

Fast Algorithms for Separable Linear Programs

Sally Dong*
University of Washington
sallyqd@uw.edu

Gramoz Goranci†
University of Vienna
gramoz.goranci@univie.ac.at

Lawrence Li‡
University of Toronto
lawrenceli@cs.toronto.edu

Sushant Sachdeva§
University of Toronto
sachdeva@cs.toronto.edu

Guanghao Ye¶
Massachusetts Institute of Technology
ghye@mit.edu

November 5, 2023

Abstract

In numerical linear algebra, considerable effort has been devoted to obtaining faster algorithms for linear systems whose underlying matrices exhibit structural properties. A prominent success story is the method of generalized nested dissection [Lipton-Rose-Tarjan'79] for separable matrices. On the other hand, the majority of recent developments in the design of efficient linear program (LP) solvers have not leveraged the ideas underlying these faster linear system solvers nor exploited the separable structure of the constraint matrix.

In this work, we consider LPs of the form $\min_{\mathbf{A}\mathbf{x}=\mathbf{b}, \ell \leq \mathbf{x} \leq \mathbf{u}} \mathbf{c}^\top \mathbf{x}$, where the graphical support of the constraint matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$ is n^α -separable. We present an $\tilde{O}((m + m^{1/2+2\alpha}) \log(1/\epsilon))$ -time algorithm for solving these LPs to ϵ relative accuracy.

Our new solver has two important implications: for the k -multicommodity flow problem on planar graphs, we obtain an $\tilde{O}(k^{5/2}m^{3/2} \log(1/\epsilon))$ -time algorithm; and when the support of \mathbf{A} is n^α -separable with $\alpha \leq 1/4$, our runtime of $\tilde{O}(m \log(1/\epsilon))$ is nearly optimal. The latter significantly improves upon the natural approach of combining interior point methods and nested dissection, whose time complexity is lower bounded by $\Omega(\sqrt{m}(m + m^{\alpha\omega})) = \Omega(m^{3/2})$, where $\omega \approx 2.373$ is the matrix multiplication exponent. Lastly, our solver can be applied to low-treewidth LPs to recover the results of [DLY21,GS22] while using significantly simpler data structure machinery.

*Sally Dong was supported by NSERC PGS-D 557770-2021, and NSF awards CCF-1749609, DMS-1839116, DMS-2023166, CCF-2105772.

†Faculty of Computer Science, University of Vienna, Vienna, Austria.

‡Lawrence Li was supported by NSERC Discovery Grant RGPIN-2018-06398 and Ontario Early Researcher Award (ERA) ER21-16-283 awarded to SS.

§Sushant Sachdeva was supported by an NSERC Discovery Grant RGPIN-2018-06398, an Ontario Early Researcher Award (ERA) ER21-16-283, and a Sloan Research Fellowship.

¶Guanghao Ye was supported by NSF awards CCF-1955217 and DMS-2022448.

Contents

1	Introduction	1
1.1	Previous work	2
1.2	Technical overview	4
2	Preliminaries	5
3	Overview of RIPM framework	6
3.1	Robust interior point method	6
3.2	Projection operators	6
3.2.1	Operators on a tree	7
3.3	Implicit representations of the solution	8
3.4	Solution approximation	9
3.5	Main theorem for the RIPM framework	9
4	From separator tree to projection operators	10
4.1	Separator tree	10
4.2	Nested dissection using a separator tree	12
4.3	Projection operators definition	12
4.4	Maintenance of projection operators	14
4.5	Projection operator complexities	17
5	Proofs of main theorems	19
5.1	Proof of Theorem 1.1	19
5.2	Proof of Theorem 1.2	21
6.1	Proof of Theorem 1.3	22
A	Robust interior point method	26
B	Maintaining the implicit representation	28
E	Maintaining vector approximation	35
E.1	Reduction to change detection	36
E.2	From change detection to sketch maintenance	38
E.3	Sketch maintenance	41
F.1	Proof of 3.3	43

1 Introduction

Linear programming (LP) is a widely used technique for solving a broad range of problems that emerge in optimization, operations research, and computer science, among others. LP solvers have been a subject of research for many years, both from theoretical as well as practical perspectives. This has led to the development of several algorithmic gems such as the Simplex algorithm [12], ellipsoid algorithm [43] and interior point method [39], to name a few.

Fast solvers for LPs via interior point methods have received considerable attention recently, especially in the theoretical computer science community. A series of improvements culminated in the recent breakthrough work of Cohen, Lee and Song [9], which shows that any linear program $\min_{\mathbf{Ax}=\mathbf{b}, \mathbf{l}\leq\mathbf{x}\leq\mathbf{u}} \mathbf{c}^\top \mathbf{x}$ with n constraints and m variables can be solved in $\tilde{O}(m^\omega \log(1/\varepsilon))$ time, where ε is the accuracy parameter and $\omega \approx 2.3715$ is the matrix multiplication exponent [18, 70]. When \mathbf{A} is a dense matrix, their running time is almost optimal as it nearly matches the $O(m^\omega)$ algorithm for solving a linear system $\mathbf{Ax} = \mathbf{b}$, which is a sub-problem of linear programming. However, the case when \mathbf{A} is a sparse matrix is equally important, since the constraint matrices of many LP instances that arise in practical applications happen to be sparse.

A widely-used method for identifying structures in a sparse matrix \mathbf{A} involves associating a graph with its non-zero pattern, which captures the interactions between the equations in the system. In this paper, we are interested in when said graph is *separable*; we use a weighted-version of the definition as is common in literature, such as [32]:

DEFINITION 1.1. (SEPARABLE GRAPHS) *A (hyper-)graph $G = (V, \mathcal{E})$ is n^α -separable for some $\alpha \in [0, 1]$ if there exists constants $b \in (0, 1)$ and $c > 0$, such that for any vertex weight assignment w , the vertices of G can be partitioned into S , A and B such that $|S| \leq c \cdot |V|^\alpha$, there are no edges between A and B , and $\max\{w(A), w(B)\} \leq b \cdot w(V)$. We call S the (b -)balanced vertex separator of H (with respect to w).*

A notable case is $\alpha = 1/2$, which includes the family of planar and bounded-genus graphs [50]. It has also been empirically observed that road networks have separators of size $n^{1/3}$ [13, 59].

Building upon the seminal work of George [25], Lipton Tarjan and Rose [49] introduced the generalized nested dissection algorithm, which solves the linear system $\mathbf{Ax} = \mathbf{b}$ in $O(m + m^{\alpha\omega})$ time when \mathbf{A} is a symmetric-positive definite matrix and the associated graph is $O(n^\alpha)$ -separable. When $\alpha < 1$, this algorithm outperforms the canonical $O(m^\omega)$ -time algorithm for general linear systems. Motivated by this, we ask the natural question of how to leverage the structures in the constraints to speed up linear programming:

Are there faster LP solvers for the class of problems where the constraint matrix \mathbf{A} can be represented by an $O(n^\alpha)$ -separable graph?

Given the constraint matrix \mathbf{A} , [49] associates with it the unique graph whose adjacency matrix has the same non-zero pattern as \mathbf{A} . In the context of linear programs, we define the *dual graph* $G_{\mathbf{A}}$ of a constraint matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$ to be the hypergraph with vertex set $\{1, \dots, n\}$ corresponding to the rows of \mathbf{A} and hyper-edges $\{e_1, \dots, e_m\}$, such that vertex i is in hyperedge e_j if $\mathbf{A}_{i,j} \neq 0$.

In this paper, we present a faster solver for LPs whose dual graph is separable.

THEOREM 1.1. *Given a linear program $\min \{\mathbf{c}^\top \mathbf{x} : \mathbf{Ax} = \mathbf{b}, \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}\}$, where $\mathbf{A} \in \mathbb{R}^{n \times m}$ is a full-rank matrix with $n \leq m$, suppose the dual graph $G_{\mathbf{A}}$ is $O(n^\alpha)$ -separable with $\alpha \in [0, 1]$, and a balanced separator is computable in $T(n)$ time.*

Suppose that r is the inner radius of the polytope, namely, there is \mathbf{x} such that $\mathbf{Ax} = \mathbf{b}$ and $\mathbf{l} + r \leq \mathbf{x} \leq \mathbf{u} - r$. Let $L = \|\mathbf{c}\|_2$ and $R = \|\mathbf{u} - \mathbf{l}\|_2$. Then, for any $0 < \varepsilon \leq 1/2$, we can find a feasible \mathbf{x} with high probability such that

$$\mathbf{c}^\top \mathbf{x} \leq \min_{\mathbf{Ax}=\mathbf{b}, \mathbf{l}\leq\mathbf{x}\leq\mathbf{u}} \mathbf{c}^\top \mathbf{x} + \varepsilon \cdot LR,$$

in time

$$\tilde{O}\left((m + m^{1/2+2\alpha}) \cdot \log(R/(r\varepsilon)) + T(n)\right).$$

Our result should be compared against the natural $\tilde{O}(m^{1/2}(m + m^{\alpha\omega}))$ runtime, which directly follows from the fact that IPM-based methods require $\tilde{O}(\sqrt{m})$ iterations, each of which can be implemented in $O(m + m^{\alpha\omega})$

time using the nested dissection [49, 2] algorithm. For linear programs whose dual graph are $O(n^\alpha)$ -separable with $\alpha \leq 1/4$, our algorithm achieves $\tilde{O}(m \log(1/\varepsilon))$ time, which is optimal up to poly-logarithmic factors.

We would like to emphasize that $\tilde{O}(m + m^{1/2+2\alpha})$ represents a natural barrier for the (robust) IPM-based approaches. At a high level, each iteration of IPM involves performing matrix operations using the inverse of an $O(m^\alpha) \times O(m^\alpha)$ matrix¹. Even if one is given access to said inverse, multiplying a vector against it takes at least $\Omega(m^{2\alpha})$, showing that improving upon the $m^{2\alpha}$ factor will require significantly new ideas in the design and analysis of robust Interior Point Methods. Obtaining an LP solver whose time complexity is $\tilde{O}(m + m^{\alpha\omega})$, which would in turn nearly match the time complexity for solving linear systems with recursively separable structure, remains an outstanding open problem [26].

An immediate application of Theorem 1.1 is a faster algorithm for solving the (fractional) k -commodity flow problem on planar graphs to high accuracy. For general sparse graphs, an $\tilde{O}((km)^\omega)$ time algorithm for this problem follows by the recent linear program solvers that run in matrix multiplication time [9, 65]. It is known that solving the k -commodity flow problem is as hard as linear programming [34, 14], suggesting that additional structural assumptions on the input graph are necessary to obtain faster algorithms. As shown in the theorem below, our result achieves a polynomial speed-up when the input graph is planar.

THEOREM 1.2. *Given a minimum-cost k -multicommodity flow problem on a planar graph $G = (V, E)$ on n vertices and m edges, with edge-vertex incidence matrix \mathbf{B} , integer edge capacities $\mathbf{u} \in \mathbb{R}_{\geq 0}^E$, integer costs $\mathbf{c}_1, \dots, \mathbf{c}_k \in \mathbb{R}^E$ and integer demands $\mathbf{d}_1, \dots, \mathbf{d}_k \in \mathbb{R}^E$ for each commodity, we can solve the LP*

$$(1.1) \quad \begin{aligned} \min \quad & \sum_{i=1}^k \mathbf{c}_i^\top \mathbf{f}_i \\ \text{s.t.} \quad & \mathbf{B}^\top \mathbf{f}_i = \mathbf{d}_i \quad \forall i \in [k] \\ & \sum_{i=1}^k \mathbf{f}_i \leq \mathbf{u} \\ & \mathbf{f}_i \geq \mathbf{0} \quad \forall i \in [k] \end{aligned}$$

to ϵ accuracy in $\tilde{O}(k^{2.5}m^{1.5} \log(M/\epsilon))$ time, where M is an upper on the absolute values of $\mathbf{u}, \mathbf{c}, \mathbf{d}$.

Our main result also has the important advantage of recovering and simplifying the recent work by Dong, Lee and Ye [16] and Gu and Song [28] who obtain fast solvers for LPs whose constraint matrix has bounded treewidth.

THEOREM 1.3. *Suppose we have a linear program with the same setup as Theorem 1.1, and we are given a tree-decomposition of the dual graph $G_{\mathbf{A}^2}$ of width τ . Then we can solve the linear program in time*

$$\tilde{O}(m\tau^2 \log(R/(\epsilon r))) \text{ or } \tilde{O}(m\tau^{(\omega+1)/2} \log(R/(\epsilon r))).$$

1.1 Previous work It is known that 2-commodity flow is as hard as linear programming [34]. Recently, [14] showed a linear-time reduction from linear programs to *sparse* k -commodity flow instance, indicating that sparse k -commodity flow instances are hard to solve. This has led to renewed interest in solving k -commodity flow in restricted settings, with the authors of [68] making progress on dense graphs.

Linear programming solvers. The quest for understanding the computational complexity of linear programming has a long and rich history in computer science and mathematics. Since the seminal works of Khachiyan [43] and later Karmarkar [39], who were the first to prove that LPs can be solved in polynomial time, the interior point method and its subsequent variants have become the central methods for efficiently solving linear programs with provable guarantees. This has led to a series of refined and more efficient IPM-based solvers [58, 64, 54, 46, 47, 9, 37], which culminated in the recent breakthrough work of Cohen, Lee, and Song [9] who showed that an LP solver whose running time essentially matches the matrix multiplication cost, up to small low-order terms. In a follow-up work, Brand [65] managed to derandomize their algorithm while retaining the same time complexity.

¹For $O(n^\alpha)$ -separable graph, where $\alpha < 1$, it is known that $m = O(n)$, see e.g., [49].

²We can view the hypergraph $G_{\mathbf{A}}$ as a graph, where we interpret each hyper-edge as a clique, and consider its treewidth as usual.

A problem closely related to this paper is solving LPs when the support of the constraint matrix has bounded treewidth τ . Dong, Lee and Ye [16] showed that such structured LPs can be solved in $\tilde{O}(m\tau^2)$, which is near-linear when τ is poly-logarithmic in the parameters of the input.

High-accuracy and approximate multi-commodity flow. As mentioned above, it is known that we can solve multicommodity flow in the high-accuracy regime using linear programming. For a graph with n nodes, m edges, and k commodities, the underlying constraint matrix has km variables and $kn + m$ equality constraints. Thus, using the best-known algorithms for solving linear programs [9, 65], one can achieve a runtime time complexity of $\tilde{O}((km)^\omega)$ for solving multi-commodity flow. In the special case of dense graphs, Brand and Zhang [68] recently showed an improved algorithm achieving $\tilde{O}(k^{2.5}\sqrt{mn}^{\omega-1/2})$ runtime.

In the approximate regime, Leighton et al. [48] show that $(1 + \epsilon)$ multi-commodity flow on *undirected graphs* can be solved in $\tilde{O}(kmn)$, albeit with a rather poor dependency on ϵ . This result led to several follow-up improvements in the low-accuracy regime [24, 21, 51]. Later on, breakthrough works in approximating single commodity max flow in nearly-linear time were also extended to the k -commodity flow problem on undirected graphs [42, 60, 56], culminating in the work of Sherman [61] who achieved an $\tilde{O}(mk\epsilon^{-1})$ time algorithm for the problem.

Multi-commodity flow on planar graphs. The multi-commodity flow problem on planar graphs was studied in the 1980s, but there has not been much interest in it until most recently. Results in the past focused on finding conditions under which solutions existed [55, 30], or finding simple algorithms in even more restricted settings, with the authors of [53] demonstrating that the problem could be solved in $O(kn + n^2(\log n)^{1/2})$ time if the sources and sinks were all on the outer face of the graph. More recently, [41] studied the all-or-nothing version of planar multi-commodity flow, where flows have to be integral, and demonstrate that an $O(1)$ -approximation could be achieved in polynomial time.

Max flow and min-cost flow on general graphs. In what follows, we will focus on surveying only *exact* algorithms for max-flow and min-cost flow on general graphs. For earlier developments on these problems, including fast approximation algorithms, we refer the reader to the following works [44, 1, 8, 60, 42, 56, 62, 4], and the references therein.

An important view, unifying almost all recent max-flow or min-cost flow developments, is interpreting max-flow as the problem of finding one unit of s - t flow that minimizes the ℓ_∞ congestion of the flow vector. Motivated by the near-linear Laplacian solver of Spielman and Teng [63] (which in turn can be used to solve the problem of finding one unit of s - t flow that minimizes the ℓ_2 congestion), and the fact that the gap between ℓ_∞ and ℓ_2 is roughly $O(\sqrt{m})$, Daich and Spielman [11] showed how to implement the IPM for solving min-cost flows in $\tilde{O}(m^{3/2})$ time.

Follow-up works initially made progress on the case of unit capacitated graphs, with the work of Madry [52] achieving an $\tilde{O}(m^{10/7})$ time algorithm for max flow and thus being the first to break the $3/2$ -exponent barrier in the runtime. The running time was later improved to $O(m^{4/3+o(1)})$ and it was generalized to the min-cost flow problem [3, 40].

For general, polynomially bounded capacities, Brand et al. [67] gave an improved algorithm for dense graphs that runs in $\tilde{O}(m + n^{3/2})$. In the sparse graph regime, Gao, Liu and Peng [23] were the first to break the $3/2$ -exponent barrier by giving an $\tilde{O}(m^{3/2-1/128})$ time algorithm, which was later improved to $\tilde{O}(m^{3/2-1/58})$ [66]. Very recently, the breakthrough work of Chen et al. [7] shows that the min-cost flow problem can be solved in $\tilde{O}(m^{1+o(1)})$, which is optimal up to the subpolynomial term.

Max flow and min-cost flow on planar graphs. The study of flows on planar graphs dates back to the celebrated work of Ford and Fulkerson [22] who showed that for the case of s, t -planar graphs³, there is an $O(n^2)$ time algorithm for max flow. This was subsequently improved to $O(n \log n)$ by Itai and Shiloach [35] and finally to $O(n)$ by Henzinger et al [32], the latter building upon a prior work of Hassin [29].

For general planar graphs, there have been two lines of work focusing on the undirected and the directed version of the problem respectively. In the first setting, Reif [57] (and later Hassin and Johnson [31]) gave an $O(n \log^2 n)$ time algorithm. The state-of-the-art algorithm is due to Italiano et al. [36] and achieves $O(n \log \log n)$ runtime. Weihe [69] gave the first speed-up for directed planar max flow running in $O(n \log n)$ time. However, his algorithm required some assumptions on the connectivity of the input graph. Later on, Borradaile and Klein [5] gave an $O(n \log n)$ algorithm for general planar directed graphs. Generalization of planar graphs, e.g., graphs of

³planar graphs where s and t lie on the same face

bounded genus have also been studied in the context of the max flow problem. The work of Chambers et al. [6] showed that these graphs also admit near-linear time max flow algorithms.

Imai and Iwano [33] obtained an $O(n^{1.594} \log M)$ min-cost flow algorithm for graphs that are $O(\sqrt{n})$ recursively separable. For the min-cost flow problem on planar graphs with unit capacities, Karczmarz and Sankowski [38] gave an $O(n^{4/3})$ algorithm. Very recently, Dong et al. [15] showed that the min-cost flow on planar directed graphs with polynomially bounded capacities admits an $\tilde{O}(n)$ time algorithm, which is optimal up to polylogarithmic factors.

1.2 Technical overview Our algorithm framework builds on the work of Dong-Gao-Goranci-Lee-Peng-Sachdeva-Ye on planar min-cost flow [15]. We solve our linear program using the robust interior point method used in [16, 15], where we maintain feasible primal and dual solutions \mathbf{x} and \mathbf{s} to the linear program that converge to the optimal solution over $\tilde{O}(\sqrt{m})$ -many steps of IPM. At every step, we want to move our solutions in the direction of steepest descent of the objective function. To stay close to the central path and avoid violating the capacity lower and upper bounds, the IPM controls the weights \mathbf{W} on the variables and the step direction \mathbf{v} , in order to limit the magnitude of the update to a variable as it approach its bounds. Both \mathbf{W} and \mathbf{v} are defined to be entry-wise dependent on the current solution \mathbf{x} and \mathbf{s} . To maintain feasibility of the solutions, we apply the weighted projection $\mathbf{P}_w \stackrel{\text{def}}{=} \mathbf{W}^{1/2} \mathbf{A}^\top (\mathbf{A} \mathbf{W} \mathbf{A}^\top)^{-1} \mathbf{A} \mathbf{W}^{1/2}$ matrix to the desired step direction \mathbf{v} , which ensures the resulting \mathbf{x} and \mathbf{s} after a step remain in their respective feasible subspaces. In robust IPMs, we also maintain entry-wise approximations $\bar{\mathbf{x}}, \bar{\mathbf{s}}$ to \mathbf{x} and \mathbf{s} , and use these approximations to compute \mathbf{w}, \mathbf{v} , and \mathbf{P}_w at every step. By limiting the updates to $\bar{\mathbf{x}}, \bar{\mathbf{s}}$, robust IPMs achieve efficient runtimes.

The key challenge in the RIPM framework is to implement each step efficiently, specifically, computing the projection $\mathbf{P}_w \mathbf{v}$, as well as updating $\bar{\mathbf{x}}, \bar{\mathbf{s}}$. Similar to [15], we use a *separator tree* to recursively factor the term $\mathbf{A} \mathbf{W} \mathbf{A}^\top$ in \mathbf{P}_w via nested dissection and recursive Schur complements. However, there are several challenges in applying the framework from [15] to general linear programs: In flow problems, the constraint matrix \mathbf{A} is the vertex-edge incidence matrix of the underlying graph, and therefore $\mathbf{A} \mathbf{W} \mathbf{A}^\top$, as well as all recursive Schur complements along the separator tree, are weighted Laplacian matrices, for which we have efficient nearly-linear-time solvers [63] and sparse approximations to Schur complements [45, 27]. This allows [15] to work with an approximate $\tilde{\mathbf{P}}_w \approx \mathbf{P}_w$ efficiently, with implicit access via a collection of approximate Schur complements that can be viewed as sparse Laplacians.

In the context of general separable linear programs, we do not have fast solvers or sparse approximate Schur complements, so instead, we must maintain the collection of Schur complements and their inverses explicitly, by computing them in a bottom-up fashion using the separator tree. To bound the update time, we show that a rank- k update to $\mathbf{A} \mathbf{W} \mathbf{A}^\top$ induced by changes in \mathbf{W} corresponds to rank- k updates to all the recursive Schur complements.

Our second contribution is the dynamic data structures to maintain the implicit representations of \mathbf{x}, \mathbf{s} , which can be viewed as a significant refinement of those from [15]. We recall the notion of *tree operators* introduced in [15] and define an analogous *inverse tree operator*, and give simplified modular data structures to maintain \mathbf{x}, \mathbf{s} using the tree and inverse tree operator. Specifically, we demonstrate more cleanly the power of nested dissection and the recursive subgraph structure in supporting efficient lazy updates to the IPM solutions.

Our third technical contribution is the definition of a fine-grained separator tree which we call the (a, b, λ) -separator tree. The parameters are defined based on the parameters of separable graphs, but they also capture important characteristics of other classes such as low-treewidth graphs. These trees guarantee that at any node, we are able to separate not only the associated graph region, but also the boundary of the region. We use them to maintain the tree operators from the implicit representations, and a careful analysis of node and boundary sizes allows us to conclude that the maintenance can be performed efficiently.

Finally, we note that this work recovers the treewidth LP result of [16] and [28] with significantly lower technical complexity. Whereas [16] constructs an *elimination tree* to directly compute the Cholesky factorization of $\mathbf{A} \mathbf{W} \mathbf{A}^\top = \mathbf{L} \mathbf{L}^\top$, we use a *separator tree* to recursively factor $\mathbf{A} \mathbf{W} \mathbf{A}^\top$. There is a key difference in the two tree constructions, which we believe this paper is correct in: To construct an elimination tree, [16] finds a balanced vertex separator S of $G_{\mathbf{A}}$, *remove S from $G_{\mathbf{A}}$ yielding two disconnected subgraphs H_1, H_2* , recursively construct the elimination tree for H_1 and H_2 , and attach them as children to a vertical path of length $|S|$ corresponding to the vertices of S . When the treewidth of $\mathbf{A} \in \mathbb{R}^{n \times m}$ is t , this process results in an elimination tree of height $\tilde{O}(t)$ where each node corresponds to a vertex of $G_{\mathbf{A}}$, which can then be used to identify explicit coordinates in

the Cholesky factor \mathbf{L} to update when \mathbf{W} changes. Next, an extremely involved transformation using *heavy-light decomposition* is needed to turn the elimination tree into a *sampling tree* of height $O(\log n)$, in order to facilitate the sampling of entries from some implicit vector of the form $\mathbf{A}^\top \mathbf{L}^{-\top} \mathbf{z}$ (required for maintaining $\bar{\mathbf{x}}, \bar{\mathbf{s}}$). In contrast, to construct our separator tree, we find a balanced vertex separator S of $G_{\mathbf{A}}$, but *include S in both subgraphs that are recursed on*, and partition the hyperedges in $\mathcal{E}(S)$ arbitrarily between the two subgraphs. The resulting separator tree has height $O(\log n)$, where each node corresponds to a subgraph of $G_{\mathbf{A}}$, and each level of the tree gives a partition of the columns of \mathbf{A} . This recursive partitioning gives a cleaner, recursive rather than brute-force factorization of $\mathbf{A}\mathbf{W}\mathbf{A}^\top$, and leads to a significant difference in the data structures. When \mathbf{W} changes at a step, [16] updates the Cholesky factorization by processing one changed coordinate at a time ([28] processes one block at a time), so that the data structure update time is linear in the number of new coordinates. On the other hand, our separator tree allows us to update \mathbf{W} in one pass through the tree and yields a sublinear dependence on the number of new coordinates. Moreover, as each node in our separator tree naturally corresponds to a subset of columns of \mathbf{A} , we can use it in a much more straight-forward manner to sample coordinates of $\mathbf{A}^\top \mathbf{L}^{-\top} \mathbf{z}$.

2 Preliminaries

General Notations. We assume all matrices and vectors in an expression have matching dimensions. That is, we will trivially pad matrices and vectors with zeros when necessary. This abuse of notation is unfortunately unavoidable as we will be considering lots of submatrices and subvectors.

An event holds with high probability if it holds with probability at least $1 - n^c$ for arbitrarily large constant c . The choice of c affects guarantees by constant factors.

We use boldface lowercase variables to denote vectors, and boldface uppercase variables to denote matrices. We use $\|\mathbf{v}\|_2$ to denote the 2-norm of vector \mathbf{v} and $\|\mathbf{v}\|_{\mathbf{M}}$ to denote $\sqrt{\mathbf{v}^\top \mathbf{M} \mathbf{v}}$. We use $\text{nnz}(\mathbf{v})$ to denote the number of non-zero entries in the vector \mathbf{v} , equivalently, it is the zero-norm. For any vector \mathbf{v} and scalar x , we define $\mathbf{v} + x$ to be the vector obtained by adding x to each coordinate of \mathbf{v} and similarly $\mathbf{v} - x$ to be the vector obtained by subtracting x from each coordinate of \mathbf{v} . We use $\mathbf{0}$ for all-zero vectors and matrices where dimensions are determined by context.

For an index set A , we use $\mathbf{1}_A$ for the vector with value 1 on coordinates in A and 0 everywhere else. We use \mathbf{I} for the identity matrix and \mathbf{I}_S for the identity matrix in $\mathbb{R}^{S \times S}$. For any vector $\mathbf{x} \in \mathbb{R}^S$, $\mathbf{x}|_C$ denotes the sub-vector of \mathbf{x} supported on $C \subseteq S$; more specifically, $\mathbf{x}|_C \in \mathbb{R}^S$, where $x_i = 0$ for all $i \notin C$.

