

ExactSim: Benchmarking Single-Source SimRank Algorithms with High-Precision Ground Truths

Hanzhi Wang · Zhewei Wei · Yu Liu · Ye Yuan · Xiaoyong Du ·
Ji-Rong Wen

Received: date / Accepted: date

Abstract *SimRank* is a popular measurement for evaluating the node-to-node similarities based on the graph topology. In recent years, single-source and top- k SimRank queries have received increasing attention due to their applications in web mining, social network analysis, and spam detection. However, a fundamental obstacle in studying SimRank has been the lack of ground truths. The only exact algorithm, Power Method, is computationally infeasible on graphs with more than 10^6 nodes. Consequently, no existing work has evaluated the actual accuracy of various single-source and top- k SimRank algorithms on large real-world graphs.

In this paper, we present ExactSim, the first algorithm that computes the exact single-source and top- k

Zhewei Wei is the corresponding author.

Hanzhi Wang
School of Information, Renmin University of China, China
E-mail: hanzhi.wang@ruc.edu.cn

Zhewei Wei
Gaoling School of Artificial Intelligence, Renmin University of China, China
E-mail: zhewei@ruc.edu.cn

Yu Liu
Wangxuan Institute of Computer Technology, Peking University, China
E-mail: dokiliu@pku.edu.cn

Ye Yuan
School of Computer Science and technology, Beijing Institute of Technology, China
E-mail: yuan-ye@bit.edu.cn

Xiaoyong Du
MOE Key Lab DEKE, Renmin University of China, China
E-mail: duyong@ruc.edu.cn

Ji-Rong Wen
Beijing Key Lab of Big Data Management and Analysis Method, Renmin University of China, China
E-mail: jrwen@ruc.edu.cn

SimRank results on large graphs. This algorithm produces ground truths with precision up to 7 decimal places with high probability. With the ground truths computed by ExactSim, we present the first experimental study of the accuracy/cost trade-offs of existing approximate SimRank algorithms on large real-world graphs and synthetic graphs. Finally, we use the ground truths to exploit various properties of SimRank distributions on large graphs.

Keywords SimRank, Single-Source, Exact computation, Ground truths, Power-Law, Benchmarks

1 Introduction

Computing link-based similarity is an overarching problem in graph analysis and mining. Amid the existing similarity measures [31, 40, 49, 48], SimRank has emerged as a popular metric for assessing structural similarities between nodes in a graph. SimRank was introduced by Jeh and Widom [13] to formalize the intuition that “two pages are similar if they are referenced by similar pages.” Given a directed graph $G = (V, E)$ with n nodes $\{v_1, \dots, v_n\}$ and m edges, the SimRank matrix S defines the similarity between any two nodes v_i and v_j as follows:

$$S(i, j) = \begin{cases} 1, & \text{for } i = j; \\ \sum_{v_{i'} \in \mathcal{I}(v_i)} \sum_{v_{j'} \in \mathcal{I}(v_j)} \frac{c \cdot S(i', j')}{d_{in}(v_i) \cdot d_{in}(v_j)}, & \text{for } i \neq j. \end{cases} \quad (1)$$

Here, c is a decay factor typically set to 0.6 or 0.8 [13, 26]. $\mathcal{I}(v_i)$ denotes the set of in-neighbors of v_i , and $d_{in}(v_i)$ denotes the in-degree of v_i . SimRank aggregates similarities of multi-hop neighbors of v_i and v_j to

produce high-quality similarity measure, and has been adopted in various applications such as recommendation systems [20], link prediction [27], and graph embeddings [35].

A fundamental obstacle for studying SimRank is the lack of ground truths on large graphs. Currently, the only methods that compute the SimRank matrix is Power Method and its variations [13, 25], which inherently takes $\Omega(n^2)$ space and at least $\Omega(n^2)$ time as there are $\Omega(n^2)$ node pairs in the graphs. This complexity is infeasible on large graphs ($n \geq 10^6$). Consequently, the majority of recent works [16, 29, 34, 10, 17, 21, 32, 44, 14, 24, 39] focus on *single-source and top- k queries*. Given a source node v_i , a single-source query asks for the SimRank similarity between every node and v_i , and a top- k query asks for the k nodes with the highest SimRank similarities to v_i . Unfortunately, computing ground truths for the single-source and top- k queries on large graphs still remains an open problem. To the best of our knowledge, Power Method is still the only way to obtain exact single-source and top- k results, which is not feasible on large graphs. Due to the hardness of exact computation, existing works on single-source and top- k queries focus on approximate computations with efficiency and accuracy guarantees.

The lack of ground truths has severely limited our understanding towards SimRank and SimRank algorithms. First of all, designing approximate algorithms without the ground truths is like shooting in the dark. Most existing works take the following approach: they evaluate the accuracy on small graphs where the ground truths can be obtained by the Power Method with $\Omega(n^2)$ space complexity. Then they report the efficiency/ scalability results on large graphs with consistent parameters. This approach is flawed for the reason that consistent parameters may still lead to unfair comparisons. For example, some of the existing methods generate a fixed number of random walks from each node, while others fix the maximum error ε and generate $\frac{\log n}{\varepsilon^2}$ random walks from each node. If we increase the graph size n , the comparison becomes unfair as the latter methods require more random walks from each node. Secondly, it is known that the structure of large real-world graphs can be very different from that of small graphs. Consequently, the accuracy results on small graphs can only serve as a rough guideline for accessing the actual error of the algorithms in real-world applications. We believe that the only right way to evaluate the effectiveness of a SimRank algorithm is to evaluate its results against the ground truths on large real-world graphs.

Second, the lack of ground truths has also prevented us from exploiting the distribution of SimRank on real-world graphs. For example, it is known [7] that the

PageRank of most real-world graphs follows the power-law distribution. The natural question is that, does SimRank also follow the power-law distribution on real-world graphs? Furthermore, the performances of some existing methods [38] depend on the *density* of the SimRank, which is defined as the percentage of node pairs with SimRank similarities larger than some threshold ε . Analyzing the distribution or density of SimRank is clearly infeasible without the ground truths.

Finally, the lack of ground truths restricts us to conduct scientific benchmarking experiments towards these existing approximation algorithms. Without insightful experimental observations, we are hard to explore the connections between algorithms' characteristics and performances. For example, what kinds of algorithms tend to show better scalabilities? Algorithms belonging to which categories can perform better trade-off lines? A comprehensive benchmarking survey is fundamentally based on the ground truths.

Exact Single-Source SimRank Computation. In this paper, we study the problem of computing the exact single-source SimRank results on large graphs. A key insight is that exactness does not imply absolutely zero error. This is because SimRank values may be infinite decimals, and we can only store these values with finite precision. Moreover, we note that the ground truths computed by Power Method also incur an error of at most c^L , where L is the number of iterations in Power Method. In most applications, L is set to be large enough such that c^L is smaller than the numerical error and thus can be ignored. In this paper, we aim to develop an algorithm that answers single-source SimRank queries with an additive error of at most $\varepsilon_{min} = 10^{-7}$. Note that the float type in various programming languages usually support precision of up to 6 or 7 decimal places. So by setting $\varepsilon_{min} = 10^{-7}$, we guarantee the algorithm returns the same answers as the ground truths in the float type. As we shall see, this precision is extremely challenging for existing methods. To make the exact computation possible, we are also going to allow a small probability to fail. We define the probabilistic exact single-source SimRank algorithm as follows.

Definition 1 With probability at least $1 - 1/n$, for every source node $v_i \in V$, a probabilistic exact single-source SimRank algorithm answers the single-source SimRank query of v_i with additive error of at most $\varepsilon_{min} = 10^{-7}$.

Our Contributions. In this paper, we propose ExactSim, the first algorithm that enables probabilistic exact single-source SimRank queries on large graphs. We show that existing single-source methods share a

common complexity term $O\left(\frac{n \log n}{\varepsilon_{\min}^2}\right)$, and thus are unable to achieve exactness on large graphs. However, ExactSim runs in $O\left(\frac{\log n}{\varepsilon_{\min}^2} + m \log \frac{1}{\varepsilon_{\min}}\right)$ time, which is feasible for both large graph size m and small error guarantee ε_{\min} . We also apply several non-trivial optimization techniques to reduce the query cost and space overhead of ExactSim. In our empirical study, we show that ExactSim is able to compute the ground truth with a precision of up to 7 decimal places within one hour on graphs with billions of edges. Hence, we believe ExactSim is an effective tool for producing the ground truths for single-source SimRank queries on large graphs.

Comparison with the conference version[36]. We make the following new contributions over the conference version.

- We conduct a comprehensive survey on all single-source SimRank algorithms which can support large graphs. We summarize the complexity of each method and analyze the reasons why these methods cannot achieve exactness on large graphs.
- Based on the ground truths provided by ExactSim, we conduct the first empirical study on the accuracy/cost trade-offs of existing approximate single-source algorithms on large real-world graphs and synthetic graphs.
- We use ExactSim to exploit various properties of SimRank on large real-world graphs. In particular, we show that the single-source SimRank values follow the power-law distribution on real-world graphs. We also study the density of SimRank values on large graphs.

