

Faster Min-Cost Flow and Approximate Tree Decomposition on Bounded Treewidth Graphs

Sally Dong ✉

University of Washington, Seattle, WA, USA

Guanghao Ye ✉ 

Massachusetts Institute of Technology, Cambridge, MA, USA

Abstract

We present an algorithm for min-cost flow in graphs with n vertices and m edges, given a tree decomposition of width τ and size S , and polynomially bounded, integral edge capacities and costs, running in $\tilde{O}(m\sqrt{\tau} + S)$ time. This improves upon the previous fastest algorithm in this setting achieved by the bounded-treewidth linear program solver of [26, 16], which runs in $\tilde{O}(m\tau^{(\omega+1)/2})$ time, where $\omega \approx 2.37$ is the matrix multiplication exponent. Our approach leverages recent advances in structured linear program solvers and robust interior point methods (IPM). In general graphs where treewidth is trivially bounded by n , the algorithm runs in $\tilde{O}(m\sqrt{n})$ time, which is the best-known result without using the Lee-Sidford barrier or ℓ_1 IPM, demonstrating the surprising power of robust interior point methods.

As a corollary, we obtain a $\tilde{O}(\text{tw}^3 \cdot m)$ time algorithm to compute a tree decomposition of width $O(\text{tw} \cdot \log(n))$, given a graph with m edges.

2012 ACM Subject Classification Theory of computation \rightarrow Approximation algorithms analysis; Theory of computation \rightarrow Parameterized complexity and exact algorithms; Theory of computation \rightarrow Network optimization

Keywords and phrases Min-cost flow, tree decomposition, interior point method, bounded treewidth graphs

Digital Object Identifier 10.4230/LIPIcs.ESA.2024.49

Related Version *Full Version:* <https://arxiv.org/abs/2308.14727> [18]

Funding *Guanghao Ye:* Supported by NSF awards CCF-1955217 and DMS-2022448.

1 Introduction

An active area of research in recent years is the advancement of interior point methods (IPM) for linear and convex programs, with its origins tracing back to the works of [30, 34]. This, along with the design of problem-specific data structures supporting the IPM, had led to breakthroughs in faster linear program solvers [13] and faster max flow algorithms [31, 10, 24, 4, 12], among others.

One line of research, inspired by nested dissection from [33] and methods in numerical linear algebra [14], focuses on exploiting separable structures in the constraint matrix of the linear program, which can be captured succinctly by associating the matrix with a graph. Several classical graph-theoretic concepts are appropriate here: Given a graph $G = (V, E)$ on n vertices, we say $S \subseteq V$ is a (b) -balanced vertex separator if there exists some constant $b \in (0, 1)$ such that every connected component of $G \setminus S$ has size at most bn . For a function f , we say G is f -separable if any subgraph H of G has a balanced separator of size $f(|V(H)|)$. The *treewidth* of a graph informally measures how close a graph is to a tree, and is closely related to separability. If G has treewidth τ , then any subgraph of G has a $1/2$ -balanced separator of size $\tau + 1$; conversely, if G is τ -separable, then G has treewidth at most $\tau \log n$ (c.f. [7]). Treewidth has been an important focus in the study of parametrized complexity [19, 8].



© Sally Dong and Guanghao Ye;
licensed under Creative Commons License CC-BY 4.0
32nd Annual European Symposium on Algorithms (ESA 2024).

Editors: Timothy Chan, Johannes Fischer, John Iacono, and Grzegorz Herman; Article No. 49; pp. 49:1–49:14
Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Formally,

► **Definition 1.** A tree decomposition of a graph G is a pair (X, T) where T is a tree, and $X : V(T) \mapsto 2^{V(G)}$ is a family of subsets of $V(G)$ labelling the nodes of T called bags, such that

1. $\cup_{t \in V(T)} X(t) = V(G)$,
2. for each $v \in V(G)$, the nodes $t \in V(T)$ with $v \in X(t)$ induces a connected subgraph of T , and
3. for each $e = uv \in E(G)$, there is a node $t \in V(T)$ such that $u, v \in X(t)$.

Given a tree decomposition (X, T) , its width is defined to be $\max\{|X(t)| - 1 : t \in T\}$; and its size is $\sum_{t \in V(T)} |X(t)|$. The treewidth of G is the minimum width over all tree decompositions of G .

By leveraging various separable properties of graphs, the sequence of papers [17, 15, 16] have iteratively refined the robust IPM framework and associated data structures for solving structured linear programs. [17] gave the first general linear program solver parametrized by treewidth: Given an LP of the form $\min\{\mathbf{c}^\top \mathbf{x} : \mathbf{A}\mathbf{x} = \mathbf{b}, \mathbf{x} \geq \mathbf{0}\}$ and a width τ decomposition of the dual graph¹ $G_{\mathbf{A}}$ of the constraint matrix \mathbf{A} , suppose the feasible region has inner radius r and outer radius R , and the costs are polynomially bounded. Then the LP can be solved to ε -accuracy in $\tilde{O}(m\tau^2 \log(R/(\varepsilon r)))$ time². The τ^2 factor arises from carefully analyzing the sparsity-pattern of the Cholesky factorization of $\mathbf{A}\mathbf{A}^\top$ and the associated matrix computation times. This runtime dependence on τ was improved from quadratic to $(\omega + 1)/2 \approx 1.68$ in [26], by means of a coordinate-grouping technique applied to the updates during the IPM. [26] further combined [17] with [36] to give the current best algorithm for semidefinite programs with bounded treewidth. The result of [17] is now subsumed by [16], which shows that under the same setup, except where $G_{\mathbf{A}}$ is known to be a κn^α -separable graph³ the LP can be solved in $\tilde{O}((\kappa^2(m + m^{2\alpha+0.5}) + m\kappa^{(\omega-1)/(1-\alpha)} + \kappa^\omega n^{\alpha\omega}) \cdot \log(\frac{mR}{\varepsilon r}))$ time. Taking $\kappa = \tau$ and $\alpha = 0$ for bounded treewidth graphs recovers the [17] result. For n^α -separable graphs, the runtime expression simplifies to $\tilde{O}((m + m^{1/2+2\alpha}) \log(\frac{mR}{\varepsilon r}))$.

In the special setting of flow problems, the constraint matrix \mathbf{A} is the vertex-edge incidence matrix of a graph and $\mathbf{A}\mathbf{A}^\top$ is corresponding graph Laplacian. Fast Laplacian solvers [35, 28] and approximate Schur complements [20] can be combined with the separable structure of the graph to further speed up computations at every IPM step. Using these ideas, [15] gave a nearly-linear time min-cost flow algorithm for planar graphs, which are $O(\sqrt{n})$ -separable. Notably, the combination of nested dissections with fast approximate Schur complements was explored before in [25], where they demonstrated how to maintain dynamic effective resistances in separable graphs.

In this work, we expand the landscape by giving the first specific result for min-cost flow on bounded treewidth graphs. Previously, [23] showed how to compute a max vertex flow in $O(\tau^2 \cdot n \log n)$ time given a directed graph with unit capacities and its tree decomposition of width τ ; their approach is combinatorial in nature.

► **Theorem 2 (Main result).** Let $G = (V, E)$ be a directed graph with n vertices and m edges. Assume that the demands \mathbf{d} , edge capacities \mathbf{u} and costs \mathbf{c} are all integers and bounded by M in absolute value. Given a tree decomposition of G with width τ and size S , there is an algorithm that computes a minimum-cost flow in $\tilde{O}(m\sqrt{\tau} \log M + S)$ expected time.

