of the Number of Commodities”. In: *SIAM Journal on Discrete Mathematics* 13.4 (2000), pp. 505–520. ISSN: 0895-4801.
- [Fri96] J. Friedman. “Computing Betti numbers via combinatorial Laplacians”. In: *Proceedings of the twenty-eighth annual ACM symposium on Theory of Computing*. 1996, pp. 386–391.

- [Fri98] J. Friedman. “Computing Betti numbers via combinatorial Laplacians”. In: *Algorithmica* 21.4 (1998), pp. 331–346.
- [Geo73] A. George. “Nested dissection of a regular finite element mesh”. In: *SIAM Journal on Numerical Analysis* 10.2 (1973), pp. 345–363.
- [Ghr08a] R. Ghrist. “Barcodes: The Persistent Topology of Data”. In: *Bulletin of the American Mathematical Society* 45.1 (2008), pp. 61–75.
- [Ghr08b] R. Ghrist. “Barcodes: the persistent topology of data”. In: *Bulletin of the American Mathematical Society* 45.1 (2008), pp. 61–75.
- [GK07] N. Garg and J. Könemann. “Faster and Simpler Algorithms for Multicommodity Flow and Other Fractional Packing Problems”. In: *SIAM Journal on Computing* 37.2 (2007), pp. 630–652. ISSN: 0097-5397.
- [GLP21] Y. Gao, Y. P. Liu, and R. Peng. “Fully Dynamic Electrical Flows: Sparse Maxflow Faster Than Goldberg-Rao”. In: *arXiv:2101.07233 [cs]* (2021). arXiv: [2101.07233 \[cs\]](https://arxiv.org/abs/2101.07233).
- [GR98] A. V. Goldberg and S. Rao. “Beyond the Flow Decomposition Barrier”. In: *Journal of the ACM* 45.5 (1998), pp. 783–797. ISSN: 0004-5411.
- [GV96] G. Golub and C. Van Loan. *Matrix Computations*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, 1996. ISBN: 978-0-8018-5414-9.
- [HS+52] M. R. Hestenes, E. Stiefel, et al. *Methods of conjugate gradients for solving linear systems*. Vol. 49. 1. NBS Washington, DC, 1952.
- [IP01] R. Impagliazzo and R. Paturi. “On the complexity of k-SAT”. In: *Journal of Computer and System Sciences* 62.2 (2001), pp. 367–375.
- [Ita78] A. Itai. “Two-Commodity Flow”. In: *Journal of the ACM (JACM)* 25.4 (1978), pp. 596–611.
- [Jia+11] X. Jiang, L.-H. Lim, Y. Yao, and Y. Ye. “Statistical ranking and combinatorial Hodge theory”. In: *Mathematical Programming* 127.1 (2011), pp. 203–244.
- [Jia+21] S. Jiang, Z. Song, O. Weinstein, and H. Zhang. “A faster algorithm for solving general LPs”. In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 2021, pp. 823–832.
- [JS21] A. Jambulapati and A. Sidford. “Ultrasparse ultrasparsifiers and faster laplacian system solvers”. In: *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2021, pp. 540–559.
- [Kar84] N. Karmarkar. “A New Polynomial-Time Algorithm for Linear Programming”. In: *Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing*. 1984, pp. 302–311.
- [Kel+13] J. A. Kelner, L. Orecchia, A. Sidford, and Z. A. Zhu. “A Simple, Combinatorial Algorithm for Solving SDD Systems in Nearly-Linear Time”. In: *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing*. 2013, pp. 911–920.

- [Kel+14] J. A. Kelner, Y. T. Lee, L. Orecchia, and A. Sidford. “An Almost-Linear-Time Algorithm for Approximate Max Flow in Undirected Graphs, and Its Multicommodity Generalizations”. In: *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2014, pp. 217–226.
- [Ken78] J. L. Kennington. “A Survey of Linear Cost Multicommodity Network Flows”. In: *Operations Research* 26.2 (1978), pp. 209–236. ISSN: 0030-364X.
- [Kha80a] L. G. Khachiyan. “Polynomial algorithms in linear programming”. In: *USSR Computational Mathematics and Mathematical Physics* 20.1 (1980), pp. 53–72.
- [Kha80b] L. G. Khachiyan. “Polynomial Algorithms in Linear Programming”. In: *USSR Computational Mathematics and Mathematical Physics* 20.1 (1980), pp. 53–72.
- [KLS20] T. Kathuria, Y. P. Liu, and A. Sidford. “Unit Capacity Maxflow in Almost $m^{4/3}$ Time”. In: *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2020, pp. 119–130.
- [KMP10] I. Koutis, G. L. Miller, and R. Peng. “Approaching Optimality for Solving SDD Linear Systems”. In: *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science*. FOCS ’10. USA: IEEE Computer Society, 2010, pp. 235–244. ISBN: 978-0-7695-4244-7.
- [KMP11] I. Koutis, G. L. Miller, and R. Peng. “A Nearly- $m \log n$ Time Solver for SDD Linear Systems”. In: *Proceedings of the 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*. FOCS ’11. Available at <http://arxiv.org/abs/1102.4842>. Washington, DC, USA: IEEE Computer Society, 2011, pp. 590–598. ISBN: 978-0-7695-4571-4.
- [KMT11] I. Koutis, G. L. Miller, and D. Tolliver. “Combinatorial preconditioners and multilevel solvers for problems in computer vision and image processing”. In: *Computer Vision and Image Understanding* 115.12 (2011), pp. 1638–1646.
- [KS16] R. Kyng and S. Sachdeva. “Approximate gaussian elimination for laplacians-fast, sparse, and simple”. In: *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*. IEEE. 2016, pp. 573–582.
- [KWZ20] R. Kyng, D. Wang, and P. Zhang. “Packing LPs Are Hard to Solve Accurately, Assuming Linear Equations Are Hard”. In: *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2020, pp. 279–296.
- [Kyn+16] R. Kyng, Y. T. Lee, R. Peng, S. Sachdeva, and D. A. Spielman. “Sparsified Cholesky and Multigrid Solvers for Connection Laplacians”. In: *Proceedings of the Forty-eighth Annual ACM Symposium on Theory of Computing*. STOC ’16. Cambridge, MA, USA: ACM, 2016, pp. 842–850. ISBN: 978-1-4503-4132-5.
- [Kyn+18] R. Kyng, R. Peng, R. Schwiertman, and P. Zhang. “Incomplete nested dissection”. In: *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. 2018, pp. 404–417.