2 Preliminaries and Related Work

In this section, we review the state-of-the-art single-source SimRank algorithms which can support large graphs. We introduce a taxonomy to classify these algorithms into three categories: Monte Carlo methods, iterative methods, and local push/sampling methods. Note that our ExactSim algorithm is largely inspired by three prior works: Linearization [29], PRSim [39] and pooling [24], and we will describe them in details. In Section 5, we will also use the ground truths provided by ExactSim to evaluate the algorithms mentioned in this section. Table 1 summarizes the notations used in this paper.

2.1 Monte Carlo Methods

A popular interpretation of SimRank is the *meeting probability* of random walks. In particular, we consider a

Table 1 Table of notations.

Notation	Description
n, m	the numbers of nodes and edges in G
$\mathcal{I}(v_i), \mathcal{O}(v_i)$	the in/out-neighbor set of node v_i
$S, S(i, j)$	the SimRank matrix and the SimRank similarity of v_i and v_j
c	the decay factor in the definition of SimRank
$\varepsilon, \varepsilon_{\min}$	additive error parameter and error required for exactness ($\varepsilon_{\min} = 10^{-7}$)
P, D	the transition matrix and the diagonal correction matrix
$\vec{\pi}_i, \vec{\pi}_i^\ell$,	the Personalized PageRank and ℓ -hop Personalized PageRank vectors of node v_i
\vec{h}_i^ℓ	the ℓ -hop Hitting Probability vector of v_i

random walk from node u that, at each step, moves to a random *in-neighbor* with probability \sqrt{c} , and stops at the current node with probability $1 - \sqrt{c}$. Such a random walk is called a \sqrt{c} -walk. Suppose we start a \sqrt{c} -walk from node v_i and a \sqrt{c} -walk from node v_j , we call the two \sqrt{c} -walks *meet* if they visit the same node at the same step. It is known [34] that

$$S(i, j) = \Pr[\text{two } \sqrt{c}\text{-walks from } v_i \text{ and } v_j \text{ meet}]. \quad (2)$$

According to Equation (2), we can employ Monte-Carlo sampling to estimate $S(i, j)$. That is, by simulating adequate pairs of \sqrt{c} -walks from nodes v_i, v_j , the percentage of the walks that meet in the walking process serves as the estimator of $S(i, j)$. Hence, we classify the approximation algorithms as Monte Carlo methods if they use the fraction of target random walks to estimate the meeting probability based on Equation (2) or its variants.

MC [9] makes use of Equation (2) to derive a Monte-Carlo algorithm for computing single-source SimRank. In the preprocessing phase, we simulate $R \sqrt{c}$ -walks from each node in V . Given a source node v_i , we compare the \sqrt{c} -walks from v_i and from each node $v_j \in V$, and use the fraction of \sqrt{c} -walks that meet as an estimator for $S(i, j)$. By standard concentration inequalities, the maximum error of estimated $S(i, j)$ is bounded by ε with high probability if we set $R = O\left(\frac{\log n}{\varepsilon^2}\right)$, leading to a preprocessing time of $O\left(\frac{n \log n}{\varepsilon^2}\right)$.

READS [14] is an optimized version of the MC-based algorithm. The key idea is to build an index of nR compressed \sqrt{c} -walks such that the algorithm only needs to generate a few more \sqrt{c} -walks in the query phase. An appealing feature of READS is that its index support efficient insertions and deletions of edges. Consequently,

Table 2 Comparison of MC-based SimRank algorithms.

Algorithm	Query Time	Preprocessing time	Index Size	Dynamic update time
MC [9]	$O(n \log n / \varepsilon^2)$	$O(n \log n / \varepsilon^2)$	$O(n \log n / \varepsilon^2)$	-
READS [14]	$O(n \log n / \varepsilon^2)$	$O(n \log n / \varepsilon^2)$	$O(n \log n / \varepsilon^2)$	$O(\log n / \varepsilon^2)$
TSF [32]	$O(n \log n / \varepsilon^2)$	$O(n \log n / \varepsilon^2)$	$O(n \log n / \varepsilon^2)$	$O(\log n / \varepsilon^2)$
Uniwalk [28]	$O(n^2 \log n / \varepsilon^2)$	0	0	$O(n^2 \log n / \varepsilon^2)$

READS is able to support approximate single-source queries on large dynamic graphs. The theoretical query cost of READS remains $O\left(\frac{n \log n}{\varepsilon^2}\right)$.

TSF [32] is a MC-based algorithm for single-source and top- k SimRank queries on both static and dynamic graphs. TSF builds an index that consists of R_g *one-way graphs*, each of which contains the coupling of random walks of length T from each node. In the query phase, TSF samples R_q more random walks for each one-way graph to provide the final estimators. TSF allows two random walks to meet multiple times and assumes that there is no cycle with a length shorter than T , leading to a lower precision in practice. The query time of TSF is bounded by $O(nR_gR_q)$, which is in turn bounded by $O\left(\frac{n \log n}{\varepsilon^2}\right)$ for ε additive error.

Uniwalk [28] is a MC-based method for single-source and top- k SimRank computation on undirected graphs. It randomly generates R unidirectional random walks from the given source node s . With the help of a rectified factor, Uniwalk regards the probability of the node s walking along the unidirectional path to the terminal node t as the SimRank value $S(s, t)$, that is, two random walks starting from s and t meet at the midpoint of the original unidirectional path. The query time of Uniwalk is bounded by $O(RL)$, where L denotes the expected length of the unidirectional path. However, the rectified factor can influence the error bound. On the graph with a hub node, R can reach $O\left(\frac{n^2 \log n}{\varepsilon^2}\right)$ for ε additive error. Hence, the query time of Uniwalk can be bounded by $O\left(\frac{n^2 \log n}{\varepsilon^2}\right)$.

2.2 Iterative Methods

Given a graph $G = (V, E)$, let P denote the (reverse) *transition matrix*, that is, $P(i, j) = 1/d_{in}(v_j)$ for $v_i \in \mathcal{I}(v_j)$, and $P(i, j) = 0$ otherwise. S denotes the SimRank matrix. Yu et al. [47] proves that the definition formula of SimRank can be expressed as

$$S = (cP^\top SP) \vee I, \quad (3)$$

where I denotes an $n \times n$ identity matrix, and \vee is an element-wise maximum operator that for any matrices $A, B \in \mathcal{R}^{n \times n}$ and $\forall i, j \in \{0, 1, \dots, n-1\}$, $(A \vee B)(i, j) = \max\{A(i, j), B(i, j)\}$. Equation (3) provides an iterative calculation method to derive SimRank results. That is, we can initialize $S = I$, and repeat the iteration to update matrix S . We classify all the SimRank algorithms as iterative methods if they calculate SimRank values via iterative updating based on Equation (3) or its variants. We list all the iterative methods which can support single-source SimRank queries on large graphs in the following.

Linearization and ParSim. It is shown in two independent works, Linearization [29] and ParSim [45], that the iterative definition equation (3) can be expressed as the following linear summation:

$$S = cP^\top SP + D = \sum_{\ell=0}^{+\infty} c^\ell (P^\ell)^\top DP^\ell, \quad (4)$$

where D is the *diagonal correction matrix* with each diagonal element $D(k, k)$ taking value from $1 - c$ to 1. Consequently, a single-source query for node v_i can be computed by

$$S \cdot \vec{e}_i = \sum_{\ell=0}^{+\infty} c^\ell (P^\ell)^\top DP^\ell \cdot \vec{e}_i, \quad (5)$$

where \vec{e}_i denotes the one-hot vector with the i -th element being 1 and all other elements being 0. Assuming the diagonal matrix D is correctly given, the single-source query for node v_i can be approximated by

$$S_L \cdot \vec{e}_i = \sum_{\ell=0}^L c^\ell (P^\ell)^\top DP^\ell \cdot \vec{e}_i, \quad (6)$$

where L is the number of iterations. After L iterations, the additive error reduces to c^L . So setting $L = O(\log \frac{1}{\varepsilon})$ is sufficient to guarantee a maximum error of ε . At the ℓ -th iterations, the algorithm performs $2\ell + 1$ matrix-vector multiplications to calculate $c^\ell (P^\ell)^\top DP^\ell \cdot \vec{e}_i$, and each matrix-vector multiplication takes $O(m)$ time. Consequently, the total query time is bounded by

Table 3 Comparison of iterative SimRank algorithms.

Algorithm	Query Time	Preprocessing Time	Index Size	Dynamic update time
Linearization [29]	$O(m \log^2 \frac{1}{\varepsilon})$	$O(n \log \frac{1}{\varepsilon} \log \frac{n}{\varepsilon} \log n / \varepsilon^2)$	$O(n)$	-
ParSim [45]	$O\{\min\{m \log \frac{1}{\varepsilon}, d^{2 \log \frac{1}{\varepsilon}}\}\}$	0	0	-

$O\left(\sum_{\ell=1}^L m\ell\right) = O(mL^2) = O(m \log^2 \frac{1}{\varepsilon})$. Maehara et al. and Yu et al. also show in [29] and [45] that if we first compute and store the transition probability vectors $\vec{u}_\ell = P^\ell \cdot \vec{e}_i$ for $\ell = 0, \dots, L$, then we can use the following equation to compute

$$\begin{aligned} S_L \cdot \vec{e}_i &= D \cdot \vec{u}_0 \\ &+ cP^\top (D \cdot \vec{u}_1 + \dots + cP^\top (D \cdot \vec{u}_{T-1} + cP^\top \cdot D \cdot \vec{u}_T) \dots). \end{aligned} \quad (7)$$

This optimization reduces the query time to $O(m \log \frac{1}{\varepsilon})$. While it requires a memory size of $O(nL) = O(n \log \frac{1}{\varepsilon})$, which is usually several times larger than the graph size m . Therefore, [29] only uses the $O(m \log^2 \frac{1}{\varepsilon})$ algorithm in the experiments.

