

Network Design for s - t Effective Resistance

PAK HAY CHAN, Google Waterloo, Canada

LAP CHI LAU, University of Waterloo, Canada

AARON SCHILD, University of Washington, USA

SAM CHIU-WAI WONG, Microsoft Research Redmond, USA

HONG ZHOU, Fuzhou University, China

We consider a new problem of designing a network with small s - t effective resistance. In this problem, we are given an undirected graph $G = (V, E)$, two designated vertices $s, t \in V$, and a budget k . The goal is to choose a subgraph of G with at most k edges to minimize the s - t effective resistance. This problem is an interpolation between the shortest path problem and the minimum cost flow problem and has applications in electrical network design.

We present several algorithmic and hardness results for this problem and its variants. On the hardness side, we show that the problem is NP-hard, and the weighted version is hard to approximate within a factor smaller than two assuming the small-set expansion conjecture. On the algorithmic side, we analyze a convex programming relaxation of the problem and design a constant factor approximation algorithm. The key of the rounding algorithm is a randomized path-rounding procedure based on the optimality conditions and a flow decomposition of the fractional solution. We also use dynamic programming to obtain a fully polynomial time approximation scheme when the input graph is a series-parallel graph, with better approximation ratio than the integrality gap of the convex program for these graphs.

CCS Concepts: • **Theory of computation** → **Routing and network design problems; Rounding techniques; Network optimization**; • **Mathematics of computing** → **Spectra of graphs; Approximation algorithms**;

Additional Key Words and Phrases: Effective resistance, network design

ACM Reference format:

Pak Hay Chan, Lap Chi Lau, Aaron Schild, Sam Chiu-Wai Wong, and Hong Zhou. 2022. Network Design for s - t Effective Resistance. *ACM Trans. Algorithms* 18, 3, Article 22 (October 2022), 45 pages.

<https://doi.org/10.1145/3522588>

Most of this work was done while Pak Hay Chan and Hong Zhou were graduate students at the University of Waterloo.

This work was supported by NSERC Discovery Grant No. 2950-120715.

Authors' addresses: P. H. Chan, Google Waterloo, Waterloo, ON, Canada; email: ph5chan@uwaterloo.ca; L. C. Lau, University of Waterloo, Waterloo, ON, Canada; email: lapchi@uwaterloo.ca; A. Schild, University of Washington, Seattle, WA, USA; email: aaron.schild@gmail.com; S. C.-W. Wong, Microsoft Research Redmond, Redmond, WA, USA; email: samwon@microsoft.com; H. Zhou (corresponding author), School of Mathematics and Statistics, Fuzhou University, Fuzhou, Fujian, China; email: hong.zhou@fzu.edu.cn.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

1549-6325/2022/10-ART22 \$15.00

<https://doi.org/10.1145/3522588>

1 INTRODUCTION

Network design problems are generally about finding a minimum cost subgraph that satisfies certain “connectivity” requirements. One of the most well studied problems is the survivable network design problem [1, 22, 25, 26, 29], where the requirement is to have a specified number $r_{u,v}$ of edge-disjoint paths between every pair of vertices u, v . Other combinatorial requirements are also well studied in the literature, including vertex connectivity [7, 9, 13, 18, 33, 34] and shortest path distances [14, 15].

Some spectral requirements are also studied, including spectral expansion [2, 32], total effective resistances [24, 40], and mixing time [5], but in general much less is known about these problems. See Section 1.1 for more discussions of previous work.

In this article, we study a basic problem in designing networks with a spectral requirement—the effective resistance between two vertices.

Definition 1.1 (The s - t Effective Resistance Network Design Problem). The input is an undirected graph $G = (V, E)$, two specified vertices $s, t \in V$, and a budget k . The goal is to find a subgraph H of G with at most k edges that minimizes $\text{Reff}_H(s, t)$, where $\text{Reff}_H(s, t)$ denotes the effective resistance between s and t in the subgraph H . See Section 2.2 for the definition of effective resistance and Section 3.1 for a mathematical formulation of the problem.

The s - t effective resistance is an interpolation between s - t shortest path distance and s - t edge connectivity. To see this, let $f \in \mathbb{R}^{|E|}$ be a unit s - t flow in G and define the ℓ_p -energy of f as $\mathcal{E}_p(f) := (\sum_e |f_e|^p)^{1/p}$, and let $\mathcal{E}_p(s, t) := \min_f \{\mathcal{E}_p(f) \mid f \text{ is a unit } s\text{-}t \text{ flow}\}$ be the minimum ℓ_p -energy of a unit s - t flow that the graph G can support. Thomson’s principle (see Section 2.2) states that $\text{Reff}_G(s, t) = \mathcal{E}_2^2(s, t)$, so that a graph of small s - t effective resistance can support a unit s - t flow with small ℓ_2 -energy. Note that the shortest path distance between s and t is $\mathcal{E}_1(s, t)$ (as the ℓ_1 -energy of a flow is just the average path length and is minimized by a shortest s - t path), and so a graph with small $\mathcal{E}_1(s, t)$ has a short path between s and t . Note also that the edge-connectivity between s and t is equal to the reciprocal of $\mathcal{E}_\infty(s, t)$ (because if there are k edge-disjoint s - t paths, we can set the flow value on each path to be $1/k$), and so a graph with small $\mathcal{E}_\infty(s, t)$ has many edge-disjoint s - t paths. As ℓ_2 is between ℓ_1 and ℓ_∞ , the objective function $\text{Reff}(s, t) = \mathcal{E}_2^2(s, t)$ takes both the s - t shortest path distance and the s - t edge-connectivity into consideration.

A simple property suggests that ℓ_2 -energy may be even more desirable than ℓ_1 and ℓ_∞ as a connectivity measure. Conceptually, adding an edge e to G would make s and t more connected. For ℓ_1 and ℓ_∞ , however, adding e would not yield a better energy if e does not improve the shortest path and the edge connectivity, respectively. In contrast, the ℓ_2 -energy would typically improve after adding an edge, and so ℓ_2 -energy provides a smoother quantitative measure that better captures our intuition how well s and t are connected in a network.

Traditionally, the effective resistance has many useful probabilistic interpretations, such as the commute time [8], the cover time [38], and the probability of an edge in a random spanning tree [31]. These interpretations suggest that the effective resistance is a useful distance function and have applications in the study of social networks. Recently, effective resistance has found surprising applications in solving problems about graph connectivity, including constructing spectral sparsifiers [47] (by using the effective resistance of an edge as the sampling probability), computing maximum flow [11], finding thin trees for ATSP [3], and generating random spanning trees [39, 45].

Thomson’s principle [30] also states that the electrical flow between s and t is the unique flow that minimizes the ℓ_2 -energy (see Section 2.2). So, designing a network with small s - t effective resistance has natural applications in designing electrical networks [16, 24, 28]. One natural formulation is to keep at most k wires in the input electrical network to minimize $\text{Reff}(s, t)$, so that

the electrical flow between s and t can still be sent with small energy while we switch off many wires in the electrical network.

Based on the above reasons, we believe that the effective resistance is a nice and natural alternative connectivity measure in network design. More generally, it is an interesting direction to develop techniques to solve network design problems with spectral requirements.

1.1 Main Results

Unlike the classical problems of shortest path and min-cost flow (corresponding to the ℓ_1 and ℓ_∞ versions of the problem), the s - t effective resistance network design problem is NP-hard.

THEOREM 1.2. *The s - t effective resistance network design problem is NP-hard.*

On the other hand, we analyze a natural convex programming relaxation for the problem (Section 3.1), and use it to design a constant factor approximation algorithm for the problem.

THEOREM 1.3. *There is a convex programming-based randomized algorithm that returns an 8-approximate solution in polynomial time with high probability for the s - t effective resistance network design problem.*

The algorithm crucially uses a nice characterization of the optimal solutions to the convex program (Lemma 3.2) to design a randomized path-rounding procedure (Section 3.2) for Theorem 1.3.

A simple example shows that the integrality gap of the convex program is at least two. When the budget k is much larger than the length of a shortest s - t path, we show how to achieve an approximation ratio close to two with a randomized “short” path rounding algorithm (Section 3.5).

THEOREM 1.4. *There is a $(2+O(\epsilon))$ -approximation algorithm for the s - t effective resistance network design problem, when $k \geq 2d_{st}/\epsilon^{10}$ where d_{st} is the length of a shortest s - t path.*

1.2 Other Results

We consider some variants of the s - t effective resistance network design problem, including the weighted version, the dual version, and the problem on special graphs.

There is a natural weighted generalization of the s - t effective resistance network design problem, where we associate a cost c_e and resistance r_e to each edge e of the input graph.

Definition 1.5 (The Weighted s - t Effective Resistance Network Design Problem). The input is an undirected graph $G = (V, E)$ where each edge e has a non-negative cost c_e and a non-negative resistance r_e , two specified vertices $s, t \in V$, and a cost budget k . The goal is to find a subgraph H of G that minimizes $\text{Reff}_H(s, t)$ subject to the constraint that the total edge cost of H is at most k . In the following, we may refer to this problem as the weighted problem for simplicity.

In the weighted problem, the integrality gap of the convex program (Section 3.1) becomes unbounded, even when the cost on the edges are the same ($c_e = 1$ for all $e \in E$). This suggests that the weighted version may be strictly harder. Indeed, we show stronger hardness result for the weighted problem assuming the small-set expansion conjecture [43, 44].

THEOREM 1.6. *Assuming the small-set expansion conjecture, it is NP-hard to approximate the weighted s - t effective resistance network design problem within a factor of $2 - \epsilon$ for any $\epsilon > 0$, even when $c_e = 1$ for every edge e .*

On the other hand, when the cost on the edges are the same, the following approximation follows from the randomized path rounding algorithm in a black box manner.

COROLLARY 1.7. *There is a convex programming-based $O(R)$ -approximation randomized algorithm for the weighted s - t effective resistance network design problem when $c_e = 1$ for every edge e , where $R = \max_e r_e / \min_e r_e$ is the ratio between the maximum and minimum resistance.*

As our problem is related to electrical network design, it is natural to consider the special case when the input graph is a series-parallel graph. In this setting, we can use dynamic programming to design an exact algorithm for the original problem, and a **fully polynomial time approximation scheme (FPTAS)** for the weighted problem.

THEOREM 1.8. *There is an exact algorithm for the s - t effective resistance network design problem with running time $O(|E| \cdot k^2)$ when the input graph is a series-parallel graph.*

There is a $(1 + \varepsilon)$ -approximation algorithm for the weighted s - t effective resistance network design problem when the input graph is a series-parallel graph. The running time of the algorithm is $O(|E|^7 R^2/\varepsilon^2)$ where $R = \max_e r_e / \min_e r_e$ is the ratio between the maximum and minimum resistance. By a simple preprocessing scaling step, we can assume that R is bounded by a polynomial, and so the algorithm is a FPTAS for the weighted problem.

We note that the integrality gap examples in Section 3.1 are actually series-parallel graphs, and so the dynamic programming algorithms go beyond the limitation of the natural convex program. We leave it as an open problem whether the weighted problem admits a constant factor approximation algorithm (possibly by combining these techniques).

We also consider the “dual” problem, where we set the effective resistance as a hard constraint, and the objective is to minimize the number of edges in the solution subgraph. We present similar results as the original problem in Section 3.6.

1.3 Related Work

In the survivable network design problem, we are given an undirected graph and a connectivity requirement $r_{u,v}$ for every pair of vertices u, v , and the goal is to find a minimum cost subgraph such that there are at least $r_{u,v}$ edge-disjoint paths for all u, v . This problem is extensively studied and captures many interesting special cases [1, 22, 25, 26]. The best approximation algorithm for this problem is due to Jain [29], who introduced the technique of iterative rounding to design a 2-approximation algorithm. His result has been extended in various directions, including element-connectivity [10, 19], directed graphs [21, 22], and with degree constraints [17, 20, 35, 36].

Other combinatorial connectivity requirements were also considered. A natural variation is to require $r_{u,v}$ internally vertex disjoint paths for every pair of vertices u, v . This problem is much harder to approximate [33, 34], but there are good approximation algorithms for global connectivity [9, 18] and when the maximum connectivity requirement is small [7, 13]. Another natural problem is to require a path of length $l_{u,v}$ between every pair of vertices u, v . This problem is also hard to approximate in general but there are better approximation algorithms when every edge has the same cost and the same length [15].

Spectral connectivity requirements were also studied, including algebraic connectivity [23, 32] (closely related to graph expansion), total effective resistances [24], and mixing time [5]. In particular, Ghosh, Boyd and Saberi [24] studied the related problem of minimizing the sum of effective resistances over all pairs of vertices. They gave a convex programming relaxation of the problem but did not provide any result for the discrete optimization setting. Most of the earlier works only proposed convex programming relaxations and heuristic algorithms, and approximation guarantees are only obtained recently for the more general experimental design problems. When every edge has the same cost, there is a $(1 + \varepsilon)$ -approximation algorithm for minimizing the total effective resistance when the budget is at least $\Omega(|V|/\varepsilon)$ [40], and there is a $(1 + \varepsilon)$ -approximation algorithm for both maximizing the algebraic connectivity and minimizing the total effective resistance when the budget is at least $\Omega(|V|/\varepsilon^2)$ [2]. For general edge costs, subsequent to our work, there is a randomized $(1 + \varepsilon)$ -approximation algorithm for both maximizing the algebraic connectivity and

minimizing the total effective resistance when the budget is at least $\Omega(c_{\max}|V|/\varepsilon^2)$ [37], where c_{\max} is the maximum edge cost in the input graph.

We remark that those methods in recent papers on experimental design [2, 37, 40] are closely related to the s - t effective resistance network design problem, where they consider each edge/vector as a unit and use convex relaxation-based rounding techniques. In Reference [40], a probability distribution on the edges/vectors is carefully designed for an dependent randomized rounding. Although this technique can be used to minimize total effective resistance and to handle some more general settings, it is highly unlikely that it can be applied to the s - t effective resistance network design problem as the objective functions are very different. In References [2, 37], a potential function as in spectral sparsification is used to guide a local search algorithm to swap two edges/vectors at a time to improve the current solution. These techniques can be applied to the s - t effective resistance network design problem. However, they only work in the case when the solutions form a spanning set so that the “contribution” of each individual edge/vector is well-defined. Thus, the techniques in References [2, 37] only apply when $k \geq \Omega(|V|/\varepsilon^2)$ for the desired precision $\varepsilon < 1$. In this article, the interesting regime is when k is much smaller than $|V|$. We have developed a set of new techniques for analyzing and rounding the solutions to the convex program that will hopefully find applications for solving related problems in the regime when k is small.

1.4 Techniques

Our main technical contribution is in designing rounding techniques for a convex programming relaxation of our problem. There is a natural convex programming relaxation, by using the conductance of the edges as variables, and writing the s - t effective resistance as the objective function and noting that it is convex with respect to the variables (Section 3.1).

We show that optimal solutions of this convex program enjoy some nice properties.¹ Given an optimal fractional solution x^* and the unit s - t electrical flow f^* supported in x^* , we derive from the KKT optimality conditions that there is a flow-conductance ratio $\alpha > 0$ such that $f_e^* = \alpha x_e^*$ for every fractional edge e with $0 < x_e^* < 1$ and $f_e^* \geq \alpha$ for every integral edge e with $x_e^* = 1$. The flow-conductance ratio α is crucial in the rounding algorithm and the analysis.

Our approach is based on a randomized rounding procedure on s - t paths. Given x^* , we compute the unit s - t electrical flow f^* supported in x^* , and decompose f^* as a convex combination of s - t paths. The rounding algorithm has $T = 1/\alpha$ iterations (recall that α is the flow-conductance ratio of the optimal solution x^*), where we pick a random path P_i from the convex combination in each iteration, and return $H := \cup_{i=1}^T P_i$ as our solution. One difference from the previous techniques in the literature is that each unit in the rounding algorithm is a s - t path, so in particular s and t are always connected in our solution. Another difference is that our problem has some extra structure, so that we can compute the electrical flow f^* to guide our rounding procedure, where the variables f_e^* are not in the convex program. These allow us to obtain a constant factor approximation algorithm for all budget $k \geq d_{st}$ (note that when $k < d_{st}$ there is no feasible integral solution).

In the analysis, we prove in Lemma 3.6 that the expected number of edges in H is at most k , and in Lemma 3.7 that the expected effective resistance is $\text{Reff}_H(s, t) \leq 2\text{Reff}_{x^*}(s, t)$. To bound the expected effective resistance, we use Thomson’s principle and construct a unit s - t flow F to show that $\text{Reff}_H(s, t) \leq \mathcal{E}_H(F) \leq 2\text{Reff}_{x^*}(s, t)$. To construct the unit s - t flow F , we keep the flow-conductance ratio and send α units of flow on each sampled path P_i (i.e., $f_e = \alpha$ and $x_e = 1$). The flow-conductance ratio plays a crucial role in the proofs of both lemmas. This is because the rounding algorithm is based on the flow variables f_e^* , and thus the performance guarantees are in

¹We can also show that there *exists* an optimal solution such that the fractional edges form a forest, but this is not included in the article, as we have not used this property in the rounding algorithm.

terms of f_e^* , but the ratio α allows us to relate them back to the variables x_e^* in the convex program. Combining the two lemmas give us a constant factor bicriteria approximation algorithm for the problem. This can be turned into a true approximation algorithm by scaling down the budget to $k/2$ and run the bicriteria approximation algorithm with some additional claims (Section 3.4).

The improvement on the approximation ratio when budget k is large comes from two observations. The first is that if k is much larger than the length of the shortest s - t path, then the number of independent iterations in the rounding scheme is large (Lemma 3.3). The second is that we can ignore some s - t paths in the flow decomposition with many fractional edges without affecting the performance much. Combining these, we can apply a Chernoff-Hoeffding bound to show that the number of edges is at most $(1 + \varepsilon)k$ with high probability. Then it is not necessary to scale down the budget by a factor of 2 and we can prove a stronger bound that the effective resistance is at most $2 + O(\varepsilon)$ times the optimal value.

1.5 Organization

In Section 2, we define the notations used in this article and cover background knowledge on effective resistances. We present the convex programming relaxation and our two rounding procedures in Section 3, and the dynamic programming algorithm in Section 4. The NP-hardness and small set expansion hardness results are provided in Section 5.

2 PRELIMINARIES

We introduce the notations and definitions for graphs and matrices in Section 2.1, and then we define electrical flow and effective resistance and state some basic results in Section 2.2.

2.1 Graphs and Matrices

Let $G = (V, E)$ be an undirected graph with edge weight $w_e \geq 0$ on each edge $e \in E$. The number of vertices and the number of edges are denoted by $n := |V|$ and $m := |E|$. For a subset of edges $F \subseteq E$, the total weight of edges in F is $w(F) := \sum_{e \in F} w_e$. For a subset of vertices $S \subseteq V$, the set of edges with one endpoint in S and one endpoint in $V - S$ is denoted by $\delta(S)$. For a vertex v , the set of edges incident on a vertex v is $\delta(v) := \delta(\{v\})$, and the weighted degree of v is $\deg(v) := w(\delta(v))$. The volume of a set $\text{vol}(S) := \sum_{v \in S} \deg(v)$ is defined as the sum of the weighted degrees of vertices in S . The conductance of a set $\phi(S) := w(\delta(S))/\text{vol}(S)$ is defined as the ratio of the total weight on the boundary of S to the total weighted degrees in S . For two disjoint subsets $S_1, S_2 \subseteq V$, the set of edges with one endpoint in S_1 and one endpoint in S_2 is denoted by $E(S_1, S_2)$.

In this article, an undirected graph $G = (V, E)$ with non-negative edge weights $w \in \mathbb{R}^E$ is interpreted as an electrical network, where each edge e is a resistor with conductance w_e (not to be confused with the conductance $\phi(S)$ of a set S as defined above), or equivalently with resistance $r_e := 1/w_e$. The adjacency matrix $A \in \mathbb{R}^{V \times V}$ of the graph is defined as $A_{u,v} = w_{u,v}$ for all $u, v \in V$. The Laplacian matrix $L \in \mathbb{R}^{V \times V}$ of the graph is defined as $L = D - A$ where $D \in \mathbb{R}^{V \times V}$ is the diagonal degree matrix with $D_{u,u} = \deg(u)$ for all $u \in V$. For each edge $e = uv \in E$, let $b_e := \chi_u - \chi_v$, where $\chi_u \in \mathbb{R}^n$ is the vector with one in the u th entry and zero otherwise (the orientation from u to v is chosen arbitrarily). The Laplacian matrix can also be written as

$$L = \sum_{e \in E} w_e b_e b_e^T = BWB^T,$$

where $W = \text{diag}(w)$ and $B \in \mathbb{R}^{V \times E}$ is the n -by- m matrix with b_e being the e th column.

Let $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ be the eigenvalues of L with corresponding orthonormal eigenvectors v_1, v_2, \dots, v_n so that $L = \sum_{i=1}^n \lambda_i v_i v_i^T$. It is well-known that the Laplacian matrix is positive semi-definite and $\lambda_1 = 0$ with $v_1 = \vec{1}/\sqrt{n}$ as the corresponding eigenvector, and $\lambda_2 > 0$ if and only if G

is connected (see, e.g., Chapter 1 of Reference [12]). The pseudo-inverse of the Laplacian matrix L of a connected graph is defined as

$$L^\dagger = \sum_{i=2}^n \frac{1}{\lambda_i} v_i v_i^T,$$

which maps every vector x orthogonal to v_1 to a vector y such that $Ly = x$.

2.2 Electrical Flow and Effective Resistance

We start with defining a unit s - t flow on an undirected graph $G = (V, E)$. We first fix an orientation of the edges of G arbitrarily. Let $B \in \mathbb{R}^{V \times E}$ be the matrix defined in Section 2.1 that is consistent with this orientation. A unit s - t flow $f : E \rightarrow \mathbb{R}$ is an m -dimensional vector that satisfies flow conservation constraints:

$$Bf = \chi_s - \chi_t \quad \text{or equivalently} \quad \sum_{e=uv : u \in \delta^+(v)} f_e - \sum_{e=uv : u \in \delta^-(v)} f_e = \begin{cases} 1 & v = s \\ -1 & v = t \\ 0 & \text{otherwise,} \end{cases}$$

where $\delta^+(v)$ and $\delta^-(v)$ are the set of outgoing and incoming neighbours of v with respect to the fixed orientation. Note that positive f_e indicates the flow on e is in the same direction as the orientation of e , while negative f_e indicates the opposite.

The unit s - t electrical flow is a unit s - t flow f that also satisfies the Ohm's law: There exists a potential vector $\varphi \in \mathbb{R}^n$ such that for all $e = uv \in E$ oriented from u to v ,

$$f_e = w_e \cdot (\varphi(u) - \varphi(v)).$$

Given a potential vector φ , we will use the orientation where all edges are pointing from the high potential endpoint to the low potential endpoint, so that f_e is nonnegative for all $e \in E$ in the rest of this article.