¹ The dual graph of a matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$ is a graph on n vertices, where each column of \mathbf{A} gives rise to a clique (or equivalently, a hyper-edge). Importantly, when the linear program is a flow problem on a graph, $G_{\mathbf{A}}$ is precisely the graph in the original problem.

² Throughout the paper, we use $\tilde{O}(\cdot)$ to hide polylog m and polylog M factors.

³ Here, κ is a constant expression, but could be a function of the size of $G_{\mathbf{A}}$ (which is considered fixed).

As a direct corollary of Theorem 2, we can solve min-cost flow on any graph with n vertices, m edges and integral polynomially-bounded costs and constraints in $\tilde{O}(m\sqrt{n})$ expected time, as the treewidth of any n -vertex graph is at most $\tau = n$, and the tree decomposition is trivially the graph itself. This result matches that of [31] obtained using the Lee-Sidford barrier for the IPM, which requires $\tilde{O}(\sqrt{n})$ iterations. In contrast, we use the standard log barrier which requires $\tilde{O}(\sqrt{m})$ iterations, and leverage the robustness of the IPM and custom data structures to reduce the amortized cost per iteration. We find it noteworthy that all other max flow results matching or beating our time require either the Lee-Sidford barrier ([10]) or significantly divergent IPM techniques ([12]).

Using our faster max-flow algorithm as a subroutine, we can efficiently compute a tree decomposition of any given graph, where the width is within a $O(\log n)$ -factor of the optimal:

► **Corollary 3** (Approximating treewidth). *Let $G = (V, E)$ be a graph with n vertices, m edges, and treewidth $\text{tw}(G)$. There is an algorithm to find a tree decomposition of G with width at most $O(\text{tw}(G) \cdot \log n)$ in $\tilde{O}(\text{tw}(G)^3 \cdot m)$ expected time.*

It is well known that computing the treewidth of a graph is NP-hard [1], and there is conditional hardness result for even constant-factor approximation algorithm [3]. For polynomial time algorithms, the best known result is a $O(\sqrt{\log \text{tw}})$ -approximation algorithm by [22], which involves solving a semidefinite program with n^2 variables.

There is a series of works focused on computing approximate treewidth for small treewidth graph in nearly-linear time; we refer the readers to [6] for a more detailed survey. Notably, [23] showed for any graph G , there is an algorithm to compute a tree decomposition of width $O(\text{tw}(G)^2)$ in $\tilde{O}(\text{tw}(G)^7 \cdot n)$ time. [11] improved the running time to $\tilde{O}(\text{tw}(G)^3 \cdot m)$ with slightly compromised approximation ratio $O(\text{tw}(G)^2 \cdot \log^{1+o(1)} n)$. More recently, [5] showed how to compute a tree decomposition of width $O(\text{tw}(G) \cdot \log^3 n)$ in $O(m^{1+o(1)})$ time.

1.1 Overview of techniques

The foundation of our algorithm is the planar min-cost flow algorithm from [15]. We begin with the identical robust IPM algorithm in abstraction, as given in Algorithm 1. [15] first defines a separator tree \mathcal{T} for the input graph, and uses it as the basis for the data structures in the IPM. We modify the separator tree construction, so that instead of recursively decomposing the input planar graph which is $O(\sqrt{n})$ -separable, we recursively decompose the bounded treewidth graph. The leaf nodes of \mathcal{T} partition the edges of the input graph; We guarantee that each leaf node contains $O(\tau)$ -many edges, compared to constantly-many in the planar case.

There are two main components to the data structures from [15]:

1. A data structure DYNAMICSC (Theorem 10) is used to maintain approximate Schur complement matrices at every node of of the separator tree \mathcal{T} . It implicitly represents the matrix $(\mathbf{B}^\top \mathbf{W} \mathbf{B})^{-1}$, where \mathbf{W} are edge weights changing at every step of the IPM. We use DYNAMICSC in exactly the same way.
2. A data structure MAINTAINSOLN (Theorem 11) using \mathcal{T} to implicitly maintain the primal and dual solutions \mathbf{f} and \mathbf{s} , and explicitly maintain approximate solutions $\bar{\mathbf{f}}$ and $\bar{\mathbf{s}}$ at the current IPM iteration. The approximate solutions are used in the subsequent iteration to compute the step direction \mathbf{v} and edge weights \mathbf{w} .

In [15], the approximate solutions $\bar{\mathbf{f}}$ and $\bar{\mathbf{s}}$ are updated coordinate-wise whenever a coordinate is sufficiently far from $\bar{\mathbf{f}}$ and $\bar{\mathbf{s}}$ respectively. An update induces changes in the data structures as follows: If $\bar{\mathbf{f}}_e$ or $\bar{\mathbf{s}}_e$ is updated for an edge e (subsequently, \mathbf{w}_e and \mathbf{v}_e

are updated), then we find the unique leaf node H in \mathcal{T} containing the edge e , and update DYNAMICSC and MAINTAIN SOLN at all nodes along the path from H to the root of \mathcal{T} . [15] shows this runtime depends on the sizes of the nodes visited.

[26] introduced a natural grouping technique for the coordinate updates, where coordinates are grouped into blocks, and updates are performed block-wise. Since the leaves of our separator tree \mathcal{T} contain $O(\tau)$ edges, it is natural for us to also incorporate a blocking scheme, where the blocks are given by the edge partition according to the leaves of \mathcal{T} . In this case, the runtime for data structure updates is the same whether we update a single coordinate or a block containing said coordinate, since they affect the same path from the leaf node to the root of \mathcal{T} . We bound the overall runtime expression using properties of the new separator tree.

Lastly, the RIPM guarantees that for each k iterations, $\tilde{O}(k^2)$ -many blocks of $\bar{\mathbf{f}}$ and $\bar{\mathbf{s}}$ are updated, meaning that running our data structures for many IPM iterations leads to superlinear runtime scaling. We can, however, restart our data structures at any point, by explicitly computing the current exact solutions \mathbf{f}, \mathbf{s} and reinitializing the data structures with these values in $\tilde{O}(m)$ total time. On balance, we choose to restart the data structures every $\sqrt{m/\tau}$ -many iterations, for a total restarting cost of $\tilde{O}(m\sqrt{\tau})$. Between each restart, we make $O(m/\tau)$ -many block updates using $\tilde{O}(\tau)$ time each, for a total update cost of $\tilde{O}(m\sqrt{\tau})$. Hence, the overall runtime is $\tilde{O}(m\sqrt{\tau})$.

2 Robust interior point method

For the sake of completion, we give the robust interior point method from [17, 16], which is a refinement of the methods in [13, 9], for solving linear programs of the form

$$\min_{\mathbf{f} \in \mathcal{F}} \mathbf{c}^\top \mathbf{f} \quad \text{where } \mathcal{F} = \{\mathbf{B}^\top \mathbf{f} = \mathbf{b}, \mathbf{l} \leq \mathbf{f} \leq \mathbf{u}\} \quad (2.1)$$

for some matrix $\mathbf{B} \in \mathbb{R}^{m \times n}$.

