



A Combinatorial Cut-Toggling Algorithm for Solving Laplacian Linear Systems

Monika Henzinger¹ · Billy Jin² · Richard Peng³ · David P. Williamson²

Received: 13 March 2023 / Accepted: 6 July 2023 / Published online: 1 August 2023

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2023

Abstract

Over the last two decades, a significant line of work in theoretical algorithms has made progress in solving linear systems of the form $\mathbf{L}\mathbf{x} = \mathbf{b}$, where \mathbf{L} is the Laplacian matrix of a weighted graph with weights $w(i, j) > 0$ on the edges. The solution \mathbf{x} of the linear system can be interpreted as the potentials of an electrical flow in which the resistance on edge (i, j) is $1/w(i, j)$. Kelner et al. (in: Proceedings of the 45th Annual ACM Symposium on the Theory of Computing, pp 911–920, 2013. <https://doi.org/10.1145/2488608.2488724>) give a combinatorial, near-linear time algorithm that maintains the Kirchoff Current Law, and gradually enforces the Kirchoff Potential Law by updating flows around cycles (*cycle toggling*). In this paper, we consider a dual version of the algorithm that maintains the Kirchoff Potential Law, and gradually enforces the Kirchoff Current Law by *cut toggling*: each iteration updates all potentials on one side of a fundamental cut of a spanning tree by the same amount. We prove that this dual algorithm also runs in a near-linear number of iterations. We show, however, that if we abstract cut toggling as a natural data structure problem, this problem can be reduced to the online vector–matrix–vector problem, which has been conjectured to be difficult for dynamic algorithms (Henzinger et al., in: Proceedings of the 47th Annual ACM Symposium on the Theory of Computing, pp 21–30, 2015. <https://doi.org/10.1145/2746539.2746609>). The conjecture implies that the data structure does not have an

✉ Billy Jin

bzj3@cornell.edu

Monika Henzinger

monika.henzinger@ista.ac.at

Richard Peng

y5peng@uwaterloo.ca

David P. Williamson

davidpwilliamson@cornell.edu

¹ Institute of Science and Technology Austria, Klosterneuburg, Austria

² School of Operations Research and Information Engineering, Cornell University, Ithaca, USA

³ Cheriton School of Computer Science, University of Waterloo, Waterloo, Canada

$O(n^{1-\epsilon})$ time algorithm for any $\epsilon > 0$, and thus a straightforward implementation of the cut-toggling algorithm requires essentially linear time per iteration. To circumvent the lower bound, we batch update steps, and perform them simultaneously instead of sequentially. An appropriate choice of batching leads to an $\tilde{O}(m^{1.5})$ time cut-toggling algorithm for solving Laplacian systems. Furthermore, we show that if we sparsify the graph and call our algorithm recursively on the Laplacian system implied by batching and sparsifying, we can reduce the running time to $O(m^{1+\epsilon})$ for any $\epsilon > 0$. Thus, the dual cut-toggling algorithm can achieve (almost) the same running time as its primal cycle-toggling counterpart.

Keywords Laplacian solver · Electrical flow · Data structure

1 Introduction

Over the last two decades, a significant line of work in theoretical algorithms has made progress in solving linear systems of the form $\mathbf{L}\mathbf{x} = \mathbf{b}$, where \mathbf{L} is the Laplacian matrix of a weighted graph with weights $w(i, j) > 0$ on the edges. Starting with the work of Spielman and Teng [1], researchers have devised a number of algorithms that run in near-linear time in the number of edges of the graph (corresponding to the number of non-zeros in the matrix \mathbf{L}). The solution \mathbf{x} of the linear system can be interpreted as the potentials of an electrical flow in which the resistance of each edge is $r(i, j) = 1/w(i, j)$, and the current supplied to each node i is $b(i)$. There have been many nice algorithmic ideas introduced using this interpretation of the linear system, as well as many applications of the fast algorithms for solving this system to other flow problems, such as the maximum flow problem [2–6].

Since the initial work of Spielman and Teng, a number of different near-linear time algorithms have been proposed. In this paper, we focus on a particular simple, combinatorial algorithm by Kelner et al. [7], hereafter referred to as the KOSZ algorithm; this algorithm has been the subject of several implementation studies [8–11]. The KOSZ algorithm uses the idea of an *electrical flow* in solving the linear system $\mathbf{L}\mathbf{x} = \mathbf{b}$. An electrical flow \mathbf{f} is one that obeys the flow conservation constraints at each node i , saying that the net flow out of i is $b(i)$ (sometimes known in this context as the Kirchoff Current Law, or KCL), and Ohm’s law, which says that there exists a vector \mathbf{x} such that the flow from i to j , $f(i, j)$, equals $(x(i) - x(j))/r(i, j)$. There exists a flow \mathbf{f} that satisfies these two properties, and the corresponding potential vector \mathbf{x} solves $\mathbf{L}\mathbf{x} = \mathbf{b}$. Ohm’s Law is known to be equivalent to the Kirchoff Potential Law (KPL), which says that the flow \mathbf{f} satisfies the property that around any directed cycle C , $\sum_{(i,j) \in C} r(i, j)f(i, j) = 0$. Given KPL, potentials satisfying Ohm’s law can be constructed by picking any spanning tree T rooted at a vertex r , setting $x(r)$ to 0, and $x(k)$ to the sum of $f(i, j)r(i, j)$ on the path in T from k to r ; these potentials are known as *tree-induced potentials*.

The KOSZ algorithm starts by picking a spanning tree T ; for the running time of the algorithm, it is important that the tree has low *stretch*; its definition is otherwise not crucial to the description of the algorithm. The algorithm starts by constructing a flow \mathbf{f} that satisfies flow conservation using only the edges in T . For a near-linear number

of iterations, the algorithm picks at random a non-tree edge (i, j) and considers the fundamental cycle C closed by adding (i, j) to T ; it then alters the flow \mathbf{f} along C to satisfy KPL (and such that KCL continues to be satisfied). By picking (i, j) with the appropriate probability, Kelner et al. show that the energy of the resulting flow decreases by a factor $1 - \frac{1}{\tau}$ in expectation, where τ is a parameter related to the stretch. The algorithm then returns the tree-induced potentials \mathbf{x} associated with T and the flow \mathbf{f} . Kelner et al. [7] show that the resulting potentials are close to the potentials of the associated electrical flow for the graph. The KOSZ algorithm has the pleasing properties that it is easy to understand both the algorithm and the analysis, and it is also close in spirit to known network flow algorithms; in particular, it resembles the primal network simplex algorithm for the minimum-cost flow problem.

The primal network simplex algorithm for the minimum-cost flow problem has a natural dual algorithm in the dual network simplex algorithm, in which node potentials are altered on one side of a fundamental cut of a tree (the cut induced by removing an edge in the tree). Similarly, polynomial-time cycle-canceling algorithms for minimum-cost flow (e.g. Goldberg and Tarjan [12]) have natural dual analogs of polynomial-time cut-cancelling algorithms (e.g. Ervolina and McCormick [13]). We refer to the first type of algorithm as a *cycle-toggling* algorithm, and the dual analog as a *cut-toggling* algorithm. Thus, the KOSZ algorithm is a cycle-toggling algorithm for solving Laplacian linear systems. However, no corresponding cut-toggling algorithm exists in the literature, leading immediately to the following question:

Does there exist a cut-toggling algorithm for solving Laplacian linear systems/computing near-minimum energy flows, and how efficiently can it be implemented?

1.1 Our Contributions

We propose a dual analog of the KOSZ algorithm which performs cut-toggling rather than cycle-toggling. We refer to this algorithm as Dual KOSZ, and we show it converges in a nearly-linear number of cut-toggling steps. Thus, Dual KOSZ would be a nearly-linear time algorithm if each cut-toggling operation could be implemented to run in polylogarithmic time.

Our next contribution is to show that the natural data structure abstraction of this cut-toggling process can be reduced to the online matrix–vector (OMv) problem, which has been conjectured to be hard [14]. This implies that it is unlikely for there to be a black-box data structure that implements a single cut-toggling operation in sublinear time, unlike cycle toggling.

This result initially seems to present an insurmountable difficulty to obtaining a nearly-linear time algorithm. However, we show that we can exploit the offline nature of the cut toggles, obtaining an exact data structure for computing a sequence of K cut-toggles in $O(K\sqrt{m})$ time total, which yields an algorithm that runs in $\tilde{O}(m^{1.5})$ time overall. Interestingly, Boman, Dewese, and Gilbert [9] explored an implementation of KOSZ that also batched its cycle-toggling updates by looking for collections of edge-disjoint cycles.

We further show that by incorporating sparsification and its associated approximations, we can reduce the running time to almost-linear, which means the cut toggling algorithm can still be implemented in almost-linear time.

Our result demonstrates that graph optimization algorithms and dynamic graph data structures can—and sometimes need to—interact in more intricate fashion than optimization in the outer loop and black-box data structure in the inner loop.

The remainder of the paper is structured as follows. In Sect. 1.2, we give a high-level overview of Dual KOSZ and our implementation ideas to obtain almost linear time. Section 2 describes some notation and concepts we will use. Section 3 gives the Dual KOSZ in detail and shows that it can be implemented in a near-linear number of cut-toggling iterations, with each iteration running in linear time. In Sect. 4, we abstract the problem of toggling a cut to a data structure problem, and show that given the OMv conjecture of [14], we cannot implement the operations needed in sublinear time. In Sect. 5, we show how to overcome this difficulty by batching the cut-toggle operations, and further speed up the algorithm through sparsification and recursion.

1.2 Technical Overview

As the KOSZ algorithm maintains a flow \mathbf{f} and implicitly tree-induced potentials \mathbf{x} , its natural dual is to maintain potentials \mathbf{x} , which implicitly define a flow \mathbf{f}^1 . The Dual KOSZ algorithm starts by choosing a low-stretch spanning tree T . It maintains a set of potentials \mathbf{x} (initially zero), and the corresponding (infeasible) flow \mathbf{f} implied by Ohm’s Law. In each iteration, we sample a fundamental cut S of the tree T and perform a cut-toggling update so that the net flow leaving S is $\sum_{i \in S} b(i)$, as required in every feasible flow. Following arguments dual to those made in Kelner et al. we show that this algorithm also performs a near-linear number of iterations in order to find a near-optimal set of potentials \mathbf{x} and flow \mathbf{f} .

Theorem 1 *Let τ be the total stretch of T . After $K = \tau \ln(\frac{\tau}{\epsilon})$ iterations, Dual KOSZ returns $\mathbf{x}^K \in \mathbb{R}^V$ and $\mathbf{f}^K \in \mathbb{R}^{\bar{E}}$ such that $\mathbb{E} \|\mathbf{x}^* - \mathbf{x}^K\|_{\mathbf{L}}^2 \leq \frac{\epsilon}{\tau} \|\mathbf{x}^*\|_{\mathbf{L}}^2$ and $\mathbb{E}[\mathcal{E}(\mathbf{f}^K)] \leq (1 + \epsilon)\mathcal{E}(\mathbf{f}^*)$, for \mathbf{f}^* and \mathbf{x}^* optimal primal and dual solutions respectively.*

Here $\|\mathbf{y}\|_{\mathbf{L}} = \sqrt{\mathbf{y}^T \mathbf{L} \mathbf{y}}$, and $\mathcal{E}(\mathbf{f})$ is the energy of flow \mathbf{f} .

However, unlike Kelner et al., we cannot show that each individual cut-toggling update can be made to run in polylogarithmic time. If we abstract the desired cut-toggling update step as a natural data structure problem, we show that such a data structure cannot be implemented in $O(n^{1-\epsilon})$ time for any $\epsilon > 0$ given a conjecture about the *online matrix–vector multiplication problem (OMv)* made by Henzinger, Krinninger, Nanongkai and Saranurak [14]. They have conjectured that this problem does not have any algorithm that can carry out an online sequence of n Boolean matrix–vector multiplications in time $O(n^{3-\epsilon})$, and show that if the conjecture is false, then various long-standing dynamic graph problems will have faster algorithms. We show that a single Boolean matrix–vector multiply can be carried out as a sequence of $O(n)$

¹ To make matters somewhat confusing, the Laplacian solver literature treats the space of potentials as *primal* due to its origins in numerical analysis.

operations of our desired data structure. Given the conjecture, then, we cannot implement the data structure operations in $O(n^{1-\epsilon})$ time. Thus there is not a straightforward near-linear time version of the Dual KOSZ algorithm.²

In a bit more detail, the data structure we define is given as input an undirected graph with a spanning tree T , edge resistances and a supply $b(v)$ and a potential value $\phi(v)$ at every vertex v . The data structure supports two operations: 1) additively update the value of all the vertices in a given subtree of T , and 2) query the value of the flow induced by the potential values across any given fundamental cut of T . We show that even if the data structure can only decide whether the value of the flow is non-zero or not, it would refute the OMV conjecture. Therefore, we obtain a lower bound for this data structure conditional on the OMV conjecture [14]. This even holds if all edge resistances are 1 and all $b(v)$ are 0.

Nevertheless, we circumvent this data structural lower bound by exploiting the fact that the sequence of cuts to be updated can be sampled in advance and, thus, the updates can be batched, circumventing the “online” (or “sequential”) requirement in OMV. This is possible because both the spanning tree T and the probability distribution over cuts of T are fixed at the beginning of the algorithm. More precisely, denote the number of iterations of Dual KOSZ by K (which is $\tilde{O}(m)$). Instead of sampling the fundamental cuts one at a time, consider sampling the next l cuts that need to be updated for some $l \ll K$. In each “block” of size $l \ll K$, we contract all the edges of T that do not correspond to one of the l fundamental cuts to be updated. In this way, we work with a contracted tree of size $O(l)$ in each block (instead of the full tree, which has size $O(n)$). This makes the updates faster. However, the price we pay is that at the end of each block, we need to propagate the updates we made (which were on the contracted tree), back to the entire tree. Overall, we show that each block takes $O(l^2 + m)$ time. Since there are $\tilde{O}(\frac{m}{l})$ blocks, the total runtime is $\tilde{O}(ml + \frac{m^2}{l})$. Choosing $l = \sqrt{m}$ thus gives a $\tilde{O}(m^{1.5})$ time algorithm.

By augmenting the batching idea with sparsification and recursion, one can further improve the running time of Dual KOSZ to $\tilde{O}(m^{1+\delta})$ for any $\delta > 0$. To do this, observe that l cut-toggling updates effectively break the spanning tree into $l + 1$ components. After contracting the components to get a graph H with $l + 1$ vertices, we can show that solving an appropriate Laplacian system on H gives a single update step that makes at least as much progress as the sequence of l updates performed by the straightforward unbatched algorithm. A natural approach is to solve this Laplacian system by recursively calling the algorithm. However, this by itself does not give an improved running time. Instead, we first spectrally sparsify H and then call the algorithm recursively to solve the *sparsified* Laplacian system. Here we use a combinatorial spectral sparsifier [16] because it does not require calling Laplacian solvers as a subroutine (e.g. [17]). By carefully analyzing the error incurred by sparsification, we are able to show that the update step using sparsification makes about as much progress as the update step without sparsification. The total running time of the recursive algorithm is then obtained by bounding the time taken at each layer of the recursion tree.

² In a personal communication, Sherman [15] said he also had worked out a dual version of the KOSZ algorithm, but was unable to solve the data structure problem for the updates to potentials. Our result explains why this might be difficult to do.

Theorem 2 For any $\delta \in (0, 1)$ and $\epsilon > 0$, Dual KOSZ with batching, sparsification, and recursion finds \mathbf{x} with $\mathbb{E} \|\mathbf{x}^* - \mathbf{x}\|_{\mathbf{L}}^2 \leq \epsilon \|\mathbf{x}^*\|_{\mathbf{L}}^2$ in $O(m^{1+\delta} (\log n)^{\frac{8}{\delta}} (\log \frac{1}{\epsilon})^{\frac{1}{\delta}})$ time.

2 Notation and Problem Statement

In this section, we give some notation and define concepts we will use in our algorithm. We are given as input an undirected graph $G = (V, E)$, with positive resistances $\mathbf{r} \in \mathbb{R}_+^E$. Although our graph is undirected, it will be helpful for us notationally to direct the edges. To that end, fix an arbitrary orientation of the edges, and denote the set of directed edges by \vec{E} .

In addition to G and the resistances \mathbf{r} , we are given a *supply vector* $\mathbf{b} \in \mathbb{R}^V$.

The *Laplacian matrix* of G with respect to the resistances \mathbf{r} is the matrix $\mathbf{L} \in \mathbb{R}^{V \times V}$ defined by

$$\mathbf{L} = \sum_{ij \in E} \frac{1}{r(i, j)} (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^\top,$$

where \mathbf{e}_i is the i th unit basis vector. We note then that $\mathbf{x}^\top \mathbf{L} \mathbf{x} = \sum_{(i, j) \in E} \frac{1}{r(i, j)} (x(i) - x(j))^2$ for any vector \mathbf{x} .

