



Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Cui, G;Wang, H;Wei, Z

Title:

Mixing Time Matters: Accelerating Effective Resistance Estimation via Bidirectional Method

Date:

2025-07-20

Citation:

Cui, G., Wang, H. & Wei, Z. (2025). Mixing Time Matters: Accelerating Effective Resistance Estimation via Bidirectional Method. Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 1, pp.177-188. Association for Computing Machinery. <https://doi.org/10.1145/3690624.3709298>.

Persistent Link:

<https://hdl.handle.net/11343/365423>

# Mixing Time Matters: Accelerating Effective Resistance Estimation via Bidirectional Method

Guanyu Cui  
Renmin University of China  
Beijing, China  
cuiguanyu@ruc.edu.cn

Hanzhi Wang  
BARC, University of Copenhagen  
Copenhagen, Denmark  
hzwang.helen@gmail.com

Zhewei Wei\*  
Renmin University of China  
Beijing, China  
zhewei@ruc.edu.cn

## Abstract

We study the problem of efficiently approximating the *effective resistance* (ER) on undirected graphs, where ER is a widely used node proximity measure with applications in graph spectral sparsification, multi-class graph clustering, network robustness analysis, graph machine learning, and more. Specifically, given any nodes  $s$  and  $t$  in an undirected graph  $G$ , we aim to efficiently estimate the ER value  $R(s, t)$  between nodes  $s$  and  $t$ , ensuring a small absolute error  $\epsilon$ . The previous best algorithm for this problem has a worst-case computational complexity of  $\tilde{O}\left(\frac{L_{\max}^3}{\epsilon^2 d^2}\right)$ , where the value of  $L_{\max}$  depends on the mixing time of random walks on  $G$ ,  $d = \min\{d(s), d(t)\}$ , and  $d(s), d(t)$  denote the degrees of nodes  $s$  and  $t$ , respectively. We improve this complexity to  $\tilde{O}\left(\min\left\{\frac{L_{\max}^{7/3}}{\epsilon^{2/3}}, \frac{L_{\max}^3}{\epsilon^2 d^2}, mL_{\max}\right\}\right)$ , achieving a theoretical improvement of  $\tilde{O}\left(\max\left\{\frac{L_{\max}^{2/3}}{\epsilon^{4/3} d^2}, 1, \frac{L_{\max}^2}{\epsilon^2 d^2 m}\right\}\right)$  over previous results. Here,  $m$  denotes the number of edges. Given that  $L_{\max}$  is often very large in real-world networks (e.g.,  $L_{\max} > 10^4$ ), our improvement on  $L_{\max}$  is significant, especially for real-world networks. We also conduct extensive experiments on real-world and synthetic graph datasets to empirically demonstrate the superiority of our method. The experimental results show that our method achieves a  $10\times$  to  $1000\times$  speedup in running time while maintaining the same absolute error compared to baseline methods.

## CCS Concepts

• **Theory of computation** → **Graph algorithms analysis; Approximation algorithms analysis; Streaming, sublinear and near linear time algorithms.**

## Keywords

effective resistance; bidirectional algorithm; computational complexity; mixing time

\*Zhewei Wei is the corresponding author. The work was partially done at Gaoling School of Artificial Intelligence, Beijing Key Laboratory of Big Data Management and Analysis Methods, MOE Key Lab of Data Engineering and Knowledge Engineering, and Pazhou Laboratory (Huangpu), Guangzhou, Guangdong 510555, China.

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 the author(s) 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](mailto:permissions@acm.org).  
KDD '25, Toronto, ON, Canada

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 979-8-4007-1245-6/25/08  
<https://doi.org/10.1145/3690624.3709298>

## ACM Reference Format:

Guanyu Cui, Hanzhi Wang, and Zhewei Wei. 2025. Mixing Time Matters: Accelerating Effective Resistance Estimation via Bidirectional Method. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.1 (KDD '25)*, August 3–7, 2025, Toronto, ON, Canada. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3690624.3709298>

## 1 Introduction

Effective Resistance (ER) is a widely adopted measure for node proximity in graphs, with applications across various scenarios. For example, ER approximation is closely related to optimal transport problems [27], the detection of low conductance sets in graph clustering [1], the maximum flow problem [8], and graph spectral sparsification [28]. Additionally, the ER metric is utilized to enhance seeding strategies in influence maximization [12] and to assess network robustness against disruptions [31]. Moreover, with the growing prominence of graph neural networks (GNNs), ER is increasingly employed in the graph rewiring process of GNNs to address the oversquashing problem [5] and to improve the expressive power of GNN models [34].

Recognizing these widespread applications, the problem of efficiently estimating ER scores between a given pair of nodes has been the subject of extensive research [18, 24, 33]. Specifically, considering an undirected graph  $G$  with two arbitrary nodes  $s$  and  $t$  in  $G$ , the ER score between  $s$  and  $t$ , denoted by  $R(s, t)$ , is defined as

$$R(s, t) = (\mathbf{e}_s - \mathbf{e}_t)^\top \mathbf{L}^+ (\mathbf{e}_s - \mathbf{e}_t) = \mathbf{L}_{ss}^+ - \mathbf{L}_{st}^+ - \mathbf{L}_{ts}^+ + \mathbf{L}_{tt}^+, \quad (1)$$

where  $\mathbf{L}^+$  denotes the Moore-Penrose pseudo-inverse of the graph  $G$ 's Laplacian matrix  $\mathbf{L}$ , and  $\mathbf{e}_s$  denotes a one-hot vector with the  $s$ -th element being 1.

To understand this definition formula more intuitively, Peng et al. [24] provide a probabilistic interpretation of ER scores, showing that Equation (1) can be rewritten as

$$R(s, t) = \sum_{\ell=0}^{\infty} \left( \frac{p^{(\ell)}(s, s)}{d(s)} - \frac{p^{(\ell)}(s, t)}{d(t)} - \frac{p^{(\ell)}(t, s)}{d(s)} + \frac{p^{(\ell)}(t, t)}{d(t)} \right), \quad (2)$$

where  $p^{(\ell)}(s, t)$  denotes the transition probability of a random walk on  $G$  moving from node  $s$  to  $t$  in its  $\ell$ -th step, and  $d(s)$  corresponds to the degree of node  $s$  in  $G$ . In other words, the ER score between nodes  $s$  and  $t$  equals the sum of bidirectional degree-normalized random walk probabilities. This interpretation has inspired a line of research [24, 33] focused on estimating ER scores by leveraging techniques for computing random walk transition probabilities on graphs. This approach involves setting a maximum length  $L_{\max}$  for random walks, estimating the  $\ell$ -hop transition probabilities between nodes  $s$  and  $t$  for each  $\ell \in [0, L_{\max}]$  (i.e., estimating  $p^{(\ell)}(s, t)$ ,  $p^{(\ell)}(t, s)$ ,  $p^{(\ell)}(s, s)$ , and  $p^{(\ell)}(t, t)$ ), then calculating the

**Table 1:**  $L_{\max}$  across various real-world datasets, assuming  $\epsilon = 1e-3$  and  $d(s) = d(t) = \lceil \bar{d} \rceil$ .

| Facebook | DBLP  | Youtube | Orkut | LiveJournal | Friendster |
|----------|-------|---------|-------|-------------|------------|
| 14,546   | 4,536 | 6,353   | 1,767 | 146,133     | 23,734     |

**Table 2:**  $L_{\max}$ -truncated ER values (with 4 significant figures) versus  $L_{\max}$  on the Facebook dataset.

| $L_{\max}$           | 100    | 1,000  | 2,000  | 3,000  | 4,000  | 5,000  |
|----------------------|--------|--------|--------|--------|--------|--------|
| $R_{L_{\max}}(s, t)$ | 0.1517 | 0.1791 | 0.1809 | 0.1814 | 0.1815 | 0.1816 |

$L_{\max}$ -truncated ER, denoted as  $R_{L_{\max}}(s, t)$ , as the estimate for  $R(s, t)$ . The definition of  $R_{L_{\max}}(s, t)$  is provided below:

$$R_{L_{\max}}(s, t) = \sum_{\ell=0}^{L_{\max}} \left( \frac{p^{(\ell)}(s, s)}{d(s)} - \frac{p^{(\ell)}(s, t)}{d(t)} - \frac{p^{(\ell)}(t, s)}{d(s)} + \frac{p^{(\ell)}(t, t)}{d(t)} \right). \quad (3)$$

Notably, due to the mixing time of random walks on graphs, the  $\ell$ -hop transition probability converges after several steps. Consequently, as shown in Equation (3), this line of research [24, 33] truncates the infinite summation in Equation (2) at a level of  $L_{\max}$  and focuses on approximating the  $\ell$ -hop transition probabilities for  $\ell \leq L_{\max}$ . As a result, to estimate the ER score for a given pair of nodes  $s$  and  $t$  while assuring an additive error  $\epsilon$  and a constant failure probability, the best-known methods [33], AMC and GEER, both achieve a worst-case computational complexity of  $\tilde{O}\left(\frac{L_{\max}^3}{\epsilon^2 d^2}\right)$ , where  $d = \min\{d(s), d(t)\}$ .

**Our Motivations.** As proved in [24], the value of  $L_{\max}$  depends on the additive error  $\epsilon$  specified in the approximation problem and the spectral radius of the graph. Although previous studies [24, 33] claim that the order of  $L_{\max}$  for real-world graphs is  $O(\log n)$ , the actual value of  $L_{\max}$  can still be very large. As reported in Table 1, the value of  $L_{\max}$  always exceeds  $10^3$  or even  $10^4$ . This large value of  $L_{\max}$  makes the  $O(L_{\max}^3)$  time-dependence in the previously best complexity result [33] infeasible on real-world large graphs. To address this challenge, existing studies often set it to a much smaller value than its true value in experiments. For example, on the Facebook dataset with an additive error  $\epsilon = 1e-3$ , the actual value of  $L_{\max}$  is 14546 as given in Table 1, while previous works set  $L_{\max}$  to 74 in [24] and 96 in [33] in experiments<sup>1</sup>. However, it is important to note that the values of  $R_{L_{\max}}(s, t)$  with different  $L_{\max}$  values vary significantly. Table 2 lists the value of  $R_{L_{\max}}(s, t)$  for different  $L_{\max}$ . We observe that  $R_{L_{\max}}(s, t)$  only converges to within an additive error of  $1e-3$  after  $L_{\max} \geq 2,000$ . In other words, to estimate the ER score  $R(s, t)$  with an additive error of  $1e-3$ ,  $L_{\max}$  must be at least 2,000. Using smaller values of  $L_{\max}$  in experiments, as done in previous methods [24, 33], can lead to unsatisfactory approximation accuracy. This highlights the urgent need to improve the time dependence on  $L_{\max}$  for ER approximation.

## 1.1 Our Contributions

To address the aforementioned challenges, we present several contributions in this paper.

<sup>1</sup>The value of  $L_{\max}$  depends on the spectral radius  $\lambda$  of  $G$  and the additive error  $\epsilon$  specified in the approximation problem. For the Facebook dataset, the true spectral radius is  $\lambda = 0.9992$ , which results in  $L_{\max} = 14,546$  with  $\epsilon = 10^{-3}$ . In contrast, previous work [24] and [33] used  $\lambda = 0.9$  and  $\lambda = 0.9192$ , respectively, leading to much smaller values of  $L_{\max}$ , specifically 74 and 96.

First, we consider the problem of estimating the ER score  $R(s, t)$  for an arbitrary pair of nodes  $s$  and  $t$  in an undirected graph  $G$ . We improve the worst-case computational complexity from the previous best result of  $\tilde{O}\left(\frac{L_{\max}^3}{\epsilon^2 d^2}\right)$  to  $\tilde{O}\left(\min\left\{\frac{L_{\max}^{7/3}}{\epsilon^{2/3}}, \frac{L_{\max}^3}{\epsilon^2 d^2}, mL_{\max}\right\}\right)$ , where  $d = \min\{d(s), d(t)\}$  and  $m$  denotes the number of edges in  $G$ . Our result is asymptotically better than previous results, achieving a factor of  $\tilde{O}\left(\max\left\{\frac{L_{\max}^{2/3}}{\epsilon^{4/3} d^2}, 1, \frac{L_{\max}^2}{\epsilon^2 d^2 m}\right\}\right)$  improvements. Notably, in real-world networks where node degrees often follow power-law distributions—with most nodes having small degrees—our theoretical improvements are further enhanced due to the typically small values of  $d = \min\{d(s), d(t)\}$ .