► **Theorem 4** ([17]). *Consider the linear program*

$$\min_{\mathbf{B}^\top \mathbf{f} = \mathbf{b}, \mathbf{l} \leq \mathbf{f} \leq \mathbf{u}} \mathbf{c}^\top \mathbf{f}$$

with $\mathbf{B} \in \mathbb{R}^{m \times n}$. Suppose there exists some interior point $\tilde{\mathbf{f}}_o$ satisfying $\mathbf{B}^\top \tilde{\mathbf{f}}_o = \mathbf{b}$ and $\mathbf{l} + r \leq \tilde{\mathbf{f}}_o \leq \mathbf{u} - r$,⁴ for some scalar $r > 0$. Let $L = \|\mathbf{c}\|_2$ and $R = \|\mathbf{u} - \mathbf{l}\|_2$. For any $0 < \epsilon \leq 1/2$, the algorithm RIPM (Algorithm 1) finds \mathbf{f} such that $\mathbf{B}^\top \mathbf{f} = \mathbf{b}$, $\mathbf{l} \leq \mathbf{f} \leq \mathbf{u}$, and

$$\mathbf{c}^\top \mathbf{f} \leq \min_{\mathbf{B}^\top \mathbf{f} = \mathbf{b}, \mathbf{l} \leq \mathbf{f} \leq \mathbf{u}} \mathbf{c}^\top \mathbf{f} + \epsilon LR.$$

Furthermore, the algorithm has the following properties:

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

We note that this algorithm only requires access to $(\bar{\mathbf{f}}, \bar{\mathbf{s}})$, but not (\mathbf{f}, \mathbf{s}) during the main while-loop in SOLVE. Hence, (\mathbf{f}, \mathbf{s}) can be implicitly maintained via any data structure.

⁴ 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} . We define $\mathbf{v} - x$ to be the vector obtained by subtracting x from each coordinate of \mathbf{v} .

■ **Algorithm 1** Robust Interior Point Method from [17].

```

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

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

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


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

  ▷ Optimize the original linear program
12:   $(\mathbf{f}, \mathbf{s}) \leftarrow \text{SOLVE}(\mathbf{B}, \phi, \mathbf{f}, \mathbf{s}, LR, \frac{\epsilon}{4m})$ 
13:  return  $\mathbf{f}$ 
14: end procedure

15: procedure SOLVE( $\mathbf{B}, \phi, \mathbf{f}, \mathbf{s}, t_{\text{start}}, t_{\text{end}}$ )
16:   Let  $\alpha = \frac{1}{2^{20}\lambda}$  and  $\lambda = 64 \log(256m^2)$  where  $m$  is the number of rows in  $\mathbf{B}$ 
17:   Let  $t \leftarrow t_{\text{start}}, \bar{\mathbf{f}} \leftarrow \mathbf{f}, \bar{\mathbf{s}} \leftarrow \mathbf{s}, \bar{t} \leftarrow t$ 
18:   while  $t \geq t_{\text{end}}$  do
19:      $t \leftarrow \max\left\{\left(1 - \frac{\alpha}{\sqrt{m}}\right)t, t_{\text{end}}\right\}$ 
20:     Update  $h = -\alpha / \|\cosh(\lambda\gamma^{\bar{t}}(\bar{\mathbf{f}}, \bar{\mathbf{s}}))\|_2$ 
21:     Update the diagonal weight matrix  $\mathbf{W} = \nabla^2\phi(\bar{\mathbf{f}})^{-1}$ 
22:     Update the direction  $\mathbf{v}$  where  $\mathbf{v}_i = \sinh(\lambda\gamma^{\bar{t}}(\bar{\mathbf{f}}, \bar{\mathbf{s}})_i)\mu^{\bar{t}}(\bar{\mathbf{f}}, \bar{\mathbf{s}})_i$ 
23:     Pick  $\mathbf{v}^\parallel$  and  $\mathbf{v}^\perp$  such that  $\mathbf{W}^{-1/2}\mathbf{v}^\parallel \in \text{Range}(\mathbf{B}), \mathbf{B}^\top \mathbf{W}^{1/2}\mathbf{v}^\perp = \mathbf{0}$  and

      
$$\|\mathbf{v}^\parallel - \mathbf{P}_w \mathbf{v}\|_2 \leq \alpha \|\mathbf{v}\|_2,$$

      
$$\|\mathbf{v}^\perp - (\mathbf{I} - \mathbf{P}_w)\mathbf{v}\|_2 \leq \alpha \|\mathbf{v}\|_2 \quad (\mathbf{P}_w \stackrel{\text{def}}{=} \mathbf{W}^{1/2}\mathbf{B}(\mathbf{B}^\top \mathbf{W}\mathbf{B})^{-1}\mathbf{B}^\top \mathbf{W}^{1/2})$$


24:     Implicitly update  $\mathbf{f} \leftarrow \mathbf{f} + h\mathbf{W}^{1/2}\mathbf{v}^\perp, \mathbf{s} \leftarrow \mathbf{s} + \bar{t}h\mathbf{W}^{-1/2}\mathbf{v}^\parallel$ 
25:     Explicitly maintain  $\bar{\mathbf{f}}, \bar{\mathbf{s}}$  such that  $\|\mathbf{W}^{-1/2}(\bar{\mathbf{f}} - \mathbf{f})\|_\infty \leq \alpha$  and  $\|\mathbf{W}^{1/2}(\bar{\mathbf{s}} - \mathbf{s})\|_\infty \leq \bar{t}\alpha$ 
26:     Update  $\bar{t} \leftarrow t$  if  $|\bar{t} - t| \geq \alpha\bar{t}$ 
27:   end while
28:   return  $(\mathbf{f}, \mathbf{s})$ 
29: end procedure

```

3 Nested dissection on bounded treewidth graphs

In this section, we show how to leverage the structural properties of bounded treewidth graphs to find a sparse Cholesky factorization of $\mathbf{L} \stackrel{\text{def}}{=} \mathbf{B}^\top \mathbf{W} \mathbf{B}$, and hence implicitly maintain an approximation of \mathbf{L}^{-1} as part of the projection matrix $\mathbf{P}_w \stackrel{\text{def}}{=} \mathbf{W}^{1/2} \mathbf{B} (\mathbf{B}^\top \mathbf{W} \mathbf{B})^{-1} \mathbf{B}^\top \mathbf{W}^{1/2}$ used in the RIPM.

3.1 Separator tree for bounded treewidth graph

The notion of using a *separator tree* to represent the recursive decomposition of a separable graph is well-established in literature, c.f [21, 27]. In [15], the authors show that the separator tree can be used to construct a sparse approximate projection matrix for RIPM. Here, we extend their result to bounded treewidth graphs.

► **Definition 5** (Separator tree). *Let G be a graph with n vertices and m edges. A separator tree \mathcal{T} of G is a rooted constant-degree tree whose nodes represent a recursive decomposition of G based on vertex separators.*

Formally, each node H of \mathcal{T} is a region (edge-induced subgraph) of G ; we denote this by $H \in \mathcal{T}$. At a node H , we store 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 vertex separator of H ; and F_H is the set of eliminated vertices at H .