The effective resistance between s and t is defined as

$$\text{Reff}(s, t) := \varphi(s) - \varphi(t),$$

which is the potential difference between s and t when one unit of electrical flow is sent from s to t . The s - t effective resistance can be interpreted as the resistance of the whole graph G as a big resistor when an electrical flow is sent from s to t .

One can write the effective resistance in terms of the Laplacian matrix. For $u, v \in V$, let $b_{uv} = \chi_u - \chi_v$. Combining the flow conservation constraint and the Ohm's law, it can be checked that the potential vector $\varphi \in \mathbb{R}^V$ of the unit s - t electrical flow is a solution to the linear system

$$L \cdot \varphi = b_{st}.$$

Note that $\varphi = L^\dagger b_{st}$ is a solution, and if G is connected then any solution is given by $\varphi + c \cdot \vec{1}$ for $c \in \mathbb{R}$. Therefore, $\text{Reff}(s, t) = \varphi(s) - \varphi(t)$ is (uniquely) well-defined, and we can write

$$\text{Reff}(s, t) = \varphi(s) - \varphi(t) = b_{st}^T L^\dagger b_{st}.$$

The effective resistance can also be characterized by the energy of a flow. The energy of an s - t flow f is defined as

$$\mathcal{E}(f) := \sum_{e \in E} \frac{f_e^2}{w_e} = \sum_{e \in E} r_e f_e^2.$$

Thomson's principle [30] states that the unit s - t electrical flow is the unique unit s - t flow that minimizes the energy. This can be verified by writing down the optimality condition of the minimization problem. Moreover, this energy is exactly the s - t effective resistance. To see this, note

that the flow value on edge uv in the unit s - t electrical flow satisfies $f(uv) = w_{uv} \cdot (\varphi(u) - \varphi(v)) = w_{uv} \cdot b_{uv}^T L^\dagger b_{st}$, and thus

$$\mathcal{E}(f) = \sum_{uv \in E} w_{uv} (b_{uv}^T L^\dagger b_{st})^2 = b_{st}^T L^\dagger \left(\sum_{uv \in E} w_{uv} b_{uv} b_{uv}^T \right) L^\dagger b_{st} = b_{st}^T L^\dagger L L^\dagger b_{st} = \text{Reff}(s, t).$$

To summarize, we will use the following result from Thomson's principle.

FACT 2.1 (THOMSON'S PRINCIPLE [30]). *Let f^* be the unit electrical s - t flow in G . Then,*

$$\text{Reff}_G(s, t) = \min_f \{ \mathcal{E}(f) \mid f \text{ is a unit } s\text{-}t \text{ flow in } G \} = \mathcal{E}(f^*).$$

A corollary of Thomson's principle is the following intuitive result known as the Rayleigh's monotonicity principle.

FACT 2.2 (RAYLEIGH'S MONOTONICITY PRINCIPLE). *The s - t effective resistance cannot increase if the resistance of an edge is decreased.*

We will also use the following result to write a convex programming relaxation of our problem.

FACT 2.3 ([24]). *The s - t effective resistance is a convex function with respect to the conductance of the edges.*

3 CONVEX PROGRAMMING ALGORITHM

In this section, we analyze a convex programming relaxation for our problem in Definition 1.1. We first describe the convex program and prove a characterization of the optimal solutions in Section 3.1. We then present a randomized rounding algorithm using flow decomposition in Section 3.2, and show that it is a constant factor bicriteria approximation algorithm in Section 3.3. Then, we show how to convert the bicriteria approximation algorithm into a true approximation algorithm in Section 3.4, and how to modify the algorithm slightly to achieve a better approximation guarantee when the budget k is large in Section 3.5. Finally, we discuss the dual problem of minimizing the cost while satisfying the effective resistance constraint in Section 3.6.

3.1 Convex Programming Relaxation

The formulation is for the weighted problem, where each edge has a weight $w_e := 1/r_e$. We introduce a variable x_e for each edge e to indicate whether e is chosen in our subgraph. Let

$$L_x := \sum_{e \in E} x_e w_e b_e b_e^T$$

be the Laplacian matrix of the fractional solution x , and $\text{Reff}_x(s, t)$ be the s - t effective resistance of the graph with conductance $x_e w_e$ on edge $e \in E$. The following is a natural convex programming relaxation for the problem:

$$\begin{aligned} & \min_{x \in \mathbb{R}^m} && \text{Reff}_x(s, t) \\ & \text{subject to} && \sum_{e \in E} c_e x_e \leq k, \\ & && 0 \leq x_e \leq 1, \quad \forall e \in E. \end{aligned} \tag{CP}$$

This is an exact formulation if $x_e \in \{0, 1\}$ for all $e \in E$. The objective function is convex in x by Fact 2.3. The convex program can be solved in polynomial time by the ellipsoid method to inverse exponential accuracy, or by the techniques described in Reference [2] to inverse polynomial accuracy, which are both sufficient for the rounding algorithm.

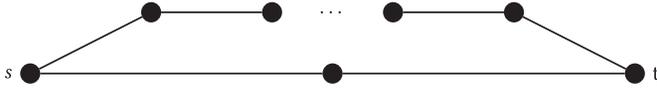


Fig. 1. Integrality gap example with arbitrary cost and unit resistance.

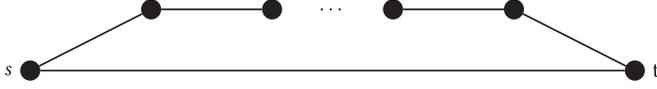


Fig. 2. Integrality gap example with arbitrary resistance and unit-cost.

3.1.1 Integrality Gap Examples. We show some limitations of the convex program for general w_e and c_e . Figure 1 shows a simple example where the integrality gap is unbounded if the cost could be arbitrary.

In this graph, the top path has length $n-2$ where each edge has cost $1/(n-2)$. The bottom path has two edges with cost 1. The resistance of each edge is 1, and the budget is $k = 1$. The integrality gap of this example is $\Omega(n)$. To see this, the integral solution can only afford the top path, and the effective resistance is $n-2$. However, the fractional solution can set $x_e = 1/2$ for each of the two bottom edges, and the effective resistance of this fractional solution is 4.

Figure 2 shows another simple example where the integrality gap is unbounded if the edge costs are the same but the resistances could be arbitrary.

In this example, the top path has length $n-1$ with each edge of resistance 1. The bottom path has only one edge with resistance R . All edges have cost 1 and the budget $k = n-2$. The integral solution can only afford the bottom path, with effective resistance R . The fractional solution can set $x_e = (n-2)/(n-1)$ for each edge in the top path, with effective resistance $O(n)$. When $R \gg n$, the integrality gap could be arbitrarily large.

Even in the unit-cost unit-resistance case, the integrality gap is unbounded if k is smaller than the s - t shortest path distance. Henceforth, in view of these observations, we assume the following in the rest of this section.

ASSUMPTION 3.1. We assume that $c_e = w_e = r_e = 1$ for every edge $e \in E$, which is the setting of the s - t effective resistance network design problem, and the budget k is at least the shortest path distance d_{st} between s and t in the input graph.

The integrality gap of the convex program is still at least two with Assumption 3.1. For a simple example, consider a graph with two vertex-disjoint s - t paths, each of length $k/2 + 1$, and the budget is k . Then the optimal integral value is $k/2 + 1$ while the optimal fractional value is close to $k/4$, and so the integrality gap gets arbitrarily close to two.

We will show that the integrality gap of the convex program is at most 8 with these assumptions. Note that just to connect s and t , then k must be at least the s - t shortest path distance. It is interesting that this small additional assumption could reduce the integrality gap from unbounded to a constant.

3.1.2 Characterization of Optimal Solutions. In the case $c_e = w_e = r_e = 1$ for all edges $e \in E$, we will prove that the electrical flow f^* supported in the optimal solution x^* to **the convex program (CP)** satisfies a crucial property about the flow-conductance ratio f_e^*/x_e^* .

LEMMA 3.2 (CHARACTERIZATION OF OPTIMAL SOLUTION). Let $G = (V, E)$ be the input graph with $c_e = w_e = 1$ for all edges $e \in E$. Let $x^* : E \rightarrow \mathbb{R}_{\geq 0}$ be an optimal solution to the convex program (CP). Let $E_F \subseteq E$ be the set of fractional edges with $0 < x_e^* < 1$, and $E_I \subseteq E$ be the set of integral edges with

$x_e^* = 1$. Let $f^* : E \rightarrow \mathbb{R}_{\geq 0}$ be the unit s - t electrical flow supported in x^* . There exists $\alpha > 0$ such that

$$f_e^* = \alpha x_e^* \forall e \in E_F \quad \text{and} \quad f_e^* \geq \alpha \forall e \in E_I.$$

PROOF. By removing edges with $x_e^* = 0$, we can assume $x_e^* > 0$ for every $e \in E$. By removing isolated vertices, we can further assume that the nonzero edges form a connected graph. So, we can write $\text{Reff}_{x^*}(s, t) = b_{st}^T L_{x^*}^\dagger b_{st}$, where L_{x^*} has rank $n - 1$ and the null space of L_{x^*} is $\text{span}(\vec{1})$. Since $b_{st} \perp \vec{1}$, we have $L_{x^*} L_{x^*}^\dagger b_{st} = b_{st}$ and $L_{x^*}^\dagger b_{st} \perp \vec{1}$, which implies that $L_{x^*}^\dagger b_{st} = (L_{x^*} + \frac{1}{n}J)^{-1} b_{st}$ where J is the all-ones matrix. Let ∂_e denote the partial derivative with respect to the variable x_e , and using the fact that $\partial A^{-1} = -A^{-1}(\partial A)A^{-1}$ (see, e.g., Reference [41]), we derive

$$\begin{aligned} \partial_e \text{Reff}_{x^*}(s, t) &= \partial_e b_{st}^T \left(L_{x^*} + \frac{1}{n}J \right)^{-1} b_{st} = -b_{st}^T \left(L_{x^*} + \frac{1}{n}J \right)^{-1} (\partial_e L_{x^*}) \left(L_{x^*} + \frac{1}{n}J \right)^{-1} b_{st} \\ &= -b_{st}^T L_{x^*}^\dagger \left(\partial_e \left(\sum_{e \in E} x_e^* w_e b_e b_e^T \right) \right) L_{x^*}^\dagger b_{st} = -b_{st}^T L_{x^*}^\dagger b_e b_e^T L_{x^*}^\dagger b_{st} = -(b_{st}^T L_{x^*}^\dagger b_e)^2, \end{aligned}$$

where we used the assumption that $w_e = 1$ for all $e \in E$. With this, we write down the KKT conditions for the convex program. Let μ be the dual variable for the budget constraint $\sum_{e \in E} c_e x_e^* \leq k$, and λ_e^+ and λ_e^- be the dual variables for the upper bound $x_e^* \leq 1$ and the nonnegative constraint $x_e^* \geq 0$, respectively. The KKT conditions states if x^* is an optimal solution to (CP), then there exist λ^+, λ^- and μ such that

$$\sum_{e \in E} x_e^* \leq k, \quad 0 \leq x_e^* \leq 1 \forall e \in E, \quad (\text{Primal feasibility})$$

$$\mu \geq 0, \quad \lambda_e^+ \geq 0 \text{ and } \lambda_e^- \geq 0 \forall e \in E, \quad (\text{Dual feasibility})$$

$$\mu \cdot \left(k - \sum_{e \in E} x_e^* \right) = 0, \quad \lambda_e^+ \cdot (x_e^* - 1) = 0 \text{ and } \lambda_e^- \cdot x_e^* = 0 \forall e \in E, \quad (\text{Complementary slackness})$$

$$(b_e^T L_{x^*}^\dagger b_{st})^2 = \lambda_e^+ - \lambda_e^- + c_e \mu = \lambda_e^+ - \lambda_e^- + \mu \forall e \in E, \quad (\text{Lagrangian optimality})$$

where we used the assumption that $c_e = 1$ for all $e \in E$. For an integral edge with $x_e^* = 1$, we have $\lambda_e^- = 0$ by the complementary slackness condition. Since $\lambda_e^+ \geq 0$, it follows from the Lagrangian optimality condition that $(b_e^T L_{x^*}^\dagger b_{st})^2 \geq \mu$. For a fractional edge with $0 < x_e^* < 1$, we have $\lambda_e^+ = \lambda_e^- = 0$ by the complementary slackness condition, and therefore $(b_e^T L_{x^*}^\dagger b_{st})^2 = \mu$ by the Lagrangian optimality condition. We can assume that $\mu > 0$. Otherwise, $\mu = 0$ implies that the flow on all fractional edges are zero, and so we can delete them from the graph without affecting the s - t effective resistance, and we have an integral solution.

Let φ be a potential vector of the electrical flow f^* supported in x^* . For an edge $e = uv \in E$,

$$\left(\frac{f_e^*}{x_e^*} \right)^2 = (\varphi(u) - \varphi(v))^2 = (b_e^T L_{x^*}^\dagger b_{st})^2,$$

where the first equality is by Ohm's law and the assumption that $w_{uv} = 1$ for all $uv \in E$, and the second equality uses that $L_{x^*} \varphi = b_{st}$ as explained in Section 2.2. The lemma then follows from the above paragraph and writing μ as α^2 . \square

The flow-conductance ratio α will be crucial in the rounding algorithm and its analysis. Intuitively, α measures the "effectiveness" of the fractional edges. A larger α means that the fractional edges make more contributions when sending electrical flows from s to t . Lemma 3.2 says all fractional edges are equally "effective", and the integral edges are at least as "effective" as those fractional edges.

The following lemma shows an upper bound on α using the budget k and the shortest path distance d_{st} between s and t .

LEMMA 3.3. *Under the conditions in Assumption 3.1, it holds that $\alpha^2 \leq d_{st}/k \leq 1$.*

PROOF. Let x^* be an optimal solution to (CP), and f^* be the unit s - t electrical flow supported in x^* . As $k \geq d_{st}$, a shortest path is a feasible solution to (CP), and thus $\text{Reff}_{x^*}(s, t) \leq d_{st}$. On the other hand, by Thomson's principle and Lemma 3.2,

$$\text{Reff}_{x^*}(s, t) = \sum_{e \in E} \frac{(f_e^*)^2}{x_e^*} = \sum_{e \in E_I} (f_e^*)^2 + \sum_{e \in E_F} \frac{(f_e^*)^2}{x_e^*} \geq \sum_{e \in E_I} \alpha^2 + \sum_{e \in E_F} \alpha^2 x_e^* = \alpha^2 \sum_{e \in E} x_e^* = \alpha^2 k,$$

where the last equality holds, since we can assume $\sum_{e \in E} x_e^* = k$ for the optimal solution x^* without loss of generality by Rayleigh's principle (or, otherwise, we have an integral optimal solution). The lemma follows by combining the upper bound and the lower bound. \square

3.2 Randomized Path-Rounding Algorithm

Our rounding algorithm uses the unit electrical flow f^* supported in the optimal solution x^* to construct an integral solution. The property of x^* and f^* in Lemma 3.2 is the key of the success of the rounding algorithm.

Remark. In this subsection, we present the rounding algorithm assuming we have an exact optimal solution x^* , to focus on the main ideas of the analysis. However, standard convex optimization techniques (e.g., ellipsoid method) can only guarantee to return an approximately optimal solution x to (CP) such that $\text{Reff}_x(s, t) \leq \text{Reff}_{x^*}(s, t) + \varepsilon$ in $O(\text{poly}(n, m) \log \frac{1}{\varepsilon})$ time. We will explain the modifications needed for rounding an approximately optimal solution in Appendix A.

The algorithm will first decompose the flow f^* as a convex combination of flow paths, and then randomly choose the flow paths and return the union of the chosen flow paths as our solution. The following lemma about flow decomposition is by the standard argument to remove one (fractional) flow path at a time, which holds for any unit directed acyclic s - t flow.

LEMMA 3.4 (FLOW DECOMPOSITION). *Given the unit s - t electrical flow f , there is a polynomial time algorithm to find a set \mathcal{P} of s - t paths with $|\mathcal{P}| \leq |E|$ such that the flow vector $f : E \rightarrow \mathbb{R}_{\geq 0}$ can be written as a convex combination of the characteristic vectors of the paths in \mathcal{P} , i.e.,*

$$f = \sum_{p \in \mathcal{P}} v_p \cdot \chi_p \quad \text{and} \quad \sum_{p \in \mathcal{P}} v_p = 1 \quad \text{and} \quad v_p > 0 \text{ for each } p \in \mathcal{P},$$

where $\chi_p \in R^{|E|}$ is the characteristic vector of the path p with one on each edge $e \in p$ and zero otherwise.

With the flow decomposition, we are ready to present the rounding algorithm.

Randomized Path Rounding Algorithm

- (1) Let x^* be an optimal solution to the convex program (CP). Let f^* be the unit s - t electrical flow supported in x^* . Let α be the flow-conductance ratio defined in Lemma 3.2.
- (2) Compute a flow decomposition \mathcal{P} of f^* as defined in Lemma 3.4.
- (3) For i from 1 to $T := \lfloor 1/\alpha \rfloor$ do
 - Let P_i be a random path from \mathcal{P} where each path $p \in \mathcal{P}$ is sampled with probability v_p .
- (4) Return the subgraph H formed by the edge set $\cup_{i=1}^T P_i$.

The following lemma shows that the rounding algorithm will always return a non-empty subgraph.

LEMMA 3.5. *Suppose the input instance satisfies the conditions in Assumption 3.1. Let x^* be an optimal solution to (CP) and $\alpha > 0$ be the flow-conductance ratio as defined in Lemma 3.2. Then*

$$\frac{1}{\alpha} \geq T \geq \frac{1}{2\alpha} > 0.$$

PROOF. Since we assumed that the budget k is at least the length d_{st} of a shortest s - t path, it follows from Lemma 3.3 that $\alpha \leq 1$. This implies that

$$\frac{1}{\alpha} \geq T = \left\lfloor \frac{1}{\alpha} \right\rfloor \geq \max \left\{ 1, \frac{1}{\alpha} - 1 \right\} \implies 1 \geq T\alpha \geq \max \{ \alpha, 1 - \alpha \} \geq \frac{1}{2}.$$

□

3.3 Bicriteria Approximation

The analysis of the approximation guarantee goes as follows. First, we show that the expected number of edges in the returned subgraph H is at most the budget k . Then, we prove that the expected effective resistance of the returned subgraph is at most two times that of the optimal fractional solution. Both of these steps use the flow-conductance ratio α crucially. These combine to show that the randomized path rounding algorithm is a constant factor bicriteria approximation algorithm.

Let x^* be an optimal solution to (CP). Let E_F and E_I be the set of fractional edges and integral edges in x^* . We assume that each edge $e \in E_I$ will be included in the subgraph H returned by the rounding algorithm. We focus on bounding the number of edges in E_F that will be included in H .

LEMMA 3.6 (EXPECTED BUDGET). *Let x^* be an optimal solution to (CP) when $w_e = 1$ for all edges $e \in E$. Let X_e be an indicator variable of whether e is included in the returned subgraph H by the rounding algorithm. Then,*

$$\mathbb{E} \left[\sum_{e \in E_F} X_e \right] \leq T\alpha \sum_{e \in E_F} x_e^* \leq \sum_{e \in E_F} x_e^*.$$

PROOF. Note that an edge e is contained in P_i with probability $\sum_{p \in \mathcal{P}: p \ni e} v_p$. By the union bound, an edge e is included in the returned subgraph H by the rounding algorithm with probability

$$\mathbb{P}(X_e = 1) \leq \sum_{i=1}^T \sum_{p \in \mathcal{P}: p \ni e} v_p = T \sum_{p \in \mathcal{P}: p \ni e} v_p = T f_e^*, \quad (3.1)$$

where the last equality holds by the property of the flow decomposition \mathcal{P} of the electrical flow f^* in Lemma 3.4.

By Lemma 3.2, $f_e^* = \alpha x_e^*$ for each fractional edge $e \in E_F$, and this implies that

$$\mathbb{P}(X_e = 1) \leq T f_e^* = T\alpha x_e^* \quad \forall e \in E_F.$$

Therefore,

$$\mathbb{E} \left[\sum_{e \in E_F} X_e \right] = \sum_{e \in E_F} \mathbb{P}(X_e = 1) \leq T\alpha \sum_{e \in E_F} x_e^* = \left\lfloor \frac{1}{\alpha} \right\rfloor \alpha \sum_{e \in E_F} x_e^* \leq \sum_{e \in E_F} x_e^*.$$

□

The key step is to show that $\mathbb{E}[\text{Reff}_H(s, t)] \leq 2\text{Reff}_{x^*}(s, t)$. To prove this, we construct a unit s - t flow F and show that $\mathbb{E}[\mathcal{E}_H(F)] \leq 2\text{Reff}_{x^*}(s, t)$, and hence by Thomson's principle $\mathbb{E}[\text{Reff}_H(s, t)] \leq \mathbb{E}[\mathcal{E}_H(F)] \leq 2\text{Reff}_{x^*}(s, t)$. To construct the flow F , the idea is to follow the ratio α in the fractional solution x^* and send α units of flow on each path P_i selected.

LEMMA 3.7 (EXPECTED EFFECTIVE RESISTANCE). *Suppose the input instance satisfies the conditions in Assumption 3.1. Let x^* be an optimal solution to (CP) and f^* be the unit s - t electrical flow supported in x^* . The expected s - t effective resistance of the subgraph H returned by the rounding algorithm is*

$$\mathbb{E}[\text{Reff}_H(s, t)] \leq \left(1 - \frac{1}{T} + \frac{1}{T\alpha}\right) \cdot \mathcal{E}_{x^*}(f^*) = \left(1 - \frac{1}{T} + \frac{1}{T\alpha}\right) \cdot \text{Reff}_{x^*}(s, t) \leq 2\text{Reff}_{x^*}(s, t).$$

PROOF. Consider the flow vector $F : E \rightarrow \mathbb{R}_{\geq 0}$ defined by sending α units of flow on each path P_i chosen by the rounding algorithm, i.e., the random variable $F = \sum_{i=1}^T \alpha \cdot \chi_{P_i}$ with $F_e = \alpha \cdot \#\{P_i \mid 1 \leq i \leq T, P_i \ni e\}$ for each edge $e \in E$. We would like to upper bound the expected energy $\mathcal{E}_H(F)$ to upper bound $\text{Reff}_H(s, t)$.