Besides the large space overhead, another problem with Linearization and ParSim is that the diagonal correction matrix D is hard to compute. Linearization [29] formulates D as the solution to a linear system, and proposes a Monte Carlo solution that takes $O\left(\frac{n \log n}{\varepsilon^2}\right)$ to derive an ε -approximation of D . On the other hand, ParSim directly sets $D = (1-c)I$, where I is the identity matrix. This approximation basically ignores the first meeting constraint and has been adopted in many other SimRank works [11, 12, 42, 19, 41, 44, 16]. It is shown that the similarities calculated by this approximation are different from the actual SimRank [16]. However, the quality of this approximation is still a myth due to the lack of ground truths on large graphs.

2.3 Local Push/Sampling Methods

Compared with Monte Carlo and iterative methods, local push/sampling methods locally restrict each SimRank update operation and omit to touch a large fraction of nodes on the graphs in each update. Hence, the time cost of each update operation is smaller than $O(n)$. This allows local push/sampling methods to outperform other methods in terms of scalability.

SLING [34] is an index-based SimRank algorithm that supports fast single-source and top- k queries on static graphs. Let $\vec{h}_i^\ell = (\sqrt{c}P)^\ell \cdot \vec{e}_i$ denote the ℓ -hop hitting

probability vector of v_i . Note that \vec{h}_i^ℓ describes the probability of an \sqrt{c} -walk from node v_i visiting each node at its ℓ -th step. [34] suggests that equation (5) can be re-written as

$$S(i, j) = \sum_{\ell=0}^{\infty} \sum_{k=1}^n \vec{h}_i^\ell(k) \cdot \vec{h}_j^\ell(k) \cdot D(k, k). \quad (8)$$

where $D(k, k)$ denote the k -th entry in the diagonal correction matrix D . It is shown [34] that $D(k, k)$ can be characterized by the meeting probability of two \sqrt{c} -walks from the same node v_k :

$$D(k, k) = \Pr[\text{two } \sqrt{c}\text{-walks from } v_k \text{ never meet}]. \quad (9)$$

This interpretation implies a simple Monte-Carlo algorithm for estimating $D(k, k)$: we simulate R pairs of \sqrt{c} -walks from v_k , and use the fraction of pairs that do not meet as the estimator for $D(k, k)$. By setting $R = O\left(\frac{\log n}{\varepsilon^2}\right)$, we can approximate each $D(k, k)$ with additive error ε . SLING precomputes each $D(k, k)$ in the preprocessing phase using $O\left(\frac{n \log n}{\varepsilon^2}\right)$ time. SLING also precomputes $\vec{h}_i^\ell(k)$ with additive error ε for each ℓ and $v_i, v_k \in V$, using a *local push* algorithm [5]. Given a single-source query for node v_i , SLING retrieves $\vec{h}_i^\ell(k)$ and $D(k, k)$ for each $v_j, v_k \in V$ from the index, and uses Equation (8) to estimate $S(i, j)$ for each $v_j \in V$. SLING answers a single-source query with time $O(\min\{n/\varepsilon, m\})$, and the index size is bounded by $O\left(\frac{n}{\varepsilon}\right)$.

ProbeSim [24] is an index-free solution based on reverse local sampling and local push. ProbeSim starts by sampling a \sqrt{c} -walk from the source node v_i . For the ℓ -th node v_k on the \sqrt{c} -walk, ProbeSim uses a *Probe algorithm* to reversely sample each node v_j at level ℓ with probability $\vec{h}_j^\ell(k)$, the hitting probability that any other node $v_j \in V$ can reach v_k at the ℓ -th step. It is shown in [24] that each sample takes $O(n)$ time, and we need $O\left(\frac{\log n}{\varepsilon^2}\right)$ samples to ensure an maximum error of ε with high probability. Consequently, the query time of ProbeSim is bounded by $O\left(\frac{n \log n}{\varepsilon^2}\right)$. ProbeSim naturally works on dynamic graphs due to its index-free nature.

Table 4 Comparison of local push/sampling SimRank algorithms.

Algorithm	Query Time	Preprocessing time	Index Size	Dynamic update time
SLING [34]	$O(n/\varepsilon)$	$O\left(\frac{m}{\varepsilon} + \frac{n \log \frac{n}{\delta}}{\varepsilon^2}\right)$	$O(n/\varepsilon)$	-
	$O(m \log^2 \frac{1}{\varepsilon})$			
ProbeSim [24]	$O(n \log n/\varepsilon^2)$	0	0	$O(n \log n/\varepsilon^2)$
PRSim [39]	$O(n \log n \cdot \ \vec{\pi}_i\ ^2/\varepsilon^2)$	$O(m/\varepsilon)$	$O(\min\{\frac{n}{\varepsilon}, m\})$	-
TopSim [17]	$O(m^{2n}/n^{2n})$	0	0	$O(m^{2n}/n^{2n})$

PRSim [39] introduces a partial indexing and a probe algorithm. Let $\vec{\pi}_i^\ell = (1 - \sqrt{c})\vec{h}_i^\ell = (1 - \sqrt{c})(\sqrt{c}P)^\ell \cdot \vec{e}_i$ denote the ℓ -hop *Personalize PageRank vector* of v_i . In particular, $\vec{\pi}_i^\ell(k)$ is the probability that a \sqrt{c} -walk from node v_i stops at node v_k in exactly ℓ steps. PRSim suggests that equation (5) can be re-written as

$$S(i, j) = \frac{1}{(1 - \sqrt{c})^2} \sum_{\ell=0}^{\infty} \sum_{k=1}^n \vec{\pi}_i^\ell(k) \cdot \vec{\pi}_j^\ell(k) \cdot D(k, k). \quad (10)$$

PRSim precomputes $\vec{\pi}_j^\ell(k)$ with additive error ε for each ℓ and $v_j, v_k \in V$, using a *local push* algorithm [5]. To avoid overwhelming index size, PRSim only precomputes $\vec{\pi}_j^\ell(k)$ for a small subset of v_k . Furthermore, PRSim computes D by estimating the product $\vec{\pi}_i^\ell(k) \cdot D(k, k)$ together with an $O\left(\frac{\log n}{\varepsilon^2}\right)$ time Monte-Carlo algorithm. Finally, PRSim proposes a new Probe algorithm that samples each node v_j according to $\vec{\pi}_j^\ell(k)$. The average query time of PRSim is bounded by $O\left(\frac{n \log n}{\varepsilon^2} \cdot \sum_{k=1}^n \vec{\pi}(k)^2\right)$, where $\vec{\pi}(k)$ denotes the PageRank of v_k . It is well-known that on scale-free networks, the PageRank vector $\vec{\pi}$ follows the power-law distribution, and thus $\|\vec{\pi}\|^2 = \sum_{k=1}^n \vec{\pi}(k)^2$ is a value much smaller than 1. However, for worst-case graphs or even some “bad” source nodes on scale-free networks, the running time of PRSim remains $O\left(\frac{n \log n}{\varepsilon^2}\right)$.

TopSim [17] is an index-free algorithm based on local exploitation. Given source node v_i , TopSim firstly finds all nodes v_k reachable from v_i within $\ell = 1, \dots, L$ steps. For each such v_k on the ℓ -th level, TopSim deterministically computes $\vec{h}_j^\ell(k)$, the probability that each v_j reaches v_k in exactly ℓ steps. [17] also proposes various optimizations to reduce the query cost. Due to the dense structures of real-world networks, TopSim is only able to exploit a few levels on large graphs, which leads to a low precision.

2.4 Other Related Work

Besides the state-of-the-art methods that we discuss above, there are several other techniques for SimRank

computation, which we review in the following. *Power method* [13] is the classic algorithm that computes all-pair SimRank similarities for a given graph. Power method recursively computes the SimRank Matrix S based on Equation (3). Several follow-up works [26, 47, 43] improve the efficiency or effectiveness of the power method in terms of either efficiency or accuracy. However, these methods still incur $O(n^2)$ space overheads, as there are $O(n^2)$ pairs of nodes in the graph. Finally, there are existing works on *SimRank similarity join* [33, 30, 51] and the variants of SimRank [6, 9, 22, 46, 50], but the proposed solutions are inapplicable for top- k and single-source SimRank queries.

Pooling. Finally, pooling [24] is an experimental method for evaluating the accuracy of top- k SimRank algorithms without the ground truths. Suppose the goal is to compare the accuracy of top- k queries for z algorithms A_1, \dots, A_z . Given a query node v_i , we retrieve the top- k nodes returned by each algorithm, remove the duplicates, and merge them into a pool. Note that there are at most ℓk nodes in the pool. Then we estimate $S(i, j)$ for each node v_j in the pool using the Monte Carlo algorithm. We set the number of random walks to be $O\left(\frac{\log n}{\varepsilon_{\min}^2}\right)$ so that we can obtain the ground truth of $S(i, j)$ with high probability. After that, we take the k nodes with the highest SimRank similarity to v_i from the pool as the ground truth of the top- k query, and use this “ground truth” to evaluate the precision of each of the ℓ algorithms. Note that the set of these k nodes is not the actual ground truth. However, it represents the best possible k nodes that can be found by the ℓ algorithms that participate in the pool and thus can be used to compare the quality of these algorithms.