The nodes of \mathcal{T} satisfy the following recursive definition:

1. *The root of \mathcal{T} is the node $H = G$, with $\partial H = \emptyset$ and $F_H = S(H)$.*
2. *A non-leaf node $H \in \mathcal{T}$ has exactly two children $H_1, H_2 \in \mathcal{T}$ that form an edge-disjoint partition of H , and $S(H) \stackrel{\text{def}}{=} V(H_1) \cap V(H_2)$ is a vertex separator of H . Define the set of eliminated vertices at H to be $F_H \stackrel{\text{def}}{=} S(H) \setminus \partial H$.
By definition of boundary vertices, we have $\partial H_1 \stackrel{\text{def}}{=} (\partial H \cup S(H)) \cap V(H_1)$, and $\partial H_2 \stackrel{\text{def}}{=} (\partial H \cup S(H)) \cap V(H_2)$.*
3. *If H is a leaf node, define $S(H) = \emptyset$ and $F_H = V(H) \setminus \partial H$.*

By construction, the leaf nodes of \mathcal{T} partition the edges of G . If H is a leaf node, let $E(H)$ denote the edges contained in H . If H is not a leaf node, let $E(H)$ denote the union of all the edges in the leaf nodes in the subtree \mathcal{T}_H . Let η denote the height of \mathcal{T} .

Next, we show how to construct an appropriate separator tree for bounded treewidth graphs.

► **Theorem 6.** *Let G be a graph with n vertices and m edges. Given a tree decomposition of G with width τ and size s , we can construct a separator tree \mathcal{T} of G in $\tilde{O}(s)$ time, such that \mathcal{T} has height $\eta = O(\log m)$, each node $H \in \mathcal{T}$ satisfies $|F_H \cup \partial H| \leq O(\tau \text{ poly log}(m))$, and each leaf node $H \in \mathcal{T}$ has $|E(H)| \leq \Theta(\tau)$.*

To prove the theorem above, we need a lemma about balanced vertex separators. Let us adopt the notation (A, S, B) to mean that S is a vertex separator such that its removal from the graph leaves two disjoint subgraphs A and B .

► **Lemma 7** (Slight modification of [17, Theorem 4.17]). *Let (X, T) be a tree decomposition of a graph G on n vertices with width τ and size s . Then in $O(s)$ time, we can find a $2/3$ -balanced vertex separator (A, S, B) of G and compute tree decompositions (X_1, T_1) of $G[A \cup S]$ with width τ and size s_1 , and (X_2, T_2) of $G[B \cup S]$ with width τ and size s_2 , such that $s_1 + s_2 \leq s + \tau$.*

Proof. By scanning through the bags of T in $O(s)$ time, we can find a node $t \in T$ such that $T \setminus t$ is two disjoint subtrees T'_1 and T'_2 , with $|\bigcup_{u \in T'_1} X(s) \setminus X(t)| \leq 2/3n$, and similarly for T'_2 . Then $S \stackrel{\text{def}}{=} X(t) \subseteq V(G)$ is a $2/3$ -balanced vertex separator of G , and $A \stackrel{\text{def}}{=} \bigcup_{u \in T'_1} X(u)$ and $B \stackrel{\text{def}}{=} \bigcup_{u \in T'_2} X(u)$ partitions the remaining vertices of G .

Observe that (X, T) restricted to the nodes $T_1 \cup t$ yields a tree decomposition of $G[A \cup S]$ (with vertices of B deleted from the bags), and similarly when restricted to $T_2 \cup t$ yields a tree decomposition of $G[B \cup S]$. Both trivially have width τ . Since only the node t is shared between the two, and t has size at most τ , we see that $s_1 + s_2 \leq s + \tau$. ◀

Proof of Theorem 6. We construct the separator tree \mathcal{T} by recursively decomposing the graph starting with subgraph $H = G$:

1. Given a subgraph H of G , if $|E(H)| \leq \Theta(\tau)$, then set H as a leaf node.
2. Else if $|V(H)| > 2\tau$, find a $2/3$ -balanced vertex separator (A, S, B) of H using the tree decomposition of H by Lemma 7. Let (X_1, T_1) and (X_2, T_2) be the resulting tree decompositions of $G[A \cup S]$ and $G[B \cup S]$.

Let H_1 and H_2 be subgraphs of H where $V(H_1) = A \cup S$ and $V(H_2) = B \cup S$. Partition $E(H)$ so that edges incident to A are in H_1 and edges incident to B are in H_2 , and edges in $S(H)$ can be arbitrarily assigned to H_1 or H_2 . Observe that (X_1, T_1) and (X_2, T_2) are indeed tree decompositions of H_1 and H_2 respectively. Set H_1 and H_2 as the children of H .

3. Else if $|V(H)| \leq 2\tau$, then we partition the edges of H into two sets each of size at most $\frac{2}{3}|E(H)|$, and let H_1 and H_2 be graphs on $V(H)$ with their respective edge sets.⁵

Consider a non-leaf node H with children H_1 and H_2 , we note that

$$|V(H_i)| \cdot |E(H_i)| \leq \frac{2}{3}|V(H)| \cdot |E(H)| \quad \text{for } i \in \{1, 2\}.$$

This directly shows the height of the separator tree \mathcal{T} is $\eta = O(\log m)$. At any node H , the cardinality of $F_H \cup \partial H$ is bounded by the number of vertices contained in any ancestor node of H in \mathcal{T} , which is at most $O(\tau \log m)$.

The runtime is bounded by the total size of the tree decompositions across node generated by option 2 in the above construction. [16] shows that there are $\tilde{O}(n/\tau)$ -many interior nodes in \mathcal{T} generated this way. Each such node contributes a term of τ to the total size, while each layer in the separator tree contributes a term of s . In total, this gives a size of $s \log m + (n/\tau) \cdot \tau = \tilde{O}(s + n)$, as desired. ◀

3.2 Nested dissection using a separator tree

► **Definition 8** (Block Cholesky decomposition). *The block Cholesky decomposition of a symmetric matrix \mathbf{L} with two blocks indexed by F and C is:*

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

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 :

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

⁵ We use 2τ here to differentiate between the second and third case, instead of the more natural τ , in order to avoid any infinite-loop edge cases in the recursive process arising from division and rounding.

We use the separator tree structure to factor the matrix $\mathbf{L}^{-1} \stackrel{\text{def}}{=} (\mathbf{B}^\top \mathbf{W} \mathbf{B})^{-1}$:

► **Theorem 9** (Approximate \mathbf{L}^{-1} factorization, c.f. [15, Theorem 33]). *Let \mathcal{T} be the separator tree of G with height η . For each node $H \in \mathcal{T}$ with edges $E(H)$, let $\mathbf{B}[H]$ denote the matrix \mathbf{B} restricted to rows indexed by $E(H)$, and define $\mathbf{L}[H] \stackrel{\text{def}}{=} \mathbf{B}[H]^\top \mathbf{W} \mathbf{B}[H]$.*

Given approximation parameter $\epsilon_{\mathbf{P}}$, suppose for each node H at level i of \mathcal{T} , we have a matrix $\mathbf{L}^{(H)}$ satisfying the $e^{\epsilon_{\mathbf{P}}}$ -spectral approximation

$$\mathbf{L}^{(H)} \approx_{i\epsilon_{\mathbf{P}}} \mathbf{Sc}(\mathbf{L}[H], \partial H \cup F_H). \quad (3.1)$$