- [Kyn+19] R. Kyng, R. Peng, S. Sachdeva, and D. Wang. “Flows in almost linear time via adaptive preconditioning”. In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2019. Phoenix, AZ, USA: Association for Computing Machinery, 2019, pp. 902–913. ISBN: 9781450367059.
- [KZ17] R. Kyng and P. Zhang. “Hardness Results for Structured Linear Systems”. In: *SIAM Journal on Computing* 49.4 (2017), FOCS17–280.
- [Lee+19] H. Lee, M. K. Chung, H. Choi, H. Kang, S. Ha, Y. K. Kim, and D. S. Lee. “Harmonic holes as the submodules of brain network and network dissimilarity”. In: *Computational Topology in Image Context: 7th International Workshop, CTIC 2019, Málaga, Spain, January 24–25, 2019, Proceedings 7*. Springer. 2019, pp. 110–122.
- [Lei+95] T. Leighton, F. Makedon, S. Plotkin, C. Stein, E. Tardos, and S. Tragoudas. “Fast Approximation Algorithms for Multicommodity Flow Problems”. en. In: *Journal of Computer and System Sciences* 50.2 (1995), pp. 228–243. ISSN: 0022-0000.
- [Lim20] L.-H. Lim. “Hodge laplacians on graphs”. In: *Siam Review* 62.3 (2020), pp. 685–715.
- [LRS13] Y. T. Lee, S. Rao, and N. Srivastava. “A New Approach to Computing Maximum Flows Using Electrical Flows”. In: *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing*. 2013, pp. 755–764.
- [LRT79] R. J. Lipton, D. J. Rose, and R. E. Tarjan. “Generalized nested dissection”. In: *SIAM journal on numerical analysis* 16.2 (1979), pp. 346–358.
- [LS13] Y. T. Lee and A. Sidford. “Efficient accelerated coordinate descent methods and faster algorithms for solving linear systems”. In: *2013 IEEE 54th annual symposium on foundations of computer science*. IEEE. 2013, pp. 147–156.
- [LS14] Y. T. Lee and A. Sidford. “Path Finding Methods for Linear Programming: Solving Linear Programs in $\tilde{O}(\sqrt{\text{rank}})$ Iterations and Faster Algorithms for Maximum Flow”. In: *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*. Available at <http://arxiv.org/abs/1312.6677> and <http://arxiv.org/abs/1312.6713>. IEEE. 2014, pp. 424–433.
- [LS20] Y. P. Liu and A. Sidford. “Faster Energy Maximization for Faster Maximum Flow”. In: *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*. 2020, pp. 803–814.
- [Ma+11a] W. Ma, J.-M. Morel, S. Osher, and A. Chien. “An L 1-based variational model for Retinex theory and its application to medical images”. In: *CVPR 2011*. IEEE. 2011, pp. 153–160.
- [Ma+11b] W. Ma, J.-M. Morel, S. Osher, and A. Chien. “An L 1-based variational model for Retinex theory and its application to medical images”. In: *CVPR 2011*. IEEE. 2011, pp. 153–160.
- [Mad10] A. Madry. “Faster Approximation Schemes for Fractional Multicommodity Flow Problems via Dynamic Graph Algorithms”. In: *Proceedings of the Forty-Second ACM Symposium on Theory of Computing*. STOC ’10. New York, NY, USA: Association for Computing Machinery, 2010, pp. 121–130. ISBN: 978-1-4503-0050-6.

- [Mad13] A. Madry. “Navigating Central Path with Electrical Flows: From Flows to Matchings, and Back”. In: *Proceedings of the 2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. FOCS ’13. Washington, DC, USA: IEEE Computer Society, 2013, pp. 253–262. ISBN: 978-0-7695-5135-7.
- [Mad16] A. Madry. “Computing Maximum Flow with Augmenting Electrical Flows”. In: *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*. 2016, pp. 593–602.
- [Mil+98] G. L. Miller, S.-H. Teng, W. Thurston, and S. A. Vavasis. “Geometric separators for finite-element meshes”. In: *SIAM Journal on Scientific Computing* 19.2 (1998), pp. 364–386.
- [MN21] W. Maxwell and A. Nayyeri. “Generalized max-flows and min-cuts in simplicial complexes”. In: *arXiv preprint arXiv:2106.14116* (2021).
- [Moh91] B. Mohar. “Eigenvalues, diameter, and mean distance in graphs”. In: *Graphs and combinatorics* 7.1 (1991), pp. 53–64.
- [MT90] G. L. Miller and W. Thurston. “Separators in two and three dimensions”. In: *Proceedings of the twenty-second annual ACM symposium on Theory of computing*. 1990, pp. 300–309.
- [Mus+19] C. Musco, P. Netrapalli, A. Sidford, S. Ubaru, and D. P. Woodruff. “Spectrum Approximation Beyond Fast Matrix Multiplication: Algorithms and Hardness”. In: *arXiv:1704.04163 [cs, math]* (2019). arXiv: [1704.04163 \[cs, math\]](#).
- [Nie22] Z. Nie. “Matrix anti-concentration inequalities with applications”. In: *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. 2022, pp. 568–581.
- [ODO13] B. Osting, J. Darbon, and S. Osher. “Statistical ranking using the l1-norm on graphs”. In: *AIMS Journal on Inverse Problems and Imaging* 7.3 (2013), pp. 907–926.
- [OMV00] A. Ouorou, P. Mahey, and J.-P. Vial. “A Survey of Algorithms for Convex Multicommodity Flow Problems”. In: *Management Science* 46.1 (2000), pp. 126–147. ISSN: 0025-1909.
- [Pen16] R. Peng. “Approximate undirected maximum flows in $O(m \text{ polylog}(n))$ time”. In: *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA ’16. Arlington, Virginia: Society for Industrial and Applied Mathematics, 2016, pp. 1862–1867. ISBN: 9781611974331.
- [PS14] R. Peng and D. A. Spielman. “An efficient parallel solver for SDD linear systems”. In: *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*. ACM. 2014, pp. 333–342.
- [PV21] R. Peng and S. Vempala. “Solving sparse linear systems faster than matrix multiplication”. In: *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2021, pp. 504–521.

- [PW17] D. Pruša and T. Werner. “LP relaxations of some NP-hard problems are as hard as any LP”. In: *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM. 2017, pp. 1372–1382.
- [Ren01] J. Renegar. *A Mathematical View of Interior-Point Methods in Convex Optimization*. Society for Industrial and Applied Mathematics, 2001.
- [Ren88] J. Renegar. “A Polynomial-Time Algorithm, Based on Newton’s Method, for Linear Programming”. In: *Mathematical programming* 40.1 (1988), pp. 59–93.
- [Ren95] J. Renegar. “Incorporating Condition Measures into the Complexity Theory of Linear Programming”. In: *SIAM Journal on Optimization* 5.3 (1995), pp. 506–524. ISSN: 1052-6234.
- [Sav+25] A. Savostianov, M. T. Schaub, N. Guglielmi, and F. Tudisco. “Efficient Sparsification of Simplicial Complexes via Local Densities of States”. In: *arXiv preprint arXiv:2502.07558* (2025).
- [Sch+20] M. T. Schaub, A. R. Benson, P. Horn, G. Lippner, and A. Jadbabaie. “Random walks on simplicial complexes and the normalized hodge 1-laplacian”. In: *SIAM Review* 62.2 (2020), pp. 353–391.
- [She13] J. Sherman. “Nearly Maximum Flows in Nearly Linear Time”. In: *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*. 2013, pp. 263–269.
- [She17] J. Sherman. “Area-Convexity, L_∞ Regularization, and Undirected Multi-commodity Flow”. In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. 2017, pp. 452–460.
- [SS08] D. A. Spielman and N. Srivastava. “Graph sparsification by effective resistances”. In: *Proceedings of the fortieth annual ACM symposium on Theory of computing*. 2008, pp. 563–568.
- [ST04] D. A. Spielman and S.-H. Teng. “Nearly-Linear Time Algorithms for Graph Partitioning, Graph Sparsification, and Solving Linear Systems”. In: *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing*. STOC ’04. New York, NY, USA: Association for Computing Machinery, 2004, pp. 81–90. ISBN: 978-1-58113-852-8.
- [ST08] G. Shklarski and S. Toledo. “Rigidity in finite-element matrices: Sufficient conditions for the rigidity of structures and substructures”. In: *SIAM Journal on Matrix Analysis and Applications* 30.1 (2008), pp. 7–40.
- [ST14] D. A. Spielman and S.-H. Teng. “Nearly Linear Time Algorithms for Preconditioning and Solving Symmetric, Diagonally Dominant Linear Systems”. In: *SIAM Journal on Matrix Analysis and Applications* 35.3 (2014), pp. 835–885.
- [Str69] V. Strassen. “Gaussian elimination is not optimal”. In: *Numerische mathematik* 13.4 (1969), pp. 354–356.
- [Tan16] M. Tancer. “Recognition of collapsible complexes is NP-complete”. In: *Discrete & Computational Geometry* 55.1 (2016), pp. 21–38.