Although pooling is proved to be effective in our scenario where ground truths are hard to obtain, it has some drawbacks. First of all, the precision results obtained by pooling are *relative* and thus cannot be used outside the pool. This is because the top- k nodes from the pool is not the actual ground truth. Consequently, an algorithm that achieves 100% precision in the pool may have a precision of 0% when compared to the ac-

tual top- k result. Secondly, the complexity of pooling z algorithms is $O\left(\frac{kz \log n}{\varepsilon_{\min}^2}\right)$, so pooling is only feasible for evaluating top- k queries with small k . In particular, we cannot use pooling to evaluate the single-sources queries on large graphs.

2.5 Limitations of Existing Methods

We now analyze the reasons why existing methods are unable to achieve exactness (a.k.a an error of at most $\varepsilon_{\min} = 10^{-7}$). First of all, ParSim and TSF ignore the first meeting constraint and thus incur large errors. For other methods that enforce the first meeting constraint, they all incur a complexity term of $O\left(\frac{n \log n}{\varepsilon^2}\right)$, either in the preprocessing phase or in the query phase. In particular, SLING and Linearization simulate $O\left(\frac{n \log n}{\varepsilon^2}\right)$ random walks to estimate the diagonal correction matrix D . For ProbeSim, MC, READS and PRSim, this complexity is caused by simulating random walks in the query phase or the preprocessing phase. The $O\left(\frac{n \log n}{\varepsilon^2}\right)$ complexity is infeasible for exact SimRank computation on large graphs, since it combines two expensive terms n and $\frac{1}{\varepsilon_{\min}^2}$. As an example, we consider the IT dataset used in our experiment, with $4 * 10^7$ nodes and over 1 billion edges. In order to achieve a maximum error of $\varepsilon_{\min} = 10^{-7}$, we need to simulate $\frac{n \log n}{\varepsilon^2} \approx 10^{23}$ random walks. This may take years, even with parallelization on a cluster of thousands of machines.

Besides, there are many works focusing on all-pairs SimRank queries [37, 19, 41, 12, 26]. As we shall show in Section 5.3, the number of node pairs whose SimRank values are more than 10^{-4} can nearly achieve n^2 . For large graphs with million nodes, like Twitter(TW) dataset with 4×10^7 nodes, this can cost 10^4 TB for storage, not to mention the exact SimRank computation for each node pair. Hence, it's may infeasible for exact all-pairs SimRank computation within reasonable time.

3 Basic ExactSim Algorithm

In this section, we present ExactSim, a probabilistic algorithm that computes the exact single-source SimRank results within reasonable running time. We first present a basic version of ExactSim. In Section 4, we will introduce some more advanced techniques to optimize the query and space cost.

Our ExactSim algorithm is largely inspired by three prior works: pooling [24], Linearization [29] and PRSim [39]. We now discuss how ExactSim extends from these

existing methods in details. These discussions will also reveal the high level ideas of the ExactSim algorithm.

1. Despite its limitations, pooling [24] provides a key insight for achieving exactness: while an $O\left(\frac{n \log n}{\varepsilon^2}\right)$ algorithm is not feasible for exact SimRank computation on large graphs, we can actually afford an $O\left(\frac{\log n}{\varepsilon^2}\right)$ algorithm. The $\frac{1}{\varepsilon^2}$ term is still expensive for $\varepsilon = \varepsilon_{\min} = 10^{-7}$, however, the new complexity reduces the dependence on the graph size n to logarithmic, and thus achieves very high scalability.
2. Linearization [29] and ParSim [45] show that if the diagonal correction matrix D is correctly given, then we can compute the exact single-source SimRank results in $O\left(m \log \frac{1}{c} \frac{1}{\varepsilon_{\min}}\right)$ time and $O\left(n \log \frac{1}{c} \frac{1}{\varepsilon_{\min}}\right)$ extra space. For typical setting of c (0.6 to 0.8), the number of iterations $\log \frac{1}{c} \frac{1}{\varepsilon_{\min}} = \log 10^7 \leq 73$ is a constant, so this complexity is essentially the same as that of performing BFS multiple times on the graphs. The scalability of the algorithm is confirmed in the experiments of [45], where D is set to be $(1 - c)I$. Moreover, the exact algorithms [31] for Personalized PageRank and PageRank also incurs a running time of $O\left(m \log \frac{1}{\varepsilon_{\min}}\right)$, and has been widely used for computing ground truths on large graphs.
3. While the $O\left(\frac{n \log n}{\varepsilon^2}\right)$ complexity seems unavoidable as we need to estimate each entry in the diagonal correction matrix D with additive error ε , PRSim [39] shows that it only takes $O\left(\frac{\log n}{\varepsilon^2}\right)$ time to estimate the product $\vec{\pi}_i^\ell(k) \cdot D(k, k)$ with additive error ε for each $k = 1, \dots, n$ and $\ell = 0, \dots, \infty$, where $\vec{\pi}_i^\ell$ is the ℓ -hop Personalized PageRank vector of v_i . This result provides two crucial observations: 1) It is possible to answer an single-source query without an ε -approximation of each $D(k, k)$; 2) The accuracy of each $D(k, k)$ should depend on $\vec{\pi}_i(k)$, the Personalized PageRank of v_k with respect to the source node v_i .

We combine the ideas of PRSim and Linearization/ ParSim to derive the basic ExactSim algorithm. Given an error parameter ε , ExactSim fixes the total number of \sqrt{c} -walk samples to be $R = O\left(\frac{\log n}{\varepsilon^2}\right)$, and distributes a fraction of $R \cdot \vec{\pi}_i(k)$ samples (note that $\sum_{k=1}^n \vec{\pi}_i(k) = 1$) to estimate $D(k, k)$. Then it performs Linearization/ ParSim with the estimated D to obtain the single-source result. The algorithm runs in $O\left(\frac{\log n}{\varepsilon^2} + m \log \frac{1}{\varepsilon}\right)$ time and uses $O\left(n \log \frac{1}{\varepsilon}\right)$ extra space. Since both complexity terms $O\left(\frac{\log n}{\varepsilon^2}\right)$ and $O\left(m \log \frac{1}{\varepsilon}\right)$ are feasible for $\varepsilon_{\min} =$

Algorithm 1: Basic ExactSim Algorithm

Input: Graph G with transition matrix P , source node v_i , maximum error ε
Output: Estimated single-source SimRank vector $S \cdot \vec{e}_i^*$

- 1 $L = \lceil \log_{\frac{1}{c}} \frac{2}{\varepsilon} \rceil$;
- 2 $\vec{\pi}_i^0, \vec{\pi}_i = (1 - \sqrt{c})\vec{e}_i^*$;
- 3 **for** ℓ from 1 to L **do**
- 4 $\vec{\pi}_i^\ell = \sqrt{c}P \cdot \vec{\pi}_i^{\ell-1}$;
- 5 $\vec{\pi}_i = \vec{\pi}_i + \vec{\pi}_i^\ell$;
- 6 $R = \frac{6 \log n}{(1 - \sqrt{c})^4 \varepsilon^2}$;
- 7 **for** k from 1 to n **do**
- 8 Invoke Algorithm 2 with $R(k) = \lceil R \cdot \vec{\pi}_i(k) \rceil$ to obtain an estimator $\hat{D}(k, k)$ for $D(k, k)$;
- 9 $\vec{s}^0 = \frac{1}{1 - \sqrt{c}} \hat{D} \cdot \vec{\pi}_i^L$;
- 10 **for** ℓ from 1 to L **do**
- 11 $\vec{s}^\ell = \sqrt{c}P^\top \cdot \vec{s}^{\ell-1} + \frac{1}{1 - \sqrt{c}} \hat{D} \cdot \vec{\pi}_i^{L-\ell}$;
- 12 Clear $\vec{s}^{\ell-1}$;
- 13 **return** \vec{s}^L ;

Algorithm 2: Basic method for estimating $D(k, k)$

Input: Graph G , node v_k , number of samples $R(k)$
Output: $\hat{D}(k, k)$ as an estimation for $D(k, k)$

- 1 $\hat{D}(k, k) = 0$;
- 2 **for** x from 1 to $R(k)$ **do**
- 3 Sample two independent \sqrt{c} -walks from v_k ;
- 4 **if** The two \sqrt{c} -walks do not meet **then**
- 5 $\hat{D}(k, k) = \hat{D}(k, k) + 1/R(k)$;
- 6 **return** $\hat{D}(k, k)$;

10^{-7} and large graph size m , we have a working algorithm for exact single-source SimRank queries on large graphs.