Then, we can approximate \mathbf{L}^{-1} by

$$\mathbf{L}^{-1} \approx_{\eta\epsilon_{\mathbf{P}}} \mathbf{\Pi}^{(0)\top} \dots \mathbf{\Pi}^{(\eta-1)\top} \mathbf{\Gamma} \mathbf{\Pi}^{(\eta-1)} \dots \mathbf{\Pi}^{(0)}, \quad (3.2)$$

where

$$\mathbf{\Gamma} \stackrel{\text{def}}{=} \begin{bmatrix} \sum_{H \in \mathcal{T}^{(0)}} \left(\mathbf{L}_{F_H, F_H}^{(H)} \right)^{-1} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \ddots & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \sum_{H \in \mathcal{T}^{(\eta)}} \left(\mathbf{L}_{F_H, F_H}^{(H)} \right)^{-1} \end{bmatrix}, \quad (3.3)$$

and for $i = 0, \dots, \eta - 1$,

$$\mathbf{\Pi}^{(i)} \stackrel{\text{def}}{=} \mathbf{I} - \sum_{H \in \mathcal{T}^{(i)}} \mathbf{X}^{(H)}, \quad (3.4)$$

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

$$\mathbf{X}^{(H)} \stackrel{\text{def}}{=} \mathbf{L}_{\partial H, F_H}^{(H)} \left(\mathbf{L}_{F_H, F_H}^{(H)} \right)^{-1}. \quad (3.5)$$

[15] gives a data structure to maintain an implicit representation of \mathbf{L}^{-1} as the weights \mathbf{w} undergoes changes in the IPM:

► **Theorem 10** ([15, Theorem 6]). *Given a graph G with m edges and its $\tilde{O}(\tau)$ -separator tree \mathcal{T} with height $\eta = O(\log m)$, there is a deterministic data structure DYNAMICSC which maintains the edge weights \mathbf{w} from the IPM, and at every node $H \in \mathcal{T}$, maintains two Laplacians $\mathbf{L}^{(H)}$ and $\tilde{\mathbf{S}}\mathbf{c}(\mathbf{L}^{(H)}, \partial H)$ dependent on \mathbf{w} . It supports the following procedures:*

- *INITIALIZE($G, \mathbf{w} \in \mathbb{R}_{>0}^m, \epsilon_{\mathbf{P}} > 0$): Given a graph G , initial weights \mathbf{w} , projection matrix approximation accuracy $\epsilon_{\mathbf{P}}$, preprocess in $\tilde{O}(\epsilon_{\mathbf{P}}^{-2}m)$ time.*
- *REWEIGHT($\mathbf{w} \in \mathbb{R}_{>0}^m$, given implicitly as a set of changed coordinates): Update the weights to \mathbf{w} in $\tilde{O}(\epsilon_{\mathbf{P}}^{-2} \sum_{H \in \mathcal{H}} |F_H \cup \partial H|)$ time, where $\mathcal{H} \stackrel{\text{def}}{=} \{H \in \mathcal{T} : (\mathbf{w} - \mathbf{w}^{(\text{prev})})|_{E(H)} \neq \mathbf{0}\}$ is the set of nodes containing an edge with updated weight. (Note that \mathcal{H} is a union of paths from leaf nodes to the root.)*
- *Access to Laplacian $\mathbf{L}^{(H)}$ at any node $H \in \mathcal{T}$ in time $\tilde{O}(\epsilon_{\mathbf{P}}^{-2}|\partial H \cup F_H|)$.*
- *Access to Laplacian $\tilde{\mathbf{S}}\mathbf{c}(\mathbf{L}^{(H)}, \partial H)$ at any node $H \in \mathcal{T}$ in time $\tilde{O}(\epsilon_{\mathbf{P}}^{-2}|\partial H|)$.*

Furthermore, with high probability, for any node H in \mathcal{T} , if H is at level i , then

$$\mathbf{L}^{(H)} \approx_{i\epsilon_{\mathbf{P}}} \mathbf{Sc}(\mathbf{L}[H], \partial H \cup F_H), \quad \text{and}$$

$$\tilde{\mathbf{S}}\mathbf{c}(\mathbf{L}^{(H)}, \partial H) \approx_{\epsilon_{\mathbf{P}}} \mathbf{Sc}(\mathbf{L}^{(H)}, \partial H).$$

4 Solution maintenance

Assuming the correct maintenance of Laplacians and Schur complements along a recursive separator tree, [15] gave detailed data structures for maintaining the exact and approximate flow and slack solutions throughout the RIPM. Recall that the leaf nodes of the separator tree \mathcal{T} form a partition of the edges of G . In [15], each leaf node of the separator tree contains $O(1)$ -many edges, whereas in our case, each leaf node contains $O(\tau)$ -many edges, and therefore we update the approximate solution in a block manner. The data structures in [15] naturally generalized from coordinate-wise updates to the block-wise case, so we use their implementation in a black-box manner. Here, we state a combined version of their main theorems.

We use $\mathbf{f}_{[i]}$ to denote the subvector of \mathbf{f} indexed by edges in the i -th leaf node of \mathcal{T} , and similarly $\mathbf{s}_{[i]}$. We use $\mathbf{f}_{[i]}^{(k)}$ and $\mathbf{s}_{[i]}^{(k)}$ to denote the vector $\mathbf{f}_{[i]}$ and $\mathbf{s}_{[i]}$ at the k -th step of the IPM. Each block contains $O(\tau)$ -many variables.

► **Theorem 11** ([15, Theorems 9, 10]). *Given a graph G with m edges and its separator tree \mathcal{T} with height η , there is a randomized data structure implicitly maintains the IPM solution pair (\mathbf{f}, \mathbf{s}) undergoing IPM changes, and explicitly maintains its approximation $(\bar{\mathbf{f}}, \bar{\mathbf{s}})$, and supports the following procedures with high probability against an adaptive adversary:*

- **INITIALIZE**($G, \mathbf{f}^{(\text{init})} \in \mathbb{R}^m, \mathbf{s}^{(\text{init})} \in \mathbb{R}^m, \mathbf{v} \in \mathbb{R}^m, \mathbf{w} \in \mathbb{R}_{>0}^m, \epsilon_{\mathbf{P}} > 0, \bar{\epsilon} > 0$): *Given a graph G , initial solutions $\mathbf{f}^{(\text{init})}, \mathbf{s}^{(\text{init})}$, initial direction \mathbf{v} , initial weights \mathbf{w} , target step accuracy $\epsilon_{\mathbf{P}}$ and target approximation accuracy $\bar{\epsilon}$, preprocess in $\tilde{O}(m\epsilon_{\mathbf{P}}^{-2})$ time, set the implicit representations $\bar{\mathbf{f}} \leftarrow \mathbf{f}^{(\text{init})}, \bar{\mathbf{s}} \leftarrow \mathbf{s}^{(\text{init})}$, and set the approximations $\bar{\mathbf{f}} \leftarrow \mathbf{f}, \bar{\mathbf{s}} \leftarrow \mathbf{s}$.*
- **REWEIGHT**($\mathbf{w} \in \mathbb{R}_{>0}^m$, given implicitly as a set of changed weights): *Set the current weights to \mathbf{w} in $\tilde{O}(\epsilon_{\mathbf{P}}^{-2}\tau K)^6$ time, where K is the number of blocks changed in \mathbf{w} .*
- **MOVE**($\alpha \in \mathbb{R}, \mathbf{v} \in \mathbb{R}^m$ given implicitly as a set of changed coordinates): *Implicitly update*

$$\begin{aligned} \mathbf{s} &\leftarrow \mathbf{s} + \alpha \mathbf{W}^{-1/2} \tilde{\mathbf{P}}_{\mathbf{w}} \mathbf{v}, \\ \mathbf{f} &\leftarrow \mathbf{f} + \alpha \mathbf{W}^{1/2} \mathbf{v} - \alpha \mathbf{W}^{1/2} \tilde{\mathbf{P}}'_{\mathbf{w}} \mathbf{v}, \end{aligned}$$