For any matrix $\mathbf{M} \in \mathbb{R}^{A \times B}$, we use the convention that $\mathbf{M}_{C,D}$ denotes the sub-matrix of \mathbf{M} supported on $C \times D$ where $C \subseteq A$ and $D \subseteq B$. When \mathbf{M} is not symmetric and only one subscript is specified, as in \mathbf{M}_D , this denotes the sub-matrix of \mathbf{M} supported on $A \times D$. To keep notations simple, \mathbf{M}^{-1} will denote the inverse of \mathbf{M} if it is an invertible matrix and the Moore-Penrose pseudo-inverse otherwise.

For any vector \mathbf{v} , we use the corresponding capitalized letter \mathbf{V} to denote the diagonal matrix with \mathbf{v} on the diagonal.

For two positive semi-definite matrices \mathbf{L}_1 and \mathbf{L}_2 , we write $\mathbf{L}_1 \approx_t \mathbf{L}_2$ if $e^{-t} \mathbf{L}_1 \preceq \mathbf{L}_2 \preceq e^t \mathbf{L}_1$, where $\mathbf{A} \preceq \mathbf{B}$ means $\mathbf{B} - \mathbf{A}$ is positive semi-definite. Similarly we define \geq_t and \leq_t for scalars, that is, $x \leq_t y$ if $e^{-t} x \leq y \leq e^t x$.

When multiplying two matrices of differing sizes, say an $m \times n$ matrix with an $n \times k$ matrix, we decompose both matrices into blocks of size $\min\{m, n, k\}$. We then perform block matrix multiplication, with fast matrix multiplication used for the multiplication of two blocks. For example, multiplying a $m \times n$ matrix with an $n \times n$ matrix, with $m \geq n$, takes $(m/n)(n^\omega)$ time.

Trees. For a tree \mathcal{T} , we write $H \in \mathcal{T}$ to mean H is a node in \mathcal{T} . We write \mathcal{T}_H to mean the complete subtree of \mathcal{T} rooted at H . We say a node A is an ancestor of H and H is a descendant of A if H is in the subtree rooted at A , and $H \neq A$. Given a set of nodes \mathcal{H} , we use $\mathcal{P}_{\mathcal{T}}(\mathcal{H})$ to denote the set of all nodes in \mathcal{T} that are ancestors of some node in \mathcal{H} unioned with \mathcal{H} .

The *level* of a node in a tree has the following properties: the root is at level 0; the maximum level is one less than the height of the tree; and the level of a node must be at least one greater than the level of its parent, but this difference does not have to be equal to one. We may assign levels to nodes arbitrarily as long as the above is satisfied. We use $\mathcal{T}(i)$ to denote the collection of all nodes at level i in tree \mathcal{T} .

IPM data structures. When we discuss data structures in the context of the IPM, step 0 means the initialization step. For $k > 0$, step k means the k -th iteration of the while-loop in SOLVE (Algorithm 2); that is, it is the k -th time we update the current solutions. For any vector or matrix \mathbf{x} used in the IPM, we use $\mathbf{x}^{(k)}$ to denote the value of \mathbf{x} at the end of the k -th step.

In all procedures in these data structures, we assume inputs are given by the set of changed coordinates and their values, compared to the previous input. Similarly, we output a vector by the set of changed coordinates and their values, compared to the previous output. This can be implemented by checking memory for changes.

3 Overview of RIPM framework

In this section, we set up the general framework for solving a linear program using a robust IPM. We show that if the projection matrix from the IPM can be maintained efficiently based on the structure of its sparsity pattern, then the overall IPM can be implemented efficiently.

3.1 Robust interior point method

THEOREM 3.1. (RIPM) *Consider the linear program*

$$\min_{\mathbf{Ax}=\mathbf{b}, \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}} \mathbf{c}^\top \mathbf{x}$$

with $\mathbf{A} \in \mathbb{R}^{n \times m}$. We are given a scalar $r > 0$ such that there exists some interior point \mathbf{x}_o satisfying $\mathbf{Ax}_o = \mathbf{b}$ and $\mathbf{l} + r \leq \mathbf{x}_o \leq \mathbf{u} - r$. Let $L = \|\mathbf{c}\|_2$ and $R = \|\mathbf{u} - \mathbf{l}\|_2$. For any $0 < \varepsilon \leq 1/2$, the algorithm RIPM (Algorithm 2) finds \mathbf{x} such that $\mathbf{Ax} = \mathbf{b}$, $\mathbf{l} \leq \mathbf{x} \leq \mathbf{u}$ and

$$\mathbf{c}^\top \mathbf{x} \leq \min_{\mathbf{Ax}=\mathbf{b}, \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}} \mathbf{c}^\top \mathbf{x} + \varepsilon LR.$$

Furthermore, the algorithm has the following properties:

- Each call of SOLVE involves $O(\sqrt{m} \log m \log(\frac{mR}{\varepsilon r}))$ -many steps, and \bar{t} is only updated $O(\log m \log(\frac{mR}{\varepsilon r}))$ -many times.
- In each step of SOLVE, the coordinate i in \mathbf{w}, \mathbf{v} changes only if \bar{x}_i or \bar{s}_i changes.
- In each step of SOLVE, $h\|\mathbf{v}\|_2 = O(\frac{1}{\log m})$.
- Algorithm 2 to Algorithm 2 takes $O(K)$ time in total, where K is the total number of coordinate changes in $\bar{\mathbf{x}}, \bar{\mathbf{s}}$.

We note that this algorithm only requires access to $(\bar{\mathbf{x}}, \bar{\mathbf{s}})$, but not (\mathbf{x}, \mathbf{s}) during the main while loop. Hence, (\mathbf{x}, \mathbf{s}) can be implicitly maintained via any data structure. We only require (\mathbf{x}, \mathbf{s}) explicitly when returning the approximately optimal solution at the end of the algorithm Algorithm 2.

3.2 Projection operators At step k of SOLVE with step direction $\mathbf{v}^{(k)}$ and weights \mathbf{w} (we drop its superscript $^{(k)}$ for convenience), recall we define the projection matrix

$$\mathbf{P}_w \stackrel{\text{def}}{=} \mathbf{W}^{1/2} \mathbf{A}^\top (\mathbf{A} \mathbf{W} \mathbf{A}^\top)^{-1} \mathbf{A} \mathbf{W}^{1/2}.$$

We want to make the primal and dual updates

$$\begin{aligned} \mathbf{x} &\leftarrow \mathbf{x} + h^{(k)} \mathbf{W}^{1/2} \mathbf{v}^{(k)} - h^{(k)} \mathbf{W}^{1/2} \mathbf{P}_w \mathbf{v}^{(k)}, \\ \mathbf{s} &\leftarrow \mathbf{s} + \bar{t} h^{(k)} \mathbf{W}^{-1/2} \mathbf{P}_w \mathbf{v}^{(k)}. \end{aligned}$$

The first term for the primal update is straightforward to maintain, so we may ignore it without loss of generality. After this reduction, we see that the primal and dual updates are analogous. In the remainder of this section, we show how to maintain \mathbf{x} undergoing the update

$$\mathbf{x} \leftarrow \mathbf{x} + h^{(k)} \mathbf{W}^{1/2} \mathbf{P}_w \mathbf{v}^{(k)}.$$

First, observe that $\mathbf{W}^{1/2} \mathbf{P}_w$ is an operator dependent on the dynamic weights \mathbf{w} , which motivates us to formalize this problem setting:

DEFINITION 3.1. (DYNAMIC LINEAR OPERATOR, UPDATE COMPLEXITY) Let \mathbf{w} be a dynamic vector. We say \mathbf{M} is a dynamic linear operator dependent on \mathbf{w} if \mathbf{M} is a function of \mathbf{w} . Let $\mathbf{w}^{(k)}$ be the value of \mathbf{w} at step k , then we use $\mathbf{M}^{(k)}$ to denote the corresponding value of \mathbf{M} at step k .

Suppose exists a data structure that dynamically maintains \mathbf{M} and \mathbf{w} , such that at every step k , if $\mathbf{w}^{(k-1)}$ and $\mathbf{w}^{(k)}$ differ on K coordinates, then the data structure can update $\mathbf{M}^{(k-1)}$ to $\mathbf{M}^{(k)}$ in $f(K)$ time. Then we say \mathbf{M} has update complexity f .

Next, we define two types of dynamic operators dependent on the weights \mathbf{w} from the IPM: the *inverse tree operator* ∇ and the *tree operator* Δ . For linear programs with separable structures, they should crucially combine so that throughout algorithm, we have

$$(3.1) \quad \mathbf{W}^{1/2} \mathbf{P}_{\mathbf{w}} = \Delta \nabla.$$

3.2.1 Operators on a tree In this section, we fix \mathcal{T} to be a *constant-degree* rooted tree with root node G , called the *operator tree*. Let each node $H \in \mathcal{T}$ be associated with a set F_H , where the F_H 's are pairwise disjoint. Let each *leaf* node $L \in \mathcal{T}$ be further associated with a *non-empty* set $E(L)$, where the $E(L)$'s are pairwise disjoint. For a non-leaf node H , define $E(H) \stackrel{\text{def}}{=} \bigcup_{\text{leaf } L \in \mathcal{T}_H} E(L)$. Finally, define $E \stackrel{\text{def}}{=} E(G) = \bigcup_{\text{leaf } L \in \mathcal{T}} E(L)$ and $V \stackrel{\text{def}}{=} \bigcup_{H \in \mathcal{T}} F_H$.

We define two special classes of linear operators that build on the structure of \mathcal{T} . The advantage of these operators lie in their decomposability, which allows them to be efficiently maintained.

DEFINITION 3.2. (INVERSE TREE OPERATOR) Let \mathcal{T} be an operator tree with the associated sets as above. We say a linear operator $\nabla : \mathbb{R}^E \mapsto \mathbb{R}^V$ is an inverse tree operator supported on \mathcal{T} if there exists a linear edge operator ∇_H for each non-root node H in \mathcal{T} , corresponding to the edge from H to its parent, such that ∇ can be decomposed as

$$\nabla = \sum_{\text{leaf } L, \text{ node } H : L \in \mathcal{T}_H} \mathbf{I}_{F_H} \nabla_{H \leftarrow L},$$

where $\nabla_{H \leftarrow L}$ is defined as follows: If $L = H$, then $\nabla_{H \leftarrow L} \stackrel{\text{def}}{=} \mathbf{I}$; otherwise, suppose the path in \mathcal{T} from leaf L to node H is given by $(H_t \stackrel{\text{def}}{=} L, H_{t-1}, \dots, H_1, H_0 \stackrel{\text{def}}{=} H)$, then

$$\nabla_{H \leftarrow L} \stackrel{\text{def}}{=} \nabla_{H_1} \cdots \nabla_{H_{t-1}} \nabla_{H_t}.$$

To maintain ∇ , it will suffice to maintain ∇_H at each non-root node H in \mathcal{T} .

Intuitively, when applying an inverse tree operator to a vector $\mathbf{v} \in \mathbb{R}^E$, \mathbf{v} is partitioned according to the leaves of \mathcal{T} , and then the edge operators are applied sequentially along the tree edges in a bottom-up fashion. It is natural to then also define the opposite process, where edge operators are applied along the tree edges in a top-down fashion.

DEFINITION 3.3. (TREE OPERATOR) Let \mathcal{T} be an operator tree with the associated sets as above. We say a linear operator $\Delta : \mathbb{R}^V \mapsto \mathbb{R}^E$ is tree operator supported on \mathcal{T} if there exists a linear edge operator Δ_H for each non-root node H in \mathcal{T} , corresponding to the edge from H to its parent, such that Δ can be decomposed as

$$\Delta \stackrel{\text{def}}{=} \sum_{\text{leaf } L, \text{ node } H : L \in \mathcal{T}_H} \Delta_{L \leftarrow H} \mathbf{I}_{F_H}.$$

where $\Delta_{L \leftarrow H}$ is defined as follows: If $L = H$, then $\Delta_{L \leftarrow H} \stackrel{\text{def}}{=} \mathbf{I}$. Otherwise, suppose the path in \mathcal{T} from node H to leaf L is given by $(H_t \stackrel{\text{def}}{=} L, H_{t-1}, \dots, H_0 \stackrel{\text{def}}{=} H)$, then

$$\Delta_{L \leftarrow H} \stackrel{\text{def}}{=} \Delta_{H_t} \cdots \Delta_{H_2} \Delta_{H_1}.$$

We define the complexity of a tree (and inverse tree) operator to be parameterized by the number of edge operators applied.

DEFINITION 3.4. (QUERY COMPLEXITY) Let $\Delta \stackrel{\text{def}}{=} \{\Delta_H : H \in \mathcal{T}\}$ be a tree (or inverse tree) operator on tree \mathcal{T} . Suppose for any set \mathcal{H} of K distinct non-root nodes in \mathcal{T} , and any two families of K vectors indexed by \mathcal{H} , $\{\mathbf{u}_H : H \in \mathcal{H}\}$ and $\{\mathbf{v}_H : H \in \mathcal{H}\}$, the total time to compute $\{\mathbf{u}_H^\top \Delta_H : H \in \mathcal{H}\}$ and $\{\Delta_H \mathbf{v}_H : H \in \mathcal{H}\}$ is bounded by $f(K)$. Then we say Δ has query complexity f for some function f .

Without loss of generality, we may assume $f(0) = 0$, $f(k) \geq k$, and f is concave.

By examining the definition of the inverse tree and tree operator, we see they are related.

LEMMA 3.1. If Δ is a tree operator on \mathcal{T} , then Δ^\top is an inverse tree operator on \mathcal{T} , where its edge operators are obtained from Δ 's edge operators by taking a transpose. Furthermore, Δ and ∇ have the same query and update complexity.

3.3 Implicit representations of the solution Assuming we have dynamic inverse tree and tree operators ∇ and Δ on tree \mathcal{T} dependent on \mathbf{w} such that $\mathbf{W}^{1/2} \mathbf{P}_\mathbf{w} = \Delta \nabla$, we can now state how to abstractly maintain the implicit representation of the solutions throughout SOLVE (Algorithm 2). Specifically, we want to maintain the solution \mathbf{x} , and at every step k , carry out an update of the form

$$(3.2) \quad \mathbf{x} \leftarrow \mathbf{x} + h^{(k)} \mathbf{W}^{1/2} \mathbf{P}_\mathbf{w} \mathbf{v}^{(k)}.$$

We design a data structure MAINTAINREP to accomplish this, by:

- At the start of SOLVE, initializing the data structure using the procedure INITIALIZE with $\mathbf{x} = \mathbf{x}^{(\text{init})}$,
- At each step k , updating the weights \mathbf{w} in the data structure using the procedure REWEIGHT, followed by updating \mathbf{x} according to Eq. (3.2) using the procedure MOVE,
- At the end of SOLVE, outputting the final \mathbf{x} using the procedure EXACT.

The key to designing an efficient data structure is to make use of the structure of the operators. Due to their decomposition along \mathcal{T} , we can update the operators and apply them to vectors without exploring all of \mathcal{T} every time.

THEOREM 3.2. (IMPLICIT REPRESENTATION MAINTENANCE) Let \mathbf{w} be the weights changing at every step of SOLVE (Algorithm 2). Suppose there exists dynamic inverse tree and tree operators ∇ and Δ on tree \mathcal{T} both dependent on \mathbf{w} such that $\mathbf{W}^{1/2} \mathbf{P}_\mathbf{w} = \Delta \nabla$ throughout the IPM. Let Q be the max of the query complexity of the tree and inverse tree operator, and let U be the max of the update complexity of the two operators. Suppose \mathcal{T} has constant degree and height η . Then there is a data structure MAINTAINREP that satisfies the following invariants at the end of step k :

- It explicitly maintains the dynamic weights \mathbf{w} and step direction \mathbf{v} from the current step.
- It explicitly maintains scalar c and vectors $\mathbf{z}^{(\text{step})}, \mathbf{z}^{(\text{sum})}$, which together represent the implicitly-maintained vector $\mathbf{z} \stackrel{\text{def}}{=} c\mathbf{z}^{(\text{step})} + \mathbf{z}^{(\text{sum})}$. At the end of step k , $\mathbf{z}^{(\text{step})} = \nabla^{(k)} \mathbf{v}^{(k)}$, and

$$\mathbf{z} = \sum_{i=1}^k h^{(i)} \nabla^{(i)} \mathbf{v}^{(i)}.$$

- It implicitly maintains \mathbf{x} so that at the end of step k ,

$$\mathbf{x} = \mathbf{x}^{(\text{init})} + \sum_{i=1}^k h^{(i)} \Delta^{(i)} \nabla^{(i)} \mathbf{v}^{(i)},$$

where $\mathbf{x}^{(\text{init})}$ is some initial value set at the start of SOLVE.

The data structure supports the following procedures and runtimes:

- INITIALIZE($\Delta, \nabla, \mathbf{v}^{(\text{init})} \in \mathbb{R}^m, \mathbf{w}^{(\text{init})} \in \mathbb{R}_{>0}^m, \mathbf{x}^{(\text{init})} \in \mathbb{R}^m$): Preprocess and set $\mathbf{x} \leftarrow \mathbf{x}^{(\text{init})}$. The procedure runs in $O(U(m) + Q(m))$ time.
- REWEIGHT($\delta_{\mathbf{w}} \in \mathbb{R}_{>0}^m$): Update the weights to $\mathbf{w} \leftarrow \mathbf{w} + \delta_{\mathbf{w}}$. The procedure runs in $O(U(\eta K) + Q(\eta K))$ total time, where $K = \text{nnz}(\delta_{\mathbf{w}})$.
- MOVE($h \in \mathbb{R}, \delta_{\mathbf{v}} \in \mathbb{R}^m$): Update the current step direction to $\mathbf{v} \leftarrow \mathbf{v} + \delta_{\mathbf{v}}$. Update the implicit representation of \mathbf{x} to reflect the following change in value:

$$\mathbf{x} \leftarrow \mathbf{x} + h\Delta\nabla\mathbf{v}.$$

The procedure runs in $O(Q(\eta K))$ time, where $K = \text{nnz}(\delta_{\mathbf{v}})$.

- EXACT: Output the current exact value of \mathbf{x} in $O(Q(m))$ time.

3.4 Solution approximation In the IPM, one key operation is to maintain the solution vector $\bar{\mathbf{x}}$ that is close to \mathbf{x} throughout the algorithm. (Analogously for the slack $\bar{\mathbf{s}}$ close to \mathbf{s} .) Since we have implicit representations of the solution \mathbf{x} from MAINTAINREP, we now show how to maintain $\bar{\mathbf{x}}$ close to \mathbf{x} . To accomplish this, we use a meta data structure that solves this in a more general setting introduced in [15].

THEOREM 3.3. (APPROXIMATE VECTOR MAINTENANCE WITH TREE OPERATOR [15]) *Let $0 < \rho < 1$ be a failure probability. Suppose $\Delta \in \mathbb{R}^{m \times n}$ is a tree operator with query complexity Q and supported on a constant-degree tree \mathcal{T} with height η . There is a randomized data structure MAINTAINAPPROX that takes as input the dynamic weights \mathbf{w} and the dynamic \mathbf{x} implicitly maintained according to Theorem 3.2 at every step, and explicitly maintains the approximation $\bar{\mathbf{x}}$ to \mathbf{x} satisfying $\|\mathbf{W}^{-1/2}(\mathbf{x} - \bar{\mathbf{x}})\|_{\infty} \leq \delta$ at every step with probability $1 - \rho$.*

Suppose $\|\mathbf{W}^{(k)-1/2}(\mathbf{x}^{(k)} - \mathbf{x}^{(k-1)})\|_2 \leq \beta$ for all steps k . Furthermore, suppose \mathbf{w} is a function of $\bar{\mathbf{x}}$ coordinate-wise. Then, for each $\ell \geq 0$, $\bar{\mathbf{x}}$ admits $2^{2\ell}$ coordinate changes every 2^ℓ steps. Over N total steps, the total cost of the data structure is

$$(3.3) \quad \tilde{O}(\eta^3(\beta/\delta)^2 \log^3(mN/\rho)) \left(Q(m) + \sum_{k=1}^N Q(S^{(k)}) + \sum_{\ell=0}^{\log N} \frac{N}{2^\ell} \cdot Q(2^{2\ell}) \right),$$

where $S^{(k)}$ is the number of nodes H where Δ_H or \mathbf{u}_H in the implicit representation of \mathbf{x} changed at step k .

3.5 Main theorem for the RIPM framework We are now ready to state and prove the main result in this framework.

THEOREM 3.4. (RIPM FRAMEWORK) *Consider an LP of the form*

$$(3.4) \quad \min_{\mathbf{x} \in \mathcal{P}} \mathbf{c}^\top \mathbf{x} \quad \text{where} \quad \mathcal{P} = \{\mathbf{A}\mathbf{x} = \mathbf{b}, \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}\}$$

where $\mathbf{A} \in \mathbb{R}^{n \times m}$. For any vector \mathbf{w} , let $\mathbf{P}_{\mathbf{w}} \stackrel{\text{def}}{=} \mathbf{W}^{1/2} \mathbf{A}^\top (\mathbf{A} \mathbf{W} \mathbf{A}^\top)^{-1} \mathbf{A} \mathbf{W}^{1/2}$, and suppose there exists dynamic tree and inverse tree operators Δ and ∇ dependent on \mathbf{w} , such that $\mathbf{W}^{1/2} \mathbf{P}_{\mathbf{w}} = \Delta \nabla$. Let U be the update complexity of Δ and ∇ , and let Q be their query complexity. Let r and $R = \|\mathbf{u} - \mathbf{l}\|_2$ be the inner and outer radius of \mathcal{P} , and let $L = \|\mathbf{c}\|_2$. Then, there is a data structure to solve Eq. (3.4) to εLR accuracy with probability $1 - 2^{-m}$ in time

$$\tilde{O} \left(\eta^4 \sqrt{m} \log\left(\frac{R}{\varepsilon r}\right) \cdot \sum_{\ell=0}^{\frac{1}{2} \log m} \frac{U(2^{2\ell}) + Q(2^{2\ell})}{2^\ell} \right).$$

Proof. [Proof of Theorem 3.4] We implement the IPM algorithm using the data structures from Sections 3.3 and 3.4, and bound the cost of each operations of the data structures. For simplicity, we only discuss the primal variables in this proof, but the slack variables are analogous. We use one copy of MAINTAINREP to maintain \mathbf{x} , and one copy of MAINTAINAPPROX to maintain $\bar{\mathbf{x}}$. At each step, we perform the implicit update of \mathbf{x} using

MOVE and update \mathbf{w} using REWEIGHT in MAINTAINREP. We construct the explicit approximations $\bar{\mathbf{x}}$ using APPROXIMATE in MAINTAINAPPROX.

Theorem 3.3 shows that throughout the IPM, for each $\ell \geq 0$, there are $2^{2\ell}$ coordinate changes to $\bar{\mathbf{x}}$ every 2^ℓ steps. Since \mathbf{w} is a function of $\bar{\mathbf{x}}$ coordinate-wise, there are also $2^{2\ell}$ coordinate changes in \mathbf{w} every 2^ℓ steps. Similarly, we observe that \mathbf{v} is defined as a function of $\bar{\mathbf{x}}$ and $\bar{\mathbf{s}}$ coordinate-wise, so there are $O(2^{2\ell})$ coordinate changes to \mathbf{v} every 2^ℓ steps. Then Theorem 3.2 shows that the total runtime over N steps for the MAINTAINREP data structure is

$$(3.5) \quad \tilde{O}(U(m) + Q(m)) + \tilde{O} \left(\sum_{\ell=0}^{\log N} \frac{N}{2^\ell} \cdot (U(\eta \cdot 2^{2\ell}) + Q(\eta \cdot 2^{2\ell})) \right).$$

Theorem 3.3 shows that the total runtime over N steps for MAINTAINAPPROX is

$$(3.6) \quad \tilde{O}(\eta^3(\beta/\delta)^2 \log^3(mN/\rho)) \left(Q(m) + \sum_{k=1}^N Q(S^{(k)}) + \sum_{\ell=0}^{\log N} \frac{N}{2^\ell} \cdot Q(2^{2\ell}) \right),$$

where the variables are defined as in the theorem statement. By examining Theorem 3.2, we see that when a coordinate of \mathbf{w} or \mathbf{v} changes, the implicit representation of \mathbf{x} admits updates at $O(\eta)$ -many nodes. Combined with the concavity of Q , we can bound

$$\sum_{k=1}^N Q(S^{(k)}) \leq O(\eta) \cdot \sum_{\ell=0}^{\log N} \frac{N}{2^\ell} \cdot Q(2^{2\ell}).$$

Theorem 3.1 guarantees that there are $N = \sqrt{m} \log m \log(\frac{mR}{\epsilon r})$ total IPM steps, and at each step k , we have $h^{(k)} \left\| \mathbf{W}^{(k)-1/2}(\mathbf{x}^{(k)} - \mathbf{x}^{(k-1)}) \right\|_2 = h^{(k)} \left\| \mathbf{v}^{(k)} - \mathbf{P}_w \mathbf{v}^{(k)} \right\|_2 \leq O(\frac{1}{\log m})$, so we can set $\beta = O(\frac{1}{\log m})$. By examining Algorithm 2, we see it suffices to set $\delta = O(\frac{1}{\log m})$. We choose the failure probability ρ to be appropriately small, e.g. 2^{-m} . Finally, we conclude that the overall runtime of the IPM framework is

$$\tilde{O} \left(\eta^4 \sqrt{m} \log\left(\frac{R}{\epsilon r}\right) \cdot \sum_{\ell=0}^{\frac{1}{2} \log m} \frac{U(2^{2\ell}) + Q(2^{2\ell})}{2^\ell} \right),$$

where the terms for initialization times have been absorbed.

□

4 From separator tree to projection operators

In this section, we explore the separable structure of the dual graph $G_{\mathbf{A}}$ of the LP constraint matrix \mathbf{A} , and use these properties to help define and maintain the tree operator and inverse tree operator as needed for the IPM framework from Section 3.

Throughout this section, we fix $\mathbf{A} \in \mathbb{R}^{n \times m}$, so that the dual graph $G_{\mathbf{A}} = (V, E)$ has n vertices, m hyperedges. Additionally, let ρ denote the max hyperedge size in $G_{\mathbf{A}}$; equivalently, ρ is the column sparsity of \mathbf{A} .

4.1 Separator tree The notion of using a *separator tree* to represent the recursive decomposition of a separable graph is well-established in literature, c.f [19, 32]. In our work, we use the following definition:

DEFINITION 4.1. (SEPARATOR TREE) *Let G be a hypergraph with n vertices, m hyperedges, and max hyperedge size ρ . A separator tree \mathcal{S} for G is a constant-degree tree whose nodes represent a recursive decomposition of G based on balanced separators.*

Formally, each node of \mathcal{S} is a region (edge-induced subgraph) H of G ; we denote this by $H \in \mathcal{S}$. At a node H , we define subsets of vertices $\partial H, S(H), F_H$, where ∂H is the set of boundary vertices of H , i.e. vertices with neighbours outside H in G ; $S(H)$ is a balanced vertex separator of H ; and F_H is the set of eliminated vertices at H . Furthermore, let $E(H)$ denote the edges contained in H .

The nodes and associated vertex sets are defined in a top-down manner as follows:




Figure 4.1: An example of a separator tree. The bold edges denote the boundary of each component, ∂H while the dotted lines denote the separators $S(H)$. Note that $F_H = S(H) \setminus \partial H$ is defined differently on the leaves.