Moreover, we conduct extensive experiments on real-world and synthetic large graphs to show the superiority of our algorithm. Experimental results show that our algorithm outperforms all baseline methods by up to an order of magnitude in both efficiency and accuracy.

In particular, we hope to draw attention to the time dependence on  $L_{\max}$  in ER approximation. We emphasize that  $L_{\max}$  is influenced by the mixing time of random walks and the spectral radius of the graph, and the true value of  $L_{\max}$  for real-world networks can be very large. Therefore, improving the dependence on  $L_{\max}$  has substantial theoretical and practical implications.

## 2 Preliminaries

**Notations.** We consider an undirected connected graph  $G = (V, E)$  with  $n$  nodes and  $m$  edges. The neighborhood of a node  $i$ , denoted as  $N(i)$ , consists of nodes that share an edge with it, i.e.,  $N(i) = \{j : \{i, j\} \in E\}$ . The degree of a node  $i$ , denoted as  $d(i)$ , is the number of nodes in its neighborhood. We use  $d_{\max}$  to represent the maximum degree,  $d_{\min}$  for the minimum degree, and  $\bar{d}$  for the average degree. The adjacency matrix  $\mathbf{A}$  is defined as  $A_{ij} = \mathbb{I}[\{i, j\} \in E]$ . The degree matrix  $\mathbf{D}$  is a diagonal matrix with  $D_{ii} = d(i)$ . The Laplacian matrix  $\mathbf{L}$  is defined as  $\mathbf{D} - \mathbf{A}$ . Additionally, the Moore-Penrose pseudo-inverse of  $\mathbf{L}$  is denoted as  $\mathbf{L}^+$ . Moreover,  $\mathbf{e}_s$  denotes the one-hot vector with the  $s$ -th element being 1. Table 3 summarizes frequently used notations for reference.

**Problem Definition.** In this paper, we consider the problem of estimating Single-Pair Effective Resistance (SPER), as defined below.

**Definition 2.1** (SPER Estimation with Absolute Error Guarantee). Given an undirected connected graph  $G = (V, E)$ , an arbitrary pair of nodes  $s, t \in V$ , an absolute error tolerance  $\epsilon > 0$ , and a failure probability  $0 < p_f \leq 1$ , the objective of a SPER query with an absolute error guarantee is to provide an estimator  $\hat{R}(s, t)$  such that

$$\Pr\left(\left|\hat{R}(s, t) - R(s, t)\right| < \epsilon\right) \geq 1 - p_f.$$

### 2.1 Probabilistic Interpretation

The ER score can be interpreted as transition probabilities of random walks on graphs. Specifically, the random walk on graphs is an important stochastic process. It operates as follows: starting from a node  $v_0$ , at each time step, when at node  $v_t$ , the next step involves moving to any neighboring node of  $v_t$  with a probability of  $1/d(v_t)$ . This random walk can also be viewed as a Markov chain with transition matrix  $\mathbf{P} = \mathbf{A}\mathbf{D}^{-1}$ . Let  $1 = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n \geq -1$

**Table 3: Frequently used notations.**

| Notation                                 | Description   |
|--|---|
| $G = (V, E)$                             | undirected graph with node set $V$ and edge set $E$ . |
| $n, m$                                   | number of nodes and edges.                            |
| $N(i)$                                   | neighborhood of node $i$ .                            |
| $d(i)$                                   | degree of node $i$ .                                  |
| $d_{\max}, d_{\min}, \bar{d}$            | maximum, minimum, and average degree.                 |
| $\mathbf{A}$                             | adjacency matrix of $G$ .                             |
| $\mathbf{D}$                             | degree matrix of $G$ .                                |
| $\mathbf{L} = \mathbf{D} - \mathbf{A}$   | Laplacian matrix of $G$ .                             |
| $\mathbf{L}^+$                           | Moore-Penrose inverse of $\mathbf{L}$ .               |
| $\mathbf{e}_s$                           | one-hot vector whose $s$ -th element is 1.            |
| $\mathbf{P} = \mathbf{A}\mathbf{D}^{-1}$ | transition matrix of $G$ .                            |
| $\lambda_i$                              | $i$ -th largest eigenvalue of $\mathbf{P}$ .          |
| $p^{(\ell)}(s, t)$                       | $\ell$ -hop transition probability from $s$ to $t$ .  |
| $h(s, t)$                                | hitting time from $s$ to $t$ .                        |
| $\kappa(s, t)$                           | commute time between $s$ and $t$ .                    |
| $R(s, t)$                                | effective resistance between $s$ and $t$ .            |
| $L_{\max}$                               | maximum random walk length (see Lemma 2.3).           |
| $\epsilon, p_f$                          | absolute error and the failure probability.           |
| $O(\cdot)$                               | big-Oh notation, asymptotic upper bound.              |
| $\tilde{O}(\cdot)$                       | soft big-Oh notation, ignoring some log factors.      |

denote the eigenvalues of  $\mathbf{P}$  sorted in descending order. The  $\ell$ -hop transition probability  $p^{(\ell)}(s, t)$  represents the probability that a random walk starting at node  $s$  visits node  $t$  after  $\ell$  hops. It's noteworthy that the  $\ell$ -hop transition probability matrix  $\mathbf{P}^{(\ell)} = (p^{(\ell)}(i, j))_{i, j}$  is equivalent to the  $\ell$ -th power of  $\mathbf{P}$ , i.e.,  $\mathbf{P}^{(\ell)} = \mathbf{P}^\ell$ . We will also highlight a well-known symmetric property concerning the multi-step transition probabilities (MSTP) on undirected graphs as follows:

**Lemma 2.2** (Symmetric Property of MSTP, [20]). *Given any undirected graph  $G = (V, E)$ , for any node  $s$  and  $t$ , and for all  $\ell$ , the following property of the  $\ell$ -hop transition probabilities holds:*

$$\frac{p^{(\ell)}(s, t)}{d(t)} = \frac{p^{(\ell)}(t, s)}{d(s)}.$$

The hitting time from  $s$  to  $t$ , denoted as  $h(s, t)$ , indicates the expected steps for a random walk starting from node  $s$  to hit node  $t$  for the first time. The commute time between  $s$  and  $t$ , denoted as  $\kappa(s, t)$ , is defined as  $\kappa(s, t) = h(s, t) + h(t, s)$ .

In particular, Peng et al. [24] provide Equation (2) to illustrate the connection between the Moore-Penrose pseudo-inverse of the Laplacian matrix  $\mathbf{L}^+$  and the MSTP matrix  $\mathbf{P}^\ell$ , given that  $(\mathbf{e}_s - \mathbf{e}_t)^\top \mathbf{L}^+ (\mathbf{e}_s - \mathbf{e}_t) = \sum_{\ell=0}^{\infty} (\mathbf{e}_s - \mathbf{e}_t)^\top \mathbf{D}^{-1} \mathbf{P}^\ell (\mathbf{e}_s - \mathbf{e}_t)$ . A line of research [24, 33] use the  $L_{\max}$ -truncated ER  $R_{L_{\max}}(s, t)$  as defined in Equation (3) to compute the estimate for  $R(s, t)$ .

Additionally, we need to bound the truncation error to satisfy the absolute error constraint. Peng et al. have proved such a bound in [24], which is later refined by Yang et al. in [33]. The bound from [33] is listed below:

**Lemma 2.3** (Maximum Steps Needed, [33]). *Given an undirected graph  $G$ ,  $|R(s, t) - R_L(s, t)| \leq \frac{\epsilon}{2}$  holds for  $s$  and  $t$  when  $L$  satisfies*

$$L \geq L_{\max} = \left\lceil \log_{1/\lambda} \frac{2(1/d(s) + 1/d(t))}{\epsilon(1-\lambda)} \right\rceil, \quad (4)$$

---

**Algorithm 1: Forward-Push<sub>s</sub>( $u, \ell$ )**


---

```

1  $\mathbf{q}_s^{(\ell)}(u) \leftarrow \mathbf{q}_s^{(\ell)}(u) + \mathbf{r}_s^{(\ell)}(u)$ 
2 for  $v \in N(u)$  do
3    $\mathbf{r}_s^{(\ell+1)}(v) \leftarrow \mathbf{r}_s^{(\ell+1)}(v) + \frac{\mathbf{r}_s^{(\ell)}(u)}{d(u)}$ 
4  $\mathbf{r}_s^{(\ell)}(u) \leftarrow 0$ 
    
```

---

where  $\lambda = \max\{\lambda_2, |\lambda_n|\}$ .

Notably, there is a strong connection between  $L_{\max}$  and the mixing time  $\tau_{\text{mix}}(\epsilon)$ , a crucial property of graphs. The mixing time of an undirected graph quantifies how quickly a random walk on the graph approaches the stationary distribution  $\boldsymbol{\pi} = \frac{\mathbf{D}\mathbf{1}}{2m}$ . It can be formally defined as  $\tau_{\text{mix}}(\epsilon) = \min \left\{ t : \max_{x \in V} \|\mathbf{P}^t \mathbf{e}_x - \boldsymbol{\pi}\|_{\text{TV}} \leq \epsilon \right\}$ ,

where  $\|\mathbf{P}^t \mathbf{e}_x - \boldsymbol{\pi}\|_{\text{TV}} = \frac{1}{2} \sum_{y \in V} |p^{(t)}(x, y) - \boldsymbol{\pi}(y)|$  is the total variation distance between two distributions. It can be shown that the mixing time is bounded by  $O\left(\frac{\log\left(\frac{m}{d_{\min}\epsilon}\right)}{1-\lambda}\right)$ . Additionally, according

to Equation (4), we have  $L_{\max} \leq \left\lceil \log_{1/\lambda} \frac{4}{d_{\min}\epsilon(1-\lambda)} \right\rceil$ . Since for any  $c > 0$ , we also have  $\lim_{\lambda \rightarrow 1} \frac{\log(A/(1-\lambda))}{\log(1/\lambda)} \Big/ \frac{A}{(1-\lambda)^{1+c}} = 0$ , we can bound

$L_{\max}$  as  $O\left(\frac{\log\left(\frac{1}{d_{\min}\epsilon}\right)}{(1-\lambda)^{1+c}}\right)$ , which shows a strong connection between

$L_{\max}$  and the mixing time  $\tau_{\text{mix}}(\epsilon)$ . In fact, in the proof of Lemma 2.3 [33], the authors merely change the definition of ‘‘close’’ by replacing the total variation distance with the  $L_\infty$  norm.


## 2.2 Key Techniques

According to the previous subsection, ER has a strong connection to MSTPs. In this section, we introduce some key techniques for estimating MSTPs, including the forward push operation, Monte Carlo sampling, and bidirectional methods. Notably, our method also utilizes some of these techniques.

**Forward Push Operation.** The forward push operation is a deterministic procedure to spread probability masses on graphs, originally proposed by Anderson et al. [2] for PageRank vector computation. Here, we introduce it in the context of MSTP computation.

The core idea of the forward push operation, with respect to a node  $s$ , involves maintaining two types of vectors for each layer (number of steps)  $0 \leq \ell \leq L_{\max}$ : reserve vectors  $\mathbf{q}_s^{(\ell)} \in \mathbb{R}^n$  and residue vectors  $\mathbf{r}_s^{(\ell)} \in \mathbb{R}^n$ . The  $u$ -th element of a reserve vector  $\mathbf{q}_s^{(\ell)}$  represents the accumulated probability mass on node  $u$  at layer  $\ell$  and serves as an underestimate of  $p^{(\ell)}(s, u)$ . The  $u$ -th element of a residue vector  $\mathbf{r}_s^{(\ell)}$  indicates the current active probability mass on node  $u$  at layer  $\ell$  that will be distributed to neighboring nodes in the next layer.

The Forward-Push procedure, detailed in Algorithm 1, operates as follows: invoking Forward-Push<sub>s</sub>( $u, \ell$ ) increases  $\mathbf{q}_s^{(\ell)}(u)$  by  $\mathbf{r}_s^{(\ell)}(u)$ , then distributes  $\mathbf{r}_s^{(\ell)}(u)$  evenly to its neighbors by increasing  $\mathbf{r}_s^{(\ell+1)}(v)$  by  $\frac{\mathbf{r}_s^{(\ell)}(u)}{d(u)}$  for each  $v \in N(u)$ , and finally set  $\mathbf{r}_s^{(\ell)}(u)$

**Figure 1: A running example of Forward-Push<sub>S</sub>(u, ℓ).**

to 0. We also present a running example in Figure 1 where node  $u$  has three neighbors.