- [Ten10] S.-H. Teng. “The Laplacian paradigm: Emerging algorithms for massive graphs”. In: *Theory and Applications of Models of Computation: 7th Annual Conference, TAMC 2010, Prague, Czech Republic, June 7-11, 2010. Proceedings 7*. Springer. 2010, pp. 2–14.
- [Ton+03a] Y. Tong, S. Lombeyda, A. N. Hirani, and M. Desbrun. “Discrete multiscale vector field decomposition”. In: *ACM transactions on graphics (TOG)* 22.3 (2003), pp. 445–452.
- [Ton+03b] Y. Tong, S. Lombeyda, A. N. Hirani, and M. Desbrun. “Discrete multiscale vector field decomposition”. In: *ACM transactions on graphics (TOG)* 22.3 (2003), pp. 445–452.
- [TX98] L. Trevisan and F. Xhafa. “The parallel complexity of positive linear programming”. In: *Parallel Processing Letters* 8.04 (1998), pp. 527–533.
- [Vai89] P. M. Vaidya. “Speeding-up Linear Programming Using Fast Matrix Multiplication”. In: *30th Annual Symposium on Foundations of Computer Science*. IEEE Computer Society, 1989, pp. 332–337.
- [van+21] J. van den Brand, Y. T. Lee, Y. P. Liu, T. Saranurak, A. Sidford, Z. Song, and D. Wang. “Minimum Cost Flows, MDPs, and L1-Regression in Nearly Linear Time for Dense Instances”. In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2021. New York, NY, USA: Association for Computing Machinery, 2021, pp. 859–869. ISBN: 978-1-4503-8053-9.
- [Wan18] I.-L. Wang. “Multicommodity Network Flows: A Survey, Part I: Applications and Formulations”. In: *International Journal of Operations Research* 15.4 (2018), pp. 145–153.
- [Wil05] R. Williams. “A new algorithm for optimal 2-constraint satisfaction and its implications”. In: *Theoretical Computer Science* 348.2-3 (2005), pp. 357–365.
- [Wil18] V. V. Williams. “On some fine-grained questions in algorithms and complexity”. In: *Proceedings of the international congress of mathematicians: Rio de janeiro 2018*. World Scientific. 2018, pp. 3447–3487.
- [WW18] V. V. Williams and R. R. Williams. “Subcubic Equivalences Between Path, Matrix, and Triangle Problems”. In: *Journal of the ACM* 65.5 (2018), 27:1–27:38. ISSN: 0004-5411.
- [Xu+12] Q. Xu, Q. Huang, T. Jiang, B. Yan, W. Lin, and Y. Yao. “HodgeRank on random graphs for subjective video quality assessment”. In: *IEEE Transactions on Multimedia* 14.3 (2012), pp. 844–857.
- [YL17a] K. Ye and L.-H. Lim. “Cohomology of cryo-electron microscopy”. In: *SIAM Journal on Applied Algebra and Geometry* 1.1 (2017), pp. 507–535.
- [YL17b] K. Ye and L.-H. Lim. “Cohomology of cryo-electron microscopy”. In: *SIAM Journal on Applied Algebra and Geometry* 1.1 (2017), pp. 507–535.
- [Zom05] A. J. Zomorodian. *Topology for Computing*. Vol. 16. Cambridge university press, 2005.

Appendix A

Missing Proofs of Linear Algebra Facts

Proof of Fact 2.1.1. By the definition of image, $\text{Im}(\mathbf{A}) \supseteq \text{Im}(\mathbf{A}\mathbf{A}^\top)$. It suffices to show $\text{Im}(\mathbf{A}) \subseteq \text{Im}(\mathbf{A}\mathbf{A}^\top)$. Let \mathbf{x} be an arbitrary vector in $\text{Im}(\mathbf{A})$. Then, $\mathbf{x} = \mathbf{A}\mathbf{y}$ for some $\mathbf{y} \in \mathbb{R}^n$. Write $\mathbf{y} = \mathbf{y}_1 + \mathbf{y}_2$ such that $\mathbf{y}_1 \in \text{Im}(\mathbf{A}^\top)$ and $\mathbf{y}_2 \in \text{Ker}(\mathbf{A})$. Let \mathbf{z} satisfy $\mathbf{A}^\top \mathbf{z} = \mathbf{y}_1$. Then,

$$\mathbf{A}\mathbf{A}^\top \mathbf{z} = \mathbf{A}\mathbf{y}_1 = \mathbf{A}\mathbf{y} = \mathbf{x}.$$

Thus, $\mathbf{x} \in \text{Im}(\mathbf{A}\mathbf{A}^\top)$. □

Proof of Fact 2.1.2. Since $\mathbf{A} \preceq \mathbf{B}$, by definition, $\mathbf{C} := \mathbf{B} - \mathbf{A}$ is PSD, i.e., for all $\mathbf{x} \in \mathbb{R}^n$, we have $\mathbf{x}^\top \mathbf{C} \mathbf{x} \geq 0$. Then, for any $\mathbf{y} \in \mathbb{R}^m$, let $\mathbf{x} = \mathbf{V}^\top \mathbf{y}$. Substituting this into the quadratic form, we get:

$$\mathbf{y}^\top (\mathbf{V} \mathbf{C} \mathbf{V}^\top) \mathbf{y} = (\mathbf{V}^\top \mathbf{y})^\top \mathbf{C} (\mathbf{V}^\top \mathbf{y}) = \mathbf{x}^\top \mathbf{C} \mathbf{x} \geq 0,$$

This shows that $\mathbf{V} \mathbf{C} \mathbf{V}^\top = \mathbf{V}(\mathbf{B} - \mathbf{A})\mathbf{V}^\top$ is PSD. Consequently, $\mathbf{V} \mathbf{A} \mathbf{V}^\top \preceq \mathbf{V} \mathbf{B} \mathbf{V}^\top$. □