1. The root of \mathcal{S} is the node $H = G$, with $\partial H = \emptyset$ and $F_H = S(H)$.
2. A non-leaf node $H \in \mathcal{S}$ has a constant number of children whose union is H . The children form a edge-disjoint partition of H , and the intersection of their vertex sets is a balanced separator $S(H)$ of H . Define the set of eliminated vertices at H to be $F_H \stackrel{\text{def}}{=} S(H) \setminus \partial H$.

The set $F_H \cup \partial H$ consists of all vertices in the boundary and separator, which can intuitively be interpreted as the skeleton of H . In later sections, we recursively construct graphs (matrices) on $F_H \cup \partial H$ which capture compressed information about all of H .

By definition of boundary vertices, for a child D of H , we have $\partial D \stackrel{\text{def}}{=} (\partial H \cup S(H)) \cap V(D)$.

3. At a leaf node H , we define $S(H) = \emptyset$ and $F_H = V(H) \setminus \partial H$. (This convention allows leaf nodes to exist at different levels in \mathcal{S} .) The leaf nodes of \mathcal{S} partition the edges of G .

We use η to denote the height of \mathcal{S} .

For a separator tree to be meaningful, the leaf node regions should be sufficiently small, to indicate that we have a good overall decomposition of the graph. Additionally, for our work, we want a more careful bound on the sizes of the skeleton of regions. This motivates the following refined definition:

DEFINITION 4.2. ((a, b, λ) -SEPARATOR TREE) Let G be a graph with n vertices, m edges, and max hyperedge size ρ . Let $a \in [0, 1]$ and $b \in (0, 1)$ be constants, and $\lambda \geq 1$ be an expression in terms of m, n, ρ . An (a, b, λ) -separator tree \mathcal{S} for G is a separator tree satisfying the following additional properties:

1. There are at most $O(b^{-i})$ nodes at level i in \mathcal{S} ,
2. any node H at level i satisfies $|F_H \cup \partial H| \leq O(\lambda \cdot b^{ai})$,
3. a node H at level i is a leaf node if and only if $|V(H)| \leq O(\rho)$.

Intuitively, a and b come from the separability parameters of G , and λ is a scaling factor for node sizes in \mathcal{S} . Since there could be hyperedges of size ρ , regions of size ρ are not necessarily separable, so we set the region as a leaf.

We make extensive use of these properties in subsequent sections when computing runtimes.

4.2 Nested dissection using a separator tree Let \mathcal{S} be any separator tree for $G_{\mathbf{A}}$. In this section, we show how to use \mathcal{S} to factor the matrix $\mathbf{L}^{-1} \stackrel{\text{def}}{=} (\mathbf{A}\mathbf{W}\mathbf{A}^\top)^{-1}$ recursively:

DEFINITION 4.3. (BLOCK CHOLESKY DECOMPOSITION) *The block Cholesky decomposition of a symmetric matrix \mathbf{L} with blocks indexed by F and C is:*

$$(4.1) \quad \mathbf{L} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{L}_{C,F}(\mathbf{L}_{F,F})^{-1} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{L}_{F,F} & \mathbf{0} \\ \mathbf{0} & \mathbf{Sc}(\mathbf{L}, C) \end{bmatrix} \begin{bmatrix} \mathbf{I} & (\mathbf{L}_{F,F})^{-1}\mathbf{L}_{F,C} \\ \mathbf{0} & \mathbf{I} \end{bmatrix},$$

where the middle matrix in the decomposition is a block-diagonal matrix with blocks indexed by F and C , with the lower-right block being the Schur complement $\mathbf{Sc}(\mathbf{L}, C)$ of \mathbf{L} onto C :

$$(4.2) \quad \mathbf{Sc}(\mathbf{L}, C) \stackrel{\text{def}}{=} \mathbf{L}_{C,C} - \mathbf{L}_{C,F}\mathbf{L}_{F,F}^{-1}\mathbf{L}_{F,C}.$$

Since $\mathbf{Sc}(\mathbf{L}, C)$ is a symmetric matrix, we can recursively apply the decomposition Eq. (4.1) to it. By choosing the index sets F, C for each recursive step according to \mathcal{S} , we get a recursive decomposition of \mathbf{L}^{-1} :

THEOREM 4.1. (\mathbf{L}^{-1} FACTORIZATION, C.F. [15] THEOREM 33) *Let \mathcal{S} be the separator tree of $G_{\mathbf{A}}$ with height η . For each node $H \in \mathcal{S}$ with hyperedges $E(H)$, let $\mathbf{A}_H \in \mathbb{R}^{n \times m}$ denote the matrix \mathbf{A} restricted to columns indexed by $E(H)$. Define*

$$(4.3) \quad \mathbf{L}[H] \stackrel{\text{def}}{=} \mathbf{A}_H \mathbf{W} \mathbf{A}_H^\top, \text{ and}$$

$$(4.4) \quad \mathbf{L}^{(H)} \stackrel{\text{def}}{=} \mathbf{Sc}(\mathbf{L}[H], F_H \cup \partial H).$$

Then, we have

$$(4.5) \quad \mathbf{L}^{-1} = \mathbf{\Pi}^{(\eta)\top} \dots \mathbf{\Pi}^{(1)\top} \mathbf{\Gamma} \mathbf{\Pi}^{(1)} \dots \mathbf{\Pi}^{(\eta)},$$

where⁴

$$(4.6) \quad \mathbf{\Gamma} \stackrel{\text{def}}{=} \begin{bmatrix} \left(\sum_{H \in \mathcal{T}(\eta)} \left(\mathbf{L}_{F_H, F_H}^{(H)} \right)^{-1} \right) & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \ddots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \left(\sum_{H \in \mathcal{T}(0)} \left(\mathbf{L}_{F_H, F_H}^{(H)} \right)^{-1} \right) \end{bmatrix},$$

and for $i = 1, \dots, \eta$,

$$(4.7) \quad \mathbf{\Pi}^{(i)} \stackrel{\text{def}}{=} \mathbf{I} - \sum_{H \in \mathcal{T}(i)} \mathbf{X}^{(H)},$$

where $\mathcal{T}(i)$ is the set of nodes at level i in \mathcal{T} , the matrix $\mathbf{\Pi}^{(i)}$ is supported on $\bigcup_{H \in \mathcal{T}(i)} F_H \cup \partial H$ and padded with zeros to n -dimensions, and for each $H \in \mathcal{S}$,

$$(4.8) \quad \mathbf{X}^{(H)} \stackrel{\text{def}}{=} \mathbf{L}_{\partial H, F_H}^{(H)} \left(\mathbf{L}_{F_H, F_H}^{(H)} \right)^{-1}.$$

4.3 Projection operators definition Suppose \mathcal{S} is a separator tree for $G_{\mathbf{A}}$. In this subsection, we define the operator tree \mathcal{T} based on \mathcal{S} , followed by the tree operator $\mathbf{\Delta}$ and inverse tree operator $\mathbf{\nabla}$ which will be supported on \mathcal{T} . Finally, we will show that our definitions indeed satisfy

$$\mathbf{W}^{1/2} \mathbf{P}_w = \mathbf{\Delta} \mathbf{\nabla}.$$

⁴We use a different definition of *level* compared to [15]. In [15], the root has level η in and leaf nodes have level 0, and in this paper, the root has level 0 and leaf nodes have level η . This is purely for notational convenience in later calculations, so this theorem is otherwise unaffected.

Recall that \mathcal{S} is a constant-degree tree. The leaf nodes of \mathcal{S} partition the hyperedges of $G_{\mathbf{A}}$, however, we do not have a bound on the number of hyperedges in a leaf node. In constructing \mathcal{T} , we simply want to modify \mathcal{S} so that each leaf contains exactly one hyperedge. Specifically, for each leaf node $H \in \mathcal{S}$ containing $|E(H)|$ hyperedges, we construct a complete binary tree \mathcal{T}_H^+ rooted at H with $|E(H)|$ leaves, assign one hyperedge from $E(H)$ to one new leaf, and attach \mathcal{T}_H^+ at the node H . This construction yields the desired operator tree \mathcal{T} whose height is within a $\log |E|$ factor of \mathcal{S} .

We define the tree operator Δ on \mathcal{T} follows: For non-root node H in \mathcal{T} , let

$$(4.9) \quad \Delta_H \stackrel{\text{def}}{=} \begin{cases} \mathbf{I}_{F_H \cup \partial H} - \mathbf{X}^{(H)\top} & \text{if } H \text{ exists in } \mathcal{S} \\ \mathbf{W}_{E(H)}^{1/2} \mathbf{A}_H^\top & \text{if } H \text{ is a leaf node in } \mathcal{T} \\ \mathbf{I} & \text{else.} \end{cases}$$

Note that the first two cases are indeed disjoint by construction. We pad zeros to all matrices in order to arrive at the correct overall dimensions.

LEMMA 4.1. (C.F. [15], LEMMA 59) *Let Δ be the tree operator as defined above. Then*

$$(4.10) \quad \Delta = \mathbf{W}^{1/2} \mathbf{A}^\top \mathbf{\Pi}^{(\eta)\top} \dots \mathbf{\Pi}^{(1)\top}.$$

Next, we establish the query complexity of the tree operator:

LEMMA 4.2. *Suppose L is the total number of leaf nodes in \mathcal{S} . The query complexity of Δ is*

$$Q(K) = O \left(\rho K + \max_{\mathcal{H}: \text{set of } K \text{ leaves in } \mathcal{S}} \sum_{H \in \mathcal{P}_{\mathcal{S}}(\mathcal{H})} |F_H \cup \partial H|^2 \right)$$

for $K \leq L$, where $\mathcal{P}_{\mathcal{S}}(\mathcal{H})$ is the set of all nodes in \mathcal{S} that are ancestors of some node in \mathcal{H} unioned with \mathcal{H} . When $K > L$, then we define $Q(K) = Q(L)$.

Proof. First, we consider the query time $Q(1)$ for a single edge. Let \mathbf{u} be any vector, and let H be a non-root node in \mathcal{T} . If H is a leaf node, then computing $\Delta_H \mathbf{u}$ and $\mathbf{u}^\top \Delta_H$ both take $O(\rho)$ time. If H exists in \mathcal{S} , then computing $\Delta_H \mathbf{u}$ takes $O(|F_H|^2 + |\partial H||F_H|) \leq O(|F_H \cup \partial H|^2)$ time, since the bottleneck is naively computing $\mathbf{L}_{\partial H, F_H}^{(H)} \left(\mathbf{L}_{F_H, F_H}^{(H)} \right)^{-1} \mathbf{u}$. Therefore, $Q(1) = O(\rho + \max_{H \in \mathcal{S}} |F_H \cup \partial H|^2)$.

For $K > 1$, we can simply bound the query time for K distinct edges by

$$Q(K) = O \left(\rho K + \max_{\mathcal{H}: \text{set of } K \text{ nodes in } \mathcal{S}} \sum_{H \in \mathcal{H}} |F_H \cup \partial H|^2 \right).$$

Finally, note that we can take the summation over $H \in \mathcal{P}_{\mathcal{S}}(\mathcal{H})$ instead of $H \in \mathcal{H}$ for an upper bound. In this case, it suffices to take the max over sets of leaf nodes. \square

By taking the transpose of Δ , we get an inverse tree operator, and together, they give the projection matrix using Eq. (4.5).

COROLLARY 4.1. *Let $\nabla \stackrel{\text{def}}{=} \Delta^\top$ be the inverse tree operator obtained from Δ by transposing the edge and leaf operators. Then*

$$(4.11) \quad \mathbf{W}^{1/2} \mathbf{P}_w \stackrel{\text{def}}{=} \mathbf{W}^{1/2} \mathbf{W}^{1/2} \mathbf{A}^\top \mathbf{L}^{-1} \mathbf{A} \mathbf{W}^{1/2} = (\mathbf{W}^{1/2} \Delta) \mathbf{\Gamma} \nabla.$$

REMARK 4.1. *Without loss of generality, we have chosen to simplify our presentation and consider $\Delta \nabla$ in place of $\mathbf{W}^{1/2} \Delta \mathbf{\Gamma} \nabla$.*

This is possible for two reasons: One, $\mathbf{W}^{1/2} \Delta$ is a tree operator, which we can in fact maintain in the same time complexity as Δ . Two, $\mathbf{\Gamma}$ is a block-diagonal matrix, with a block for each $H \in \mathcal{T}$ that is indexed by F_H . It is straightforward to show we can maintain and apply $\mathbf{\Gamma} \nabla$ in the same time complexity as ∇ .

4.4 Maintenance of projection operators So far, we have defined the separator tree \mathcal{S} for the graph $G_{\mathbf{A}}$, which we then used to define the operator tree \mathcal{T} , which supports the tree operator Δ needed for the IPM framework. In this subsection, we discuss how to maintain $\mathbf{L}^{(H)}$, $\mathbf{Sc}(\mathbf{L}^{(H)}, \partial H)$, and $(\mathbf{L}_{F_H, F_H}^{(H)})^{-1}$ at each node $H \in \mathcal{S}$ using the data structure DYNAMICSC (Algorithm 1), as the weight vector \mathbf{w} undergoes changes throughout the IPM. This will in turn allow us to maintain the tree operator Δ .

We begin with a lemma showing that given a symmetric matrix and a low-rank update, we can compute its new inverse and Schur complement quickly.

LEMMA 4.3. *Let $\mathbf{L}' = \mathbf{L} + \mathbf{U}\mathbf{V} \in \mathbb{R}^{n \times n}$ be a symmetric matrix plus a rank- K update, where \mathbf{U} and \mathbf{V}^\top both have dimensions $n \times K$. Given $\mathbf{L}', \mathbf{U}, \mathbf{V}$, we can compute \mathbf{L}'^{-1} in $O(n^2 K^{\omega-2})$ time.*

Additionally, suppose we are also given \mathbf{L}^{-1} and $\mathbf{Sc}(\mathbf{L}, S)$ for an index set S . Then we can compute $\mathbf{Sc}(\mathbf{L}', S)$, \mathbf{U}', \mathbf{V}' in $O(n^2 K^{\omega-2})$ time, so that $\mathbf{Sc}(\mathbf{L}, S) + \mathbf{U}'\mathbf{V}' = \mathbf{Sc}(\mathbf{L}', S)$, and $\mathbf{U}', \mathbf{V}'^\top$ both have K columns.

Proof. The Sherman-Morrison formula states

$$\mathbf{L}'^{-1} = \mathbf{L}^{-1} - \mathbf{L}^{-1}\mathbf{U}(\mathbf{I}_K + \mathbf{V}\mathbf{L}^{-1}\mathbf{U})^{-1}\mathbf{V}\mathbf{L}^{-1}.$$

The time to compute this update is dominated by the time required to multiply an $n \times n$ matrix with an $n \times K$ matrix, which is $O(n^2 K^{\omega-2})$.

For the second part of the lemma, recall that the Schur complement is defined to be:

$$(4.12) \quad \mathbf{Sc}(\mathbf{L}, C) \stackrel{\text{def}}{=} \mathbf{L}_{C,C} - \mathbf{L}_{C,F}\mathbf{L}_{F,F}^{-1}\mathbf{L}_{F,C}.$$

If we were to naively use this definition of the Schur complement to perform the updates and construct \mathbf{U}' and \mathbf{V}'^\top , we will run into an issue where the rank of the new update blows up by a factor of 8, leading to an exponential blowup in the rank as we go up the levels recursively. Instead, we make use of the fact that the inverse of the Schur complement, $\mathbf{Sc}(\mathbf{L}, S)^{-1}$ is exactly the S, S submatrix of \mathbf{L}^{-1} to control the rank of the updates.

We first apply the definition of Schur complement and then use the Sherman-Morrison formula to get

$$\begin{aligned} \mathbf{Sc}(\mathbf{L}', S)^{-1} &= \mathbf{L}'^{-1}_{S,S} \\ &= \mathbf{L}^{-1}_{S,S} - (\mathbf{L}^{-1}\mathbf{U}(\mathbf{I}_K + \mathbf{V}\mathbf{L}^{-1}\mathbf{U})^{-1}\mathbf{V}\mathbf{L}^{-1})_{S,S} \\ &= \mathbf{Sc}(\mathbf{L}, S)^{-1} - \mathbf{I}_S\mathbf{L}^{-1}\mathbf{U}(\mathbf{I}_K + \mathbf{V}\mathbf{L}^{-1}\mathbf{U})^{-1}\mathbf{V}\mathbf{L}^{-1}\mathbf{I}_S. \end{aligned}$$

This gives us the new rank- K update $\mathbf{Sc}(\mathbf{L}', S)^{-1} = \mathbf{Sc}(\mathbf{L}, S)^{-1} + \mathbf{U}^*\mathbf{V}^*$ with

$$\begin{aligned} \mathbf{U}^* &= -\mathbf{I}_S\mathbf{L}^{-1}\mathbf{U} \\ \mathbf{V}^* &= (\mathbf{I}_K + \mathbf{V}\mathbf{L}^{-1}\mathbf{U})^{-1}\mathbf{V}\mathbf{L}^{-1}\mathbf{I}_S. \end{aligned}$$

We can now determine the Schur complement update by applying Sherman-Morrison again:

$$\mathbf{Sc}(\mathbf{L}', S) = \mathbf{Sc}(\mathbf{L}, S) - \mathbf{Sc}(\mathbf{L}, S)\mathbf{U}^*(\mathbf{I}_K + \mathbf{V}^*\mathbf{Sc}(\mathbf{L}, S)\mathbf{U}^*)^{-1}\mathbf{V}^*\mathbf{Sc}(\mathbf{L}, S).$$

This is a rank- K update $\mathbf{Sc}(\mathbf{L}', S) = \mathbf{Sc}(\mathbf{L}, S) + \mathbf{U}'\mathbf{V}'$ with

$$\begin{aligned} \mathbf{U}' &= -\mathbf{Sc}(\mathbf{L}, S)\mathbf{U}^* \\ \mathbf{V}' &= (\mathbf{I}_K + \mathbf{V}^*\mathbf{Sc}(\mathbf{L}, S)\mathbf{U}^*)^{-1}\mathbf{V}^*\mathbf{Sc}(\mathbf{L}, S). \end{aligned}$$

The time to compute $\mathbf{U}^*, \mathbf{V}^*, \mathbf{U}', \mathbf{V}'$ are all dominated by the time to multiply an $n \times n$ matrix with an $n \times K$ matrix, which is $O(n^2 K^{\omega-2})$. \square

Now, we are ready to present the data structure for maintaining the Schur complement matrices along a separator tree.

Algorithm 1 Data structure to maintain dynamic Schur complements

```

1: data structure DYNAMICSC
2: private: member
3:   Hypergraph  $G_{\mathbf{A}}$  with incidence matrix  $\mathbf{A}$ 
4:    $\mathbf{w} \in \mathbb{R}^m$ : Dynamic weight vector
5:    $\mathcal{S}$ : Separator tree of height  $\eta$ . Every node  $H$  of  $\mathcal{S}$  stores:
6:      $F_H, \partial H$ : Sets of eliminated vertices and boundary vertices of region  $H$ 
7:      $E(H)$ : Set of hyperedges of region  $H$ 
8:      $\mathbf{L}^{(H)}, (\mathbf{L}_{F_H, F_H}^{(H)})^{-1}, \mathbf{Sc}(\mathbf{L}^{(H)}, \partial H)$ : Matrices to maintain as a function of  $\mathbf{w}$ 
9:      $\mathbf{L}^{(H)^{-1}}$ : Additional inverse matrix to maintain as a function of  $\mathbf{w}$ 
10:     $\mathbf{U}_H, \mathbf{V}_H$ : Low-rank update at  $H$ , used in REWEIGHT
11:
12: procedure INITIALIZE( $\mathcal{S}, \mathbf{w}^{(\text{init})} \in \mathbb{R}^m$ )
13:    $\mathcal{S} \leftarrow \mathcal{S}, \mathbf{w} \leftarrow \mathbf{w}^{(\text{init})}$ 
14:   for level  $i = \eta$  to 0 do
15:     for each node  $H$  at level  $i$  do
16:        $\mathbf{L}^{(H)}, (\mathbf{L}_{F_H, F_H}^{(H)})^{-1}, \mathbf{Sc}(\mathbf{L}^{(H)}, \partial H) \leftarrow \mathbf{0}, \mathbf{0}, \mathbf{0}$ 
17:       SCHURNODE( $H, \mathbf{w}$ )
18:     end for
19:   end for
20: end procedure
21:
22: procedure REWEIGHT( $\delta_{\mathbf{w}} \in \mathbb{R}^m$ )
23:    $\mathcal{H} \leftarrow$  set of nodes  $H$  in  $\mathcal{S}$  where  $\delta_{\mathbf{w}}|_{E(H)} \neq \mathbf{0}$ 
24:   for level  $i = \eta$  to 0 do
25:     for each node  $H \in \mathcal{H}$  at level  $i$  do
26:       SCHURNODE( $H, \delta_{\mathbf{w}}$ )
27:     end for
28:   end for
29:    $\mathbf{w} \leftarrow \mathbf{w} + \delta_{\mathbf{w}}$ 
30: end procedure
31:
32: procedure SCHURNODE( $H \in \mathcal{S}, \delta_{\mathbf{w}} \in \mathbb{R}^m$ )
33:   if  $H$  is a leaf node then ▷ rank of update  $\leq \min\{\text{nnz}(\delta_{\mathbf{w}}|_{E(H)}), |F_H \cup \partial H|\}$ 
34:      $\mathbf{L}^{(H)} \leftarrow \mathbf{L}^{(H)} + \mathbf{A}_H \text{diag}(\delta_{\mathbf{w}}|_{E(H)}) \mathbf{A}_H^{\top}$ 
35:   else if  $\text{nnz}(\delta_{\mathbf{w}}|_{E(H)}) \leq |F_H \cup \partial H|$  then ▷ rank of update  $\leq \sum_{\text{child } D} K_D \leq \text{nnz}(\delta_{\mathbf{w}}|_{E(H)})$ 
36:      $\mathbf{L}^{(H)} \leftarrow \mathbf{L}^{(H)} + \sum_{\text{child } D \text{ of } H} \mathbf{U}_D \mathbf{V}_D$ 
37:   else ▷ rank of update  $\leq |F_H \cup \partial H|$ 
38:      $\mathbf{L}^{(H)} \leftarrow \sum_{\text{child } D \text{ of } H} \mathbf{Sc}(\mathbf{L}^{(D)}, \partial D)$ 
39:   end if
40:   Let  $K_H \stackrel{\text{def}}{=} \min\{\text{nnz}(\delta_{\mathbf{w}}|_{E(H)}), |F_H \cup \partial H|\}$  ▷ upper bound on the rank of update to  $\mathbf{L}^{(H)}$ 
41:   Compute  $(\mathbf{L}_{F_H, F_H}^{(H)})^{-1}$  and  $\mathbf{L}^{(H)^{-1}}$  by Lemma 4.3
42:   Compute  $\mathbf{Sc}(\mathbf{L}^{(H)}, \partial H)$  and its rank- $K_H$  update factorization  $\mathbf{U}_H, \mathbf{V}_H$  by Lemma 4.3
43: end procedure

```

LEMMA 4.4. Let \mathbf{w} be the weights changing at every step of the IPM. Let \mathcal{S} be any separator tree for $G_{\mathbf{A}}$. Recall $G_{\mathbf{A}}$ has n vertices, m hyperedges, and max hyperedge size ρ . Then the data structure DYNAMICSC (Algorithm 1) correctly maintains the matrices $\mathbf{L}^{(H)}, (\mathbf{L}_{F_H, F_H}^{(H)})^{-1}, \mathbf{Sc}(\mathbf{L}^{(H)}, \partial H)$ at every node $H \in \mathcal{S}$ dependent on \mathbf{w} throughout the IPM. The data structure supports the following procedures and runtimes:

- INITIALIZE($\mathcal{S}, \mathbf{w}^{(\text{init})} \in \mathbb{R}^m$): Set $\mathbf{w} \leftarrow \mathbf{w}^{(\text{init})}$, and compute all matrices with respect to \mathbf{w} , in time

$$O\left(\sum_{\text{leaf } H \in \mathcal{S}} |E(H)| \cdot |F_H \cup \partial H|^{\omega-1} + \sum_{H \in \mathcal{S}} |F_H \cup \partial H|^{\omega}\right).$$

- REWEIGHT($\delta_{\mathbf{w}} \in \mathbb{R}^m$): Update the weight vector to $\mathbf{w} \leftarrow \mathbf{w} + \delta_{\mathbf{w}}$, and update all the maintained matrices with respect to the new weights, in time

$$O\left(\sum_{\text{leaf } H \in \mathcal{H}} \text{nnz}(\delta_{\mathbf{w}}|_{E(H)}) \cdot |F_H \cup \partial H|^{\omega-1} + \sum_{H \in \mathcal{H}} |F_H \cup \partial H|^2 \cdot K_H^{\omega-2}\right).$$

where \mathcal{H} is the set of nodes H with $\delta_{\mathbf{w}}|_{E(H)} \neq \mathbf{0}$, $K_H \stackrel{\text{def}}{=} \min\{\text{nnz}(\delta_{\mathbf{w}}|_{E(H)}), |F_H \cup \partial H|\}$.

Proof. INITIALIZE is a special case of REWEIGHT, where the change in the weight vector is from $\mathbf{0}$ to $\mathbf{w}^{(\text{init})}$, so we focus on a single call of REWEIGHT.

It suffices for REWEIGHT visits only nodes in \mathcal{H} , since if none of the edges in a region admits a weight update, then the matrices stored at the node remain the same by definition. Also note that $H \in \mathcal{H}$ implies all ancestors of H are also in \mathcal{H} .

Correctness. We use the superscript $^{(\text{new})}$ on $\mathbf{L}^{(H)}$ to indicate that it is computed with respect to the new weights, and $^{(\text{old})}$ otherwise. Recall that $\mathbf{L}^{(H)}$ is supported on $F_H \cup \partial H$.

We maintain some additional matrices at each node, in order to efficiently compute low-rank updates. Specifically, we use helper matrices $\mathbf{U}_H, \mathbf{V}_H$ at H , and guarantee that during a single REWEIGHT($\delta_{\mathbf{w}}$) call, after SCHURNODE($H, \delta_{\mathbf{w}}$) is run, they satisfy $\mathbf{Sc}(\mathbf{L}^{(H)(\text{old})}, \partial H) = \mathbf{Sc}(\mathbf{L}^{(H)(\text{new})}, \partial H) + \mathbf{U}_H \mathbf{V}_H$, and $\mathbf{U}_H, \mathbf{V}_H^{\top}$ both have at most K_H -many columns.

Now, we show inductively that after SCHURNODE($H, \delta_{\mathbf{w}}$) is run, all matrices at H , as well as all matrices at all descendants of H , are updated correctly: When H is leaf node, recall $\mathbf{L}^{(H)}$ is defined to be $\mathbf{L}[H] \stackrel{\text{def}}{=} \mathbf{A}_H \mathbf{W}_{E(H)} \mathbf{A}_H^{\top}$, so clearly SCHURNODE updates $\mathbf{L}^{(H)}$ correctly, and the rank of the update is at most K_H . The remaining matrices at H are computed correctly by Lemma 4.3.