for some $\tilde{\mathbf{P}}_{\mathbf{w}}$ satisfying $\|(\tilde{\mathbf{P}}_{\mathbf{w}} - \mathbf{P}_{\mathbf{w}})\mathbf{v}\|_2 \leq \eta\epsilon_{\mathbf{P}}\|\mathbf{v}\|_2$ and $\tilde{\mathbf{P}}_{\mathbf{w}}\mathbf{v} \in \text{Range}(\mathbf{B})$, and some other $\tilde{\mathbf{P}}'_{\mathbf{w}}$ satisfying $\|\tilde{\mathbf{P}}'_{\mathbf{w}}\mathbf{v} - \mathbf{P}_{\mathbf{w}}\mathbf{v}\|_2 \leq O(\eta\epsilon_{\mathbf{P}})\|\mathbf{v}\|_2$ and $\mathbf{B}^{\top} \mathbf{W}^{1/2} \tilde{\mathbf{P}}'_{\mathbf{w}} \mathbf{v} = \mathbf{B}^{\top} \mathbf{W}^{1/2} \mathbf{v}$. The total runtime is $\tilde{O}(\epsilon_{\mathbf{P}}^{-2}\tau K)$, where K is the number of blocks changed in \mathbf{v} .

- **APPROXIMATE**() $\rightarrow \mathbb{R}^{2m}$: *Return the vector pair $\bar{\mathbf{f}}, \bar{\mathbf{s}}$ implicitly as a set of changed coordinates, satisfying $\|\mathbf{W}^{-1/2}(\bar{\mathbf{f}} - \mathbf{f})\|_{\infty} \leq \bar{\epsilon}$ and $\|\mathbf{W}^{1/2}(\bar{\mathbf{s}} - \mathbf{s})\|_{\infty} \leq \bar{\epsilon}$, for the current weight \mathbf{w} and the current solutions \mathbf{f}, \mathbf{s} .*
- **EXACT**() $\rightarrow \mathbb{R}^{2m}$: *Output the current vector \mathbf{f}, \mathbf{s} in $\tilde{O}(m\epsilon_{\mathbf{P}}^{-2})$ time.*

Suppose $\alpha\|\mathbf{v}\|_2 \leq \beta$ for some β for all calls to **MOVE**. Suppose in each step, **REWEIGHT**, **MOVE** and **APPROXIMATE** are called in order. Let K denote the total number of blocks changed in \mathbf{v} and \mathbf{w} between the $(k-1)$ -th and k -th **REWEIGHT** and **MOVE** calls. Then at the k -th **APPROXIMATE** call,

- the data structure sets $\bar{\mathbf{f}}_{[i]} \leftarrow \mathbf{f}_{[i]}^{(k)}, \bar{\mathbf{s}}_{[i]} \leftarrow \mathbf{s}_{[i]}^{(k)}$ for $O(N_k \stackrel{\text{def}}{=} 2^{2\ell_k} (\frac{\beta}{\bar{\epsilon}})^2 \log^2 m)$ blocks i , where ℓ_k is the largest integer ℓ with $k \equiv 0 \pmod{2^\ell}$ when $k \neq 0$ and $\ell_0 = 0$, and
- the amortized time for the k -th **APPROXIMATE** call is $\tilde{O}(\epsilon_{\mathbf{P}}^{-2}\tau(K + N_{k-2^{\ell_k}}))$.

⁶ The original bound here is $\tilde{O}(\epsilon_{\mathbf{P}}^{-2}\sqrt{mK})$, where the \sqrt{mK} factor comes from the fact that they can bound $\sum_{H \in \mathcal{H}} |F_H \cup \partial H| \leq \sqrt{mK}$ and $\mathcal{H} = \{H \in \mathcal{T} : (\mathbf{w} - \mathbf{w}^{(\text{prev})})|_{E(H)} \neq \mathbf{0}\}$. See [15, Section 9] for more details. Here, we replace it by $\tilde{O}(\tau K)$ by the guarantees of Theorem 6.

An alternative derivation of the running time for APPROXIMATE shows that 2^ℓ steps take $\tilde{O}(2^{2^\ell}\tau)$ time. For $2^\ell = \Omega(\sqrt{m})$ steps of IPM, this would take $\tilde{O}(m\tau)$ total time, which is too expensive for our use case. Since initialization only takes $\tilde{O}(m)$ time, we reinitialize the data structure every $2^\ell = \Omega(\sqrt{m}/\tau)$ steps, allowing us to avoid running the same data structure for too many steps and accruing large superlinear costs.

5 Proof of main theorems

Proof of Theorem 2. We apply RIPM from Theorem 4 combined with the data structures from Theorem 11.

First, we reduce the min-cost flow problem to min-cost circulation, for the sake of simplicity in showing the existence of an interior point in the polytope. We begin by adding extra vertices s and t to the input graph G , which increases the treewidth of G by at most 2. For every vertex v with demand $\mathbf{d}_v < 0$, we add a directed edge from s to v with capacity $-\mathbf{d}_v$ and cost 0. For every vertex v with $\mathbf{d}_v > 0$, we add a directed edge from v to t with capacity \mathbf{d}_v and cost 0. Then, we add a directed edge from t to s with capacity $4nM$ and cost $-4nM$. The cost and capacity on the $t \rightarrow s$ edge is chosen such that the min-cost flow problem on the original graph is equivalent to the min-cost circulation on this new graph.

We apply Algorithm 1 in Theorem 4 to find a circulation \mathbf{f} in the new graph such that $(\mathbf{c}^{\text{new}})^\top \mathbf{f} \leq \text{OPT} + \frac{1}{2}$ by setting $\epsilon = \frac{1}{CM^2m^2}$ for some large constant C . The other parameters L, R, r in the theorem can be bounded by $L = \|\mathbf{c}^{\text{new}}\|_2 = O(Mm)$, $R = \|\mathbf{u}^{\text{new}} - \mathbf{l}^{\text{new}}\|_2 = O(Mm)$, and $r \geq \frac{1}{4m}$, as shown in [15]. Now, the solution \mathbf{f} , when restricted to the original graph, is almost a flow routing the required demand, with flow value off by at most $\frac{1}{2nM}$ from the optimal. This is because sending extra k units of fractional flow from s to t gives extra negative cost $\leq -knM$. Now we can round \mathbf{f} to an integral flow \mathbf{f}^{int} with same or better flow value using no more than $\tilde{O}(m)$ time [29]. Since \mathbf{f}^{int} is integral with flow value at least the total demand minus $\frac{1}{2}$, \mathbf{f}^{int} routes the demand completely. Again, since \mathbf{f}^{int} is integral with cost at most $\text{OPT} + \frac{1}{2}$, \mathbf{f}^{int} must have the minimum cost.

The RIPM given in Algorithm 1 runs the subroutine SOLVE twice. In the first run, the constraint matrix is the incidence matrix of the input graph G copied three times. Copying edges does not affect treewidth, and our data structures allow for duplicate edges, so we may treat the two runs identically. We construct a separator tree \mathcal{T} based on the input graph in $\tilde{O}(S)$ time using Theorem 6, and implement the IPM data structures on top of \mathcal{T} using Theorem 11.