Each P_i is a random s - t path sampled from the flow decomposition \mathcal{P} of the flow vector $f^* : E \rightarrow \mathbb{R}_{\geq 0}$ of the unit s - t electrical flow supported in x^* , and $\chi_{P_i} \in \mathbb{R}^m$ is its characteristic vector with expected value

$$\mathbb{E}[\chi_{P_i}] = \sum_{p \in \mathcal{P}} v_p \cdot \chi_p = f^*.$$

Since each edge in H is of conductance one, the expected energy of F in H is

$$\mathbb{E}[\mathcal{E}_H(F)] = \mathbb{E}\left[\sum_{e \in E} F_e^2\right] = \mathbb{E}[\langle F, F \rangle] = \mathbb{E}\left[\left\langle \sum_{i=1}^T \alpha \cdot \chi_{P_i}, \sum_{j=1}^T \alpha \cdot \chi_{P_j} \right\rangle\right] = \sum_{i=1}^T \sum_{j=1}^T \alpha^2 \cdot \mathbb{E}[\langle \chi_{P_i}, \chi_{P_j} \rangle].$$

As each path P_i is sampled independently, for $i \neq j$,

$$\mathbb{E}[\langle \chi_{P_i}, \chi_{P_j} \rangle] = \langle \mathbb{E}[\chi_{P_i}], \mathbb{E}[\chi_{P_j}] \rangle = \langle f^*, f^* \rangle = \sum_{e \in E} (f_e^*)^2.$$

For $i = j$,

$$\mathbb{E}[\langle \chi_{P_i}, \chi_{P_i} \rangle] = \sum_{p \in \mathcal{P}} v_p \langle \chi_p, \chi_p \rangle = \sum_{p \in \mathcal{P}} v_p \sum_{e \in p} 1 = \sum_{e \in E} \sum_{p \in \mathcal{P}: p \ni e} v_p = \sum_{e \in E} f_e^*,$$

where the last equality follows from the property of the flow decomposition in Lemma 3.4. Combining these two terms, it follows that

$$\mathbb{E}[\mathcal{E}_H(F)] = \alpha^2 T \sum_{e \in E} f_e^* + \alpha^2 T(T-1) \sum_{e \in E} (f_e^*)^2.$$

Thomson's principle states that the $\text{Reff}_H(s, t)$ is upper bounded by the energy of any one unit s - t flow. Note that F is an s - t flow of $T\alpha$ units, and $T\alpha > 0$ by Lemma 3.5. Scaling F to a one unit s - t flow by dividing the flow on each edge by $T\alpha$ gives an upper bound on

$$\begin{aligned} \mathbb{E}[\text{Reff}_H(s, t)] &\leq \frac{\mathbb{E}[\mathcal{E}_H(F)]}{T^2 \alpha^2} = \frac{1}{T} \sum_{e \in E} f_e^* + \left(1 - \frac{1}{T}\right) \sum_{e \in E} (f_e^*)^2 \\ &\leq \frac{1}{T\alpha} \sum_{e \in E} \frac{(f_e^*)^2}{x_e^*} + \left(1 - \frac{1}{T}\right) \sum_{e \in E} \frac{(f_e^*)^2}{x_e^*} \\ &= \left(1 - \frac{1}{T} + \frac{1}{T\alpha}\right) \cdot \mathcal{E}_{x^*}(f^*) \\ &= \left(1 - \frac{1}{T} + \frac{1}{T\alpha}\right) \cdot \text{Reff}_{x^*}(s, t), \end{aligned} \tag{3.2}$$

where the second inequality follows from Lemma 3.2 that $f_e^*/x_e^* \geq \alpha$ for every edge $e \in E$ and also $x_e^* \leq 1$ for every edge $e \in E$, and the last equality is from Thomson's principle that $\text{Reff}_{x^*}(s, t) = \mathcal{E}_{x^*}(f^*)$. Finally, notice that $1 - 1/T + 1/(T\alpha) \leq 2$ as $1/\alpha - 1 \leq \lfloor 1/\alpha \rfloor = T$. \square

Combining Lemmas 3.6 and 3.7, it follows from a simple application of Markov's inequality that there is an outcome of the randomized path-rounding algorithm that uses at most $2k$ edges with s - t effective resistance at most $4\text{Reff}_{x^*}(s, t)$. In the following, we apply Markov's inequality more carefully to show that the success probability is at least $\Omega(\alpha)$. In the next subsection, we will argue that α can be assumed to be $\Omega(1/m)$ and so the path-rounding algorithm is a randomized polynomial time algorithm.

THEOREM 3.8 (BICRITERIA APPROXIMATION). *Suppose the input instance satisfies the conditions in Assumption 3.1. Let x^* be an optimal solution to (CP). Given x^* , the randomized path rounding algorithm will return a subgraph H with at most $2k$ edges and $\text{Reff}_H(s, t) \leq 4\text{Reff}_{x^*}(s, t)$ with probability at least $\Omega(\alpha)$.*

PROOF. First, we bound the probability that the subgraph H has more than $2k$ edges. Let X_e be an indicator variable of whether the edge e is included in the returned subgraph H . Recall that E_F and E_I denote the set of fractional edges and integral edges in x^* , respectively. We assume pessimistically that all edges in E_I will be included in the subgraph H returned by the rounding algorithm. Then, by Markov's inequality and Lemma 3.6,

$$\Pr\left(\sum_{e \in E} X_e > 2k\right) \leq \Pr\left(\sum_{e \in E_F} X_e > 2k - |E_I|\right) \leq \frac{\mathbb{E}\left[\sum_{e \in E_F} X_e\right]}{2k - |E_I|} \leq \frac{T\alpha \sum_{e \in E_F} x_e^*}{2k - |E_I|} \leq \frac{T\alpha}{2},$$

where the last inequality is by $\sum_{e \in E_F} x_e^* \leq k - |E_I|$.

Next, we bound the probability that $\text{Reff}_H(s, t) > 4\text{Reff}_{x^*}(s, t)$. By Markov's inequality and Lemma 3.7,

$$\Pr\left(\text{Reff}_H(s, t) > 4\text{Reff}_{x^*}(s, t)\right) \leq \frac{1}{4} \left(1 - \frac{1}{T} + \frac{1}{T\alpha}\right) = \frac{T\alpha + 1}{4T\alpha} - \frac{1}{4T} \leq \frac{T\alpha + 1}{4T\alpha} - \Omega(\alpha),$$

where the last inequality is because $T = \lfloor 1/\alpha \rfloor \leq 1/\alpha$.

To prove the theorem, it remains to show that

$$\frac{T\alpha}{2} + \frac{T\alpha + 1}{4T\alpha} \leq 1 \iff 2(T\alpha)^2 - 3(T\alpha) + 1 = (2T\alpha - 1)(T\alpha - 1) \leq 0,$$

which follows from Lemma 3.5. \square

3.4 Constant Factor Approximation

We showed that the randomized path rounding algorithm is a bicriteria approximation algorithm. To achieve a true approximation algorithm, a natural idea is to scale down the budget from k to $k/2$ and apply the randomized path rounding algorithm. The following lemma takes care of the case of $k/2 < d_{st}$, when the shortest path assumption does not hold after scaling, by showing that simply returning a shortest s - t path is already a good enough approximation.

LEMMA 3.9. *When the budget k is at least the length d_{st} of a shortest s - t path, any s - t shortest path is a (k/d_{st}) -approximate solution for the s - t effective resistance network design problem.*

PROOF. When $k \geq d_{st}$, an s - t shortest path is a feasible solution to the problem with s - t effective resistance at most d_{st} . To prove the lemma, we will show that $\text{Reff}_x(s, t) \geq d_{st}^2/k$ for any feasible solution x to (CP), and so an s - t shortest path is already a (k/d_{st}) -approximation.

Let G_x be the graph G with fractional conductance x_e on each edge $e \in E$. To show a lower bound on $\text{Reff}_x(s, t)$, we identify the vertices in G_x to form a path graph P_x as follows: For each $i \geq 0$, let U_i be the set of vertices in G with shortest path distance i to s , where the shortest path distance is defined where each edge in G is of length one. First, for each $0 \leq i \leq d_{st} - 1$, we identify the vertices in U_i to a single vertex u_i . Then, we identify all the vertices in $\cup_{i \geq d_{st}} U_i$ to a single vertex $u_{d_{st}}$. The path graph P_x has vertex set $\{u_0, \dots, u_{d_{st}}\}$ and edge set $\{ab \in E \mid a \in U_i \text{ and } b \in U_{i+1} \text{ for } 0 \leq i \leq d_{st} - 1\}$. For each edge e in P_x , its conductance x_e in P_x is the same as that in G_x . As an electrical network, identifying two vertices uv is equivalent to adding an edge of resistance zero between u and v . So, it follows from Rayleigh's monotonicity principle (Fact 2.2) that $\text{Reff}_{G_x}(s, t) \geq \text{Reff}_{P_x}(u_0, u_{d_{st}})$ as $s \in U_0$ and $t \in U_{d_{st}}$.

As P_x is a series-parallel graph, we can compute $\text{Reff}_{P_x}(s, t)$ directly. For each $1 \leq i \leq d_{st}$, let E_i be the set of parallel edges connecting u_{i-1} and u_i in P_x , and $c_i = \sum_{e \in E_i} x_e$ be the effective conductance between u_{i-1} and u_i in P_x . Then, by Fact 4.3,

$$\text{Reff}_{P_x}(u_{i-1}, u_i) = \frac{1}{c_i} \quad \text{and} \quad \text{Reff}_{P_x}(u_0, u_{d_{st}}) = \sum_{i=1}^{d_{st}} \text{Reff}_{P_x}(u_{i-1}, u_i) = \sum_{i=1}^{d_{st}} \frac{1}{c_i}.$$

Note that $\sum_{i=1}^{d_{st}} c_i = \sum_{i=1}^{d_{st}} \sum_{e \in E_i} x_e \leq \sum_{e \in E} x_e \leq k$ for any feasible solution x . Using Cauchy-Schwarz inequality,

$$d_{st} = \sum_{i=1}^{d_{st}} \sqrt{c_i} \cdot \frac{1}{\sqrt{c_i}} \leq \sqrt{\sum_{i=1}^{d_{st}} c_i} \cdot \sqrt{\sum_{i=1}^{d_{st}} \frac{1}{c_i}} \leq \sqrt{k} \cdot \sqrt{\text{Reff}_{P_x}(u_0, u_{d_{st}})}.$$

Therefore, we conclude that

$$\text{Reff}_{G_x}(s, t) \geq \text{Reff}_{P_x}(u_0, u_{d_{st}}) \geq d_{st}^2/k. \quad (3.3)$$

□

We are ready to prove our main approximation result.

THEOREM 3.10. *Suppose the input instance satisfies the conditions in Assumption 3.1. There is a polynomial time 8-approximation algorithm for the s - t effective resistance network design problem.*

PROOF. If the budget $k \leq 2d_{st}$, then Lemma 3.9 shows that simply returning an s - t shortest path would give a 2-approximation. Henceforth, we assume $k \geq 2d_{st}$.

Let $\text{opt}(k)$ be the objective value of an optimal solution x^* to the convex program (CP) with budget k , so $\text{Reff}_{x^*}(s, t) = \text{opt}(k)$. As $\frac{1}{2}x^*$ is a feasible solution to (CP) with budget $\frac{1}{2}k$, by Thomson's principle,

$$\text{opt}\left(\frac{k}{2}\right) \leq \text{Reff}_{\frac{1}{2}x^*}(s, t) = b_{st}^T \left(\sum_{e \in E} \frac{x_e^*}{2} b_e b_e^T \right)^\dagger b_{st} = 2b_{st}^T \left(\sum_{e \in E} x_e^* b_e b_e^T \right)^\dagger b_{st} = 2\text{Reff}_{x^*}(s, t) = 2\text{opt}(k).$$

Given the original budget $k \geq 2d_{st}$, our algorithm is to find an optimal solution z^* to (CP) with budget $k/2 \geq d_{st}$, and use the path-rounding algorithm with input z^* to return a subgraph H . By Theorem 3.8, with probability $\Omega(\alpha)$, the subgraph H satisfies

$$|E(H)| \leq 2 \sum_{e \in E} z_e^* \leq 2 \left(\frac{k}{2} \right) = k \quad \text{and} \quad \text{Reff}_H(s, t) \leq 4\text{opt}\left(\frac{k}{2}\right) \leq 8\text{opt}(k),$$

and so H is an 8-approximate solution to the s - t effective resistance network design problem.

Finally, we consider the time complexity of the algorithm. The number of iterations in the path rounding algorithm is $O(1/\alpha)$, and we need to run the path rounding algorithm $O(1/\alpha)$ times to

boost the success probability to a constant. This is a randomized polynomial time algorithm when $\alpha = \Omega(1/m)$.

In the following, we show that when $\alpha \leq 1/(4m)$, it is easy to obtain a 2-approximate solution without running the path-rounding algorithm. Let x^* be an optimal solution to (CP) with budget k , and f^* be the unit s - t electrical flow supported in x^* . Let \mathcal{P} be the flow decomposition of f^* as in Lemma 3.4. We call a path $p \in \mathcal{P}$ an integral path if every edge $e \in p$ has $x_e^* = 1$; otherwise, we call p a fractional path. When $\alpha \leq 1/(4m)$, we simply return the union of all integral paths as our solution H . Clearly, H has at most k edges as it only contains integral edges. Next, we bound $\text{Reff}_H(s, t)$ by the energy of the flow supported in the integral paths. By Lemma 3.2, an edge e with $x_e^* < 1$ has $f_e^* = \alpha x_e^* < \alpha \leq 1/(4m)$. This implies that each fractional path p has flow value $v_p \leq 1/(4m)$ in the flow decomposition. Since \mathcal{P} has at most m paths (Lemma 3.4), the total flow in the fractional paths is at most $1/4$, and thus the total flow in the integral paths is at least $3/4$. By scaling the flow supported in the integral paths to a one unit s - t flow, we see that

$$\text{Reff}_H(s, t) \leq \frac{\mathcal{E}_{x^*}(f^*)}{(3/4)^2} \leq 2\mathcal{E}_{x^*}(f^*) = 2\text{Reff}_{x^*}(s, t).$$

To summarize, in all cases including $k < 2d_{st}$ and $\alpha \leq 1/(4m)$, there is a polynomial time algorithm to return an 8-approximate solution to the s - t effective resistance network design problem. \square

We make two remarks about improvements of Theorem 3.10.

Remark 3.11 (Approximation Ratio). The analysis of the 8-approximation algorithm is not tight. By a more careful analysis of the expected energy in Lemma 3.7 and the short path idea used in Section 3.5, we can show that the approximation guarantee of the same algorithm in Theorem 3.10 is less than 5. However, the analysis is quite involved and not very insightful, so we have decided to omit those details and only keep the current analysis.

Remark 3.12 (Deterministic Algorithm). Using the standard pessimistic estimator technique [42], we can derandomize the path-rounding algorithm to obtain a deterministic 8-approximation algorithm. The analysis is standard, and we omit the details that would take a few pages.

3.5 The Large Budget Case

In this subsection, we show how to modify the algorithm in Theorem 3.10 to achieve a better approximation ratio when the budget is much larger than the s - t shortest path distance.

The observation is that when $k \gg d_{st}$, then α is small by Lemma 3.3, and so there are many iterations in the path-rounding algorithm. Since each iteration is independent, we can use Chernoff-Hoeffding's bound to prove a stronger bound on the probability that the number of edges in the returned solution is significantly more than k (which outperforms the bound proved in Lemma 3.6 using Markov's inequality). We can then show that the expected s - t effective resistance is close to two times the optimal value by arguments similar to the proof of Lemma 3.7.

Modified Rounding Algorithm. For our analysis, we slightly modify the path-rounding algorithm to ignore "long" paths in the flow decomposition, so that we have a worst case bound to apply Chernoff-Hoeffding's bound. Unlike the flow decomposition in Lemma 3.4, the short-path flow decomposition definition is specific to the electrical flow of an optimal solution to (CP). In the following definition, c is a parameter (not to be confused with the edge cost c_e in the weighted problem) that will be set to be $1/\varepsilon > 1$ to achieve a $(2 + O(\varepsilon))$ -approximation.

Definition 3.13 (Short-Path Decomposition of Electrical Flow of Optimal Solution). Let x^* be an optimal solution to the convex program (CP). Let f^* be the unit s - t electrical flow supported in x^* . Let α be the flow-conductance ratio defined in Lemma 3.2.

Let \mathcal{P}^* be a flow decomposition of f^* as defined in Lemma 3.4. Let $x_F^* := \sum_{e \in E_F} x_e^*$ be the total fractional value on the fractional edges E_F in the optimal solution x^* .

We call a path $p \in \mathcal{P}^*$ a long path if p has at least $c\alpha x_F^*$ edges in E_F , i.e., $|p \cap E_F| \geq c\alpha x_F^*$. Otherwise, we call a path $p \in \mathcal{P}^*$ a short path.

Let $\mathcal{P} := \{p \in \mathcal{P}^* \mid p \text{ is a short path}\}$ be the collection of short paths in \mathcal{P}^* . Let $f_{\mathcal{P}} := \sum_{p \in \mathcal{P}} v_p \chi_p$ be the s - t flow defined by the short paths, and $v_{\mathcal{P}} := \sum_{p \in \mathcal{P}} v_p$ be the total flow value of $f_{\mathcal{P}}$.

The modified algorithm is very similar to the randomized path-rounding algorithm in Section 3.3. The only difference is that we only sample the paths in the short-path flow decomposition in Definition 3.13, and we adjust the sampling probability of a path p to $v_p/v_{\mathcal{P}}$ so that the sum is one.

Randomized Short Path Rounding Algorithm

- (1) Let x^* be an optimal solution to the convex program (CP). Let f^* be the unit s - t electrical flow supported in x^* . Let α be the flow-conductance ratio defined in Lemma 3.2.
- (2) Compute a short-path flow decomposition \mathcal{P} of f^* as described in Definition 3.13.
- (3) For i from 1 to $T = \lfloor 1/\alpha \rfloor$ do
 - Let P_i be a random path from \mathcal{P} where each path $p \in \mathcal{P}$ is sampled with probability $v_p/v_{\mathcal{P}}$.
- (4) Return the subgraph H formed by the edge set $\cup_{i=1}^T P_i$.

The following simple lemma shows that the total flow on the long paths is negligible when c is large, which will be useful in the analysis.

LEMMA 3.14. *For the short-path flow decomposition in Definition 3.13, $v_{\mathcal{P}} \geq 1 - \frac{1}{c}$.*

PROOF. Using $\alpha x_e^* = f_e^*$ for $e \in E_F$ from Lemma 3.2 and the properties of the flow decomposition \mathcal{P}^* of f^* in Lemma 3.4,

$$\alpha x_F^* = \sum_{e \in E_F} f_e^* = \sum_{p \in \mathcal{P}^*} v_p \cdot |p \cap E_F| \geq \sum_{p \in \mathcal{P}^* - \mathcal{P}} v_p \cdot |p \cap E_F| \geq c\alpha x_F^* \sum_{p \in \mathcal{P}^* - \mathcal{P}} v_p = c\alpha x_F^* (1 - v_{\mathcal{P}}),$$

where the last inequality is by the definition of long paths and the last equality is because f^* is a unit s - t flow. \square

Analysis of Approximation Guarantee. First, we consider the expected s - t effective resistance of the returned subgraph H . Essentially the same proof as in Lemma 3.7 can give us the following lemma. We only provide a sketched proof here.

LEMMA 3.15. *Suppose the input instance satisfies the conditions in Assumption 3.1. Let x^* be an optimal solution to (CP) and f^* be the unit s - t electrical flow supported in x^* . The expected s - t effective resistance of the subgraph H returned by the randomized short path rounding algorithm is*

$$\mathbb{E}[\text{Reff}_H(s, t)] \leq \frac{2}{v_{\mathcal{P}}^2} \mathcal{E}_{x^*}(f^*) = \frac{2}{v_{\mathcal{P}}^2} \text{Reff}_{x^*}(s, t),$$

where \mathcal{P} is the short-path flow decomposition of f^* as described in Definition 3.13.

PROOF SKETCH. With the same idea in the proof of Lemma 3.7, we can define a random s - t flow F with expectation $\mathbb{E}[F] = f_{\mathcal{P}}/v_{\mathcal{P}}$. Then, going through the same analysis until Equation (3.2), it holds that

$$\mathbb{E}[\text{Reff}_H(s, t)] \leq \frac{1}{T} \sum_{e \in E} \frac{f_{\mathcal{P}}(e)}{v_{\mathcal{P}}} + \left(1 - \frac{1}{T}\right) \sum_{e \in E} \frac{f_{\mathcal{P}}(e)^2}{v_{\mathcal{P}}^2}.$$

We notice that $f_{\mathcal{P}} \leq f^*$ entry-wise by the definition of the short-path decomposition and $v_{\mathcal{P}} \leq 1$. Thus, the expected s - t effective resistance is further upper bounded by

$$\mathbb{E}[\text{Reff}_H(s, t)] \leq \frac{1}{v_{\mathcal{P}}^2} \left(\frac{1}{T} \sum_{e \in E} f_e^* + \left(1 - \frac{1}{T}\right) \sum_{e \in E} (f_e^*)^2 \right) \leq \frac{2}{v_{\mathcal{P}}^2} \mathcal{E}_{x^*}(f^*),$$

where the last inequality follows by the same remaining analysis as in Lemma 3.7. \square

The main difference in the analysis is to apply the following Hoeffding's inequality (instead of Markov's inequality) to bound the probability that the returned subgraph has significantly more than k edges.

FACT 3.16 (HOEFFDING'S INEQUALITY). *Let $X_1, \dots, X_n \in [0, M]$ be independent random variables. Let $X = \sum_{i=1}^n X_i$, and $\mu = \mathbb{E}[X]$, then for any $\delta > 0$,*

$$\Pr(X \geq (1 + \delta)\mu) \leq \exp\left(-\frac{2\delta^2\mu^2}{nM^2}\right).$$

LEMMA 3.17. *Suppose the input instance satisfies the conditions in Assumption 3.1. Let x^* be an optimal solution to (CP) and f^* be the unit s - t electrical flow supported in x^* . Let H be the subgraph returned by the randomized short path rounding algorithm given x^* as input, and $|E(H)|$ be the number of edges in H . Then, for any $\delta > 0$,*

$$\Pr(|E(H)| \geq (1 + \delta)k) \leq \exp\left(-\frac{2\delta^2}{c^2\alpha}\right),$$

where c is the parameter in the short-path flow decomposition in Definition 3.13 and α is the flow-conductance ratio of f^* and x^* as defined in Lemma 3.2.

PROOF. As in Lemma 3.6, we assume pessimistically that all integral edges E_I will be included in H , and so we focus on the fractional edges E_F . Let $X_{i,e}$ be the indicator variable of whether the edge e is sampled in the i th iteration of the short path rounding algorithm, and $X_{i,F} := \sum_{e \in E_F} X_{i,e}$ be the total number of fractional edges sampled in the i th iteration. Let X_F be the total number of fractional edges in H . Note that $X_F \leq \sum_{i=1}^T X_{i,F}$, since if some fractional edge was sampled in different iterations, then we only count it once in X_F . By linearity of expectation, $\mathbb{E}[X_F] \leq \sum_{i=1}^T \mathbb{E}[X_{i,F}]$.