Our goal is to solve the system of linear equations $\mathbf{L} \mathbf{x} = \mathbf{b}$ for \mathbf{x} . However, we will not be able to solve $\mathbf{L} \mathbf{x} = \mathbf{b}$ exactly, so we will solve it approximately instead. It is usual to measure the quality of a solution \mathbf{x} in terms of the *matrix norm* induced by \mathbf{L} . In other words, if \mathbf{x} is the vector of potentials returned by our algorithm and \mathbf{x}^* is an actual solution to $\mathbf{L} \mathbf{x}^* = \mathbf{b}$, then the *error* of our solution \mathbf{x} is

$$\|\mathbf{x}^* - \mathbf{x}\|_{\mathbf{L}}^2 := (\mathbf{x}^* - \mathbf{x})^\top \mathbf{L} (\mathbf{x}^* - \mathbf{x}).$$

Hence, our objective is to find $\mathbf{x} \in \mathbb{R}^V$ that minimizes $\|\mathbf{x}^* - \mathbf{x}\|_{\mathbf{L}}^2$. A precise statement of this is given below.

Goal: Given $\epsilon > 0$, find potentials $\mathbf{x} \in \mathbb{R}^V$ that satisfy $\|\mathbf{x}^* - \mathbf{x}\|_{\mathbf{L}}^2 \leq \epsilon \|\mathbf{x}^*\|_{\mathbf{L}}^2$.

Of course, the algorithm does not know the actual solution \mathbf{x}^* . The place where \mathbf{x}^* appears is in the analysis.

Equations of the form $\mathbf{L} \mathbf{x} = \mathbf{b}$, where \mathbf{L} is the Laplacian matrix of a graph, are called Laplacian systems, and are found in a wide variety of applications in computer science and other fields. When interpreted as an optimization problem, solving a Laplacian linear system has the following nice interpretation: it is the dual of the problem of finding an electrical flow. The primal problem below is that of finding an electrical flow; here \mathbf{A} is the vertex-arc incidence matrix of (V, \vec{E}) . The optimal solution to the primal is called the *electrical flow* in G defined by the resistances \mathbf{r} . The dual problem is equivalent to solving $\mathbf{L} \mathbf{x} = \mathbf{b}$; this fact can be seen by setting the gradient of the

dual objective to equal 0 since the dual is concave.

$$\begin{aligned}
 (P) \quad & \min \quad \frac{1}{2} \sum_{(i,j) \in \bar{E}} r(i,j) f(i,j)^2 & (D) \quad & \max \quad \mathbf{b}^\top \mathbf{x} - \frac{1}{2} \mathbf{x}^\top \mathbf{L} \mathbf{x} \\
 \text{s.t.} \quad & \mathbf{A} \mathbf{f} = \mathbf{b} & \text{s.t.} \quad & \mathbf{x} \in \mathbb{R}^V \text{ unconstrained}
 \end{aligned}$$

Let $\mathcal{E}(\mathbf{f})$ denote the primal objective and $\mathcal{B}(\mathbf{x})$ denote the dual objective. The primal objective is sometimes referred to as the *energy* of the electrical flow \mathbf{f} . Also, let \mathbf{f}^* denote the optimal primal solution and let \mathbf{x}^* denote an optimal dual solution. Note that there are infinitely many dual solutions, because the dual objective is invariant under adding a constant to every component of \mathbf{x} . By strong duality (note that Slater’s condition holds), $\mathcal{E}(\mathbf{f}^*) = \mathcal{B}(\mathbf{x}^*)$. Moreover, the KKT conditions give a primal-dual characterization of optimality.

Fact 1 (KKT Conditions for Electrical Flow) *Consider $\mathbf{f} \in \mathbb{R}^{\bar{E}}$ and $\mathbf{x} \in \mathbb{R}^V$. Then \mathbf{f} is optimal for the primal and \mathbf{x} is optimal for the dual if and only if the following conditions hold:*

1. \mathbf{f} is a feasible \mathbf{b} -flow;
2. (Ohm’s Law) $f(i,j) = \frac{x(i)-x(j)}{r(i,j)}$ for all $(i,j) \in \bar{E}$.

Thus solving $\mathbf{L} \mathbf{x} = \mathbf{b}$ and finding the electrical \mathbf{b} -flow in G are equivalent: Given a solution \mathbf{x} to $\mathbf{L} \mathbf{x} = \mathbf{b}$, we can calculate the electrical flow \mathbf{f} using $f(i,j) = \frac{x(i)-x(j)}{r(i,j)}$. On the other hand, given the electrical flow \mathbf{f} , we can recover corresponding potentials \mathbf{x} by setting $x(v) = 0$ for some arbitrary vertex v , and using the equation $f(i,j) = \frac{x(i)-x(j)}{r(i,j)}$ to solve for the potentials on every other vertex.

A \mathbf{b} -flow satisfies Ohm’s Law if and only if it satisfies the Kirchoff Potential Law (KPL): KPL states that for every directed cycle C , $\sum_{(i,j) \in C} f(i,j)r(i,j) = 0$.

Both the Kelner et al. algorithm and our algorithm use a low-stretch spanning tree T . Given resistances \mathbf{r} , the *stretch* of a tree is defined as

$$\text{st}_T(G) = \sum_{(i,j) \in \bar{E}} \text{st}_T(i,j) = \sum_{(i,j) \in \bar{E}} \frac{1}{r(i,j)} \sum_{(k,l) \in P(i,j)} r(k,l),$$

where $P(i,j)$ is the unique path from i to j in T . We can find a spanning tree T with stretch $\text{st}_T(G) = O(m \log n \log \log n)$ in $O(m \log n \log \log n)$ time [18].

We use the notation $\mathbb{1}$ to stand for the vector of all 1 s, and $\mathbb{1}_X$ to be the characteristic vector of a set X that has 1 s in the entries corresponding to the elements of X and 0 s elsewhere.

3 A Cut-Toggling Algorithm for Solving Laplacian Linear Systems

We present a cut-toggling algorithm for computing an approximate solution to $\mathbf{L} \mathbf{x} = \mathbf{b}$, and also an approximate minimum-energy \mathbf{b} -flow. The goal of this section is to show that the cut-toggling algorithm converges in a near-linear number of iterations, and

Algorithm 1 The KOSZ algorithm for solving $Lx = b$.

- 1: Compute a tree T with low stretch with respect to resistances r
- 2: Find flow f^0 in T satisfying supplies b
- 3: Let x^0 be tree-defined potentials for f^0 with respect to tree T
- 4: **for** $t \leftarrow 1$ to K **do**
- 5: Pick an $(i, j) \in E - T$ with probability P_{ij}
- 6: Update f^{t-1} to satisfy KPL on the fundamental cycle closed by (i, j)
- 7: Let f^t be resulting flow
- 8: Let x^t be tree-defined potentials for f^t
- 9: **end for**
- 10: **return** f^K, x^K

that each iteration runs in linear time. Later in Sect. 5, we will show how to speed up the algorithm to an almost-linear total running time. Since our algorithm is dual to the cycle-toggling algorithm of Kelner et al. [7] (which we call KOSZ in this paper), we will begin by describing the KOSZ algorithm.

The KOSZ algorithm works by maintaining a feasible b -flow f , and iteratively updates f along cycles to satisfy Kirchkoff’s Potential Law on the cycle. It starts by choosing a spanning tree T that has low stretch, and computes a b -flow f^0 that uses only edges in the tree T . Then for a number of iterations K that depends on the stretch of the tree, it chooses a non-tree edge $(i, j) \in E - T$ according to a probability distribution, and for the fundamental cycle closed by adding edge (i, j) to T , it modifies the flow f so that Kirchkoff’s Potential Law is satisfied on the cycle. The probability P_{ij} that edge (i, j) gets chosen is proportional to the total resistance around the cycle closed by (i, j) divided by $r(i, j)$. Given the tree T with root r and the current flow f^t in iteration t , there is a standard way to define a set of potentials x^t (called the *tree-induced* or *tree-defined* potentials): set $x(r)$ to 0, and $x(k)$ to the sum of $f(i, j)r(i, j)$ on the path in T from k to r . We summarize KOSZ in Algorithm 1.

Our algorithm, which we will call Dual KOSZ, works by maintaining a set of potentials x . It iteratively samples cuts in the graph, updating potentials on one side of the cut to satisfy flow conservation across that cut. Following KOSZ, we choose a spanning tree T of low stretch. Then for a number of iterations K that depends on the stretch of tree T , we repeatedly sample a fundamental cut from the spanning tree (i.e. a cut induced by removing one of the tree edges). We update all of the potentials on one side of the cut by an amount Δ so that the amount of flow crossing the cut via Ohm’s Law is what is required by the supply vector. We summarize Dual KOSZ in Algorithm 2. The main result of this section is a bound on the iteration complexity of Dual KOSZ.

Theorem 1 *Let τ be the total stretch of T . After $K = \tau \ln(\frac{\tau}{\epsilon})$ iterations, Dual KOSZ returns $x^K \in \mathbb{R}^V$ and $f^K \in \mathbb{R}^{\bar{E}}$ such that $\mathbb{E} \|x^* - x^K\|_L^2 \leq \frac{\epsilon}{\tau} \|x^*\|_L^2$ and $\mathbb{E}[\mathcal{E}(f^K)] \leq (1 + \epsilon)\mathcal{E}(f^*)$, for f^* and x^* optimal primal and dual solutions respectively.*

Next we give the algorithm in somewhat more detail. Let $R(C) = (\sum_{(k,l) \in \delta(C)} \frac{1}{r(k,l)})^{-1}$ for $C \subset V$, where $\delta(C)$ is the set of edges with exactly one endpoint in C . Note that $R(C)$ has units of resistance. For every tree edge (i, j) , let $C(i, j)$ be the set of vertices on one side of the fundamental cut defined by (i, j) , such that $i \in C(i, j)$ and $j \notin C(i, j)$. We set up a probability distribution P_{ij} on edges (i, j) in the spanning

Algorithm 2 Algorithm Dual KOSZ for solving $\mathbf{Lx} = \mathbf{b}$.

- 1: Compute a spanning tree T with low stretch with respect to resistances \mathbf{r}
- 2: Set $x^0(i) = 0$ for all $i \in V$
- 3: **for** $t \leftarrow 1$ to K **do**
- 4: Pick edge $(i, j) \in T$ with probability $P_{ij} \propto \frac{r(i,j)}{R(C(i,j))}$ and let $C = C(i, j)$
- 5: $\Delta^t \leftarrow (b(C) - f^t(C)) \cdot R(C)$
- 6: $x^{t+1}(v) \leftarrow \begin{cases} x^t(v) + \Delta^t, & \text{if } v \in C, \\ x^t(v), & \text{if } v \notin C. \end{cases}$
- 7: **end for**
- 8: Let \mathbf{f}^K be the tree-defined flow with respect to \mathbf{x}^K and T
- 9: **return** $\mathbf{x}^K, \mathbf{f}^K$

tree T , where $P_{ij} \propto \frac{r(i,j)}{R(C(i,j))}$. We initialize potentials $x^0(i)$ to 0 for all nodes $i \in V$. In each iteration, we sample edge $(i, j) \in T$ according to the probabilities P_{ij} . Let $b(C) = \mathbf{b}^\top \mathbb{1}_C$ be the total supply of the nodes in C . Note that $b(C)$ is also the amount of flow that should be flowing out of C in any feasible \mathbf{b} -flow.

Let $f^t(C)$ be the total amount of flow going out of C in the flow induced by \mathbf{x}^t . That is,

$$f^t(C) = \sum_{\substack{ij \in E \\ i \in C, j \notin C}} \frac{x^t(i) - x^t(j)}{r(i, j)}.$$

Note that $f^t(C)$ can be positive or negative. In any feasible \mathbf{b} -flow, the amount of flow leaving C should be equal to $\mathbf{b}^\top \mathbb{1}_C = b(C)$. Hence, we define $\Delta^t = (b(C) - f^t(C)) \cdot R(C)$. Observe that Δ^t is precisely the quantity by which we need to increase the potentials of every node in C so that flow conservation is satisfied on $\delta(C)$. We then update the potentials, so that

$$p^{t+1}(v) = \begin{cases} p^t(v) + \Delta^t, & \text{if } v \in C, \\ p^t(v), & \text{if } v \notin C. \end{cases}$$

After K iterations, we return the final potentials \mathbf{x}^K . The last step is to convert \mathbf{x}^K to a feasible flow by taking a *tree-defined flow* with respect to T : $f^K(i, j) = \frac{x^K(i) - x^K(j)}{r(i,j)}$ on all non-tree edges, and \mathbf{f}^K routes the unique flow on T to make \mathbf{f}^K a feasible \mathbf{b} -flow.

3.1 Analysis of Dual KOSZ

Recall that $\mathcal{E}(\mathbf{f}) = \frac{1}{2} \sum_{e \in E} f(e)^2$ and $\mathcal{B}(\mathbf{x}) = \mathbf{b}^\top \mathbf{x} - \frac{1}{2} \mathbf{x}^\top \mathbf{Lx}$. By convex duality, we have $\mathcal{B}(\mathbf{x}) \leq \mathcal{E}(\mathbf{f})$ for any $\mathbf{x} \in \mathbb{R}^V$ and \mathbf{b} -flow \mathbf{f} . Moreover, \mathbf{x} maximizes $\mathcal{B}(\mathbf{x})$ if and only if $\mathbf{Lx} = \mathbf{b}$. (See e.g. [19, Lemma 8.9]). Thus solving the Laplacian system $\mathbf{Lx} = \mathbf{b}$ is equivalent to finding a vector of potentials that maximizes the dual objective. In what follows, we present the lemmas that form the bulk of the analysis. Their proofs are deferred to the ‘‘Appendix’’. These lemmas (and their proofs) are similar to their

counterparts in Kelner et al. [7], because everything that appears here is dual to what appears there.

First, we show that each iteration of the algorithm increases $\mathcal{B}(\mathbf{x})$.

Lemma 1 *Let $\mathbf{x} \in \mathbb{R}^V$ be a vector of potentials and let $C \subset V$. Let \mathbf{x}' be the potentials obtained from \mathbf{x} as in the algorithm (that is, by adding Δ to the potential of every vertex in C so that flow conservation is satisfied across $\delta(C)$). Then*

$$\mathcal{B}(\mathbf{x}') - \mathcal{B}(\mathbf{x}) = \frac{\Delta^2}{2R(C)}.$$

The second ingredient in the analysis is to introduce an upper bound on how large the potential bound $\mathcal{B}(\mathbf{x})$ can become. This will allow us to bound the number of iterations the algorithm takes.

Definition 1 (Gap) *Let \mathbf{f} be a feasible \mathbf{b} -flow and let \mathbf{x} be any vertex potentials. Define*

$$\text{gap}(\mathbf{f}, \mathbf{x}) := \mathcal{E}(\mathbf{f}) - \mathcal{B}(\mathbf{x}) = \frac{1}{2} \sum_{e \in E} r(e) f(e)^2 - \left(\mathbf{b}^\top \mathbf{x} - \frac{1}{2} \mathbf{x}^\top \mathbf{L} \mathbf{x} \right).$$

This same notion of a gap was introduced in the analysis of the Kelner et al. algorithm, and was also used to bound the number of iterations of the algorithm.

The electrical flow \mathbf{f}^* minimizes $\mathcal{E}(\mathbf{f})$ over all \mathbf{b} -flows \mathbf{f} , and the corresponding vertex potentials \mathbf{x}^* maximize $\mathcal{B}(\mathbf{x}^*)$ over all vertex potentials \mathbf{x} . Moreover, $\mathcal{E}(\mathbf{f}^*) = \mathcal{B}(\mathbf{x}^*)$. Therefore, for any feasible flow \mathbf{f} , $\text{gap}(\mathbf{f}, \mathbf{x})$ is an upper bound on optimality:

$$\text{gap}(\mathbf{f}, \mathbf{x}) \geq \mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}).$$

The lemma below gives us another way to write $\text{gap}(\mathbf{f}, \mathbf{x})$, and will be useful to us later. This relation is shown in Kelner et al. [7, Lemma 4.4].

Lemma 2 *We have $\text{gap}(\mathbf{f}, \mathbf{x}) = \frac{1}{2} \sum_{(i,j) \in \bar{E}} r(i,j) \left(f(i,j) - \frac{x(i)-x(j)}{r(i,j)} \right)^2$.*

The analysis of Kelner et al. [7] relies on measuring progress in terms of the above-defined duality gap between primal flow energy and dual potential bound. The high-level idea of the analysis is that one can show that the duality gap decreases by a constant factor each iteration, which implies a linear convergence rate. In the analysis of their algorithm, they maintain a feasible \mathbf{b} -flow \mathbf{f} at each iteration, and measure $\text{gap}(\mathbf{f}, \mathbf{x})$ against corresponding tree-defined potentials \mathbf{x} .

One difference between their algorithm and ours is that we do not maintain a feasible \mathbf{b} -flow at each iteration. However, for $\text{gap}(\mathbf{f}, \mathbf{x})$ to be a valid bound on distance to optimality, we need \mathbf{f} to be a feasible \mathbf{b} -flow. To this end, we introduce the definition of “tree-defined flow” below.