**Monte Carlo Sampling.** Monte Carlo sampling is a classic technique for estimating probabilistic quantities. To estimate a quantity  $r$ , we design a random variable  $X$  such that  $\mathbb{E}[X] = r$  and  $\text{Var}[X] < \infty$ . By sampling  $N$  independent instances  $(X_i)_{i=1}^N$ , the empirical mean  $\frac{1}{N} \sum_{i=1}^N X_i$  serves as an estimate for  $r$ . In the context of SPER estimation, Monte Carlo sampling involves generating multiple  $L_{\max}$  random walks to estimate MSTPs, which then provides an estimator for  $R(s, t)$ .

**Bidirectional Methods.** Banerjee et al. [4] introduce a bidirectional algorithm named Bidirectional-MSTP (BiMSTP for short) for approximating MSTPs; and similar techniques have also been applied to (Personalized) PageRank approximation [20, 21, 30, 32]. Banerjee et al. [4] prove an invariant (Lemma 2.4), and integrate the forward push operation with Monte Carlo sampling to estimate the  $\ell$ -hop transition probability  $p^{(\ell)}(s, t)$ .

**Lemma 2.4** (Invariant, [4]). *If we initialize the reserve and residue vectors as  $\mathbf{q}_s^{(\ell)} = \mathbf{0}$  for  $\ell = 0, 1, \dots, L_{\max}$ , and  $\mathbf{r}_s^{(0)} = \mathbf{e}_s, \mathbf{r}_s^{(\ell)} = \mathbf{0}$  for  $\ell = 1, 2, \dots, L_{\max}$ , then after any sequence of Forward-Push<sub>S</sub>(u, ℓ) operations, for any  $0 \leq \ell \leq L_{\max}$ , the following invariant holds:*

$$p^{(\ell)}(s, t) = \mathbf{q}_s^{(\ell)}(t) + \sum_{k=0}^{\ell} \sum_{v \in V} \mathbf{r}_s^{(\ell-k)}(v) p^{(k)}(v, t).$$

### 3 Related Works

Existing works that focus on SPER estimation can be categorized into four types: estimating multi-step transition probabilities, utilizing landmark nodes, estimating commute time, and solving the Laplacian system of equations. In this section, we begin by introducing transition-probabilities-based methods, which are fast, have an absolute error guarantee, and are easy to implement. Following this, we provide a concise overview of other relevant lines of research.

#### 3.1 Transition-Probabilities-Based Methods

In their work [24], Peng et al. introduce the EstEff-TranProb algorithm, which employs random walk simulations to approximate the transition probabilities, subsequently deriving an estimator for  $R(s, t)$  with a computational complexity of  $\tilde{O}\left(\frac{L_{\max}^4}{\epsilon^2}\right)$ . Building on this, Yang et al. [33] suggest an approach that adaptively samples random walks, applying Bernstein's inequality for early termination. This innovation led to the development of the AMC algorithm and an enhanced variant, GEER. Both algorithms achieve a time complexity of  $\tilde{O}\left(\frac{L_{\max}^3}{\epsilon^2 d^2}\right)$ , where  $d = \min\{d(s), d(t)\}$ .

**Table 4: Time complexity of the algorithms.**

| Method                      | Query Time   |
|-----------------------------|--|
| EstEff-TranProb [24]        | $\tilde{O}\left(\frac{L_{\max}^4}{\epsilon^2}\right)$  |
| AMC / GEER [33]             | $\tilde{O}\left(\frac{L_{\max}^3}{\epsilon^2 d^2}\right)$  |
| EstEff-MC [24]              | $\tilde{O}\left(\frac{m}{(1-\lambda_2)^2 \epsilon^2 d}\right)$   |
| Lap. Solvers [9, 13–15, 29] | $\tilde{O}(m)$   |
| BiSPER (Ours)               | $\tilde{O}\left(\min\left\{\frac{L_{\max}^{7/3}}{\epsilon^{2/3}}, \frac{L_{\max}^3}{\epsilon^2 d^2}, mL_{\max}\right\}\right)$ |

### 3.2 Landmark-Based Methods

In their recent study, Liao et al. [18] introduce a novel technique for computing the SPER, proposing four distinct algorithms that leverage the concept of a landmark node. Subsequently, Liao et al. [19] expand the algorithms by moving from the use of a single landmark node to incorporating a set of landmark nodes. However, for all proposed algorithms except one, it is not possible to directly set the algorithm parameters to get an estimator with an absolute error guarantee, making these methods heuristic and lacking guaranteed error approximation. For the exception LocalTree algorithm, it has a worst-case running time of  $\tilde{O}\left(\frac{n^2(h(s,v)+h(t,v))}{\epsilon^2}\right)$ , where  $v$  is the landmark node. Additionally, the efficiency of all four algorithms is dependent on the choice of the landmark node, a decision that largely relies on heuristic approaches.

### 3.3 Commute-Time-Based Methods

Peng et al. [24] capitalized on the relation between the ER and the commute time, as detailed in [22], to develop their EstEff-MC algorithm. This algorithm provides an estimation of  $R(s, t)$  achieving a relative error margin within  $\epsilon$ . The algorithm's worst-case expected running time is formulated as  $\tilde{O}\left(\frac{m}{(1-\lambda_2)^2 d \epsilon^2}\right)$ , showcasing a direct proportionality to  $m$  and a quadratic dependence on  $\frac{1}{\epsilon}$ .

### 3.4 Laplacian-Solver-Based Methods

Per Equation (1), computing  $R(s, t)$  involves solving the linear system  $\mathbf{L}\mathbf{x} = \mathbf{e}_s - \mathbf{e}_t$  to obtain  $\mathbf{L}^+(\mathbf{e}_s - \mathbf{e}_t)$ , and then subtracting the  $t$ -th element from its  $s$ -th element. The advancements in [9, 13–15, 29] have achieved a nearly-linear time complexity. This enables the computation of an estimator for  $R(s, t)$  within  $\tilde{O}\left(m \log \frac{1}{\epsilon}\right)$  time.

## 4 Our Approach

In this section, we present our methodology for estimating the ( $L_{\max}$ -truncated) SPER  $R_{L_{\max}}(s, t)$ . We introduce the BiSPER algorithm, a novel approach that combines a redesigned forward push operation with adaptive Monte Carlo sampling techniques.

### 4.1 High Level Idea

By integrating Equation (3), Lemma 2.4, and the symmetry property of MSTP in Lemma 2.2, we can derive a direct estimator for  $R_{L_{\max}}(s, t)$ , using a number ( $N$ ) of random walks:

$$\begin{aligned} \hat{R}_{L_{\max}}(s, t) &= \sum_{\ell=0}^{L_{\max}} \left( \frac{\mathbf{q}_s^{(\ell)}(s)}{d(s)} - \frac{\mathbf{q}_s^{(\ell)}(t)}{d(t)} \right) + \sum_{\ell=0}^{L_{\max}} \left( \frac{\mathbf{q}_t^{(\ell)}(t)}{d(t)} - \frac{\mathbf{q}_t^{(\ell)}(s)}{d(s)} \right) \\ &+ \frac{1}{N} \sum_{i=1}^N \sum_{\ell=0}^{L_{\max}} \sum_{v \in V} X_{s,i}^{(\ell)}(v) \left( \sum_{k=0}^{L_{\max}-\ell} \frac{\mathbf{r}_s^{(k)}(v)}{d(v)} - \sum_{k=0}^{L_{\max}-\ell} \frac{\mathbf{r}_t^{(k)}(v)}{d(v)} \right) \\ &+ \frac{1}{N} \sum_{i=1}^N \sum_{\ell=0}^{L_{\max}} \sum_{v \in V} X_{t,i}^{(\ell)}(v) \left( \sum_{k=0}^{L_{\max}-\ell} \frac{\mathbf{r}_t^{(k)}(v)}{d(v)} - \sum_{k=0}^{L_{\max}-\ell} \frac{\mathbf{r}_s^{(k)}(v)}{d(v)} \right), \end{aligned} \quad (5)$$

where  $X_{s,i}^{(\ell)}(v) = \mathbb{I}[V_i^{(\ell)} = v \mid V_i^{(0)} = s]$  is an indicator random variable indicating whether the  $\ell$ -th node in the  $i$ -th random walk starting from  $s$ ,  $V_i^{(\ell)}$ , is  $v$ . Similarly,  $X_{t,i}^{(\ell)}(v)$  is defined for node  $t$ .

Our algorithm's core idea involves initially performing carefully designed forward push operations from nodes  $s$  and  $t$  until all degree-normalized residue values  $\frac{\mathbf{r}^{(k)}(v)}{d(v)}$  are below a specified parameter  $r_{\max}$ . Following this, we adaptively sample a designated number of random walks from both  $t$  and  $s$ , using the described estimator to compute an approximation of the  $L_{\max}$ -truncated SPER. This approach forms the basis of our Bidirectional Single-Pair Effective Resistance (BiSPER) algorithm.

The pseudo-code of our BiSPER algorithm is outlined in Algorithm 2. It operates in two phases: the *push phase* and the *adaptive Monte Carlo phase*. It invokes two procedures named BiSPER-Forward-Push and BiSPER-AMC, which will be elaborated in subsequent sections.

## 4.2 Push Phase and Maintaining Prefix Sums

In the push phase, our basic goal is to ensure all degree-normalized residues are below a threshold  $r_{\max}$ . This can be easily achieved by calling Forward-Push( $u, \ell$ ) for each  $\ell = 0, 1, \dots, L_{\max}$  and node  $u \in V$  where  $\frac{\mathbf{r}^{(\ell)}(u)}{d(u)} > r_{\max}$ .

However, according to Equation (5), the estimator  $\hat{R}_{L_{\max}}(s, t)$  involves querying  $\sum_{k=0}^{L_{\max}-\ell} \frac{\mathbf{r}^{(k)}(v)}{d(v)}$ , which can be seen as a prefix sum

of degree-normalized residue values  $\left( \frac{\mathbf{r}^{(k)}(v)}{d(v)} \right)_{k=0}^{L_{\max}}$ . If we query the prefix sum only when node  $v$  is sampled, it introduces  $O(L_{\max})$  additional cost, making each sample of random walk of length  $L_{\max}$  takes  $O(L_{\max}^2)$  time. Pre-computing all prefix sums after the push phase costs  $O(nL_{\max})$  time, which is linear to the number of nodes. Both solutions are computationally expensive.

To address this problem, we use Binary Indexed Trees (BITs) [10], also known as Fenwick Trees, to efficiently support prefix sum queries while dynamically maintaining the values in an array. Given an array  $(a_i)_{i=0}^{L-1}$  of length  $L$ , a BIT is a data structure associated with it, designed to support two fundamental operations:

- Update( $\ell, r$ ): Increases the  $\ell$ -th element of the original array by  $r$  and also maintain the data structure.
- Query( $\ell$ ): Computes and returns the prefix sum up to the  $\ell$ -th element, i.e.,  $\sum_{k=0}^{\ell} a_k$ .

<sup>2</sup>Hereafter, we may occasionally omit the subscript when the statement applies to both  $\mathbf{r}_s^{(k)}(v)$  and  $\mathbf{r}_t^{(k)}(v)$ , or both  $\mathbf{q}_s^{(k)}(v)$  and  $\mathbf{q}_t^{(k)}(v)$ , for simplicity. If multiple subscripts are omitted in the same formula, they represent the same value.