Let \mathcal{P}^* be the flow path decomposition of f^* in Lemma 3.4, and \mathcal{P} be the short-path flow decomposition of f^* as described in Definition 3.13. For an edge e , recall that $(f_{\mathcal{P}})_e := \sum_{p \in \mathcal{P}: p \ni e} v_p$ is the total flow value on e from the short paths in \mathcal{P} . As we scaled the probability of each path by $1/v_{\mathcal{P}}$ in the rounding algorithm, the probability that edge e is sampled in the i th iteration is $(f_{\mathcal{P}})_e/v_{\mathcal{P}}$. Let $\bar{f}_e := f_e^* - (f_{\mathcal{P}})_e$ be the total flow value on e from the long paths in $\mathcal{P}^* - \mathcal{P}$. The expected value of $X_{i,F}$ is

$$\mathbb{E}[X_{i,F}] = \sum_{e \in E_F} \mathbb{E}[X_{i,e}] = \sum_{e \in E_F} \frac{(f_{\mathcal{P}})_e}{v_{\mathcal{P}}} = \sum_{e \in E_F} \frac{f_e^* - \bar{f}_e}{v_{\mathcal{P}}} = \sum_{e \in E_F} \frac{\alpha x_e^* - \bar{f}_e}{v_{\mathcal{P}}}.$$

By the definition of the long paths,

$$\sum_{e \in E_F} \bar{f}_e = \sum_{e \in E_F} \sum_{p \in \mathcal{P}^* - \mathcal{P}} v_p = \sum_{p \in \mathcal{P}^* - \mathcal{P}} v_p \cdot |p \cap E_F| \geq c\alpha x_F^* \sum_{p \in \mathcal{P}^* - \mathcal{P}} v_p = c\alpha x_F^*(1 - v_{\mathcal{P}}),$$

where we recall that $x_F^* = \sum_{e \in E_F} x_e^*$. Therefore,

$$\mathbb{E}[X_{i,F}] = \sum_{e \in E_F} \frac{\alpha x_e^* - \bar{f}_e}{v_{\mathcal{P}}} \leq \alpha x_F^* \cdot \frac{1 - c + cv_{\mathcal{P}}}{v_{\mathcal{P}}} = \alpha x_F^* \cdot \left(c - \frac{c-1}{v_{\mathcal{P}}}\right) \leq \alpha x_F^*,$$

where the last inequality uses that $v_{\mathcal{P}} \leq 1$ and $c > 1$. It follows that $\mathbb{E}[X_F] \leq T\alpha x_F^* \leq x_F^*$.

As each iteration is independent, the random variables $X_{i,F}$ for $1 \leq i \leq T$ are independent. Since we only use short paths, the maximum value of each $X_{i,F}$ is at most $c\alpha x_F^*$. So, we can apply Hoeffding's inequality to show that

$$\Pr\left(X_F \geq (1 + \delta)x_F^*\right) \leq \exp\left(-\frac{2\delta^2(x_F^*)^2}{Tc^2\alpha^2(x_F^*)^2}\right) \leq \exp\left(-\frac{2\delta^2}{c^2\alpha}\right).$$

Let X_I be the total number of integral edges in H . As $X_I \leq |E_I|$, we conclude that

$$\Pr\left(|E(H)| \geq (1 + \delta)k\right) = \Pr\left(X_I + X_F \geq (1 + \delta)(|E_I| + x_F^*)\right) \leq \Pr\left(X_F \geq (1 + \delta)x_F^*\right) \leq \exp\left(-\frac{2\delta^2}{c^2\alpha}\right). \quad \square$$

As in Section 3.3, we can combine Lemmas 3.17 and 3.15 to show that the randomized short path rounding algorithm is a bicriteria approximation algorithm.

THEOREM 3.18. *Suppose the input instance satisfies the conditions in Assumption 3.1. Suppose further that $k \geq d_{st}/\varepsilon^{10}$, where $\varepsilon > 0$ is an error parameter satisfying $\varepsilon \leq \eta$ for a small constant η . Let x^* be an optimal solution to (CP). Given x^* , the randomized short path rounding algorithm with $c = 1/\varepsilon$ will return a subgraph H with at most $(1 + \varepsilon)k$ edges and $\text{Reff}_H(s, t) \leq (2 + 10\varepsilon) \cdot \text{Reff}_{x^*}(s, t)$ with probability at least ε .*

PROOF. The additional assumption $k \geq d_{st}/\varepsilon^{10}$ implies that $\alpha \leq \varepsilon^5$ by Lemma 3.3. Setting $c = 1/\varepsilon$ and $\delta = \varepsilon$, it follows from Lemma 3.17 that

$$\Pr(|E(H)| \geq (1 + \varepsilon)k) \leq \exp\left(-\frac{2\varepsilon^2}{c^2\alpha}\right) \leq \exp\left(-\frac{2}{\varepsilon}\right) < \varepsilon,$$

where the last inequality holds for $\varepsilon > 0$.

Since $c = 1/\varepsilon$, Lemma 3.14 implies that $v_\varphi \geq 1 - \varepsilon$ for the short-path flow decomposition in Definition 3.13. Using Markov's inequality and Lemma 3.15, for sufficiently small ε , we have

$$\begin{aligned} \Pr\left(\text{Reff}_H(s, t) \geq (2 + 10\varepsilon) \cdot \text{Reff}_{x^*}(s, t)\right) &\leq \frac{\mathbb{E}[\text{Reff}_H(s, t)]}{(2 + 10\varepsilon) \cdot \text{Reff}_{x^*}(s, t)} \\ &\leq \frac{2}{v_\varphi^2(2 + 10\varepsilon)} \leq \frac{2}{(1 - \varepsilon)^2(2 + 10\varepsilon)} < 1 - 2\varepsilon. \end{aligned}$$

Therefore, with probability at least ε , the subgraph H returned by the randomized short path rounding algorithm satisfies both properties. \square

Using the same arguments as in Section 3.4, we can turn the above bicriteria approximation algorithm into a true approximation algorithm.

THEOREM 3.19. *Suppose the input instance satisfies the conditions in Assumption 3.1. Suppose further that $k \geq 2d_{st}/\varepsilon^{10}$, where $\varepsilon > 0$ is an error parameter satisfying $\varepsilon \leq \eta$ for a small constant η . There is a polynomial time $(2 + O(\varepsilon))$ -approximation algorithm for the s - t effective resistance network design problem.*

PROOF. As in the proof of Theorem 3.10, we apply the bicriteria approximation algorithm in Theorem 3.18 with input x^* , an optimal solution to (CP) with the scaled-down budget $k/(1 + \varepsilon)$, to return a subgraph H . As the new budget $k/(1 + \varepsilon)$ is still greater than d_{st}/ε^{10} , by Theorem 3.18, with probability at least ε the subgraph H satisfies $|E(H)| \leq (1 + \varepsilon)k/(1 + \varepsilon) = k$ and

$$\text{Reff}_H(s, t) \leq (2 + O(\varepsilon)) \cdot \text{opt}\left(\frac{k}{1 + \varepsilon}\right) \leq (2 + O(\varepsilon))(1 + \varepsilon) \cdot \text{opt}(k) \leq (2 + O(\varepsilon)) \cdot \text{opt}(k),$$

where we used the notations and arguments in Theorem 3.10.

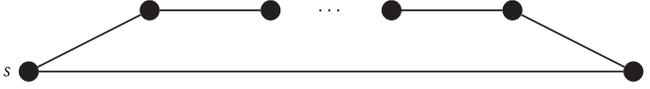


Fig. 3. $\Omega(n)$ integrality gap example.

For the time complexity, note that $\alpha \leq \varepsilon^5$ by Lemma 3.3 and the large budget assumption, and so we can assume that $\varepsilon^5 \geq \alpha \geq 1/(4m)$, as otherwise there is a simple 2-approximation algorithm in the case $\alpha \leq 1/(4m)$ described in Theorem 3.10. Therefore, the success probability can be boosted to a constant in polynomial number of executions of the bicriteria algorithm in Theorem 3.18. \square

3.6 Cost Minimization with s - t Effective Resistance Constraint

In this subsection, we consider a “dual” problem of the s - t effective resistance minimization problem. In the dual problem, we are given a graph $G = (V, E)$ and a target effective resistance R , and the objective is to find a subgraph H of minimum number of edges such that $\text{Reff}_H(s, t) \leq R$. The same NP-hardness proof in Section 5.1 can be used to show that the dual problem is NP-complete.

Using the same techniques for the s - t effective resistance minimization problem, we can obtain a constant factor bicriteria approximation algorithm for this problem. As the proofs are very similar, we will just state the results and highlight the differences. The main difference is that the convex program has unbounded integrality gap, and as a consequence, we cannot turn the bicriteria approximation algorithm into a true approximation algorithm as in the s - t effective resistance network design problem. Using the same technique as in Theorem 3.10, however, we can return an 8-approximation to the optimal number of edges without violating the effective resistance constraint, if we are allowed to buy up to four copies of the same edge (see Theorem 3.20).

Convex Programming Relaxation. We consider the following natural convex programming relaxation for the dual problem:

$$\begin{aligned} \min_{x \in \mathbb{R}^m} \quad & \sum_{e \in E} x_e \\ \text{subject to} \quad & \text{Reff}_x(s, t) = b_{st}^T L_x^\dagger b_{st} \leq R, \\ & 0 \leq x_e \leq 1 \quad \forall e \in E. \end{aligned} \tag{DCP}$$

Integrality Gap Examples. Unlike the s - t effective resistance network design problem, the **dual convex program (DCP)** has unbounded integrality gap. Consider the following example in Figure 3, where the top path has length $n-1$, and the bottom path has only one edge. The target effective resistance is $R = (n-1)^2 / ((n-1)^2 + \varepsilon)$ for some constant $\varepsilon > 0$. Since $R < 1$, to satisfy the effective resistance constraint, any integral solution must contain both paths and thus has cost n . However, the fractional solution can set $x_e = \varepsilon / (n-1)$ for each edge in the top path and set $x_e = 1$ for the bottom edge. It can be checked that this fractional solution satisfies the constraint, and the total cost is $1 + \varepsilon$. Therefore, the integrality gap of this example is $\Omega(n)$.

Optimal Solutions. Although the convex program (DCP) has a large integrality gap, the same rounding technique can be used to obtain a constant factor bicriteria approximation algorithm. Exactly the same characterization of the optimality conditions as in the s - t effective resistance network design problem holds, such that any optimal solution satisfies the flow-conductance ratio $\alpha > 0$ as described in Lemma 3.2.

Analogous to Lemma 3.3, we can prove an upper bound on α that

$$\alpha^2 \leq \frac{R}{d_{st}}.$$

Analogous to Lemma 3.9, we can prove a lower bound on any optimal solution x that

$$\text{opt} := \sum_{e \in E} x_e \geq \frac{d_{st}^2}{R}.$$

We can assume that $R < d_{st}$, as otherwise a shortest s - t path is an optimal solution, and so we can assume that $0 < \alpha < 1$.

Rounding Algorithm. The rounding algorithm is exactly the same as in Section 3.3. The same proofs as in Lemmas 3.6 and 3.7 will imply that, with probability $\Omega(\alpha)$, the subgraph H returned by the randomized path rounding algorithm satisfies

$$|E(H)| \leq 2 \sum_{e \in E} x_e^* \quad \text{and} \quad \text{Reff}_H(s, t) \leq 4 \text{Reff}_{x^*}(s, t),$$

where x^* is an optimal solution to (DCP) and so $|E(H)| \leq 2\text{opt}$. The same lower bound on $\alpha = \Omega(1/m)$ as described in Theorem 3.10 applies, and so this is a randomized polynomial time algorithm.

An Alternative Bicriteria Approximation Algorithm. In the s - t effective resistance network design problem, we turn a bicriteria approximation algorithm into a true approximation algorithm, by scaling down the budget k by a factor of two and running the bicriteria approximation algorithm. For the proof, we argue $\text{opt}(k/2) \leq 2 \cdot \text{opt}(k)$ by scaling down an optimal solution x^* with budget k to a solution $x^*/2$ with budget $k/2$.

In the dual problem, we can also try a similar approach, by scaling down the target effective resistance R by a factor of 4 and run the bicriteria approximation algorithm. However, we cannot argue that $\text{opt}(R/4) \leq 4 \cdot \text{opt}(R)$, as an optimal solution x^* with effective resistance R may not be able to scale up to $4x^*$ with effective resistance $R/4$ because of the capacity constraints $0 \leq x_e \leq 1$ for $e \in E$. This approach would work if we are allowed to violate the capacity constraint by a factor of 4.

THEOREM 3.20. *Given an weighted input graph $G = (V, E)$, there is a polynomial time algorithm for the dual problem that returns a multi-subgraph H with $|E(H)| \leq 8\text{opt}$ and $\text{Reff}_H(s, t) \leq R$ where there are at most four parallel copies of each edge.*

4 DYNAMIC PROGRAMMING ALGORITHMS FOR SERIES-PARALLEL GRAPHS

In this section, we will present the dynamic programming algorithms for solving the weighted s - t effective resistance network design problem on series-parallel graphs. We first review the definitions of series-parallel graphs in Section 4.1. Then, we present the exact algorithm in Theorem 1.8 when every edge has the same cost in Section 4.2, and the fully polynomial time approximation scheme in Theorem 1.8 in Section 4.3.

4.1 Series-Parallel Graphs

Definition 4.1 (Two-terminal Series-parallel Graph). A two-terminal **series-parallel graph (SP graph)** is a graph with two distinguished vertices (the source s and the target t) that can be constructed recursively as follows:

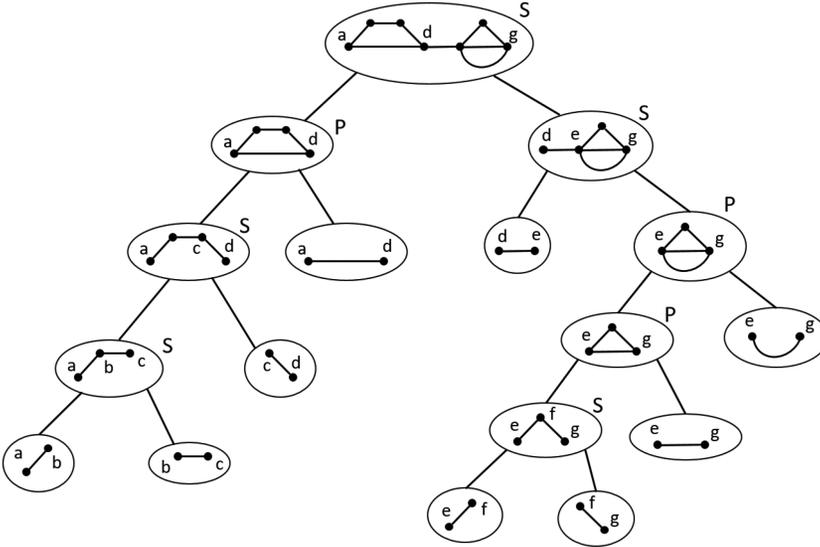


Fig. 4. An example of a SP-tree.

- Base case: A single edge (s, t)
- Compose step: If G_1 and G_2 are two series parallel graphs with source s_i and target t_i ($i = 1, 2$), then we can combine them in two ways:
 - Series-composition: We identify t_1 with s_2 as the same vertex, the source of the new graph is s_1 and the target is t_2 .
 - Parallel-composition: We identify s_1 with s_2 as the same vertex and t_1 with t_2 as the same vertex, the new source is $s_1 = s_2$ and the new target is $t_1 = t_2$.

Given the sequence of steps of constructing a series-parallel graph G , we can define a tree T (SP-tree) as follows.

Definition 4.2 (SP-tree). An SP-tree of a given two-terminal series-parallel graph G is a binary tree that is recursively defined as follows (see Figure 4 for an example):

- Leaf node: If G is a single edge, then T is a single node containing the edge.
- Recurse step: G is either a series-composition (S) or a parallel-composition (P) of G_1 and G_2 , then T is a S-node (P-node) containing G , and its children are roots of the SP-trees of G_1 and G_2 .

For a tree node v in a SP-tree T , let G_v be the subgraph that v represents, s_v, t_v be the two terminals of G_v , and v_l, v_r be its left and right child if v is an internal node. Note that the SP-tree is a fully binary tree with $2m-1$ nodes where m is the number of edges in the graph.

Given a two-terminal SP graph, the corresponding SP-tree can be computed in $O(n+m)$ time. The linear time SP-graph recognition algorithm in Reference [48] will give us the construction sequence of G , and we can build the SP-tree bottom-up.

4.2 Exact Algorithm for Unit-Cost

The following fact shows that the weighted s - t effective resistance can be computed easily from the SP-tree.

FACT 4.3 (RESISTANCE OF SERIES-PARALLEL NETWORK). *Let G be a two-terminal SP graph and each edge e has a non-negative resistance r_e . Let T be the corresponding SP-tree. For every tree node v , we can compute the source-target effective resistance as follows:*

$$\begin{aligned} \text{Leaf node: } \text{Reff}_{G_v}(s_v, t_v) &= r_e \text{ if } v \text{ is a leaf node with a single edge } e, \\ \text{S-node: } \text{Reff}_{G_v}(s_v, t_v) &= \text{Reff}_{G_{v_l}}(s_{v_l}, t_{v_l}) + \text{Reff}_{G_{v_r}}(s_{v_r}, t_{v_r}), \\ \text{P-node: } \text{Reff}_{G_v}(s_v, t_v) &= \frac{\text{Reff}_{G_{v_l}}(s_{v_l}, t_{v_l}) \cdot \text{Reff}_{G_{v_r}}(s_{v_r}, t_{v_r})}{\text{Reff}_{G_{v_l}}(s_{v_l}, t_{v_l}) + \text{Reff}_{G_{v_r}}(s_{v_r}, t_{v_r})}. \end{aligned}$$

We design the dynamic programming algorithm by defining the subproblems using the SP-tree T . For every tree node v and $b = 0, 1 \dots k$, we define the subproblem

$$R(v, b) := \min_{H \subseteq G_v} \left\{ \text{Reff}_H(s_v, t_v) \mid \sum_{e \in H} c_e \leq b \right\}.$$

Since we assume that every edge e has cost $c_e = 1$, there are at most $2mk$ subproblems, as the SP-tree has at most $2m$ nodes and there are at most k possibilities for the cost of a subgraph.

It follows from the definition that $R(v_{\text{root}}, k)$ would be the optimal s - t effective resistance for our problem. To compute $R(v, b)$, with Fact 4.3, we can use the following recurrence, which exhausts all possible distributions of the budget among the two children:

$$R(v, b) = \begin{cases} \infty & \text{if } v \text{ is a leaf node and } b < c_e, \\ r_e & \text{if } v \text{ is a leaf node and } b \geq c_e, \\ \min_{b'=0 \dots b} R(v_l, b') + R(v_r, b - b') & \text{if } v \text{ is a S-node,} \\ \min_{b'=0 \dots b} \frac{R(v_l, b') \cdot R(v_r, b - b')}{R(v_l, b') + R(v_r, b - b')} & \text{if } v \text{ is a P-node.} \end{cases}$$

As there are $O(mk)$ subproblems and each subproblem can be computed in $O(k)$ time, the time complexity of this dynamic programming algorithm is $O(mk^2)$.

4.3 Fully Polynomial Time Approximation Scheme

In this subsection, we use dynamic programming to design a fully polynomial time approximation scheme to prove Theorem 1.8. In the previous subsection, we assume that every edge has the same cost to obtain an exact algorithm, by having a bounded number of subproblems in dynamic programming. When the cost could be arbitrary, the number of subproblems can no longer be bounded by a polynomial. Since the cost constraint must be satisfied, we do not change the cost of the edges, but instead discretize the resistance of the edges and optimize over the cost. We show that this gives an arbitrarily good approximation provided that the discretization is fine enough.

Rescaling: First, by rescaling, we assume that $\min_e r_e = 1$ and $\max_e r_e = U$ in G . Let $m = |E|$ and $L = \varepsilon/m^2$ where $\varepsilon > 0$ is the error in the approximation guarantee. We further rescale the resistance by setting $r_e \leftarrow r_e/L$. This rescaling ensures that for any subgraph of G in which s - t is connected, the s - t effective resistance is upper bounded by Um/L (when all the edges are in series) and is lower bounded by $1/(mL)$ (when all the edges are in parallel).

Subproblems and Recurrence: Let T be the SP-tree of G and let v_{root} be the root of T . We define two similar sets of subproblems. For every tree node v and a value $R \in [1/(mL), Um/L]$, we define the subproblem

$$C(v, R) := \min_{H \subseteq G_v} \left\{ \sum_{e \in H} c_e \mid \text{Reff}_H(s_v, t_v) \leq R \right\}.$$

Similar to the reasoning in the previous subsection, the subproblems satisfy the following recurrence relation:

$$C(v, R) = \begin{cases} c_e & \text{if } v \text{ is a leaf node with a single edge } e \text{ and } R \geq r_e, \\ \infty & \text{if } v \text{ is a leaf node with a single edge } e \text{ and } R < r_e, \\ \min_{R_1, R_2 \in [1/(mL), Um/L]} \{C(v_l, R_1) + C(v_r, R_2) \mid R_1 + R_2 \leq R\} & \text{if } v \text{ is a S-node,} \\ \min_{R_1, R_2 \in [1/(mL), Um/L]} \{C(v_l, R_1) + C(v_r, R_2) \mid \frac{R_1 R_2}{R_1 + R_2} \leq R\} & \text{if } v \text{ is a P-node.} \end{cases}$$

Discretized subproblems: We cannot use dynamic programming to solve the above recurrence relation efficiently as there are unbounded number of subproblems. Instead, we use dynamic programming to compute the solution of all the “discretized” subproblems using the same recurrence relation. For every integer R from $\lceil 1/(mL) \rceil$ to $\lceil Um/L \rceil$, we define

$$\bar{C}(v, R) := \begin{cases} c_e & \text{if } v \text{ is a leaf node with a single edge } e \text{ and } R \geq \lceil r_e \rceil, \\ \infty & \text{if } v \text{ is a leaf node with a single edge } e \text{ and } R < \lceil r_e \rceil, \\ \min_{R_1, R_2 \in \{\lceil 1/(mL) \rceil, \dots, \lceil Um/L \rceil\}} \{\bar{C}(v_l, R_1) + \bar{C}(v_r, R_2) \mid R_1 + R_2 \leq R\} & \text{if } v \text{ is a S-node,} \\ \min_{R_1, R_2 \in \{\lceil 1/(mL) \rceil, \dots, \lceil Um/L \rceil\}} \{\bar{C}(v_l, R_1) + \bar{C}(v_r, R_2) \mid \left\lceil \frac{R_1 R_2}{R_1 + R_2} \right\rceil \leq R\} & \text{if } v \text{ is a P-node.} \end{cases}$$

We can think of $\bar{C}(v, R)$ as the minimum cost required to select a subset of edges such that the effective resistance between s_v and t_v is at most R , when the effective resistance is rounded up to an integer during each step of the computation in the recurrence relation.