Proof of Fact 2.3.2. To minimize $\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$, we compute the gradient with respect to \mathbf{x} and set it to zero. This gives the normalized equation:

$$\mathbf{A}^\top (\mathbf{A}\mathbf{x}^* - \mathbf{b}) = \mathbf{0}.$$

Solving it gives $\mathbf{x}^* = (\mathbf{A}^\top \mathbf{A})^\dagger \mathbf{A}^\top \mathbf{b}$, thus

$$\mathbf{A}\mathbf{x}^* = \mathbf{A}(\mathbf{A}^\top \mathbf{A})^\dagger \mathbf{A}^\top \mathbf{b} = \Pi_{\text{Im}(\mathbf{A})} \mathbf{b}.$$

Substituting $\mathbf{A}\mathbf{x}^* = \Pi_{\text{Im}(\mathbf{A})} \mathbf{b}$ back to the ℓ_2 norm of residual gives

$$\|\mathbf{A}\mathbf{x}^* - \mathbf{b}\|_2^2 = \|(I - \Pi_{\text{Im}(\mathbf{A})})\mathbf{b}\|_2^2.$$

□

Proof of Fact 2.3.4. For the first equality, we have

$$\begin{aligned} \|\mathbf{A}^\top \mathbf{A}\mathbf{x} - \mathbf{A}^\top \mathbf{b}\|_{(\mathbf{A}^\top \mathbf{A})^\dagger} &= (\mathbf{A}^\top \mathbf{A}\mathbf{x} - \mathbf{A}^\top \mathbf{b})^\top (\mathbf{A}^\top \mathbf{A})^\dagger (\mathbf{A}^\top \mathbf{A}\mathbf{x} - \mathbf{A}^\top \mathbf{b}) \\ &= (\mathbf{A}\mathbf{x} - \mathbf{b})^\top \Pi_{\text{Im}(\mathbf{A})} (\mathbf{A}\mathbf{x} - \mathbf{b}) \\ &= (\Pi_{\text{Im}(\mathbf{A})} \mathbf{A}\mathbf{x} - \Pi_{\text{Im}(\mathbf{A})} \mathbf{b})^\top (\Pi_{\text{Im}(\mathbf{A})} \mathbf{A}\mathbf{x} - \Pi_{\text{Im}(\mathbf{A})} \mathbf{b}) \\ &= \|\Pi_{\text{Im}(\mathbf{A})} \mathbf{A}\mathbf{x} - \Pi_{\text{Im}(\mathbf{A})} \mathbf{b}\|_2^2. \end{aligned}$$

For the second equality, we have

$$\begin{aligned}
\|\mathbf{x} - \mathbf{x}^*\|_{\mathbf{A}^\top \mathbf{A}} &= (\mathbf{x} - \mathbf{x}^*)^\top \mathbf{A}^\top \mathbf{A} (\mathbf{x} - \mathbf{x}^*) \\
&= \mathbf{x}^\top \mathbf{A}^\top \mathbf{A} \mathbf{x} - 2\mathbf{x}^\top \mathbf{A}^\top \mathbf{A} \mathbf{x}^* + (\mathbf{x}^*)^\top \mathbf{A}^\top \mathbf{A} \mathbf{x}^* \\
&= \mathbf{x}^\top \mathbf{A}^\top \mathbf{A} \mathbf{x} - 2\mathbf{x}^\top \mathbf{A}^\top \Pi_{\text{Im}(\mathbf{A})} \mathbf{b} + \mathbf{b}^\top \Pi_{\text{Im}(\mathbf{A})}^\top \Pi_{\text{Im}(\mathbf{A})} \mathbf{b} \\
&= \|\mathbf{A} \mathbf{x} - \Pi_{\text{Im}(\mathbf{A})} \mathbf{b}\|_2.
\end{aligned}$$

□

Proof of Fact 2.3.5. We have

$$\begin{aligned}
\|\mathbf{A} \mathbf{x} - \mathbf{b}\|_2^2 &\leq \|\mathbf{A} \mathbf{x} - \Pi_{\text{Im}(\mathbf{A})} \mathbf{b}\|_2^2 + \|(\mathbf{I} - \Pi_{\text{Im}(\mathbf{A})}) \mathbf{b}\|_2^2 \\
&= \|\mathbf{A} \mathbf{x} - \Pi_{\text{Im}(\mathbf{A})} \mathbf{b}\|_2^2 + \|\mathbf{A} \mathbf{x}^* - \mathbf{b}\|_2^2 && \text{by Fact 2.3.2} \\
&\leq \epsilon^2 \|\Pi_{\text{Im}(\mathbf{A})} \mathbf{b}\|_2^2 + \|\mathbf{A} \mathbf{x}^* - \mathbf{b}\|_2^2, && \text{by Definition 2.3.3}
\end{aligned}$$

□

Proof of Fact 2.3.6. We start with proving the first statement. Since \mathbf{A} is PSD, eigen-decomposition gives $\mathbf{A} = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^\top$, where \mathbf{U} is an orthogonal matrix and $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_n)$ with $\lambda_i \geq 0$. Since $\mathbf{b} \in \text{Im}(\mathbf{A})$ and by Eq. (2.4), we have

$$\mathbf{b} = \Pi_{\text{Im}(\mathbf{A})} \mathbf{b} = \mathbf{A}^{1/2} \mathbf{A}^{\dagger/2} \mathbf{b}.$$

Then we rewrite

$$\begin{aligned}
\|\mathbf{A} \mathbf{Z} \mathbf{b} - \mathbf{b}\|_2 &= \left\| \mathbf{A}^{1/2} \left(\mathbf{A}^{1/2} \mathbf{Z} \mathbf{A}^{1/2} - \Pi_{\text{Im}(\mathbf{A})} \right) \mathbf{A}^{\dagger/2} \mathbf{b} \right\|_2 \\
&\leq \left\| \mathbf{A}^{1/2} \left(\mathbf{A}^{1/2} \mathbf{Z} \mathbf{A}^{1/2} - \Pi_{\text{Im}(\mathbf{A})} \right) \mathbf{A}^{\dagger/2} \right\|_2 \|\mathbf{b}\|_2 \\
&= \left\| \mathbf{U} \mathbf{A}^{1/2} \mathbf{U}^\top \left(\mathbf{A}^{1/2} \mathbf{Z} \mathbf{A}^{1/2} - \Pi_{\text{Im}(\mathbf{A})} \right) \mathbf{U} \mathbf{A}^{-1/2} \mathbf{U}^\top \right\|_2 \|\mathbf{b}\|_2 \\
&= \left\| \mathbf{A}^{1/2} \mathbf{U}^\top \left(\mathbf{A}^{1/2} \mathbf{Z} \mathbf{A}^{1/2} - \Pi_{\text{Im}(\mathbf{A})} \right) \mathbf{U} \mathbf{A}^{-1/2} \right\|_2 \|\mathbf{b}\|_2 \\
&\hspace{15em} \text{since } \mathbf{U} \text{ is orthogonal} \\
&\leq \left\| \mathbf{A}^{1/2} \right\|_2 \left\| \mathbf{A}^{-1/2} \right\|_2 \left\| \mathbf{A}^{1/2} \mathbf{Z} \mathbf{A}^{1/2} - \Pi_{\text{Im}(\mathbf{A})} \right\|_2 \|\mathbf{b}\|_2 \\
&= \frac{\sqrt{\lambda_{\max}}}{\sqrt{\lambda_{\min}}} \left\| \mathbf{A}^{1/2} \mathbf{Z} \mathbf{A}^{1/2} - \Pi_{\text{Im}(\mathbf{A})} \right\|_2 \|\mathbf{b}\|_2 \\
&\leq \sqrt{\kappa(\mathbf{A})} \left\| \mathbf{A}^{1/2} \mathbf{Z} \mathbf{A}^{1/2} - \Pi_{\text{Im}(\mathbf{A})} \right\|_2 \|\mathbf{b}\|_2 \\
&\leq \epsilon \sqrt{\kappa(\mathbf{A})} \|\mathbf{b}\|_2,
\end{aligned}$$