Inductively, when H is a non-leaf node, the recursive property of Schur complements (c.f. [15, Lemma 18]) allows us to write $\mathbf{L}^{(H)(\text{new})} = \sum_{\text{child } D \text{ of } H} \mathbf{Sc}(\mathbf{L}^{(D)(\text{new})}, \partial D)$ at every node $H \in \mathcal{S}$. This formula trivially shows that the update $\mathbf{L}^{(H)(\text{new})} - \mathbf{L}^{(H)(\text{old})}$ has rank $|F_H \cup \partial H|$ (ie. full-rank). Alternatively, if $\text{nnz}(\delta_{\mathbf{w}}|_{E(H)}) \leq |F_H \cup \partial H|$, then by the guarantees on the helper matrices, we have

$$\begin{aligned} \mathbf{L}^{(H)(\text{new})} &= \sum_{\text{child } D \text{ of } H} \mathbf{Sc}(\mathbf{L}^{(D)(\text{new})}, \partial D) \\ &= \sum_{\text{child } D \text{ of } H} \mathbf{Sc}(\mathbf{L}^{(D)(\text{old})}, \partial D) + \mathbf{U}_D \mathbf{V}_D \\ &= \mathbf{L}^{(H)(\text{old})} + \sum_{\text{child } D \text{ of } H} \mathbf{U}_D \mathbf{V}_D. \end{aligned}$$

This gives a low-rank factorization of the update $\mathbf{L}^{(H)(\text{new})} - \mathbf{L}^{(H)(\text{old})}$ with rank at most $\sum_{\text{child } D \text{ of } H} K_D$, which we can show by induction is at most $\text{nnz}(\delta_{\mathbf{w}}|_{E(H)})$. Since we have the correct low-rank update to $\mathbf{L}^{(H)}$, the remaining matrices at H again are computed correctly by Lemma 4.3.

This completes the correctness proof.

Runtime. Consider the runtime of the procedure $\text{SCHURNODE}(H, \delta_{\mathbf{w}})$ at a node H : If H is a leaf node, then computing the update to $\mathbf{L}^{(H)}$ involves multiplying a $|F_H \cup \partial H| \times \text{nnz}(\delta_{\mathbf{w}}|_{E(H)})$ -sized matrix with its transpose (Algorithm 1). Note that if $|F_H \cup \partial H| > \text{nnz}(\delta_{\mathbf{w}}|_{E(H)})$, then this runtime can be absorbed into the runtime expression for the remaining steps of the procedure, since $K_H = \text{nnz}(\delta_{\mathbf{w}}|_{E(H)})$. Otherwise, we use fast matrix multiplication which takes $O(\text{nnz}(\delta_{\mathbf{w}}|_{E(H)}) \cdot |F_H \cup \partial H|^{\omega-1})$ time. If H is a non-leaf node, there are two cases for the update to $\mathbf{L}^{(H)}$ in the algorithm. The first case (Algorithm 1) takes $O(|F_H \cup \partial H| \cdot (\sum K_D)) \leq O(|F_H \cup \partial H| \cdot K_H)$ time, and the second case (Algorithm 1) takes $O(|F_H \cup \partial H|^2)$ time. Computing the other matrices at any node H takes $O(|F_H \cup \partial H|^2 \cdot K_H^{\omega-2})$ time by Lemma 4.3.

The runtime of $\text{REWEIGHT}(\delta_{\mathbf{w}})$ is therefore given by

$$\begin{aligned} & \sum_{H \in \mathcal{H}} \text{SCHURNODE}(H, \delta_{\mathbf{w}}) \text{ time} \\ &= O\left(\sum_{\text{leaf } H \in \mathcal{H}} \text{nnz}(\delta_{\mathbf{w}}|_{E(H)}) \cdot |F_H \cup \partial H|^{\omega-1} + \sum_{H \in \mathcal{H}} |F_H \cup \partial H|^2 \cdot K_H^{\omega-2}\right). \end{aligned}$$

For INITIALIZE, we further simplify the expression using $\text{nnz}(\delta_{\mathbf{w}}|_{E(H)}) = |E(H)|$ and $K_H \leq |F_H \cup \partial H|$. \square

4.5 Projection operator complexities In this subsection, we summarize the runtime complexities for the tree operator, in the special case when \mathcal{S} is a (a, b, λ) -separator tree for $G_{\mathbf{A}}$. Parametrizing the separator tree this way allows us to write the runtime expressions using geometric series. For non-negative x , we use the standard bound $\sum_{i=\ell}^u x^i \leq O(x^\ell + x^u)$. When it is clear $x < 1$, we bound $\sum_{i=\ell}^u x^i \leq O(x^\ell)$.

LEMMA 4.5. *Suppose \mathcal{S} is an (a, b, λ) -separator tree for $G_{\mathbf{A}}$ on n vertices, m edges, with max hyperedge size ρ , where $a \in [0, 1]$ and $b \in (0, 1)$. Let η denote the height of \mathcal{S} , and let L denote the number of leaf nodes. Let Δ be the tree operator on \mathcal{T} as defined in Section 4.3. Then there is a data structure to maintain Δ as a function of the weights \mathbf{w} throughout SOLVE, so that:*

- The data structure can be initialize in time

$$(4.13) \quad O(\rho^{\omega-1}m + \lambda^\omega \cdot (1 + (b^{a\omega-1})^\eta)).$$

- The query complexity of Δ is

$$(4.14) \quad Q(K) = O(\rho K + \lambda^2 (1 + (\min\{K, L\})^{1-2a}))$$

- When $a < 1$, the update complexity of Δ is $U(K) =$

$$(4.15) \quad \rho^{\omega-1}K + \lambda^2 \min\{K, \lambda\}^{\omega-2} + \begin{cases} \lambda^2 K^{1-2a} & \text{if } K \leq \lambda \\ \lambda^2 K^{1-2a} + \lambda^{\frac{\omega-1}{1-\alpha}} K^{\frac{1-\alpha\omega}{1-\alpha}} & \text{if } \lambda < K \leq \lambda \cdot b^{(a-1)\eta} \\ \lambda^\omega \cdot b^{(a\omega-1)\eta} & \text{if } K > \lambda \cdot b^{(a-1)\eta}. \end{cases}$$

When $a = 1$, the update complexity is $U(K) = \rho^{\omega-1}K + \lambda^2 \min\{K, \lambda\}^{\omega-2}$.

Proof. The data structure we use to maintain Δ is precisely the data structure DYNAMICSC with respect to \mathcal{S} .

Initialization time. We use the runtime expression for INITIALIZE in DYNAMICSC (Lemma 4.4) combined with the parameters of the (a, b, λ) -separator tree. For any H , we have $|F_H \cup \partial H| \leq \rho$, so $\sum_{\text{leaf } H \in \mathcal{S}} |E(H)| \cdot |F_H \cup \partial H|^{\omega-1} \leq \rho^{\omega-1}m$. Moreover,

$$\sum_{H \in \mathcal{S}} |F_H \cup \partial H|^\omega \leq \sum_{i=0}^{\eta} b^{-i} (\lambda \cdot b^{ai})^\omega \leq O(\lambda^\omega) \cdot [1 + (b^{a\omega-1})^\eta].$$

Query complexity. We substitute the (a, b, λ) -separator tree bounds in Lemma 4.2, to conclude that the query complexity of Δ is

$$Q(K) = O\left(\rho K + \sum_{H \in \mathcal{P}_S(\mathcal{H})} (\lambda \cdot b^{ai})^2\right),$$

where \mathcal{H} is any set of K leaf nodes in \mathcal{S} . We group terms according to their node level, and note that there are $\min\{K, b^{-i}\}$ -many terms at any level i , so

$$\begin{aligned} &= O\left(\rho K + \sum_{i=0}^{\eta} \min\{K, b^{-i}\} \cdot (\lambda \cdot b^{ai})^2\right) \\ &= O(\rho K) + O(\lambda^2) \cdot \left(\sum_{i=0}^{-\log_b K} b^{-i} \cdot b^{2ai} + K \cdot \sum_{i=-\log_b K}^{\eta} b^{2ai}\right) \end{aligned}$$

Note that K can be at most L in the summation, so we have

$$= O(\rho K + \lambda^2(1 + (\min\{K, L\})^{1-2a})).$$

Update complexity. When \mathbf{w} changes, we update Δ by invoking $\text{REWEIGHT}(\delta_{\mathbf{w}})$ in DYNAMICSC , where $\delta_{\mathbf{w}}$ denotes the change in \mathbf{w} . By Lemma 4.4, the runtime for the fixed $\delta_{\mathbf{w}}$ is

$$(4.16) \quad \sum_{\text{leaf } H: \delta_{\mathbf{w}}|_{E(H)} \neq \mathbf{0}} \text{nnz}(\delta_{\mathbf{w}}|_{E(H)}) \cdot |F_H \cup \partial H|^{\omega-1} + \sum_{\text{node } H: \delta_{\mathbf{w}}|_{E(H)} \neq \mathbf{0}} |F_H \cup \partial H|^2 \cdot K_H^{\omega-2}.$$

Hence, the update complexity of Δ is the max of the above expression taken over all choices of $\delta_{\mathbf{w}}$. For any leaf node H , we upper bound $|F_H \cup \partial H|^{\omega-1} \leq \rho^{\omega-1}$, and therefore the first summation is at most $\rho^{\omega-1}K$.

For the second summation, we substitute in the (a, b, λ) -separator tree bounds, and group terms according to their node level. Let $\mathcal{S}(i)$ denotes all nodes at level i in \mathcal{S} . Then for any $\delta_{\mathbf{w}}$, we have

$$(4.17) \quad \sum_{\text{node } H: \delta_{\mathbf{w}}|_{E(H)} \neq \mathbf{0}} |F_H \cup \partial H|^2 \cdot K_H^{\omega-2} \leq \sum_{i=0}^{\eta} \left((\lambda \cdot b^{ai})^2 \cdot \sum_{H \in \mathcal{S}(i)} K_H^{\omega-2} \right),$$

where the K_H 's are non-negative integers satisfying $\sum_{H \in \mathcal{S}(i)} K_H \leq K$ and $K_H \leq |F_H \cup \partial H| \leq \lambda \cdot b^{ai}$ for $H \in \mathcal{S}(i)$. We are interested in upper bounding Eq. (4.17). At any level i , there are b^{-i} nodes, and the sum is maximized when all the K_H 's are equal. Depending on the relationship between K and the level i , we have the following three cases:

- If $K \leq b^{-i}$, that is, the total update rank is less than the number of nodes at the level, then the sum is maximized if $K_H = 1$ for K -many nodes, and $K_H = 0$ for the rest.
- If $b^{-i} < K \leq b^{-i} \cdot (\lambda \cdot b^{ai})$, the sum is upper bounded by setting $K_H = K/b^{-i}$.
- If $K > O(b^{-i}) \cdot (\lambda \cdot b^{ai})$, the sum is upper bounded by setting $K_H = \lambda \cdot b^{ai}$.

Then, we can bound the summation term in Eq. (4.17) by

$$\begin{aligned} &\sum_{\substack{0 \leq i \leq \eta \\ K \leq b^{-i}}} K(\lambda \cdot b^{ai})^2 + \sum_{\substack{0 \leq i \leq \eta: \\ b^{-i} < K \leq \lambda \cdot b^{(a-1)i}}} (\lambda \cdot b^{ai})^2 \cdot b^{-i} \cdot (Kb^i)^{\omega-2} + \sum_{\substack{0 \leq i \leq \eta: \\ K > \lambda \cdot b^{(a-1)i}}} b^{-i} \cdot (\lambda \cdot b^{ai})^{\omega} \\ &\leq \lambda^2 K \sum_{i=-\log_b K}^{\eta} b^{2ai} + \lambda^2 K^{\omega-2} \sum_{i=\frac{\log_b(K/\lambda)}{a-1}}^{-\log_b K} b^{(2a+\omega-3)i} + \lambda^{\omega} \sum_{i=0}^{\frac{\log_b(K/\lambda)}{a-1}} b^{(a\omega-1)i}. \end{aligned}$$

We need to further consider different cases for the possible values of K , which affects the summation indices. If $\log_b(K/\lambda) < 0$, i.e. $K < \lambda$, the expression simplifies to

$$\lambda^2 K^{1-2a} + \lambda^2 K^{\omega-2}.$$

If $0 \leq \log_b(K/\lambda)/(a-1) \leq \eta$, i.e. $\lambda \leq K \leq \lambda \cdot b^{(a-1)\eta}$, the expression simplifies to

$$\lambda^2 K^{1-2a} + \lambda^{\frac{\omega-1}{1-\alpha}} K^{\frac{1-\alpha\omega}{1-\alpha}} + \lambda^\omega.$$

And lastly, if $\log_b(K/\lambda)/(a-1) > \eta$, i.e. $K > \lambda \cdot b^{(a-1)\eta}$, the expression simplifies to

$$\lambda^\omega + \lambda^\omega \cdot b^{(a\omega-1)\eta}.$$

We combine the cases to arrive at the overall update complexity, having implicitly assumed that $\alpha < 1$. When $\alpha = 1$, the summation in Eq. (4.17) is maximized when $K_H = \min\{K, \lambda\} \cdot b^i$ for H at level i . Then we can upper bound the summation term by

$$\sum_{i=0}^{\eta} (\lambda \cdot b^i)^2 \cdot b^{-i} \cdot (\min\{K, \lambda\} \cdot b^i)^{\omega-2} \leq O(\lambda^2 \min\{K, \lambda\}^{\omega-2}).$$

□

5 Proofs of main theorems

For our main theorems, it remains to show that we can construct an appropriate (a, b, λ) -separator tree for $G_{\mathbf{A}}$ in each of the scenarios: when $G_{\mathbf{A}}$ is n^α -separable; when \mathbf{A} is the constraint matrix for a planar k -multicommodity flow instance; and when $G_{\mathbf{A}}$ has a tree decomposition of width τ . Then, we apply Lemma 4.5 to the separator tree get the complexity of the tree operator, which we combine with Theorem 3.4 to conclude the overall IPM running times.

5.1 Proof of Theorem 1.1 First, we show how to construct a separator tree for an n^α -separable graph, by modifying the proof from [20].

LEMMA 5.1. *Suppose $G_{\mathbf{A}}$ is a graph on n vertices and m edges. If $G_{\mathbf{A}}$ is n^α -separable for $\alpha < 1$, then $G_{\mathbf{A}}$ admits an (α, b, cn^α) -separator tree, where $b \in (0, 1)$ and $c > 0$ are some constants. Furthermore, if a balanced vertex separator for $G_{\mathbf{A}}$ can be computed in $T(n)$ time, then the separator tree can be computed in $\tilde{O}(T(n))$ time.*

Proof. Let $b' \in (0, 1)$ and $c' = 1$ (without loss of generality) be the parameters for $G_{\mathbf{A}}$ being n^α -separable. In the separator tree construction process, assume inductively that we have constants $b \in (0, 1)$ and $c > 0$, both to be chosen later, such that for any node H at level i , we have $|V(H)| \leq b^i n$ and $|\partial H| \leq cn^\alpha \cdot b^{\alpha i}$. In the base case at the root node, we have $i = 0$, and $|V(G_{\mathbf{A}})| \leq n$ and $|\partial G_{\mathbf{A}}| = 0 \leq cn^\alpha$.

We show how to construct the nodes at level $i + 1$. Let H be an already-constructed node at level i . There are three cases:

1. If H satisfies $|V(H)| \leq b^{i+1}n$ and $|\partial H| \leq cn^\alpha \cdot b^{\alpha(i+1)}$, put a copy of H as its only child at level $i + 1$.
2. If $|V(H)| \geq b^{i+1}n$, then assign a weight of 1 to all vertices, find a balanced vertex separator $S(H)$, and partition H accordingly into H_1 and H_2 . Let us consider H_1 ; the analogous holds for H_2 .

By definition of separability, we know $|V(H_1)| \leq b' \cdot |V(H)| + |V(H)|^\alpha \leq b \cdot |V(H)| \leq b^{i+1}n$ as long as $b \in (b', 1)$. If $|\partial H_1| \leq c \cdot |V(H_1)|^\alpha$, then we can upper bound this expression by $cn^\alpha \cdot b^{\alpha(i+1)}$, and we are done.

On the other hand, if $|\partial H_1| > c \cdot |V(H_1)|^\alpha$, then by definition of boundary, we have $|\partial H_1| \leq |\partial H| + |S(H)| \leq (c + 1)n^\alpha \cdot b^{\alpha i}$ using the guarantees at H . Next, we assign a weight of 1 to vertices in ∂H_1 and 0 to all other vertices, find a balanced separator $S(H_1)$ of H_1 with respect to these weights, and create two children

D_1, D_2 of H_1 accordingly. Then, for $j = 1, 2$, we have

$$\begin{aligned} |\partial D_j| &\leq b \cdot |\partial H_1| + |V(H_1)|^\alpha \\ &\leq b(c+1)n^\alpha \cdot b^{\alpha i} + n^\alpha \cdot b^{\alpha(i+1)} \\ &\leq \left(b^{1-\alpha} \cdot \frac{c+1}{c} + \frac{1}{c} \right) \cdot cn^\alpha \cdot b^{\alpha(i+1)}, \end{aligned}$$

As long as c is large enough so the expression in the parentheses to be less than 1. In this case, observe that we can add $S(H_1)$ to the balanced separator $S(H)$, and set D_1 and D_2 directly as the children of H .

3. If $|V(H)| \leq b^{(i+1)n}$ and $|\partial H| \geq cn^\alpha \cdot b^{\alpha(i+1)}$, then we apply case 2 with H_1 being H .

So we have shown inductively that at the end of this construction, any node H at level i satisfies $|V(H)| \leq b^i n$ and $|\partial H| \leq cn^\alpha \cdot b^{\alpha i}$. It follows that $|F_H \cup \partial H| \leq |S(H)| + |\partial H| = O(cn^\alpha \cdot b^{\alpha i})$.

Next, we show that there are only $O(b^{-i})$ nodes at level i . Let $L_i(n)$ denote the total number of boundary vertices with multiplicities, when carrying out the construction starting on a graph of size n and ending when each leaf node H satisfies the level- i assumptions. We can recursively write

$$\begin{aligned} L_i(k) &= \sum_{j=1}^4 L_i(b_j k + 3ck^\alpha), & \text{if } k > Cb^i n \\ B_i(k) &= 1 & \text{else.} \end{aligned}$$

where $\sum b_j = 1$, each $b_j \leq b'$, and C is a positive constant we choose. To see this, note that a node of size k has at most four children in the construction; the separator is of size $3ck^\alpha$ since we may need to compute up to three separators each of size ck^α and take their union; and child j has at most $b_j k$ vertices that are not from the separator. Solving the recursion yields $L_i(k) \leq k/(Cb^i n) - \gamma k^\alpha$ for some constant $\gamma > 0$. Therefore, there are at most $L_i(n) \leq O(b^{-i})$ nodes at level i .

Finally, it is straightforward to see that the separator tree can be computed in $\tilde{O}(T(n))$ time, since the node sizes decrease by a geometric factor as we proceed down the tree during construction. \square

Proof. [Proof of Theorem 1.1] We consider the cases when $\alpha = 1$ and $\alpha < 1$ separately.

All hypergraphs are trivially n -separable with max hyperedge size $\rho = n$. In this case, let \mathcal{S} be the separator tree consisting of simply one node representing $G_{\mathbf{A}}$, which is a $(1, 1/2, n)$ -separator tree. By Lemma 4.5, the tree operator data structure can be initialized in $O(m^\omega)$ time; the query complexity is $Q(K) = O(nK + n^2)$, and the update complexity is $U(K) = O(n^{\omega-1}K + n^2K^{\omega-2})$.

We apply Theorem 3.4 to get the overall runtime:

$$\tilde{O} \left(\sqrt{m} \log\left(\frac{R}{\epsilon r}\right) \cdot \sum_{\ell=0}^{\frac{1}{2} \log m} \frac{n2^{2\ell} + n^2 + n^2 2^{2\ell(\omega-2)}}{2^\ell} \right) = \tilde{O} \left(\sqrt{mn^2} \log\left(\frac{R}{\epsilon r}\right) \right).$$

If $G_{\mathbf{A}}$ is n^α -separable for $\alpha < 1$, then by Lemma 5.1, $G_{\mathbf{A}}$ admits a (α, b, cn^α) -separator tree computable in $\tilde{O}(n)$ time. In this case, $\rho = O(1)$, and $\eta = O(\log_{1/b} n)$. Plugging the parameters into Lemma 4.5, we get the following tree operator runtimes:

- The data structure can be initialize in $O(m + n^{\alpha\omega}(1 + n^{1-\alpha\omega})) \leq O(m + m^{\alpha\omega})$ time.
- The query complexity is $Q(K) \leq O(K + n^{2\alpha}(1 + K^{1-2\alpha}))$.
- The update complexity is

$$U(K) \leq O \left(K + n^{2\alpha} \min\{K, n^\alpha\}^{\omega-2} + n^{2\alpha} K^{1-2\alpha} + n^{\frac{\alpha(\omega-1)}{1-\alpha}} K^{\frac{1-\alpha\omega}{1-\alpha}} \cdot \mathbb{1}_{K \geq n^\alpha} \right).$$

We apply Theorem 3.4 to get the overall runtime.

$$\begin{aligned} & \tilde{O}\left(\sqrt{m} \log\left(\frac{R}{\varepsilon r}\right)\right) \cdot \sum_{\ell=0}^{\frac{1}{2} \log m} \frac{2^{2\ell} + n^{2\alpha} + n^{\alpha\omega} \cdot \mathbb{1}_{2^{2\ell} > n^\alpha} + n^{2\alpha} 2^{(1-2\alpha)2\ell} + n^{\frac{\alpha(\omega-1)}{1-\alpha}} 2^{\frac{1-\alpha\omega}{1-\alpha} 2\ell} \cdot \mathbb{1}_{2^{2\ell} > n^\alpha}}{2^\ell} \\ &= \tilde{O}\left(\sqrt{m} \log\left(\frac{R}{\varepsilon r}\right)\right) \cdot \left(\sqrt{m} + n^{2\alpha} + n^{\alpha\omega - \frac{\alpha}{2}} + n^{2\alpha} m^{1-2\alpha - \frac{1}{2}} + n^{\frac{\alpha(\omega-1)}{1-\alpha}} \left(n^{\frac{\alpha(1-\alpha\omega)}{1-\alpha} - \frac{\alpha}{2}} + m^{\frac{1-\alpha\omega}{1-\alpha} - \frac{1}{2}}\right)\right) \\ &= \tilde{O}\left(\left(m + m^{\frac{1}{2} + 2\alpha}\right) \cdot \log\left(\frac{R}{\varepsilon r}\right)\right), \end{aligned}$$

where in the last step, we used the fact $\alpha\omega - \frac{\alpha}{2} \leq 2\alpha$. \square

5.2 Proof of Theorem 1.2 Let $G = (V, E)$ denote the planar graph for the original problem, with $V = \{v_1, \dots, v_n\}$ and $E = \{e_1, \dots, e_m\}$. First, we write the LP in Eq. (1.1) in standard form by adding slack variables $\mathbf{s} \in \mathbb{R}^E$:

$$(P') \quad \begin{aligned} & \min \sum_{i=1}^k \mathbf{c}_i^\top \mathbf{f}_i \\ & \text{s.t.} \quad \mathbf{B}^\top \mathbf{f}_i = \mathbf{d}_i \quad \forall i \in [k] \\ & \sum_{i=1}^k \mathbf{f}_i + \mathbf{s} = \mathbf{u} \\ & \mathbf{f}_i \geq \mathbf{0} \quad \forall i \in [k] \\ & \mathbf{s} \geq \mathbf{0} \end{aligned}$$

Let \mathbf{A} denote the full constraint matrix of P' . Then

$$(5.1) \quad \mathbf{A} = \left[\begin{array}{cccc|c} \mathbf{B}^\top & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{B}^\top & & \vdots & \vdots \\ & & \ddots & & \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{B}^\top & \mathbf{0} \\ \hline \mathbf{I} & \mathbf{I} & \dots & \mathbf{I} & \mathbf{I} \end{array} \right] \in \mathbb{R}^{(kn+m) \times (k+1)m}$$

where the top left part of \mathbf{A} contains k copies of \mathbf{B}^\top in block-diagonal fashion, and all the identity matrices are of dimension $m \times m$. The dual graph of \mathbf{B}^\top is precisely G . Let $G_{\mathbf{A}}$ be the dual graph of \mathbf{A} .

First, we describe $G_{\mathbf{A}}$: It contains k independent copies of the vertices V , which we label with $V^i = (v_1^i, \dots, v_n^i)$, so that v_j^i is a copy of $v_j \in V$. Additionally, $G_{\mathbf{A}}$ contains m vertices u_1, \dots, u_m , where the vertex u_i is identified with edge $e_i \in E$. For each edge $e_i \in E$ with endpoints v_{i_1}, v_{i_2} , there are k hyper-edges in $G_{\mathbf{A}}$ of the form $\{v_{i_1}^\ell, v_{i_2}^\ell, u_i\}$ for $\ell = 1, \dots, k$. Additionally, there are m hyper-edges f_1, \dots, f_m where f_i contains only the vertex u_i .

Next, we show how to construct an appropriate separator tree efficiently.

CLAIM 6. $G_{\mathbf{A}}$ admits a $(\frac{1}{2}, b, kn^{1/2})$ -separator tree that can be computed in $O(kn \log n)$ time.

Proof. Let G be the original planar graph which is \sqrt{n} -separable, and let $\tilde{\mathcal{S}}$ be the $(\frac{1}{2}, b, n^{1/2})$ -separator tree for G constructed using Lemma 5.1 in $O(n \log n)$ time by [49]. We show how to construct a $(\frac{1}{2}, b, kn^{1/2})$ -separator tree \mathcal{S} for $G_{\mathbf{A}}$ based on $\tilde{\mathcal{S}}$. Without loss of generality, we ignore the hyper-edges f_1, \dots, f_m in this construction.

Intuitively, \mathcal{S} will have the same tree structure as $\tilde{\mathcal{S}}$, but each node will be larger by a factor of $O(k)$ due to the k copies of G in $G_{\mathbf{A}}$. For each $\tilde{H} \in \tilde{\mathcal{S}}$, we construct a corresponding $H \in \mathcal{S}$ as follows: if $v_j \in \tilde{H}$, then $v_j^i \in H$ for all $i \in [k]$; if $e_j \in E(\tilde{H})$, i.e. both endpoints of e_j are in \tilde{H} , add u_j to H . Since the k copies v_j^1, \dots, v_j^k are always grouped together, we will refer to them together as v_j in $G_{\mathbf{A}}$ as well.

Let us show that this is indeed a $(\frac{1}{2}, b, kn^{1/2})$ -separator tree. Suppose H is a node with children D_1 and D_2 in \mathcal{S} , corresponding to nodes $\tilde{H}, \tilde{D}_1, \tilde{D}_2$ in $\tilde{\mathcal{S}}$. Let $S(H) \stackrel{\text{def}}{=} V(D_1) \cap V(D_2)$, then $v_j \in S(H)$ iff $v_j \in S(\tilde{H})$, and $u_j \in S(H)$ iff $e_j \in E(S(\tilde{H}))$ for all values of j . It is straightforward to see that $S(H)$ is indeed a separator of H . When it comes to the set of boundary vertices, we see $v_j \in \partial H$ iff $v_j \in \partial \tilde{H}$, and if $u_j \in \partial H$ with v_{j_1}, v_{j_2} being the two endpoints of e_j , then v_{j_1}, v_{j_2} are both in ∂H . Since G is a planar graph, the number of edges in \tilde{H} is on the same order as the number of vertices, so we conclude that $|V(H)| \leq O(k) \cdot |V(\tilde{H})|$, and similarly, $|F_H \cup \partial H| \leq O(k) \cdot |F_{\tilde{H}} \cup \partial \tilde{H}|$. Since node sizes in \mathcal{S} have increased by a factor of $O(k)$ compared to $\tilde{\mathcal{S}}$, we conclude \mathcal{S} is a $(\frac{1}{2}, b, kn^{1/2})$ -separator tree.