Algorithm and Complexity: After computing all $\bar{C}(v, R)$, the algorithm will return

$$\min\{R \mid \bar{C}(v_{\text{root}}, R) \leq k\}$$

as the approximate minimum s - t effective resistance. Given a tree node v , by trying all possible integral values of R_1 and R_2 , we can compute the values of $\bar{C}(v, R)$ for each possible R in $O((Um/L)^2)$ time. Therefore, the total running time of computing all $\bar{C}(v, R)$ is $O(m) \cdot ((Um/L)^2) = O(m^7 U^2 / \varepsilon^2)$. To output the optimal edge set, we can store the optimal values of R_1, R_2 for each pair of (v, R) to reconstruct the edge set.

Correctness and Approximation Guarantee: Since we have not changed the edge cost, the solution returned by the algorithm will have total cost at most k . It remains to show that the s - t effective resistance is at most $(1 + \varepsilon)$ times the optimal s - t effective resistance. For every tree node v and every $b \in [0, k]$, we define

$$R(v, b) := \min\{R \mid C(v, R) \leq b, R \in [1/(mL), Um/L]\}, \\ \bar{R}(v, b) := \min\{R \mid \bar{C}(v, R) \leq b, R \in \{\lceil 1/(mL) \rceil, \dots, \lceil Um/L \rceil\}\}.$$

It follows from the definitions that the optimal s - t effective resistance is $R(v_{\text{root}}, k)$, and the output of our algorithm will be $\bar{R}(v_{\text{root}}, k)$. The following lemma establishes the approximation guarantee.

LEMMA 4.4. *For every tree node v and for every $b \in [0, k]$, it holds that*

$$\bar{R}(v, b) \leq \left(1 + \frac{\varepsilon |E(G_v)|}{m}\right) R(v, b).$$

PROOF. We prove the lemma by induction on the tree node of the SP-tree.

Base Case: Suppose v is a leaf node of T and G_v is a graph of a single edge e .

- For $b < c_e$, we have $\bar{R}(v, b) = R(v, b) = \infty$,
- For $b \geq c_e$, we have $R(v, b) = r_e$, and

$$\bar{R}(v, b) = \lceil r_e \rceil \leq r_e + 1 = r_e + \left(\frac{\varepsilon}{m}\right) \left(\frac{1}{mL}\right) \leq r_e + \frac{\varepsilon}{m} r_e = \left(1 + \frac{\varepsilon |E(G_v)|}{m}\right) R(v, b),$$

where the second inequality uses the fact that every resistance is at least $1/(mL)$, and the last equality uses $|E(G_v)| = 1$ and $r_e = R(v, b)$.

S-node: Suppose v is a S-node. For every $b \in [0, k]$, we have

$$\begin{aligned} \bar{R}(v, b) &= \min_{b_1, b_2 | b_1 + b_2 = b} \{ \bar{R}(v_l, b_1) + \bar{R}(v_r, b_2) \} \\ &\leq \min_{b_1, b_2 | b_1 + b_2 = b} \left\{ \left(1 + \frac{\varepsilon |E(G_{v_l})|}{m}\right) R(v_l, b_1) + \left(1 + \frac{\varepsilon |E(G_{v_r})|}{m}\right) R(v_r, b_2) \right\} \\ &\leq \min_{b_1, b_2 | b_1 + b_2 = b} \left\{ \left(1 + \frac{\varepsilon |E(G_v)|}{m}\right) (R(v_l, b_1) + R(v_r, b_2)) \right\} \\ &= \left(1 + \frac{\varepsilon |E(G_v)|}{m}\right) \min_{b_1, b_2 | b_1 + b_2 = b} \{(R(v_l, b_1) + R(v_r, b_2))\} \\ &= \left(1 + \frac{\varepsilon |E(G_v)|}{m}\right) R(v, b), \end{aligned}$$

where the first inequality follows from the induction hypothesis, and the second inequality follows from the fact that $\max(|E(G_{v_l})|, |E(G_{v_r})|) \leq |E(G_v)| - 1$.

P-node: Suppose v is a P-node. For every $b \in [0, B]$, we have

$$\begin{aligned} \bar{R}(v, b) &= \min_{b_1, b_2 | b_1 + b_2 = b} \left\{ \left[\frac{1}{1/\bar{R}(v_l, b_1) + 1/\bar{R}(v_r, b_2)} \right] \right\} \\ &\leq \min_{b_1, b_2 | b_1 + b_2 = b} \left\{ \left[\left(1 + \frac{\varepsilon (|E(G_v)| - 1)}{m}\right) \frac{1}{1/R(v_l, b_1) + 1/R(v_r, b_2)} \right] \right\} \\ &= \left[\left(1 + \frac{\varepsilon (|E(G_v)| - 1)}{m}\right) R(v, b) \right] \\ &\leq \left(1 + \frac{\varepsilon (|E(G_v)| - 1)}{m}\right) R(v, b) + 1 \\ &= \left(1 + \frac{\varepsilon (|E(G_v)| - 1)}{m}\right) R(v, b) + \frac{\varepsilon}{m} \frac{1}{mL} \\ &\leq \left(1 + \frac{\varepsilon (|E(G_v)| - 1)}{m}\right) R(v, b) + \frac{\varepsilon}{m} R(v, b) \\ &= \left(1 + \frac{\varepsilon |E(G_v)|}{m}\right) R(v, b), \end{aligned}$$

where the first inequality follows from the induction hypothesis and the fact that $\max(|E(G_{v_l})|, |E(G_{v_r})|) \leq |E(G_v)| - 1$, and the last inequality holds as the minimum resistance of any subgraph is at least $1/(mL)$.

Therefore, the lemma follows by induction on the SP-tree. \square

By substituting $v = v_{\text{root}}$ and $b = k$, we have

$$\bar{R}(v_{\text{root}}, k) \leq \left(1 + \frac{m\varepsilon}{m}\right) R(v_{\text{root}}, k) = (1 + \varepsilon) R(v_{\text{root}}, k),$$

which completes the proof of first part of Theorem 1.8.

The following lemma shows that we actually do not lose too much by focusing on the input graphs with polynomially bounded resistance ratio.

LEMMA 4.5. *Given any input graph $G = (V, E, \{r_e\}, \{c_e\})$ of the weighted problem with the designated vertices s and t , within polynomial time, we can transform G into a graph $G' = (V', E', \{r'_e\}, \{c'_e\})$ containing both s and t such that the following three conditions are satisfied.*

- (1) *The graph G' is a subgraph of G , $V' \subseteq V$ and $E' \subseteq E$. Moreover, for any edge $e \in E'$, the resistance does not decrease $r'_e \geq r_e$ and the cost does not change $c'_e = c_e$.*
- (2) *The ratio between maximum resistance and minimum resistance in G' is polynomially bounded, in particular $\max_{e \in E'} r'_e / \min_{e \in E'} r'_e \in O(n^7)$.*
- (3) *The s - t effective resistance of an optimal solution in G' does not increase too much, i.e., $\text{opt}(G') \leq (1 + O(\varepsilon))\text{opt}(G)$, where $\text{opt}(\cdot)$ is the optimal value of an input graph and ε is an arbitrary small constant.*

PROOF. Let $P_G^*(k)$ be the minimum effective resistance of an s - t path in G with cost at most k . We consider the following transformation procedure, given some small constant ε ,

- Delete all edges with resistance larger than $mP_G^*(k)/\varepsilon$ from E , i.e., $E' = \{e \in E : r_e \leq mP_G^*(k)/\varepsilon\}$.
- Round up the minimum resistance to $\varepsilon P_G^*(k)/m^2n$, i.e., for any $e \in E'$, set

$$r'_e = \begin{cases} \frac{\varepsilon}{m^2n} \cdot P_G^*(k), & \text{if } r_e \leq \frac{\varepsilon}{m^2n} \cdot P_G^*(k), \\ r_e, & \text{otherwise.} \end{cases}$$

The first condition is satisfied trivially according to the transformation procedure. The second condition is also satisfied as the ratio $\max_{e \in E'} r'_e / \min_{e \in E'} r'_e \leq m^3n/\varepsilon^2$, which is in $O(n^7)$ for small constant ε .

It remains to prove the last condition. We start with a claim that $P_G^*(k)$ is a polynomially bounded approximation to the optimal s - t effective resistance $\text{opt}(G)$.

CLAIM. $P_G^*(k)/mn \leq \text{opt}(G) \leq P_G^*(k)$.

PROOF. The second inequality is easy as any s - t path with cost at most k is a feasible solution to the weighted problem. We just focus on the first inequality.

Let H be an optimal solution with input G , then any s - t path in H has total resistance at least $P_G^*(k)$. As any simple s - t path has at most n edges, there must be an edge with resistance at least $P_G^*(k)/n$ in any simple s - t path. This implies there must be an s - t cut (S, T) in H such that all edges $e \in \delta(S)$ have resistance at least $P_G^*(k)/n$.

Now, we identify all vertices in S and T , respectively, to form a new graph H' with two vertices and multiple parallel edges between these two vertices. Note that there are at most m such parallel edges and each edge has resistance at least $P_G^*(k)/n$ as it belongs to the cut $\delta(S)$. Thus, the s - t effective resistance in H' is lower bounded by $\text{Reff}_{H'}(s, t) \geq P_G^*(k)/mn$.

According to Rayleigh's monotonicity principle (Fact 2.2), identifying vertices never increases the effective resistance, hence $\text{opt}(G) = \text{Reff}_H(s, t) \geq \text{Reff}_{H'}(s, t) \geq P_G^*(k)/mn$. \square

With the above claim, we are ready to show that deleting edges with large resistances and rounding up edges with small resistances do not change the optimal value too much.

Let H be an optimal solution of G , i.e., $\text{Reff}_H(s, t) = \text{opt}(G)$ and $\text{cost}(H) \leq k$. To prove a lower bound on $\text{opt}(G)$, we first perform the deletion step on H to get $H_1 \subseteq H$. Then, with the rounding

up step on H_1 , we obtain $H_2 \subseteq G'$. We are going to show each of these steps does not increase the s - t effective resistance too much:

$$\text{Reff}_{H_1}(s, t) \leq (1 + 2\varepsilon) \cdot \text{opt}(G) \quad \text{and} \quad \text{Reff}_{H_2}(s, t) \leq \varepsilon \cdot \text{opt}(G) + \text{Reff}_{H_1}(s, t).$$

As H_2 has the same cost as H_1 , which is less than cost of H , H_2 is a feasible solution for G' . Combining together, we establish the third condition of the Lemma, $\text{opt}(G') \leq (1 + 3\varepsilon) \cdot \text{opt}(G)$.

We prove $\text{Reff}_{H_1}(s, t) \leq (1 + 2\varepsilon) \cdot \text{opt}(G)$ first. Let f be the unit s - t electrical flow on H . We are going to modify f into a unit s - t flow supported on H_1 . Consider an edge e in H with resistance $r_e \geq mP_G^*(k)/\varepsilon$, the potential difference between the two endpoints of e is $r_e \cdot f_e$, which is at most the potential difference between s and t , i.e., $\text{Reff}_H(s, t) \leq P_G^*(k)$. Thus, the flow f_e on edge e is at most $P_G^*(k)/r_e \leq \varepsilon/m$. Since there are at most m edges with resistance $r_e \geq mP_G^*(k)/\varepsilon$, with standard argument to remove one fractional s - t flow path at a time, we can obtain an s - t flow $f' \leq f$ with flow value at least $1 - \varepsilon$ and $f'_e = 0$ on any edge e with resistance $r_e \geq mP_G^*(k)/\varepsilon$. Note that f' is supported on H_1 and the energy of f' on H_1 is the same as the energy on H , which is less than the energy of f . After rescaling f' to a unit s - t flow, the $\mathcal{E}_{H_1}(f')$ is at most $\mathcal{E}_H(f)/(1 - \varepsilon)^2 \leq (1 + 2\varepsilon) \cdot \text{opt}(G)$. By Thomson's principle (Fact 2.1), $\text{Reff}_{H_1}(s, t) \leq \mathcal{E}_{H_1}(f') \leq (1 + 2\varepsilon) \cdot \text{opt}(G)$.

Then, we move on to show that $\text{Reff}_{H_2}(s, t) \leq \varepsilon \cdot \text{opt}(G) + \text{Reff}_{H_1}(s, t)$. Let f'' be the unit s - t electrical flow on H_1 . Again, by Thomson's principle (Fact 2.1), the s - t effective resistance of H_2 is bounded by the energy of f'' on H_2 :

$$\begin{aligned} \text{Reff}_{H_2}(s, t) &\leq \mathcal{E}_{H_2}(f'') = \sum_{e \in H_2} r'_e \cdot (f''_e)^2 = \sum_{r_e \leq \varepsilon \cdot P_G^*(k)/m^2 n} \frac{\varepsilon \cdot P_G^*(k)}{m^2 n} \cdot (f''_e)^2 + \sum_{r_e > \varepsilon \cdot P_G^*(k)/m^2 n} r_e \cdot (f''_e)^2 \\ &\leq m \cdot \frac{\varepsilon \cdot P_G^*(k)}{m^2 n} + \sum_{e \in H_1} r_e \cdot (f''_e)^2 \leq \varepsilon \cdot \text{opt}(G) + \text{Reff}_{H_1}(s, t). \end{aligned}$$

The first inequality on the second line holds as $f''_e \leq 1$ in the unit s - t electrical flow. The second inequality on the second line is because of the claim $P_G^*(k)/mn \leq \text{opt}(G)$ and f'' is the unit s - t electrical flow on H_1 .

We have established $\text{opt}(G') \leq (1 + 3\varepsilon) \cdot \text{opt}(G)$, and almost finish the proof of the lemma except that we do not know how to calculate the exact value of $P_G^*(k)$ efficiently. Fortunately, finding an affordable s - t path with effective resistance $P_G^*(k)$ is equivalent to the so called Restricted Shortest Path problem, which admits an FPTAS [27]. By using a $(1 \pm \varepsilon)$ -estimation of $P_G^*(k)$, the same analysis works only with a slightly worse dependence on ε . As finding approximate $P_G^*(k)$, deleting and rounding up edges all run in polynomial time, we can transform G to G' efficiently. \square

Given an input graph G of the weighted problem and a fixed arbitrary small constant ε , with a polynomial time preprocessing as suggested in Lemma 4.5, we end up with a graph G' with polynomially bounded resistance ratio. Then, with the discretized dynamic programming algorithm, we return a feasible $H' \subseteq G'$ in polynomial time, such that $\text{Reff}_{H'}(s, t) \leq (1 + \varepsilon) \cdot \text{opt}(G')$. Let H be a graph the same as H' , but with the original resistance (from G) on each edge. Since we did not increase any edge resistance from H' to H (Lemma 4.5), $\text{Reff}_H(s, t) \leq \text{Reff}_{H'}(s, t)$ by monotonicity principle (Fact 2.2). Again, according to Lemma 4.5, we have

$$\text{Reff}_H(s, t) \leq (1 + \varepsilon) \cdot \text{opt}(G') \leq (1 + O(\varepsilon))\text{opt}(G).$$

As a conclusion, we have an FPTAS for the weighted problem on SP-graphs, completed the whole proof of Theorem 1.8.

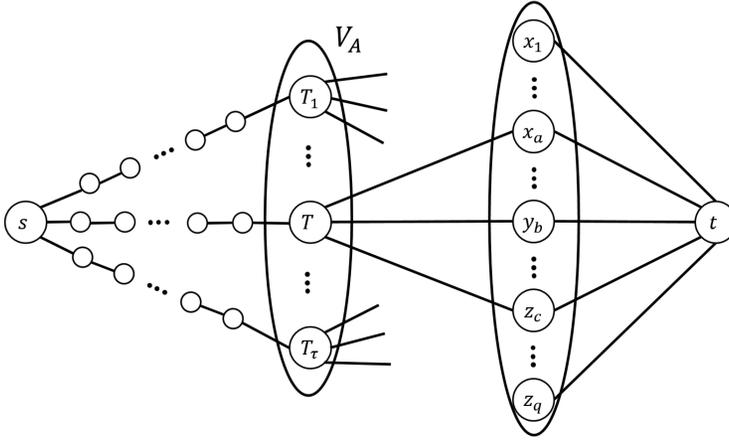


Fig. 5. An illustration of the construction of the graph G from a 3DM instance.

5 HARDNESS RESULTS

In this section, we first prove that the s - t effective resistance network design problem is NP-hard in Section 5.1. Then, we prove that the weighted problem is APX-hard assuming the small-set expansion conjecture in Section 5.2.

5.1 NP-Hardness

We will prove Theorem 1.2 in this subsection. The following is the decision version of the problem.

PROBLEM 5.1 (s - t EFFECTIVE RESISTANCE NETWORK DESIGN).

Input: An undirected graph $G = (V, E)$, two vertices $s, t \in V$, and two parameters k and R .

Question: Does there exist a subgraph H of G with at most k edges and $\text{Reff}_H(s, t) \leq R$?

We will show that this problem is NP-complete by a reduction from the **3-Dimensional Matching (3DM)** problem.

PROBLEM 5.2 (3-DIMENSIONAL MATCHING).

Input: Three disjoint sets of elements $X = \{x_1, \dots, x_q\}$, $Y = \{y_1, \dots, y_q\}$, $Z = \{z_1, \dots, z_q\}$; a set of triples $\mathcal{T} \subseteq X \times Y \times Z$ where each triple contains exactly one element in X, Y, Z .

Question: Does there exist a subset of q pairwise disjoint triples in \mathcal{T} ?

Reduction: Given an instance of 3DM with $\{(X, Y, Z), \mathcal{T}\}$, let $\tau = |\mathcal{T}|$ and denote the triples by $\mathcal{T} = \{T_1, \dots, T_\tau\}$.

We construct a graph $G = (V, E)$ as follows (see Figure 5 for an illustration):

Vertex Set: The vertex set V is the disjoint union of five sets $\{s\}$, $\{t\}$, V_A , V_B , and D . Each vertex in V_A corresponds to a triple in \mathcal{T} , that is $V_A = \{T_1, \dots, T_\tau\}$. Each vertex in V_B corresponds to an element in $X \cup Y \cup Z$, that is $V_B = \{x_1, \dots, x_q, y_1, \dots, y_q, z_1, \dots, z_q\}$. Let $l = 3\tau + 3q$. The set D consists of $\tau \cdot l$ “dummy” vertices $\{d_{i,j} \mid 1 \leq i \leq \tau, 1 \leq j \leq l\}$. So, there are totally $\tau + 3q + 2 + \tau(3\tau + 3q)$ vertices in G , which is polynomial in the input size of the 3DM instance.

Edge Set: The edge set E is the disjoint union of three edge sets F_1, F_2 and P defined as follows. There are 3τ edges in F_1 , where we have three edges (T, x_a) , (T, y_b) and (T, z_c) for each triple

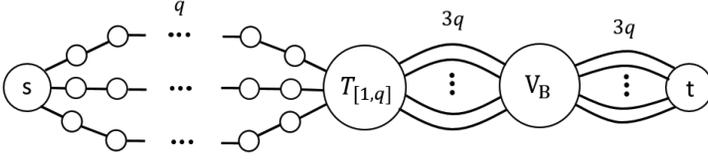


Fig. 6. The subgraph H when the 3DM instance has q disjoint triples.

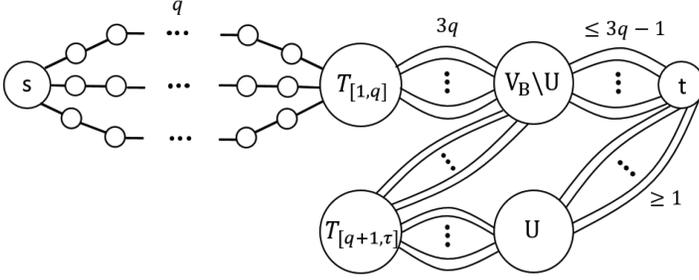


Fig. 7. The subgraph H when U is non-empty.

$T = (x_a, y_b, z_c) \in \mathcal{T}$. There are $3q$ edges in F_2 , where there is an edge from each vertex in V_B to t . There are $\tau(l+1)$ edges in P , where there is a path $P_i := (s, d_{i,1}, d_{i,2}, \dots, d_{i,l}, T_i)$ for each triple $T_i \in \mathcal{T}$, $1 \leq i \leq \tau$. So, there are totally $3\tau + 3q + \tau(3\tau + 3q + 1)$ edges in E , which is polynomial in the input size of the 3DM instance.

The following claim completes the proof of Theorem 1.2.

LEMMA 5.3. *Let $k = q(l+1) + 3\tau + 3q$ and $R = (3(l+1) + 2)/3q$. The 3DM instance has q disjoint triples if and only if the graph G has a subgraph H with at most k edges and $\text{Reff}_H(s, t) \leq R$.*

PROOF. One direction is easy. If there are q disjoint triples in the 3DM instance, say $\{T_1, \dots, T_q\}$, then H will consist of the q paths P_1, \dots, P_q , the $3q$ edges in F_1 incident on T_1, \dots, T_q , and all the $3q$ edges in F_2 . There are $(l+1)q + 3q + 3q \leq k$ edges in H , and $\text{Reff}_H(s, t) = (l+1)/q + 1/3q + 1/3q = (3(l+1) + 2)/3q = R$, as in the graph in Figure 6.

The other direction is more interesting. If there do not exist q disjoint triples in the 3DM instance, then we need to argue that $\text{Reff}_H(s, t) > R$ for any H with at most k edges. First, note that $k < (q+1)(l+1)$, and so the budget is not enough for us to buy more than q paths. As it is useless to buy only a proper subset of a path, we can thus assume that H consists of q paths and all the edges in F_1, F_2 . H has a total of exactly $q(l+1) + 3\tau + 3q = k$ edges. For any such H , we will argue that $\text{Reff}_H(s, t) > R$. Without loss of generality, assume that H consists of P_1, \dots, P_q and all edges in F_1 and F_2 . As T_1, \dots, T_q are not disjoint, there are some vertices in V_B that are not neighbors of $T_1 \cup \dots \cup T_q$. Call those vertices U .

We consider the following modifications of H to obtain H' , and use $\text{Reff}_{H'}(s, t)$ to lower bound $\text{Reff}_H(s, t)$. For every pair of vertices in V_B , we add an edge of zero resistance. For each edge incident on T_{q+1}, \dots, T_τ , we decrease its resistance to zero. By the monotonicity principle, the modifications will not increase the s - t effective resistance, as we either add edges with zero resistance or decrease the resistance of existing edges. The modifications are equivalent to contracting the vertices with zero resistance edges in between, and so H' is equivalent to the graph in Figure 6. Therefore, we have $\text{Reff}_H(s, t) \geq \text{Reff}_{H'}(s, t) \geq R$.