Definition 2 (Tree-defined flow) *Let T be a spanning tree, $\mathbf{x} \in \mathbb{R}^V$ vertex potentials, and $\mathbf{b} \in \mathbb{R}^V$ satisfying $\mathbb{1}^\top \mathbf{b} = 0$ be a supply vector. The **tree-defined flow** with*

respect to T , \mathbf{x} and \mathbf{b} is the flow $\mathbf{f}_{T,\mathbf{x}}$ defined by

$$\mathbf{f}_{T,\mathbf{x}}(i, j) = \frac{x(i) - x(j)}{r(i, j)} \quad \text{if } (i, j) \notin T,$$

and for $(i, j) \in T$, $f_{T,\mathbf{x}}(i, j)$ is the unique value such that the resulting $\mathbf{f}_{T,\mathbf{x}}$ is a feasible \mathbf{b} -flow. That is, for $(i, j) \in T$, if $C = C(i, j)$ is the fundamental cut defined by (i, j) and $b(C) = \mathbf{b}^\top \mathbb{1}_C$ is the amount of flow that should be flowing out of C in a feasible \mathbf{b} -flow, then

$$f_{T,\mathbf{x}}(i, j) = b(C) - \sum_{\substack{k \in C, l \notin C \\ kl \in E - ij}} f_{T,\mathbf{x}}(k, l) = b(C) - \sum_{\substack{k \in C, l \notin C \\ kl \in E - ij}} \frac{x(k) - x(l)}{r(k, l)}.$$

In other words, $\mathbf{f}_{T,\mathbf{x}}$ is a potential-defined flow outside of the tree T , and routes the unique flow on T to make it a feasible \mathbf{b} -flow.

The below lemma expresses $\text{gap}(\mathbf{f}_{T,\mathbf{x}}, \mathbf{x})$ in a nice way.

Lemma 3 *Let T be a spanning tree, \mathbf{x} vertex potentials, and \mathbf{b} a supply vector. Let $\mathbf{f}_{T,\mathbf{x}}$ be the associated tree-defined flow. Then*

$$\text{gap}(\mathbf{f}_{T,\mathbf{x}}, \mathbf{x}) = \frac{1}{2} \sum_{(i,j) \in T} r(i, j) \cdot \frac{\Delta(C(i, j))^2}{R(C(i, j))^2}.$$

Suppose we have a probability distribution $(P_{ij} : (i, j) \in T)$ on the edges in T . If the algorithm samples an edge $(i, j) \in T$ from this distribution, then by Lemma 1 the expected increase in the dual objective is

$$\mathbb{E}[\mathcal{B}(\mathbf{x}')] - \mathcal{B}(\mathbf{x}) = \frac{1}{2} \sum_{(i,j) \in T} P_{ij} \cdot \Delta(C(i, j))^2 / R(C(i, j)).$$

We want to set the P_{ij} to cancel terms appropriately so that the right-hand side is a multiple of the gap. Looking at Lemma 3, we see that an appropriate choice is to set

$$P_{ij} := \frac{1}{\tau} \cdot \frac{r(i, j)}{R(C(i, j))},$$

where $\tau := \sum_{(i,j) \in T} \frac{r(i,j)}{R(C(i,j))}$ is the normalizing constant. For this choice of probabilities,

$$\mathbb{E}[\mathcal{B}(\mathbf{x}')] - \mathcal{B}(\mathbf{x}) = \frac{1}{2\tau} \sum_{(i,j) \in T} r(i, j) \cdot \frac{\Delta(C(i, j))^2}{R(C(i, j))^2} = \frac{1}{\tau} \text{gap}(\mathbf{f}_{T,\mathbf{x}}, \mathbf{x}),$$

where $\mathbf{f}_{T,\mathbf{x}}$ is the tree-defined flow associated with potentials \mathbf{x} . As a consequence, we have:

Lemma 4 *If each iteration of the algorithm samples an edge $(i, j) \in T$ according to the probabilities $P_{ij} = \frac{1}{\tau} \cdot \frac{r(i,j)}{R(C(i,j))}$, then we have*

$$\mathcal{B}(\mathbf{x}^*) - \mathbb{E}[\mathcal{B}(\mathbf{x}^{t+1})] \leq \left(1 - \frac{1}{\tau}\right) (\mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}^t)).$$

Corollary 1 *After $K = \tau \ln(\frac{1}{\epsilon})$ iterations, $\mathcal{B}(\mathbf{x}^*) - \mathbb{E}[\mathcal{B}(\mathbf{x}^K)] \leq \epsilon \cdot \mathcal{B}(\mathbf{x}^*)$.*

We now use the previous lemmas to bound the number of iterations Dual KOSZ takes. Lemma 4 shows that the quantity $\mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}^t)$ decreases multiplicatively by $(1 - \frac{1}{\tau})$ each iteration. Thus, a smaller value of τ gives faster progress. Moreover, it is not difficult to show that $\tau = \text{st}_T(G, \mathbf{r})$ (see Lemma 12 in the ‘‘Appendix’’), which is why T is chosen to be a low-stretch spanning tree.

We also need to argue that rounding \mathbf{x}^K to \mathbf{f}^K via a tree-defined flow preserves approximate optimality. One can show that for any distribution over \mathbf{x} such that $\mathbb{E}_{\mathbf{x}}[\mathcal{B}(\mathbf{x})] \geq (1 - \frac{\epsilon}{\tau})\mathcal{B}(\mathbf{x}^*)$, we have $\mathbb{E}_{\mathbf{x}}[\mathcal{E}(\mathbf{f}_T, \mathbf{x})] \leq (1 + \epsilon)\mathcal{E}(\mathbf{f}^*)$. Combining everything together, we conclude:

Theorem 1 *Let τ be the total stretch of T . After $K = \tau \ln(\frac{\tau}{\epsilon})$ iterations, Dual KOSZ returns $\mathbf{x}^K \in \mathbb{R}^V$ and $\mathbf{f}^K \in \mathbb{R}^{\bar{E}}$ such that $\mathbb{E} \|\mathbf{x}^* - \mathbf{x}^K\|_{\mathcal{L}}^2 \leq \frac{\epsilon}{\tau} \|\mathbf{x}^*\|_{\mathcal{L}}^2$ and $\mathbb{E}[\mathcal{E}(\mathbf{f}^K)] \leq (1 + \epsilon)\mathcal{E}(\mathbf{f}^*)$, for \mathbf{f}^* and \mathbf{x}^* optimal primal and dual solutions respectively.*

We end this section with a naïve bound on the running time of Dual KOSZ.

Lemma 5 *Dual KOSZ can be implemented to run in $\tilde{O}(mn \log \frac{1}{\epsilon})$ time.*

Proof We can find a spanning tree T with total stretch $\tau = O(m \log n \log \log n)$ in $O(m \log n \log \log n)$ time [18].

For concreteness, fix an arbitrary vertex to be the root of T , and direct all edges in T towards the root. The set of fundamental cuts we consider will be the vertex sets of subtrees of T .

To compute $b(C)$ for these $n - 1$ fundamental cuts C , we can work our way from the leaves up to the root. If $C = \{v\}$ is a leaf of T , then $b(C) = b(v)$. Otherwise, C is a subtree rooted at v , and $b(C) = b(v) + \sum_{C'} b(C')$, where the sum is over the subtrees of C . Hence we can compute $b(C)$ for all fundamental cuts C in $O(n)$ time.

To compute $R(C)$ for the fundamental cuts C , we can maintain $n - 1$ variables, one for each fundamental cut. The variable corresponding to cut C will represent $R(C)^{-1} = \sum_{e \in \delta(C)} \frac{1}{r(e)}$, and the variables are initialized to 0. We then iterate through all the edges in the graph, and for each such edge e , add $\frac{1}{r(e)}$ to the value of each variable that represents a cut C such that $e \in \delta(C)$. Although this naive implementation takes $O(mn)$ time in the worst-case, it is possible to improve this running time to $O(m \log n)$ using link-cut trees [20]. One can also achieve this running time using the same data structure as the one used in [7].

The last part of the running time is the time it takes to run a single iteration of the algorithm. In each iteration of the algorithm, we need to compute $\Delta = (b(C) - f(C)) \cdot R(C)$, where C is the fundamental cut selected at that iteration. In the above

two paragraphs, we described how to precompute the values of $b(C)$ and $R(C)$ for every fundamental cut C ; note that these values are fixed at the beginning and do not change during the course of the algorithm. Hence, it remains to compute $f(C)$. One way to compute $f(C)$ is to simply iterate over all the edges in $\delta(C)$, summing each edge's contribution to $f(C)$. This takes time proportional to $|\delta(C)|$, which could be $O(m)$ in the worst case. We can get this down to $O(n)$ per iteration by maintaining the values of $f(C)$ for every fundamental cut C , and updating these values each time the algorithm updates potentials. Since there are $n - 1$ cuts, to do this in $O(n)$ time requires us to be able to update $f(C)$ for a single cut in $O(1)$ time. To do this, we can precompute an $(n - 1) \times (n - 1)$ table with a row/column for each fundamental cut, where the (C_1, C_2) entry is the amount by which the flow out of C_2 increases if we add 1 to the potential of every node in C_1 . Let $H(C_1, C_2)$ denote this value. With this table, updating the value of $f(C)$ after a potential update step essentially reduces to a single table lookup, which takes $O(1)$ time.

Finally, note that one can construct the $H(C_1, C_2)$ table in $O(n^2)$ time using results from [21]. In the language of Definitions 5.3 and 5.5 in that paper, we are trying to compute $C(v^\downarrow, w^\downarrow)$ for all vertices v, w in the tree, where the edge weights are the reciprocals of the resistances. At the bottom of page 11, it states that the n^2 values $C(v, w^\downarrow)$ can be computed in $O(n^2)$ time. At the top of Page 12, it then says that we get the values of $C(v^\downarrow, w^\downarrow)$ using n tree prefix sums. (Each tree prefix sum is the procedure described in Lemma 5.8, and takes $O(n)$ time.)

To summarize, we can run each iteration of the algorithm in $O(m)$ time, which can be reduced to $O(n)$ time if we precompute the $H(C_1, C_2)$ table, which incurs an overhead of $O(n^2)$ storage and $O(n^2)$ preprocessing time.

Suppose each iteration of the algorithm takes $O(I)$ time, and the algorithm uses $O(L)$ preprocessing time (not including the time needed to compute the low-stretch spanning tree). Then the total running time of the algorithm is $O(L + I\tau \ln(\frac{\tau}{\epsilon}) + m \log n \log \log n) = O(L + mI \log n \log \log n \log \frac{\tau}{\epsilon})$.

If we use the version which uses $O(n^2)$ preprocessing time and $O(n)$ time per iteration, then $L = O(n^2)$ and $I = O(n)$. This gives the running time of Dual KOSZ to be $O(mn \log n \log \log n \log \frac{\tau}{\epsilon})$. \square

In Sect. 4, we argue that given a natural abstraction of the data structure problem we use in computing $f(C)$ and updating potentials, it appears unlikely that we can implement each iteration in $o(n^{1-\epsilon})$ time, if each iteration is to be processed one-by-one in an online fashion. In Sect. 5, we show how to overcome this data structure lower bound by taking advantage of the fact that the sequence of updates that we perform can be generated in advance.

4 Lower Bound on the Per-Iteration Complexity of the Algorithm

Recall that each single iteration of KOSZ can be implemented in logarithmic time. In this section we show that assuming the OMv conjecture (see below) each single iteration of Dual KOSZ cannot be implemented in linear time. This implies that in order to speed up our algorithm we need to “batch-up” iterations, which is the approach we

use in the next section. We first present a natural data structure, called the *TreeFlow* data structure, such that each iteration of the algorithm requires only two operations of the *TreeFlow* data structure and then prove that assuming the OMv conjecture [14] it is impossible to implement the *TreeFlow* data structure such that each operation of the data structure takes $O(n^{1-\epsilon})$ time. To simplify the reduction we reduce from a closely related problem called the *Online Vector–Matrix–Vector Multiplication Problem* (*OuMv*).

Definition 3 (Online Vector–Matrix–Vector Multiplication Problem) We are given a positive integer n , and a Boolean $n \times n$ matrix \mathbf{M} . At each time step $t = 1, \dots, n$, we are shown a pair of Boolean vectors $(\mathbf{u}_t, \mathbf{v}_t)$, each of length n . Our task is to output $\mathbf{u}_t^\top \mathbf{M} \mathbf{v}_t$ using Boolean matrix–vector operations. Specifically, “addition” is replaced by the OR operation, so that $0 + 0 = 0$, and $0 + 1 = 1 + 0 = 1 + 1 = 1$. Hence, $\mathbf{u}_t^\top \mathbf{M} \mathbf{v}_t$ is always either 0 or 1.

The OMv conjecture implies that no algorithm for the OuMv problem can do substantially better than naively multiplying $\mathbf{u}_t^\top \mathbf{M} \mathbf{v}_t$ at time step t . Specifically, it says the following:

Lemma 6 ([14]) *Let $\epsilon > 0$ be any constant. Assuming the OMv conjecture, there is no algorithm for the online vector–matrix–vector multiplication problem that uses preprocessing time $O(n^{3-\epsilon})$ and takes total time $O(n^{3-\epsilon})$ with error probability at most $1/3$ in the word-RAM model with $O(\log n)$ bit words.*

Thus we will reduce the OuMv problem to the *TreeFlow* data structure such that computing $\mathbf{u}_t^\top \mathbf{M} \mathbf{v}_t$ requires two operations in the *TreeFlow* data structure. The lower bound then follows from Lemma 6.

4.1 The *TreeFlow* Data Structure

The *TreeFlow* data structure is given as input (1) an undirected graph $G = (V, E)$ with $n = |V|$, (2) a spanning tree T of G that is rooted at a fixed vertex x , (3) a value $r(u, v)$ for each edge $(u, v) \in E$ (representing the resistance of (u, v)), and (4) a value $b(v)$ for each vertex $v \in V$ (representing the supply at v). The quantities $r(u, v)$ and $b(v)$ are given at the beginning and will remain unchanged throughout the operations. For any set $C \subset V$, let $b(C) := \sum_{v \in C} b(v)$.

Furthermore, each vertex v has a non-negative value, denoted $\text{value}(v)$, which can be seen as the “potential” of v . It is initially 0 and can be modified. For any set $C \subset V$ we define the *flow out of C* to be the quantity

$$f(C) := \sum_{(u,v) \in E, u \in C, v \notin C} (\text{value}(u) - \text{value}(v)) / r(u, v).$$

The *TreeFlow* data structure supports the following operations.

- **addvalue(vertex v , real x):** Add x to the value of every vertex in the subtree of T rooted at v .

- **findflow(vertex v):** Return $b(C) - f(C)$, where C is the set of vertices in the subtree of T rooted at v .

The *TreeFlow* data structure implements exactly the operations we require for each iteration of Dual KOSZ: The **addvalue** operation allows us to update the potentials on a fundamental cut, and **findflow** computes $b(C) - f(C)$, thereby allowing us to compute Δ at each iteration. Note that if all $b(v)$ -values are zero, the *TreeFlow* data structure simply returns $-f(C)$, which gives it its name.

We even show the lower bound for a “relaxed” version defined as follows: In an α -approximate *TreeFlow* data structure the operation **addvalue** remains as above and the operation **findflow**(v) returns a value that is within a multiplicative factor $\alpha \geq 1$ (that can be a function of n) of the correct answer, i.e., a value between $(b(C) - f(C))/\alpha$ and $(b(C) - f(C)) \cdot \alpha$. The hardness of approximation is interesting, because it turns out that even an approximation of this quantity is sufficient to obtain an algorithm for $\mathbf{Lx} = \mathbf{b}$. (Albeit with a convergence rate that deteriorates with the approximation factor.)

Lemma 7 *Let $\epsilon > 0$ be any constant and let $\alpha \geq 1$ be any value. Assuming the OMv conjecture, no implementation of the α -approximate *TreeFlow* data structure exists that uses preprocessing time $O(n^{3-\epsilon})$ and where the two operations **addvalue** and **findflow** both take $O(n^{1-\epsilon})$ time, such that over a polynomial number of operations the error probability is at most $1/3$ in the word-RAM model with $O(\log n)$ bit words. This even holds if all $r(u, v)$ values are 1 and if all $b(v)$ are 0.*

Proof Given an $n \times n$ Boolean matrix \mathbf{M} , we create the following *TreeFlow* data structure. The graph contains $2n + 1$ nodes, namely a special node x , one node c_j for each column j with $1 \leq j \leq n$ and one node d_i for each row i with $1 \leq i \leq n$. There is an edge (d_i, c_j) if entry $M_{ij} = 1$. Additionally, every node c_j and every node d_i has an edge to x . These edges are added to guarantee that the graph is connected. We set $r(c, d) = 1$ for every edge (c, d) and denote this graph by G . Let T be the spanning tree of G that is rooted at x and consists of all the edges incident to x . Note that the subtree of T rooted at any node $y \neq x$ consists of a single node y .

Now consider the sequence of n vector pairs $(\mathbf{u}_t, \mathbf{v}_t)$ of the OuMv problem. Let (\mathbf{u}, \mathbf{v}) be any such pair. We show below how to compute $\mathbf{u}^\top \mathbf{M} \mathbf{v}$ with $O(n)$ operations in the *TreeFlow* data structure. Thus the sequence of n vector pairs leads to $O(n^2)$ operations. It then follows from the OMv conjecture and Lemma 6 that this sequence of $O(n^2)$ operations in the *TreeFlow* data structure cannot take time $O(n^{3-\epsilon})$, i.e., that it is not possible that the complexity of both the **addvalue** operation and the **findflow** operation are $O(n^{1-\epsilon})$.