Finally, we can compute $\tilde{\mathcal{S}}$ for G in $O(n \log n)$ time, so we can compute \mathcal{S} in $O(kn \log n)$ time. \square

We reduce our problem to minimum cost multi-commodity circulation problem in order to establish the existence of an interior point in the polytope, before invoking the RIPM in Theorem 3.1. For each commodity $i \in [k]$, we add extra vertices s_i and t_i . Let \mathbf{d}_i be the demand vector of the i -th commodity. For every vertex v with $\mathbf{d}_{i,v} < 0$, we add a directed edge from s_i to v with capacity $-\mathbf{d}_{i,v}$ and cost 0. For every vertex v with $\mathbf{d}_{i,v} > 0$, we add a directed edge from v to t_i with capacity $\mathbf{d}_{i,v}$ and cost 0. Then, we add a directed edge from t_i to s_i with capacity $4kmM$ and cost $-4kmM$. The modified graph G' has only $2k$ extra vertices of the form s_i and t_i compared to $G_{\mathbf{A}}$, so we can construct a $(\frac{1}{2}, b, kn^{1/2} + 2k)$ -separator tree for G' based on the $(\frac{1}{2}, kn^{1/2})$ -separator tree for $G_{\mathbf{A}}$, where we include the extra vertices at every node of the tree.

To show the existence of the interior point, we remove all the directed edges that no single commodity flow from s_i to t_i can pass for any $i \in [k]$. This can be done by run BFS for k times which takes $O(km)$ time. For the interior point \mathbf{f} , we construct this finding a circulation $\mathbf{f}^{(e)}$ that passing through e and s_i, t_i for some i with flow value $1/(10km)$ for all the remaining edge e . Then, since the capacities are integers, we find a feasible \mathbf{f}, \mathbf{s} with value at least $1/(10km)$. This shows the inner radius r of the polytope is at least $1/(10km)$. For the L and R , we note we can bound it by $O(kmM)$.

Let \mathbf{A}' be the constraint matrix of the reduced problem with dual graph G' . The RIPM in Theorem 3.1 invokes the subroutine SOLVE twice. In the first run, we make a new constraint matrix by concatenating \mathbf{A}' three times. One can check that the dual graph is G' with each edge duplicated three times, so the corresponding separator tree is straightforward to construct.

Now, we bounding the running time. The tree operator complexities are similar to the analysis in the previous section with an additional factor of k in the expression for λ . The initialization time is $O(km + (kn^{1/2})^\omega)$. The query complexity is $Q(K) = O(K + k^2n)$. After simplifying, the update complexity is

$$U(K) = K + \begin{cases} k^2nK^{\omega-2} & \text{if } K \leq kn^{1/2} \\ (kn^{1/2})^\omega & \text{else.} \end{cases}$$

Note that the number of variables is km . Plugging our choice of L, R , and r , by Theorem 3.4, the total runtime simplifies to

$$\tilde{O}(k^{2.5}m^{1.5} \log(M/\varepsilon)).$$

6.1 Proof of Theorem 1.3 First, we show how to construct a $(0, 1/2, O(\tau \log n))$ -separator tree \mathcal{S} for $G_{\mathbf{A}}$ when we have a tree decomposition of $G_{\mathbf{A}}$ of width τ . At the root of \mathcal{S} , we can use the tree decomposition to compute a balanced separator S of $G_{\mathbf{A}}$ of size $O(\tau)$ in $\tilde{O}(n\tau)$ time (c.f. [16, Theorem 4.17]), so that the two parts A and B of $G_{\mathbf{A}} \setminus S$ each have size at most $\frac{2}{3}n$. We construct two children of the root node on the vertex sets $A \cup S$ and $B \cup S$ respectively, and apply this procedure recursively until the nodes are of size at most 9τ .

CLAIM 7. *There are $O(n/\tau)$ -many leaves at the end of this construction.*

Proof. Let $L(k)$ denote the number of leaves when starting the construction with a size k subgraph. We know $L(k) = 1$ if $k \leq 9\tau$, and $L(k) = L(k_1 + \tau) + L(k_2 + \tau)$ if $k > 9\tau$, where $k_1 + k_2 + \tau = k$ and $k_1, k_2 \leq 2/3k$. By induction, we can show that $L(k) \leq 2(k/\tau - 1)$ when $k > 2\tau$, where the balanced separator crucially ensures that the recursion does not reach the base case of $k \leq 2\tau$. \square

The resulting separator tree is binary, so there are at most 2^i nodes at level i . Since there are $L = O(n/\tau)$ -many leaves, the height η is at most $\eta \leq \log_2(n/\tau)$. The boundary of a node H is contained in the union of balanced separators over its ancestors, so $|F_H \cup \partial H| \leq \tau\eta \leq O(\tau \log n)$. The max hyperedge size of $G_{\mathbf{A}}$ is $\rho = \tau$.

Using these values, we simplify the complexities in Lemma 4.5: The initialization time for the tree operator data structure is $\tilde{O}(\tau^{\omega-1}m + \tau^\omega(1 + n/\tau)) = \tilde{O}(\tau^{\omega-1}m)$. The query complexity of Δ is $Q(K) = \tilde{O}(\tau K + \tau^2 \min\{K, L\})$. The update complexity of Δ is

$$U(K) \leq \tau^{\omega-1}K + \begin{cases} \tau^2 K & \text{if } K \leq n \\ \tau^\omega & \text{if } K > n \end{cases}$$

Finally, we apply Theorem 3.4 to get the overall runtime, which is clearly bounded by

$$\tilde{O}\left(\sqrt{m} \log\left(\frac{R}{\varepsilon r}\right)\right) \cdot \sum_{\ell=0}^{\frac{1}{2} \log m} \frac{\tau^{2\ell} 2^{2\ell}}{2^\ell} = \tilde{O}(m\tau^2 \log(R/(\varepsilon r))).$$

To obtain the faster runtime given in [28], we use the data structure restarting trick: Recall MAINTAINAPPROX guarantees there are $2^{2\ell}$ -many coordinate updates to \bar{x} and \bar{s} every 2^ℓ steps, i.e. the number of coordinate updates grows superlinearly with respect to the total number of steps taken. By reinitializing MAINTAINAPPROX with the exact solution once in a while, we limit the total number of coordinate updates. In the proof of Theorem 3.4, we showed that running M steps of the RIPM takes

$$\tilde{O}\left(U(m) + Q(m) + \eta^4 M \log\left(\frac{R}{\varepsilon r}\right) \cdot \sum_{\ell=0}^{\log M} \frac{U(2^{2\ell}) + Q(2^{2\ell})}{2^\ell}\right)$$

time, where $U(m) + Q(m)$ is the time to initialize the data structures and obtain the final exact solutions. There are $N = \sqrt{m} \log m \log(\frac{mR}{\varepsilon r})$ -many total IPM steps, and we reinitialize the data structures every M steps. Then the total running time is (ignoring the big-O notation and log factors)

$$\begin{aligned} & \frac{N}{M} \left(U(m) + Q(m) + M \sum_{\ell=0}^{\log M} \frac{U(2^{2\ell}) + Q(2^{2\ell})}{2^\ell} \right) \\ &= \frac{\sqrt{m}}{M} (\tau^{\omega-1}m + \tau^2 M^2). \end{aligned}$$

The expression is minimized by taking $M = \sqrt{m\tau^{\frac{\omega-3}{2}}}$, which gives an overall runtime of

$$\tilde{O}\left(m\tau^{(\omega+1)/2} \log(R/(\varepsilon r))\right).$$

Acknowledgements We thank the anonymous reviewers for helpful feedback.

References

- [1] AHUJA, R. K., MAGNANTI, T. L., AND ORLIN, J. B. *Network Flows*. Prentice Hall, 1988.
- [2] ALON, N., AND YUSTER, R. Solving linear systems through nested dissection. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science (2010)*, IEEE, pp. 225–234.
- [3] AXIOTIS, K., MADRY, A., AND VLADU, A. Faster sparse minimum cost flow by electrical flow localization. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS) (2022)*, pp. 528–539.
- [4] BERNSTEIN, A., GUTENBERG, M. P., AND SARANURAK, T. Deterministic decremental SSSP and approximate min-cost flow in almost-linear time. In *62st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021 (2021)*, IEEE.
- [5] BORRADAILE, G. *Exploiting Planarity for Network Flow and Connectivity Problems*. Brown University, 2008.
- [6] CHAMBERS, E. W., ERICKSON, J., FOX, K., AND NAYYERI, A. Minimum cuts in surface graphs. *SIAM J. Comput.* 52, 1 (2023), 156–195.
- [7] CHEN, L., KYNG, R., LIU, Y. P., PENG, R., GUTENBERG, M. P., AND SACHDEVA, S. Maximum flow and minimum-cost flow in almost-linear time. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 - November 3, 2022 (2022)*, IEEE, pp. 612–623.

- [8] CHRISTIANO, P., KELNER, J. A., MADRY, A., SPIELMAN, D. A., AND TENG, S.-H. 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.
- [9] COHEN, M. B., LEE, Y. T., AND SONG, Z. Solving linear programs in the current matrix multiplication time. *Journal of the ACM (JACM)* 68, 1 (2021), 1–39.
- [10] CORMEN, T. H., LEISERSON, C. E., RIVEST, R. L., AND STEIN, C. *Introduction to algorithms*. MIT press, 2009.
- [11] DAITCH, S. I., AND SPIELMAN, D. A. Faster approximate lossy generalized flow via interior point algorithms. In *Proceedings of the 40th annual ACM symposium on Theory of computing* (2008), pp. 451–460.
- [12] DANTZIG, G. B. Maximization of a linear function of variables subject to linear inequalities. *Activity analysis of production and allocation* 13 (1951), 339–347.
- [13] DIBBELT, J., STRASSER, B., AND WAGNER, D. Customizable contraction hierarchies. In *Experimental Algorithms* (Cham, 2014), J. Gudmundsson and J. Katajainen, Eds., Springer International Publishing, pp. 271–282.
- [14] DING, M., KYNG, R., AND ZHANG, P. Two-commodity flow is equivalent to linear programming under nearly-linear time reductions. In *49th International Colloquium on Automata, Languages, and Programming, ICALP 2022, July 4-8, 2022, Paris, France* (2022), M. Bojanczyk, E. Merelli, and D. P. Woodruff, Eds., vol. 229 of *LIPICs*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 54:1–54:19.
- [15] DONG, S., GAO, Y., GORANCI, G., LEE, Y. T., PENG, R., SACHDEVA, S., AND YE, G. Nested dissection meets ipms: Planar min-cost flow in nearly-linear time. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* (2022), SIAM, pp. 124–153.
- [16] DONG, S., LEE, Y. T., AND YE, G. A nearly-linear time algorithm for linear programs with small treewidth: A multiscale representation of robust central path. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing* (2021), STOC 2021, ACM, pp. 1784–1797.
- [17] DONG, S., LEE, Y. T., AND YE, G. A nearly-linear time algorithm for linear programs with small treewidth: A multiscale representation of robust central path. *arXiv preprint arXiv:2011.05365v2* (2021).
- [18] DUAN, R., WU, H., AND ZHOU, R. Faster matrix multiplication via asymmetric hashing. *arXiv preprint arXiv:2210.10173* (2022).
- [19] EPPSTEIN, D., GALIL, Z., ITALIANO, G. F., AND SPENCER, T. H. Separator based sparsification: I. planarity testing and minimum spanning trees. *journal of computer and system sciences* 52, 1 (1996), 3–27.
- [20] FEDERICKSON, G. N. Fast algorithms for shortest paths in planar graphs, with applications. *SIAM Journal on Computing* 16, 6 (1987), 1004–1022.
- [21] FLEISCHER, L. K. Approximating fractional multicommodity flow independent of the number of commodities. *SIAM Journal on Discrete Mathematics* 13, 4 (2000), 505–520.
- [22] FORD, L. R., AND FULKERSON, D. R. Maximal flow through a network. *Canadian journal of Mathematics* 8 (1956), 399–404.
- [23] GAO, Y., LIU, Y. P., AND PENG, R. Fully dynamic electrical flows: Sparse maxflow faster than Goldberg-Rao. In *62st IEEE Annual Symposium on Foundations of Computer Science, FOCS2021* (2021), IEEE.
- [24] GARG, N., AND KÖNEMANN, J. Faster and simpler algorithms for multicommodity flow and other fractional packing problems. *SIAM Journal on Computing* 37, 2 (2007), 630–652.
- [25] GEORGE, A. Nested dissection of a regular finite element mesh. *SIAM journal on numerical analysis* 10, 2 (1973), 345–363.
- [26] GOHBERG, I., KAILATH, T., AND KOLTRACHT, I. Efficient solution of linear systems of equations with recursive structure. *Linear Algebra and its Applications* 80 (1986), 81–113.
- [27] GORANCI, G., HENZINGER, M., AND PENG, P. Dynamic effective resistances and approximate Schur Complement on separable graphs. In *26th Annual European Symposium on Algorithms, ESA 2018* (2018), vol. 112 of *LIPICs*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 40:1–40:15.
- [28] GU, Y., AND SONG, Z. A faster small treewidth SDP solver. *CoRR abs/2211.06033* (2022).
- [29] HASSIN, R. Maximum flow in (s, t) planar networks. *Information Processing Letters* 13, 3 (1981), 107.
- [30] HASSIN, R. On multicommodity flows in planar graphs. *Networks* 14, 2 (1984), 225–235.
- [31] HASSIN, R., AND JOHNSON, D. B. An $O(n \log^2 n)$ algorithm for maximum flow in undirected planar networks. *SIAM Journal on Computing* 14, 3 (1985), 612–624.
- [32] HENZINGER, M. R., KLEIN, P., RAO, S., AND SUBRAMANIAN, S. Faster shortest-path algorithms for planar graphs. *Journal of Computer and System Sciences* 55, 1 (1997), 3–23.
- [33] IMAI, H., AND IWANO, K. Efficient sequential and parallel algorithms for planar minimum cost flow. In *Algorithms, International Symposium SIGAL '90, Tokyo, Japan* (1990), vol. 450 of *Lecture Notes in Computer Science*, Springer, pp. 21–30.
- [34] ITAI, A. Two-commodity flow. *J. ACM* 25, 4 (1978), 596–611.
- [35] ITAI, A., AND SHILOACH, Y. Maximum flow in planar networks. *SIAM Journal on Computing* 8, 2 (1979), 135–150.
- [36] ITALIANO, G. F., NUSSBAUM, Y., SANKOWSKI, P., AND WULFF-NILSEN, C. Improved algorithms for min cut and

- max flow in undirected planar graphs. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011* (2011), ACM, pp. 313–322.
- [37] JIANG, S., SONG, Z., WEINSTEIN, O., AND ZHANG, H. A faster algorithm for solving general lps. In *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21–25, 2021* (2021), S. Khuller and V. V. Williams, Eds., ACM, pp. 823–832.
- [38] KARCZMARZ, A., AND SANKOWSKI, P. Min-cost flow in unit-capacity planar graphs. In *27th Annual European Symposium on Algorithms, ESA 2019, Munich/Garching, Germany* (2019), vol. 144 of *LIPICs*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 66:1–66:17.
- [39] KARMAKAR, N. A new polynomial-time algorithm for linear programming. In *Proceedings of the sixteenth annual ACM symposium on Theory of computing* (1984), pp. 302–311.
- [40] KATHURIA, T., LIU, Y. P., AND SIDFORD, A. Unit capacity maxflow in almost $o(m^{4/3})$ time. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16–19, 2020* (2020), pp. 119–130.
- [41] KAWARABAYASHI, K.-I., AND KOBAYASHI, Y. All-or-nothing multicommodity flow problem with bounded fractionality in planar graphs. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science* (2013), pp. 187–196.
- [42] KELNER, J. A., LEE, Y. T., ORECCHIA, L., AND SIDFORD, A. 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* (2014), SIAM, pp. 217–226.
- [43] KHACHIYAN, L. G. Polynomial algorithms in linear programming. *USSR Computational Mathematics and Mathematical Physics* 20, 1 (1980), 53–72.
- [44] KING, V., RAO, S., AND TARJAN, R. A faster deterministic maximum flow algorithm. *Journal of Algorithms* 17, 3 (1994), 447–474.
- [45] KYNG, R., AND SACHDEVA, S. Approximate gaussian elimination for laplacians-fast, sparse, and simple. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)* (2016), IEEE, pp. 573–582.
- [46] LEE, Y. T., AND SIDFORD, A. Solving linear programs with $\sqrt{\text{rank}}$ linear system solves. *CoRR abs/1910.08033* (2019).
- [47] LEE, Y. T., SONG, Z., AND ZHANG, Q. Solving empirical risk minimization in the current matrix multiplication time. In *Conference on Learning Theory* (2019), PMLR, pp. 2140–2157.
- [48] LEIGHTON, T., STEIN, C., MAKEDON, F., TARDOS, É., PLOTKIN, S., AND TRAGOUDAS, S. Fast approximation algorithms for multicommodity flow problems. In *Proceedings of the twenty-third annual ACM symposium on Theory of Computing* (1991), pp. 101–111.
- [49] LIPTON, R. J., ROSE, D. J., AND TARJAN, R. E. Generalized nested dissection. *SIAM journal on numerical analysis* 16, 2 (1979), 346–358.
- [50] LIPTON, R. J., AND TARJAN, R. E. A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics* 36, 2 (1979), 177–189.
- [51] MADRY, A. Faster approximation schemes for fractional multicommodity flow problems via dynamic graph algorithms. In *Proceedings of the forty-second ACM symposium on Theory of computing* (2010), pp. 121–130.
- [52] MADRY, A. Navigating central path with electrical flows: From flows to matchings, and back. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science* (2013), IEEE, pp. 253–262.
- [53] MATSUMOTO, K., NISHIZEKI, T., AND SAITO, N. An efficient algorithm for finding multicommodity flows in planar networks. *SIAM Journal on Computing* 14, 2 (1985), 289–302.
- [54] NESTEROV, Y., AND NEMIROVSKY, A. Acceleration and parallelization of the path-following interior point method for a linearly constrained convex quadratic problem. *SIAM Journal on Optimization* 1, 4 (1991), 548–564.
- [55] OKAMURA, H., AND SEYMOUR, P. Multicommodity flows in planar graphs. *Journal of Combinatorial Theory, Series B* 31, 1 (1981), 75–81.
- [56] PENG, R. Approximate undirected maximum flows in $o(m \text{ polylog}(n))$ time. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms* (2016), SIAM, pp. 1862–1867.
- [57] REIF, J. H. Minimum s - t cut of a planar undirected network in $O(n \log^2 n)$ time. *SIAM Journal on Computing* 12, 1 (1983), 71–81.
- [58] RENEGAR, J. A polynomial-time algorithm, based on newton’s method, for linear programming. *Mathematical programming* 40, 1-3 (1988), 59–93.
- [59] SCHILD, A., AND SOMMER, C. On balanced separators in road networks. In *Experimental Algorithms* (Cham, 2015), E. Bampis, Ed., Springer International Publishing, pp. 286–297.
- [60] SHERMAN, J. Nearly maximum flows in nearly linear time. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science* (2013), IEEE, pp. 263–269.
- [61] SHERMAN, J. Area-convexity, linf regularization, and undirected multicommodity flow. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing* (2017), pp. 452–460.
- [62] SIDFORD, A., AND TIAN, K. Coordinate methods for accelerating linf regression and faster approximate maximum

flow. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)* (2018), IEEE, pp. 922–933.

[63] SPIELMAN, D. A., AND TENG, S.-H. 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* (2004), pp. 81–90.

[64] VAIDYA, P. M. A new algorithm for minimizing convex functions over convex sets. *Mathematical programming* 73, 3 (1996), 291–341.

[65] VAN DEN BRAND, J. A deterministic linear program solver in current matrix multiplication time. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms* (2020), SIAM, pp. 259–278.

[66] VAN DEN BRAND, J., GAO, Y., JAMBULAPATI, A., LEE, Y. T., LIU, Y. P., PENG, R., AND SIDFORD, A. Faster maxflow via improved dynamic spectral vertex sparsifiers. *CoRR abs/2112.00722* (2021).

[67] VAN DEN BRAND, J., LEE, Y. T., LIU, Y. P., SARANURAK, T., SIDFORD, A., SONG, Z., AND WANG, D. Minimum cost flows, MDPs, and ℓ_1 -regression in nearly linear time for dense instances. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing* (2021), pp. 859–869.

[68] VAN DEN BRAND, J., AND ZHANG, D. Faster high accuracy multi-commodity flow from single-commodity techniques. *arXiv e-prints* (2023), arXiv–2304.

[69] WEIHE, K. Maximum (s, t) -flows in planar networks in $O(|v| \log |v|)$ time. *Journal of Computer and System Sciences* 55, 3 (1997), 454–475.

[70] WILLIAMS, V. V., XU, Y., XU, Z., AND ZHOU, R. New bounds for matrix multiplication: from alpha to omega. *arXiv preprint arXiv:2307.07970* (2023).

A Robust interior point method

For completeness, we include the robust interior point method from [15], developed in [17], which is a refinement of the methods in [9, 65]. Although there are many other robust interior point methods, we simply refer to this method as RIPM. Consider a linear program of the form

$$(A.1) \quad \min_{\mathbf{x} \in \mathcal{P}} \mathbf{c}^\top \mathbf{x} \quad \text{where} \quad \mathcal{P} = \{\mathbf{A}\mathbf{x} = \mathbf{b}, \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}\}$$

for some matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$. As with many other IPMs, RIPM follows the central path $\mathbf{x}(t)$ from an interior point ($t \gg 0$) to the optimal solution ($t = 0$):

$$\mathbf{x}(t) \stackrel{\text{def}}{=} \arg \min_{\mathbf{x} \in \mathcal{P}} \mathbf{c}^\top \mathbf{x} - t\phi(\mathbf{x}) \quad \text{where} \quad \phi(\mathbf{x}) \stackrel{\text{def}}{=} - \sum_i \log(\mathbf{x}_i - \mathbf{l}_i) - \sum_i \log(\mathbf{u}_i - \mathbf{x}_i),$$

where the term ϕ controls how close the solution \mathbf{x}_i can be to the constraints \mathbf{u}_i and \mathbf{l}_i . Following the central path exactly is expensive. Instead, RIPM maintains feasible primal and dual solution $(\mathbf{x}, \mathbf{s}) \in \mathcal{P} \times \mathcal{S}$, where \mathcal{S} is the dual space given by $\mathcal{S} = \{\mathbf{s} : \mathbf{A}^\top \mathbf{y} + \mathbf{s} = \mathbf{c} \text{ for some } \mathbf{y}\}$, and ensures $\mathbf{x}(t)$ is an approximate minimizer. Specifically, the optimality condition for $\mathbf{x}(t)$ is given by

$$(A.2) \quad \begin{aligned} \mu^t(\mathbf{x}, \mathbf{s}) &\stackrel{\text{def}}{=} \mathbf{s}/t + \nabla \phi(\mathbf{x}) = \mathbf{0} \\ (\mathbf{x}, \mathbf{s}) &\in \mathcal{P} \times \mathcal{S} \end{aligned}$$

where $\mu^t(\mathbf{x}, \mathbf{s})$ measures how close \mathbf{x} is to the minimizer $\mathbf{x}(t)$. RIPM maintains (\mathbf{x}, \mathbf{s}) such that

$$(A.3) \quad \|\gamma^t(\mathbf{x}, \mathbf{s})\|_\infty \leq \frac{1}{C \log m} \quad \text{where} \quad \gamma^t(\mathbf{x}, \mathbf{s})_i = \frac{\mu^t(\mathbf{x}, \mathbf{s})_i}{(\nabla^2 \phi(\mathbf{x}))_{ii}^{1/2}},$$

for some universal constant C . The normalization term $(\nabla^2 \phi)_{ii}^{1/2}$ makes the centrality measure $\|\gamma^t(\mathbf{x}, \mathbf{s})\|_\infty$ scale-invariant in \mathbf{l} and \mathbf{u} .

The key subroutine SOLVE takes as input a point close to the central path $(\mathbf{x}(t_{\text{start}}), \mathbf{s}(t_{\text{start}}))$, and outputs another point on the central path $(\mathbf{x}(t_{\text{end}}), \mathbf{s}(t_{\text{end}}))$. Each step of the subroutine decreases t by a multiplicative factor of $(1 - \frac{1}{\sqrt{m \log m}})$ and moves (\mathbf{x}, \mathbf{s}) within $\mathcal{P} \times \mathcal{S}$ such that $\mathbf{s}/t + \nabla \phi(\mathbf{x})$ is smaller for the current t . [17] proved that even if each step is computed approximately, IPM still outputs a point close to $(\mathbf{x}(t_{\text{end}}), \mathbf{s}(t_{\text{end}}))$ using $\tilde{O}(\sqrt{m} \log(t_{\text{end}}/t_{\text{start}}))$ steps. See Algorithm 2 for a simplified version.