---

### Algorithm 2: BiSPER( $G, s, t, L_{\max}, \epsilon$ )

---

**Input:** Graph  $G = (V, E)$ , node  $s$  and  $t$ , maximum length  $L_{\max}$ , absolute error  $\epsilon$

**Output:**  $\hat{R}_{L_{\max}}(s, t)$

1  $d \leftarrow \min\{d(s), d(t)\}$ ,  $r_{\max} \leftarrow$  Equation (7)

2  $\mathbf{q}_s^{(\ell)}, \mathbf{q}_t^{(\ell)} \leftarrow \mathbf{0}$ ,  $\ell = 0, 1, \dots, L_{\max}$

3  $\mathbf{r}_s^{(0)} \leftarrow \mathbf{e}_s$ ,  $\mathbf{r}_t^{(0)} \leftarrow \mathbf{e}_t$ ,  $\mathbf{r}_s^{(\ell)}, \mathbf{r}_t^{(\ell)} \leftarrow \mathbf{0}$ ,  $\ell = 1, 2, \dots, L_{\max}$

4  $Q_s[s].\text{Update}(0, 1/d(s))$ ,  $Q_t[t].\text{Update}(0, 1/d(t))$

5 **for**  $\ell = 0, 1, \dots, L_{\max}$  **do**

6     **while**  $\exists u$ , such that  $\mathbf{r}_s^{(\ell)}(u)/d(u) > r_{\max}$  **do**

7         BiSPER-Forward-Push $_s(u, \ell)$

8     **while**  $\exists u$ , such that  $\mathbf{r}_t^{(\ell)}(u)/d(u) > r_{\max}$  **do**

9         BiSPER-Forward-Push $_t(u, \ell)$

/\* Provide upper bounds for sampling random walks. \*/

10  $T_{B_1} \leftarrow (L_{\max} + 1)(L_{\max} + 2)r_{\max}$

11  $T_{B_2} \leftarrow 2(L_{\max} + 1) - \sum_{\ell=0}^{L_{\max}} \sum_{v \in V} (\mathbf{q}_s^{(\ell)}(v) + \mathbf{q}_t^{(\ell)}(v))$

12 **if**  $r_{\max} \geq 1/d$  **then**

13      $N \leftarrow \lceil 8(L_{\max} + 1)^2 \log(2/p_f) / (\epsilon^2 d^2) \rceil$

14 **else**

15      $N \leftarrow \lceil 2 \min\{T_{B_1}^2, T_{B_2}^2\} \log(2/p_f) / \epsilon^2 \rceil$

16  $\hat{R}_{L_{\max}} \leftarrow$  BiSPER-AMC( $N$ )

17 **return**  $\hat{R}_{L_{\max}}$

---

Remarkably, both the Update( $\ell, r$ ) and Query( $\ell$ ) operations have a logarithmic time complexity of  $O(\log L)$ , where  $L$  is the array length. This logarithmic efficiency makes BITs an almost optimal choice for dynamically maintaining prefix sums in our BiSPER algorithm [25].

In the BiSPER algorithm, we employ two arrays of BITs,  $Q_s$  and  $Q_t$ , where each entry,  $Q_s[v]$  or  $Q_t[v]$ , represents a BIT associated with the sequences  $\left( \frac{\mathbf{r}_s^{(\ell)}(v)}{d(v)} \right)_{\ell=0}^{L_{\max}}$  or  $\left( \frac{\mathbf{r}_t^{(\ell)}(v)}{d(v)} \right)_{\ell=0}^{L_{\max}}$ , respectively. Using these BITs, we design a novel forward-push operation named BiSPER-Forward-Push, detailed in Algorithm 3. This procedure redistributes residue values layer by layer, ensuring the degree-normalized values do not exceed  $r_{\max}$ . At the same time, it dynamically maintains both the residue values and their prefix sums, which are essential for the subsequent sampling process.

## 4.3 Adaptive Monte Carlo Phase

Once all degree-normalized residues are below the threshold  $r_{\max}$ , we move to the adaptive Monte Carlo phase. This phase involves the BiSPER-AMC procedure, detailed in Algorithm 4.

The BiSPER-AMC procedure samples  $N$  random walks and uses the BITs, dynamically maintained during the push phase, to construct an estimator  $T(s, t)$ , which corresponds to the last two terms in Equation (5) (formally defined in Definition A.1).

In the  $i$ -th iteration, we sample two random walks of length  $L_{\max}$  starting from  $t$  and  $s$ . Let  $V_{t,i}^{(\ell)}$  and  $V_{s,i}^{(\ell)}$  denote the  $\ell$ -th sampled nodes in the  $i$ -th random walk. We construct the  $i$ -th sample of

**Algorithm 3:** BiSPER-Forward-Push<sub>s</sub>( $u, \ell$ )

---

```

1  $q_s^{(\ell)}(u) \leftarrow q_s^{(\ell)}(u) + r_s^{(\ell)}(u)$ 
2 for  $v \in N(u)$  do
   /*  $Q_s[v]$  is a BIT maintaining  $\sum_{k=0}^{\ell} \frac{r_s^{(k)}(v)}{d(v)}$ . */
3    $r_s^{(\ell+1)}(v) \leftarrow r_s^{(\ell+1)}(v) + \frac{r_s^{(\ell)}(u)}{d(u)}$ 
4    $Q_s[v].\text{Update}(\ell + 1, \frac{r_s^{(\ell)}(u)}{d(u)d(v)})$ 
5  $Q_s[u].\text{Update}(\ell, -\frac{r_s^{(\ell)}(u)}{d(u)})$ 
6  $r_s^{(\ell)}(u) \leftarrow 0$ 

```

---

**Algorithm 4:** BiSPER-AMC( $N$ )

---

```

1  $\hat{T}_{\text{sum}} \leftarrow 0, \hat{\sigma}_{\text{sum}}^2 \leftarrow 0$ 
2 for  $i = 1, 2, \dots, N$  do
   /* Sample two  $L_{\text{max}}$ -length random walks. */
3   Generate  $i$ -th  $L_{\text{max}}$ -length random walk from  $s$  and  $t$ :
    $\mathcal{W}_{s,i} = (V_{s,i}^{(0)} = s, V_{s,i}^{(1)}, \dots, V_{s,i}^{(L_{\text{max}})})$ ,
    $\mathcal{W}_{t,i} = (V_{t,i}^{(0)} = t, V_{t,i}^{(1)}, \dots, V_{t,i}^{(L_{\text{max}})})$ 
4    $\hat{T}_i \leftarrow \text{Equation (6)}$ 
5    $\hat{T}_{\text{sum}} \leftarrow \hat{T}_{\text{sum}} + \hat{T}_i, \hat{\sigma}_{\text{sum}}^2 \leftarrow \hat{\sigma}_{\text{sum}}^2 + \hat{T}_i^2$ 
6    $\hat{T} \leftarrow \hat{T}_{\text{sum}}/i, \hat{\sigma}^2 \leftarrow \hat{\sigma}_{\text{sum}}^2/i - \hat{T}^2$ 
7   if  $\sqrt{\frac{2\hat{\sigma}^2 \log(3/p_f)}{i} + \frac{6 \min\{T_{B_1}, T_{B_2}\} \log(3/p_f)}{i}} \leq \epsilon$  then
8     break
9    $\hat{R}_{L_{\text{max}}} \leftarrow \hat{T}$ 
10 for  $\ell = 0, 1, \dots, L_{\text{max}}$  do
11    $\hat{R}_{L_{\text{max}}} \leftarrow \hat{R}_{L_{\text{max}}} + \frac{q_s^{(\ell)}(s)}{d(s)} - \frac{q_s^{(\ell)}(t)}{d(t)} + \frac{q_t^{(\ell)}(t)}{d(t)} - \frac{q_t^{(\ell)}(s)}{d(s)}$ 
12 return  $\hat{R}_{L_{\text{max}}}$ 

```

---

$T(s, t)$  as follows:

$$\hat{T}_i(s, t) = \sum_{\ell=0}^{L_{\text{max}}} (Q_s[V_{s,i}^{(\ell)}].\text{Query}(L_{\text{max}} - \ell) - Q_t[V_{s,i}^{(\ell)}].\text{Query}(L_{\text{max}} - \ell) + Q_t[V_{t,i}^{(\ell)}].\text{Query}(L_{\text{max}} - \ell) - Q_s[V_{t,i}^{(\ell)}].\text{Query}(L_{\text{max}} - \ell)). \quad (6)$$

We then use the empirical mean of all  $\hat{T}_i(s, t)$ , denoted as  $\hat{T}(s, t)$ , as an estimator of  $T(s, t)$ . Early termination is facilitated using Bernstein's inequality for empirical variance (Lemma A.4). After the adaptive Monte Carlo phase,  $\hat{T}(s, t)$  is utilized to compute  $\hat{R}_{L_{\text{max}}}(s, t)$ , according to Equation (5) and the definition of  $T(s, t)$ .

#### 4.4 Theoretical Analysis

The correctness and time complexity of our BiSPER algorithm are established through the following theorems. First, we prove that the BiSPER algorithm provides an estimator with an absolute error guarantee, as stated in Theorem 4.1.

**Theorem 4.1** (Correctness of Approximation). *In our BiSPER algorithm, we have  $\Pr\left(|\hat{R}_{L_{\text{max}}}(s, t) - R_{L_{\text{max}}}(s, t)| < \epsilon\right) \geq 1 - p_f$ .*

Theorem 4.1 suggests that our BiSPER algorithm produces an unbiased estimator of the  $L_{\text{max}}$ -truncated Effective Resistance (ER). For the SPER value  $R(s, t)$ , we can obtain an approximation with an absolute error guarantee by setting  $L_{\text{max}}$  in BiSPER to the value specified in Lemma 2.3.

Next, we bound the worst-case time complexity of the BiSPER algorithm in Theorem 4.2.

**Theorem 4.2** (Time Complexity). *If we set  $r_{\text{max}}$  as Equation (7), the time complexity of BiSPER is  $\tilde{O}\left(\min\left\{\frac{L_{\text{max}}^{7/3}}{\epsilon^{2/3}}, \frac{L_{\text{max}}^3}{\epsilon^2 d^2}, mL_{\text{max}}\right\}\right)$ .*

$$r_{\text{max}} = \begin{cases} 0, & L_{\text{max}} \geq \max\left\{\frac{m^{1/2} \epsilon d}{2 \log^{1/2}(2/p_f)}, \frac{2m^{3/4} \epsilon^{1/2}}{3^{3/4} \log^{1/4}(2/p_f)}\right\}, \\ \frac{1}{d}, & d \geq \max\left\{\frac{2^{5/3} (L_{\text{max}}+1)^{1/3} \log^{1/3}(2/p_f)}{3^{1/2} \epsilon^{2/3}}, \frac{2(L_{\text{max}}+1) \log^{1/2}(2/p_f)}{m^{1/2} \epsilon}\right\}, \\ \frac{\epsilon^{2/3}}{2^{2/3} (L_{\text{max}}+1)^{4/3} \log^{1/3}(2/p_f)}, & \text{otherwise.} \end{cases} \quad (7)$$

Due to space limitations, we only state the main theorems in this section and defer the proofs to Appendices A and B.

#### 4.5 Discussions on Algorithm Lower Bounds and Comparison with Other Methods

From our theoretical analysis, we conclude that the BiSPER algorithm significantly improves by  $\tilde{O}\left(\frac{L_{\text{max}}^{2/3}}{\epsilon^{4/3} d^2}\right)$  over the previous best algorithms, AMC and GEER. In this subsection, we also discuss the computational complexity lower bound of SPER algorithms. We show that our BiSPER algorithm is near-optimal on hard instance graphs, and performs the best among all known algorithms on graphs other than hard instances.

In the full version [6] of their paper [7], Cai et al. establish a theorem that delineates the lower bound applicable to any algorithm designed for estimating the SPER:

**Theorem 4.3** (Theorem A.1, [6]). *There are  $c_0 > 0$  and infinitely many  $n$  such that given any  $d \in [4, n]$  and any  $\ell \in [4, n]$ , for graphs of  $n$  vertices and degree  $d$ , any local algorithm to approximate  $R_G(s, t)$  with success probability 0.6 and (relative) approximation ratio  $1 + c_0 \min\{d, \ell\}$  needs  $\Omega(dn/\ell)$  queries.*

Setting  $\ell = 4$  and  $d \geq \ell$ , we conclude that there exists a family of graphs where all algorithms estimating SPER within an additive error of  $\epsilon = 4c_0 R(s, t)$  require  $\Omega(dn/\ell) = \Omega(m)$  time. On these hard instances, nearly-linear time Laplacian solvers [9, 13–15, 29] are nearly optimal due to only poly-logarithmic overhead. However, this worst-case near-optimality allows for improved performance on other graph families while still ensuring nearly optimal running times for hard instances. Note that  $L_{\text{max}} = O\left(\log_{\frac{1}{\epsilon}} \frac{1}{\epsilon(1-\lambda)}\right)$ . For families of graphs with a spectral radius  $\lambda$  bounded by a constant less than 1 – commonly observed in real-world power-law graphs [26] – we have  $L_{\text{max}} = O\left(\log \frac{1}{\epsilon}\right)$  in such cases. Under this assumption, the time complexity is  $\tilde{O}\left(\frac{1}{\epsilon^2}\right)$  for the EstEff-TranProb algorithm [24],  $\tilde{O}\left(\frac{1}{\epsilon^2 d^2}\right)$  for AMC and GEER [33],  $\tilde{O}\left(\frac{m}{\epsilon^2 d}\right)$  for EstEff-MC [24],  $\tilde{O}(m)$  for Laplacian solvers [9, 13–15, 29], and

$\tilde{O}\left(\min\left\{\frac{1}{\epsilon^{2/3}}, \frac{1}{\epsilon^2 d^2}, m\right\}\right)$  for our BiSPER algorithm. Thus, under the  $1 - \lambda = \Theta(1)$  assumption, BiSPER performs the best among all candidate algorithms.