Algorithm 1 illustrates the pseudocode of the basic ExactSim algorithm. Note that to cope with Personalized PageRank, we use the fact that $\vec{\pi}_i^\ell = (1 - \sqrt{c}) \cdot (\sqrt{c}P)^\ell \cdot \vec{e}_i^*$ and re-write Equation (5) as

$$S \cdot \vec{e}_i^* = \frac{1}{1 - \sqrt{c}} \sum_{\ell=0}^{\infty} (\sqrt{c}P^\top)^\ell D \cdot \vec{\pi}_i^\ell. \quad (11)$$

Given a source node v_i and a maximum error ε , we first set the number of iterations L to be $L = \lceil \log_{\frac{1}{c}} \frac{2}{\varepsilon} \rceil$ (line 1). We then iteratively compute the ℓ -hop Personalized PageRank vector $\vec{\pi}_i^\ell = (\sqrt{c}P)^\ell \cdot \vec{e}_i^*$ for $\ell = 0, \dots, L$, as well as the Personalized PageRank vector $\vec{\pi}_i = \sum_{\ell=0}^L \vec{\pi}_i^\ell$ (lines 2-5). To obtain an estimator \hat{D} for the diagonal correction matrix D , we set the total number of samples to be $R = \frac{6 \log n}{(1 - \sqrt{c})^4 \varepsilon^2}$ (line 6). For each $D(k, k)$, we set $R(k) = \lceil R \vec{\pi}_i(k) \rceil$ and invoke Algorithm 2 to estimate $D(k, k)$ (lines 7-8). Algorithm 2 essentially simulates

$R(k)$ pairs of \sqrt{c} -walks from node v_k and uses the fraction of pairs that do not meet as an estimator $\hat{D}(k, k)$ for $D(k, k)$. Finally, we use Equation (11) to iteratively compute $\vec{s}^0 = \frac{1}{1 - \sqrt{c}} \hat{D} \cdot \vec{\pi}_i^L$,

$$\begin{aligned} \vec{s}^1 &= \sqrt{c}P^\top \cdot \vec{s}^0 + \frac{1}{1 - \sqrt{c}} \hat{D} \cdot \vec{\pi}_i^{L-1} \\ &= \frac{1}{1 - \sqrt{c}} (\sqrt{c}P^\top \cdot \hat{D} \cdot \vec{\pi}_i^L + \hat{D} \cdot \vec{\pi}_i^{L-1}) \end{aligned} \quad (12)$$

(lines 9-12), ..., and

$$\begin{aligned} \vec{s}^L &= \frac{(\sqrt{c}P^\top (\dots (\sqrt{c}P^\top \cdot \hat{D} \cdot \vec{\pi}_i^L + \hat{D} \cdot \vec{\pi}_i^{L-1}) + \dots) + \hat{D} \cdot \vec{\pi}_i^0)}{1 - \sqrt{c}} \\ &= \frac{1}{1 - \sqrt{c}} \sum_{\ell=0}^L (\sqrt{c}P^\top)^\ell \hat{D} \cdot \vec{\pi}_i^\ell. \end{aligned} \quad (13)$$

We return \vec{s}^L as the single-source query result (line 13).

Analysis. To derive the running time and space overhead of the basic ExactSim algorithm, note that computing and storing each ℓ -hop Personalized PageRank vector $\vec{\pi}_i^\ell$ takes $O(m)$ time and $O(n)$ space. This results a running time of $O(mL)$ and a space overhead of $O(nL)$. To estimate the diagonal correction matrix D , the algorithm simulates R pairs of \sqrt{c} -walks, each of which takes $\frac{1}{\sqrt{c}} = O(1)$ time. Therefore, the running time for estimating D can be bounded by $O(R)$. Finally, computing each \vec{s}^ℓ also takes $O(m)$ time, resulting an additional running time of $O(mL)$. Summing up all costs, and we have the total running time is bounded by $O(mL + R) = O\left(\frac{\log n}{\varepsilon^2} + m \log \frac{1}{\varepsilon}\right)$, and the space overhead is bounded by $O(nL) = O(n \log \frac{1}{\varepsilon})$.

We now analyze the error of the basic ExactSim algorithm. Recall that ExactSim returns $\vec{s}^L(j)$ as the estimator for $S(i, j)$, the SimRank similarity between the source node v_i and any other node v_j . We have the following Theorem.

Theorem 1 *With probability at least $1 - 1/n$, for any source node $v_i \in V$, the basic ExactSim provide an single-source SimRank vector \vec{s}^L such that, for any node $v_j \in V$, we have $|\vec{s}^L(j) - S(i, j)| \leq \varepsilon$.*

Theorem 1 essentially states that with high probability, the basic ExactSim algorithm can compute any single-source SimRank query with additive ε . The proof of Theorem 1 is fairly technical. However, the basic idea is to show that the variance of the estimator $\vec{s}^L(j)$ can be bounded by $O(\frac{1}{R}) = O(\varepsilon^2)$. In particular, we first note that by Equation (13), $\vec{s}^L(j)$ can be expressed as

$$\begin{aligned} \vec{s}^L(j) &= \vec{e}_j^\top \cdot \vec{s}^L = \frac{1}{1 - \sqrt{c}} \vec{e}_j^\top \cdot \sum_{\ell=0}^L (\sqrt{c}P^\top)^\ell \hat{D} \cdot \vec{\pi}_i^\ell \\ &= \frac{1}{(1 - \sqrt{c})^2} \sum_{\ell=0}^L \left((1 - \sqrt{c}) (\sqrt{c}P)^\ell \cdot \vec{e}_j \right)^\top \cdot \hat{D} \cdot \vec{\pi}_i^\ell. \end{aligned}$$

Since $(1 - \sqrt{c})(\sqrt{c}P)^\ell \cdot \vec{e}_j = \vec{\pi}_j^\ell$, we have

$$\vec{s}^L(j) = \frac{1}{(1 - \sqrt{c})^2} \sum_{\ell=0}^L (\vec{\pi}_j^\ell)^\top \cdot \hat{D} \cdot \vec{\pi}_i^\ell. \quad (14)$$

Summing up over the diagonal elements of D follows that

$$\vec{s}^L(j) = \frac{1}{(1 - \sqrt{c})^2} \sum_{\ell=0}^L \sum_{k=1}^n \vec{\pi}_i^\ell(k) \cdot \vec{\pi}_j^\ell(k) \cdot \hat{D}(k, k). \quad (15)$$

We observe that there are two discrepancies between $\vec{s}^L(j)$ and the actual SimRank value $S(i, j)$ (10): 1) We change the number of iterations from ∞ to L , and 2) we use the estimator \hat{D} to replace actual diagonal correction matrix D . For the first approximation, we can bound the error by $c^L \leq \varepsilon/2$ if ExactSim sets $L = \lceil \log_{\frac{1}{c}} \frac{2}{\varepsilon} \rceil$. Consequently, we only need to bound the error from replacing D with \hat{D} . In particular, we will make use of the following Bernstein Inequality.

Lemma 1 (Bernstein Inequality [8]) *Let X_1, \dots, X_R be independent random variables with $|X_i| < b$ for $i = 1, \dots, R$. Let $X = \frac{1}{R} \cdot \sum_{i=1}^R X_i$, we have*

$$\Pr[|X - \mathbb{E}[X]| \geq \lambda] \leq 2 \cdot \exp\left(-\frac{\lambda^2 \cdot R}{2R \cdot \text{Var}[X] + 2b\lambda/3}\right), \quad (16)$$

where $\text{Var}[X]$ is the variance of X .

To make use of Lemma 1, we need to express $\vec{s}^L(j)$ as the average of independent random variables. In particular, let $\hat{D}_r(k, k)$, $r = 1, \dots, R(k)$ denote the r -th estimator of $D(k, k)$ by Algorithm 2. We observe that each $\hat{D}_r(k, k)$ is a Bernoulli random variable, that is, $\hat{D}_r(k, k) = 1$ with probability $D(k, k)$ and $\hat{D}_r(k, k) = 0$ with probability $1 - D(k, k)$. We have

$$\begin{aligned} \vec{s}^L(j) &= \frac{1}{(1 - \sqrt{c})^2} \sum_{\ell=0}^L \sum_{k=1}^n \vec{\pi}_i^\ell(k) \cdot \vec{\pi}_j^\ell(k) \cdot \frac{\sum_{r=1}^{R(k)} \hat{D}_r(k, k)}{R(k)} \\ &= \frac{1}{(1 - \sqrt{c})^2} \sum_{k=1}^n \sum_{r=1}^{R(k)} \frac{\sum_{\ell=0}^L \vec{\pi}_i^\ell(k) \cdot \vec{\pi}_j^\ell(k)}{R(k)} \cdot \hat{D}_r(k, k). \end{aligned}$$

Let $\rho(k) = R(k)/R$ be the fraction of pairs of \sqrt{c} -walks assigned to v_k , it follows that

$$\vec{s}^L(j) = \frac{1}{R} \cdot \frac{1}{(1 - \sqrt{c})^2} \sum_{k=1}^n \sum_{r=1}^{R(k)} \frac{\sum_{\ell=0}^L \vec{\pi}_i^\ell(k) \cdot \vec{\pi}_j^\ell(k)}{\rho(k)} \cdot \hat{D}_r(k, k). \quad (17)$$

We will treat each $\frac{\sum_{\ell=0}^L \vec{\pi}_i^\ell(k) \cdot \vec{\pi}_j^\ell(k)}{\rho(k)} \cdot \hat{D}_r(k, k)$ as an independent random variable. The number of such random

variables is $\sum_{k=1}^n R\rho(k) = R$, so we have expressed $\vec{s}^L(j)$ as the average of R independent random variables. To utilize Lemma 1, we first bound the variance of $\vec{s}^L(j)$.