Algorithm 2 Robust Primal-Dual Interior Point Method from [17]

- 1: **procedure** RIPM($\mathbf{A} \in \mathbb{R}^{n \times m}, \mathbf{b}, \mathbf{c}, \mathbf{l}, \mathbf{u}, \epsilon$)
- 2: Define $L \stackrel{\text{def}}{=} \|\mathbf{c}\|_2$ and $R \stackrel{\text{def}}{=} \|\mathbf{u} - \mathbf{l}\|_2$
- 3: Define $\phi_i(x) \stackrel{\text{def}}{=} -\log(\mathbf{u}_i - x) - \log(x - \mathbf{l}_i)$
- 4: Define $\mu_i^t(\mathbf{x}, \mathbf{s}) \stackrel{\text{def}}{=} \mathbf{s}_i/t + \nabla\phi_i(\mathbf{x}_i)$

- \triangleright Modify the linear program and obtain an initial (\mathbf{x}, \mathbf{s}) for modified linear program
- 5: Let $t = 2^{21}m^5 \cdot \frac{LR}{128} \cdot \frac{R}{r}$
- 6: Compute $\mathbf{x}_c = \arg \min_{\mathbf{l} \leq \mathbf{x} \leq \mathbf{u}} \mathbf{c}^\top \mathbf{x} + t\phi(\mathbf{x})$ and $\mathbf{x}_o = \arg \min_{\mathbf{A}\mathbf{x}=\mathbf{b}} \|\mathbf{x} - \mathbf{x}_c\|_2$
- 7: Let $\mathbf{x} = (\mathbf{x}_c, 3R + \mathbf{x}_o - \mathbf{x}_c, 3R)$ and $\mathbf{s} = (-t\nabla\phi(\mathbf{x}_c), \frac{t}{3R + \mathbf{x}_o - \mathbf{x}_c}, \frac{t}{3R})$
- 8: Let the new matrix $\mathbf{A}^{\text{new}} \stackrel{\text{def}}{=} [\mathbf{A}; \mathbf{A}; -\mathbf{A}]$, the new barrier

$$\phi_i^{\text{new}}(x) = \begin{cases} \phi_i(x) & \text{if } i \in [m], \\ -\log x & \text{else.} \end{cases}$$

- \triangleright Find an initial (\mathbf{x}, \mathbf{s}) for the original linear program
- 9: $((\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \mathbf{x}^{(3)}), (\mathbf{s}^{(1)}, \mathbf{s}^{(2)}, \mathbf{s}^{(3)})) \leftarrow \text{SOLVE}(\mathbf{A}^{\text{new}}, \phi^{\text{new}}, \mathbf{x}, \mathbf{s}, t, LR)$
- 10: $(\mathbf{x}, \mathbf{s}) \leftarrow (\mathbf{x}^{(1)} + \mathbf{x}^{(2)} - \mathbf{x}^{(3)}, \mathbf{s}^{(1)})$

\triangleright Optimize the original linear program

- 11: $(\mathbf{x}, \mathbf{s}) \leftarrow \text{SOLVE}(\mathbf{A}, \phi, \mathbf{x}, \mathbf{s}, LR, \frac{\epsilon}{4m})$
- 12: **return** \mathbf{x}
- 13: **end procedure**

14: **procedure** SOLVE($\mathbf{A}, \phi, \mathbf{x}, \mathbf{s}, t_{\text{start}}, t_{\text{end}}$)

- 15: Define $\alpha \stackrel{\text{def}}{=} \frac{1}{2^{20}\lambda}$ and $\lambda \stackrel{\text{def}}{=} 64 \log(256m^2)$
- 16: Let $t \leftarrow t_{\text{start}}, \bar{\mathbf{x}} \leftarrow \mathbf{x}, \bar{\mathbf{s}} \leftarrow \mathbf{s}, \bar{t} \leftarrow t$
- 17: **while** $t \geq t_{\text{end}}$ **do**
- 18: $t \leftarrow \max((1 - \frac{\alpha}{\sqrt{m}})t, t_{\text{end}})$
- 19: Update step size $h = -\alpha / \|\cosh(\lambda\gamma^{\bar{t}}(\bar{\mathbf{x}}, \bar{\mathbf{s}}))\|_2$ where γ is defined in Eq. (A.3)
- 20: Update diagonal weight matrix $\mathbf{W} = \nabla^2\phi(\bar{\mathbf{x}})^{-1}$
- 21: Update step direction \mathbf{v} where $\mathbf{v}_i = \sinh(\lambda\gamma^{\bar{t}}(\bar{\mathbf{x}}, \bar{\mathbf{s}})_i) \cdot \mu^{\bar{t}}(\bar{\mathbf{x}}, \bar{\mathbf{s}})_i$
- 22: Implicitly update \mathbf{x}, \mathbf{s} , with $\mathbf{P}_w \stackrel{\text{def}}{=} \mathbf{W}^{1/2} \mathbf{A}^\top (\mathbf{A} \mathbf{W} \mathbf{A}^\top)^{-1} \mathbf{A} \mathbf{W}^{1/2}$

$$\begin{aligned} \mathbf{x} &\leftarrow \mathbf{x} + h\mathbf{W}^{1/2}(\mathbf{v} - \mathbf{P}_w\mathbf{v}), \\ \mathbf{s} &\leftarrow \mathbf{s} + \bar{t}h\mathbf{W}^{-1/2}\mathbf{P}_w\mathbf{v} \end{aligned}$$

- 23: Explicitly update $\bar{\mathbf{x}}, \bar{\mathbf{s}}$ such that

$$\begin{aligned} \|\mathbf{W}^{-1/2}(\bar{\mathbf{x}} - \mathbf{x})\|_\infty &\leq \alpha, \\ \|\mathbf{W}^{1/2}(\bar{\mathbf{s}} - \mathbf{s})\|_\infty &\leq \bar{t}\alpha \end{aligned}$$

- 24: If $|\bar{t} - t| \geq \alpha\bar{t}$, update $\bar{t} \leftarrow t$
 - 25: **end while**
 - 26: **return** (\mathbf{x}, \mathbf{s})
 - 27: **end procedure**
-

RIPM calls SOLVE twice. The first call to SOLVE finds a feasible point by following the central path of the following modified linear program

$$\min_{\substack{\mathbf{A}(\mathbf{x}^{(1)} + \mathbf{x}^{(2)} - \mathbf{x}^{(3)}) = \mathbf{b} \\ \mathbf{l} \leq \mathbf{x}^{(1)} \leq \mathbf{u}, \mathbf{x}^{(2)} \geq \mathbf{0}, \mathbf{x}^{(3)} \geq \mathbf{0}}} \mathbf{c}^{(1)\top} \mathbf{x}^{(1)} + \mathbf{c}^{(2)\top} \mathbf{x}^{(2)} + \mathbf{c}^{(3)\top} \mathbf{x}^{(3)}$$

where $\mathbf{c}^{(1)} = \mathbf{c}$, and $\mathbf{c}^{(2)}, \mathbf{c}^{(3)}$ are some positive large vectors. The above modified linear program is chosen so that we know an explicit point on its central path, and any approximate minimizer to this new linear program gives an approximate central path point for the original problem. The second call to SOLVE finds an approximate solution by following the central path of the original linear program.

THEOREM A.1. (RIPM) *Consider the linear program*

$$\min_{\mathbf{Ax}=\mathbf{b}, \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}} \mathbf{c}^\top \mathbf{x}$$

with $\mathbf{A} \in \mathbb{R}^{n \times m}$. We are given a scalar $r > 0$ such that there exists some interior point \mathbf{x}_o satisfying $\mathbf{Ax}_o = \mathbf{b}$ and $\mathbf{l} + r \leq \mathbf{x}_o \leq \mathbf{u} - r$. Let $L = \|\mathbf{c}\|_2$ and $R = \|\mathbf{u} - \mathbf{l}\|_2$. For any $0 < \varepsilon \leq 1/2$, the algorithm RIPM (Algorithm 2) finds \mathbf{x} such that $\mathbf{Ax} = \mathbf{b}$, $\mathbf{l} \leq \mathbf{x} \leq \mathbf{u}$ and

$$\mathbf{c}^\top \mathbf{x} \leq \min_{\mathbf{Ax}=\mathbf{b}, \mathbf{l} \leq \mathbf{x} \leq \mathbf{u}} \mathbf{c}^\top \mathbf{x} + \varepsilon LR.$$

Furthermore, the algorithm has the following properties:

- Each call of SOLVE involves $O(\sqrt{m} \log m \log(\frac{mR}{\varepsilon}))$ -many steps, and \bar{t} is only updated $O(\log m \log(\frac{mR}{\varepsilon}))$ -many times.
- In each step of SOLVE, the coordinate i in \mathbf{w}, \mathbf{v} changes only if \bar{x}_i or \bar{s}_i changes.
- In each step of SOLVE, $h\|\mathbf{v}\|_2 = O(\frac{1}{\log m})$.
- Algorithm 2 to Algorithm 2 takes $O(K)$ time in total, where K is the total number of coordinate changes in $\bar{\mathbf{x}}, \bar{\mathbf{s}}$.

Proof. The number of steps follows from Theorem A.1 in [17], with the parameter $\mathbf{w}_i = \nu_i = 1$ for all i . The number of coordinate changes in \mathbf{W}, \mathbf{v} and the runtime of Algorithm 2 to Algorithm 2 follows directly from the formula of $\mu^t(\mathbf{x}, \mathbf{s})_i$ and $\gamma^t(\mathbf{x}, \mathbf{s})_i$. For the bound for $h\|\mathbf{v}\|_2$, it follows from

$$h\|\mathbf{v}\|_2 \leq \alpha \frac{\|\sinh(\lambda \gamma^{\bar{t}}(\bar{\mathbf{x}}, \bar{\mathbf{s}}))\|_2}{\|\cosh(\lambda \gamma^{\bar{t}}(\bar{\mathbf{x}}, \bar{\mathbf{s}}))\|_2} \leq \alpha = O\left(\frac{1}{\log m}\right).$$

□

B Maintaining the implicit representation

In this section, we give the general data structure MAINTAINREP, which implicitly maintains a vector \mathbf{x} throughout a call of SOLVE of Algorithm 2. We break up the representation into two parts, the first using the inverse tree operator, and the second using the tree operator.

First, we present some of the alternative decomposition properties of the tree operator.

DEFINITION B.1. (SUBTREE OPERATOR) *Let Δ be a tree operator on \mathcal{T} . Recall \mathcal{T}_H is the complete subtree of \mathcal{T} rooted at H . We define the subtree operator $\Delta^{(H)}$ at each node H to be*

$$(B.1) \quad \Delta^{(H)} \stackrel{\text{def}}{=} \sum_{\text{leaf } L \in \mathcal{T}_H} \Delta_{L \leftarrow H}.$$

COROLLARY B.1. *Based on the above definitions, we have*

$$(B.2) \quad \Delta = \sum_{H \in \mathcal{T}} \Delta^{(H)} \mathbf{I}_{F_H}.$$

Furthermore, if H has children H_1, H_2 , then

$$(B.3) \quad \Delta^{(H)} = \Delta^{(H_1)} \Delta_{H_1} + \Delta^{(H_2)} \Delta_{H_2}.$$

The output of Δ when restricted to $E(H)$ for a node $H \in \mathcal{T}$ can be written in two parts, which is useful for our data structures. The first part involves summing over all nodes in \mathcal{T}_H , ie. descendants of H and H itself, and the second part involves a sum over all ancestors of H .

LEMMA B.1. *At any node $H \in \mathcal{T}$, we have*

$$\mathbf{I}_{E(H)} \Delta = \sum_{D \in \mathcal{T}_H} \Delta^{(D)} \mathbf{I}_{F_D} + \Delta^{(H)} \sum_{\text{ancestor } A \text{ of } H} \Delta_{H \leftarrow A} \mathbf{I}_{F_A}.$$

Proof. We consider the terms in the sum for Δ that map into to $E(H)$, which is precisely the set of leaf nodes in the subtree rooted at H .

$$\mathbf{I}_{E(H)} \Delta = \sum_{\text{leaf } L \in \mathcal{T}_H} \sum_{A: L \in \mathcal{T}_A} \Delta_{L \leftarrow A} \mathbf{I}_{F_A}.$$

The right hand side involves a sum over the set $\{(L, A) : \text{leaf } L \in \mathcal{T}_H, L \in \mathcal{T}_A\}$. Observe that (L, A) is in this set if and only if A is a descendant of H , or $A = H$, or A is an ancestor of H . Hence, the summation can be written as

$$\sum_{\text{leaf } L \in \mathcal{T}_H} \sum_{\text{node } A \in \mathcal{T}_H} \Delta_{L \leftarrow A} \mathbf{I}_{F_A} + \sum_{\text{leaf } L \in \mathcal{T}_H} \sum_{\text{ancestor } A \text{ of } H} \Delta_{L \leftarrow A} \mathbf{I}_{F_A}.$$

The first term is precisely the first term in the lemma statement. For the second term, we can use the fact that A is an ancestor of H to expand $\Delta_{L \leftarrow A} = \Delta_{L \leftarrow H} \Delta_{H \leftarrow A}$. Then, the second term is

$$\begin{aligned} & \sum_{\text{leaf } L \in \mathcal{T}_H} \sum_{\text{ancestor } A \text{ of } H} \Delta_{L \leftarrow H} \Delta_{H \leftarrow A} \mathbf{I}_{F_A} \\ &= \sum_{\text{leaf } L \in \mathcal{T}_H} \Delta_{L \leftarrow H} \left(\sum_{\text{ancestor } A \text{ of } H} \Delta_{H \leftarrow A} \mathbf{I}_{F_A} \right) \\ &= \Delta^{(H)} \left(\sum_{\text{ancestor } A \text{ of } H} \Delta_{H \leftarrow A} \mathbf{I}_{F_A} \right), \end{aligned}$$

by definition of $\Delta^{(H)}$. \square

Now, we consider the cost of applying the inverse tree operator and the tree operator.

LEMMA B.2. *Let $\nabla : \mathbb{R}^E \mapsto \mathbb{R}^V$ be an inverse tree operator on \mathcal{T} with query complexity Q . Given $\mathbf{v} \in \mathbb{R}^E$, we can compute $\nabla \mathbf{v}$ as well as $\mathbf{y}_H \stackrel{\text{def}}{=} \sum_{\text{leaf } L \in \mathcal{T}_H} \nabla_{H \leftarrow L} \mathbf{v}$ for all $H \in \mathcal{T}$ in $O(Q(\eta K))$ time, where $K = \text{nnz}(\mathbf{v})$ and η is the height of \mathcal{T} .*

Proof. Recall the definition

$$\nabla \mathbf{v} \stackrel{\text{def}}{=} \sum_{\text{leaf } L} \left(\sum_{H: L \in \mathcal{T}_H} \mathbf{I}_{F_H} \nabla_{H \leftarrow L} \right) \mathbf{v}.$$

At a leaf node L , if we have $\mathbf{v}_e = 0$ for all $e \in E(L)$, then we can ignore the term for L in the outer sum. So we can reduce $\nabla \mathbf{v}$ to consist of at most K terms in the outer sum. We can further rearrange the order of applying the edge operators so that each edge operator is applied at most once, and this naturally gives the values for all non-zero \mathbf{y}_H 's. We bound the overall runtime loosely by $O(Q(\eta K))$. \square

Unlike the inverse tree operator, the tree operator is applied downwards along a tree, and therefore we do not have non-trivial bounds on total number of edge operators applied. Instead, we have a more general bound:

LEMMA B.3. Let $\Delta : \mathbb{R}^V \mapsto \mathbb{R}^E$ be a tree operator on \mathcal{T} with query complexity Q . Given $\mathbf{z} \in \mathbb{R}^V$, we can compute $\Delta \mathbf{v}$ in $O(Q(|E|))$ time.

Proof. We simply observe that we can compute $\Delta \mathbf{v}$ by applying each edge operator at most once. Since the leaf nodes partition the set E , we know in \mathcal{T} , there are $O(|E|)$ edge operators in total, so the overall time is at most $O(Q(|E|))$. \square

With the appropriate partial computations taking advantage of the decomposition of ∇ , we can maintain $\nabla \mathbf{v}$ efficiently for dynamic ∇ and \mathbf{v} . Specifically, we use the following property:

LEMMA B.4. Given a vector $\mathbf{v} \in \mathbb{R}^E$, let $\mathbf{y}_H \stackrel{\text{def}}{=} \sum_{\text{leaf } L \in \mathcal{T}_H} \nabla_{H \leftarrow L} \mathbf{v}$ for each $H \in \mathcal{T}$. If H has children H_1, H_2 , then

$$(B.4) \quad \mathbf{y}_H = \nabla_{H_1} \mathbf{y}_{H_1} + \nabla_{H_2} \mathbf{y}_{H_2}.$$

Furthermore,

$$(B.5) \quad \sum_{H \in \mathcal{T}} \mathbf{I}_{F_H} \mathbf{y}_H = \nabla \mathbf{v}.$$

LEMMA B.5. Let $\nabla : \mathbb{R}^E \mapsto \mathbb{R}^V$ be an inverse tree operator with query complexity Q . Let $\nabla^{(\text{new})}$ be ∇ with K updated edge operators. Suppose we know $\nabla \mathbf{v}$, and we know $\mathbf{y}_H \stackrel{\text{def}}{=} \sum_{\text{leaf } L \in \mathcal{T}_H} \nabla_{H \leftarrow L} \mathbf{v}$ at all nodes H , then we can compute $(\nabla^{(\text{new})} - \nabla) \mathbf{v}$ and the $\mathbf{y}_H^{(\text{new})}$'s in $O(Q(\eta K))$ time.

Proof. Observe that for a node $H \in \mathcal{T}$, if no edge operator in \mathcal{T}_H was updated, then \mathbf{y}_H remains the same. We use Eq. (B.4) to compute $\mathbf{y}_H^{(\text{new})}$ up the tree for the $O(\eta K)$ -many nodes that admit changes, and then Eq. (B.5) to extract the change $(\nabla^{(\text{new})} - \nabla) \mathbf{v}$. \square

Now we are ready for the complete data structure involving the inverse tree operator.

THEOREM B.1. (INVERSE TREE OPERATOR DATA STRUCTURE) Let $\mathbf{w} \in \mathbb{R}^m$ be the weights changing at every step of SOLVE, and let $\mathbf{v} \in \mathbb{R}^n$ be a dynamic vector. Suppose $\nabla : \mathbb{R}^m \mapsto \mathbb{R}^n$ is an inverse tree operator dependent on \mathbf{w} supported on \mathcal{T} with query complexity Q and update complexity U . Let η be the height of \mathcal{T} . Then the data structure INVERSE TREE OP (Algorithm 3) maintains $\mathbf{z}^{(k)} \stackrel{\text{def}}{=} \sum_{i=1}^k h^{(i)} \nabla^{(i)} \mathbf{v}^{(i)}$ so that at the end of each step k , the variables in the algorithm satisfy

- $\mathbf{z} = c\mathbf{z}^{(\text{step})} + \mathbf{z}^{(\text{sum})}$,
- $\mathbf{z}^{(\text{step})} = \nabla \mathbf{v}$, and
- $\mathbf{y}_H = \sum_{\text{leaf } L \in \mathcal{T}_H} \nabla_{H \leftarrow L} \mathbf{v}$ for all nodes H .

The data structure is initialized via INITIALIZE in $O(U(m) + Q(m))$ time. At step k , there is one call REWEIGHT($\delta_{\mathbf{w}}$) taking $O(U(K) + Q(\eta K))$ time, where $K = \text{nnz}(\delta_{\mathbf{w}})$, followed by one call of MOVE($h, \delta_{\mathbf{v}}$) taking $O(Q(\eta \cdot \text{nnz}(\delta_{\mathbf{v}})))$ time.

Proof. In the data structure, we always maintain $\mathbf{z}^{(\text{step})}$ and the \mathbf{y}_H 's together. Specifically, at every step, we update the \mathbf{y}_H 's up the tree using the recursive property Eq. (B.4) only at the necessary nodes, and from the \mathbf{y}_H 's, we get $\mathbf{z}^{(\text{step})} = \sum_H \mathbf{I}_{F_H} \mathbf{y}_H$.

Consider INITIALIZE. At the end of the function, the variables satisfy

$$\mathbf{z} \stackrel{\text{def}}{=} c\mathbf{z}^{(\text{step})} + \mathbf{z}^{(\text{sum})} = 0 \cdot \nabla \mathbf{v} + \mathbf{0} = \mathbf{0},$$

Algorithm 3 Dynamic data structure to maintain cumulative ∇v

1: **dynamic data structure** **INVERSE TREE OP**

2: **member:**

3: \mathcal{T} : tree supporting ∇ with edge operators on the edges

4: $\mathbf{w} \in \mathbb{R}^m$: dynamic weight vector

5: $\mathbf{v} \in \mathbb{R}^n$: dynamic vector

6: $c, \mathbf{z}^{(\text{step})}, \mathbf{z}^{(\text{sum})} \in \mathbb{R}^n$: coefficient, result vectors

7: $\mathbf{y}_H \in \mathbb{R}^n$ for each $H \in \mathcal{T}$: sparse partial computations

8:

9: **procedure** **INITIALIZE**($\mathcal{T}, \mathbf{w}^{(\text{init})}, \mathbf{v}^{(\text{init})}$)

10: $\mathbf{w} \leftarrow \mathbf{w}^{(\text{init})}, \mathbf{v} \leftarrow \mathbf{v}^{(\text{init})}, c \leftarrow 0, \mathbf{z}^{(\text{sum})} \leftarrow \mathbf{0}$

11: Initialize ∇ on \mathcal{T} based on \mathbf{w}

12: Compute ∇v and \mathbf{y}_H 's, set $\mathbf{z}^{(\text{step})} \leftarrow \nabla v$

13: **end procedure**

14:

15: **procedure** **REWEIGHT**(δ_w)

16: $\mathbf{w}^{(\text{new})} \leftarrow \mathbf{w} + \delta_w$

17: Let $\nabla^{(\text{new})}$ be the new tree operator using $\mathbf{w}^{(\text{new})}$

18: $\mathbf{z}' \leftarrow (\nabla^{(\text{new})} - \nabla)v$, and update \mathbf{y}_H 's ▷ Lemma B.5

19: $\mathbf{z}^{(\text{step})} \leftarrow \mathbf{z}^{(\text{step})} + \mathbf{z}'$

20: $\mathbf{z}^{(\text{sum})} \leftarrow \mathbf{z}^{(\text{sum})} - c \cdot \mathbf{z}'$

21: $\mathbf{w} \leftarrow \mathbf{w}^{(\text{new})}, \nabla \leftarrow \nabla^{(\text{new})}$

22: **end procedure**

23:

24: **procedure** **MOVE**(h, δ_v)

25: Compute $\mathbf{z}' \stackrel{\text{def}}{=} \nabla \delta_v$ and the $\mathbf{y}'_H \stackrel{\text{def}}{=} \sum_{\text{leaf } L \in \mathcal{T}_H} \nabla_{H \leftarrow L} \delta_v$ for each node H ▷ Lemma B.2

26: $\mathbf{z}^{(\text{step})} \leftarrow \mathbf{z}^{(\text{step})} + \mathbf{z}'$, and $\mathbf{y}_H \leftarrow \mathbf{y}_H + \mathbf{y}'_H$ for each node H

27: $\mathbf{z}^{(\text{sum})} \leftarrow \mathbf{z}^{(\text{sum})} - c \mathbf{z}'$

28: $c \leftarrow c + h$

29: $\mathbf{v} \leftarrow \mathbf{v} + \delta_v$

30: **end procedure**

and $\mathbf{z}^{(\text{step})} = \nabla \mathbf{v}$, as required.

Let us consider REWEIGHT. Let the superscript (new) denote the value of an algorithm variable at the end of the function, and let no superscript denote the value at the start.

$$\begin{aligned} \mathbf{z}^{(\text{new})} &= c^{(\text{new})} \mathbf{z}^{(\text{step})^{(\text{new})}} + \mathbf{z}^{(\text{sum})^{(\text{new})}} \\ &= c(\mathbf{z}^{(\text{step})} + \mathbf{z}') + \mathbf{z}^{(\text{sum})} - c\mathbf{z}' \\ &= c\mathbf{z}^{(\text{step})} + \mathbf{z}^{(\text{sum})}, \\ \text{and } \mathbf{z}^{(\text{step})^{(\text{new})}} &= \mathbf{z}^{(\text{step})} + (\nabla^{(\text{new})} - \nabla)\mathbf{v} \\ &= \nabla^{(\text{new})}\mathbf{v}, \end{aligned}$$

as required. Similarly, let us consider MOVE:

$$\begin{aligned} \mathbf{z}^{(\text{new})} &= c^{(\text{new})} \mathbf{z}^{(\text{step})^{(\text{new})}} + \mathbf{z}^{(\text{sum})^{(\text{new})}} \\ &= (c + h)(\mathbf{z}^{(\text{step})} + \mathbf{z}') + \mathbf{z}^{(\text{sum})} - c\mathbf{z}' \\ &= c\mathbf{z}^{(\text{step})} + \mathbf{z}^{(\text{sum})} + h\mathbf{z}^{(\text{step})}, \\ \text{and } \mathbf{z}^{(\text{step})^{(\text{new})}} &= \mathbf{z}^{(\text{step})} + \nabla(\mathbf{v}^{(\text{new})} - \mathbf{v}) \\ &= \nabla\mathbf{v} + \nabla(\mathbf{v}^{(\text{new})} - \mathbf{v}) \\ &= \nabla\mathbf{v}^{(\text{new})}, \end{aligned}$$

which is exactly the update we want to make to \mathbf{z} , and the invariant we want to maintain.

The runtimes follow directly from Lemmas B.2 and B.5. \square

Next, we present the tree operator data structure, which is significantly more involved compared to the inverse tree operator. Applying the tree operator involves going down the tree to the leaves, which is too costly to do at every step. To circumvent the issue, we use lazy computations.

THEOREM B.2. (TREE OPERATOR DATA STRUCTURE) *Let $\mathbf{w} \in \mathbb{R}^m$ be the weights changing at every step of SOLVE. Suppose $\Delta : \mathbb{R}^n \mapsto \mathbb{R}^m$ is a tree operator dependent on \mathbf{w} supported on \mathcal{T} with query complexity Q and update complexity U . Let $\mathbf{z} \in \mathbb{R}^n$ be the vector maintained by Algorithm 3, so that at the end of step k , $\mathbf{z} = \sum_{i=1}^k h^{(i)} \nabla^{(i)} \mathbf{v}^{(i)}$. Then the data structure TREEOP (Algorithm 4) implicitly maintains \mathbf{x} so that at the end of step k ,*

$$\mathbf{x}^{(k)} = \mathbf{x}^{(\text{init})} + \sum_{i=1}^k \Delta^{(i)} \nabla^{(i)} \mathbf{v}^{(i)}.$$

The data structure is initialized via INITIALIZE in $O(U(m))$ time. At step k , there is one call to REWEIGHT($\delta_{\mathbf{w}}$) taking $O(U(K) + Q(\eta K))$ time, where $K = \text{nnz}(\delta_{\mathbf{w}})$, followed by one call to MOVE($\delta_{\mathbf{z}}$) taking $O(\text{nnz}(\delta_{\mathbf{z}}))$ time. At the end of SOLVE, \mathbf{x} is returned via EXACT in $O(Q(m))$ time.

Proof. We will show that the data structure maintains the implicit representation via the identity

$$(B.6) \quad \mathbf{x} = c\Delta\mathbf{z} + \sum_{H \in \mathcal{T}} \Delta^{(H)} \mathbf{u}_H,$$

where the RHS expression refers to the state of the variables at the end of step k during the algorithm.

At a high level, the variables Δ and \mathbf{z} in the data structure at step k represent the latest $\Delta^{(k)}$ and $\mathbf{z}^{(k)}$. We need to introduce additional vectors \mathbf{u}_H at every node H which intuitively stores lazy computations at node H , in order to take advantage of the tree structure of Δ . The function PUSHDOWN performs the accumulated computation at H , and moves the result to its children nodes to be computed lazily at a later point. The next claim describes this process formally.

Algorithm 4 Dynamic data structure to maintain cumulative Δz

```

1: dynamic data structure TREEOP
2: member:
3:    $\mathcal{T}$ : tree supporting  $\Delta$ 
4:    $\mathbf{w} \in \mathbb{R}^m$ : dynamic weight vector
5:    $\mathbf{z} \in \mathbb{R}^n$ : dynamic vector
6:    $\mathbf{u}_H$  for each  $H \in \mathcal{T}$ : lazy pushdown computation vectors
7:
8: procedure INITIALIZE( $\mathcal{T}, \mathbf{w}^{(\text{init})}, \mathbf{z}^{(\text{init})}, \mathbf{x}^{(\text{init})}$ )
9:    $\mathbf{w} \leftarrow \mathbf{w}^{(\text{init})}, \mathbf{z} \leftarrow \mathbf{z}^{(\text{init})}$ 
10:  Initialize  $\Delta$  on  $\mathcal{T}$  based on  $\mathbf{w}$ 
11:   $\mathbf{u}_H \leftarrow \mathbf{0}$  for each non-leaf  $H \in \mathcal{T}$ 
12:   $\mathbf{u}_H \leftarrow \mathbf{x}^{(\text{init})}|_{E(H)}$  for each leaf  $H \in \mathcal{T}$ 
13: end procedure
14:
15: procedure REWEIGHT( $\delta_{\mathbf{w}}$ )
16:    $\mathbf{w} \leftarrow \mathbf{w} + \delta_{\mathbf{w}}$ 
17:   Let  $\Delta^{(\text{new})}$  be the new tree operator wrt the new weights
18:   Let  $\mathcal{H}$  be all nodes  $H$  where  $\Delta_H$  changed
19:   for  $H \in \mathcal{P}_{\mathcal{T}}(\mathcal{H})$  going down the tree level by level do
20:     PUSHDOWN( $H$ )
21:   end for
22:   for  $H \in \mathcal{P}_{\mathcal{T}}(\mathcal{H})$  going down the tree level by level do
23:      $\mathbf{u}_H \leftarrow c\mathbf{z}|_{F_H}$ 
24:     PUSHDOWN( $H$ )
25:   end for
26:    $\Delta \leftarrow \Delta^{(\text{new})}$ 
27:   for  $H \in \mathcal{P}_{\mathcal{T}}(\mathcal{H})$  going down the tree level by level do
28:      $\mathbf{u}_H \leftarrow -c\mathbf{z}|_{F_H}$ 
29:     PUSHDOWN( $H$ )
30:   end for
31: end procedure
32:
33: procedure MOVE( $\delta_{\mathbf{z}}$ )
34:    $\mathbf{z} \leftarrow \mathbf{z} + \delta_{\mathbf{z}}$ 
35: end procedure
36:
37: procedure EXACT
38:   for  $H \in \mathcal{T}$  going down the tree level by level do
39:      $\mathbf{u}_H \leftarrow \mathbf{u}_H + \mathbf{z}|_{F_H}$ 
40:     PUSHDOWN( $H$ )
41:   end for
42:   return  $\mathbf{x}$  defined by  $\mathbf{x}|_{E(H)} \stackrel{\text{def}}{=} \mathbf{u}_H$  at each leaf  $H \in \mathcal{T}$ 
43: end procedure

```

Algorithm 5 Dynamic data structure to maintain cumulative $\Delta \mathbf{z}$, con't

```

1: dynamic data structure TREEOP
2: procedure PUSHDOWN( $H \in \mathcal{T}$ )
3:   for each child  $D$  of  $H$  do
4:      $\mathbf{u}_D \leftarrow \mathbf{u}_D + \Delta_D \mathbf{u}_H$ 
5:   end for
6:    $\mathbf{u}_H \leftarrow \mathbf{0}$ 
7: end procedure

```

CLAIM 3. Let $H \in \mathcal{T}$ be a non-leaf node. PUSHDOWN(H) does not change the value of the implicit representation in Eq. (B.6). Also, at the end of the procedure, $\mathbf{u}_H = \mathbf{0}$.