where the last step follows from the assumption $(1 - \epsilon) \mathbf{A}^\dagger \preceq \mathbf{Z} \preceq (1 + \epsilon) \mathbf{A}^\dagger$, and multiplying $\mathbf{A}^{1/2}$ on both sides gives

$$(1 - \epsilon) \Pi_{\text{Im}(\mathbf{A})} \preceq \mathbf{A}^{1/2} \mathbf{Z} \mathbf{A}^{1/2} \preceq (1 + \epsilon) \Pi_{\text{Im}(\mathbf{A})} \quad \Rightarrow \quad \left\| \mathbf{A}^{1/2} \mathbf{Z} \mathbf{A}^{1/2} - \Pi_{\text{Im}(\mathbf{A})} \right\|_2 \leq \epsilon.$$

To prove the second statement, again, since $\mathbf{b} \in \text{Im}(\mathbf{A})$, we can write $\mathbf{b} = \mathbf{A} \mathbf{x}$ for some \mathbf{x} :

$$\|\mathbf{A} \mathbf{Z} \mathbf{b} - \mathbf{b}\|_2 = \|\mathbf{A} \mathbf{Z} \mathbf{A} \mathbf{x} - \mathbf{A} \mathbf{x}\|_2 \leq \epsilon \|\mathbf{A} \mathbf{x}\|_2.$$

Then for any \mathbf{x} ,

$$(1 - \epsilon) \|\mathbf{A}\mathbf{x}\|_2 \leq \|\mathbf{AZA}\mathbf{x}\|_2 \leq (1 + \epsilon) \|\mathbf{A}\mathbf{x}\|_2.$$

As a result,

$$(1 - \epsilon)\mathbf{A} \preceq \mathbf{AZA} \preceq (1 + \epsilon)\mathbf{A}.$$

Multiplying \mathbf{A}^\dagger on both sides gives

$$(1 - \epsilon)\mathbf{A}^\dagger \mathbf{A} \mathbf{A}^\dagger \preceq \mathbf{A}^\dagger \mathbf{AZA} \mathbf{A}^\dagger \preceq (1 + \epsilon)\mathbf{A}^\dagger \mathbf{A} \mathbf{A}^\dagger.$$

By $\mathbf{A}^\dagger \mathbf{A} \mathbf{A}^\dagger = \mathbf{A}^\dagger$ and $\mathbf{A}^\dagger \mathbf{A} = \mathbf{A} \mathbf{A}^\dagger = \Pi_{\text{Im}(\mathbf{A})}$, we have

$$(1 - \epsilon)\mathbf{A}^\dagger \preceq \Pi_{\text{Im}(\mathbf{A})} \mathbf{Z} \Pi_{\text{Im}(\mathbf{A})} \preceq (1 + \epsilon)\mathbf{A}^\dagger.$$

□

Proof of Fact 2.3.7. We expand the left hand side,

$$\begin{bmatrix} \mathbf{x}^\top & \mathbf{y}^\top \end{bmatrix} \mathbf{A} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} = \mathbf{x}^\top \mathbf{A}(F, F) \mathbf{x} + 2\mathbf{x}^\top \mathbf{A}(F, C) \mathbf{y} + \mathbf{y}^\top \mathbf{A}(C, C) \mathbf{y}. \quad (\text{A.1})$$

Taking the derivative w.r.t. \mathbf{x} and setting it to be 0 gives that

$$2\mathbf{A}(F, F) \mathbf{x} + 2\mathbf{A}(F, C) \mathbf{y} = \mathbf{0}.$$

Plugging $\mathbf{x} = -\mathbf{A}(F, F)^\dagger \mathbf{A}(F, C) \mathbf{y}$ into Eq. (A.1),

$$\min_{\mathbf{x}} \begin{bmatrix} \mathbf{x}^\top & \mathbf{y}^\top \end{bmatrix} \mathbf{A} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} = \mathbf{y}^\top \mathbf{A}(C, C) \mathbf{y} - \mathbf{y}^\top \mathbf{A}(F, C) \mathbf{A}(F, F)^\dagger \mathbf{A}(F, C)^\top \mathbf{y} = \mathbf{y}^\top \text{Sc}(\mathbf{A})_C \mathbf{y}.$$

This completes the proof. □

Proof of Fact 2.3.9. We multiply the right-hand side:

$$\begin{aligned} & \begin{pmatrix} \mathbf{I} & \\ \mathbf{A}[C, F] \mathbf{A}[F, F]^\dagger & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{A}[F, F] & \\ & \text{Sc}[\mathbf{A}]_C \end{pmatrix} \begin{pmatrix} \mathbf{I} & \mathbf{A}[F, F]^\dagger \mathbf{A}[F, C] \\ & \mathbf{I} \end{pmatrix} \\ &= \begin{pmatrix} \mathbf{I} & \\ \mathbf{A}[C, F] \mathbf{A}[F, F]^\dagger & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{A}[F, F] & \mathbf{A}[F, C] \\ & \text{Sc}[\mathbf{A}]_C \end{pmatrix} \\ &= \begin{pmatrix} \mathbf{A}[F, F] & \mathbf{A}[F, C] \\ \mathbf{A}[C, F] & \mathbf{A}[C, F] \mathbf{A}[F, F]^\dagger \mathbf{A}[F, C] + \text{Sc}[\mathbf{A}]_C \end{pmatrix} \\ &= \begin{pmatrix} \mathbf{A}[F, F] & \mathbf{A}[F, C] \\ \mathbf{A}[C, F] & \mathbf{A}[C, C] \end{pmatrix} = \mathbf{A}. \end{aligned} \quad \square$$

Proof of Fact 2.3.10. By the definition of the Schur complement,

$$\begin{aligned} \text{Sc}[\mathbf{A}]_C &= \mathbf{B}_C \mathbf{B}_C^\top - \mathbf{B}_C \mathbf{B}_F^\top (\mathbf{B}_F \mathbf{B}_F^\top)^\dagger \mathbf{B}_F \mathbf{B}_C^\top \\ &= \mathbf{B}_C \left(\mathbf{I} - \mathbf{B}_F^\top (\mathbf{B}_F \mathbf{B}_F^\top)^\dagger \mathbf{B}_F \right) \mathbf{B}_C^\top = \mathbf{B}_C \Pi_{\text{Ker}(\mathbf{B}_F)} \mathbf{B}_C^\top. \end{aligned}$$

□

Appendix B

Missing Proofs of Chapter 4

B.1 Reducing 2-Complex Boundary LE to Combinatorial Laplacian LE

In this section, we formally state Theorem 1.4.5 as below and provide a proof. Recall that we use $\sigma_{\min}(\mathbf{A})$ to denote the smallest non-zero eigenvalue.