It remains to show how to compute $\mathbf{u}^\top \mathbf{M} \mathbf{v}$ with $O(n)$ operations in the *TreeFlow* data structure. Initially the value $\text{value}(v)$ of all nodes v is 0. Let Z be a large enough constant that we will specify later. First, increase the value of all nodes to Z by calling **addvalue**(x, Z).

When given (\mathbf{u}, \mathbf{v}) we decrease the value of each row node d_i with $u_i = 1$ to 0 by calling **addvalue**($d_i, -Z$). Then, we perform a **findflow**(c_j) operation for each column node c_j with $v_j = 1$. Afterwards we undo these operations again, so that every node has value Z again. (Alternatively, we could also increase the value of every

node c_j with $v_j = 1$ and every node d_i with $u_i = 0$ to $2Z$, and instead of undoing the operations, we increase after the query the value of every node to $2Z$. In which way we never execute an `addvalue` operation with negative second parameter.)

Note that $\mathbf{u}^\top \mathbf{M} \mathbf{v} = 1$ iff there exists an edge between a column node c_j with $v_j = 1$ (i.e. $\text{value}(c_j) = Z$) and a row node d_i with $u_i = 1$ (i.e. $\text{value}(d_i) = 0$).

We now show that $\mathbf{u}^\top \mathbf{M} \mathbf{v} = 1$ iff $f(c_j) > 0$ for some column node c_j with $v_j = 1$.
 (a) Assume first that $\mathbf{u}^\top \mathbf{M} \mathbf{v} = 1$ and let c^* denote a node c_j and d^* denote a node d_i such that $v_j = 1$, $u_i = 1$ and $M_{ij} = 1$. We will show that $f(c^*) > 0$. Recall that the subtree of c^* consists only of c^* . The edge (c^*, d^*) leaves the subtree of c^* , contributing a positive amount to $f(c^*)$ because $\text{value}(c^*) = Z$ and $\text{value}(d^*) = 0$. All other edges leaving the subtree of c^* contribute a non-negative amount to $f(c^*)$, since $\text{value}(c^*) = Z$ and $\text{value}(d_k)$ for other $k \neq i$ is either Z or 0 . Thus $f(c^*) > 0$.
 (b) Assume next that $\mathbf{u}^\top \mathbf{M} \mathbf{v} = 0$. In this case every node c_j with $u_j = 1$ (and value Z) only has edges to nodes d_i with $v_i = 0$ (and value Z). As before the subtree of every node c_j only consists of c_j and, thus, all edges leaving the subtree of c_j contribute 0 to the flow out of the subtree. Thus, for every node c_j with $u_j = 1$ we have $f(c_j) = 0$.

To summarize we have shown above that $\mathbf{u}^\top \mathbf{M} \mathbf{v} = 1$ iff $f(c_j) > 0$ for some column node c_j with $\text{value}(c_j) = Z$. We will now show how to use the results of the `findflow` queries returned by an α -approximate `TreeFlow` data structure to determine if $f(c_j)$ is positive or zero.

Here is where we will choose the value of Z . The idea is to make Z large enough so that if $f(c_j) > 0$, then $f(c_j)$ is very large. The idea is that this will allow us to distinguish between $f(c_j) = 0$ versus $f(c_j) > 0$, even if we only have access to an α -approximation of $S(c_j) - f(c_j) = b(c_j) - f(c_j)$.

It will suffice to choose Z large enough so that if $f(c_j) > 0$, then $f(c_j) > \max\{b(c_j), b(c_j)(1 - \alpha^2)\}$ (As $(1 - \alpha^2) < 0$, the second term makes sense if $b(c_j) < 0$.) The value of Z depends on α , the supplies \mathbf{b} , and the resistances \mathbf{r} . For instance, it suffices to choose $Z > \|\mathbf{r}\|_\infty \|\mathbf{b}\|_\infty \alpha^2$. For this choice of Z , we have that if $f(c_j) > 0$ then (since it must have an edge to some d_i with $\text{value}(d_i) = 0$),

$$f(c_j) \geq \frac{\text{value}(c_j) - \text{value}(d_i)}{r(c_j, d_i)} = \frac{Z - 0}{r(c_j, d_i)} > |b(c_j)| \alpha^2 > \max\{b(c_j), b(c_j)(1 - \alpha^2)\}.$$

Having chosen Z this way, we have the following:

- If $b(c_j) \geq 0$, then $b(c_j) - f(c_j)$ is non-negative if $f(c_j) = 0$, and negative otherwise (because $f(c_j) > b(c_j)$ when $f(c_j) > 0$.) Any α -approximation of $b(c_j) - f(c_j)$ allows us to correctly deduce the sign of $b(c_j) - f(c_j)$, hence also whether $b(c_j) - f(c_j) \geq 0$ or whether $b(c_j) - f(c_j) < 0$. From this we can deduce whether $f(c_j) = 0$ or $f(c_j) > 0$.
- Suppose $b(c_j) < 0$. If $f(c_j) = 0$, the approximate data structure returns an answer in the interval $[b(c_j) \cdot \alpha, \frac{b(c_j)}{\alpha}]$. If $f(c_j) > 0$, it returns an answer in the interval $[(b(c_j) - f(c_j)) \cdot \alpha, \frac{b(c_j) - f(c_j)}{\alpha}]$. Note that the left endpoint of the first interval is to the right of the right endpoint of the second interval as $f(c_j) > b(c_j) (1 - \alpha^2)$

implies that

$$\implies b(c_j) \cdot \alpha > \frac{b(c_j) - f(c_j)}{\alpha}.$$

Since the two intervals for $f(c_j) = 0$ and $f(c_j) > 0$ do not overlap, we can correctly distinguish the two cases using the approximate data structure.

To summarize, each `findflow` query on c_j allows us to determine if $f(c_j) > 0$ or $f(c_j) = 0$. If the flow is positive for some c_j , then the answer is $\mathbf{u}^\top \mathbf{M}\mathbf{v} = 1$, otherwise it is 0. Note that it requires $O(n)$ operations in the *TreeFlow* data structure to determine one $\mathbf{u}^\top \mathbf{M}\mathbf{v}$ value, which completes the proof. \square

Remark 1 Note that the proof can be modified to be more similar to the update sequence generated by Dual KOSZ which alternates between `addvalue` and `findflow` operations by inserting after each `addvalue` operation a `findflow` operation (whose answer might be ignored for the logic of the proof).

As mentioned, the proof can also be adapted so that the values stored at the nodes are only increased, but this is not necessary for our application.

5 Speeding Up Dual KOSZ

We now show how to surmount the OMv lower bound by taking advantage of the fact that the sequence of updates that Dual KOSZ performs can be generated in advance. In Sect. 5.1, we show that batching the updates yields a modification of the algorithm that runs in $\tilde{O}(m^{1.5})$ time. Then in Sect. 5.2, we use sparsification and recursion to further improve the runtime to $\tilde{O}(m^{1+\alpha})$ for any $\alpha > 0$.

5.1 A Faster Algorithm Using Batching

First, we show that it is possible to speed up the running time to $\tilde{O}(m^{1.5})$ time by batching the updates performed by Dual KOSZ. In Lemma 5, we showed that the algorithm can be implemented to run in time $\tilde{O}(mn)$. (Here the tilde hides a factor of $\log n \log \log n \log \frac{1}{\epsilon}$.) This running time essentially comes from $\tilde{O}(m)$ iterations, and $O(n)$ time per iteration. Recall that each iteration of Dual KOSZ involves sampling a fundamental cut C of the low-stretch spanning tree T from a fixed probability distribution P , and then adding a constant to the potential of every vertex in C so that the resulting potential-defined flow satisfies flow conservation across C .

The main idea of batching is as follows. Denote the number of iterations by K (which is $\tilde{O}(m)$). Instead of sampling the fundamental cuts one at a time, consider sampling the next l cuts that need to be updated for some $l \ll K$. We can perform this sampling in advance because both the tree T and the probability distribution over cuts of T are fixed over the entire course of the algorithm. In each “block” of size $l \ll K$, we contract all the edges of T that do not correspond to one of the l fundamental cuts to be updated. In this way, we work with a contracted tree of size $O(l)$ in each block (instead of the full tree, which has size $O(n)$). This makes the updates faster. However,

the price we pay is that at the end of each block, we need to propagate the updates we made (which were on the contracted tree), back to the entire tree. We will show that by choosing $l = \sqrt{m}$, we can balance this tradeoff and get an improved running time of $\tilde{O}(m^{1.5})$. Pseudocode for Dual KOSZ with batching is given in Algorithm 3. Note that the correctness of this algorithm follows directly from the correctness of Dual KOSZ: Algorithm 3 samples cuts from exactly the same distribution as Dual KOSZ, and if we fix the same sequence of cuts to be used by both algorithms, then the output of the two algorithms is identical.

Algorithm 3 Dual KOSZ with batching

```

1: Compute a tree  $T$  with low stretch with respect to resistances  $\mathbf{r}$ .
2: Compute  $b(C)$  and  $R(C)$  for all fundamental cuts  $C$  of  $T$ .
3: Set  $\mathbf{x}^0(i) = 0$  for all  $i \in V$ . Set  $f(C) = 0$  for all fundamental cuts  $C$  of  $T$ .
4: for  $t \leftarrow 1$  to  $\lceil \frac{K}{l} \rceil$  do
5:   Sample  $l$  edges  $(i_1, j_1), \dots, (i_l, j_l)$  with replacement from  $T$ , according to the distribution  $P$ .
6:   Contract all edges in  $T$  that were not sampled in step 5.
7: Let  $\tilde{G}$  be the resulting graph and  $\tilde{T}$  be the resulting tree.
8:   For each  $1 \leq k \leq l$ , let  $C_k$  denote the fundamental cut in  $T$  determined by edge  $(i_k, j_k)$ . Let  $\tilde{C}_k$ 
   denote the fundamental cut in  $\tilde{T}$  determined by  $(i_k, j_k)$ .
9:    $\mathbf{y}(\tilde{v}) \leftarrow 0$  for all  $\tilde{v} \in V(\tilde{G})$ .
10:  for  $k \leftarrow 1$  to  $l$  do
11:    Compute  $\Delta_k = (b(C_k) - f(C_k)) \cdot R(C_k)$ . ▷ Requires  $f(C_k)$  to be
12:  already computed
13:     $\mathbf{y}(\tilde{v}) \leftarrow \mathbf{y}(\tilde{v}) + \Delta_k$  for all  $\tilde{v} \in \tilde{C}_k$ .
14:    Update values of  $f(C_j)$  for all  $j \in \{k + 1, \dots, l\}$ .
15:  end for
16:  for  $i \in V$  do
17:    Let  $\tilde{v}(i)$  be the vertex in  $\tilde{G}$  that  $i$  was contracted to.
18:     $\mathbf{x}^t(i) \leftarrow \mathbf{x}^{t-1}(i) + \mathbf{y}(\tilde{v}(i))$ .
19:  end for
20:  Recompute  $f(C)$  for all fundamental cuts  $C$  of  $T$ .
21: end for
22: Let  $\mathbf{f}^{\lceil K/l \rceil}$  be the tree-defined flow with respect to  $\mathbf{x}^{\lceil K/l \rceil}$  and  $T$ .
23: return  $\mathbf{x}^{\lceil K/l \rceil}, \mathbf{f}^{\lceil K/l \rceil}$ 

```

Theorem 3 *The running time of Dual KOSZ with batching is $O(m^{1.5} \log n \log \log n \log \frac{1}{\epsilon})$. This is achieved by choosing $l = \sqrt{m}$.*

Proof At the beginning of the algorithm, we compute the values of $b(C)$ ($O(n)$ time via a dynamic program) and $R(C)$ ($O(m \log n)$ time via link-cut trees), just as in the proof of Lemma 5.

Consider a batch of l updates. Note that the contracted tree \tilde{T} has at most $l + 1$ vertices. After contracting, we need to perform l updates. This involves, for each $k \in \{1, 2, \dots, l\}$:

- Computing $\Delta_k := (b(C_k) - f(C_k)) \cdot R(C_k)$,
 - This takes $O(1)$ time assuming $f(C_k)$ has already been computed. (Recall that the values $b(C_k)$ and $R(C_k)$ are computed at the very beginning of the algorithm.)

- Adding Δ_k to $\mathbf{y}(\tilde{v})$ for every $\tilde{v} \in \tilde{C}_k$.
 - This takes $O(l)$ time, because the contracted tree has size $O(l)$.
- Updating the values $f(C_{k+1}), f(C_{k+2}), \dots, f(C_l)$ so they can be used in the later iterations of the inner loop.
 - If each $f(C_j)$ can be updated in $O(1)$ time, this takes $O(l)$ time.
 - To update each $f(C_j)$ in $O(1)$ time, we can precompute at the beginning of the block the $H(C_i, C_j)$ table for $i, j \in \{1, 2, \dots, l\}$, just as in the proof of Lemma 5. The difference now is that we only need to compute the table for the cuts that will be updated in the block. There are l such cuts, so the total time to compute the table is $O(l^2)$, again using Karger’s method.

At the end of each block, we propagate the updates we made on the contracted graph back to the original graph. This involves

- Determining the new potential of each node in G .
 - This takes $O(n)$ time by iterating over the nodes of G .
- Determining the value of $f(C)$ for each fundamental cut determined by T . (By convention, assume the edges of T are directed toward the root, and that the fundamental cuts we consider are the vertex sets of the subtrees of T .)
 - This can be done in $O(m)$ time using a dynamic program that works from the leaves to the root. First, we compute $f(C)$ at each leaf of the tree. Next, suppose we have are at a non-leaf node v , and let C be the set of vertices in the subtree rooted at v . Suppose we have already computed $f(D_1), f(D_2), \dots, f(D_k)$, where D_1, \dots, D_k are the proper subtrees of v . Then we can compute $f(C)$ as follows:

$$f(C) = \sum_{i=1}^k f(D_i) + \sum_{w:vw \in E} \frac{p(v) - p(w)}{r(v, w)}.$$

This sum correctly counts the flow leaving C . This is because any edge leaving C is counted once. On the other hand, if an edge is between D_i and D_j , then it is counted once in the $f(D_i)$ term, and once with the *opposite sign* in the $f(D_j)$ term, so it zeros out. Similarly, if an edge is between D_i and v , it also zeros out.

The running time of this dynamic program is $O(m)$, because the time taken at each node is proportional to its degree, and the sum of all the node degrees is equal to $2m$.

To summarize, there are K iterations, divided into blocks of size l . In each block, we pay the following.

- **Start of block** $O(m)$ time to contract the tree, and $O(l^2)$ time to compute the $H(C, C')$ table for the cuts that will be updated in the block.
- **During the block** $O(l)$ time per iteration. Since each block consists of l iterations, this is $O(l^2)$ in total.

- **End of block** $O(m)$ time to propagate the changes from the contracted tree to the original tree.

Hence, each block takes $O(m + l^2)$ time. Multiplying by the number of blocks, which is K/l , this gives a running time of $O(K(\frac{m}{l} + l))$. Choosing $l = \sqrt{m}$ to minimize this quantity, we get $O(K\sqrt{m})$.

The final running time is therefore $O(K\sqrt{m})$ plus the preprocessing time. Preprocessing consists of finding a low-stretch spanning tree ($O(m \log n \log \log n)$), plus computing the values of $R(C)$ ($O(m \log n)$), and $f(C)$ ($O(n)$).

Thus, the preprocessing time is dominated by the time it takes to run the iterations. So, the total running time is now: $O(K\sqrt{m}) = O(m^{1.5} \log n \log \log n \log \frac{1}{\epsilon})$. \square

5.2 A Still Faster Algorithm via Batching, Sparsification, and Recursion

We now show that we can further speed up the algorithm using sparsification and recursion. The goal is to show that we can obtain a running time of the form $O(A^{\frac{1}{\delta}} m^{1+\delta} (\log n)^{\frac{B}{\delta}} (\log \frac{1}{\epsilon})^{\frac{1}{\delta}})$ for any $\delta > 0$, where A and B are numerical constants.

Consider batching the iterations of the algorithm as follows. Pick a positive integer d , and repeat K times:

- Sample the next d updates to be performed by the algorithm. These correspond to d edges of the spanning tree T .
- Let V_0, V_1, \dots, V_d be the vertex sets that T is partitioned into by the d tree edges.
- Add $\Delta(i)$ to every vertex in V_i . We will choose the values $\Delta(0), \Delta(1), \dots, \Delta(d)$ to greedily maximize the increase in the dual bound.

Note that our original algorithm corresponds to the case when $d = 1$. The lemma below quantifies the increase of the dual objective after one step of the above update.