Lemma 2 *The variance of $\vec{s}^L(j)$ is bounded by*

$$\text{Var}[\vec{s}^L(j)] \leq \frac{1}{(1 - \sqrt{c})^4 R} \sum_{k=1}^n \frac{\vec{\pi}_i(k)^2 \vec{\pi}_j(k)^2}{\rho(k)} \cdot D(k, k). \quad (18)$$

In particular, by setting $\rho(k) = R(k)/R = \lceil R\vec{\pi}_i(k) \rceil / R$ in the basic ExactSim algorithm, we have

$$\text{Var}[\vec{s}^L(j)] \leq \frac{1}{(1 - \sqrt{c})^4 R}. \quad (19)$$

Note that we only need Inequality (19) to derive the error bound for the basic ExactSim algorithm. The more complex Inequality (18) will be used to design various optimization techniques.

Proof [Proof of Lemma 2] Note that $\hat{D}_r(k, k)$ is a Bernoulli random variable with expectation $D(k, k)$, and thus has variance $\text{Var}[\hat{D}_r(k, k)] = D(k, k)(1 - D(k, k)) \leq D(k, k)$. Since $\hat{D}_r(k, k)$'s are independent random variables, we have

$$\begin{aligned} \text{Var}[\vec{s}^L(j)] &= \frac{1}{(1 - \sqrt{c})^4 R^2} \sum_{k=1}^n \sum_{r=1}^{R(k)} \left(\frac{\sum_{\ell=0}^L \vec{\pi}_i^\ell(k) \cdot \vec{\pi}_j^\ell(k)}{\rho(k)} \right)^2 \cdot \text{Var}[\hat{D}_r(k, k)] \\ &= \frac{1}{(1 - \sqrt{c})^4 R} \sum_{k=1}^n \frac{\left(\sum_{\ell=0}^L \vec{\pi}_i^\ell(k) \cdot \vec{\pi}_j^\ell(k) \right)^2}{\rho(k)} \cdot D(k, k)(1 - D(k, k)). \end{aligned}$$

By the Cauchy-Schwarz inequality, we have

$$\begin{aligned} \left(\sum_{\ell=0}^L \vec{\pi}_i^\ell(k) \cdot \vec{\pi}_j^\ell(k) \right)^2 &\leq \left(\sum_{\ell=0}^L \vec{\pi}_i^\ell(k) \right)^2 \left(\sum_{\ell=0}^L \vec{\pi}_j^\ell(k) \right)^2 \\ &\leq \vec{\pi}_i(k)^2 \vec{\pi}_j(k)^2. \end{aligned}$$

Combining with the fact that $1 - D(k, k) \leq 1$, we have

$$\text{Var}[\vec{s}^L(j)] \leq \frac{1}{(1 - \sqrt{c})^4 R} \sum_{k=1}^n \frac{\vec{\pi}_i(k)^2 \vec{\pi}_j(k)^2}{\rho(k)} \cdot D(k, k). \quad (20)$$

and the first part of the Lemma follows.

Plugging $\rho(k) = R(k)/R = \lceil R\vec{\pi}_i(k) \rceil / R \geq \vec{\pi}_i(k)$ into Lemma 2, we have

$$\begin{aligned} \text{Var}[\vec{s}^L(j)] &\leq \frac{1}{(1 - \sqrt{c})^4 R} \sum_{k=1}^n \frac{\vec{\pi}_i(k)^2 \vec{\pi}_j(k)^2}{\vec{\pi}_i(k)} \cdot D(k, k) \\ &\leq \frac{1}{(1 - \sqrt{c})^4 R} \sum_{k=1}^n \vec{\pi}_i(k). \end{aligned}$$

For the last inequality, we use the fact that $D(k, k) \leq 1$ and $\vec{\pi}_j(k) \leq 1$. Finally, since $\sum_{k=1}^n \vec{\pi}_i(k) = 1$, we have $\text{Var}[\vec{s}^L(j)] \leq \frac{1}{(1-\sqrt{c})^4 R}$, and the second part of the Lemma follows.

Proof (Proof of Theorem 1) We are now ready to prove Theorem 1. To utilize Bernstein Inequality given in Lemma 1, we also need to bound b , the maximum value of the random variables $\sum_{\ell=0}^L \frac{\vec{\pi}_i^\ell(k) \cdot \vec{\pi}_j^\ell(k)}{\rho(k)} \cdot \hat{D}_r(k, k)$. We have

$$\frac{\sum_{\ell=0}^L \vec{\pi}_i^\ell(k) \cdot \vec{\pi}_j^\ell(k)}{\vec{\pi}_i(k)} \cdot \hat{D}_r(k, k) \leq \frac{\sum_{\ell=0}^L \vec{\pi}_i^\ell(k)}{\vec{\pi}_i(k)} \leq \frac{\vec{\pi}_i(k)}{\vec{\pi}_i(k)} = 1.$$

Applying Bernstein Inequality with $b=1$ and $\text{Var}[\vec{s}^L(j)] \leq \frac{1}{(1-\sqrt{c})^4 R}$, where $R = \frac{6 \log n}{(1-\sqrt{c})^4 \varepsilon^2}$, we have $\Pr[|\vec{s}^L(j) - E[\vec{s}^L(j)]| > \varepsilon/2] < 1/n^3$. Combining with the $\varepsilon/2$ error introduced by the truncation L , we have $\Pr[|\vec{s}^L(j) - S(i, j)| > \varepsilon] < 1/n^3$. By union bound over all possible target nodes $j = 1, \dots, n$ and all possible source nodes $i = 1, \dots, n$, we ensure that for all n possible source node and n target nodes,

$$\Pr[\forall i, j, |\vec{s}^L(j) - S(i, j)| > \varepsilon] < 1/n,$$

and the Theorem follows.

4 Optimizations

Although the basic ExactSim algorithm is a working algorithm for exact single-source SimRank computation on large graphs, it suffers from some drawbacks. First of all, the $O(n \log \frac{1}{\varepsilon})$ space overhead can be several times larger than the actual graph size m . Secondly, we still need to simulate $R = O\left(\frac{\log n}{\varepsilon^2}\right)$ of pairs of \sqrt{c} -walks, which is a significant cost for $\varepsilon_{\min} = 10^{-7}$. Although parallelization can help, we are still interested in developing algorithmic techniques that reduces the number of random walks. In this section, we provide three optimization techniques that address these drawbacks.

Sparse Linearization. We design a sparse version of Linearization that significantly reduces the $O(n \log \frac{1}{\varepsilon})$ space overhead while retaining the $O(\varepsilon)$ error guarantee. Recall that this space overhead is causing by storing the ℓ -hop Personalized PageRank vectors $\vec{\pi}_i^\ell$ for $\ell = 0, \dots, L$. We propose to make the following simple modification: Instead of storing the dense vector $\vec{\pi}_i^\ell$, we sparsify the vector by removing all entries of with $\vec{\pi}_i^\ell(k) \leq (1 - \sqrt{c})^2 \varepsilon$. To understand the effectiveness of this approach, recall that a nice property of the ℓ -hop Personalized PageRank vectors is that all possible entries sum up to $\sum_{\ell=0}^{\infty} \sum_{k=1}^n \vec{\pi}_i^\ell(k) = \sum_{k=1}^n \vec{\pi}_i^\ell(k) = 1$. By the Pigeonhole principle, the number of $\vec{\pi}_i^\ell(k)$'s that

are larger than $(1 - \sqrt{c})^2 \varepsilon$ is bounded by $\frac{1}{(1-\sqrt{c})^2 \varepsilon}$. Thus the space overhead is reduced to $O\left(\frac{1}{\varepsilon}\right)$. This overhead is acceptable for exact computations where we set $\varepsilon = \varepsilon_{\min} = 10^{-7}$, as it does not scale with the graph size.

Sampling according to $\vec{\pi}_i(k)^2$. Recall that in the basic ExactSim algorithm, we simulate R pairs of \sqrt{c} -walks in total, and distribute $\vec{\pi}_i(k)$ fraction of the R samples to estimate $D(k, k)$. A natural question is that, is there a better scheme to distribute these R samples? It turns out if we distribute the samples according to $\vec{\pi}_i(k)^2$, we can further reduce the variance of the estimator and hence achieve a better running time. More precisely, we will set $R(k) = R \left\lceil \frac{\vec{\pi}_i(k)^2}{\|\vec{\pi}_i\|^2} \right\rceil$, where $\|\vec{\pi}_i\|^2 = \sum_{k=1}^n \vec{\pi}_i(k)^2$ is the squared norm of the Personalized PageRank vector $\vec{\pi}_i$.