We will prove that one of the inequalities in $\text{Reff}_H(s, t) \geq \text{Reff}_{H'}(s, t) \geq R$ must be strict when $U \neq \emptyset$ (Figure 7). To argue the strict inequality, we look at the unit s - t electrical flow f in H and consider two cases.

- If there exists some vertex $u \in U$ with no incoming electrical flow, then we can delete such a vertex without changing $\text{Reff}_H(s, t)$. But then in the modified graph H' , the number of parallel edges to t is now strictly smaller than $3q$, and therefore $\text{Reff}_{H'}(s, t) > R$.
- If there exists some vertex $u \in U$ with some incoming electrical flow, then $f(T_j u) > 0$ for some $j \geq q + 1$. Since we have decreased the resistance of such an edge $T_j u$ to 0, the energy of f in H' is strictly smaller than the energy of f in H . By Thomson's principle, we have $\text{Reff}_{H'}(s, t) \leq \mathcal{E}_{H'}(f) < \mathcal{E}_H(f) = \text{Reff}_H(s, t)$.

Since the 3DM instance has no q disjoint triples, it follows that $U \neq \emptyset$ and thus one of the above two cases must apply. In either case, we have $\text{Reff}_H(s, t) > R$ and this completes the proof of the other direction. \square

5.2 Improved Hardness Assuming Small-Set Expansion Conjecture

In this subsection, we will prove Theorem 1.6 that it is NP-hard to approximate the weighted s - t effective resistance network design problem within a factor smaller than 2. First, we will state the small-set expansion conjecture and its variant on bipartite graphs, and present an overview of the proof in Section 5.2.1. Next, we will reduce the bipartite small-set expansion problem to the weighted s - t effective resistance network design problem in Section 5.2.2, and then reduce the small-set expansion problem to the bipartite small-set expansion problem in Section 5.2.3 to complete the proof.

5.2.1 The Small-Set Expansion Conjecture and Proof Overview. The gap small-set expansion problem was formulated by Raghavendra and Steurer [43]. We use the version stated in Reference [44].

Definition 5.4 (Gap Small-Set Expansion Problem [43, 44]). Given an undirected graph $G = (V, E)$, two parameters $0 < \beta < \alpha < 1$ and $\delta > 0$, the (α, β) -gap δ -small-set expansion problem, denoted by $\text{SSE}_\delta(\alpha, \beta)$, is to distinguish between the following two cases:

- YES: There exists a subset $S \subseteq V$ with $\text{vol}(S) = \delta \text{vol}(V)$ and $\phi(S) \leq \beta$.
- No: Every subset $S \subseteq V$ with $\text{vol}(S) = \delta \text{vol}(V)$ has $\phi(S) \geq \alpha$.

It is conjectured in Reference [43] that the gap small-set expansion problem becomes harder when δ becomes smaller.

CONJECTURE 5.5 (SMALL-SET EXPANSION CONJECTURE [43, 44]). *For any $\varepsilon \in (0, \frac{1}{2})$, there exists sufficiently small $\delta > 0$ such that $\text{SSE}_\delta(1 - \varepsilon, \varepsilon)$ is NP-hard even for regular graphs.*

It is known that the small-set expansion conjecture implies the **Unique Game Conjecture (UGC)** [43] and is equivalent to some variant of the UGC, i.e., the conjecture that UGC still holds when the input constraint graphs are required to be small-set expanders (see Reference [44] for more details).

We will show the SSE-hardness of the weighted s - t effective resistance network design problem in two steps, and use the small-set expansion problem on regular *bipartite* graphs as an intermediate problem.

PROPOSITION 5.6. *For any $\varepsilon > 0$, there is a polynomial time reduction from $\text{SSE}_\delta(1 - \varepsilon, \varepsilon)$ on d -regular graphs to $\text{SSE}_\delta(1 - 16\varepsilon, \varepsilon)$ on d -regular bipartite graphs.*

PROPOSITION 5.7. *Given an instance of $SSE_\delta(\alpha, \beta)$ on a d -regular bipartite graph B , there is a polynomial time algorithm to construct an instance of the weighted s - t effective resistance network design problem with graph G and cost budget k satisfying the following properties.*

- *If B is a YES-instance, then there is a subgraph H of G with cost at most k and*

$$\text{Reff}_H(s, t) \leq \frac{2}{(1 - \beta)dk}.$$

- *if B is a NO-instance, then every subgraph H of G with cost at most k has*

$$\text{Reff}_H(s, t) \geq \frac{2}{(1 - \frac{\alpha}{2})dk}.$$

The two propositions immediately lead to the following theorem.

THEOREM 5.8. *For any $\epsilon' > 0$, it is NP-hard to approximate the weighted s - t effective resistance network design problem to within a factor of $2 - \epsilon'$, assuming that $SSE_\delta(1 - \epsilon, \epsilon)$ is NP-hard on regular graphs for sufficiently small $\epsilon > 0$.*

PROOF. First, given a d -regular instance of $SSE_\delta(1 - \epsilon, \epsilon)$, we apply Proposition 5.6 to obtain a d -regular bipartite instance of $SSE_\delta(1 - 16\epsilon, \epsilon)$. Then, we apply Proposition 5.7 with $\alpha = 1 - 16\epsilon$ and $\beta = \epsilon$ and see that the ratio between the s - t effective resistance of the NO-case and the YES-case is at least

$$\frac{(1 - \beta)dk}{(1 - \frac{\alpha}{2})dk} = \frac{1 - \epsilon}{\frac{1}{2} + 8\epsilon} = \frac{2(1 - \epsilon)}{1 + 16\epsilon} > 2 - \epsilon',$$

for sufficiently small ϵ . □

The hardness result of the general weighted s - t effective resistance network design problem in Theorem 1.6 follows from Theorem 5.8 directly. For the hardness of the special setting where $c_e = 1$ for all $e \in E$, we provide more details in Remark 5.10 after we presented the reduction to the weighted s - t effective resistance network design problem.

In the remaining of this section, we will prove Proposition 5.7 in Section 5.2.2 and Proposition 5.6 in Section 5.2.3.

5.2.2 From Bipartite Small-set Expansion to Weighted s - t Effective Resistance Network Design.

We prove Proposition 5.7 in this subsection. In the YES-case of bipartite **Small-set Expansion (SSE)**, we use the small dense subgraph (from the small low conductance set) to construct a small subgraph with small s - t effective resistance. In the NO-case of bipartite SSE, we argue that every small subgraph has considerably larger s - t effective resistance.

Construction: Given an $SSE_\delta(\alpha, \beta)$ instance with a d -regular bipartite graph $B = (V_X, V_Y; E_B)$, we construct an instance of the weighted s - t effective resistance network design problem with graph $G = (V, E)$ as follows. See Figure 8 for an illustration.

Vertex Set: The vertex set V of G is simply the disjoint union of $\{s\}, V_X, V_Y, \{t\}$.

Edge Set: The edge set E of G is the disjoint union of three edge sets E_s, E_B, E_t . The edge set E_s has $|V_X|$ edges, where there is an edge from s to each vertex $v \in V_X$. The edge set E_t has $|V_Y|$ edges, where there is an edge from each vertex $v \in V_Y$ to t .

Costs and Resistances: Every edge e in E_B has $c_e = 0$ and $r_e = 1$. Every edge $e \in E_s \cup E_t$ has $c_e = 1$ and $r_e = 0$.

Budget: The cost budget k is $\delta|V_X \cup V_Y|$.

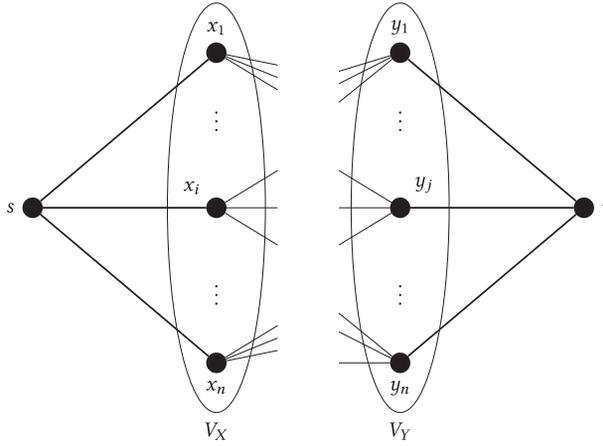


Fig. 8. Reduction from bipartite small set expansion to weighted s - t effective resistance network design.

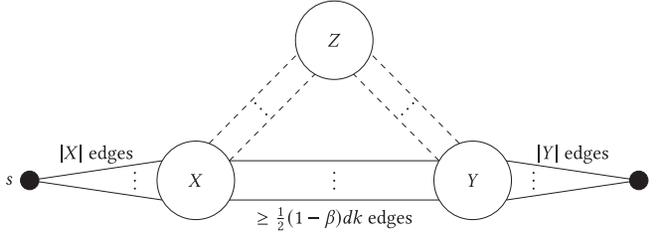


Fig. 9. In the YES-case, the solid edges are included in H and the dashed edges are deleted.

YES-case: Suppose B is a YES-instance of $\text{SSE}_\delta(\alpha, \beta)$. Since B is regular, there exist subsets $X \subseteq V_X$ and $Y \subseteq V_Y$ such that $|X \cup Y| = \delta|V_X \cup V_Y| = k$ and $\phi_B(X \cup Y) \leq \beta$. We construct the subgraph H of G as follows.

Subgraph H : The subgraph H includes all the edges from s to X , all the edges from X to Y , and all the edges from Y to t . Since edges from X to Y are of cost zero, the total cost in H is equal to $|X| + |Y| = k$.

The following claim will complete the proof of the first item of Proposition 5.7.

LEMMA 5.9. $\text{Reff}_H(s, t) \leq 2/((1 - \beta)dk)$.

PROOF. Since B is a d -regular bipartite graph, we have

$$d(|X| + |Y|) = \text{vol}_B(X \cup Y) = |\delta_B(X \cup Y)| + 2|E_B(X, Y)|,$$

where $E_B(X, Y)$ denotes the set of edges with one endpoint in X and one endpoint in Y . Since $\phi_B(X \cup Y) \leq \beta$, we have $|\delta_B(X \cup Y)| \leq \beta \cdot \text{vol}_B(X \cup Y) = d\beta(|X| + |Y|)$. Hence, the number of edges between X and Y is

$$|E_B(X, Y)| = \frac{d(|X| + |Y|) - |\delta_B(X \cup Y)|}{2} \geq \frac{1}{2}(1 - \beta)d(|X| + |Y|) = \frac{1}{2}(1 - \beta)dk.$$

In terms of s - t effective resistance, H is equivalent to the graph in Figure 9, where $Z = (V_X \setminus X) \cup (V_Y \setminus Y)$ is the set of vertices not in X and Y . Notice that the edges from s to X and from Y to t have

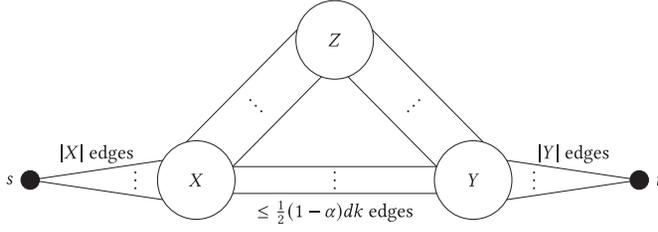


Fig. 10. The subgraph H' is obtained by identifying the subsets X, Y, Z into single vertices.

zero resistance and edges from X and Y have resistance one. By removing the edges and vertices associated to Z , we can upper bound $\text{Reff}_H(s, t) \leq 2/((1 - \beta)dk)$. \square

No-case: We will prove the second item of Proposition 5.7 by arguing that every subgraph of B with total cost at most k has considerably larger s - t effective resistance. Since all the edges between V_X and V_Y have zero cost and adding edges never increases s - t effective resistance (by Rayleigh's monotonicity principle), we can assume without loss of generality that any solution H to the weighted s - t effective resistance network design problem takes all edges between V_X and V_Y and also takes exactly k edges from $E_s \cup E_t$. Consider an arbitrary subgraph H with the above properties. Let $X \subseteq V_X$ be the set of neighbors of s and $Y \subseteq V_Y$ be the set of neighbors of t , with $|X| + |Y| = k$. Let $\phi := \phi_B(X \cup Y)$. Note that $\phi \geq \alpha$ as we are in the No-case where $\phi_B(X \cup Y) \geq \alpha$ for every $|X \cup Y| = k$. Using the same calculation as above, we have

$$|E_B(X, Y)| = \frac{1}{2}(1 - \phi_B(X \cup Y))dk = \frac{1}{2}(1 - \phi)dk.$$

The subgraph H is shown in Figure 10, where $Z = (V_X \setminus X) \cup (V_Y \setminus Y)$ is the set of vertices not in X and Y , and the edges within Z are not shown. To lower bound $\text{Reff}_H(s, t)$, we modify H to obtain H' and argue that $\text{Reff}_H(s, t) \geq \text{Reff}_{H'}(s, t)$ and then show a lower bound on $\text{Reff}_{H'}(s, t)$.

To obtain H' from H , we simply identify the three subsets of vertices X, Y, Z to three vertices, which is equivalent to adding a clique of zero resistance edges to each of these three subsets. By Rayleigh's monotonicity principle, this could only decrease the s - t effective resistance, and so we have $\text{Reff}_H(s, t) \geq \text{Reff}_{H'}(s, t)$.

In terms of s - t effective resistance, the subgraph H' is equivalent to the graph with two paths between X and Y (with parallel edges): one path P_1 of length one with $|E_B(X, Y)|$ parallel edges between X and Y , another path P_2 of length two with $|E_B(X, Z)|$ parallel edges between X and Z and $|E_B(Z, Y)|$ parallel edges between Z and Y . To lower bound $\text{Reff}_{H'}(s, t)$, we lower bound the resistance of P_1 and P_2 , denoted by $r(P_1)$ and $r(P_2)$. Note that

$$r(P_1) = \frac{1}{|E_B(X, Y)|} = \frac{2}{(1 - \phi)dk}.$$

For $r(P_2)$, let $x = |\delta_B(X, Z)|$ and $y = |\delta_B(Y, Z)|$, then

$$r(P_2) = \frac{1}{x} + \frac{1}{y} = \frac{1}{x+y} \cdot \frac{(x+y)^2}{xy} = \frac{1}{x+y} \cdot \left(\frac{x}{y} + \frac{y}{x} + 2 \right) \geq \frac{4}{x+y} = \frac{4}{\phi dk},$$

where the inequality holds, since $a + 1/a \geq 2$ for any $a > 0$, and the last equality holds, because $x + y = |\delta_B(X \cup Y)| = \phi dk$. Finally, by Fact 4.3,

$$\text{Reff}_H(s, t) \geq \text{Reff}_{H'}(s, t) = \frac{1}{1/r(P_1) + 1/r(P_2)} \geq \frac{1}{\frac{1}{2}(1 - \phi)dk + \frac{1}{4}\phi dk} = \frac{2}{(1 - \phi/2)dk} \geq \frac{2}{(1 - \alpha/2)dk},$$

where the last inequality is because we are in the No-case. This completes the proof of the second item of Proposition 5.7.

Remark 5.10. In this subsection, we show the hardness of the weighted s - t effective resistance network design problem, when the edge cost and the edge resistance could be arbitrary. Using a similar argument as in the proof of Theorem 1.2, the reduction can be modified to the unit-cost case if we replace the edges from s to V_X and V_Y to t by sufficiently long paths (so that the cost of connecting s to a vertex in V_X is much larger than the cost of connecting a vertex in V_X to a vertex in V_Y). Therefore, the same $(2 - \epsilon)$ -SSE-hardness also holds in the case when every edge has the same cost.

5.2.3 From Small Set Expansion to Bipartite Small Set Expansion. We prove Proposition 5.6 in this subsection.

Construction: Given an instance $\text{SSE}_\delta(1 - \epsilon, \epsilon)$ on a d -regular graph $G = (V, E)$, we construct a d -regular bipartite graph $B = (V_X, V_Y; E_B)$ as follows. For each vertex v in V , we create a vertex $v_X \in V_X$ and a vertex $v_Y \in V_Y$, so that $|V_X| = |V_Y| = |V|$. For each edge $uv \in E$, we add two edges $u_X v_Y$ and $u_Y v_X$ to E_B . It is clear from the construction that B is d -regular.

Correctness: To prove Proposition 5.6, we will establish the following two claims.

- (1) **YES-case:** If there is a set $S \subseteq V$ with $|S| = \delta|V|$ and $\phi_G(S) \leq \epsilon$ in G , then there exist $X \subseteq V_X$ and $Y \subseteq V_Y$ with $|X| + |Y| = \delta(|V_X| + |V_Y|)$ and $\phi_B(X \cup Y) \leq \epsilon$ in B .
- (2) **NO-case:** If every set $S \subseteq V$ with $|S| = \delta|V|$ has $\phi_G(S) \geq 1 - \epsilon$ in G , then every sets $X \subseteq V_X$ and $Y \subseteq V_Y$ with $|X| + |Y| = \delta(|V_X| + |V_Y|)$ has $\phi_B(X \cup Y) \geq 1 - 16\epsilon$ in B .

YES-case: Let $S \subseteq V$ be a subset with $|S| = \delta|V|$ and $\phi_G(S) \leq \epsilon$ in G . Let $S_X := \{v_X \mid v \in S\}$ and $S_Y := \{v_Y \mid v \in S\}$, with $|S| = |S_X| = |S_Y|$. By construction, an edge $uv \in \delta_G(S)$ if and only if both $u_X v_Y$ and $v_X u_Y$ are in $\delta_B(S_X \cup S_Y)$, and thus $|\delta_B(S_X \cup S_Y)| = 2|\delta_G(S)|$. Since $|S_X \cup S_Y| = |S_X| + |S_Y| = 2|S|$ and B is d -regular, we have

$$\phi_B(S_X \cup S_Y) = \frac{|\delta_B(S_X \cup S_Y)|}{\text{vol}_B(S_X \cup S_Y)} = \frac{|\delta_B(S_X \cup S_Y)|}{d(|S_X| + |S_Y|)} = \frac{2|\delta_G(S)|}{2d|S|} = \phi_G(S) \leq \epsilon.$$

NO-case: Consider arbitrary subsets $X \subseteq V_X$ and $Y \subseteq V_Y$ with $|X| + |Y| = \delta(|V_X| + |V_Y|) = 2\delta|V|$. To lower bound $\phi_B(X \cup Y)$, we will upper bound $|E_B(X, Y)|$. We partition X into groups X_1, \dots, X_a where every group except the last group is of size $\delta|V|/2$ and the last group is of size at most $\delta|V|/2$. We partition Y into groups Y_1, \dots, Y_b in a similar way. The following claim uses the small-set expansion property in G to show that there is no small dense subset in B .

LEMMA 5.11. *Suppose G is a NO-instance of $\text{SSE}_\delta(1 - \epsilon, \epsilon)$. Then, for any $1 \leq i \leq a$ and $1 \leq j \leq b$,*

$$|E_B(X_i, Y_j)| \leq \epsilon\delta d|V|.$$

PROOF. We first argue that there is no small dense subset in G , and then we will use it to bound $|E_B(X_i, Y_j)|$. Suppose $S \subseteq V$ with $|S| = \delta|V|$. As G is a NO-instance, we know that $\phi_G(S) \geq 1 - \epsilon$ and thus $|\delta_G(S)| \geq (1 - \epsilon)\text{vol}_G(S) = (1 - \epsilon)d|S|$. Since $d|S| = \text{vol}_G(S) = |\delta_G(S)| + 2|E_G(S, S)|$, it follows that $|E_G(S, S)| \leq \epsilon d|S|/2 = \epsilon\delta d|V|/2$. Note that this also implies trivially that $|E_G(Z, Z)| \leq \epsilon\delta d|V|/2$ for any Z with $|Z| \leq \delta|V|$.

Given X_i and Y_j , let $Z := \{v \in G \mid v_X \in X_i \text{ or } v_Y \in Y_j\}$. In other words, Z is the set of vertices in G that have at least one copy in $X_i \cup Y_j$ in B . Since each X_i and Y_j is of size at most $\delta|V|/2$, it follows that $|Z| \leq \delta|V|$. Also, note that $|E_B(X_i, Y_j)| \leq 2|E_G(Z, Z)|$, as each edge in $E_B(X_i, Y_j)$ corresponds to one edge in $E_G(Z, Z)$ while each edge in $E_G(Z, Z)$ is corresponded to at most two edges in $E_B(X_i, Y_j)$. Therefore, we can apply the bound in the previous paragraph to conclude that $|E(X_i, Y_j)| \leq 2|E_G(Z, Z)| \leq \epsilon\delta d|V|$. \square

We now use the lemma to bound $|E_B(X, Y)|$. Since $|X| + |Y| = 2\delta|V|$, it follows that $a \leq 4$ and $b \leq 4$, and therefore

$$|E_B(X, Y)| \leq \sum_{i=1}^a \sum_{j=1}^b |E_B(X_i, Y_j)| \leq ab\epsilon\delta d|V| \leq 16\epsilon\delta d|V|.$$

As B is bipartite,

$$|\delta_B(X \cup Y)| = \text{vol}_B(X \cup Y) - 2|E_B(X, Y)| \geq 2\delta d|V| - 32\epsilon\delta d|V| = 2(1 - 16\epsilon)\delta d|V|.$$

Therefore, we have

$$\phi_B(X \cup Y) = \frac{|\delta_B(X \cup Y)|}{\text{vol}_B(X \cup Y)} \geq \frac{2(1 - 16\epsilon)\delta d|V|}{2\delta d|V|} = 1 - 16\epsilon.$$

This completes the proof of Proposition 5.6. We remark that a more careful argument gives $|E_B(X, Y)| \leq 6\epsilon\delta d|V|$ and thus $\phi_B(X \cup Y) \geq 1 - 6\epsilon$, but this constant does not matter for the proof of Theorem 5.8.

6 CONCLUDING REMARKS

We have formulated a new and natural network design problem and presented some hardness and algorithmic results. It opens up a number of interesting problems to be studied.

- (1) For the s - t effective resistance network design problem, we conjecture that the integrality gap of the convex program is exactly two. As mentioned in Remark 3.11, the analysis of the 8-approximation is not tight, and we can show that the same algorithm achieves an approximation ratio strictly smaller than 5. It would be good to close the gap completely.
- (2) The weighted case of arbitrary costs and arbitrary resistances is wide open. It will be interesting if there are stronger convex programming relaxations for the problem (perhaps adding some knapsack constraint as suggested by the dynamic programming algorithms for series-parallel graphs).
- (3) As in survivable network design, one could study the general problem when there are multiple source-sink pairs and each pair has a different effective resistance requirement. It will be very interesting if it is still possible to achieve a constant factor approximation in this general setting.
- (4) An interesting intermediate problem is to find a minimum cost network so that the maximum effective resistance over pairs (the resistance diameter) is minimized. This is an analog of the global connectivity problem in traditional network design.