Finally, we bound the IPM runtime. We initialize the data structures for flow and slack by INITIALIZE. Here, the data structures are given the first IPM step direction \mathbf{v} for preprocessing; the actual step is taken in the first iteration of the main while-loop. At each step, we perform the implicit update of \mathbf{f} and \mathbf{s} using MOVE; we update \mathbf{W} in the data structures using REWEIGHT; and we construct the explicit approximations $\bar{\mathbf{f}}$ and $\bar{\mathbf{s}}$ using APPROXIMATE; each in the respective flow and slack data structures. We return the true (\mathbf{f}, \mathbf{s}) by EXACT. The total cost of is dominated by MOVE, REWEIGHT, and APPROXIMATE.

Since we call MOVE, REWEIGHT and APPROXIMATE in order in each step and the runtime for MOVE, REWEIGHT are both dominated by the runtime for APPROXIMATE, it suffices to bound the runtime for APPROXIMATE only. Theorem 4 guarantees that there are $T = O(\sqrt{m} \log n \log(nM))$ total APPROXIMATE calls. We implement reinitialize the data structures every \sqrt{m}/τ steps.

At the k -th step after any initialization, the number of blocks changed in \mathbf{w} and \mathbf{v} is bounded by $K \stackrel{\text{def}}{=} O(2^{2^{\ell_k-1}} \log^2 m)$, where ℓ_k is the largest integer ℓ with $k \equiv 0 \pmod{2^\ell}$, or equivalently, the number of trailing zeros in the binary representation of k . Theorem 4

guarantees we can apply Theorem 11 with parameter $\beta = O(1/\log m)$, which in turn shows the amortized time for the k -th call is $\tilde{O}(\epsilon_{\mathbf{P}}^{-2\tau}(K + N_{k-2^{\ell_k}}))$, where $N_k \stackrel{\text{def}}{=} 2^{2^{\ell_k}}(\beta/\alpha)^2 \log^2 m = O(2^{2^{\ell_k}} \log^2 m)$, with $\alpha = O(1/\log m)$ and $\epsilon_{\mathbf{P}} = O(1/\log m)$. Observe that $K + N_{k-2^{\ell_k}} = O(N_{k-2^{\ell_k}})$. Now, summing over $T = \sqrt{m}/\tau$ steps, the total runtime between each restart is

$$O(\sqrt{m}/\tau) \sum_{k=1}^T \tilde{O}(\tau N_{k-2^{\ell_k}}) = O(\sqrt{m}/\tau) \sum_{\ell=0}^{\log T} \frac{T}{2^{\ell}} \cdot \tilde{O}(2^{2^{\ell}} \tau) = \tilde{O}(m).$$

Theorem 4 guarantees $O(\sqrt{m} \log n \log(nM))$ -many IPM steps in total. The data structure restarts $\tilde{O}(\sqrt{\tau})$ -many times, and each initialization time is $\tilde{O}(m)$. Hence, the total runtime for the RIPM is $\tilde{O}(m\sqrt{\tau} \log M)$. ◀

► **Corollary 3** (Approximating treewidth). *Let $G = (V, E)$ be a graph with n vertices, m edges, and treewidth $\text{tw}(G)$. There is an algorithm to find a tree decomposition of G with width at most $O(\text{tw}(G) \cdot \log n)$ in $\tilde{O}(\text{tw}(G)^3 \cdot m)$ expected time.*

Our algorithm for Corollary 3 requires some tree decomposition of the graph as input. We use the following lemma to construct the initial tree decomposition.

► **Lemma 12** ([11]). *For any $\frac{2}{3} < \alpha < 1$ and $0 < \epsilon < 1 - \alpha$, given a graph G with n vertices and m edges, if the graph G contains an α -balanced vertex separator of size K , then there is a randomized algorithm that finds a balanced vertex separator of size $\tilde{O}(K^2/\epsilon)$ in $\tilde{O}(mK^3/\epsilon)$ expected time. The algorithm does not require knowledge of K .*

Next, the lemma below establishes the relationship between max flow and balanced edge separators. We first give the relevant definitions. For a given constant $c \leq 1/2$, a directed edge-cut (S, \bar{S}) is called a c -balanced edge separator if both $|S| \geq cn$ and $|\bar{S}| \geq cn$. The capacity of the cut (S, \bar{S}) is the total capacity of all edges crossing the cut. The *minimum c -balanced edge separator* is the c -balanced edge separator with minimum capacity. A λ *pseudo-approximation* to the minimum c -balanced edge separator is a c' -balanced cut (S, \bar{S}) for some other constant c' , whose capacity is within a factor of λ of that of the minimum c -balanced edge separator.

► **Lemma 13** ([2]). *An $O(\log n)$ pseudo-approximation to the minimum c -balanced edge separator in directed graphs can be computed using polylog n single-commodity flow computations on the same graph.*

Proof of Corollary 3. It is well known that given a $O(\log n)$ approximation algorithm for finding a balanced vertex separator, one can construct a tree decomposition of width $O(\text{tw}(G) \log n)$. Specifically, the algorithm of [7] finds such a tree decomposition by recursively using a balanced vertex separator algorithm and requires only an additional log factor in the runtime.

Now, it suffices to show we can find a $\log(n)$ pseudo-approximation balanced vertex separator in $\tilde{O}(m \cdot \text{tw}(G)^3)$ expected time. Using the reduction from [32], we reduce the balanced vertex separator to directed edge separator on graph $G^* = (V^*, E^*)$, where

$$V^* = \{v \mid v \in V\} \cup \{v' \mid v \in V\},$$

and

$$E^* = \{(v, v') \mid v \in V\} \cup \{(u', v) \mid (u, v) \in E\} \cup \{(v', u) \mid (u, v) \in E\}.$$

We note that $\text{tw}(G^*) = O(\text{tw}(G))$. This shows G^* has a $2/3$ -balanced vertex separator of size $O(\text{tw}(G))$. We first use Lemma 12 to construct a $\tilde{O}(\text{tw}(G)^2)$ -separator tree for G^* . Then, we use the algorithm in [2] combined with our flow algorithm to find a balanced edge separator in $\tilde{O}(m \cdot \text{tw}(G))$ expected time. Hence, we can find a balanced vertex separator in $\tilde{O}(m \cdot \text{tw}(G)^3)$ expected time. ◀