Lemma 8 *Let (V_0, \dots, V_d) be a partition of V . Let $\mathbf{x} \in \mathbb{R}^V$ be a vector of potentials, and let $\Delta = (\Delta(0), \dots, \Delta(d))$ be any vector in \mathbb{R}^{d+1} . Let $\tilde{\mathbf{x}}$ be obtained from \mathbf{x} by adding $\Delta(i)$ to the potential of every node in V_i . Then, the increase in the dual bound is given by the formula*

$$\mathcal{B}(\tilde{\mathbf{x}}) - \mathcal{B}(\mathbf{x}) = \mathbf{b}_H^T \Delta - \frac{1}{2} \Delta^T \mathbf{L}_H \Delta,$$

where

- H is the contracted graph with vertices V_0, V_1, \dots, V_d and resistances $r(V_k, V_l) = \left(\sum_{ij \in \delta(V_k, V_l)} \frac{1}{r(i, j)} \right)^{-1}$,
- \mathbf{L}_H is the Laplacian matrix of H , and
- $b_H(k) = b(V_k) - f(V_k)$ for $k = 0, 1, \dots, d$.

In particular, the choice of Δ that maximizes $\mathcal{B}(\tilde{\mathbf{x}}) - \mathcal{B}(\mathbf{x})$ is given by the solution to $\mathbf{L}_H \Delta = \mathbf{b}_H$.

Proof We write the increase in the dual potential bound. Recall that $f(V_k)$ is the amount of flow leaving V_k in the flow $f(i, j) = \frac{x(i)-x(j)}{r(i, j)}$. We let $f(V_k, V_l)$ be the amount of flow going from V_k to V_l .

$$\begin{aligned}
 & 2(\mathcal{B}(\tilde{\mathbf{x}}) - \mathcal{B}(\mathbf{x})) \\
 &= (2\mathbf{b}^T \tilde{\mathbf{x}} - \tilde{\mathbf{x}}^T \mathbf{L} \tilde{\mathbf{x}}) - (2\mathbf{b}^T \mathbf{x} - \mathbf{x}^T \mathbf{L} \mathbf{x}) \\
 &= 2 \sum_k b(V_k) \Delta(k) + \sum_{(i, j) \in \bar{E}} \frac{1}{r(i, j)} \left[(x(i) - x(j))^2 - (\tilde{x}(i) - \tilde{x}(j))^2 \right] \\
 &= 2 \sum_k b(V_k) \Delta(k) + \sum_{(i, j) \in \bar{E}} \frac{1}{r(i, j)} \left[(x(i) - x(j) + \tilde{x}(i) - \tilde{x}(j))(x(i) - x(j) - \tilde{x}(i) + \tilde{x}(j)) \right] \\
 &= 2 \sum_k b(V_k) \Delta(k) + \sum_{k < l} \sum_{(i, j) \in \delta(V_k, V_l)} \frac{1}{r(i, j)} [(2x(i) - 2x(j) + \Delta(k) - \Delta(l))(\Delta(l) - \Delta(k))] \\
 &= 2 \sum_k b(V_k) \Delta(k) + 2 \sum_{k < l} (\Delta(l) - \Delta(k)) \sum_{(i, j) \in \delta(V_k, V_l)} \frac{1}{r(i, j)} (x(i) - x(j)) \\
 &\quad - \sum_{k < l} (\Delta(k) - \Delta(l))^2 \sum_{(i, j) \in \delta(V_k, V_l)} \frac{1}{r(i, j)} \\
 &= 2 \sum_k b(V_k) \Delta(k) + 2 \sum_{k < l} (\Delta(l) - \Delta(k)) f(V_k, V_l) - \Delta^T \mathbf{L}_H \Delta \\
 &= 2 \sum_k b(V_k) \Delta(k) - 2 \sum_k \Delta(k) f(V_k) - \Delta^T \mathbf{L}_H \Delta \\
 &= 2 \sum_k (b(V_k) - f(V_k)) \Delta(k) - \Delta^T \mathbf{L}_H \Delta \\
 &= 2\mathbf{b}_H^T \Delta - \Delta^T \mathbf{L}_H \Delta
 \end{aligned}$$

Note that this is a concave function of Δ , because \mathbf{L}_H is positive semidefinite. Therefore, maximizing this expression is equivalent to setting its gradient to 0. Taking its gradient and setting to 0 yields $\mathbf{L}_H \Delta = \mathbf{b}_H$, as claimed. \square

Remark 2 Another interpretation of the Δ that maximizes $\mathcal{B}(\tilde{\mathbf{x}}) - \mathcal{B}(\mathbf{x})$ in the Lemma above is as follows: $(\Delta(0), \dots, \Delta(d))$ are the values such that if one adds $\Delta(i)$ to the potential of every vertex in V_i , the resulting potential-induced flow satisfies the flow constraints $f(V_k) = b(V_k)$ for all $k = 0, \dots, d$.

5.3 The Sparsify and Recurse Algorithm

Next we give the algorithm with sparsification and recursion in more detail. Observe that d cut-toggling updates effectively break the spanning tree into $d + 1$ components. After contracting the components to get a graph H with $d + 1$ vertices, Lemma 8 shows that solving the Laplacian system $\mathbf{L}_H \Delta = \mathbf{b}_H$ gives the update that maximizes the increase in $\mathcal{B}(\tilde{\mathbf{x}}) - \mathcal{B}(\mathbf{x})$ among all updates that increment the potential of all vertices in V_i by the same amount. In particular, the progress made by this update step is at least as large as the progress made by the sequence of d updates performed by the straightforward unbatched algorithm.

A natural approach is to solve $\mathbf{L}_H \Delta = \mathbf{b}$ recursively. However, this by itself does not give an improved running time. Instead, we will first spectrally sparsify H to get a sparsified approximation $\tilde{\mathbf{L}}_H$ of \mathbf{L}_H , satisfying $(1 - \gamma)\mathbf{L}_H \preceq \tilde{\mathbf{L}}_H \preceq (1 + \gamma)\mathbf{L}_H$ for an appropriate constant $\gamma \in (0, 1)$. Such a matrix $\tilde{\mathbf{L}}_H$ is known as a γ -spectral sparsifier of \mathbf{L}_H . We then call the algorithm recursively on H to solve $\tilde{\mathbf{L}}_H \tilde{\Delta} = \mathbf{b}_H$. This approach is akin to recursive sparsification based Laplacian solvers [22–24], but whose progress is bounded via the effect of single cut toggles. A main task of the analysis is to bound the error incurred by solving the sparsified system instead of the exact one.

For the spectral sparsification, we need to use an algorithm that does not require calling Laplacian solvers as a subroutine. For examples of such algorithms, see [16, 25–28]. We will use the sparsifier given in Theorem 1.2 of [16], which returns a γ -spectral sparsifier with $O(n \log^2 n \log \log n / \gamma^2)$ edges in $O(m \log^2 n \log \log n / \gamma^2)$ time, with probability at least $1 - \frac{1}{\text{poly}(n)}$. Here, m, n are the number of edges and vertices, respectively, in the graph before sparsifying, and the $1 - \frac{1}{\text{poly}(n)}$ probability means that one can make the probability $1 - \frac{1}{n^k}$ for any constant k , with the k appearing as a multiplicative factor in the big- O expressions. We will take $k = 2$ when calling the sparsifier in our algorithm, so it succeeds with probability at least $1 - \frac{1}{n^2}$.

Pseudocode for Dual KOSZ with batching, sparsification, and recursion is given in Algorithm 4. The base case of the recursive algorithm is when $|V| \leq n_0$, where n_0 is a constant that the algorithm can choose. For the base case, we simply use Gaussian elimination to solve $\mathbf{L}_G \mathbf{x} = \mathbf{b}$, which takes $O(1)$ time since n_0 is a constant. For every t , contracting G down to H^t and computing the new resistances takes $O(m)$ time; this is because each edge in G contributes to the resistance between exactly one pair of nodes in H^t .

5.4 Analysis of the Sparsify and Recurse Algorithm

We now analyze Algorithm 4. We first bound the convergence rate, then analyze the running time.

5.4.1 Error Analysis

The lemma below bounds the expected rate of convergence of \mathbf{x}^t to \mathbf{x}^* .

Lemma 9 *For all $t \geq 0$, we have $\mathbb{E} \|\mathbf{x}^* - \mathbf{x}^t\|_{L_G}^2 \leq \left(1 - \beta + \beta e^{-\frac{d}{\tau}}\right)^t \|\mathbf{x}^*\|_{L_G}^2$. Here,*

- $\tau = O(m \log n \log \log n)$ is the stretch of the spanning tree,
- d is the number of updates in each batch,
- $\beta = \left(1 - \frac{1}{n_0}\right) \left(1 - \left(4\epsilon' \cdot \frac{1+\gamma}{1-\gamma} \cdot \left(1 + \frac{\gamma^2}{(1-\gamma)^2}\right) + \frac{2\gamma^2}{(1-\gamma)^2}\right)\right)$.

In particular, if we choose $n_0 = 10$, $\gamma = \frac{1}{100}$, and $\epsilon' = \frac{1}{100}$, then $\beta \geq \frac{4}{5}$, so that

$$\mathbb{E} \|\mathbf{x}^* - \mathbf{x}^t\|_{L_G}^2 \leq \left(\frac{1}{5} + \frac{4}{5} e^{-\frac{d}{\tau}}\right)^t \|\mathbf{x}^*\|_{L_G}^2 .$$

Algorithm 4 Dual KOSZ with batching, sparsification, and recursion

- 1: If $|V| \leq n_0$, solve $\mathbf{L}_G \mathbf{x} = \mathbf{b}$ using Gaussian elimination and **return** \mathbf{x} .
- 2: ▷ \mathbf{L}_G is the Laplacian matrix of G .
- 3: Compute a low-stretch spanning tree T of G .
- 4: Initialize $\mathbf{x}^0 = \mathbf{0}$.
- 5: **for** $t = 0$ to K **do**
- 6: Generate the next d updates. These correspond to d tree edges $e_1^t, \dots, e_d^t \in T$.
- 7: Let V_0^t, \dots, V_d^t be the vertex sets of the connected components of
- 8: $T - \{e_1^t, \dots, e_d^t\}$.
- 9: Contract G to H^t with $d + 1$ vertices: Each V_i^t is a vertex in H^t , and
- 10: resistances in H^t are

$$r_{H^t}(V_k^t, V_l^t) = \left(\sum_{ij \in \delta(V_k, V_l)} \frac{1}{r(i, j)} \right)^{-1}.$$

- 11: Compute $\tilde{\mathbf{L}}_{H^t}$, a γ -spectral sparsifier of \mathbf{L}_{H^t} , where $\gamma \in (0, 1)$ will be
- 12: a parameter that we will determine later.
- 13: Let $\mathbf{b}_G^t = \mathbf{b} - \mathbf{L}_G \mathbf{x}^t$.
- 14: Let $\mathbf{b}_{H^t}^t \in \mathbb{R}^{V(H^t)}$ be defined as follows: $b_{H^t}^t(V_i^t) = \sum_{u \in V_i^t} b_G^t(u)$ for all $i = 0, 1, \dots, d$.
- 15: Call the algorithm recursively to solve the Laplacian system $\tilde{\mathbf{L}}_{H^t} \tilde{\Delta}^t = \mathbf{b}_{H^t}^t$ for $\tilde{\Delta}^t$. This will return an approximate solution $\tilde{\Delta}_{\epsilon'}^t$ that satisfies $\|\tilde{\Delta}_{\epsilon'}^t - \tilde{\Delta}^t\|_{\tilde{\mathbf{L}}_{H^t}}^2 \leq \epsilon' \|\tilde{\Delta}^t\|_{\tilde{\mathbf{L}}_{H^t}}^2$. Here, ϵ' is the error parameter that we will input to the recursive call, to be determined later.
- 16: Update \mathbf{x}^t using $\tilde{\Delta}_{\epsilon'}^t$ to get the next iterate \mathbf{x}^{t+1} . For every vertex $u \in V$,

$$x^{t+1}(u) \leftarrow x^t(u) + \tilde{\Delta}_{\epsilon'}^t(V_i^t),$$

where V_i^t is the set in V_0^t, \dots, V_d^t such that $u \in V_i^t$. In other words, we update \mathbf{x}^t to \mathbf{x}^{t+1} by adding $\tilde{\Delta}_{\epsilon'}^t(V_i^t)$ to the potential of every vertex in V_i^t .

- 17: **if** $\mathcal{B}(\mathbf{x}^{t+1}) \leq \mathcal{B}(\mathbf{x}^t)$, revert $\mathbf{x}^{t+1} \leftarrow \mathbf{x}^t$.
- 18: **end for**
- 19: Return \mathbf{x}^K and the corresponding tree-defined flow \mathbf{f}^K .

Proof Define the random variable $D^t := \mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}^t)$. We will show in Lemma 11 that for every possible realization \mathbf{x}^t , we have $\mathbb{E}[D^{t+1} | \mathbf{x}^t] \leq (1 - \beta + \beta e^{-d/\tau}) \mathbb{E}[D^t | \mathbf{x}^t]$. This implies that $\mathbb{E}[D^{t+1}] \leq (1 - \beta + \beta e^{-d/\tau}) \mathbb{E}[D^t]$ unconditionally.

It then follows that

$$\mathbb{E}[D^t] \leq (1 - \beta + \beta e^{-d/\tau})^t \mathbb{E}[D^0] = (1 - \beta + \beta e^{-d/\tau})^t \mathcal{B}(\mathbf{x}^*).$$

Thus, $\mathcal{B}(\mathbf{x}^*) - \mathbb{E}[\mathcal{B}(\mathbf{x}^t)] \leq (1 - \beta + \beta e^{-d/\tau})^t \mathcal{B}(\mathbf{x}^*)$. □

As in the original analysis of Dual KOSZ, we will study the duality gap and analyze its decrease at each step of the algorithm. Consider some iteration t of the algorithm. Recall that \mathbf{x}^t is the iterate at the start of iteration t . For every possible sequence of e_1^t, \dots, e_d^t (the trees edges chosen in iteration t), define the following:

- Let $\hat{\mathbf{x}}^{t+1}$ be the vector obtained from \mathbf{x}^t by adding $\Delta^t(V_i^t)$ to every vertex in V_i^t , where $\Delta^t := \mathbf{L}_{H^t}^\dagger \mathbf{b}_{H^t}^t$.
- Let $\tilde{\mathbf{x}}^{t+1}$ be obtained from \mathbf{x}^t by applying the updates for the sequence of tree edges e_1^t, \dots, e_d^t , one by one. (i.e. Exactly as in the original, unbatched version of Dual KOSZ described in Sect. 3.)

Lemma 10 Fix any choice of e_1^t, \dots, e_d^t , and assume that $\tilde{\mathbf{L}}_{H^t}$ is a γ -approximate sparsifier of \mathbf{L}_{H^t} . Then

$$\mathcal{B}(\mathbf{x}^{t+1}) - \mathcal{B}(\mathbf{x}^t) \geq (1 - \alpha) \left(\mathcal{B}(\hat{\mathbf{x}}^{t+1}) - \mathcal{B}(\mathbf{x}^t) \right)$$

where $\alpha = 4\epsilon' \cdot \frac{1+\gamma}{1-\gamma} \cdot \left(1 + \frac{\gamma^2}{(1-\gamma)^2} \right) + \frac{2\gamma^2}{(1-\gamma)^2}$.

If we further assume that $\gamma \in (0, \frac{1}{2})$, we can simplify to get

$$\mathcal{B}(\mathbf{x}^{t+1}) - \mathcal{B}(\mathbf{x}^t) \geq (1 - 12\epsilon' - 8\gamma^2 - 48\epsilon'\gamma^2) \left(\mathcal{B}(\hat{\mathbf{x}}^{t+1}) - \mathcal{B}(\mathbf{x}^t) \right).$$

Proof To simplify notation, in this proof we will use

- \mathbf{b}_H to denote $\mathbf{b}_{H^t}^t$,
- $\tilde{\Delta}_{\epsilon'}$ to denote $\tilde{\Delta}_{e'}^t$,
- $\tilde{\Delta}$ to denote $\tilde{\Delta}^t$
- Δ to denote Δ^t ,
- H to denote H^t ,

Later in this proof, we will show that

$$\left\| \tilde{\Delta}_{\epsilon'} - \Delta \right\|_{\mathbf{L}_H}^2 \leq \alpha \|\Delta\|_{\mathbf{L}_H}^2, \tag{1}$$

for a constant α that depends on ϵ' and γ . Assuming (1) holds, by the definition of the matrix norm it follows that

$$\left(\tilde{\Delta}_{\epsilon'} - \Delta \right)^T \mathbf{L}_H \left(\tilde{\Delta}_{\epsilon'} - \Delta \right) \leq \alpha \Delta^T \mathbf{L}_H \Delta.$$

Expanding the left-hand side and rearranging, we get

$$2\tilde{\Delta}_{\epsilon'}^T \mathbf{L}_H \Delta - \tilde{\Delta}_{\epsilon'}^T \mathbf{L}_H \tilde{\Delta}_{\epsilon'} \geq (1 - \alpha) \Delta^T \mathbf{L}_H \Delta.$$

Using $\mathbf{L}_H \Delta = \mathbf{b}_H$, this becomes

$$2\tilde{\Delta}_{\epsilon'}^T \mathbf{b}_H - \tilde{\Delta}_{\epsilon'}^T \mathbf{L}_H \tilde{\Delta}_{\epsilon'} \geq (1 - \alpha) \Delta^T \mathbf{L}_H \Delta.$$

Recall that \mathbf{x}^{t+1} is obtained from \mathbf{x}^t by adding $\tilde{\Delta}_{\epsilon'}(V_i^t)$ to every vertex in V_i^t . Using Lemma 8 with $\Delta(i) = \tilde{\Delta}_{\epsilon'}(V_i^t)$, $\mathbf{x} = \mathbf{x}^t$, $\tilde{\mathbf{x}} = \mathbf{x}^{t+1}$, $\tilde{\mathbf{L}} = \mathbf{L}_H$, and $\tilde{\mathbf{b}} = \mathbf{b}_H$ it follows that the left-hand side is equal to $\mathcal{B}(\mathbf{x}^{t+1}) - \mathcal{B}(\mathbf{x}^t)$. On the other hand, $\mathcal{B}(\hat{\mathbf{x}}^{t+1}) - \mathcal{B}(\mathbf{x}^t) = 2\mathbf{b}_H^T \Delta - \Delta^T \mathbf{L}_H \Delta = \Delta^T \mathbf{L}_H \Delta$. (Since $\mathbf{L}_H \Delta = \mathbf{b}_H$.) Thus, the right-hand side is equal to $(1 - \alpha) (\mathcal{B}(\hat{\mathbf{x}}^{t+1}) - \mathcal{B}(\mathbf{x}^t))$. Thus we have

$$\mathcal{B}(\mathbf{x}^{t+1}) - \mathcal{B}(\mathbf{x}^t) \geq (1 - \alpha) (\mathcal{B}(\hat{\mathbf{x}}^{t+1}) - \mathcal{B}(\mathbf{x}^t)),$$

as claimed.