It should be noted that for landmark-based algorithms, it is not feasible to set the push threshold,  $r_{\max}$ , and the number of samples,  $T$ , to guarantee that an estimator meets the desired error requirement. Consequently, comparing the complexity of landmark-based algorithms with our BiSPER algorithm is not straightforward. In Experiments II and III in Section 5, we observe comparable empirical performance, indicating that BiSPER is at least as effective as landmark-based approaches. However, a critical advantage of our BiSPER algorithm is its capability to handle  $L_{\max}$ -truncated SPER queries, a feature not supported by the landmark-based algorithms.

Although the idea of using a bidirectional method in our BiSPER algorithm is inspired by the BiMSTP algorithm in [4], our BiSPER algorithm distinguishes itself from related works, including BiMSTP, in several key aspects:

- **Innovative Forward Push Procedure:** Our algorithm is the first to introduce a modified forward push procedure using Binary Indexed Trees to dynamically maintain the prefix sums of normalized residue values. This adaptation is non-trivial and essential for designing a faster bidirectional SPER estimation algorithm.
- **Align Bidirectional Methods with Related Works:** The BiMSTP algorithm guarantees expected running time only for randomly selected target nodes  $t$ . Additionally, it only ensures relative error guarantees for transition probabilities above a threshold  $\delta$ . Without positive lower bounds for MSTPs, direct application in SPER estimation with absolute error guarantees is not feasible. We resolve these issues by providing a detailed analysis that establishes worst-case time complexity for any nodes  $s$  and  $t$ . This aligns our algorithm with related works while achieving more favorable complexity bounds, highlighting its theoretical and practical superiority.

## 5 Experiments

In this section, we evaluate the efficiency of our BiSPER algorithm and compare it with other algorithms through experiments:

- First, we compare our BiSPER algorithm with other algorithms on real-world graphs of various sizes. Due to large  $L_{\max}$  required for accurate SPER approximation, obtaining the ground truth of SPERs on large graphs (such as Friendster) using the Power Iteration algorithm is infeasible. Therefore, we compare our BiSPER algorithm with other transition-probabilities-based algorithms for estimating  $L_{\max}$ -truncated ER using a fixed maximum length of  $L_{\max} = 100^3$ . This setting also ensures a fair comparison with methods designed for  $L_{\max}$ -truncated ER queries; and it is also similar to the choices made in prior works [24, 33].
- Then we compare our BiSPER algorithm with other algorithms on the relative small graphs. We use Power Iteration to get the ground truths on the Facebook dataset and a synthetic Erdős-Rényi graph, where obtaining ground truths is feasible, and then compare the efficiency of representative transition-probabilities-based and landmark-based algorithms on them.

<sup>3</sup>Even when  $L_{\max}$  is set to 100 it still takes about one week to compute the truncated SPER values on the Friendster dataset.

**Table 5: Statistics of datasets.**

| Name        | $n$        | $m$           | $d_{\min}$ | $d_{\max}$ | $\bar{d}$ | $\lambda$ |
|-------------|------------|---------------|------------|------------|-----------|-----------|
| Facebook    | 4,039      | 88,234        | 1          | 1045       | 43.69     | 0.9992    |
| DBLP        | 317,080    | 1,049,866     | 1          | 343        | 6.62      | 0.9973    |
| Youtube     | 1,134,890  | 2,987,624     | 1          | 28754      | 5.27      | 0.9980    |
| Orkut       | 3,072,441  | 117,185,083   | 1          | 33313      | 76.28     | 0.9948    |
| LiveJournal | 3,997,962  | 34,681,189    | 1          | 14815      | 17.35     | 0.9999    |
| Friendster  | 65,608,366 | 1,806,067,135 | 1          | 5214       | 55.06     | 0.9995    |

### 5.1 Datasets and Implementation Details

Our experiments are conducted on six real-world SNAP datasets [17], varying in size as outlined in Table 5, and a synthetic Erdős-Rényi graph, detailed in Experiment III. All our experiments are conducted on a Linux server equipped with an Intel(R) Xeon(R) Silver 4114 CPU @ 2.20GHz 40-core processor and 692GB of RAM. Reading, loading of graphs, and allocation of space for data structures are treated as pre-processing steps and excluded from the computation of running time. We implement all algorithms in C++ and compile them using g++ 7.5 with the -O3 optimization flag. No parallelism techniques are used in our code except in the ground truth generation with Power Iteration. The eigenvalues  $\lambda$  of all graphs are obtained via the `scipy.sparse.linalg.eigs` function, which utilizes the ARPACK [16] algorithm. Our code is available at: <https://github.com/GuanyuCui/BiSPER>.

### 5.2 Experiment I: Query Efficiency for $R_{L_{\max}}(s, t)$ on Real-World Graphs

**Competitors and Parameters.** In this experiment, we aim to assess the performance of our BiSPER algorithm compared to other algorithms that use transition probabilities, specifically AMC and GEER, as introduced by Yang et al. [33]. We exclude EstEff-TranProb [24] as a baseline since Yang et al. [33] have shown that AMC and GEER outperform it.

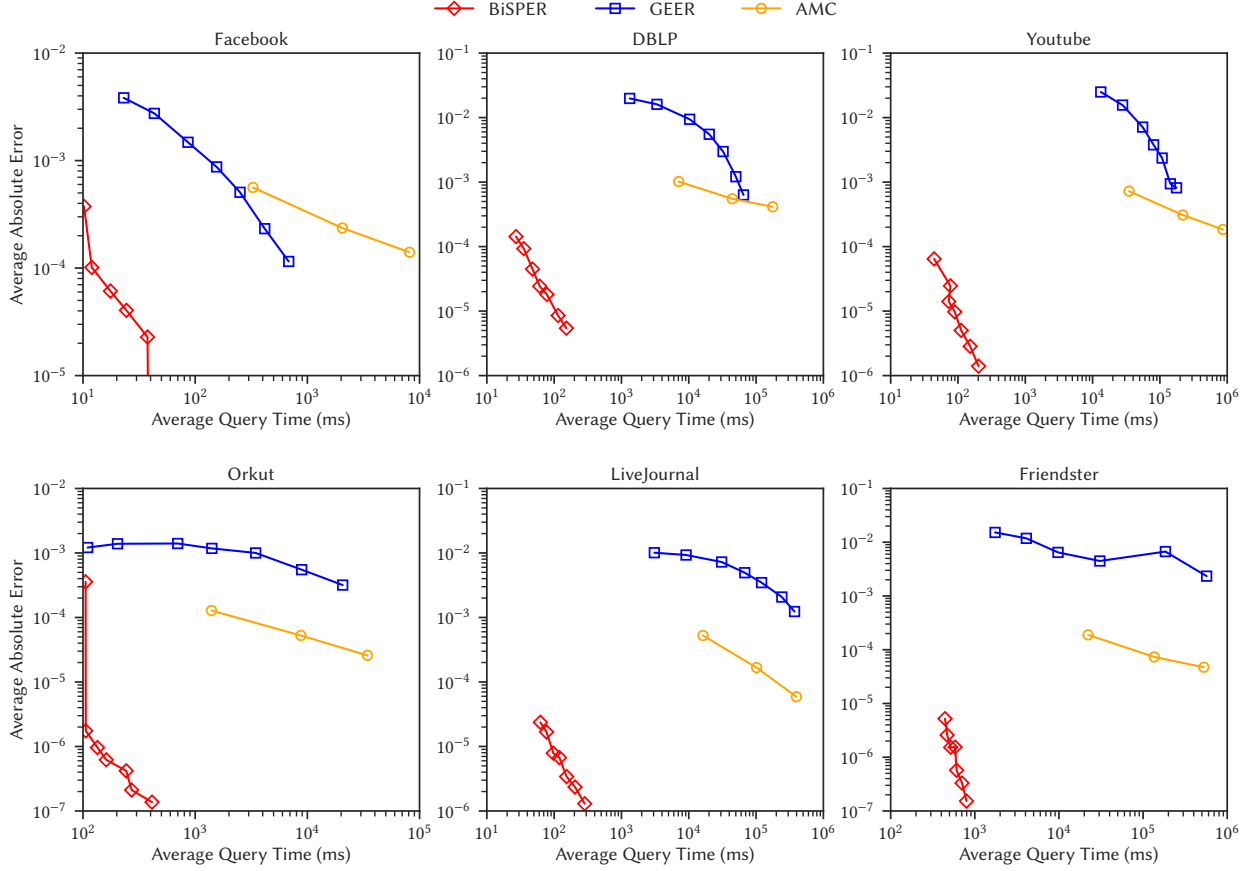
We will select 100 pairs of nodes from each graph dataset through uniform sampling, keeping these pairs constant across evaluations to ensure fair comparisons. The maximum random walk length,  $L_{\max}$ , will be standardized to 100 for all algorithms. To determine the ground truth values for ER, we will use the Power Iteration method to compute transition probabilities, followed by Equation (3) to calculate the  $L_{\max}$ -truncated ER values.

The failure probability  $p_f$  is set to 0.01 for all algorithms that allow this parameter, following the setting in [33]. Other parameters for the algorithms will be varied as follows:

- For BiSPER and GEER, the error parameter  $\epsilon$  will be varied within  $\{1e-3, 2e-3, 5e-3, 1e-2, 2e-2, 5e-2, 1e-1\}$ , except for GEER on the Friendster dataset, where it will be within  $\{1e-2, 2e-2, 5e-2, 1e-1, 2e-1, 5e-1\}$  due to the rapidly increasing running time of GEER as  $\epsilon$  decreases.
- For AMC,  $\epsilon$  will be varied within  $\{1e-1, 2e-1, 5e-1\}$ , due to similar reason.

**Results and Analysis.** We run the algorithms under various parameters, recording their average running times in milliseconds and the average absolute errors across all 100 queries. The relationship between the average running time and the average absolute error for each dataset is visualized through line charts in Figure 2. The figure shows that BiSPER not only achieves the shortest running

**Figure 2: Results of Experiment I. The average error of the two outliers on the Facebook dataset is less than  $1e-15$ .**

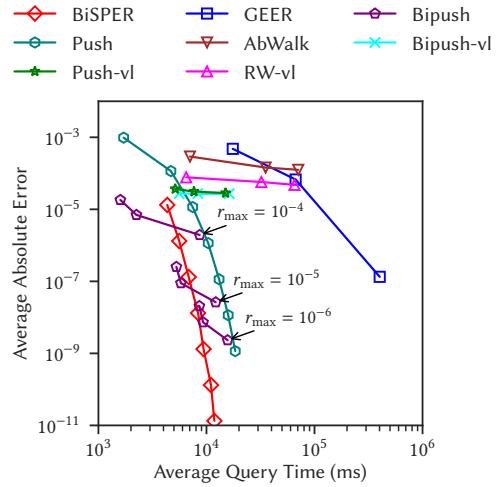


times, consistently below  $10^3$  ms, but also maintains the lowest average absolute errors, not exceeding  $10^{-3}$ . On datasets like DBLP, Youtube, and LiveJournal, BiSPER shows a substantial speed improvement, exceeding a 10x enhancement over GEER. Across all datasets, BiSPER outperforms AMC with more than a 10x increase in speed. Notably, BiSPER’s efficiency on graphs with lower average degrees supports our complexity analysis.