References

- 1 Stefan Arnborg, Derek G Corneil, and Andrzej Proskurowski. Complexity of finding embeddings in a k -tree. *SIAM Journal on Algebraic Discrete Methods*, 8(2):277–284, 1987. doi:10.1137/0608024.
- 2 Sanjeev Arora and Satyen Kale. A combinatorial, primal-dual approach to semidefinite programs. *Journal of the ACM (JACM)*, 63(2):1–35, 2016. doi:10.1145/2837020.
- 3 Per Austrin, Toniann Pitassi, and Yu Wu. Inapproximability of treewidth, one-shot pebbling, and related layout problems. In *International Workshop on Approximation Algorithms for Combinatorial Optimization (APPROX)*, pages 13–24. Springer, 2012. doi:10.1007/978-3-642-32512-0_2.
- 4 Jan van den Band, Yu Gao, Arun Jambulapati, Yin Tat Lee, Yang P. Liu, Richard Peng, and Aaron Sidford. Faster maxflow via improved dynamic spectral vertex sparsifiers. *CoRR*, abs/2112.00722, 2021. doi:10.48550/arXiv.2112.00722.
- 5 Aaron Bernstein, Maximilian Probst Gutenberg, and Thatchaphol Saranurak. Deterministic decremental sssp and approximate min-cost flow in almost-linear time. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1000–1008. IEEE, 2022. doi:10.1109/FOCS52979.2021.00100.
- 6 Hans L Bodlaender, Pål Grónås Drange, Markus S Dregi, Fedor V Fomin, Daniel Lokshtanov, and Michał Pilipczuk. A $c^k n$ 5-approximation algorithm for treewidth. *SIAM Journal on Computing*, 45(2):317–378, 2016. doi:10.1137/130947374.
- 7 Hans L Bodlaender, John R Gilbert, Hjálmtýr Hafsteinsson, and Ton Kloks. Approximating treewidth, pathwidth, frontsize, and shortest elimination tree. *Journal of Algorithms*, 18(2):238–255, 1995. doi:10.1006/jagm.1995.1009.
- 8 Hans L Bodlaender and Arie MCA Koster. Combinatorial optimization on graphs of bounded treewidth. *The Computer Journal*, 51(3):255–269, 2008. doi:10.1093/comjnl/bxm037.
- 9 Jan van den Brand. A deterministic linear program solver in current matrix multiplication time. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 259–278. SIAM, 2020. doi:10.1137/1.9781611975994.16.
- 10 Jan van den Brand, Yin Tat Lee, Yang P Liu, Thatchaphol Saranurak, Aaron Sidford, Zhao Song, and Di Wang. Minimum cost flows, MDPs, and ℓ_1 -regression in nearly linear time for dense instances. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 859–869, 2021. doi:10.1145/3406325.3451108.
- 11 Sebastian Brandt and Roger Wattenhofer. Approximating small balanced vertex separators in almost linear time. *Algorithmica*, 81:4070–4097, 2019. doi:10.1007/s00453-018-0490-x.
- 12 Li Chen, Rasmus Kyng, Yang P Liu, Richard Peng, Maximilian Probst Gutenberg, and Sushant Sachdeva. Maximum flow and minimum-cost flow in almost-linear time. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 612–623. IEEE, 2022. doi:10.1109/FOCS54457.2022.00064.
- 13 Michael B Cohen, Yin Tat Lee, and Zhao Song. Solving linear programs in the current matrix multiplication time. *Journal of the ACM (JACM)*, 68(1):1–39, 2021. doi:10.1145/3424305.
- 14 Timothy A Davis. *Direct methods for sparse linear systems*. SIAM, 2006.
- 15 Sally Dong, Yu Gao, Gramoz Goranci, Yin Tat Lee, Richard Peng, Sushant Sachdeva, and Guanghao Ye. 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)*, pages 124–153. SIAM, 2022. doi:10.1137/1.9781611977073.7.

- 16 Sally Dong, Gramoz Goranci, Lawrence Li, Sushant Sachdeva, and Guanghai Ye. Fast algorithms for separable linear programs. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3558–3604. SIAM, 2024. doi:10.1137/1.9781611977912.127.
- 17 Sally Dong, Yin Tat Lee, and Guanghai Ye. 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, STOC 2021*, pages 1784–1797. ACM, 2021. doi:10.1145/3406325.3451056.
- 18 Sally Dong and Guanghai Ye. Faster min-cost flow and approximate tree decomposition on bounded treewidth graphs. *CoRR*, abs/2308.14727, 2023. doi:10.48550/arXiv.2308.14727.
- 19 Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*, volume 4 of *Texts in Computer Science*. Springer, 2013. doi:10.1007/978-1-4471-5559-1.
- 20 David Durfee, Rasmus Kyng, John Peebles, Anup B. Rao, and Sushant Sachdeva. Sampling random spanning trees faster than matrix multiplication. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017*, pages 730–742, 2017. doi:10.1145/3055399.3055499.
- 21 David Eppstein, Zvi Galil, Giuseppe F Italiano, and Thomas H Spencer. Separator based sparsification: I. planarity testing and minimum spanning trees. *journal of computer and system sciences*, 52(1):3–27, 1996. doi:10.1006/jcss.1996.0002.
- 22 Uriel Feige, MohammadTaghi Hajiaghayi, and James R Lee. Improved approximation algorithms for minimum weight vertex separators. *SIAM Journal on Computing*, 38(2):629, 2008. doi:10.1137/05064299X.
- 23 Fedor V Fomin, Daniel Lokshantov, Saket Saurabh, Michał Pilipczuk, and Marcin Wrochna. Fully polynomial-time parameterized computations for graphs and matrices of low treewidth. *ACM Transactions on Algorithms (TALG)*, 14(3):1–45, 2018. doi:10.1145/3186898.
- 24 Yu Gao, Yang P. Liu, and Richard Peng. Fully dynamic electrical flows: Sparse maxflow faster than Goldberg-Rao. In *62st IEEE Annual Symposium on Foundations of Computer Science, FOCS2021*. IEEE, 2021. doi:10.1109/FOCS52979.2021.00058.
- 25 Gramoz Goranci, Monika Henzinger, and Pan Peng. Dynamic effective resistances and approximate Schur Complement on separable graphs. In *26th Annual European Symposium on Algorithms, ESA 2018*, volume 112 of *LIPICs*, pages 40:1–40:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPICs.ESA.2018.40.
- 26 Yuzhou Gu and Zhao Song. A faster small treewidth SDP solver. *arXiv preprint arXiv:2211.06033*, 2022. doi:10.48550/arXiv.2211.06033.
- 27 Monika R Henzinger, Philip Klein, Satish Rao, and Sairam Subramanian. Faster shortest-path algorithms for planar graphs. *Journal of Computer and System Sciences*, 55(1):3–23, 1997. doi:10.1006/jcss.1997.1493.
- 28 Arun Jambulapati and Aaron Sidford. Ultrasparse ultrasparsifiers and faster laplacian system solvers. In Dániel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 540–559. SIAM, 2021. doi:10.1137/1.9781611976465.33.
- 29 Donggu Kang and James Payor. Flow rounding. *CoRR*, abs/1507.08139, 2015. doi:10.48550/arXiv.1507.08139.
- 30 N. Karmarkar. A new polynomial-time algorithm for linear programming. *Combinatorica*, 4(4):373–395, December 1984. doi:10.1007/BF02579150.
- 31 Yin Tat Lee and Aaron Sidford. Path finding methods for linear programming: Solving linear programs in $\tilde{O}(\sqrt{rank})$ iterations and faster algorithms for maximum flow. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 424–433. IEEE, 2014. doi:10.1109/FOCS.2014.52.
- 32 Tom Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM (JACM)*, 46(6):787–832, 1999. doi:10.1145/331524.331526.

49:14 Faster Min-Cost Flow on Bounded Treewidth Graphs

- 33 Richard J Lipton, Donald J Rose, and Robert Endre Tarjan. Generalized nested dissection. *SIAM journal on numerical analysis*, 16(2):346–358, 1979. doi:10.1137/0716027.
- 34 James Renegar. A polynomial-time algorithm, based on Newton’s method, for linear programming. *Mathematical programming*, 40(1-3):59–93, 1988. doi:10.1007/BF01580724.
- 35 Daniel A Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 81–90, 2004. doi:10.1145/1007352.1007372.
- 36 Richard Y Zhang and Javad Lavaei. Sparse semidefinite programs with near-linear time complexity. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 1624–1631. IEEE, 2018. doi:10.1109/CDC.2018.8619478.