It remains to prove (1). To prove (1), note that we have

1. $\|\tilde{\Delta}_{\epsilon'} - \tilde{\Delta}\|_{\tilde{\mathbf{L}}_H}^2 \leq \epsilon' \|\tilde{\Delta}\|_{\tilde{\mathbf{L}}_H}^2$ (This is the error from the recursive solve).
2. $\|\tilde{\Delta} - \Delta\|_{\mathbf{L}_H}^2 \leq h(\gamma) \|\Delta\|_{\mathbf{L}_H}^2$ (follows by part 2 of Proposition 1 in the ‘‘Appendix’’).
This is the error from sparsification).

The first inequality, together with $(1 - \gamma)\mathbf{L}_H \preceq \tilde{\mathbf{L}}_H \preceq (1 + \gamma)\mathbf{L}_H$ and part 1 of Proposition 1 in the ‘‘Appendix’’, implies that

$$\|\tilde{\Delta}_{\epsilon'} - \tilde{\Delta}\|_{\mathbf{L}_H}^2 \leq \frac{1}{1 - \gamma} \|\tilde{\Delta}_{\epsilon'} - \tilde{\Delta}\|_{\tilde{\mathbf{L}}_H}^2 \leq \frac{\epsilon'}{1 - \gamma} \|\tilde{\Delta}\|_{\tilde{\mathbf{L}}_H}^2 \leq \epsilon' \cdot \frac{1 + \gamma}{1 - \gamma} \cdot \|\tilde{\Delta}\|_{\mathbf{L}_H}^2.$$

Now, using the inequality $\|a + b\|^2 \leq 2\|a\|^2 + 2\|b\|^2$ (which holds for any norm), we note that

$$\|\tilde{\Delta}\|_{\mathbf{L}_H}^2 \leq 2\|\Delta\|_{\mathbf{L}_H}^2 + 2\|\tilde{\Delta} - \Delta\|_{\mathbf{L}_H}^2 \leq 2\|\Delta\|_{\mathbf{L}_H}^2 + 2h(\gamma) \|\Delta\|_{\mathbf{L}_H}^2.$$

Hence,

$$\|\tilde{\Delta}_{\epsilon'} - \tilde{\Delta}\|_{\mathbf{L}_H}^2 \leq 2\epsilon' \cdot \frac{1 + \gamma}{1 - \gamma} \cdot (1 + h(\gamma)) \|\Delta\|_{\mathbf{L}_H}^2.$$

Again using $\|a + b\|^2 \leq 2\|a\|^2 + 2\|b\|^2$, we have

$$\begin{aligned} \|\tilde{\Delta}_{\epsilon'} - \Delta\|_{\mathbf{L}_H}^2 &\leq 2 \left(\|\tilde{\Delta}_{\epsilon'} - \tilde{\Delta}\|_{\mathbf{L}_H}^2 + \|\tilde{\Delta} - \Delta\|_{\mathbf{L}_H}^2 \right) \\ &\leq 2 \left(2\epsilon' \cdot \frac{1 + \gamma}{1 - \gamma} \cdot (1 + h(\gamma)) \|\Delta\|_{\mathbf{L}_H}^2 + h(\gamma) \|\Delta\|_{\mathbf{L}_H}^2 \right) \\ &= \left(4\epsilon' \cdot \frac{1 + \gamma}{1 - \gamma} \cdot (1 + h(\gamma)) + 2h(\gamma) \right) \|\Delta\|_{\mathbf{L}_H}^2. \end{aligned}$$

Therefore, (1) holds with $\alpha = 4\epsilon' \cdot \frac{1 + \gamma}{1 - \gamma} \cdot (1 + h(\gamma)) + 2h(\gamma)$. □

Lemma 11 For any vector \mathbf{x}^t , we have

$$\mathcal{B}(\mathbf{x}^*) - \mathbb{E}[\mathcal{B}(\mathbf{x}^{t+1})] \leq \left(1 - \beta + \beta e^{-d/\tau}\right) (\mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}^t)),$$

where $\beta = (1 - \frac{1}{n_0^2})(1 - \alpha)$.

Here, the expectation is taken over the random choices of e_1^t, \dots, e_d^t , and also over the randomness of the sparsification step. (Recall that we use the sparsifier given in Theorem 1.2 of [16], which successfully returns a γ -approximate sparsifier with probability at least $1 - \frac{1}{|V(H^t)|^2}$.)

Proof By Lemma 4, we know that

$$\mathcal{B}(\mathbf{x}^*) - \mathbb{E}[\mathcal{B}(\bar{\mathbf{x}}^{t+1})] \leq \left(1 - \frac{1}{\tau}\right)^d (\mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}^t)) \leq e^{-\frac{d}{\tau}} (\mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}^t)).$$

Rearranging, this is equivalent to

$$\mathbb{E}[\mathcal{B}(\bar{\mathbf{x}}^{t+1})] - \mathcal{B}(\mathbf{x}^t) \geq \left(1 - e^{-\frac{d}{\tau}}\right) (\mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}^t)).$$

Observe that for every realization of e_1^t, \dots, e_d^t , we have $\mathcal{B}(\hat{\mathbf{x}}^{t+1}) \geq \mathcal{B}(\bar{\mathbf{x}}^{t+1})$. This is because $\hat{\mathbf{x}}^{t+1} - \mathbf{x}^t = \Delta^t$, where Δ^t by definition is the vector that maximizes the increase $\mathcal{B}(\hat{\mathbf{x}}^{t+1}) - \mathcal{B}(\mathbf{x}^t)$ while subject to being incremented by the same amount on each of the components V_0^t, \dots, V_d^t . On the other hand, the vector $\bar{\mathbf{x}}^{t+1} - \mathbf{x}^t$ is also incremented by the same amount on each of the components V_0^t, \dots, V_d^t by the way our original algorithm works.

Since $\mathcal{B}(\hat{\mathbf{x}}^{t+1}) \geq \mathcal{B}(\bar{\mathbf{x}}^{t+1})$ holds for every realization of e_1^t, \dots, e_d^t , it follows that $\mathbb{E}[\mathcal{B}(\hat{\mathbf{x}}^{t+1})] \geq \mathbb{E}[\mathcal{B}(\bar{\mathbf{x}}^{t+1})]$, where the expectation is taken over the random choices of e_1^t, \dots, e_d^t made by the algorithm. Hence,

$$\mathbb{E}[\mathcal{B}(\hat{\mathbf{x}}^{t+1})] - \mathcal{B}(\mathbf{x}^t) \geq \left(1 - e^{-\frac{d}{\tau}}\right) (\mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}^t)). \tag{2}$$

To conclude, we will use Lemma 10 to translate the above inequality (which is in terms of $\hat{\mathbf{x}}^{t+1}$), to an inequality in terms of \mathbf{x}^{t+1} . We have

- With probability $\geq 1 - \frac{1}{n_0^2}$, the sparsifier is successful and by Lemma 10,

$$\mathbb{E}[\mathcal{B}(\mathbf{x}^{t+1})] - \mathcal{B}(\mathbf{x}^t) \geq (1 - \alpha) \left(\mathbb{E}[\mathcal{B}(\hat{\mathbf{x}}^{t+1})] - \mathcal{B}(\mathbf{x}^t)\right).$$

- With probability $\leq \frac{1}{n_0^2}$, the sparsifier is unsuccessful and $\mathbb{E}[\mathcal{B}(\mathbf{x}^{t+1})] - \mathcal{B}(\mathbf{x}^t) \geq 0$.

This is because in the algorithm, we evaluate $\mathcal{B}(\mathbf{x}^{t+1})$ and only update \mathbf{x}^t to \mathbf{x}^{t+1} if $\mathcal{B}(\mathbf{x}^{t+1}) \geq \mathcal{B}(\mathbf{x}^t)$. Otherwise, we make $\mathbf{x}^{t+1} = \mathbf{x}^t$.

Note that the above expectations are with respect to the random choices of e_1^t, \dots, e_d^t , conditioned on the sparsifier being successful/unsuccessful. Now, taking another expectation with respect to the randomness of the sparsifier, we get

$$\begin{aligned} \mathbb{E}[\mathcal{B}(\mathbf{x}^{t+1})] - \mathcal{B}(\mathbf{x}^t) &\geq \left(1 - \frac{1}{n_0^2}\right) (1 - \alpha) \left(\mathbb{E}[\mathcal{B}(\hat{\mathbf{x}}^{t+1})] - \mathcal{B}(\mathbf{x}^t)\right) \\ &\geq \left(1 - \frac{1}{n_0^2}\right) (1 - \alpha) \left(1 - e^{-\frac{d}{\tau}}\right) (\mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}^t)) \quad \text{(by (2))} \end{aligned}$$

Rearranging the above inequality gives

$$\mathcal{B}(\mathbf{x}^*) - \mathbb{E}[\mathcal{B}(\mathbf{x}^{t+1})] \leq \left(1 - \left(1 - \frac{1}{n_0^2}\right) (1 - \alpha) \left(1 - e^{-\frac{d}{\tau}}\right)\right) (\mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}^t)),$$

as claimed. □

5.4.2 Running Time Analysis

The following theorem bounds the running time of the algorithm.

Theorem 2 *For any $\delta \in (0, 1)$ and $\epsilon > 0$, Dual KOSZ with batching, sparsification, and recursion finds \mathbf{x} with $\mathbb{E} \|\mathbf{x}^* - \mathbf{x}\|_L^2 \leq \epsilon \|\mathbf{x}^*\|_L^2$ in $O(m^{1+\delta}(\log n)^{\frac{8}{\delta}}(\log \frac{1}{\epsilon})^{\frac{1}{\delta}})$ time.*

Proof By Lemma 9, it suffices to run the algorithm for K iterations, for any $K \geq \frac{\ln \epsilon}{\ln\left(\frac{1}{5} + \frac{4}{3}e^{-d/\tau}\right)}$.

Using the inequalities $e^{-x} \leq 1 - \frac{x}{2}$ and $\ln(1 - x) \leq -\frac{x}{2}$ which hold for $x \in (0, 1)$, we see that it suffices to choose $K = \frac{5\tau}{d} \ln(1/\epsilon)$.

Recall $\tau \leq c_3 m \log^2 n$, for some constant c_3 . Thus, we choose $d = c_3 \bar{m} (\log^2 n) \cdot m^{-\delta}$. Here, \bar{m} is the number of edges in the *current iteration*, while m is the number of edges in the *topmost* iteration (i.e. in the original graph G). With this choice of d , we have $K = 5c_3 m^\delta \ln(1/\epsilon)$. (Note that K is the same at every level of the recursion tree.)

The work involved in each call to the algorithm consists of

- Computing T ,
- Doing K times
 - Contracting and sparsifying to a graph with d vertices and $a_1 d \log^c n / \gamma^2$ edges, for some constant a_1 .
 - Doing a recursive call.

If m is the number of edges in the graph at one level of the recursion tree, then at the next level, the number of edges is

$$a_1 d \log^2 n \log \log n / \gamma^2 \leq a_1 c_3 \bar{m} (\log^2 n) \cdot m^{-\delta} \log^3 n / \gamma^2 = a_1 c_3 \bar{m} \cdot m^{-\delta} (\log n)^5 / \gamma^2.$$

Therefore, the number of edges in a graph at level l of the recursion tree is

$$\text{edges}_l \leq m^{1-l\delta} (a_1 c_3)^l (\log n)^{5l} / \gamma^{2l}.$$

The total work required by a node at level l of the recursion tree is dominated by the sparsifier, which takes time

$$\text{work}_l = K \cdot a_2 \cdot \text{edges}_l \cdot \log^2 n \log \log n / \gamma^2 \leq a_2 \cdot \text{edges}_l \cdot \log^3 n / \gamma^2$$

for some constant a_2 .

Finally, the total number of recursion tree nodes at level l is equal to K^l . This implies that the total work required by all the nodes at level l of the recursion tree is equal to

$$\begin{aligned} \text{total work}_l &= \text{work}_l \cdot K^l \\ &\leq K \cdot a_2 \cdot \text{edges}_l \cdot \log^3 n \cdot \frac{1}{\gamma^2} \cdot K^l \\ &\leq K \cdot a_2 \cdot \left(m^{1-l\delta} (a_1 c_3)^l (\log n)^{5l} / \gamma^{2l} \right) \cdot \log^3 n \cdot \frac{1}{\gamma^2} \cdot K^l \\ &= 5^{l+1} c_3^{l+1} m^{1+\delta} (\ln 1/\epsilon)^{l+1} a_2 (a_1 c_3)^l (\log n)^{(5l+3)} \cdot \frac{1}{\gamma^{2l+2}} \\ &\leq A^l m^{1+\delta} (\log n)^{8l} (\ln 1/\epsilon)^{l+1}, \end{aligned}$$

for some numerical constant A .

To conclude, we note that the total work summed across all the levels is at most a constant factor times the total work at the maximum level, which is $l = \frac{1}{\delta}$. \square

Remark 3 By setting $\delta = \sqrt{\frac{8 \log \log n + \log \log \frac{1}{\epsilon}}{\log m}}$ (which is the choice of δ that minimizes $m^{1+\delta} (\log n)^{\frac{8}{\delta}} (\log \frac{1}{\epsilon})^{\frac{1}{\delta}}$), the running time in Theorem 2 becomes $O\left(m \exp\left(2\sqrt{\log m (8 \log \log n + \log \log \frac{1}{\epsilon})}\right)\right)$, which is $O(m^{1+o(1)})$ if ϵ is a constant.

6 Conclusion

We propose a cut-based combinatorial algorithm to solve Laplacian systems approximately. This algorithm is dual to the cycle-based algorithm by Kelner et al. [7]. We show that our algorithm converges in a near-linear number of iterations.

To achieve a near-linear running time, we would further need each iteration to run in polylogarithmic time. We give evidence against this, by presenting a reduction from the OMv conjecture. This is in contrast to the algorithm in [7], which uses a data structure such that each iteration of the algorithm runs in $O(\log n)$ time. In order to obtain a better running time, one would need to show it is possible take advantage of the

particular structure of updates in the algorithm to implement the data structure. Note that our reduction crucially needs that a very specific spanning tree (albeit with very small stretch) is chosen. Is it possible for the algorithm to choose a different small-stretch spanning tree that is amendable to a polylogarithmic time implementation?



Acknowledgements Monika Henzinger was supported by funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme Grant agreement No. 101019564 “The Design of Modern Fully Dynamic Data Structures (MoDynStruct)” and from the Austrian Science Fund (FWF) project “Fast Algorithms for a Reactive Network Layer (ReactNet)”, P 33775-N, with additional funding from the *netidee SCIENCE Stiftung*, 2020–2024. Billy Jin was Supported in part by NSERC fellowship PGSD3-532673-2019 and NSF grant CCF-2007009. Richard Peng was supported in part by an NSERC Discovery Grant and NSF grant CCF-1846218. David P. Williamson was supported in part by NSF grant CCF-2007009.