Proof. For any variable in the algorithm, we add the superscript ^(new) to mean its state at the end of PUSHDOWN; if there is no superscript, then it refers to the state at the start.

We show the claim for when H has two children H_1, H_2 . Note that Δ and \mathbf{z} are not touched by PUSHDOWN, so we may ignore the term $c\Delta \mathbf{z}$ in Eq. (B.6). Then,

$$\begin{aligned}
& \sum_{H' \in \mathcal{T}} \Delta^{(H')} \mathbf{u}_{H'}^{(\text{new})} \\
(\text{expand terms}) \quad &= \Delta^{(H)} \mathbf{u}_H^{(\text{new})} + \sum_{i=1,2} \Delta^{(H_i)} \mathbf{u}_{H_i}^{(\text{new})} + \sum_{H' \in \mathcal{T}, H' \neq H, H_1, H_2} \Delta^{(H')} \mathbf{u}_{H'} \\
(\text{substitute values}) \quad &= \sum_{i=1,2} \Delta^{(H_i)} (\mathbf{u}_{H_i} + \Delta_{H_i} \mathbf{u}_H) + \sum_{H' \in \mathcal{T}, H' \neq H, H_1, H_2} \Delta^{(H')} \mathbf{u}_{H'} \\
&= \sum_{i=1,2} \Delta^{(H_i)} \Delta_{H_i} \mathbf{u}_H + \sum_{H' \in \mathcal{T}, H' \neq H} \Delta^{(H')} \mathbf{u}_{H'} \\
(\text{by Eq. (B.3)}) \quad &= \Delta^{(H)} \mathbf{u}_H + \sum_{H' \in \mathcal{T}, H' \neq H} \Delta^{(H')} \mathbf{u}_{H'} \\
&= \sum_{H' \in \mathcal{T}} \Delta^{(H')} \mathbf{u}_{H'},
\end{aligned}$$

so the implicit representation of \mathbf{x} has not changed in value. \square

This claim can be generalized from $H \in \mathcal{T}$ to $\mathcal{H} \subseteq \mathcal{T}$; we omit the full details. Next, we show that the implicit representation of \mathbf{x} by Eq. (B.6) is correctly maintained after reweight.

CLAIM 4. After the k -th call REWEIGHT, the value of \mathbf{x} is unchanged, while the value of Δ is updated to $\Delta^{(k)}$ which is a function of $\mathbf{w}^{(k)}$.

Proof. We begin by observing that if $H \notin \mathcal{P}_{\mathcal{T}}(\mathcal{H})$, then $\Delta^{(H)(\text{new})} = \Delta^{(H)}$ by definition, as there are no edges in \mathcal{T}_H with updated operators.

At a high level, we traverse the subtree $\mathcal{P}_{\mathcal{T}}(\mathcal{H})$ three rounds and perform PUSHDOWN at every node. During the first round, we simply push down the current \mathbf{u}_H values at each node H . By Claim 3, we know this does not change the value of the implicit representation.

During the second round, we first initialize $\mathbf{u}_H \leftarrow c\mathbf{z}|_{F_H}$ at each node $H \in \mathcal{P}_{\mathcal{T}}(\mathcal{H})$, and then perform PUSHDOWN. Since PUSHDOWN does not affect the value of the implicit representation, we can use the initial change in \mathbf{u}_H to determine the overall change in the implicit representation. Crucially, note that we perform PUSHDOWN using the old tree operator. So, the change in value of the implicit representation is given by

$$+c \sum_{H \in \mathcal{P}_{\mathcal{T}}(\mathcal{H})} \Delta^{(H)} \mathbf{z}|_{F_H}.$$

After the second round of PUSHDOWN, we update the tree operator Δ to $\Delta^{(\text{new})}$. Note that $\Delta^{(H)}$ changes if and only if $H \in \mathcal{P}_{\mathcal{T}}(\mathcal{H})$, and in this case, $\mathbf{u}_H = \mathbf{0}$. So, updating the tree operator at this point induces a change

in the value of the implicit representation of

$$c\Delta^{(\text{new})}\mathbf{z} - c\Delta\mathbf{z} = c \sum_{H \in \mathcal{T}} \left(\Delta^{(H)(\text{new})} - \Delta^{(H)} \right) \mathbf{z}|_{F_H} = c \sum_{H \in \mathcal{P}_{\mathcal{T}}(\mathcal{H})} \left(\Delta^{(H)(\text{new})} - \Delta^{(H)} \right) \mathbf{z}|_{F_H}.$$

During the third round, we initialize $\mathbf{u}_H \leftarrow -c\mathbf{z}|_{F_H}$ at each node $H \in \mathcal{P}_{\mathcal{T}}(\mathcal{H})$ and perform PUSHDOWN. Similar to the first round, the change to the value of the implicit representation induced by this round is given by

$$-c \sum_{H \in \mathcal{P}_{\mathcal{T}}(\mathcal{H})} \Delta^{(H)(\text{new})} \mathbf{z}|_{F_H}.$$

The sum of the changes from each of the three rounds is exactly $\mathbf{0}$, so we conclude the value of the implicit representation did not change. \square

Finally, we consider the other functions.

For INITIALIZE, we see that by substituting the values assigned during INITIALIZE and applying the definition from Eq. (B.2), we have

$$c\Delta\mathbf{z} + \sum_{H \in \mathcal{T}} \Delta^{(H)} \mathbf{u}_H = \mathbf{x}^{(\text{init})} + \Delta\mathbf{z},$$

where Δ is the initial $\Delta^{(\text{init})}$ and \mathbf{z} is the initial $\mathbf{z}^{(\text{init})}$, which is exactly how we want to initialize \mathbf{x} .

For MOVE, we see the value of \mathbf{x} is incremented by $\Delta^{(k)}(\mathbf{z}^{(k)} - \mathbf{z}^{(k-1)})$ after the step k . By definition of \mathbf{z} , we know $\mathbf{z}^{(k)} - \mathbf{z}^{(k-1)} = h^{(k)} \nabla^{(k)} \mathbf{v}^{(k)}$, so we conclude MOVE correctly makes the update $h^{(k)} \Delta^{(k)} \nabla^{(k)} \mathbf{v}^{(k)}$.

For EXACT, we calculate the value of \mathbf{x} explicitly by performing the computation $\sum_{H \in \mathcal{T}} \Delta^{(H)}(\mathbf{u}_H + \mathbf{z}|_{F_H})$ using a sequence of PUSHDOWN's down the tree. The final answer \mathbf{x} is stored in parts in the \mathbf{u}_H 's along the leaf nodes.

Note that by definition of the query complexity of Δ , PUSHDOWN uses $O(Q(1))$ time. The remaining runtimes are straightforward. \square

Finally, we combine Algorithm 3 and Algorithm 4 to get the overall data structure MAINTAINREP for maintaining \mathbf{x} throughout SOLVE as given by Eq. (3.2). We omit the pseudocode implementation.

Proof. [Proof of Theorem 3.2] We use one copy of INVERSE TREEOP, which maintains $\mathbf{z} \stackrel{\text{def}}{=} c\mathbf{z}^{(\text{step})} + \mathbf{z}^{(\text{sum})}$. We want to use TREEOP to maintain \mathbf{z} which is given in two terms by INVERSE TREEOP. To resolve this, we can simply use two copies of the data structure and track the two terms in \mathbf{z} separately; then we correctly maintain \mathbf{x} . During SOLVE, at step k , we first call REWEIGHT and MOVE in INVERSE TREEOP, followed by REWEIGHT and MOVE in each copy of TREEOP. The runtimes follow in a straightforward manner. \square

E Maintaining vector approximation

We include this section for completeness; all techniques are from [15].

Recall at every step of the IPM, we want to maintain approximate vectors $\bar{\mathbf{x}}, \bar{\mathbf{s}}$ so that

$$\left\| \mathbf{W}^{-1/2}(\bar{\mathbf{x}} - \mathbf{x}) \right\|_{\infty} \leq \delta \quad \text{and} \quad \left\| \mathbf{W}^{1/2}(\bar{\mathbf{s}} - \mathbf{s}) \right\|_{\infty} \leq \delta'$$

for some error tolerances δ and δ' .

In the previous section, we showed how to use MAINTAINREP to maintain \mathbf{x} implicitly throughout SOLVE in the IPM. In this section, we give a data structure to efficiently maintain an approximate vector $\bar{\mathbf{x}}$ to the \mathbf{x} from MAINTAINREP, so that at every step,

$$\left\| \mathbf{W}^{-1/2}(\bar{\mathbf{x}} - \mathbf{x}) \right\|_{\infty} \leq \delta.$$

In the remainder of this section, we crucially assume that \mathbf{w} is a function of $\bar{\mathbf{x}}$ coordinate-wise, which is indeed satisfied by the RIPM framework.

REMARK E.1. *In our problem setting, we do not have full access to the exact vector \mathbf{x} . The algorithms in the next two subsections however will refer to \mathbf{x} for readability and modularity. We observe that access to \mathbf{x} is limited to two types: accessing the JL-sketches of specific subvectors, and accessing exact coordinates and other specific subvectors of sufficiently small size.*

In Section E.1, we reduce the problem of maintaining $\bar{\mathbf{x}}$ to detecting coordinates of \mathbf{x} with large changes. In Section E.2, we detect coordinates of \mathbf{x} with large changes using a sampling technique on a constant-degree tree, where Johnson-Lindenstrauss sketches of subvectors of \mathbf{x} are maintained at each node of the tree. In Section E.3, we show how to compute and maintain the necessary collection of JL-sketches on the operator tree \mathcal{T} ; in particular, we do this efficiently with only an implicit representation of \mathbf{x} . Finally, we put the three parts together to prove Theorem 3.3.

For notational simplicity, we use $\mathbf{D} \stackrel{\text{def}}{=} \mathbf{W}^{-1/2}$. Recall we use the superscript (k) to denote the variable at the end of step k ; that is, $\mathbf{D}^{(k)}$ and $\mathbf{x}^{(k)}$ are the values of \mathbf{D} and \mathbf{x} at the end of step k . Step 0 is the initialization step.

E.1 Reduction to change detection In this section, we show that in order to maintain an approximation $\bar{\mathbf{x}}$ to some vector \mathbf{x} , it suffices to detect coordinates of \mathbf{x} that change a lot.

We make use of dyadic intervals. At step k , for each ℓ such that $k \equiv 0 \pmod{2^\ell}$, we find the index set $I_\ell^{(k)}$ that contains all coordinates i of \mathbf{x} such that $\mathbf{x}_i^{(k)}$ changed significantly compared to $\mathbf{x}_i^{(k-2^\ell)}$, that is, compared to 2^ℓ steps ago. Formally:

DEFINITION E.1. *At step k of the IPM, for each ℓ such that $k \equiv 0 \pmod{2^\ell}$, we define*

$$I_\ell^{(k)} \stackrel{\text{def}}{=} \{i \in [n] : \mathbf{D}_{ii}^{(k)} \cdot |\mathbf{x}_i^{(k)} - \mathbf{x}_i^{(k-2^\ell)}| \geq \frac{\delta}{2 \lceil \log m \rceil},$$

and $\bar{\mathbf{x}}_i$ has not been updated after the $(k - 2^\ell)$ -th step $\}$.

We show how to find the sets $I_\ell^{(k)}$ with high probability in the next section. Assuming the correct implementation, we have the following data structure for maintaining the desired approximation $\bar{\mathbf{x}}$:

LEMMA E.1. (APPROXIMATE VECTOR MAINTENANCE) *Suppose FINDLARGECOORDINATES(ℓ) is a procedure in ABSTRACTMAINTAINAPPROX that correctly computes the set $I_\ell^{(k)}$ at the k -th step. Then the deterministic data structure ABSTRACTMAINTAINAPPROX in Algorithm 6 maintains an approximation $\bar{\mathbf{x}}$ of \mathbf{x} with the following procedures:*

- INITIALIZE($\mathcal{T}, \mathbf{x}, \mathbf{D}, \rho > 0, \delta > 0$): *Initialize the data structure at step 0 with tree \mathcal{T} , initial vector \mathbf{x} , initial diagonal scaling matrix \mathbf{D} , target additive approximation error δ , and success probability $1 - \rho$.*
- UPDATE($\mathbf{x}^{(\text{new})}, \mathbf{D}^{(\text{new})}$): *Increment the step counter and update vector \mathbf{x} and diagonal scaling matrix \mathbf{D} .*
- APPROXIMATE: *Output a vector $\bar{\mathbf{x}}$ such that $\|\mathbf{D}(\mathbf{x} - \bar{\mathbf{x}})\|_\infty \leq \delta$ for the latest \mathbf{x} and \mathbf{D} with probability at least $1 - \rho$.*

At step k , the procedure UPDATE is called, followed by APPROXIMATE. Suppose $\|\mathbf{D}^{(k)}(\mathbf{x}^{(k)} - \mathbf{x}^{(k-1)})\|_2 \leq \beta$ for all steps k , and \mathbf{D} is a function of $\bar{\mathbf{x}}$ coordinate-wise. Then, for each $\ell \geq 0$, the data structure updates $O(2^{2\ell}(\beta/\delta)^2 \log^2 m)$ coordinates of $\bar{\mathbf{x}}$ every 2^ℓ steps.

Proof. [Proof of Lemma E.1] The failure case arises from FINDLARGECOORDINATES. Assuming FINDLARGECOORDINATES returns the correct set of coordinates, we prove the correctness of APPROXIMATE.

Fix some coordinate $i \in [m]$ and fix some step k . Suppose the latest update to $\bar{\mathbf{x}}_i$ is $\bar{\mathbf{x}}_i \leftarrow \mathbf{x}_i^{(k')}$ at step k' . So $\mathbf{D}_{ii}^{(d)}$ is the same for all $k \geq d > k'$, and i is not in the set $I_\ell^{(d)}$ returned by FINDLARGECOORDINATES for all $k \geq d > k'$. Since we set $\bar{\mathbf{x}} \leftarrow \mathbf{x}$ every $2^{\lceil \log m \rceil}$ steps by Algorithm 6, we know $k - 2^{\lceil \log m \rceil} \leq k' < k$. Using dyadic intervals, we can define a sequence k_0, k_1, \dots, k_s with $s \leq 2 \lceil \log m \rceil$, where $k' = k_0 < k_1 < k_2 < \dots < k_s = k$, each $k_{j+1} - k_j$ is a power of 2, and $(k_{j+1} - k_j) \mid k_{j+1}$. Hence, we have

$$\mathbf{x}_i^{(k)} - \bar{\mathbf{x}}_i^{(k)} = \mathbf{x}_i^{(k_s)} - \bar{\mathbf{x}}_i^{(k_0)} = \mathbf{x}_i^{(k_s)} - \mathbf{x}_i^{(k_0)} = \sum_{j=0}^{s-1} \left(\mathbf{x}_i^{(k_{j+1})} - \mathbf{x}_i^{(k_j)} \right).$$

We know that $\mathbf{D}_{ii}^{(d)}$ is the same for all $k \geq d > k'$. By the guarantees of FINDLARGECOORDINATES, we have

$$\mathbf{D}_{ii}^{(k)} \cdot |\mathbf{x}_i^{(k_{j+1})} - \mathbf{x}_i^{(k_j)}| = \mathbf{D}_{ii}^{(k_{j+1})} \cdot |\mathbf{x}_i^{(k_{j+1})} - \mathbf{x}_i^{(k_j)}| \leq \frac{\delta}{2 \lceil \log m \rceil}$$

Algorithm 6 Data structure ABSTRACTMAINTAINAPPROX, Part 1

```

1: data structure ABSTRACTMAINTAINAPPROX
2: private : member
3:    $\mathcal{T}$ : constant-degree rooted tree with height  $\eta$  and  $m$  leaves ▷ leaf  $i$  corresponds to  $\mathbf{x}_i$ 
4:    $w \stackrel{\text{def}}{=} \Theta(\eta^2 \log(\frac{m}{\rho}))$ : sketch dimension
5:    $\Phi \sim \mathbf{N}(0, \frac{1}{w})^{w \times m}$ : JL-sketch matrix
6:    $\delta > 0$ : additive approximation error
7:    $k$ : current step
8:    $\bar{\mathbf{x}} \in \mathbb{R}^m$ : current valid approximate vector
9:    $\{\mathbf{x}^{(j)} \in \mathbb{R}^m\}_{j=0}^k$ : list of previous inputs
10:   $\{\mathbf{D}^{(j)} \in \mathbb{R}^{m \times m}\}_{j=0}^k$ : list of previous diagonal scaling matrices
11:
12: procedure INITIALIZE( $\mathcal{T}, \mathbf{x} \in \mathbb{R}^m, \mathbf{D} \in \mathbb{R}_{>0}^{m \times m}, \rho > 0, \delta > 0$ )
13:    $\mathcal{T} \leftarrow \mathcal{T}, \delta \leftarrow \delta, k \leftarrow 0$ 
14:    $\bar{\mathbf{x}} \leftarrow \mathbf{x}, \mathbf{x}^{(0)} \leftarrow \mathbf{x}, \mathbf{D}^{(0)} \leftarrow \mathbf{D}$ 
15:   sample  $\Phi \sim \mathbf{N}(0, \frac{1}{w})^{w \times m}$ 
16: end procedure
17:
18: procedure UPDATE( $\mathbf{x}^{(\text{new})} \in \mathbb{R}^m, \mathbf{D}^{(\text{new})} \in \mathbb{R}_{>0}^{m \times m}$ )
19:    $k \leftarrow k + 1, \mathbf{x}^{(k)} \leftarrow \mathbf{x}^{(\text{new})}, \mathbf{D}^{(k)} \leftarrow \mathbf{D}^{(\text{new})}$ 
20: end procedure
21:
22: procedure APPROXIMATE
23:    $I \leftarrow \emptyset$ 
24:   for all  $0 \leq \ell < \lceil \log m \rceil$  such that  $k \equiv 0 \pmod{2^\ell}$  do
25:      $I_\ell^{(k)} \leftarrow \text{FINDLARGECOORDINATES}(\ell)$ 
26:      $I \leftarrow I \cup I_\ell^{(k)}$ 
27:   end for
28:   if  $k \equiv 0 \pmod{2^{\lceil \log m \rceil}}$  then
29:      $I \leftarrow [m]$  ▷ Update  $\bar{\mathbf{x}}$  in full every  $2^{\lceil \log m \rceil}$  steps
30:   end if
31:    $\bar{\mathbf{x}}_i \leftarrow \mathbf{x}_i^{(k)}$  for all  $i \in I$ 
32:   return  $\bar{\mathbf{x}}$ 
33: end procedure

```

for all $0 \leq j < s$. Summing over all $j = 0, 1, \dots, s - 1$ gives

$$\mathbf{D}_{ii}^{(k)} \cdot |\mathbf{x}_i^{(k)} - \bar{\mathbf{x}}_i^{(k)}| \leq \delta.$$

Hence, we have $\|\mathbf{D}(\mathbf{x} - \bar{\mathbf{x}})\|_\infty \leq \delta$.

Fix some ℓ with $k \equiv 0 \pmod{2^\ell}$. We bound the number of coordinates in $I_\ell^{(k)}$. For any $i \in I_\ell^{(k)}$, we know $\mathbf{D}_{ii}^{(j)} = \mathbf{D}_{ii}^{(k)}$ for all $j > k - 2^\ell$ because $\bar{\mathbf{x}}_i$ did not change in the meanwhile. By definition of $I_\ell^{(k)}$, we have

$$\mathbf{D}_{ii}^{(k)} \cdot \sum_{j=k-2^\ell}^{k-1} |\mathbf{x}_i^{(j+1)} - \mathbf{x}_i^{(j)}| \geq \mathbf{D}_{ii}^{(k)} \cdot |\mathbf{x}_i^{(k)} - \mathbf{x}_i^{(k-2^\ell)}| \geq \frac{\delta}{2 \lceil \log m \rceil}.$$

Using $\mathbf{D}_{ii}^{(j)} = \mathbf{D}_{ii}^{(k)}$ for all $j > k - 2^\ell$ again, the above inequality yields

$$\begin{aligned} \frac{\delta}{2 \lceil \log m \rceil} &\leq \sum_{j=k-2^\ell}^{k-1} \mathbf{D}_{ii}^{(j+1)} |\mathbf{x}_i^{(j+1)} - \mathbf{x}_i^{(j)}| \\ &\leq \sqrt{2^\ell \sum_{j=k-2^\ell}^{k-1} \mathbf{D}_{ii}^{(j+1)2} |\mathbf{x}_i^{(j+1)} - \mathbf{x}_i^{(j)}|^2}. \end{aligned}$$

(by Cauchy-Schwarz)

Squaring and summing over all $i \in I_\ell^{(k)}$ gives

$$\begin{aligned} \Omega \left(\frac{2^{-\ell} \delta^2}{\log^2 m} \right) |I_\ell^{(k)}| &\leq \sum_{i \in I_\ell^{(k)}} \sum_{j=k-2^\ell}^{k-1} \mathbf{D}_{ii}^{(j+1)2} |\mathbf{x}_i^{(j+1)} - \mathbf{x}_i^{(j)}|^2 \\ &\leq \sum_{i=1}^m \sum_{j=k-2^\ell}^{k-1} \mathbf{D}_{ii}^{(j+1)2} |\mathbf{x}_i^{(j+1)} - \mathbf{x}_i^{(j)}|^2 \\ &\leq 2^\ell \beta^2, \end{aligned}$$

where we use $\|\mathbf{D}^{(j+1)}(\mathbf{x}^{(j+1)} - \mathbf{x}^{(j)})\|_2 \leq \beta$ at the end. Hence, we have

$$|I_\ell^{(k)}| = O(2^{2\ell} (\beta/\delta)^2 \log^2 m).$$

In other words, for each $\ell \geq 0$, we update $|I_\ell^{(k)}|$ -many coordinates of $\bar{\mathbf{x}}$ at step k when $k \equiv 0 \pmod{2^\ell}$. So we conclude that for each $\ell \geq 0$, we update $O(2^{2\ell} (\beta/\delta)^2 \log^2 m)$ -many coordinates of $\bar{\mathbf{x}}$ every 2^ℓ steps. \square

E.2 From change detection to sketch maintenance Now we discuss the implementation of FINDLARGE-COORDINATES(ℓ) to find the set $I_\ell^{(k)}$ in Algorithm 6 of Algorithm 6. We accomplish this by repeatedly sampling a coordinate i with probability proportional to $\left(\mathbf{D}_{ii}^{(k)} (\mathbf{x}_i^{(k)} - \mathbf{x}_i^{(k-2^\ell)})\right)^2$, among all coordinates i where $\bar{\mathbf{x}}_i$ has not been updated since 2^ℓ steps ago. With high probability, we can find all indices in $I_\ell^{(k)}$ in this way efficiently. To implement the sampling procedure, we make use of a data structure based on segment trees [10] along with sketching based on the Johnson-Lindenstrauss lemma.

Formally, we define the vector $\mathbf{q} \in \mathbb{R}^m$ where $\mathbf{q}_i \stackrel{\text{def}}{=} \mathbf{D}_{ii}^{(k)} (\mathbf{x}_i^{(k)} - \mathbf{x}_i^{(k-2^\ell)})$ if $\bar{\mathbf{x}}_i$ has not been updated after the $(k - 2^\ell)$ -th step, and $\mathbf{q}_i = 0$ otherwise. Our goal is precisely to find all large coordinates of \mathbf{q} .

Let \mathcal{T} be a constant-degree rooted tree with m leaves, where leaf i represents coordinate \mathbf{q}_i , which we call a *sampling tree*. For each node $u \in \mathcal{T}$, we define $E(u) \subseteq [m]$ to be the set of indices of leaves in the subtree rooted at u . We make a random descent down \mathcal{T} , in order to sample a coordinate i with probability proportional to \mathbf{q}_i^2 . At a node u , for each child u' of u , the total probability of the leaves under u' is given precisely by $\|\mathbf{q}|_{E(u')}\|_2^2$. We can estimate this by the Johnson-Lindenstrauss lemma using a sketching matrix Φ . Then we randomly move from

Algorithm 6 Data structure ABSTRACTMAINTAINAPPROX, Part 2

 34: **procedure** FINDLARGECOORDINATES(ℓ)

 35: $\triangleright \bar{\mathbf{D}}$: diagonal matrix such that

$$\bar{\mathbf{D}}_{ii} = \begin{cases} \mathbf{D}_{ii}^{(k)} & \text{if } \bar{\mathbf{x}}_i \text{ has not been updated after the } (k - 2^\ell)\text{-th step} \\ 0 & \text{otherwise.} \end{cases}$$

 36: $\triangleright \mathbf{q} \stackrel{\text{def}}{=} \bar{\mathbf{D}}(\mathbf{x}^{(k)} - \mathbf{x}^{(k-2^\ell)})$ \triangleright vector to sample coordinates from

37:

 38: $I \leftarrow \emptyset$ \triangleright set of candidate coordinates

 39: **for** $M_\ell \stackrel{\text{def}}{=} \Theta(2^{2\ell}(\beta/\delta)^2 \log^2 m \log(m/\rho))$ iterations **do**

 40: \triangleright Sample coordinate i of \mathbf{q} w.p. proportional to \mathbf{q}_i^2 by random descent down \mathcal{T} to a leaf

 41: **while true do**

 42: $u \leftarrow \text{root}(\mathcal{T}), p_u \leftarrow 1$

 43: **while** u is not a leaf node **do**

 44: Sample a child u' of u with probability

$$\mathbf{P}(u \rightarrow u') \stackrel{\text{def}}{=} \frac{\|\Phi_{E(u')} \mathbf{q}\|_2^2}{\sum_{\text{child } u'' \text{ of } u} \|\Phi_{E(u'')} \mathbf{q}\|_2^2}$$

 \triangleright let $\Phi_{E(u)} \stackrel{\text{def}}{=} \Phi \mathbf{I}_{E(u)}$ for each node u

 45: $p_u \leftarrow p_u \cdot \mathbf{P}(u \rightarrow u')$

 46: $u \leftarrow u'$

 47: **end while**

 48: **break** with probability $p_{\text{accept}} \stackrel{\text{def}}{=} \|\mathbf{q}|_{E(u)}\|^2 / (2 \cdot p_u \cdot \|\Phi \mathbf{q}\|_2^2)$

 49: **end while**

 50: $I \leftarrow I \cup E(u)$

 51: **end for**

 52: **return** $\{i \in I : \mathbf{q}_i \geq \frac{\delta}{2 \lceil \log m \rceil}\}$.