A more open-ended direction is to unify and extend the techniques for network design problems with spectral requirements. The subsequent work [37] has made some progress towards this direction, where we provide a spectral approach to incorporate not only effective resistance constraints but also other spectral requirements (e.g., algebraic connectivity, graph expansion, etc.) and unstructured non-negative linear constraints into network design. Nevertheless, the result in this article is not subsumed by the one in Reference [37], as we have discussed in Section 1.3. The effective resistance network design problems in the regime of small k are still very interesting problems to investigate.

APPENDIX

A RANDOMIZED ROUNDING WITH APPROXIMATELY OPTIMAL SOLUTION

In this Appendix, we show how to use an approximately optimal solution of the convex program (CP) to design the rounding algorithm. We use Assumption 3.1 ($c_e = w_e = 1$ for all $e \in E$) in this section.

Ideally, we would like to argue that if $x : E \rightarrow \mathbb{R}_{\geq 0}$ is an approximately optimal solution of (CP), then there exist dual variables $\lambda_e^+, \lambda_e^-, \mu \geq 0$ such that, together with x , the KKT conditions are approximately satisfied. Thus, the crucial properties about the flow-conductance ratios in Lemma 3.2 are also approximately satisfied. However, if x_e (or $1-x_e$) is small, the approximate complementary slackness condition does not directly imply that λ_e^- (or λ_e^+) is small, and we cannot conclude that all fractional edges have similar flow-conductance ratios. Due to this technical issue, we need some extra efforts to handle those edges with x_e close to the 0/1 boundary. We modify the randomized path rounding algorithm as follows.

Randomized Path Rounding Algorithm with Approximately Optimal Solution

- (1) Set $\varepsilon \leftarrow 1/\text{poly}(n, m)$ and $\delta \leftarrow 1/\text{poly}(n, m) \gg \varepsilon$.
- (2) Let x be an approximately optimal solution to the convex program (CP) such that $\text{Reff}_x(s, t) \leq \text{opt}(k) + \varepsilon$ and $\sum_{e \in E} x_e \leq k$ and $x_e \in [0, 1]$ for all $e \in E$. Let f be the unit s - t electrical flow supported in x .
- (3) Let $e \in E$ be an arbitrary edge with $\delta < x_e < 1 - \delta$. We use $\alpha_{\text{est}} \leftarrow f_e/x_e$ to estimate the flow-conductance ratio α defined in Lemma 3.2.
- (4) Compute a flow decomposition \mathcal{P} of f as defined in Lemma 3.4.
- (5) For i from 1 to $T := \max\{1, \lfloor 1/\alpha_{\text{est}} \rfloor\}^2$ do
 - Let P_i be a random path from \mathcal{P} where each path $p \in \mathcal{P}$ is sampled with probability v_p .
- (6) Return the subgraph H formed by the edge set $\cup_{i=1}^T P_i$.

The approximately optimal solution x can be obtained in $O(\text{poly}(m, n) \log \frac{1}{\varepsilon})$ time with standard convex optimization techniques, e.g., Ellipsoid method (see, e.g., Theorem 5.2.1 in Reference [4]). The whole algorithm runs in polynomial time if $\alpha_{\text{est}} = \Omega(1/\text{poly}(n, m))$.

To analyze the modified algorithm, we first establish some properties of the approximately optimal solutions in Section A.1 as an analog of Lemma 3.2. Based on these properties, we provide a sketched analysis of the modified randomized path rounding algorithm in Section A.2.

A.1 Characterizing Approximately Optimal Solutions

LEMMA A.1 (CHARACTERIZATION OF APPROXIMATELY OPTIMAL SOLUTION). *Given an accuracy parameter $0 < \varepsilon \ll 1$, let $x : E \rightarrow \mathbb{R}_{\geq 0}$ be an approximately optimal solution of the convex program (CP) such that $\text{Reff}_x(s, t) \leq \text{opt}(k) + \varepsilon$, $\sum_{e \in E} x_e \leq k$, and $x_e \in [0, 1]$ for all $e \in E$, where $\text{opt}(k)$ is the optimal value of (CP).*

Let $1 \gg \delta \gg \varepsilon$ be a parameter that measures the distance from the 0/1 boundaries. Let $E_1 \subseteq E$ be the set of almost-one edges with $x_e \in [1 - \delta, 1]$, let $E_F \subseteq E$ be the set of fractional edges that are staying away from the 0/1 boundaries with $\delta < x_e < 1 - \delta$, and let $E_Z \subseteq E$ be the set of close-to-zero edges with $x_e \in (0, \delta]$. Let $f : E \rightarrow \mathbb{R}_{\geq 0}$ be the unit s - t electrical flow supported in x . There exists

²We take the maximum as α_{est} might slightly larger than 1 due to the estimation error.

$\alpha \geq 0$ such that

$$\frac{f_e^2}{x_e^2} \geq \alpha^2 - 20n\sqrt{\varepsilon} \quad \forall e \in E_I, \quad \left| \frac{f_e^2}{x_e^2} - \alpha^2 \right| \leq 20n\sqrt{\frac{\varepsilon}{\delta}} \quad \forall e \in E_F, \quad \text{and } f_e \leq 2n\delta \quad \forall e \in E_Z.$$

Intuitively, the above characterization says that the edges in E_I and E_F approximately satisfy the properties about the flow-conductance ratios in Lemma 3.2. Meanwhile, those tiny fractional edges in E_Z only support very small amount of electrical flow, hence they will not have significant impact in the analysis of the approximation ratio.

To prove Lemma A.1, we need to refine the proof of the necessity of KKT conditions. We first briefly recall the theory of Lagrangian duality (see, e.g., Reference [6]). Given the primal convex program in (CP), we introduce nonnegative dual variables $\lambda^+, \lambda^- : E \rightarrow \mathbb{R}_{\geq 0}$ and $\mu \geq 0$. By adding some penalty terms according to the constraints, the Lagrange function with respect to (CP) is defined as

$$L(x, \lambda^+, \lambda^-, \mu) := \text{Reff}_x(s, t) + \sum_{e \in E} \lambda_e^+ (x_e - 1) - \sum_{e \in E} \lambda_e^- x_e + \mu \left(\sum_{e \in E} x_e - k \right).$$

The Lagrangian dual problem is defined as

$$\max_{\lambda^+, \lambda^-, \mu \geq 0} \inf_x L(x, \lambda^+, \lambda^-, \mu).$$

Notice that the primal (CP) satisfies Slater's condition, thus strong duality holds and there exist an optimal primal solution x^* and an dual optimal solution $\lambda^+, \lambda^-, \mu \geq 0$ such that

$$L(x^*, \lambda^+, \lambda^-, \mu) = \text{Reff}_{x^*}(s, t) = \inf_x L(x, \lambda^+, \lambda^-, \mu). \quad (\text{A.1})$$

In the following, we fix the optimal solution $x^* : E \rightarrow \mathbb{R}_{\geq 0}$ and the optimal dual solution $\lambda^+, \lambda^-, \mu \geq 0$ that attain the strong duality. We assume that we are given an approximately optimal solution $x : E \rightarrow \mathbb{R}_{\geq 0}$ of (CP) such that

$$\text{Reff}_x(s, t) \leq \text{Reff}_{x^*}(s, t) + \varepsilon \leq n + \varepsilon, \quad \sum_{e \in E} x_e \leq k, \quad \text{and } x_e \in [0, 1] \text{ for all } e \in E, \quad (\text{A.2})$$

where $\text{Reff}_{x^*}(s, t) \leq n$ follows as x^* is an optimal solution of (CP) and any shortest s - t path (of length at most n) is a feasible solution.

The general idea of deriving an approximate version of KKT conditions for x is relatively simple. However, we need to deal with some technical issues such as differentiability of $\text{Reff}_x(s, t)$ and estimations on the second derivatives of $\text{Reff}_x(s, t)$. In the remaining of this subsection, we first state some useful facts in Section A.1.1. Then, we prove approximate Lagrangian optimality condition and approximate complementary slackness condition in Section A.1.2. Finally, we finish the proof of Lemma A.1 in Section A.1.3.

A.1.1 Useful Facts. We start with a lemma about the differentiability of $\text{Reff}_x(s, t)$.

LEMMA A.2 (SEE, E.G., LEMMA 2.4.6 IN REFERENCE [49] FOR A PROOF). *If s and t are connected in the graph formed by the support of x and $x_e > 0$ for some given edge $e \in E$, then $\partial_e \text{Reff}_x(s, t)$ exists and $\partial_e \text{Reff}_x(s, t) = -(b_{st}^T L_x^\dagger b_e)^2$.*

Then, we provide an upper bound on the second derivative $\partial_e^2 \text{Reff}_y(s, t)$ for any $y : E \rightarrow \mathbb{R}_{\geq 0}$.

LEMMA A.3. *For any $y : E \rightarrow \mathbb{R}_{\geq 0}$ such that s and t are connected in the graph formed by the support of y , if $y_e > 0$, then $\partial_e^2 \text{Reff}_y(s, t)$ exists and*

$$\partial_e^2 \text{Reff}_y(s, t) = 2 \left(b_{st}^T L_y^\dagger b_e \right)^2 \left(b_e^T L_y^\dagger b_e \right) \leq 2y_e^{-1} \left(b_{st}^T L_y^\dagger b_{st} \right)^2.$$

PROOF. By Lemma A.2, $\partial_e^2 \text{Reff}_y(s, t) = -\partial_e (b_{st}^T L_y^\dagger b_e)^2$. Let the two endpoints of e be u and v . Since $y_e > 0$, u and v are either contained in the same component as s and t , or both of them are in another component.

If u and v are in a different component than the component that contains s and t , then $b_{st}^T L_y^\dagger b_e = 0$ as L_y^\dagger can be rearranged as a block diagonal matrix with all four entries (s, u) , (s, v) , (t, u) , (t, v) being 0. Thus, $\partial_e^2 \text{Reff}_y(s, t) = -\partial_e (b_{st}^T L_y^\dagger b_e)^2 = 0$, and the lemma holds.

If u and v are within the same component as s and t , then we can restrict to this component and assume the support of y is a connected graph without loss of generality. With similar argument as in the proof of Lemma 3.2, using the fact that $\partial A^{-1} = A^{-1}(\partial A)A^{-1}$, it follows that

$$\begin{aligned} -\partial_e (b_{st}^T L_y^\dagger b_e)^2 &= -2 (b_{st}^T L_y^\dagger b_e) (b_{st}^T (\partial_e L_y^\dagger) b_e) = 2 (b_{st}^T L_y^\dagger b_e) (b_{st}^T L_x^\dagger (b_e b_e^T) L_x^\dagger b_e) \\ &= 2 (b_{st}^T L_y^\dagger b_e)^2 (b_e^T L_y^\dagger b_e). \end{aligned}$$

To upper bound $b_e^T L_y^\dagger b_e$, observe that, for any edge $e \in E$ with $y_e > 0$, $b_e^T L_y^\dagger b_e$ is the effective resistance between the two endpoints of e , which is upper bounded by the resistance of the edge e , and thus

$$b_e^T L_y^\dagger b_e \leq y_e^{-1}. \quad (\text{A.3})$$

Finally, we observe that $|b_{st}^T L_y^\dagger b_e|$ is the potential difference between the two endpoints of e when we send the unit electrical flow from s to t , which is less than the potential difference between s and t , and thus

$$|b_{st}^T L_y^\dagger b_e| \leq \text{Reff}_y(s, t) = b_{st}^T L_y^\dagger b_{st}. \quad (\text{A.4})$$

The lemma follows by combining the above two bounds. \square

A.1.2 Approximate Optimality Conditions. Now, we are ready to derive the approximate Lagrangian optimality condition and approximate complementary slackness condition.

LEMMA A.4 (APPROXIMATE LAGRANGIAN OPTIMALITY CONDITION). *Any primal approximately optimal solution x and the dual optimal solution $\lambda^+, \lambda^-, \mu$ satisfy the inequality*

$$(\partial_e L(x, \lambda^+, \lambda^-, \mu))^2 = \left((b_{st}^T L_x^\dagger b_e)^2 - \lambda_e^+ + \lambda_e^- - \mu \right)^2 \leq 2\epsilon M_e \quad \text{for all } e \in E \text{ with } x_e > 0,$$

where $M_e := 10x_e^{-1}(2n^2 + |\lambda_e^+ - \lambda_e^-|)$.

PROOF. According to the strong duality Equation (A.1), it holds that $\text{Reff}_{x^*}(s, t) = L(x^*, \lambda^+, \lambda^-, \mu)$, i.e., the optimal solution x^* of (CP) is also an optimal solution of $\inf_x L(x, \lambda^+, \lambda^-, \mu)$. Thus,

$$\begin{aligned} L(x^*, \lambda^+, \lambda^-, \mu) &\leq L(x, \lambda^+, \lambda^-, \mu) = \text{Reff}_x(s, t) + \sum_{e \in E} \lambda_e^+(x_e - 1) - \sum_{e \in E} \lambda_e^- x_e + \mu \left(\sum_{e \in E} x_e - k \right) \\ &\leq \text{Reff}_x(s, t) \leq \text{Reff}_{x^*}(s, t) + \epsilon = L(x^*, \lambda^+, \lambda^-, \mu) + \epsilon, \end{aligned} \quad (\text{A.5})$$

where the first inequality in the second line follows by the fact that the penalty terms in the Lagrange function are non-positive for the given primal solution x and dual solution $\lambda^+, \lambda^-, \mu \geq 0$. This implies that the approximately optimal solution x of (CP) is also an approximately optimal solution of $\inf_x L(x, \lambda^+, \lambda^-, \mu)$. Intuitively, we expect the derivative $\partial_e L(x, \lambda^+, \lambda^-, \mu)$ to be small, and we quantify this rigorously in the following.

In the following, we fix an edge $e \in E$ with $x_e > 0$. Let $g(x_e) := L(x, \lambda^+, \lambda^-, \mu)$, where we treat $L(x, \lambda^+, \lambda^-, \mu)$ as a univariate function in x_e with all other variables fixed. Notice that $g'(x_e) = \partial_e L(x, \lambda^+, \lambda^-, \mu)$.

For any given $h \in \mathbb{R}$, if g is twice differentiable on the interval between x_e and $x_e + h$, then by Taylor's theorem, it holds that

$$g(x_e + h) = g(x_e) + g'(x_e)h + \frac{1}{2}g''(a)h^2,$$

for some a in the interval between x_e and $x_e + h$. Taking $h = -\frac{g'(x_e)}{M_e}$, it follows that

$$g(x_e + h) = g(x_e) - \frac{(g'(x_e))^2}{M_e} + \frac{g''(a)(g'(x_e))^2}{2M_e^2} \implies \frac{1}{M_e} \left(1 - \frac{g''(a)}{2M_e}\right) (g'(x_e))^2 \leq \varepsilon, \quad (\text{A.6})$$

where the implication follows as $g(x_e) = L(x, \lambda^+, \lambda^-, \mu)$ is ε -close to the optimal value of $\inf_x L(x, \lambda^+, \lambda^-, \mu)$.

To make sure that g is twice differentiable and the second derivative is bounded on the interval between x_e and $x_e + h$, we first show that the whole interval is bounded away from 0. In particular, we claim that $x_e + h \geq \frac{x_e}{2} > 0$.

Since $\text{Reff}_x(s, t) \leq n + \varepsilon$ is bounded, s and t are connected in the support of x . Thus, by Lemma A.3, for $e \in E$ with $x_e > 0$, $g'(x_e)$ exists and

$$g'(x_e) = -\left(b_{st}^T L_x^\dagger b_e\right)^2 + \lambda_e^+ - \lambda_e^- + \mu.$$

Hence, it follows that

$$|h| = \frac{|g'(x_e)|}{M_e} \leq \frac{(b_{st}^T L_x^\dagger b_e)^2 + \mu + |\lambda_e^+ - \lambda_e^-|}{M_e} \leq \frac{(b_{st}^T L_x^\dagger b_{st})^2 + \mu + |\lambda_e^+ - \lambda_e^-|}{M_e},$$

where the last inequality follows by Equation (A.4). Notice that $b_{st}^T L_x^\dagger b_{st} \leq n + \varepsilon$, and $\mu \leq 1$ by the characterization of α in Lemma 3.3. Together with the definition of M_e , $|h| \leq \frac{x_e}{2}$ for $\varepsilon < 1$. Thus, we have established the claim that $x_e + h \geq \frac{x_e}{2} > 0$.

Since the whole interval between x_e and $x_e + h$ is at least $\frac{x_e}{2} > 0$, g is twice differentiable over the whole interval by Lemma A.3. By the same lemma, for any $a \in \mathbb{R}$ within the interval, it follows that

$$g''(a) = \partial_e^2 \text{Reff}_{x'}(s, t) \leq 2(x'_e)^{-1} \left(b_{st}^T L_{x'}^\dagger b_{st}\right)^2,$$

where x' is the same as x except that $x'_e = a \geq \frac{x_e}{2}$.

It remains to upper bound $b_{st}^T L_{x'}^\dagger b_{st}$. Since L_x and $L_{x'}$ has the same null space, by Sherman-Morrison formula [46],

$$L_{x'}^\dagger = \left(L_x + (a - x_e)b_e b_e^T\right)^\dagger = L_x^\dagger - \frac{(a - x_e)L_x^\dagger b_e b_e^T L_x^\dagger}{1 + (a - x_e)b_e^T L_x^\dagger b_e}.$$

Thus,

$$b_{st}^T L_{x'}^\dagger b_{st} = b_{st}^T L_x^\dagger b_{st} - \frac{(a - x_e)(b_{st}^T L_x^\dagger b_e)^2}{1 + (a - x_e)b_e^T L_x^\dagger b_e} = b_{st}^T L_x^\dagger b_{st} + \frac{(x_e - a)(b_{st}^T L_x^\dagger b_e)^2}{1 - (x_e - a)b_e^T L_x^\dagger b_e}.$$

If $a - x_e > 0$, then $b_{st}^T L_{x'}^\dagger b_{st} \leq b_{st}^T L_x^\dagger b_{st} \leq n + \varepsilon$. Otherwise, we upper bound the additive term by

$$\frac{(x_e - a)(b_{st}^T L_x^\dagger b_e)^2}{1 - (x_e - a)b_e^T L_x^\dagger b_e} \leq \frac{\frac{x_e}{2}(b_{st}^T L_x^\dagger b_e)^2}{1 - \frac{x_e}{2}b_e^T L_x^\dagger b_e} \leq x_e (b_{st}^T L_x^\dagger b_e)^2 \leq (b_{st}^T L_x^\dagger b_{st})^2,$$

where the first inequality follows by $a \geq \frac{x_e}{2}$, the second inequality follows by $b_e^T L_x^\dagger b_e \leq x_e^{-1}$ as in Equation (A.3), and the last inequality follows by $x_e \leq 1$ and Equation (A.4).

Hence, in both cases, it holds that $b_{st}^T L_x^\dagger b_{st} \leq 2b_{st}^T L_x^\dagger b_{st}$. Plugging back to the upper bound of $g''(a)$, for ε less than some small constant, it follows that

$$g''(a) \leq 2(x'_e)^{-1} \left(b_{st}^T L_x^\dagger b_{st} \right)^2 \leq 4x_e^{-1} \left(2b_{st}^T L_x^\dagger b_{st} \right)^2 \leq 4x_e^{-1} (2n + 2\varepsilon)^2 \leq 20x_e^{-1} n^2.$$

Therefore, $\frac{g''(a)}{2M_e} \leq \frac{20x_e^{-1} n^2}{20x_e^{-1} (2n^2 + |\lambda_e^+ - \lambda_e^-|)} \leq \frac{1}{2}$. The lemma follows by Equation (A.6). \square

LEMMA A.5 (APPROXIMATE COMPLEMENTARY SLACKNESS). *Given the primal approximately optimal solution x , and the dual optimal solution $\lambda^+, \lambda^-, \mu$, the following conditions hold:*

$$0 \leq \mu \left(k - \sum_{e \in E} x_e \right) \leq \varepsilon \quad \text{and} \quad 0 \leq \lambda_e^+ (1 - x_e) \leq \varepsilon, \quad 0 \leq \lambda_e^- x_e \leq \varepsilon, \quad \text{for all } e \in E.$$

PROOF. By Equation (A.5) and strong duality Equation (A.1), it follows that

$$\begin{aligned} \text{Reff}_{x^*}(s, t) &\leq L(x, \lambda^+, \lambda^-, \mu) = \text{Reff}_x(s, t) + \sum_{e \in E} \lambda_e^+ (x_e - 1) - \sum_{e \in E} \lambda_e^- x_e + \mu \left(\sum_{e \in E} x_e - k \right) \\ &\leq \text{Reff}_{x^*}(s, t) + \varepsilon + \sum_{e \in E} \lambda_e^+ (x_e - 1) - \sum_{e \in E} \lambda_e^- x_e + \mu \left(\sum_{e \in E} x_e - k \right), \end{aligned}$$

where the last inequality follows by $\text{Reff}_x(s, t) \leq \text{Reff}_{x^*}(s, t) + \varepsilon$. Thus,

$$\sum_{e \in E} \lambda_e^+ (1 - x_e) + \sum_{e \in E} \lambda_e^- x_e + \mu \left(k - \sum_{e \in E} x_e \right) \leq \varepsilon.$$

Notice that $\lambda_e^+ (1 - x_e) \geq 0$, $\lambda_e^- x_e \geq 0$ and $\mu(k - \sum_e x_e) \geq 0$ as x is a feasible solution. Hence, the lemma follows. \square

A.1.3 Proof of Lemma A.1.

PROOF. By Lemma A.4, it holds that

$$-\sqrt{2\varepsilon M_e} \leq (b_{st}^T L_x^\dagger b_e)^2 - \lambda_e^+ + \lambda_e^- - \mu \leq \sqrt{2\varepsilon M_e} \quad \text{for all } e \in E \text{ with } x_e > 0,$$

where $M_e = 10x_e^{-1} (2n^2 + |\lambda_e^+ - \lambda_e^-|)$.

For any given $\delta \in [\varepsilon, \frac{1}{2})$, we handle edges in the three categories E_I, E_F and E_Z separately.