### 5.3 Experiment II: Query Efficiency for $R(s, t)$ on Real-World Graphs

**Competitors and Parameters.** In this experiment, we aim to evaluate the performance of our BiSPER algorithm against transition-probabilities-based and landmark-based methods. This includes the GEER algorithm by Yang et al. [33], Bipush, Push and AbWalk by Liao et al. [18], and Bipush-vl, Push-vl and RW-vl by Liao et al. [19], using the Facebook dataset as our testbed. AMC is excluded from this comparison due to its longer execution time, as revealed in Experiment I. Among the landmark-based algorithms, only Bipush, Push, and AbWalk are selected for their superior performance in [18]. We also exclude EstEff-MC [24], following Yang et al.’s findings that AMC and GEER surpass it. Laplacian solvers are omitted due to the lack of practical implementations.

**Figure 3: Results of Experiment II.**



For consistency, we use the same 100 node pairs sampled from the Facebook graph as in the previous experiment, with the failure probability  $p_f$  keeping 0.01. To determine the ground truth ER

values, we set  $\epsilon = 1e-17$  and use Equation (4) to ascertain the maximum random walk length,  $L_{\max, \epsilon=1e-17}$ . The  $L_{\max}$ -truncated ER values, serving as our ground truth, are calculated using Power Iteration and Equation (3).

For the three single-landmark algorithms, the node with the largest degree is selected as the landmark. For the three multi-landmark algorithms, the top 100 nodes with the highest degrees are chosen as landmarks, and the number of sampled random walks is set to  $1e5$  to construct the index structure. Other parameters for the algorithms are varied as follows:

- For BiSPER,  $\epsilon$  in  $\{1e-7, 1e-6, 1e-5, \dots, 1e-1\}$ .
- For GEER,  $\epsilon$  in  $\{1e-3, 1e-2, 1e-1\}$ , due to its rapidly increasing running time as  $\epsilon$  decreases.
- For Bipush, we vary the error parameter  $r_{\max}$  and the number of samples  $T$  in  $\{1e-6, 1e-5, 1e-4\}$  and  $\{1e3, 1e4, 1e5\}$ .
- For Push, we vary  $r_{\max}$  in  $\{1e-10, 1e-9, \dots, 1e-4\}$ .
- For AbWalk, we vary  $T$  in  $\{1e5, 5e5, 1e6\}$ .
- For Bipush-vl, we vary  $r_{\max}$  in  $\{1e-7, 2e-7, 3e-7\}$  and set  $T$  to be 1,000.
- For Push-vl, we vary  $r_{\max}$  in  $\{1e-7, 2e-7, 3e-7\}$ .
- For RW-vl, we vary  $T$  in  $\{1e6, 5e6, 1e7\}$ .

**Results and Analysis.** As Experiment I, the relationship between the average running time and the average absolute error for each dataset is visualized through line charts in Figure 3. It reveals that BiSPER, Bipush, and Push demonstrate comparable efficiencies, with BiSPER showing superior performance, particularly when the average absolute error is below  $10^{-7}$ . This highlights BiSPER’s effectiveness in achieving lower errors. It’s also noteworthy that the landmark-based algorithms, Bipush and Push, cannot guarantee an absolute error below a pre-determined threshold  $\epsilon$  through direct setting of the error parameters  $r_{\max}$  or sample size  $T$ .

#### 5.4 Experiment III: Query Efficiency for $R(s, t)$ on Synthetic Graphs

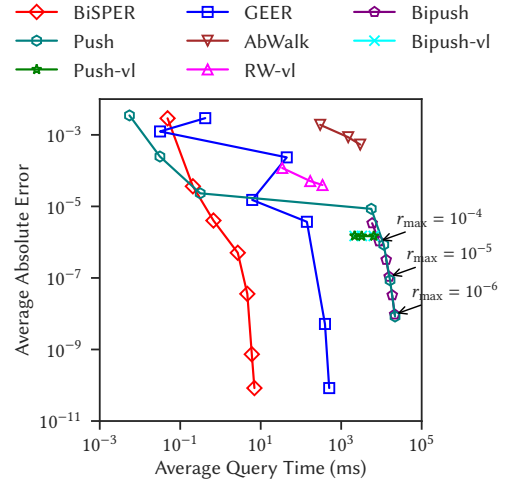
Real-world networks often follow power-law distributions. To examine how various ER algorithms perform on graphs that don’t follow these distributions, we conducted a series of experiments using synthetic Erdős-Rényi random graphs.

**Dataset Generation.** We create an Erdős-Rényi random graph with parameters  $(n, p) = (5000, 0.005)$  and focused on its largest connected component for our synthetic dataset. This component consists of 5,000 nodes and 62,619 edges, featuring a minimum degree  $d_{\min} = 8$ , a maximum degree  $d_{\max} = 44$ , an average degree  $\bar{d} = 25.05$ , and  $\lambda = 0.3900$ .

**Competitors and Parameters.** Our comparison of the BiSPER algorithm uses the baseline algorithms from Experiment II as benchmarks. We varied their parameters as follows:

- For BiSPER and GEER, we vary  $\epsilon$  in  $\{1e-7, 1e-6, \dots, 1e-1\}$ .
- For Bipush, we vary the error parameter  $r_{\max}$  and the number of samples  $T$  in  $\{1e-6, 1e-5, 1e-4\}$  and  $\{1e3, 1e4\}$ .
- For Push, we vary  $r_{\max}$  in  $\{1e-7, 1e-6, \dots, 1e-1\}$ .
- For AbWalk, we vary  $T$  in  $\{1e3, 5e3, 1e4\}$ .
- For Bipush-vl, we vary  $r_{\max}$  in  $\{1e-7, 2e-7, 3e-7\}$  and set  $T$  to be 1,000.
- For Push-vl, we vary  $r_{\max}$  in  $\{1e-7, 2e-7, 3e-7\}$ .
- For RW-vl, we vary  $T$  in  $\{1e4, 5e4, 1e5\}$ .

**Figure 4: Results of Experiment III.**



**Results and Analysis.** Consistent with earlier experiments, Figure 4 presents the trade-off between average absolute error and average running time across different algorithms on synthetic graphs. As shown in the figure, our BiSPER algorithm consistently performs well. In contrast, landmark-based algorithms, typically strong performers on real-world graphs, struggle in regions of small absolute error. The lack of high-degree hub nodes in Erdős-Rényi graphs likely hinders effective landmark selection, leading to unguaranteed running times for these algorithms.

## 6 Conclusion

Effective resistance is a crucial measure of proximity in graph theory. However, current Single-Pair Effective Resistance (SPER) estimation algorithms often suffer from high worst-case complexity, lack error guarantees, or are too theoretical and complex for practical use. In this paper, we introduce the BiSPER algorithm, a refined bidirectional approach. Our theoretical analysis and extensive experimental evaluations demonstrates BiSPER’s superiority over methods based on transition probabilities and commute times.

## Acknowledgments

This research was supported in part by National Natural Science Foundation of China (No. 92470128, No. U2241212, No. 61932001), by National Science and Technology Major Project (2022ZD0114802), by Beijing Natural Science Foundation (No. 4222028), by Beijing Outstanding Young Scientist Program No.BJJWZYJH012019100020098, by Huawei-Renmin University joint program on Information Retrieval. Hanzhi Wang is part of BARC, Basic Algorithms Research Copenhagen, supported by the VILLUM Foundation grant 54451. We also wish to acknowledge the support provided by the fund for building world-class universities (disciplines) of Renmin University of China, by Engineering Research Center of Next-Generation Intelligent Search and Recommendation, Ministry of Education, by Intelligent Social Governance Interdisciplinary Platform, Major Innovation & Planning Interdisciplinary Platform for the “Double-First Class” Initiative, Public Policy and Decision-making Research Lab, and Public Computing Cloud, Renmin University of China.