Theorem B.1.1. *Let $\mathbf{L}_1 = \partial_1^\top \partial_1 + \partial_2 \partial_2^\top \in \mathbb{R}^{m \times m}$ be the combinatorial Laplacian of a 2-complex. Let $\mathbf{d} \in \mathbb{Z}^m$. Suppose we can solve LEA $(\mathbf{L}_1, \mathbf{d}, \epsilon)$ in time $\tilde{O}(\text{nnz}(\mathbf{L}_1)^c)$ where $c \geq 1$ is a constant. Then, we can solve LEA $(\partial_2, \mathbf{d}, \delta)$ in time $\tilde{O}(\text{nnz}(\partial_2)^c)$ by choosing*

$$\epsilon < \delta \frac{\sigma_{\min}(\mathbf{L}_1)^{1/2}}{\sigma_{\max}(\partial_2)^2} \frac{1}{\|\mathbf{d}\|_2}.$$

Proof. Suppose \mathbf{x}_1 satisfies

$$\|\mathbf{L}_1 \mathbf{x}_1 - \Pi_{\text{Im}(\mathbf{L}_1)} \mathbf{d}\|_2 \leq \epsilon \|\Pi_{\text{Im}(\mathbf{L}_1)} \mathbf{d}\|_2.$$

By our assumption, we can compute \mathbf{x}_1 in time $\tilde{O}(\text{nnz}(\mathbf{L}_1)^c)$. We choose

$$\mathbf{f} = \partial_2^\top \mathbf{x}_1.$$

We claim that \mathbf{f} solves LEA $(\partial_2, \mathbf{d}, \delta)$. Since $\partial_1 \partial_2 = \mathbf{0}$, we have $\mathbf{L}_1^\dagger = (\partial_1^\top \partial_1)^\dagger + (\partial_2 \partial_2^\top)^\dagger$ and $\Pi_{\text{Im}(\partial_2)} (\partial_1^\top \partial_1)^\dagger \mathbf{d} = \mathbf{0}$. Then,

$$\begin{aligned} \|\partial_2 \mathbf{f} - \Pi_{\text{Im}(\partial_2)} \mathbf{d}\|_2 &= \|\mathbf{x}_1 - (\partial_2 \partial_2^\top)^\dagger \mathbf{d}\|_{(\partial_2 \partial_2^\top)^2} \\ &\leq \sigma_{\max}(\partial_2) \|\mathbf{x}_1 - (\partial_2 \partial_2^\top)^\dagger \mathbf{d}\|_{\partial_2 \partial_2^\top} \\ &= \sigma_{\max}(\partial_2) \|\Pi_{\text{Im}(\partial_2)} (\mathbf{x}_1 - (\partial_2 \partial_2^\top)^\dagger \mathbf{d})\|_{\partial_2 \partial_2^\top} \\ &= \sigma_{\max}(\partial_2) \|\Pi_{\text{Im}(\partial_2)} (\mathbf{x}_1 - \mathbf{L}_1^\dagger \mathbf{d})\|_{\partial_2 \partial_2^\top} \\ &= \sigma_{\max}(\partial_2) \|\mathbf{x}_1 - \mathbf{L}_1^\dagger \mathbf{d}\|_{\partial_2 \partial_2^\top} \\ &\leq \sigma_{\max}(\partial_2) \|\mathbf{x}_1 - \mathbf{L}_1^\dagger \mathbf{d}\|_{\mathbf{L}_1} \\ &\leq \frac{\sigma_{\max}(\partial_2)}{\sigma_{\min}(\mathbf{L}_1)^{1/2}} \|\mathbf{L}_1 \mathbf{x}_1 - \Pi_{\text{Im}(\mathbf{L}_1)} \mathbf{d}\|_2 \end{aligned}$$

$$\begin{aligned}
&\leq \frac{\sigma_{\max}(\partial_2)}{\sigma_{\min}(\mathbf{L}_1)^{1/2}} \epsilon \|\Pi_{\text{Im}(\mathbf{L}_1)} \mathbf{d}\|_2 \\
&\leq \frac{\delta}{\sigma_{\max}(\partial_2)} \\
&\leq \delta \|\Pi_{\text{Im}(\partial_2)} \mathbf{d}\|_2, \quad \text{by Claim 4.6.3}
\end{aligned}$$

□

B.2 Connections With Interior Point Methods

We now show that in order to solve a generalized maxflow problem in a 2-complex flow network using an Interior Point Method (IPM), it suffices to be able to apply the pseudo-inverse of $\partial_2 \mathbf{W} \partial_2^\top$ for diagonal positive weight matrices \mathbf{W} (and this problem is essentially equivalent to applying the pseudo-inverse of the combinatorial Laplacian of the complex, c.f. Appendix B.1). We sketch how these pseudo-inverse problems arise when solving a generalized maxflow using IPM, which is motivated by [Mad16]. For the more curious readers, we recommend the book [Ren01] for a complete view of general IPM algorithms.

Given a 2-complex flow network \mathcal{K} with m edges and t triangles, a non-negative capacity vector $\mathbf{c} \in \mathbb{R}_{\geq 0}^t$, and a demand vector $\gamma \in \mathbb{R}^m$ such that $\gamma \in \text{Im}(\partial_2)$. The γ -maxflow problem is formulated by the following linear programming:

$$\begin{aligned}
&\max_{F, \mathbf{f}} F \\
&\text{s.t. } \partial_2 \mathbf{f} = F \gamma \\
&\quad -\mathbf{c} \leq \mathbf{f} \leq \mathbf{c}
\end{aligned} \tag{B.1}$$

The γ -maxflow in 2-complex flow networks is a generalization of s - t maxflow in graphs. The first constraint encodes the conservation of flows for edges in \mathcal{K} . And the second constraint forces the flow on triangles to satisfy the capacity constraints.

We call F the flow value of \mathbf{f} when $\partial_2 \mathbf{f} = F \gamma$. We assume that the optimal flow value F^* is known by IPM algorithms, which can be estimated by the binary search.

The main idea of IPM is to get rid of inequality constraints by using barrier functions, and then apply Newton's method to a sequence of equality constrained problems. The most widely used barrier function is the logarithmic barrier function, which in the γ -maxflow problem gives

$$V(\mathbf{f}) = \sum_{\Delta \in [t]} -\log(\mathbf{c}(\Delta) - \mathbf{f}(\Delta)) - \log(\mathbf{c}(\Delta) + \mathbf{f}(\Delta)).$$

Then for a given $0 \leq \alpha < 1$, we define the following Barrier Problem:

$$\begin{aligned}
&\min_{\mathbf{f}} V(\mathbf{f}) \\
&\text{s.t. } \partial_2 \mathbf{f} = \alpha F^* \gamma
\end{aligned} \tag{B.2}$$

We start with zero flow, i.e., $\alpha_0 = 0$, and then increase $\alpha_{i+1} = \alpha_i + \alpha'$ gradually in each iteration to make progress. Given a small enough α' , each iteration is composed of a *progress step* and a *centering step*.