Author Contributions All authors contributed to the paper.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

A Omitted Proofs from Sect. 3

Lemma 1 *Let $\mathbf{x} \in \mathbb{R}^V$ be a vector of potentials and let $C \subset V$. Let \mathbf{x}' be the potentials obtained from \mathbf{x} as in the algorithm (that is, by adding Δ to the potential of every vertex in C so that flow conservation is satisfied across $\delta(C)$). Then*

$$\mathcal{B}(\mathbf{x}') - \mathcal{B}(\mathbf{x}) = \frac{\Delta^2}{2R(C)}.$$

Proof The way we update \mathbf{x} is by adding a constant Δ to the potentials of every vertex in C , where

$$\Delta = (b(C) - f(C)) \cdot R(C)$$

Recall that $f(C)$ is the net amount of flow going out of C in the flow induced by \mathbf{x} . That is,

$$f(C) = \sum_{\substack{ij \in E \\ i \in C, j \notin C}} \frac{x(i) - x(j)}{r(i, j)}$$

Note that the new potentials \mathbf{x}' can be expressed as $\mathbf{x}' = \mathbf{x} + \Delta \mathbb{1}_C$. We have

$$\begin{aligned} 2(\mathcal{B}(\mathbf{x}') - \mathcal{B}(\mathbf{x})) &= 2\mathbf{b}^\top \mathbf{x}' - (\mathbf{x}')^\top \mathbf{L} \mathbf{x}' - (2\mathbf{b}^\top \mathbf{x} - \mathbf{x}^\top \mathbf{L} \mathbf{x}) \\ &= 2\mathbf{b}^\top (\mathbf{x} + \Delta \cdot \mathbb{1}_C) - 2\mathbf{b}^\top \mathbf{x} - (\mathbf{x}')^\top \mathbf{L} \mathbf{x}' + \mathbf{x}^\top \mathbf{L} \mathbf{x} \end{aligned}$$

$$\begin{aligned}
 &= 2\Delta \cdot \mathbf{b}^\top \mathbb{1}_C - \sum_{ij \in E} \frac{1}{r(i, j)} [(x'(i) - x'(j))^2 - (x(i) - x(j))^2] \\
 &= 2\Delta \cdot \mathbf{b}^\top \mathbb{1}_C - \sum_{ij \in \delta(C)} \frac{1}{r(i, j)} [(x'(i) - x'(j))^2 - (x(i) - x(j))^2] \\
 &= 2\Delta \cdot \mathbf{b}^\top \mathbb{1}_C - \sum_{\substack{i \in C, j \notin C \\ ij \in \delta(C)}} \frac{1}{r(i, j)} [(x(i) + \Delta - x(j))^2 - (x(i) - x(j))^2] \\
 &= 2\Delta \cdot \mathbf{b}^\top \mathbb{1}_C - \sum_{\substack{i \in C, j \notin C \\ ij \in \delta(C)}} \frac{1}{r(i, j)} [2\Delta \cdot (x(i) - x(j)) + \Delta^2] \\
 &= 2\Delta \cdot \mathbf{b}^\top \mathbb{1}_C - 2\Delta \cdot f(C) - \Delta^2 \sum_{(i, j) \in \delta(C)} \frac{1}{r(i, j)} \\
 &= 2\Delta \cdot \mathbf{b}^\top \mathbb{1}_C - 2\Delta \cdot f(C) - \Delta^2 \cdot R(C)^{-1} \\
 &= 2\Delta \cdot b(C) - 2\Delta \cdot f(C) - \Delta^2 R(C)^{-1} \\
 &= 2\Delta^2 R(C)^{-1} - \Delta^2 R(C)^{-1} \\
 &= \Delta^2 / R(C).
 \end{aligned}$$

□

Lemma 2 We have $\text{gap}(\mathbf{f}, \mathbf{x}) = \frac{1}{2} \sum_{(i, j) \in \bar{E}} r(i, j) \left(f(i, j) - \frac{x(i) - x(j)}{r(i, j)} \right)^2$.

Proof By definition, we have

$$2 \text{gap}(\mathbf{f}, \mathbf{x}) = \sum_{e \in E} r(e) f(e)^2 - (2\mathbf{b}^\top \mathbf{x} - \mathbf{x}^\top \mathbf{L} \mathbf{x}).$$

Note that

$$\begin{aligned}
 \mathbf{b}^\top \mathbf{x} &= \sum_{i \in V} b(i)x(i) \\
 &= \sum_{i \in V} x(i) \left(\sum_{j: (i, j) \in \bar{E}} f(i, j) - \sum_{j: (j, i) \in \bar{E}} f(j, i) \right) \\
 &= \sum_{(i, j) \in \bar{E}} f(i, j)(x(i) - x(j))
 \end{aligned}$$

and

$$\mathbf{x}^\top \mathbf{L} \mathbf{x} = \sum_{(i, j) \in \bar{E}} \frac{(x(i) - x(j))^2}{r(i, j)}.$$

Plugging these into our expression for $\text{gap}(\mathbf{f}, \mathbf{x})$, we obtain

$$\begin{aligned} 2 \text{gap}(\mathbf{f}, \mathbf{x}) &= \sum_{(i,j) \in \bar{E}} \left[r(i, j) f(i, j)^2 - 2f(i, j)(x(i) - x(j)) + \frac{(x(i) - x(j))^2}{r(i, j)} \right] \\ &= \sum_{(i,j) \in \bar{E}} r(i, j) \left(f(i, j) - \frac{x(i) - x(j)}{r(i, j)} \right)^2 \end{aligned}$$

which is what we wanted to show. □

Lemma 3 *Let T be a spanning tree, \mathbf{x} vertex potentials, and \mathbf{b} a supply vector. Let $f_{T,\mathbf{x}}$ be the associated tree-defined flow. Then*

$$\text{gap}(f_{T,\mathbf{x}}, \mathbf{x}) = \frac{1}{2} \sum_{(i,j) \in T} r(i, j) \cdot \frac{\Delta(C(i, j))^2}{R(C(i, j))^2}.$$

Proof Recall that $C(i, j)$, $\Delta(C(i, j))$ and $R(C(i, j))$ were defined as follows:

- $C(i, j)$ is the set of vertices on the side of the fundamental cut of T determined by (i, j) containing i . In other words, $C(i, j)$ consists of the vertices in the component of $T - ij$ with $i \in C(i, j)$ and $j \notin C(i, j)$.
- $R(C(i, j)) = \left(\sum_{ij \in \delta(C)} \frac{1}{r(i, j)} \right)^{-1}$.
- $\Delta(C(i, j)) = (b(C(i, j)) - f(C(i, j)))R(C(i, j))$, where
 - $b(C(i, j)) = \mathbf{b}^\top \mathbb{1}_{C(i, j)}$, and
 - $f(C(i, j)) = \sum_{\substack{k \in C(i, j), l \notin C(i, j) \\ kl \in E}} \frac{x(k) - x(l)}{r(k, l)}$

We have

$$\begin{aligned} 2 \text{gap}(f_{T,\mathbf{x}}, \mathbf{x}) &= \sum_{(i,j) \in E} r(i, j) \left(f_{T,\mathbf{x}}(i, j) - \frac{x(i) - x(j)}{r(i, j)} \right)^2 \\ &= \sum_{(i,j) \in T} r(i, j) \left(f_{T,\mathbf{x}}(i, j) - \frac{x(i) - x(j)}{r(i, j)} \right)^2 \\ &= \sum_{(i,j) \in T} r(i, j) \left[\left(b(C(i, j)) - \sum_{\substack{k \in C(i, j), l \notin C(i, j) \\ kl \in E - ij}} \frac{x(k) - x(l)}{r(k, l)} \right) - \frac{x(i) - x(j)}{r(i, j)} \right]^2 \\ &= \sum_{(i,j) \in T} r(i, j) \left[b(C(i, j)) - \sum_{\substack{k \in C(i, j), l \notin C(i, j) \\ kl \in E}} \frac{x(k) - x(l)}{r(k, l)} \right]^2 \\ &= \sum_{(i,j) \in T} r(i, j) [b(C(i, j)) - f(C(i, j))]^2 \end{aligned}$$

$$= \sum_{(i,j) \in T} r(i,j) \cdot \frac{\Delta(C(i,j))^2}{R(C(i,j))^2}$$

□

Lemma 4 *If each iteration of the algorithm samples an edge $(i, j) \in T$ according to the probabilities $P_{ij} = \frac{1}{\tau} \cdot \frac{r(i,j)}{R(C(i,j))}$, then we have*

$$\mathcal{B}(\mathbf{x}^*) - \mathbb{E}[\mathcal{B}(\mathbf{x}^{t+1})] \leq \left(1 - \frac{1}{\tau}\right) (\mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}^t)).$$

Proof We know from the discussion above that

$$\mathbb{E}[\mathcal{B}(\mathbf{x}^{t+1})] - \mathcal{B}(\mathbf{x}^t) = \frac{1}{\tau} \text{gap}(\mathbf{f}_{T,\mathbf{x}}, \mathbf{x}^t),$$

where $\mathbf{f}_{T,\mathbf{x}}$ is the tree-defined flow associated with potentials \mathbf{x}^t . Since $\text{gap}(\mathbf{f}_{T,\mathbf{x}}, \mathbf{x}^t) \geq \mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}^t)$, we get

$$\mathbb{E}[\mathcal{B}(\mathbf{x}^{t+1})] - \mathcal{B}(\mathbf{x}^t) \geq \frac{1}{\tau} (\mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}^t)).$$

Rearranging gives

$$\mathcal{B}(\mathbf{x}^*) - \mathbb{E}[\mathcal{B}(\mathbf{x}^{t+1})] \leq \left(1 - \frac{1}{\tau}\right) (\mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}^t)),$$

as desired. □

Corollary 1 *After $K = \tau \ln(\frac{1}{\epsilon})$ iterations, $\mathcal{B}(\mathbf{x}^*) - \mathbb{E}[\mathcal{B}(\mathbf{x}^K)] \leq \epsilon \cdot \mathcal{B}(\mathbf{x}^*)$.*

Proof Define the random variable $D_t := \mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}^t)$. By Lemma 4, we know that

$$\mathbb{E}[D^{t+1} \mid \mathbf{x}^t] \leq \left(1 - \frac{1}{\tau}\right) \mathbb{E}[D^t \mid \mathbf{x}^t]$$

for all possible vectors of potentials \mathbf{x}^t . This implies that $\mathbb{E}[D^{t+1}] \leq (1 - \frac{1}{\tau}) \mathbb{E}[D^t]$ unconditionally.

By induction on t , it then follows that

$$\mathbb{E}[D^K] \leq \left(1 - \frac{1}{\tau}\right)^K \mathbb{E}[D^0] = \left(1 - \frac{1}{\tau}\right)^K (\mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}^0)) = \left(1 - \frac{1}{\tau}\right)^K \mathcal{B}(\mathbf{x}^*).$$

Thus,

$$\mathcal{B}(\mathbf{x}^*) - \mathbb{E}[\mathcal{B}(\mathbf{x}^K)] \leq \left(1 - \frac{1}{\tau}\right)^K \mathcal{B}(\mathbf{x}^*).$$

Using the inequality $1 - x \leq e^{-x}$, we obtain

$$\mathcal{B}(\mathbf{x}^*) - \mathbb{E}[\mathcal{B}(\mathbf{x}^K)] \leq e^{-K/\tau} \mathcal{B}(\mathbf{x}^*).$$

Hence, if $K \geq \tau \ln(\frac{1}{\epsilon})$, then we will have $\mathcal{B}(\mathbf{x}^*) - \mathbb{E}[\mathcal{B}(\mathbf{x}^K)] \leq \epsilon \cdot \mathcal{B}(\mathbf{x}^*)$, as desired. \square

Theorem 1 *Let τ be the total stretch of T . After $K = \tau \ln(\frac{\tau}{\epsilon})$ iterations, Dual KOSZ returns $\mathbf{x}^K \in \mathbb{R}^V$ and $\mathbf{f}^K \in \mathbb{R}^{\vec{E}}$ such that $\mathbb{E} \|\mathbf{x}^* - \mathbf{x}^K\|_{\mathbf{L}}^2 \leq \frac{\epsilon}{\tau} \|\mathbf{x}^*\|_{\mathbf{L}}^2$ and $\mathbb{E}[\mathcal{E}(\mathbf{f}^K)] \leq (1 + \epsilon)\mathcal{E}(\mathbf{f}^*)$, for \mathbf{f}^* and \mathbf{x}^* optimal primal and dual solutions respectively.*

Proof of Theorem 1 By Corollary 1, after $K = \tau \ln(\frac{\tau}{\epsilon})$ iterations, the algorithm returns potentials \mathbf{x}^K such that $\mathcal{B}(\mathbf{x}^*) - \mathbb{E}[\mathcal{B}(\mathbf{x}^K)] \leq \frac{\epsilon}{\tau} \mathcal{B}(\mathbf{x}^*)$. Combining with Lemma 13, we get that $\mathbb{E} \|\mathbf{x}^* - \mathbf{x}^K\|_{\mathbf{L}}^2 \leq \frac{\epsilon}{\tau} \|\mathbf{x}^*\|_{\mathbf{L}}^2$. Finally, Lemma 14 gives $\mathbb{E}[\mathcal{E}(\mathbf{f}^K)] \leq (1 + \epsilon)\mathcal{E}(\mathbf{f}^*)$. \square

Lemma 12 *We have $\tau = st_T(G, \mathbf{r})$.*

Proof We write out the definitions of τ and $st_T(G, \mathbf{r})$:

$$\tau = \sum_{(i,j) \in T} \frac{r(i, j)}{R(C(i, j))} = \sum_{(i,j) \in T} r(i, j) \sum_{(k,l) \in \delta(C(i, j))} \frac{1}{r(k, l)}$$

and

$$st_T(G, \mathbf{r}) = \sum_{(i,j) \in \vec{E}} st_T((i, j), \mathbf{r}) = \sum_{(i,j) \in \vec{E}} \frac{1}{r(i, j)} \sum_{(k,l) \in P(i, j)} r(k, l),$$

where $P(i, j)$ is the unique path from i to j in T .

It turns out that the expressions for τ and $st_T(G)$ are summing exactly the same terms, just in different ways. Indeed, we have

$$\begin{aligned} \tau &= \sum_{(i,j) \in T} \sum_{(k,l) \in \delta(C(i, j))} \frac{r(i, j)}{r(k, l)} \\ &= \sum_{(k,l) \in \vec{E}} \sum_{(i,j) \in P(k, l)} \frac{r(i, j)}{r(k, l)} \\ &= st_T(G, \mathbf{r}). \end{aligned}$$

To switch the order of summation from the first line to the second line, we used the fact that for an edge $(k, l) \in \vec{E}$, we have $(k, l) \in \delta(C(i, j))$ if and only if $(i, j) \in P(k, l)$. This is because T is a spanning tree. \square

By Corollary 1, we know that the potentials \mathbf{x}^t found by the algorithm satisfy the property that $\mathcal{B}(\mathbf{x}^t)$ converges to $\mathcal{B}(\mathbf{x}^*)$ at a linear rate, in expectation. The following lemma shows that if \mathbf{x} is a set of potentials such that $\mathcal{B}(\mathbf{x})$ is close to $\mathcal{B}(\mathbf{x}^*)$, then \mathbf{x} is close to \mathbf{x}^* as a vector (measured in the matrix norm defined by the Laplacian \mathbf{L}).

Lemma 13 *Let \mathbf{x} be any vector of potentials. Then $\frac{1}{2} \|\mathbf{x}^* - \mathbf{x}\|_L^2 = \mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x})$. In particular, if $\mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}) \leq \epsilon \cdot \mathcal{B}(\mathbf{x}^*)$, then $\|\mathbf{x}^* - \mathbf{x}\|_L^2 \leq \epsilon \|\mathbf{x}^*\|_L^2$.*

Proof We have

$$\begin{aligned} \|\mathbf{x}^* - \mathbf{x}\|_L^2 &= (\mathbf{x}^* - \mathbf{x})^\top \mathbf{L}(\mathbf{x}^* - \mathbf{x}) \\ &= (\mathbf{x}^*)^\top \mathbf{L}\mathbf{x}^* - 2\mathbf{x}^\top \mathbf{L}\mathbf{x}^* + \mathbf{x}^\top \mathbf{L}\mathbf{x} \\ &= 2\mathcal{B}(\mathbf{x}^*) - 2\mathbf{x}^\top \mathbf{b} + \mathbf{x}^\top \mathbf{L}\mathbf{x} \\ &= 2\mathcal{B}(\mathbf{x}^*) - 2\mathcal{B}(\mathbf{x}). \end{aligned}$$

In particular, if $\mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}) \leq \epsilon \cdot \mathcal{B}(\mathbf{x}^*)$, then $\|\mathbf{x}^* - \mathbf{x}\|_L^2 \leq 2\epsilon \cdot \mathcal{B}(\mathbf{x}^*) = \epsilon \|\mathbf{x}^*\|_L^2$. This is because

$$2\mathcal{B}(\mathbf{x}^*) = 2\mathbf{b}^\top \mathbf{x}^* - (\mathbf{x}^*)^\top \mathbf{L}\mathbf{x}^* = (\mathbf{x}^*)^\top \mathbf{L}\mathbf{x}^* = \|\mathbf{x}^*\|_L^2.$$

□

Next, we show that if $\mathcal{B}(\mathbf{x})$ is sufficiently close to $\mathcal{B}(\mathbf{x}^*)$, then the associated tree-defined flow $\mathbf{f}_{T,\mathbf{x}}$ has energy sufficiently close to $\mathcal{E}(\mathbf{f}^*)$.