## References

- [1] Vedat Levi Alev, Nima Anari, Lap Chi Lau, and Shayan Oveis Gharan. 2018. Graph Clustering using Effective Resistance. In *9th Innovations in Theoretical Computer Science Conference, ITCs 2018, January 11–14, 2018, Cambridge, MA, USA (LIPIcs, Vol. 94)*, Anna R. Karlin (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 41:1–41:16. doi:10.4230/LIPIcs.ITCS.2018.41
- [2] Reid Andersen, Fan R. K. Chung, and Kevin J. Lang. 2006. Local Graph Partitioning using PageRank Vectors. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), 21–24 October 2006, Berkeley, California, USA, Proceedings*. IEEE Computer Society, 475–486. doi:10.1109/FOCS.2006.44
- [3] Jean-Yves Audibert, Rémi Munos, and Csaba Szepesvári. 2007. Tuning Bandit Algorithms in Stochastic Environments. In *Algorithmic Learning Theory, 18th International Conference, ALT 2007, Sendai, Japan, October 1–4, 2007, Proceedings (Lecture Notes in Computer Science, Vol. 4754)*, Marcus Hutter, Rocco A. Servedio, and Eiji Takimoto (Eds.). Springer, 150–165. doi:10.1007/978-3-540-75225-7\_15
- [4] Siddhartha Banerjee and Peter Lofgren. 2015. Fast Bidirectional Probability Estimation in Markov Models. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7–12, 2015, Montreal, Quebec, Canada, Corinna Cortes, Neil D. Lawrence, Daniel D. Lee, Masashi Sugiyama, and Roman Garnett (Eds.)*. 1423–1431. <https://proceedings.neurips.cc/paper/2015/hash/ede7e2b6d13a41ddf9f4bdef84f4c737-Abstract.html>
- [5] Mitchell Black, Zhengchao Wan, Amir Nayyeri, and Yusu Wang. 2023. Understanding Oversquashing in GNNs through the Lens of Effective Resistance. In *International Conference on Machine Learning, ICML 2023, 23–29 July 2023, Honolulu, Hawaii, USA (Proceedings of Machine Learning Research, Vol. 202)*, Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (Eds.). PMLR, 2528–2547. <https://proceedings.mlr.press/v202/black23a.html>
- [6] Dongrun Cai, Xue Chen, and Pan Peng. 2023. Effective Resistances in Non-Expander Graphs. *CoRR* abs/2307.01218 (2023). doi:10.48550/ARXIV.2307.01218 arXiv:2307.01218
- [7] Dongrun Cai, Xue Chen, and Pan Peng. 2023. Effective Resistances in Non-Expander Graphs. In *31st Annual European Symposium on Algorithms, ESA 2023, September 4–6, 2023, Amsterdam, The Netherlands (LIPIcs, Vol. 274)*, Inge Li Gørtz, Martin Farach-Colton, Simon J. Puglisi, and Grzegorz Herman (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 29:1–29:18. doi:10.4230/LIPIcs.ESA.2023.29
- [8] Paul F. Christiano, Jonathan A. Kelner, Aleksander Madry, 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 ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6–8 June 2011*, Lance Fortnow and Salil P. Vadhan (Eds.). ACM, 273–282. doi:10.1145/1993636.1993674
- [9] Michael B. Cohen, Rasmus Kyng, Gary L. Miller, Jakob W. Pachocki, Richard Peng, Anup B. Rao, and Shen Chen Xu. 2014. Solving SDD linear systems in nearly  $m \log^{1/2} n$  time. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 – June 03, 2014*, David B. Shmoys (Ed.). ACM, 343–352. doi:10.1145/2591796.2591833
- [10] Peter M. Fenwick. 1994. A New Data Structure for Cumulative Frequency Tables. *Softw. Pract. Exp.* 24, 3 (1994), 327–336. doi:10.1002/SPE.4380240306
- [11] Wassily Hoeffding. 1994. Probability inequalities for sums of bounded random variables. *The collected works of Wassily Hoeffding* (1994), 409–426.
- [12] Seok-Hee Hong, Juan Pablo Bonilla Atades, Rowena Kok, Amyra Meidiana, and Kunsu Park. 2023. New Seeding Strategies for the Influence Maximization Problem. In *Complex Networks & Their Applications XII - Proceedings of The Twelfth International Conference on Complex Networks and their Applications: COMPLEX NETWORKS 2023, Menton, France, 28–30 November, 2023, Volume 2 (Studies in Computational Intelligence, Vol. 1142)*, Hocine Cherifi, Luis M. Rocha, Chantal Cherifi, and Murat Donduran (Eds.). Springer, 289–299. doi:10.1007/978-3-031-53499-7\_23
- [13] Arun Jambulapati and Aaron Sidford. 2021. Ultrasparse Ultrasparsifiers and Faster Laplacian System Solvers. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 – 13, 2021*, Daniel Marx (Ed.). SIAM, 540–559. doi:10.1137/1.9781611976465.33
- [14] Ioannis Koutis, Gary L. Miller, and Richard Peng. 2010. Approaching Optimality for Solving SDD Linear Systems. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23–26, 2010, Las Vegas, Nevada, USA*. IEEE Computer Society, 235–244. doi:10.1109/FOCS.2010.29
- [15] Ioannis Koutis, Gary L. Miller, and Richard Peng. 2011. A Nearly- $m \log n$  Time Solver for SDD Linear Systems. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22–25, 2011*, Rafail Ostrovsky (Ed.). IEEE Computer Society, 590–598. doi:10.1109/FOCS.2011.85
- [16] Richard B. Lehoucq, Danny C. Sorensen, and Chao Yang. 1998. *ARPACK users' guide - solution of large-scale eigenvalue problems with implicitly restarted Arnoldi methods*. SIAM.
- [17] Jure Leskovec and Andrej Krevl. 2014. SNAP Datasets: Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data>.
- [18] Meihao Liao, Rong-Hua Li, Qiangqiang Dai, Hongyang Chen, Hongchao Qin, and Guoren Wang. 2023. Efficient Resistance Distance Computation: The Power of Landmark-based Approaches. *Proc. ACM Manag. Data* 1, 1 (2023), 68:1–68:27. doi:10.1145/3588922
- [19] Meihao Liao, Junjie Zhou, Rong-Hua Li, Qiangqiang Dai, Hongyang Chen, and Guoren Wang. 2024. Efficient and Provable Effective Resistance Computation on Large Graphs: An Index-based Approach. *Proc. ACM Manag. Data* 2, 3 (2024), 133. doi:10.1145/3654936
- [20] Peter Lofgren, Siddhartha Banerjee, and Ashish Goel. 2015. Bidirectional PageRank Estimation: From Average-Case to Worst-Case. In *Algorithms and Models for the Web Graph - 12th International Workshop, WAW 2015, Eindhoven, The Netherlands, December 10–11, 2015, Proceedings (Lecture Notes in Computer Science, Vol. 9479)*, David F. Gleich, Júlia Komjáthy, and Nelly Litvak (Eds.). Springer, 164–176. doi:10.1007/978-3-319-26784-5\_13
- [21] Peter Lofgren, Siddhartha Banerjee, and Ashish Goel. 2016. Personalized PageRank Estimation and Search: A Bidirectional Approach. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining, San Francisco, CA, USA, February 22–25, 2016*, Paul N. Bennett, Vanja Josifovski, Jennifer Neville, and Filip Radlinski (Eds.). ACM, 163–172. doi:10.1145/2835776.2835823
- [22] László Lovász. 1993. Random walks on graphs. *Combinatorics, Paul erdos is eighty* 2, 1–46 (1993), 4.
- [23] Volodymyr Mnih, Csaba Szepesvári, and Jean-Yves Audibert. 2008. Empirical Bernstein stopping. In *Machine Learning, Proceedings of the Twenty-Fifth International Conference (ICML 2008), Helsinki, Finland, June 5–9, 2008 (ACM International Conference Proceeding Series, Vol. 307)*, William W. Cohen, Andrew McCallum, and Sam T. Roweis (Eds.). ACM, 672–679. doi:10.1145/1390156.1390241
- [24] Pan Peng, Daniel Lopatta, Yuichi Yoshida, and Gramoz Goranci. 2021. Local Algorithms for Estimating Effective Resistance. In *KDD '21: The 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, Singapore, August 14–18, 2021*, Feida Zhu, Beng Chin Ooi, and Chunyan Miao (Eds.). ACM, 1329–1338. doi:10.1145/3447548.3467361
- [25] Giulio Ermanno Pibiri and Rossano Venturini. 2021. Practical trade-offs for the prefix-sum problem. *Softw. Pract. Exp.* 51, 5 (2021), 921–949. doi:10.1002/SPE.2918
- [26] Yi Qi, Wanyue Xu, Liwang Zhu, and Zhongzhi Zhang. 2021. Real-World Networks Are Not Always Fast Mixing. *Comput. J.* 64, 2 (2021), 236–244. doi:10.1093/COMJNL/BXAA150
- [27] Sawyer Jack Robertson, Zhengchao Wan, and Alexander Cloninger. 2024. All You Need is Resistance: On the Equivalence of Effective Resistance and Certain Optimal Transport Problems on Graphs. *CoRR* abs/2404.15261 (2024). doi:10.48550/ARXIV.2404.15261 arXiv:2404.15261
- [28] Daniel A. Spielman and Nikhil Srivastava. 2008. Graph sparsification by effective resistances. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17–20, 2008*, Cynthia Dwork (Ed.). ACM, 563–568. doi:10.1145/1374376.1374456
- [29] Daniel A. Spielman and Shang-Hua Teng. 2004. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing, Chicago, IL, USA, June 13–16, 2004*, László Babai (Ed.). ACM, 81–90. doi:10.1145/1007352.1007372
- [30] Hanzhi Wang, Zhewei Wei, Ji-Rong Wen, and Mingji Yang. 2024. Revisiting Local Computation of PageRank: Simple and Optimal. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada, June 24–28, 2024*, Bojan Mohar, Igor Shinkar, and Ryan O'Donnell (Eds.). ACM, 911–922. doi:10.1145/3618260.3649661
- [31] Kazuyuki Yamashita and Hiroyuki Ohsaki. 2021. Effective Node Resistance and its Implication on Network Robustness. In *International Conference on Information Networking, ICOIN 2021, Jeju Island, South Korea, January 13–16, 2021*. IEEE, 149–153. doi:10.1109/ICOIN50884.2021.9333923
- [32] Mingji Yang, Hanzhi Wang, Zhewei Wei, Sibow Wang, and Ji-Rong Wen. 2024. Efficient Algorithms for Personalized PageRank Computation: A Survey. *IEEE Trans. Knowl. Data Eng.* 36, 9 (2024), 4582–4602. doi:10.1109/TKDE.2024.3376000
- [33] Renchi Yang and Jing Tang. 2023. Efficient Estimation of Pairwise Effective Resistance. *Proc. ACM Manag. Data* 1, 1 (2023), 16:1–16:27. doi:10.1145/3588696
- [34] Bohang Zhang, Shengjie Luo, Liwei Wang, and Di He. 2023. Rethinking the Expressive Power of GNNs via Graph Biconnectivity. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1–5, 2023*. OpenReview.net. <https://openreview.net/forum?id=r9hNv76Kot3>

## A Proof of Theorem 4.1

We will first prove the unbiasedness of the estimators, and then prove the correctness of the approximation. We begin with the definition of the estimators.

**Definition A.1** (Estimator). Let  $\hat{R}_{L_{\max}}(s, t)$  be the  $L_{\max}$ -truncated effective resistance estimator defined in Equation (5). For convenience, we define another value  $T(s, t)$  as follows:

$$T(s, t) = \sum_{\ell=0}^{L_{\max}} \sum_{v \in V} p^{(\ell)}(s, v) \left( \sum_{k=0}^{L_{\max}-\ell} \frac{\mathbf{r}_s^{(k)}(v)}{d(v)} - \sum_{k=0}^{L_{\max}-\ell} \frac{\mathbf{r}_t^{(k)}(v)}{d(v)} \right) + \sum_{\ell=0}^{L_{\max}} \sum_{v \in V} p^{(\ell)}(t, v) \left( \sum_{k=0}^{L_{\max}-\ell} \frac{\mathbf{r}_t^{(k)}(v)}{d(v)} - \sum_{k=0}^{L_{\max}-\ell} \frac{\mathbf{r}_s^{(k)}(v)}{d(v)} \right).$$

We also define each sample  $\hat{T}_i(s, t)$  as follows:

$$\hat{T}_i(s, t) = \sum_{\ell=0}^{L_{\max}} \sum_{v \in V} X_{s,i}^{(\ell)}(v) \left( \sum_{k=0}^{L_{\max}-\ell} \frac{\mathbf{r}_s^{(k)}(v)}{d(v)} - \sum_{k=0}^{L_{\max}-\ell} \frac{\mathbf{r}_t^{(k)}(v)}{d(v)} \right) + \sum_{\ell=0}^{L_{\max}} \sum_{v \in V} X_{t,i}^{(\ell)}(v) \left( \sum_{k=0}^{L_{\max}-\ell} \frac{\mathbf{r}_t^{(k)}(v)}{d(v)} - \sum_{k=0}^{L_{\max}-\ell} \frac{\mathbf{r}_s^{(k)}(v)}{d(v)} \right),$$

and the estimator  $\hat{T}(s, t)$ , or  $\hat{T}$  for short, as follows:

$$\hat{T}(s, t) = \frac{1}{N} \sum_{i=1}^N T_i(s, t).$$

Next, we prove the unbiasedness of the estimators defined above.

**Lemma A.2** (Unbiasedness of  $\hat{T}$  and  $\hat{R}_{L_{\max}}$ ). *In Algorithm 4, we have  $\mathbf{E}[\hat{T}(s, t)] = T(s, t)$  and  $\mathbf{E}[\hat{R}_{L_{\max}}(s, t)] = R_{L_{\max}}(s, t)$ .*

**PROOF.** According to the definition of  $\hat{T}$ , we have:

$$\mathbf{E}[\hat{T}(s, t)] = \frac{1}{N} \sum_{i=1}^N \mathbf{E}[T_i(s, t)] = T(s, t).$$

And for  $\hat{R}_{L_{\max}}(s, t)$ , we have

$$\mathbf{E}[\hat{R}_{L_{\max}}(s, t)] = \mathbf{E}\left[\hat{T}(s, t) + \sum_{\ell=0}^{L_{\max}} \left( \frac{\mathbf{q}_s^{(\ell)}(s)}{d(s)} - \frac{\mathbf{q}_s^{(\ell)}(t)}{d(t)} \right) + \sum_{\ell=0}^{L_{\max}} \left( \frac{\mathbf{q}_t^{(\ell)}(t)}{d(t)} - \frac{\mathbf{q}_t^{(\ell)}(s)}{d(s)} \right)\right].$$

The conclusion follows if we expand the right side using the linearity of the expectation and then apply  $\mathbf{E}[\hat{T}(s, t)] = T(s, t)$ .  $\square$

Before we continue to the correctness of the approximation, we need the following two lemmas: Hoeffding's inequality and Bernstein's inequality for empirical variance.

**Lemma A.3** (Hoeffding's Inequality, [11]). *Let  $Z_1, Z_2, \dots, Z_n$  be independent random variables bounded in the interval  $[a, b]$  with length  $B = b - a$  for all  $1 \leq i \leq n$ . Then the following inequality holds:*

$$\Pr\left(\left|\frac{1}{n} \sum_{i=1}^n Z_i - \frac{1}{n} \sum_{i=1}^n \mathbf{E}[Z_i]\right| \geq \epsilon\right) \leq 2 \exp\left(-\frac{2n\epsilon^2}{B^2}\right).$$

**Lemma A.4** (Bernstein's Inequality for Empirical Variance [3, 23]). *Let  $Z_1, Z_2, \dots, Z_n$  be independent and identically distributed random variables bounded in the interval  $[a, b]$  with length  $B = b - a$  for*

*all  $1 \leq i \leq n$ , and let  $\mu = \mathbf{E}[Z_i]$  be their common expectation. Consider the empirical expectation  $\bar{Z}$  and empirical variance  $\bar{\sigma}^2$  defined respectively by  $\bar{Z} = \frac{1}{n} \sum_{i=1}^n Z_i$  and  $\bar{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n (Z_i - \bar{Z})^2$ . Then we have*

$$\Pr\left(|\bar{Z} - \mathbf{E}[Z]| \leq \sqrt{\frac{2\bar{\sigma}^2 \log(3/p_f)}{n}} + \frac{3B \log(3/p_f)}{n}\right) \geq 1 - p_f.$$

We also need the following bound for the sum of all residues in our proof of correctness of approximation.