Progress Step. We first take a progress step by making a Newton step to Problem (B.2) at the current point \mathbf{f} , while increasing the flow value by α' , which gives

$$\begin{aligned} \min_{\boldsymbol{\delta}} \quad & \mathbf{g}^\top(\mathbf{f})\boldsymbol{\delta} + \frac{1}{2}\boldsymbol{\delta}^\top \mathbf{H}(\mathbf{f})\boldsymbol{\delta} \\ \text{s.t.} \quad & \partial_2 \boldsymbol{\delta} = \alpha' F^* \gamma \end{aligned} \tag{B.3}$$

where $\mathbf{g}(\mathbf{f})$ and $\mathbf{H}(\mathbf{f})$ are the gradient and Hessian of V at the current point \mathbf{f} , respectively.

Problem (B.3) has the Lagrangian

$$\mathcal{L}(\boldsymbol{\delta}, \mathbf{x}) = \mathbf{g}^\top(\mathbf{f})\boldsymbol{\delta} + \frac{1}{2}\boldsymbol{\delta}^\top \mathbf{H}(\mathbf{f})\boldsymbol{\delta} + \mathbf{x}^\top (\alpha' F^* \gamma - \partial_2 \boldsymbol{\delta}).$$

Using the optimality condition, we have

$$\nabla_{\boldsymbol{\delta}} \mathcal{L}(\boldsymbol{\delta}, \mathbf{x}) = \mathbf{g}(\mathbf{f}) + \mathbf{H}(\mathbf{f})\boldsymbol{\delta} - \partial_2^\top \mathbf{x} = \mathbf{0},$$

which gives

$$\boldsymbol{\delta} = \mathbf{H}^{-1}(\mathbf{f})(\partial_2^\top \mathbf{x} - \mathbf{g}(\mathbf{f})).$$

Multiplying ∂_2 in both sides and using the constraint $\partial_2 \boldsymbol{\delta} = \alpha' F^* \gamma$, we obtain

$$\partial_2 \mathbf{H}^{-1}(\mathbf{f}) \partial_2^\top \mathbf{x} = \partial_2 \mathbf{H}^{-1}(\mathbf{f}) \mathbf{g}(\mathbf{f}) + \alpha' F^* \gamma.$$

Thus, we have shown that it suffices to apply the pseudo-inverse of $\partial_2 \mathbf{H}^{-1}(\mathbf{f}) \partial_2^\top$ to solve \mathbf{x} and $\boldsymbol{\delta}$:

$$\begin{aligned} \mathbf{x} &= (\partial_2 \mathbf{H}^{-1}(\mathbf{f}) \partial_2^\top)^\dagger (\partial_2 \mathbf{H}^{-1}(\mathbf{f}) \mathbf{g}(\mathbf{f}) + \alpha' F^* \gamma), \\ \boldsymbol{\delta} &= \mathbf{H}^{-1}(\mathbf{f}) \partial_2^\top (\partial_2 \mathbf{H}^{-1}(\mathbf{f}) \partial_2^\top)^\dagger (\partial_2 \mathbf{H}^{-1}(\mathbf{f}) \mathbf{g}(\mathbf{f}) + \alpha' F^* \gamma) - \mathbf{H}^{-1}(\mathbf{f}) \mathbf{g}(\mathbf{f}). \end{aligned}$$

Centering Step. We then take a centering step by making a Newton step to Problem (B.2) at the updated point of $\tilde{\mathbf{f}} \stackrel{\text{def}}{=} \mathbf{f} + \boldsymbol{\delta}$ without increasing the flow value, which gives

$$\begin{aligned} \min_{\tilde{\boldsymbol{\delta}}} \quad & \mathbf{g}^\top(\tilde{\mathbf{f}})\tilde{\boldsymbol{\delta}} + \frac{1}{2}\tilde{\boldsymbol{\delta}}^\top \mathbf{H}(\tilde{\mathbf{f}})\tilde{\boldsymbol{\delta}} \\ \text{s.t.} \quad & \partial_2 \tilde{\boldsymbol{\delta}} = \mathbf{0} \end{aligned} \tag{B.4}$$

Similar to the progress step, it suffices to apply the pseudo-inverse of $\partial_2 \mathbf{H}^{-1}(\tilde{\mathbf{f}}) \partial_2^\top$ to solve $\tilde{\boldsymbol{\delta}}$:

$$\tilde{\boldsymbol{\delta}} = \left(\mathbf{H}^{-1}(\tilde{\mathbf{f}}) \partial_2^\top (\partial_2 \mathbf{H}^{-1}(\tilde{\mathbf{f}}) \partial_2^\top)^\dagger \partial_2 - \mathbf{I} \right) \mathbf{H}^{-1}(\tilde{\mathbf{f}}) \mathbf{g}(\tilde{\mathbf{f}}).$$

Appendix C

Missing Proofs of Chapter 6

Proof of Lemma 6.5.5. Since $\mathbf{b} \in \text{Im}(\mathbf{L}_1^{\text{up}})$, we know $\mathbf{b}_F \in \text{Im}(\mathbf{L}_1^{\text{up}}(F, F))$ and $\mathbf{b}_F - \mathbf{L}_1^{\text{up}}(F, C)\tilde{\mathbf{x}}_C \in \text{Im}(\mathbf{L}_1^{\text{up}}(F, F))$. We can apply the solver UPLAPFSOLVER to these two vectors. By the statement assumption, we can write $\text{UPLAPFSOLVER}(\mathbf{b}_F) = \mathbf{L}_1^{\text{up}}(F, F)^\dagger \mathbf{b}_F + \mathbf{y}$, where $\partial_2[F, :]^\top \mathbf{y} = \mathbf{0}$. Then,

$$\mathbf{h} = \mathbf{b}_C - \mathbf{L}_1^{\text{up}}(C, F) \cdot \text{UPLAPFSOLVER}(\mathbf{b}_F) = \mathbf{b}_C - \mathbf{L}_1^{\text{up}}(C, F) \cdot \mathbf{L}_1^{\text{up}}(F, F)^\dagger \mathbf{b}_F.$$