Local deterministic exploitation for D . The inequality (18) in Lemma 2 also suggests that we can reduce the variance of the estimator $\vec{s}^L(j)$ by refining the Bernoulli estimator $\hat{D}(k, k)$. Recall that we sample $R(k) = \lceil R \vec{\pi}_i(k) \rceil$ or $R(k) = R \left\lceil \frac{\vec{\pi}_i(k)^2}{\|\vec{\pi}_i\|^2} \right\rceil$ pairs of \sqrt{c} -walks to estimate $D(k, k)$. If $\vec{\pi}_i(k)$ is large, we will simulate a large number of \sqrt{c} -walks from v_k to estimate $D(k, k)$. In that case, the first few steps of these random walks will most likely visit the same local structures around v_k , so it makes sense to exploit these local structures deterministically, and use the random walks to approximate the global structures. More precisely, let $Z_\ell(k)$ denote the probability that two \sqrt{c} -walks from v_k first meet at the ℓ -th step. Since these events are mutually exclusive for different ℓ 's, we have

$$\begin{aligned} D(k, k) &= 1 - \Pr[\text{two } \sqrt{c}\text{-walks from } v_k \text{ meet}] \\ &= 1 - \sum_{\ell=1}^{\infty} Z_\ell(k). \end{aligned}$$

The idea is to deterministically compute $\sum_{\ell=1}^{\ell(k)} Z_\ell(k)$ for some tolerable step $\ell(k)$, and using random walks to estimate the other part $\sum_{\ell=\ell(k)+1}^{\infty} Z_\ell(k)$. It is easy to see that by deterministically computing $Z_\ell(k)$ for the first $\ell(k)$ levels, we reduce the variance $\text{Var}(D(k, k))$ by at least $c^{\ell(k)}$.

A simple algorithm to compute $Z_\ell(k)$ is to list all possible paths of length ℓ from v_k and aggregate all meeting probabilities of any two paths. However, the number of paths increases rapidly with the length ℓ , which makes this algorithm inefficient on large graphs. Instead, we will derive the close form for $Z_\ell(k)$ in terms of the transition probabilities. In particular, let $Z_\ell(k, q)$ denote the probability that two \sqrt{c} -walks first meet at node v_q at their ℓ -th steps. We have $Z_\ell(k) = \sum_{q=1}^n Z_\ell(k, q)$,

Algorithm 3: Improved method for estimating $D(k, k)$

```

Input: Graph  $G$ , node  $v_k$ , sample number  $R(k)$ 
Output: An estimator  $\hat{D}(k, k)$  for  $D(k, k)$ 
1 if  $d_{in}(v_k) = 0$  then
2   return  $\hat{D}(k, k) = 1$ ;
3 else if  $d_{in}(v_k) = 1$  then
4   return  $\hat{D}(k, k) = 1 - c$ ;
5  $P^\ell(x, k) = 0$  for  $\ell \geq 0, x \in V$ ;
6  $P^0(k, k) = 1$ ;
7  $E_k = 0$ ;
8 for  $\ell$  from 0 to  $\infty$  do
9   for each  $v_q$  with non-zero  $(P^\top)^\ell(k, q)$  do
10    Calculate  $Z_\ell(k, q)$  using equation (22);
11   for  $\ell'$  from 0 to  $\ell$  do
12    for each  $v_{q'}$  with non-zero  $(P^\top)^{\ell-\ell'}(k, q')$  do
13      for each  $v_x$  with non-zero  $(P^\top)^{\ell'}(q', x)$  do
14        for each  $v_q \in \mathcal{I}(v_x)$  do
15           $(P^\top)^{\ell'+1}(q', q) += \frac{(P^\top)^{\ell'}(q', x)}{d_{in}(v_x)}$ ;
16           $E_k += 1$ ;
17          if  $E_k \geq \frac{2R(k)}{\sqrt{c}}$  then
18             $\ell(k) = \ell$  and goto OUTLOOP;
19    $\ell = \ell + 1$ ;
20 OUTLOOP;
21  $\hat{D}(k, k) = 1 - \sum_{\ell=1}^{\ell(k)} \sum_{q=1}^n Z_\ell(k, q)$ ;
22 for  $z$  from 1 to  $R(k)$  do
23   Sample two independent non-stop random walks
      from  $v_k$ ;
24   if Two random walks reaches nodes  $v_x$  and  $v_y$  at
      the  $\ell(k)$  steps without meeting then
25     Sample a  $\sqrt{c}$ -walks from  $v_x$  and  $v_y$ ;
26     if the two  $\sqrt{c}$ -walks meet then
27        $\hat{D}(k, k) = \hat{D}(k, k) - c^{\ell(k)}/R(k)$ ;
28 return  $\hat{D}(k, k)$ ;

```

and hence

$$D(k, k) = 1 - \sum_{\ell=1}^{\infty} \sum_{q=1}^n Z_\ell(k, q). \quad (21)$$

Recall that P^ℓ (the ℓ -th power of the (reverse) transition matrix P) is the ℓ -step (reverse) transition matrix. We have the following Lemma that relates $Z_\ell(k, q)$ with the transition probabilities.

Lemma 3 $Z_\ell(k, q)$ satisfies the following recursive form:

$$\begin{aligned} Z_\ell(k, q) = & c^\ell (P^\top)^\ell(k, q)^2 \\ & - \sum_{\ell'=1}^{\ell-1} \sum_{q'=1}^n c^{\ell'} (P^\top)^{\ell'}(q', q)^2 Z_{\ell-\ell'}(k, q'). \end{aligned} \quad (22)$$

Proof Note that $(\sqrt{c})^\ell (P^\top)^\ell(k, q)$ is the probability that a \sqrt{c} -walk from v_k visits v_q at its ℓ -th step. Consequently, $c^\ell (P^\top)^\ell(k, q)^2$ is the probability that two \sqrt{c} -walks from node v_k visit node v_q at their ℓ -th step simultaneously. To ensure this is the first time that the two \sqrt{c} -walks meet, we subtract the probability mass that the two \sqrt{c} -walks have met before. In particular, recall that $Z_{\ell'}(k, q')$ is the probability that two \sqrt{c} -walks from node v_k first meet at $v_{q'}$ in exactly ℓ' steps. Due to the memoryless property of the \sqrt{c} -walk, the two \sqrt{c} -walks will behave as two new \sqrt{c} -walks from $v_{q'}$ after their ℓ' -th step. The probability that these two new \sqrt{c} -walks visit v_q in exact $\ell - \ell'$ steps is $c^{\ell-\ell'} (P^\top)^{\ell-\ell'}(q', q)^2$. Summing up q' from 1 to n and ℓ' from 1 to $\ell - 1$, and the Lemma follows.

Given a node v_k and a pre-determined target level $\ell(k)$, Lemma 3 also suggests a simple algorithm to compute $Z_\ell(k, q)$ for all $\ell \leq \ell(k)$. We start by performing BFS from node v_k for up to $\ell(k)$ levels to calculate the transition probabilities $(P^\top)^\ell(k, q)$ for $\ell = 0, \dots, \ell(k)$ and $v_q \in V$. For each node q' visited at the ℓ' -th level, we start a BFS from q' for $\ell(k) - \ell'$ levels to calculate $(P^\top)^{\ell(k)-\ell'}(q', q)$ for $\ell = 1, \dots, \ell(k)$ and $v_q \in V$. Then we use equation (22) to calculate $Z_\ell(k, q)$ for $\ell = 0, \dots, \ell(k)$ and $q \in V$. Note that this approach exploits strictly less edges than listing all possible paths of length $\ell(k)$, as some of the paths are combined in the computation of the transition probabilities.

However, a major problem with the above method is that the target level $\ell(k)$ has to be predetermined, which makes the running time unpredictable. An improper value of $\ell(k)$ could lead to the explosion of the running time. Instead, we will use an adaptive algorithm to compute $Z_\ell(k)$.

Algorithm 3 illustrates the new method for estimating $D(k, k)$. Given a node v_k and a sample number $R(k)$, the goal is to give an estimator for $D(k, k)$. For the two trivial case $d_{in}(k) = 0$ and $d_{in}(k) = 1$, we return $D(k, k) = 1$ and $1 - c$ accordingly (lines 1-4). For other cases, we iteratively compute all possible transition probabilities $(P^\top)^{\ell'+1}(q', q)$ for all $v_{q'}$ that is reachable from k with $\ell - \ell'$ steps (lines 5-10). Note that these $v_{q'}$'s are the nodes with $(P^\top)^{\ell-\ell'}(k, q') > 0$. To ensure the deterministic exploitation stops in time, we use a counter E_k to record the total number of edges traversed so far (line 11). If E_k exceeds $\frac{2R(k)}{\sqrt{c}}$, the expected number of steps for simulating $R(k)$ pairs of \sqrt{c} -walks, we terminate the deterministic exploitation and set $\ell(k)$ as the current target level for v_k (lines 12-13). After we fix $\ell(k)$ and compute $\sum_{\ell=1}^{\ell(k)} Z_\ell(k)$ (lines 14-17), we will use random walk sampling to es-

timate $\sum_{\ell=\ell(k)+1}^{\infty} Z_{\ell}(k)$ (lines 18-23). In particular, we start two special random walks from v_k . The random walks do not stop in its first $\ell(k)$ steps; after the $\ell(k)$ -th step, each random walk stops with probability \sqrt{c} at each step. It is easy to see that the probability of the two special random walks meet after $\ell(k)$ steps is $\frac{1}{c^{\ell(k)}} \sum_{\ell=\ell(k)+1}^{\infty} Z_{\ell}(k)$. Consequently, we can use the fraction of the random walks that meet multiplied by $c^{\ell(k)}$ as an unbiased estimator for $\sum_{\ell=\ell(k)+1}^{\infty} Z_{\ell}(k)$.

Parallelization. The ExactSim algorithm is highly parallelizable as it only uses two primitive operations: matrix-(sparse) vector multiplication and random walk simulation. Both operations are embarrassingly parallelizable on GPUs or multi-core CPUs. The only exception is the local deterministic exploitation for $D(k, k)$. To parallelize this operation, we can apply Algorithm 3 to multiple v_k simultaneously. Furthermore, we can balance the load of each thread by applying Algorithm 3 to nodes v_k 's with similar number of samples $R(k)$ in each epoch.