Lemma 14 *For any distribution over \mathbf{x} such that $\mathbb{E}_x[\mathcal{B}(\mathbf{x})] \geq (1 - \frac{\epsilon}{\tau})\mathcal{B}(\mathbf{x}^*)$, we have $\mathbb{E}_x[\mathcal{E}(\mathbf{f}_{T,\mathbf{x}})] \leq (1 + \epsilon)\mathcal{E}(\mathbf{f}^*)$.*

Proof For ease of notation, in this proof let $\mathbf{f} = \mathbf{f}_{T,\mathbf{x}}$. (Note that \mathbf{f} is a random vector that is a function of \mathbf{x} .) We have $\mathbb{E}_x[\mathcal{E}(\mathbf{f}) - \mathcal{E}(\mathbf{f}^*)] = \mathbb{E}_x[\text{gap}(\mathbf{f}, \mathbf{x}^*)]$.

For a fixed choice of \mathbf{x} , consider running the algorithm for one more iteration starting from \mathbf{x} to obtain a vector \mathbf{x}' . Then we have $\mathbb{E}[\mathcal{B}(\mathbf{x}')] - \mathcal{B}(\mathbf{x}) = \frac{1}{\tau} \text{gap}(\mathbf{f}, \mathbf{x})$. This implies $\mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}) \geq \frac{1}{\tau} \text{gap}(\mathbf{f}, \mathbf{x})$. Taking expectations with respect to \mathbf{x} , we get $\mathbb{E}_x[\mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x})] \geq \frac{1}{\tau} \mathbb{E}_x[\text{gap}(\mathbf{f}, \mathbf{x})]$. Thus,

$$\begin{aligned} \mathbb{E}_x[\mathcal{E}(\mathbf{f}) - \mathcal{E}(\mathbf{f}^*)] &= \mathbb{E}_x[\text{gap}(\mathbf{f}, \mathbf{x}^*)] \\ &= \mathbb{E}_x[\text{gap}(\mathbf{f}, \mathbf{x}) - (\mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x}))] \\ &\leq (\tau - 1)\mathbb{E}_x[\mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x})] \\ &\leq \tau \mathbb{E}_x[\mathcal{B}(\mathbf{x}^*) - \mathcal{B}(\mathbf{x})] \\ &\leq \epsilon \mathcal{B}(\mathbf{x}^*) \\ &= \epsilon \mathcal{E}(\mathbf{f}^*). \end{aligned}$$

□

B Spectral Approximations

Let \mathbf{A}, \mathbf{B} be $n \times n$ symmetric, positive semidefinite matrices. We say that \mathbf{B} is a γ -spectral sparsifier of \mathbf{A} if

$$(1 - \gamma) \mathbf{A} \preceq \mathbf{B} \preceq (1 + \gamma) \mathbf{A}.$$

Proposition 1 (Spectral Approximations) *Suppose $\mathbf{A}, \mathbf{B} \in \mathbb{S}_+^n$ and $(1 - \gamma)\mathbf{A} \preceq \mathbf{B} \preceq (1 + \gamma)\mathbf{A}$. Let $\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{b} \in \mathbb{R}^n$. Then the following hold:*

1. $(1 - \gamma) \|\mathbf{x}\|_{\mathbf{A}}^2 \leq \|\mathbf{x}\|_{\mathbf{B}}^2 \leq (1 + \gamma) \|\mathbf{x}\|_{\mathbf{A}}^2$
2. If $\mathbf{A}\mathbf{x} = \mathbf{b}$ and $\mathbf{B}\mathbf{y} = \mathbf{b}$, then $\|\mathbf{x} - \mathbf{y}\|_{\mathbf{A}}^2 \leq h(\gamma) \|\mathbf{x}\|_{\mathbf{A}}^2$, where $h(\gamma) = \frac{\gamma^2}{(1-\gamma)^2}$.

Proof The first one is by definition of $\|\mathbf{x}\|_{\mathbf{A}}^2 = \mathbf{x}^\top \mathbf{A} \mathbf{x}$.

For the second one, first we claim that it is sufficient to prove that $\|\mathbf{A}^\dagger \mathbf{b} - \mathbf{B}^\dagger \mathbf{b}\|_{\mathbf{A}}^2 \leq h(\gamma) \|\mathbf{A}^\dagger \mathbf{b}\|_{\mathbf{A}}^2$. This is because in general, we have $\mathbf{x} = \mathbf{A}^\dagger \mathbf{b} + \mathbf{u}$, and $\mathbf{y} = \mathbf{B}^\dagger \mathbf{b} + \mathbf{v}$, for some $\mathbf{u} \in \text{Null}(\mathbf{A})$ and $\mathbf{v} \in \text{Null}(\mathbf{B})$. Moreover, the condition $(1 - \gamma)\mathbf{A} \preceq \mathbf{B} \preceq (1 + \gamma)\mathbf{A}$ implies that $\text{Null}(\mathbf{A}) = \text{Null}(\mathbf{B})$. Hence, $\|\mathbf{x} - \mathbf{y}\|_{\mathbf{A}}^2 = \|\mathbf{A}^\dagger \mathbf{b} - \mathbf{B}^\dagger \mathbf{b}\|_{\mathbf{A}}^2$, and $\|\mathbf{x}\|_{\mathbf{A}}^2 = \|\mathbf{A}^\dagger \mathbf{b}\|_{\mathbf{A}}^2$.

Next, we expand $\|\mathbf{A}^\dagger \mathbf{b} - \mathbf{B}^\dagger \mathbf{b}\|_{\mathbf{A}}^2 \leq h(\gamma) \|\mathbf{A}^\dagger \mathbf{b}\|_{\mathbf{A}}^2$ into

$$(\mathbf{A}^\dagger \mathbf{b} - \mathbf{B}^\dagger \mathbf{b})^\top \mathbf{A} (\mathbf{A}^\dagger \mathbf{b} - \mathbf{B}^\dagger \mathbf{b}) \leq h(\gamma) \mathbf{b}^\top \mathbf{A}^\dagger \mathbf{b},$$

or equivalently,

$$\mathbf{b}^\top (\mathbf{A}^\dagger - \mathbf{B}^\dagger) \mathbf{A} (\mathbf{A}^\dagger - \mathbf{B}^\dagger) \mathbf{b} \leq h(\gamma) \mathbf{b}^\top \mathbf{A}^\dagger \mathbf{b}.$$

To prove the above inequality, it suffices to prove that

$$(\mathbf{A}^\dagger - \mathbf{B}^\dagger) \mathbf{A} (\mathbf{A}^\dagger - \mathbf{B}^\dagger) \preceq h(\gamma) \mathbf{A}^\dagger. \tag{3}$$

Multiplying the left and right sides of Eq. 3 by $\mathbf{A}^{\frac{1}{2}}$, we get that (3) is implied by

$$\mathbf{A}^{\frac{1}{2}} (\mathbf{A}^\dagger - \mathbf{B}^\dagger) \mathbf{A} (\mathbf{A}^\dagger - \mathbf{B}^\dagger) \mathbf{A}^{\frac{1}{2}} \preceq h(\gamma) \mathbf{A}^{\frac{1}{2}} \mathbf{A}^\dagger \mathbf{A}^{\frac{1}{2}}. \tag{4}$$

Let $\Pi := \mathbf{A}^{\frac{1}{2}} \mathbf{A}^\dagger \mathbf{A}^{\frac{1}{2}}$ be the projection map onto the row space of \mathbf{A} . Note that $\Pi = \mathbf{A}^\dagger \mathbf{A} = \mathbf{A} \mathbf{A}^\dagger$. Also, $\Pi = \mathbf{A}^{\frac{1}{2}} \mathbf{A}^\dagger \mathbf{A}^{\frac{1}{2}} = \mathbf{A}^{\frac{1}{2}} \mathbf{A}^\dagger \mathbf{A}^{\frac{1}{2}}$. These can be seen using the spectral decomposition. Now, the reason why Eq. 4 implies Eq. 3 is because if we can multiply both sides of (4) with one copy of $\mathbf{A}^{\frac{1}{2}}$ on the left and one copy of $\mathbf{A}^{\frac{1}{2}}$ on the right. Then Eq. 4 becomes $\Pi (\mathbf{A}^\dagger - \mathbf{B}^\dagger) \mathbf{A} (\mathbf{A}^\dagger - \mathbf{B}^\dagger) \Pi \preceq h(\gamma) \Pi \mathbf{A}^\dagger \Pi$. We have $\Pi (\mathbf{A}^\dagger - \mathbf{B}^\dagger) = (\mathbf{A}^\dagger - \mathbf{B}^\dagger) \Pi = \mathbf{A}^\dagger - \mathbf{B}^\dagger$ because \mathbf{A} and \mathbf{B} have the same null space. Similarly, $\Pi \mathbf{A}^\dagger = \mathbf{A}^\dagger \Pi = \mathbf{A}^\dagger$.

To prove (4), first rewrite it as

$$\left(\mathbf{A}^{\frac{1}{2}} (\mathbf{A}^\dagger - \mathbf{B}^\dagger) \mathbf{A}^{\frac{1}{2}} \right)^2 \preceq h(\gamma) \Pi,$$

or equivalently

$$\left(\Pi - \mathbf{A}^{\frac{1}{2}} \mathbf{B}^\dagger \mathbf{A}^{\frac{1}{2}} \right)^2 \preceq h(\gamma) \Pi.$$

From the spectral approximation $(1 - \gamma) \mathbf{A} \preceq \mathbf{B} \preceq (1 + \gamma) \mathbf{A}$, we deduce that

$$\frac{1}{1 + \gamma} \mathbf{A}^\dagger \preceq \mathbf{B}^\dagger \preceq \frac{1}{1 - \gamma} \mathbf{A}^\dagger,$$

which, when multiplying on the left and right by $\mathbf{A}^{\frac{1}{2}}$, implies that

$$\frac{1}{1 + \gamma} \Pi \preceq \mathbf{A}^{\frac{1}{2}} \mathbf{B}^\dagger \mathbf{A}^{\frac{1}{2}} \preceq \frac{1}{1 - \gamma} \Pi.$$

This in turn gives

$$\frac{-\gamma}{1 + \gamma} \Pi \preceq \Pi - \mathbf{A}^{\frac{1}{2}} \mathbf{B}^\dagger \mathbf{A}^{\frac{1}{2}} \preceq \frac{\gamma}{1 - \gamma} \Pi.$$

Observe that any eigenvector of $\Pi - \mathbf{A}^{\frac{1}{2}} \mathbf{B}^\dagger \mathbf{A}^{\frac{1}{2}}$ is also an eigenvector of Π (they share eigenspaces because \mathbf{A} and \mathbf{B} have the same null spaces). Moreover, the eigenvalues of Π are 0 or 1. This implies that the eigenvalues of $\Pi - \mathbf{A}^{\frac{1}{2}} \mathbf{B}^\dagger \mathbf{A}^{\frac{1}{2}}$ are all between $\frac{-\gamma}{1 + \gamma}$ and $\frac{\gamma}{1 - \gamma}$. Hence, the eigenvalues of $(\Pi - \mathbf{A}^{\frac{1}{2}} \mathbf{B}^\dagger \mathbf{A}^{\frac{1}{2}})^2$ are all between 0 and $\frac{\gamma^2}{(1 - \gamma)^2}$, and thus $(\Pi - \mathbf{A}^{\frac{1}{2}} \mathbf{B}^\dagger \mathbf{A}^{\frac{1}{2}})^2 \preceq \frac{\gamma^2}{(1 - \gamma)^2} \Pi$. □

References

1. Spielman, D.A., Teng, S.-H.: Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In: Proceedings of the 36th Annual ACM Symposium on the Theory of Computing, pp. 81–90 (2004)
2. Chen, L., Kyng, R., Liu, Y.P., Peng, R., Gutenberg, M.P., Sachdeva, S.: Maximum Flow and Minimum-Cost Flow in Almost-Linear Time (2022). <https://doi.org/10.48550/ARXIV.2203.00671>. [arXiv:2203.00671](https://arxiv.org/abs/2203.00671)
3. Christiano, P., Kelner, J.A., Mądry, A., Spielman, D., Teng, S.-H.: Electrical flows, Laplacian systems, and faster approximation of maximum flow in undirected graphs. In: Proceedings of the 42nd Annual ACM Symposium on the Theory of Computing, pp. 273–282 (2011)
4. Kathuria, T., Liu, Y.P., Sidford, A.: Unit capacity maxflow in almost $m^{4/3}$ time. *SIAM J. Comput.* (2022). <https://doi.org/10.1137/20M1383525>
5. Mądry, A.: Computing maximum flow with augmenting electrical flows. In: Proceedings of the 57th IEEE Symposium on the Foundations of Computer Science, pp. 593–602 (2016)
6. Gao, Y., Liu, Y.P., Peng, R.: Fully dynamic electrical flows: sparse maxflow faster than Goldberg-Rao. In: 2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS), pp. 516–527 (2022). <https://doi.org/10.1109/FOCS52979.2021.00058>
7. Kelner, J.A., Orecchia, L., Sidford, A., Zhu, Z.A.: A simple, combinatorial algorithm for solving SDD systems in nearly-linear time. In: Proceedings of the 45th Annual ACM Symposium on the Theory of Computing, pp. 911–920 (2013). <https://doi.org/10.1145/2488608.2488724>
8. Boman, E.G., Dewese, K., Gilbert, J.R.: Evaluating the potential of a dual randomized Kaczmarz Laplacian linear solver. *Informatica* **40**, 95–107 (2016)
9. Boman, E.G., Dewese, K., Gilbert, J.R.: An empirical comparison of graph Laplacian solvers. In: Proceedings of the 18th Workshop on Algorithm Engineering and Experiments, pp. 174–188 (2016)
10. Dewese, K., Gilbert, J.R., Miller, G., Peng, R., Xu, H., Xu, S.: An empirical study of cycle toggling based Laplacian solvers. In: Proceedings of the 7th SIAM Workshop on Combinatorial Scientific Computing, pp. 33–41 (2016)

11. Hoske, D., Lukarski, D., Meyerhenke, H., Wegner, M.: Engineering a combinatorial Laplacian solver: Lessons learned. *Algorithms* **9** (2016). Article 72
12. Goldberg, A.V., Tarjan, R.E.: Finding minimum-cost circulations by canceling negative cycles. *J. ACM* **36**, 873–886 (1989)
13. Ervolina, T.R., McCormick, S.T.: Two strongly polynomial cut cancelling algorithms for minimum cost network flow. *Discrete Appl. Math.* **46**, 133–165 (1993)
14. Henzinger, M., Krinninger, S., Nanongkai, D., Saranurak, T.: Unifying and strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture. In: *Proceedings of the 47th Annual ACM Symposium on the Theory of Computing*, pp. 21–30 (2015). <https://doi.org/10.1145/2746539.2746609>
15. Sherman, J.: Personal communication (2017)
16. Kyng, R., Pachocki, J., Peng, R., Sachdeva, S.: A framework for analyzing resparsification algorithms. In: *Proceedings of the 2017 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 2032–2043 (2017). <https://doi.org/10.1137/1.9781611974782.132>
17. Batson, J.D., Spielman, D.A., Srivastava, N., Teng, S.: Spectral sparsification of graphs: theory and algorithms. *Commun. ACM* **56**(8), 87–94 (2013)
18. Abraham, I., Neiman, O.: Using petal-decompositions to build a low stretch spanning tree. In: *Proceedings of the 44th Symposium on the Theory of Computing*, pp. 395–406 (2012). <https://doi.org/10.1145/2213977.2214015>
19. Williamson, D.P.: *Network Flow Algorithms*. Cambridge University Press, Cambridge (2019)
20. Sleator, D.D., Tarjan, R.E.: A data structure for dynamic trees. *J. Comput. Syst. Sci.* **26**(3), 362–391 (1983). [https://doi.org/10.1016/0022-0000\(83\)90006-5](https://doi.org/10.1016/0022-0000(83)90006-5)
21. Karger, D.R.: Minimum cuts in near-linear time. *J. ACM* **47**(1), 46–76 (2000). <https://doi.org/10.1145/331605.331608>
22. Koutis, I., Miller, G.L., Peng, R.: A fast solver for a class of linear systems. *Commun. ACM* **55**(10), 99–107 (2012). <https://doi.org/10.1145/2347736.2347759>
23. Jambulapati, A., Sidford, A.: Ultrasparse ultrasparsifiers and faster laplacian system solvers. In: Marx, D. (ed.) *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10–13, 2021*, pp. 540–559. SIAM (2021). [arXiv:2011.08806](https://arxiv.org/abs/2011.08806)
24. Cohen, M.B., Kyng, R., Miller, G.L., Pachocki, J.W., Peng, R., Rao, A.B., Xu, S.C.: Solving SDD linear systems in nearly $m \log^{1/2} n$ time. In: *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing*, pp. 343–352 (2014)
25. Koutis, I., Xu, S.C.: Simple parallel and distributed algorithms for spectral graph sparsification. *ACM Trans. Parallel Comput.* (2016). <https://doi.org/10.1145/2948062>
26. Spielman, D.A., Teng, S.: Spectral sparsification of graphs. *SIAM J. Comput.* **40**(4), 981–1025 (2011). [arXiv:0808.4134](https://arxiv.org/abs/0808.4134)
27. Kapralov, M., Panigrahy, R.: Spectral sparsification via random spanners. In: *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference. ITCS '12*, pp. 393–398. Association for Computing Machinery, New York, NY, USA (2012). <https://doi.org/10.1145/2090236.2090267>
28. Peng, R., Spielman, D.A.: An Efficient Parallel Solver for SDD Linear Systems (2013)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.