 53: **end procedure**

u down to child u' with probability proportional to the estimated value. To tolerate the estimation error, when reaching some leaf node representing coordinate i , we accept with probability proportional to the ratio between the exact probability of i and the estimated probability of i . If i is rejected, we repeat the process from the root again independently.

LEMMA E.2. Assume that $\|\mathbf{D}^{(k+1)}(\mathbf{x}^{(k+1)} - \mathbf{x}^{(k)})\|_2 \leq \beta$ for all IPM steps k . Let $\rho < 1$ be any given failure probability, and let $M_\ell \stackrel{\text{def}}{=} \Theta(2^{2\ell}(\beta/\delta)^2 \log^2 m \log(m/\rho))$ be the number of samples Algorithm 6 takes. Then with probability $\geq 1 - \rho$, during the k -th call of APPROXIMATE, Algorithm 6 finds the set $I_\ell^{(k)}$ correctly. Furthermore, the while-loop in Algorithm 6 happens only $O(1)$ times in expectation per sample.

Proof. The proof is similar to Lemma 6.17 in [17]. We include it for completeness. For a set S of indices, let \mathbf{I}_S be the $m \times m$ diagonal matrix that is one on S and zero otherwise.

We first prove that Algorithm 6 breaks with probability at least $\frac{1}{4}$. By the choice of w , Johnson–Lindenstrauss lemma shows that $\|\Phi_{E(u)}\mathbf{q}\|_2^2 = (1 \pm \frac{1}{9\eta})\|\mathbf{I}_{E(u)}\mathbf{q}\|_2^2$ for all $u \in \mathcal{T}$ with probability at least $1 - \rho$. Therefore, the probability we move from a node u to its child node u' is given by

$$\mathbf{P}(u \rightarrow u') = \left(1 \pm \frac{1}{3\eta}\right) \frac{\|\mathbf{I}_{E(u')}\mathbf{q}\|_2^2}{\sum_{u'' \text{ is a child of } u} \|\mathbf{I}_{E(u'')}\mathbf{q}\|_2^2} = \left(1 \pm \frac{1}{3\eta}\right) \frac{\|\mathbf{I}_{E(u')}\mathbf{q}\|_2^2}{\|\mathbf{I}_{E(u)}\mathbf{q}\|_2^2}.$$

Hence, the probability the walk ends at a leaf $u \in \mathcal{T}$ is given by

$$p_u = \left(1 \pm \frac{1}{3\eta}\right)^\eta \frac{\|\mathbf{I}_u\mathbf{q}\|_2^2}{\|\mathbf{q}\|_2^2} = \left(1 \pm \frac{1}{3\eta}\right)^\eta \frac{\|\mathbf{q}|_{E(u)}\|_2^2}{\|\mathbf{q}\|_2^2}.$$

Now, p_{accept} on Algorithm 6 is at least

$$p_{\text{accept}} = \frac{\|\mathbf{q}|_{E(u)}\|_2^2}{2 \cdot p_u \cdot \|\Phi\mathbf{q}\|_2^2} \geq \frac{\|\mathbf{q}|_{E(u)}\|_2^2}{2 \cdot \left(1 + \frac{1}{3\eta}\right)^\eta \frac{\|\mathbf{q}|_{E(u)}\|_2^2}{\|\mathbf{q}\|_2^2} \cdot \|\Phi\mathbf{q}\|_2^2} \geq \frac{\|\mathbf{q}\|_2^2}{2 \cdot \left(1 + \frac{1}{3\eta}\right)^\eta \|\Phi\mathbf{q}\|_2^2} \geq \frac{1}{4}.$$

On the other hand, we have that $p_{\text{accept}} \leq \frac{\|\mathbf{q}\|_2^2}{2\left(1 - \frac{1}{3\eta}\right)^\eta \|\Phi\mathbf{q}\|_2^2} < 1$ and hence this is a valid probability.

Next, we note that u is accepted on Algorithm 6 with probability

$$p_{\text{accept}}p_u = \frac{\|\mathbf{q}|_{E(u)}\|_2^2}{2 \cdot \|\Phi\mathbf{q}\|_2^2}.$$

Since $\|\Phi\mathbf{q}\|_2^2$ remains the same in all iterations, this probability is proportional to $\|\mathbf{q}|_{E(u)}\|_2^2$. Since the algorithm repeats when u is rejected, on Algorithm 6, u is chosen with probability exactly $\|\mathbf{q}|_{E(u)}\|_2^2 / \|\mathbf{q}\|_2^2$.

Now, we want to show the output set is exactly $\{i \in [n] : |\mathbf{q}_i| \geq \frac{\delta}{2\lceil \log m \rceil}\}$. Let S denote the set of indices where $\bar{\mathbf{x}}$ did not update between the $(k - 2^\ell)$ -th step and the current k -th step. Then

$$\begin{aligned} \|\mathbf{q}\|_2 &= \|\mathbf{I}_S \mathbf{D}^{(k)}(\mathbf{x}^{(k)} - \mathbf{x}^{(k-2^\ell)})\|_2 \\ &\leq \sum_{i=k-2^\ell}^{k-1} \|\mathbf{I}_S \mathbf{D}^{(k)}(\mathbf{x}^{(i+1)} - \mathbf{x}^{(i)})\|_2 \\ &= \sum_{i=k-2^\ell}^{k-1} \|\mathbf{I}_S \mathbf{D}^{(i+1)}(\mathbf{x}^{(i+1)} - \mathbf{x}^{(i)})\|_2 \\ &\leq \sum_{i=k-2^\ell}^{k-1} \|\mathbf{D}^{(i+1)}(\mathbf{x}^{(i+1)} - \mathbf{x}^{(i)})\|_2 \\ &\leq 2^\ell \beta, \end{aligned}$$

where we used $\mathbf{I}_S \mathbf{D}^{(i+1)} = \mathbf{I}_S \mathbf{D}^{(k)}$, because \bar{x}_i changes whenever \mathbf{D}_{i_i} changes at a step. Hence, each leaf u is sampled with probability at least $\frac{\|\mathbf{q}|_{E(u)}\|^2}{(2^\ell \beta)^2}$. If $|\mathbf{q}_i| \geq \frac{\delta}{2 \lceil \log m \rceil}$, and $i \in E(u)$ for a leaf node u , then the coordinate i is not in I with probability at most

$$\left(1 - \frac{\|\mathbf{q}|_{E(u)}\|^2}{(2^\ell \beta)^2}\right)^{M_\ell} \leq \left(1 - \frac{1}{2^{2\ell+2}(\beta/\delta)^2 \lceil \log m \rceil^2}\right)^{M_\ell} \leq \frac{\rho}{m},$$

by our choice of M_ℓ . Hence, all i with $|\mathbf{q}_i| \geq \frac{\delta}{2 \lceil \log m \rceil}$ lies in I with probability at least $1 - \rho$. This proves that the output set is exactly $I_\ell^{(k)}$ with probability at least $1 - \rho$. \square

REMARK E.2. In Algorithm 6, we only need to compute $\|\Phi_{E(u)} \mathbf{q}\|_2^2$ for $O(M_\ell)$ many nodes $u \in \mathcal{T}$. Furthermore, the randomness of the sketch is not leaked and we can use the same random sketch Φ throughout the algorithm. This allows us to efficiently maintain $\Phi_{E(u)} \mathbf{q}$ for each $u \in \mathcal{T}$ throughout the IPM.

E.3 Sketch maintenance In FINDLARGECOORDINATES in the previous subsection, we assumed the existence of a constant degree sampling tree \mathcal{T} , and for the dynamic vector \mathbf{q} , the ability to access $\Phi_{E(u)} \mathbf{q}$ at each node $u \in \mathcal{T}$ and $\mathbf{q}|_{E(u)}$ at each leaf node u .

In this section, we consider when the required sampling tree is the operator tree \mathcal{T} supporting a tree operator Δ , and the vector \mathbf{q} is $\mathbf{x} \stackrel{\text{def}}{=} \Delta \mathbf{z} + \sum_{H \in \mathcal{T}} \Delta^{(H)} \mathbf{u}_H$, where each of Δ, \mathbf{z} and the \mathbf{u}_H 's undergo changes at every IPM step. We present a data structure that implements two features efficiently on \mathcal{T} :

- access $\mathbf{x}|_{E(H)}$ at every leaf node H ,
- access $\Phi_{E(H)} \mathbf{x}$ at every node H , where $\Phi_{E(H)}$ is Φ restricted to columns given by $E(H)$.

LEMMA E.3. Let \mathcal{T} be a constant degree rooted tree with height η supporting tree operator Δ with query complexity Q . Let $w = \Theta(\eta^2 \log(\frac{m}{\rho}))$ be as defined in Algorithm 6, and let $\Phi \in \mathbb{R}^{w \times m}$ be a JL-sketch matrix. Then MAINTAINSKETCH (Algorithm 7) is a data structure that maintains $\Phi \mathbf{x}$, where \mathbf{x} is implicitly represented by

$$\mathbf{x} \stackrel{\text{def}}{=} \Delta \mathbf{z} + \sum_{H \in \mathcal{T}} \Delta^{(H)} \mathbf{u}_H.$$

The data structure supports the following procedures:

- INITIALIZE(operator tree \mathcal{T} , implicit \mathbf{x}): Initialize the data structure and compute the initial sketches in $O(Q(wm))$ time.
- UPDATE($\mathcal{H} \subseteq \mathcal{T}$): Update all the necessary sketches in $O(w \cdot Q(\eta|\mathcal{H}|))$ time, where \mathcal{H} is the set of all nodes H where \mathbf{u}_H or Δ_H changed.
- ESTIMATE($H \in \mathcal{T}$): Return $\Phi_{E(H)} \mathbf{x}$.
- QUERY($H \in \mathcal{T}$): Return $\mathbf{x}|_{E(H)}$.

If we call QUERY on N nodes, the total runtime is $O(Q(w\eta N))$.

If we call ESTIMATE along a sampling path (by which we mean starting at the root, calling estimate at both children of a node, and then recursively descending to one child until reaching a leaf), and then we call QUERY on the resulting leaf, and we repeat this N times with no updates during the process, then the total runtime of these calls is $O(Q(w\eta N))$.

We note that $\Delta \mathbf{z} = \sum_{H \in \mathcal{T}} \Delta^{(H)} \mathbf{z}|_{F_H}$. For simplicity, it suffices to give the algorithm for sketching the simpler $\mathbf{x} \stackrel{\text{def}}{=} \sum_{H \in \mathcal{T}} \Delta^{(H)} \mathbf{u}_H$.

Proof. Let us consider the correctness of the data structure, starting with the helper function SUMANCESTORS. We implement it using recursion and memoization as it is crucial for bounding subsequent runtimes.

Algorithm 7 Data structure for maintaining $\Phi \mathbf{x}$, Part 1

```

1: data structure MAINTAINSKETCH
2: private : member
3:    $\mathcal{T}$  : rooted constant degree tree, where at every node  $H$ , there is
4:      $\mathbf{S}^{(H)} \in \mathbb{R}^{w \times |E_H \cup \partial H|}$  : sketched subtree operator  $\Phi \Delta^{(H)}$ 
5:      $\mathbf{t}^{(H)} \in \mathbb{R}^w$  : sketched vector  $\Phi \sum_{H' \in \mathcal{T}_H} \Delta^{(H')} \mathbf{u}_{H'}$ 
6:      $\Phi \in \mathbb{R}^{w \times m}$  : JL-sketch matrix
7:      $\Delta \in \mathbb{R}^{m \times n}$  : dynamic tree operator on  $\mathcal{T}$ 
8:      $\mathbf{u}_H$  at every  $H \in \mathcal{T}$  : dynamic vectors
9:
10: procedure INITIALIZE(tree  $\mathcal{T}$ ,  $\Phi \in \mathbb{R}^{w \times m}$ , tree operator  $\Delta$ ,  $\mathbf{u}_H$  for each  $H \in \mathcal{T}$ )
11:    $\Phi \leftarrow \Phi, \mathcal{T} \leftarrow \mathcal{T}, \Delta \leftarrow \Delta, \mathbf{u}_H \leftarrow \mathbf{u}_H$  for each  $H \in \mathcal{T}$ 
12:    $\mathbf{S}^{(H)} \leftarrow \mathbf{0}, \mathbf{t}^{(H)} \leftarrow \mathbf{0}$  for each  $H \in \mathcal{T}$ 
13:   UPDATE( $V(\mathcal{T})$ )
14: end procedure
15:
16: procedure UPDATE( $\mathcal{H} \stackrel{\text{def}}{=} \text{set of nodes admitting implicit representation changes}$ )
17:   for  $H \in \mathcal{P}_{\mathcal{T}}(\mathcal{H})$  going up the tree level by level do
18:      $\mathbf{S}^{(H)} \leftarrow \sum_{\text{child } D \text{ of } H} \mathbf{S}^{(D)} \Delta_D$ 
19:      $\mathbf{t}^{(H)} \leftarrow \mathbf{S}^{(H)} \mathbf{u}_H + \sum_{\text{child } D \text{ of } H} \mathbf{t}^{(D)}$ 
20:   end for
21: end procedure
22:
23: procedure SUMANCESTORS( $H \in \mathcal{T}$ )
24:   if UPDATE has not been called since the last call to SUMANCESTORS( $H$ ) then
25:     return the result of the last SUMANCESTORS( $H$ )
26:   end if
27:   if  $H$  is the root then return  $\mathbf{0}$ 
28:   end if
29:   return  $\Delta_H(\mathbf{u}_P + \text{SUMANCESTORS}(P))$   $\triangleright P$  is the parent of  $H$ 
30: end procedure
31:
32: procedure ESTIMATE( $H \in \mathcal{T}$ )
33:   Let  $\mathbf{y}$  be the result of SUMANCESTORS( $H$ )
34:   return  $\mathbf{S}^{(H)} \mathbf{y} + \mathbf{t}^{(H)}$ 
35: end procedure
36:
37: procedure QUERY(leaf  $H \in \mathcal{T}$ )
38:   return  $\mathbf{u}_H + \text{SUMANCESTORS}(H)$ 
39: end procedure

```

CLAIM 6. $\text{SUMANCESTORS}(H \in \mathcal{T})$ returns $\sum_{\text{ancestor } A \text{ of } H} \Delta_{H \leftarrow A} \mathbf{u}_A$.

Proof. At the root, there are no ancestors, hence we return the zero matrix. When H is not the root, suppose P is the parent of H . Then we can recursively write

$$\sum_{\text{ancestor } A \text{ of } H} \Delta_{H \leftarrow A} \mathbf{u}_A = \Delta_H \left(\mathbf{u}_P + \sum_{\text{ancestor } A \text{ of } P} \Delta_{P \leftarrow A} \mathbf{u}_A \right).$$

The procedure implements the right hand side, and is therefore correct. \square

Assuming we correctly maintain $\mathbf{S}^{(H)} \stackrel{\text{def}}{=} \Phi \Delta^{(H)}$ and $\mathbf{t}^{(H)} \stackrel{\text{def}}{=} \Phi \sum_{H' \in \mathcal{T}_H} \Delta^{(H')} \mathbf{u}_{H'}$ at every node H , ESTIMATE and QUERY return the correct answers by the tree operator decomposition given in Lemma B.1.

For UPDATE, note that if a node H is not in \mathcal{H} and it has no descendants in \mathcal{H} , then by definition, the sketches at H are not changed. Hence, it suffices to update the sketches only at all nodes in $\mathcal{P}_{\mathcal{T}}(\mathcal{H})$. We update the nodes from the bottom of \mathcal{T} upwards, so that when we are at a node H , all the sketches at its descendant nodes are correct. Therefore, by definition, the sketches at H is also correct.

Now we consider the runtimes:

INITIALIZE: It sets the sketches to $\mathbf{0}$ in $O(wm)$ time, and then calls UPDATE to update the sketches everywhere on \mathcal{T} . By the correctness runtime of UPDATE, this step is correct and runs in $\tilde{O}(Q(wm))$ time.

UPDATE(set of nodes \mathcal{H} admitting implicit representation changes): First note that $|\mathcal{P}_{\mathcal{T}}(\mathcal{H})| \leq \eta|\mathcal{H}|$. For each node $H \in \mathcal{H}$ with children D_1, D_2 , Algorithm 7 multiplies each row of $\mathbf{S}^{(D_1)}$ with $\Delta_{(D_1, H)}$, each row of $\mathbf{S}^{(D_2)}$ with Δ_{D_2} , and sums the results. Summing over w -many rows and over all nodes in $\mathcal{P}_{\mathcal{T}}(\mathcal{H})$, we see the total runtime of Algorithm 7 is $O(Q(w\eta|\mathcal{H}|))$.

Algorithm 7 multiply each row of $\mathbf{S}^{(H)}$ with a vector and then performs a constant number of additions of w -length vectors. Since $\mathbf{S}^{(H)}$ is computed for all $H \in \mathcal{P}_{\mathcal{T}}(\mathcal{H})$ in $O(Q(w\eta|\mathcal{H}|))$ total time, this must also be a bound on their number of total non-zero entries. Since each $\mathbf{S}^{(H)}$ is used once in Algorithm 7 for a matrix-vector multiplication, the total runtime of Algorithm 7 is $O(Q(w\eta|\mathcal{H}|))$.

All other lines are not bottlenecks.

Overall ESTIMATE and QUERY time along N sampling paths: We show that if we call ESTIMATE along N sampling paths each from the root to a leaf, and we call QUERY on the leaves, the total cost is $O(Q(w\eta N))$:

Suppose the set of nodes visited is given by \mathcal{H} , then $|\mathcal{H}| \leq \eta N$. Since there is no update, and ESTIMATE is called for a node only after it is called for its parent, we know that $\text{SUMANCESTORS}(H)$ is called exactly once for each $H \in \mathcal{H}$. Each $\text{SUMANCESTOR}(H)$ multiplies a unique edge operator $\Delta_{(H, P)}$ with a vector. Hence, the total runtime of SUMANCESTORS is $O(|\mathcal{H}|)$.

Finally, each QUERY applies a leaf operator to the output of a unique SUMANCESTORS call, so the overall runtime is certainly bounded by $O(Q(|\mathcal{H}|))$. Similarly, each ESTIMATE multiplies $\mathbf{S}^{(H)}$ with the output of a unique SUMANCESTORS call. This can be computed as w -many vectors each multiplied with the SUMANCESTORS output. Then two vectors of length w are added. Summing over all nodes in \mathcal{H} , the overall runtime is $O(Q(w|\mathcal{H}|)) = O(Q(w\eta N))$.

QUERY time on N leaves: Since this is a subset of the work described above, the runtime must also be bounded by $O(Q(w\eta N))$.

\square

F.1 Proof of 3.3 We combine the previous three subsections for the overall approximation procedure. It is essentially $\text{ABSTRACTMAINTAINAPPROX}$ in Algorithm 6, with the abstractions replaced by a data structure implementation. We omit the pseudocode and simply describe the functions.

THEOREM F.1. (APPROXIMATE VECTOR MAINTENANCE WITH TREE OPERATOR [15]) *Let $0 < \rho < 1$ be a failure probability. Suppose $\Delta \in \mathbb{R}^{m \times n}$ is a tree operator with query complexity Q and supported on a constant-degree tree \mathcal{T} with height η . There is a randomized data structure MAINTAINAPPROX that takes as input the dynamic weights \mathbf{w} and the dynamic \mathbf{x} implicitly maintained according to Theorem 3.2 at every step, and explicitly maintains the approximation $\bar{\mathbf{x}}$ to \mathbf{x} satisfying $\|\mathbf{W}^{-1/2}(\mathbf{x} - \bar{\mathbf{x}})\|_{\infty} \leq \delta$ at every step with probability $1 - \rho$.*

Suppose $\|\mathbf{W}^{(k)-1/2}(\mathbf{x}^{(k)} - \mathbf{x}^{(k-1)})\|_2 \leq \beta$ for all steps k . Furthermore, suppose \mathbf{w} is a function of $\bar{\mathbf{x}}$ coordinate-wise. Then, for each $\ell \geq 0$, $\bar{\mathbf{x}}$ admits $2^{2\ell}$ coordinate changes every 2^ℓ steps. Over N total steps, the total cost of

the data structure is

$$(3.3) \quad \tilde{O}(\eta^3(\beta/\delta)^2 \log^3(mN/\rho)) \left(Q(m) + \sum_{k=1}^N Q(S^{(k)}) + \sum_{\ell=0}^{\log N} \frac{N}{2^\ell} \cdot Q(2^{2\ell}) \right),$$

where $S^{(k)}$ is the number of nodes H where Δ_H or \mathbf{u}_H in the implicit representation of \mathbf{x} changed at step k .

Proof. We apply Lemma E.1 using \mathbf{x} maintained by MAINTAINREP and $\mathbf{D} \stackrel{\text{def}}{=} \mathbf{W}^{-1/2}$ from SOLVE. We create $O(\log m)$ copies of MAINTAINSKETCH (Lemma E.3), so that for each $0 \leq \ell \leq O(\log m)$, we have one copy $\text{sketch}_{\ell,x}$ which maintains sketches of $\Phi \bar{\mathbf{D}} \mathbf{x}^{(k)}$ at step k , and one copy sketch_ℓ which maintains sketches of $\Phi \bar{\mathbf{D}} \mathbf{x}^{(k-2^\ell)}$ at step $k \geq 2^\ell$, where $\bar{\mathbf{D}}$ is defined so $\bar{\mathbf{D}}_{i,i} = \mathbf{D}_{i,i}$ if $\bar{\mathbf{x}}_i$ has not been updated after the $k - 2^\ell$ -th step, and $\bar{\mathbf{D}}_{i,i} = 0$ otherwise (as needed in Algorithm 6). Note that $\bar{\mathbf{D}}$ can be absorbed into the tree operator in the implicit representation of \mathbf{x} , so Lemma E.3 does indeed apply.

To access sketches of the vector $\mathbf{q} \stackrel{\text{def}}{=} \bar{\mathbf{D}}(\mathbf{x}^{(k)} - \mathbf{x}^{(k-2^\ell)})$ as needed in FINDLARGECOORDINATES in Algorithm 6, we can simply access the corresponding sketch in $\text{sketch}_{\ell,x}$ and sketch_ℓ , and then take the difference.

We now describe each procedure in words, and then prove their correctness and runtime.

INITIALIZE($\mathcal{T}, \mathbf{x}, \mathbf{D}, \rho, \delta$): This procedure implements the initialization of ABSTRACTMAINTAINAPPROX to approximate the dynamic vector \mathbf{x} which is given implicitly. The initialization steps described in Algorithm 6 takes $O(wm)$ time. Then, we initialize the $O(\log m)$ copies of MAINTAINSKETCH in $O(Q(wm) \log m)$ time by Lemma E.3.

UPDATE($\mathbf{x}^{(\text{new})}, \mathbf{D}^{(\text{new})}$): To implement UPDATE, it suffices to update all the sketching data structures. Let us fix ℓ , and consider the update time for $\text{sketch}_{\ell,x}$ and sketch_ℓ .

Lemma E.1 shows that throughout SOLVE, there are $O(2^{2\ell}(\beta/\delta)^2 \log^2 m)$ -many coordinate updates to $\bar{\mathbf{x}}$ every 2^ℓ steps. Since \mathbf{D} is a function of $\bar{\mathbf{x}}$ coordinate-wise, $\bar{\mathbf{x}}_i = \mathbf{x}_i^{(k-1)}$ for all i where $\mathbf{D}_{ii}^{(k)} \neq \mathbf{D}_{ii}^{(k-1)}$ by Algorithm 6. The diagonal matrix $\bar{\mathbf{D}}$ is the same as \mathbf{D} , except $\bar{\mathbf{D}}_{ii}$ is temporarily zeroed out for 2^ℓ steps after $\bar{\mathbf{x}}_i$ changes at a step. So, the overall number of coordinate changes to $\bar{\mathbf{D}}$ is $O(2^{2\ell})$ -many every 2^ℓ steps.

Let $S^{(k)}$ denote the number of nodes H where Δ_H or \mathbf{u}_H in the implicit representation of \mathbf{x} changed at step k . Additionally, since the sketching data structures maintain some variant of $\mathbf{D}\mathbf{x}$ (where $\bar{\mathbf{D}}$ is viewed as absorbed in the tree operator), every coordinate change in $\bar{\mathbf{D}}$ implies an edge operator update. Now we apply Lemma E.3 to conclude that the total time for all UPDATE calls for $\text{sketch}_{\ell,x}$ and sketch_ℓ over N steps is:

$$O(1) \cdot \left(\sum_{k=1}^N Q(w\eta S^{(k)}) + \frac{N}{2^\ell} \cdot Q(w\eta \cdot 2^{2\ell}) \right) \leq O(w\eta) \cdot \left(\sum_{k=1}^N Q(S^{(k)}) + \frac{N}{2^\ell} \cdot Q(2^{2\ell}) \right).$$

We then sum this over all ℓ to get the total update time for the sketching data structures.

APPROXIMATE: There are two operations to be implemented in the subroutine FINDLARGECOORDINATES(ℓ): Accessing $\Phi_{E(u)} \mathbf{q}$ at a node u , and accessing $\mathbf{q}|_{E(u)}$ at a leaf node u . For the first, we call $\text{sketch}_{\ell,x}.\text{ESTIMATE}(u) - \text{sketch}_\ell.\text{ESTIMATE}(u)$. For the second, we call $\text{sketch}_{\ell,x}.\text{QUERY}(u) - \text{sketch}_\ell.\text{QUERY}(u)$.

To set $\bar{\mathbf{x}}_i$ as $\mathbf{x}_i^{(k)}$ for a single coordinate at step k as needed in Algorithm 6, we find the leaf node H containing the edge e , and call $\text{sketch}_{0,x}.\text{QUERY}(H)$. This returns the sub-vector $\mathbf{x}^{(k)}|_{E(H)}$, from which we can extract $\mathbf{x}_i^{(k)}$ and set $\bar{\mathbf{x}}_i$ to be the value. This line is not a bottleneck in the runtime.

We compute the total runtime over N APPROXIMATE calls. For every $\ell \geq 0$, we call FINDLARGECOORDINATES(ℓ) once every 2^ℓ steps, for a total of $N/2^\ell$ calls. In a single call, $M_\ell \stackrel{\text{def}}{=} \Theta(2^{2\ell}(\beta/\delta)^2 \log^2 m \log(mN/\rho))$ sampling paths are explored in the sketch_ℓ and $\text{sketch}_{\ell,x}$ data structures by Lemma E.2, where a sampling path correspond to one iteration of the while-loop. This takes a total of $O(Q(w\eta M_\ell))$ time by Lemma E.3. Therefore, for every fixed ℓ , the total time for all FINDLARGECOORDINATES(ℓ) calls is

$$\frac{N}{2^\ell} \cdot O(Q(w\eta M_\ell)).$$

The total time for all LARGECOORDINATES calls is obtained by summing over all values of $\ell = 0, \dots, \log N$. To achieve overall failure probability at most ρ , it suffices to set the failure probability of each call to be $O(\rho/N)$. \square

We sum up the initialization time, update and approximate time for all values of $\ell = 0, \dots, \log N$ and over N total steps of SOLVE, to get the overall runtime of the data structure:

$$\begin{aligned} & \tilde{O}(Q(wm)) + O(w\eta) \sum_{k=1}^N Q(S^{(k)}) + O(w\eta) \sum_{\ell=0}^{\log N} \frac{N}{2^\ell} (Q(2^{2^\ell}) + Q(M_\ell)) \\ &= \tilde{O}(\eta^3(\beta/\delta)^2 \log^3(mN/\rho)) \left(Q(m) + \sum_{k=1}^N Q(S^{(k)}) + \sum_{\ell=0}^{\log N} \frac{N}{2^\ell} \cdot Q(2^{2^\ell}) \right). \end{aligned}$$