We first show that $\mathbf{h} \in \text{Im}(\text{Sc}(\mathbf{L}_1^{\text{up}})_C)$ so that we can apply the solver SCHURSOLVER to \mathbf{h} and obtain a vector $\tilde{\mathbf{x}}_C$ satisfying $\|\text{Sc}(\mathbf{L}_1^{\text{up}})_C \tilde{\mathbf{x}}_C - \mathbf{h}\|_2 \leq \delta \|\mathbf{h}\|_2$. Since $\mathbf{b} \in \text{Im}(\mathbf{L}_1^{\text{up}})$, there exists $\mathbf{x} = \begin{bmatrix} \mathbf{x}_F \\ \mathbf{x}_C \end{bmatrix}$ such that

$$\begin{aligned} \begin{bmatrix} \mathbf{b}_F \\ \mathbf{b}_C \end{bmatrix} &= \begin{bmatrix} \mathbf{I} & \\ \mathbf{L}_1^{\text{up}}(C, F)\mathbf{L}_1^{\text{up}}(F, F)^\dagger & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{L}_1^{\text{up}}(F, F) & \\ & \text{Sc}(\mathbf{L}_1^{\text{up}})_C \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{L}_1^{\text{up}}(F, F)^\dagger \mathbf{L}_1^{\text{up}}(F, C) \\ & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{x}_F \\ \mathbf{x}_C \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{I} & \\ \mathbf{L}_1^{\text{up}}(C, F)\mathbf{L}_1^{\text{up}}(F, F)^\dagger & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{L}_1^{\text{up}}(F, F) & \\ & \text{Sc}(\mathbf{L}_1^{\text{up}})_C \end{bmatrix} \begin{bmatrix} \mathbf{x}_F + \mathbf{L}_1^{\text{up}}(F, F)^\dagger \mathbf{L}_1^{\text{up}}(F, C)\mathbf{x}_C \\ \mathbf{x}_C \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{I} & \\ \mathbf{L}_1^{\text{up}}(C, F)\mathbf{L}_1^{\text{up}}(F, F)^\dagger & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{L}_1^{\text{up}}(F, F)\mathbf{x}_F + \mathbf{L}_1^{\text{up}}(F, C)\mathbf{x}_C \\ \text{Sc}(\mathbf{L}_1^{\text{up}})_C \mathbf{x}_C \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{L}_1^{\text{up}}(F, F)\mathbf{x}_F + \mathbf{L}_1^{\text{up}}(F, C)\mathbf{x}_C \\ \mathbf{L}_1^{\text{up}}(C, F)\mathbf{x}_F + \mathbf{L}_1^{\text{up}}(C, F)\mathbf{L}_1^{\text{up}}(F, F)^\dagger \mathbf{L}_1^{\text{up}}(F, C)\mathbf{x}_C + \text{Sc}(\mathbf{L}_1^{\text{up}})_C \mathbf{x}_C \end{bmatrix}. \end{aligned} \tag{C.1}$$

Here, the third equality holds since

$$\mathbf{L}_1^{\text{up}}(F, F)\mathbf{L}_1^{\text{up}}(F, F)^\dagger \mathbf{L}_1^{\text{up}}(F, C) = \mathbf{\Pi}_{\text{Im}(\partial_2[F, :])} \mathbf{L}_1^{\text{up}}(F, C) = \mathbf{L}_1^{\text{up}}(F, C),$$

and the fourth equality holds similarly by using symmetry. Thus,

$$\begin{aligned} \mathbf{b}_C &= \mathbf{L}_1^{\text{up}}(C, F)\mathbf{x}_F + \mathbf{L}_1^{\text{up}}(C, F)\mathbf{L}_1^{\text{up}}(F, F)^\dagger (\mathbf{b}_F - \mathbf{L}_1^{\text{up}}(F, F)\mathbf{x}_F) + \text{Sc}(\mathbf{L}_1^{\text{up}})_C \mathbf{x}_C \\ &= \mathbf{L}_1^{\text{up}}(C, F)\mathbf{x}_F + \mathbf{L}_1^{\text{up}}(C, F)\mathbf{L}_1^{\text{up}}(F, F)^\dagger \mathbf{b}_F - \mathbf{L}_1^{\text{up}}(C, F)\mathbf{x}_F + \text{Sc}(\mathbf{L}_1^{\text{up}})_C \mathbf{x}_C \\ &= \mathbf{L}_1^{\text{up}}(C, F)\mathbf{L}_1^{\text{up}}(F, F)^\dagger \mathbf{b}_F + \text{Sc}(\mathbf{L}_1^{\text{up}})_C \mathbf{x}_C. \end{aligned}$$

That is, $\mathbf{h} \in \text{Im}(\text{Sc}(\mathbf{L}_1^{\text{up}})_C)$.

Next we look at $\mathbf{L}_1^{\text{up}} \tilde{\mathbf{x}}$. Let $\boldsymbol{\delta} = \text{Sc}(\mathbf{L}_1^{\text{up}})_C \tilde{\mathbf{x}}_C - \mathbf{h}$. We replace \mathbf{x}_F and \mathbf{x}_C in Equation

(C.1) with $\tilde{\mathbf{x}}_F$ and $\tilde{\mathbf{x}}_C$:

$$\begin{aligned} \mathbf{L}_1^{\text{up}} \tilde{\mathbf{x}} &= \begin{bmatrix} \mathbf{I} & \\ \mathbf{L}_1^{\text{up}}(C, F) \mathbf{L}_1^{\text{up}}(F, F)^\dagger & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{L}_1^{\text{up}}(F, F) \tilde{\mathbf{x}}_F + \mathbf{L}_1^{\text{up}}(F, C) \tilde{\mathbf{x}}_C \\ \text{Sc}(\mathbf{L}_1^{\text{up}})_C \tilde{\mathbf{x}}_C \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{b}_F \\ \mathbf{L}_1^{\text{up}}(C, F) \mathbf{L}_1^{\text{up}}(F, F)^\dagger \mathbf{b}_F + \text{Sc}(\mathbf{L}_1^{\text{up}})_C \tilde{\mathbf{x}}_C \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{b}_F \\ \mathbf{b}_C + \boldsymbol{\delta} \end{bmatrix}. \end{aligned}$$

Then,

$$\begin{aligned} \|\mathbf{L}_1^{\text{up}} \tilde{\mathbf{x}} - \mathbf{b}\|_2 &= \|\boldsymbol{\delta}\|_2 \leq \delta \|\mathbf{h}\|_2 \leq \delta (\|\mathbf{b}_C\|_2 + \|\mathbf{L}_1^{\text{up}}(C, F) \mathbf{L}_1^{\text{up}}(F, F)^\dagger\|_2 \|\mathbf{b}_F\|_2) \\ &\leq \delta (1 + \|\mathbf{L}_1^{\text{up}}(C, F) \mathbf{L}_1^{\text{up}}(F, F)^\dagger\|_2) \|\mathbf{b}\|_2 \leq \epsilon \|\mathbf{b}\|_2, \end{aligned}$$

where the last inequality is by our setting of δ .

We compute $\tilde{\mathbf{x}}$ by two calls of UPLAPFSOLVER and one call of SCHURSOLVER and $O(1)$ matrix-vector multiplications and vector-vector additions. Thus, the total runtime is $O(t_1(m_F) + t_2(m_C) + \text{the number of non-zeros in } \mathbf{L}_1^{\text{up}})$. \square

Curriculum Vitae

PERSONAL DATA

Name: Ming Ding
Date of Birth: February 16, 1995
Place of Birth: Anhui, China
Citizen of: China

EDUCATION

2020 – 2025	ETH Zurich, Zurich, Switzerland Ph.D. in Theoretical Computer Science
2017 – 2020	Shanghai Jiao Tong University, Shanghai, China M.Sc. in Information and Communication Engineering
2015 – 2017	Ecole Centrale Paris, Paris, France Diplôme d'Ingénieur
2013 – 2017	Shanghai Jiao Tong University, Shanghai, China B.Sc. in Information Engineering