**Lemma A.5** (Bound of Sum of Residues). *The following bounds for the sum of residues holds in Algorithm 2:*

$$\sum_{v \in V} \sum_{k=0}^{\ell} \mathbf{r}^{(k)}(v) = 1 - \sum_{v \in V} \mathbf{q}^{(\ell)}(v).$$

Hereafter, we may omit the subscript when the result applies to both  $\mathbf{r}_s^{(k)}(v)$  and  $\mathbf{r}_t^{(k)}(v)$ , or both  $\mathbf{q}_s^{(k)}(v)$  and  $\mathbf{q}_t^{(k)}(v)$ , for simplicity.

**PROOF.** According to the property of transition probability, we have  $1 = \sum_{v \in V} p^{(\ell)}(s, v)$ . Substituting  $p^{(\ell)}(s, v)$  with the invariant in Lemma 2.4, we obtain:

$$1 = \sum_{v \in V} \mathbf{q}_s^{(\ell)}(v) + \sum_{k=0}^{\ell} \sum_{u \in V} \mathbf{r}_s^{(\ell-k)}(u) \sum_{v \in V} p^{(k)}(u, v) = \sum_{v \in V} \mathbf{q}_s^{(\ell)}(v) + \sum_{k=0}^{\ell} \sum_{v \in V} \mathbf{r}_s^{(\ell-k)}(v).$$

Moving the first term to the left side of the equation yields the desired conclusion. A similar conclusion holds for node  $t$ .  $\square$

With all the lemmas in place, we can proceed to prove the correctness of approximation provided by our BiSPER algorithm.

**PROOF.** We will abbreviate  $L_{\max}$  to  $L$  for simplicity in our proof.

The randomness of the estimator  $\hat{R}_L(s, t)$  arises from the  $\hat{T}(s, t)$  part. We first bound the number of sampled random walks ( $N$  in Algorithm 2) using Hoeffding's inequality. We will do this in the following two ways:

- (1) On one hand, note that after the push phase, for any node  $v$ ,  $0 \leq \frac{\mathbf{r}^{(k)}(v)}{d(v)} \leq r_{\max}$  holds. Thus, we have  $0 \leq \sum_{k=0}^{L-\ell} \frac{\mathbf{r}^{(k)}(v)}{d(v)} \leq (L+1-\ell)r_{\max}$ . Therefore, if we denote  $T_{B_1} = (L+1)(L+2)r_{\max}$ , we can conclude that both the first term and second term in  $\hat{T}_i$  range in  $\left[-\frac{T_{B_1}}{2}, \frac{T_{B_1}}{2}\right]$ , and thus  $\hat{T}_i \in [-T_{B_1}, T_{B_1}]$ . According to Hoeffding's inequality, when we set  $N = \left\lceil \frac{2T_{B_1}^2 \log(2/p_f)}{\epsilon^2} \right\rceil$ , we have the conclusion.
- (2) On the other hand, according to Lemma A.5, we have  $0 \leq \sum_{k=0}^{L-\ell} \sum_{v \in V} \frac{\mathbf{r}^{(k)}(v)}{d(v)} \leq \sum_{k=0}^{L-\ell} \sum_{v \in V} \mathbf{r}^{(k)}(v) = 1 - \sum_{v \in V} \mathbf{q}^{(L-\ell)}(v)$ . Therefore, it leads to

$$\hat{T}_i(s, t) \leq \sum_{\ell=0}^L \sum_{k=0}^{L-\ell} \sum_{v \in V} \left( \frac{\mathbf{r}_s^{(k)}(v)}{d(v)} + \frac{\mathbf{r}_t^{(k)}(v)}{d(v)} \right) \leq 2(L+1) - \sum_{\ell=0}^L \sum_{v \in V} \left( \mathbf{q}_s^{(\ell)}(v) + \mathbf{q}_t^{(\ell)}(v) \right) =: T_{B_2},$$

and analogously,  $\hat{T}_i(s, t) \geq -T_{B_2}$ , where  $T_{B_2} := 2(L+1) - \sum_{\ell=0}^L \sum_{v \in V} (\mathbf{q}_s^{(\ell)}(v) + \mathbf{q}_t^{(\ell)}(v))$ . Using Hoeffding's inequality again, when we set  $N = \left\lceil \frac{2T_{B_2}^2 \log(2/p_f)}{\epsilon^2} \right\rceil$ , we have the conclusion.

Another special case occurs when  $r_{\max} \geq \frac{1}{d}$ , where  $d = \min\{d(s), d(t)\}$ . In this scenario, no push operation is performed, and the only two non-zero residue values are  $\mathbf{r}_s^{(0)}(s)$  and  $\mathbf{r}_t^{(0)}(t)$ . Consequently, we have  $0 \leq \sum_{k=0}^{L-\ell} \frac{\mathbf{r}_s^{(k)}(v)}{d(v)} \leq \frac{1}{d}$  and  $0 \leq \sum_{k=0}^{L-\ell} \frac{\mathbf{r}_t^{(k)}(v)}{d(v)} \leq \frac{1}{d}$ . This implies  $\hat{T}_i \in \left[-\frac{2(L+1)}{d}, \frac{2(L+1)}{d}\right]$ . Once again, utilizing Hoeffding's inequality, we conclude that when  $r_{\max} \geq \frac{1}{d}$ , it suffices to set the number of random walks to be  $N = \left\lceil \frac{8(L+1)^2 \log(2/p_f)}{\epsilon^2 d^2} \right\rceil$ .

We observe that the upper bound  $N$  on the number of random walks to be sampled often significantly exceeds what is actually required. To address this, we incorporate Bernstein's inequality (Lemma A.4) in Lines 5–8 of Algorithm 4 for early stopping, following the approach proposed by [33].  $\square$

## B Proof of Theorem 4.2

Theorem 4.2 characterizes the time complexity of our BiSPER algorithm. Before we analyze the time complexity of the BiSPER algorithm, we will prove the following lemma.

**Lemma B.1** (Upper Bound of  $\mathbf{r}_s^{(\ell)}(u)$ ). *For any  $0 \leq \ell \leq L_{\max}$ , the inequalities  $\mathbf{r}_s^{(\ell)}(u) \leq p^{(\ell)}(s, u)$  and  $\mathbf{r}_t^{(\ell)}(u) \leq p^{(\ell)}(t, u)$  hold.*

PROOF. According to Lemma 2.4, we have

$$p^{(\ell)}(s, u) \geq \sum_{v \in V} \mathbf{r}_s^{(\ell)}(v) p^{(0)}(v, u) = \mathbf{r}_s^{(\ell)}(u).$$

Similar results can be derived for  $\mathbf{r}_t^{(\ell)}(u)$ .  $\square$

Now we can present the proof of Theorem 4.2.

PROOF. We will abbreviate  $L_{\max}$  to  $L$  for simplicity in our proof.

We first bound the cost of the push phase. On one hand, the cost of the push phase can be bounded as

$$\begin{aligned} \text{Cost}_{\text{Push}} &= \sum_{\ell=0}^L \sum_{u \in V} \mathbf{I} \left[ \frac{\tilde{\mathbf{r}}^{(\ell)}(u)}{d(u)} > r_{\max} \right] d(u) \log(L+1) \\ &\leq \log(L+1) \sum_{\ell=0}^L \sum_{u \in V} \frac{\tilde{\mathbf{r}}^{(\ell)}(u)}{r_{\max} d(u)} d(u) \\ &\leq \frac{\log(L+1)}{r_{\max}} \sum_{\ell=0}^L \sum_{u \in V} p^{(\ell)}(s, u) = \frac{(L+1) \log(L+1)}{r_{\max}}. \end{aligned}$$

where  $\tilde{\mathbf{r}}^{(\ell)}(u)$  is the residue before  $u$  is pushed. The inequality holds because  $\frac{x}{a}$  is an upper bound for  $\mathbf{I}[x \geq a]$ , and  $p^{(\ell)}(s, u)$  is an upper bound for  $\tilde{\mathbf{r}}^{(\ell)}(u)$  (Lemma B.1). On the other hand, the

cost can also be bounded by

$$\begin{aligned} \text{Cost}_{\text{Push}} &= \sum_{\ell=0}^L \sum_{u \in V} \mathbf{I} \left[ \frac{\tilde{\mathbf{r}}^{(\ell)}(u)}{d(u)} \geq r_{\max} \right] d(u) \log(L+1) \\ &\leq \log(L+1) \sum_{\ell=0}^L \sum_{u \in V} d(u) = 2m(L+1) \log(L+1). \end{aligned}$$

The inequality holds because 1 is a naive upper bound for  $\mathbf{I}[x \geq a]$ . We also note that no push operation will be performed when  $r_{\max} \geq \frac{1}{d} = \frac{1}{\min\{d(s), d(t)\}}$ . Therefore, the push cost can be summarized as

$$\text{Cost}_{\text{Push}} = \begin{cases} 0, & r_{\max} \geq \frac{1}{d} \\ \frac{(L+1) \log(L+1)}{r_{\max}}, & \frac{1}{2m} \leq r_{\max} < \frac{1}{d} \\ 2m(L+1) \log(L+1), & r_{\max} < \frac{1}{2m} \end{cases}.$$

Meanwhile, the cost of the adaptive Monte Carlo phase can be easily bounded by

$$\text{Cost}_{\text{AMC}} = \begin{cases} \frac{8(L+1)^3 \log(L+1) \log(2/p_f)}{\epsilon^2 d^2}, & r_{\max} \geq \frac{1}{d} \\ c = \frac{2(L+1)^3 (L+2)^2 r_{\max}^2 \log(L+1) \log(2/p_f)}{\epsilon^2}, & r_{\max} < \frac{1}{d} \end{cases}.$$

Considering all combinations, we get the total cost

$$\text{Cost}_{\text{Total}} = \begin{cases} \frac{8(L+1)^3 \log(L+1) \log(2/p_f)}{\epsilon^2 d^2}, & r_{\max} \geq \frac{1}{d} \\ \frac{(L+1) \log(L+1)}{r_{\max}} + c, & \frac{1}{2m} \leq r_{\max} < \frac{1}{d} \\ 2m(L+1) \log(L+1) + c, & r_{\max} < \frac{1}{2m} \end{cases}.$$

Then we need to set the threshold  $r_{\max}$  to minimize the total cost. In the first case, the total cost is irrelevant to  $r_{\max}$ . In the second case, the total cost can be expressed as  $\frac{A}{r_{\max}} + Br_{\max}^2$ . To minimize the upper bound of the total cost, we set  $r_{\max} = (\frac{A}{2B})^{1/3} = \frac{\epsilon^{2/3}}{2^{2/3}(L+1)^{2/3}(L+2)^{2/3} \log^{1/3}(2/p_f)} = \tilde{O}\left(\frac{\epsilon^{2/3}}{L^{4/3}}\right)$ . By doing so, the minimized upper bound of the total cost is  $\frac{3}{2^{2/3}} \cdot A^{2/3} B^{1/3} = \tilde{O}\left(\frac{L^{7/3}}{\epsilon^{2/3}}\right)$ . In the third case, the cost is minimized to  $2m(L+1) \log(L+1) = \tilde{O}(mL)$  if we set  $r_{\max} = 0$ . Therefore, we have the following observations:

- (1) When the maximum random walk length  $L$  is large enough, or  $L \geq \max\left\{\frac{m^{1/2} \epsilon d}{2 \log^{1/2}(2/p_f)}, \frac{2m^{3/4} \epsilon^{1/2}}{3^{3/4} \log^{1/4}(2/p_f)}\right\}$ , the third bound of the total cost is the minimum, we can set  $r_{\max} = 0$ .
- (2) When  $d = \min\{d(s), d(t)\}$  is large enough, or  $d \geq \max\left\{\frac{2^{5/3}(L+1)^{1/3} \log^{1/3}(2/p_f)}{3^{1/2} \epsilon^{2/3}}, \frac{2(L+1) \log^{1/2}(2/p_f)}{m^{1/2} \epsilon}\right\}$ , the first bound of the total cost is the minimum, we can set  $r_{\max} = \frac{1}{d}$ .
- (3) In other cases, we set  $r_{\max} = \frac{\epsilon^{2/3}}{2^{2/3}(L+1)^{2/3}(L+2)^{2/3} \log^{1/3}(2/p_f)} = \tilde{O}\left(\frac{\epsilon^{2/3}}{L^{4/3}}\right)$ .

With the above setting of  $r_{\max}$ , we can get the bound of the total cost  $\tilde{O}\left(\min\left\{\frac{L^3}{\epsilon^2 d^2}, \frac{L^{7/3}}{\epsilon^{2/3}}, mL\right\}\right)$ .  $\square$