We first consider those edges with $\delta < x_e < 1 - \delta$, i.e., $e \in E_F$. For any $e \in E_F$, we have $0 \leq \lambda_e^+, \lambda_e^- \leq \varepsilon/\delta$ by Lemma A.5. Thus, $|\lambda_e^+ - \lambda_e^-| \leq \varepsilon/\delta \leq 1$ and $M_e \leq 10\delta^{-1} (2n^2 + \varepsilon/\delta)$, which implies

$$|(b_{st}^T L_x^\dagger b_e)^2 - \mu| \leq \frac{\varepsilon}{\delta} + \sqrt{\frac{40\varepsilon n^2}{\delta} + \frac{20\varepsilon^2}{\delta^2}} \leq 2\sqrt{\frac{80\varepsilon n^2}{\delta}} \leq 20n\sqrt{\frac{\varepsilon}{\delta}}.$$

Then, we consider those almost-one edges in E_I with $x_e \in [1 - \delta, 1]$. For any edge $e \in E_I$, by Lemma A.5, we have $0 \leq \lambda_e^- \leq \varepsilon(1 - \delta)^{-1} \leq 2\varepsilon$ (note that we do not have control over the value of $\lambda_e^+ \geq 0$). Since $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for $a, b \geq 0$,

$$\sqrt{2\varepsilon M_e} \leq \sqrt{40\varepsilon(2n^2 + |\lambda_e^+ - \lambda_e^-|)} \leq 9n\sqrt{\varepsilon} + \sqrt{40\varepsilon\lambda_e^+} + \sqrt{40\varepsilon\lambda_e^-} \leq 10n\sqrt{\varepsilon} + \sqrt{40\varepsilon\lambda_e^+},$$

where the last inequality holds as $\lambda_e^- \leq 2\varepsilon$ and $\varepsilon \ll 1$. Thus, it follows that

$$(b_{st}^T L_x^\dagger b_e)^2 \geq \mu + \lambda_e^+ - \lambda_e^- - 10n\sqrt{\varepsilon} - \sqrt{40\varepsilon\lambda_e^+} \geq \mu - 20n\sqrt{\varepsilon},$$

where the last inequality follows as $\lambda_e^- \leq 2\varepsilon$ and $\lambda_e^+ - \sqrt{40\varepsilon\lambda_e^+} \geq -10\varepsilon$ for any $\lambda_e^+ \geq 0$.

Finally, we consider those almost-zero edges in E_Z with $0 < x_e \leq \delta$. We use a simpler argument to bound the value of $\frac{f_e}{x_e}$. Notice that $\frac{f_e}{x_e} = |b_{st}^T L_x^\dagger b_e| \leq \text{Reff}_x(s, t)$ as in Equation (A.4). Since $\text{Reff}_x(s, t) \leq n + \varepsilon$ for the approximately optimal solution x , $f_e \leq (n + \varepsilon)x_e \leq 2n\delta$, as desired.

We finish the proof of the lemma by taking $\alpha = \sqrt{\mu}$. \square

A.2 Analysis of the Randomized Path Rounding Algorithm with Approximately Optimal Solution

We use the properties of the approximately optimal solution in Lemma A.1 to analyze the modified randomized path rounding algorithm. As the analysis essentially follows the same flow as in Sections 3.3 and 3.4, we only provide sketched arguments in this subsection.

We first assume $E_F \neq \emptyset$ (i.e., α_{est} is well-defined) and $\alpha_{\text{est}} = \Omega(1/m)$. In this case, α_{est} is a good estimation of the flow-conductance ratio α defined in Lemma 3.2. In particular, by Lemma A.1 and the definition of α_{est} , it holds that

$$\alpha^2 - 20n\sqrt{\varepsilon/\delta} \leq \alpha_{\text{est}}^2 \leq \alpha^2 + 20n\sqrt{\varepsilon/\delta}.$$

We set the parameters δ, ε to be inverse polynomially small, e.g.,

$$\delta = \frac{1}{nm^3} \quad \text{and} \quad \varepsilon = \frac{1}{n^3m^{11}}.$$

Then, the additive error term $20n\sqrt{\varepsilon/\delta} = O(\frac{1}{m^4})$. Thus, when $\alpha_{\text{est}} = \Omega(\frac{1}{m})$, it holds that

$$\left(1 - O\left(\frac{1}{m}\right)\right) \alpha \leq \alpha_{\text{est}} \leq \left(1 + O\left(\frac{1}{m}\right)\right) \alpha. \quad (\text{A.7})$$

Therefore, by Lemma A.1, it holds that

$$\begin{cases} f_e \geq \left(1 - O\left(\frac{1}{m}\right)\right) \alpha_{\text{est}} x_e, & \forall e \in E_I, \\ f_e = \left(1 \pm O\left(\frac{1}{m}\right)\right) \alpha_{\text{est}} x_e, & \forall e \in E_F, \\ f_e \leq 2n\delta = O\left(\frac{1}{m^3}\right), & \forall e \in E_Z. \end{cases} \quad (\text{A.8})$$

With the above observations, we can upper bound the expected number of edges and the expected s - t effective resistance of the graph returned by the rounding algorithm.

LEMMA A.6. *Let x be an approximately optimal solution to (CP) defined in Equation (A.2). Let X_e be an indicator variable of whether e is included in the returned subgraph H by the rounding algorithm. Then, when $\alpha_{\text{est}} = \Omega(1/m)$, it holds that*

$$\mathbb{E} \left[\sum_{e \in E_F} X_e \right] \leq \left(1 + O\left(\frac{1}{m}\right)\right) \sum_{e \in E_F} x_e \quad \text{and} \quad \mathbb{E} \left[\sum_{e \in E_Z} X_e \right] \leq mT \cdot \max_{e \in E_Z} f_e \leq O\left(\frac{1}{m}\right).$$

Furthermore,

$$\mathbb{E} \left[\sum_{e \in E} X_e \right] \leq \left(1 + O\left(\frac{1}{m}\right)\right) k.$$

PROOF SKETCH. Following the same argument as in Lemma 3.6, it holds that $\Pr(X_e = 1) \leq T f_e$ for all $e \in E$ (see (3.1)). Then, we can use the bounds of f_e in Equation (A.8) to derive the first part of the lemma. Since all edges in E_I have $x_e \geq 1 - \delta \geq 1 - 1/m$, $\mathbb{E}[\sum_{e \in E_I} X_e] \leq (1 + O(1/m)) \sum_{e \in E_I} x_e$. Thus, the second part of the lemma follows. \square

LEMMA A.7. *Let x^* be an optimal solution to (CP), x be an approximately optimal solution to (CP) defined in Equation (A.2), and f be the unit s - t electrical flow supported in x . When $\alpha_{\text{est}} = \Omega(1/m)$, the expected s - t effective resistance of the subgraph H returned by the rounding algorithm is*

$$\mathbb{E} [\text{Reff}_H(s, t)] \leq 2 \left(1 + O\left(\frac{1}{m}\right)\right) \text{Reff}_{x^*}(s, t).$$

PROOF SKETCH. Following the same idea as in the proof of Lemma 3.7, we upper bound the expected s - t effective resistance of the returned subgraph H by constructing a random s - t flow (i.e., sending α_{est} units of flow along each sampled path). The same analysis leads us to Equation (3.2), i.e.,

$$\mathbb{E}[\text{Reff}_H(s, t)] \leq \frac{1}{T} \sum_{e \in E} f_e + \left(1 - \frac{1}{T}\right) \sum_{e \in E} f_e^2.$$

The second term can still be upper bounded by

$$\left(1 - \frac{1}{T}\right) \sum_{e \in E} f_e^2 \leq \left(1 - \frac{1}{T}\right) \sum_{e \in E} \frac{f_e^2}{x_e} \leq \left(1 - \frac{1}{T}\right) \mathcal{E}_x(f).$$

The first term needs some special treatment, where we handle the edges in $E_I \cup E_F$ and E_Z separately:

$$\begin{aligned} \frac{1}{T} \sum_{e \in E} f_e &= \frac{1}{T} \left(\sum_{e \in E_I \cup E_F} f_e + \sum_{e \in E_Z} f_e \right) \leq \frac{(1 + O(\frac{1}{m}))}{T\alpha_{\text{est}}} \sum_{e \in E_I \cup E_F} \frac{f_e^2}{x_e} + O\left(\frac{1}{Tm^2}\right) \\ &\leq \left(1 + O\left(\frac{1}{m}\right)\right) \frac{\mathcal{E}_x(f)}{T\alpha_{\text{est}}} + O\left(\frac{1}{m}\right) \mathcal{E}_x(f), \end{aligned}$$

where the first inequality follows Equation (A.8) that $f_e \geq (1 - O(1/m))\alpha_{\text{est}}x_e \geq \alpha_{\text{est}}x_e/(1 + O(1/m))$ for $e \in E_I \cup E_F$ and $f_e \leq O(1/m^3)$ for $e \in E_Z$, and the second inequality follows by $T \geq 1$ and Equation (3.3) that $\mathcal{E}_x(f) = \text{Reff}_x(s, t) \geq 1/k \geq 1/m$.

By combining the upper bounds of the two terms and noticing that $\frac{1}{2} \leq T\alpha_{\text{est}} \leq 1 + O(1/m)$ (which follows by a similar argument as in Lemma 3.5),

$$\mathbb{E}[\text{Reff}_H(s, t)] \leq \left(1 + \frac{1}{T\alpha_{\text{est}}} - \frac{1}{T} + O\left(\frac{1}{m}\right)\right) \mathcal{E}_x(f).$$

The lemma then follows by a similar argument as in the proof of Lemma 3.7 and the fact that $\mathcal{E}_x(f) \leq \text{Reff}_{x^*}(s, t) + \varepsilon$. \square

Then, following similar arguments as in the proof of Theorems 3.8 and 3.10, we have a randomized polynomial time constant approximation algorithm when $E_F \neq \emptyset$ and $\alpha_{\text{est}} = \Omega(1/m)$.

It remains to deal with the cases where $E_F = \emptyset$ or $\alpha_{\text{est}} \leq 1/4m$. In both cases, we can simply return all edges in E_I . We first note that the number of edges in E_I is at most k as $x_e \geq 1 - \delta > 1 - 1/m$ and $\sum_{e \in E_I} x_e \leq k$. Then, by Lemma A.1, we can upper bound $f_e \leq 1/2m$ for all edges $e \in E \setminus E_I$ when $E_F = \emptyset$ or $\alpha_{\text{est}} \leq 1/4m$. With a similar argument that deals with the case of $\alpha \leq 1/4m$ in the proof of Theorem 3.10, returning all edges in E_I gives us a constant approximate solution.

ACKNOWLEDGMENTS

The authors thank the anonymous referees for comments that improve the presentation of this article.

REFERENCES

- [1] Ajit Agrawal, Philip Klein, and R. Ravi. 1995. When trees collide: An approximation algorithm for the generalized Steiner problem on networks. *SIAM J. Comput.* 24, 3 (1995), 440–456. <https://doi.org/10.1137/S0097539792236237>
- [2] Zeyuan Allen-Zhu, Yuanzhi Li, Aarti Singh, and Yining Wang. 2021. Near-optimal discrete optimization for experimental design: A regret minimization approach. *Math. Program.* 186, 1–2, Ser. A (2021), 439–478. <https://doi.org/10.1007/s10107-019-01464-2>
- [3] Nima Anari and Shayan Oveis Gharan. 2015. Effective-resistance-reducing flows, spectrally thin trees, and asymmetric TSP. In *Proceedings of the 56th Annual Symposium on Foundations of Computer Science (FOCS'15)*. IEEE Computer Society, Washington, DC, 20–39. <https://doi.org/10.1109/FOCS.2015.11>

- [4] Aharon Ben-Tal and Arkadi Nemirovski. 2001. *Lectures on Modern Convex Optimization*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA; Mathematical Programming Society (MPS), Philadelphia, PA. xvi+488 pages. <https://doi.org/10.1137/1.9780898718829>
- [5] Stephen Boyd, Persi Diaconis, and Lin Xiao. 2004. Fastest mixing Markov chain on a graph. *SIAM Rev.* 46, 4 (2004), 667–689. <https://doi.org/10.1137/S0036144503423264>
- [6] Stephen Boyd and Lieven Vandenbergh. 2004. *Convex Optimization*. Cambridge University Press, Cambridge. xiv+716 pages. <https://doi.org/10.1017/CBO9780511804441>
- [7] Tanmoy Chakraborty, Julia Chuzhoy, and Sanjeev Khanna. 2008. Network design for vertex connectivity. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC'08)*. ACM, New York, NY, 167–176. <https://doi.org/10.1145/1374376.1374403>
- [8] Ashok K. Chandra, Prabhakar Raghavan, Walter L. Ruzzo, Roman Smolensky, and Prasoos Tiwari. 1996/97. The electrical resistance of a graph captures its commute and cover times. *Comput. Complex.* 6, 4 (1996/97), 312–340. <https://doi.org/10.1007/BF01270385>
- [9] Joseph Cheriyan and László A. Végh. 2014. Approximating minimum-cost k -node connected subgraphs via independence-free graphs. *SIAM J. Comput.* 43, 4 (2014), 1342–1362. <https://doi.org/10.1137/120902847>
- [10] Joseph Cheriyan, Santosh Vempala, and Adrian Vetta. 2006. Network design via iterative rounding of setpair relaxations. *Combinatorica* 26, 3 (2006), 255–275. <https://doi.org/10.1007/s00493-006-0016-z>
- [11] Paul Christiano, Jonathan A. Kelner, Aleksander Mądry, Daniel A. Spielman, and Shang-Hua Teng. 2011. Electrical flows, laplacian systems, and faster approximation of maximum flow in undirected graphs. In *Proceedings of the 43rd Annual ACM Symposium on Theory of Computing (STOC'11)*. ACM, New York, NY, 273–282. <https://doi.org/10.1145/1993636.1993674>
- [12] Fan R. K. Chung. 1997. *Spectral Graph Theory*. CBMS Regional Conference Series in Mathematics, Vol. 92. Published for the Conference Board of the Mathematical Sciences, Washington, DC; by the American Mathematical Society, Providence, RI. xii+207 pages.
- [13] Julia Chuzhoy and Sanjeev Khanna. 2009. An $O(k^3 \log n)$ -approximation algorithm for vertex-connectivity survivable network design. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS'09)*. IEEE Computer Society, Washington, DC, 437–441. Retrieved from <http://dl.acm.org/citation.cfm?id=1747597.1748059>.
- [14] Michael Dinitz and Zeyu Zhang. 2016. Approximating low-stretch spanners. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'16)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 821–840. Retrieved from <http://dl.acm.org/citation.cfm?id=2884435.2884494>.
- [15] Yevgeniy Dodis and Sanjeev Khanna. 1999. Design networks with bounded pairwise distance. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing (STOC'99)*. ACM, New York, NY, 750–759. <https://doi.org/10.1145/301250.301447>
- [16] Jeremy Elson, Richard M. Karp, Christos H. Papadimitriou, and Scott Shenker. 2004. Global synchronization in sensor networks. In *Proceedings of the Latin American Symposium on Theoretical Informatics (LATIN'04)*, Martin Farach-Colton (Ed.). Springer, Berlin, 609–624.
- [17] Alina Ene and Ali Vakilian. 2014. Improved approximation algorithms for degree-bounded network design problems with node connectivity requirements. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC'14)*. ACM, New York, NY, 754–763. <https://doi.org/10.1145/2591796.2591837>
- [18] Jittat Fakcharoenphol and Bundit Laekhanukit. 2008. An $O(\log^2 k)$ -approximation algorithm for the k -vertex connected spanning subgraph problem. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC'08)*. ACM, New York, NY, 153–158. <https://doi.org/10.1145/1374376.1374401>
- [19] Lisa Fleischer, Kamal Jain, and David P. Williamson. 2001. An iterative rounding 2-approximation algorithm for the element connectivity problem. In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science (FOCS'01)*. IEEE, IEEE Computer Society, Las Vegas, NV, 339–347.
- [20] Takuro Fukunaga, Zeev Nutov, and R. Ravi. 2015. Iterative rounding approximation algorithms for degree-bounded node-connectivity network design. *SIAM J. Comput.* 44, 5 (2015), 1202–1229. <https://doi.org/10.1137/13094503X>
- [21] Harold N. Gabow. 2005. On the L_∞ -norm of extreme points for crossing supermodular directed network LPs. In *Proceedings of the 11th International Conference on Integer Programming and Combinatorial Optimization (IPCO'05)*. Springer-Verlag, Berlin, 392–406. https://doi.org/10.1007/11496915_29
- [22] Harold N. Gabow, Michel X. Goemans, Éva Tardos, and David P. Williamson. 2009. Approximating the smallest k -edge connected spanning subgraph by LP-rounding. *Networks* 53, 4 (2009), 345–357. <https://doi.org/10.1002/net.20289>
- [23] Arpita Ghosh and Stephen Boyd. 2006. Growing well-connected graphs. In *Proceedings of the 45th IEEE Conference on Decision and Control (CDC'06)*. IEEE, 6605–6611.
- [24] Arpita Ghosh, Stephen Boyd, and Amin Saberi. 2008. Minimizing effective resistance of a graph. *SIAM Rev.* 50, 1 (2008), 37–66. <https://doi.org/10.1137/050645452>

- [25] M. X. Goemans, A. V. Goldberg, S. Plotkin, D. B. Shmoys, É. Tardos, and D. P. Williamson. 1994. Improved approximation algorithms for network design problems. In *Proceedings of the 5th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'94)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 223–232. <http://dl.acm.org/citation.cfm?id=314464.314497>.
- [26] Michel X. Goemans and David P. Williamson. 1995. A general approximation technique for constrained forest problems. *SIAM J. Comput.* 24, 2 (1995), 296–317. <https://doi.org/10.1137/S0097539793242618>
- [27] Refael Hassin. 1992. Approximation schemes for the restricted shortest path problem. *Math. Oper. Res.* 17, 1 (1992), 36–42. <https://doi.org/10.1287/moor.17.1.36>
- [28] Rabih A. Jabr, Ravindra Singh, and Bikash C. Pal. 2012. Minimum loss network reconfiguration using mixed-integer convex programming. *IEEE Trans. Power Syst.* 27, 2 (2012), 1106–1115.
- [29] Kamal Jain. 2001. A factor 2 approximation algorithm for the generalized Steiner network problem. *Combinatorica* 21, 1 (2001), 39–60. <https://doi.org/10.1007/s004930170004>
- [30] William Thomson Baron Kelvin and Peter Guthrie Tait. 1867. *Treatise on Natural Philosophy*. Vol. 1. Clarendon Press.
- [31] Gustav Kirchhoff. 1847. Ueber die Auflösung der Gleichungen, auf welche man bei der Untersuchung der linearen Vertheilung galvanischer Ströme geführt wird. *Annalen der Physik* 148, 12 (1847), 497–508.
- [32] Alexandra Kolla, Yury Makarychev, Amin Saberi, and Shang-Hua Teng. 2010. Subgraph sparsification and nearly optimal ultrasparsifiers. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC'10)*. ACM, New York, NY, 57–66. <https://doi.org/10.1145/1806689.1806699>
- [33] Guy Kortsarz, Robert Krauthgamer, and James R. Lee. 2004. Hardness of approximation for vertex-connectivity network design problems. *SIAM J. Comput.* 33, 3 (2004), 704–720. <https://doi.org/10.1137/S0097539702416736>
- [34] Bundit Laekhanukit. 2014. Parameters of two-prover-one-round game and the hardness of connectivity problems. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'14)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1626–1643. Retrieved from <http://dl.acm.org/citation.cfm?id=2634074.2634192>.
- [35] Lap Chi Lau, Joseph Naor, Mohammad R. Salavatipour, and Mohit Singh. 2009. Survivable network design with degree or order constraints. *SIAM J. Comput.* 39, 3 (2009), 1062–1087. <https://doi.org/10.1137/070700620>
- [36] Lap Chi Lau and Hong Zhou. 2015. A unified algorithm for degree bounded survivable network design. *Math. Program.* 154, 1–2, Ser. B (2015), 515–532. <https://doi.org/10.1007/s10107-015-0858-5>
- [37] Lap Chi Lau and Hong Zhou. 2020. A spectral approach to network design. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC'20)*. Association for Computing Machinery, New York, NY, 826–839. <https://doi.org/10.1145/3357713.3384313>
- [38] Peter Matthews. 1988. Covering problems for Brownian motion on spheres. *Ann. Probab.* 16, 1 (1988), 189–199. Retrieved from [http://links.jstor.org/sici?sici=0091-1798\(198801\)16:1<189:CPFBMO>2.0.CO;2-6&origin=MSN](http://links.jstor.org/sici?sici=0091-1798(198801)16:1<189:CPFBMO>2.0.CO;2-6&origin=MSN).
- [39] Aleksander Mądry, Damian Straszak, and Jakub Tarnawski. 2015. Fast generation of random spanning trees and the effective resistance metric. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'15)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2019–2036.
- [40] Aleksandar Nikolov, Mohit Singh, and Uthaiapon Tao Tantipongpipat. 2019. Proportional volume sampling and approximation algorithms for A-optimal design. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'19)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1369–1386. Retrieved from <http://dl.acm.org/citation.cfm?id=3310435.3310519>.
- [41] Kaare Brandt Petersen and Michael Syskind Pedersen. 2012. *The Matrix Cookbook*. Retrieved from <http://cw.matrixcookbook.com/>.
- [42] Prabhakar Raghavan. 1988. Probabilistic construction of deterministic algorithms: Approximating packing integer programs. *J. Comput. System Sci.* 37, 2 (1988), 130–143. [https://doi.org/10.1016/0022-0000\(88\)90003-7](https://doi.org/10.1016/0022-0000(88)90003-7)
- [43] Prasad Raghavendra and David Steurer. 2010. Graph expansion and the unique games conjecture. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC'10)*. ACM, New York, NY, 755–764. <https://doi.org/10.1145/1806689.1806792>
- [44] Prasad Raghavendra, David Steurer, and Madhur Tulsiani. 2012. Reductions between expansion problems. In *Proceedings of the IEEE Conference on Computational Complexity (CCC'12)*. IEEE Computer Society, Washington, DC, 64–73. <https://doi.org/10.1109/CCC.2012.43>
- [45] Aaron Schild. 2018. An almost-linear time algorithm for uniform random spanning tree generation. In *Proceedings of the 50th Annual ACM Symposium on Theory of Computing (STOC'18)*. ACM, New York, NY, 214–227. <https://doi.org/10.1145/3188745.3188852>
- [46] Jack Sherman and Winifred J. Morrison. 1950. Adjustment of an inverse matrix corresponding to a change in one element of a given matrix. *Ann. Math. Stat.* 21 (1950), 124–127. <https://doi.org/10.1214/aoms/1177729893>

- [47] Daniel A. Spielman and Nikhil Srivastava. 2011. Graph sparsification by effective resistances. *SIAM J. Comput.* 40, 6 (2011), 1913–1926. <https://doi.org/10.1137/080734029>
- [48] Jacobo Valdes, Robert E. Tarjan, and Eugene L. Lawler. 1982. The recognition of series parallel digraphs. *SIAM J. Comput.* 11, 2 (1982), 298–313. <https://doi.org/10.1137/0211023>
- [49] Hong Zhou. 2020. *A Spectral Approach to Network Design and Experimental Design*. Ph.D. Dissertation. University of Waterloo, Waterloo.

Received December 2020; revised September 2021; accepted February 2022