4.1 Analysis

Recall that Algorithm 3 provides an improved method for estimating $D(k, k)$. By invoking Algorithm 3 into the whole ExactSim structure (line 8 in Algorithm 1), we can derive the optimized version of ExactSim. The following theorem presents the complexity analysis of the optimized ExactSim in terms of time cost and space overhead.

Theorem 2 *Let $\vec{\pi}_i$ denote the Personalized PageRank vector with regards to node v_i . Then with probability at least $1 - \frac{1}{n}$, for any source node $v_i \in V$, the optimized ExactSim can return a single-source SimRank vector \vec{s}^L with $O\left(\frac{\|\vec{\pi}_i\|^2 \log n}{\varepsilon^2} + m \log \frac{1}{\varepsilon}\right)$ time cost and $O\left(\frac{1}{\varepsilon}\right)$ space overhead, such that for any node $v_j \in V$, we have $\|\vec{s}^L(j) - S(i, j)\| \leq \varepsilon$.*

Concerning the three optimization techniques mentioned above, sparse Linearization may influence the space overhead; Sampling according to $\vec{\pi}_i(k)^2$ reduces the number of random walks, which can impact the time cost of estimating D . Local deterministic exploitation can reduce the variance $\text{Var}(D(k, k))$, while the level of time and space complexity remains the same due to the setting of $\ell(k)$. Consequently, to prove Theorem 2, we can only analysis sparse Linearization for space bound, and sampling according to $\vec{\pi}_i(k)^2$ for time cost, respectively.

Firstly, as for the space overhead, the following lemma proves that the sparse Linearization will only introduce

an extra additive error of ε . If we scale down ε by a factor of 2, the total error guarantee and the asymptotic running time of ExactSim will remain the same, and the space overhead is reduced to $O\left(\frac{1}{\varepsilon}\right)$.

Lemma 4 *The sparse Linearization introduces an extra additive error of ε and reduces the space overhead to $O\left(\frac{1}{\varepsilon}\right)$.*

Proof We note that the sparse Linearization introduces an extra error of $(1 - \sqrt{c})^2 \varepsilon$ to each $\vec{\pi}_i^{\ell}(k)$, $k = 1, \dots, n$, $\ell = 0, \dots, \infty$. According to equation (15), the estimator $\vec{s}^L(j)$ can be expressed as

$$\vec{s}^L(j) = \frac{1}{(1 - \sqrt{c})^2} \sum_{\ell=0}^L \sum_{k=1}^n (\vec{\pi}_i^{\ell}(k) \pm (1 - \sqrt{c})^2 \varepsilon) \cdot \vec{\pi}_j^{\ell}(k) \cdot \hat{D}(k, k). \quad (23)$$

Thus, the error introduced by sparse Linearization can be bounded by

$$\frac{1}{(1 - \sqrt{c})^2} \sum_{\ell=0}^{\infty} \sum_{k=1}^n (1 - \sqrt{c})^2 \varepsilon \cdot \vec{\pi}_j^{\ell}(k) \cdot \hat{D}(k, k). \quad (24)$$

Using the facts that $\sum_{\ell=0}^{\infty} \sum_{k=1}^n \vec{\pi}_j^{\ell}(k) = 1$ and $\hat{D}(k, k) \leq 1$, the above error can be bounded by $\frac{1}{(1 - \sqrt{c})^2} \cdot (1 - \sqrt{c})^2 \varepsilon = \varepsilon$, and the lemma follows.

Then we analysis the time cost of Algorithm 3. The following Lemma shows that by sampling according to $\vec{\pi}_i(k)^2$, we can reduce the number of sample R by a factor of $\|\vec{\pi}_i\|^2$.

Lemma 5 *By sampling according to $\vec{\pi}_i(k)^2$, the number of random samples required is reduced to $O\left(\frac{\|\vec{\pi}_i\|^2 \log n}{\varepsilon^2}\right)$.*

Proof Recall that $\rho(k)$ is the fraction of sample assigned to $D(k, k)$. We have $\rho(k) = \left\lceil \frac{R \vec{\pi}_i(k)^2}{\|\vec{\pi}_i\|^2} \right\rceil / R \geq \frac{\vec{\pi}_i(k)^2}{\|\vec{\pi}_i\|^2}$. By the inequality (18) in Lemma 2, we can bound the variance of estimator $\vec{s}^L(j)$ as

$$\begin{aligned} \text{Var}[\vec{s}^L(j)] &\leq \frac{1}{(1 - \sqrt{c})^4 R} \sum_{k=1}^n \frac{\vec{\pi}_i(k)^2 \vec{\pi}_j(k)^2}{\rho(k)} \cdot D(k, k) \\ &\leq \frac{1}{(1 - \sqrt{c})^4 R} \|\vec{\pi}_i\|^2 \sum_{k=1}^n \vec{\pi}_j(k)^2 = \frac{1}{(1 - \sqrt{c})^4 R} \|\vec{\pi}_i\|^2 \|\vec{\pi}_j\|^2. \end{aligned}$$

Here, we use the facts that $\|\vec{\pi}_j\|^2 = \sum_{k=1}^n \vec{\pi}_j(k)^2$ and $D(k, k) \leq 1$. Since we need to bound the variance for all possible nodes v_j (and hence all possible $\|\vec{\pi}_j\|^2$), we make the relaxation that $\|\vec{\pi}_j\|^2 \leq \|\vec{\pi}_j\|_1^2 = 1$, where $\|\vec{\pi}_j\|_1^2 = (\sum_{k=1}^n |\vec{\pi}_j(k)|)^2$. And thus

$$\text{Var}[\vec{s}^L(j)] \leq \frac{1}{(1 - \sqrt{c})^4 R} \|\vec{\pi}_i\|^2.$$

This suggest that by sampling according to $\vec{\pi}_i(k)^2$, we reduce the variance of the estimators by a factor $\|\vec{\pi}_i\|^2$.

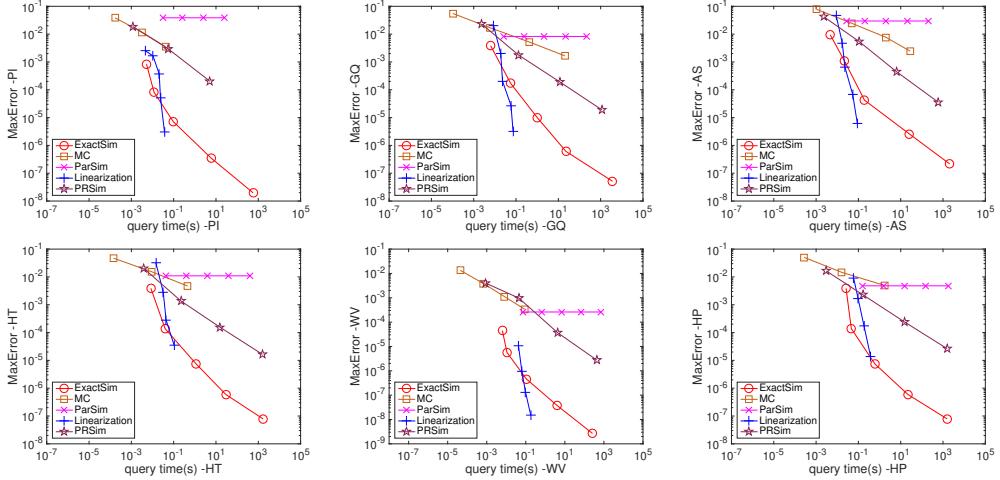


Fig. 1 MaxError v.s. Query time on small graphs

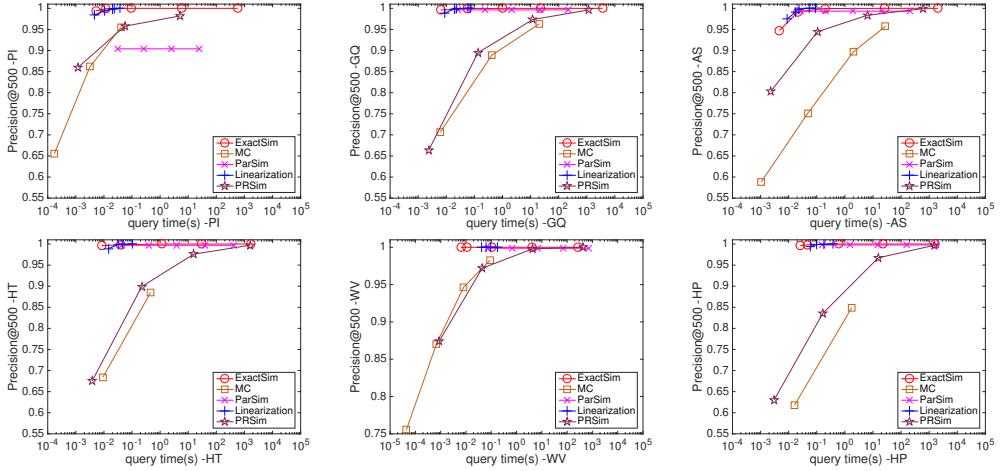


Fig. 2 Precision@500 v.s. Query time on small graphs

Recall that the ExactSim algorithm computes the Personalized PageRank vector $\vec{\pi}_i$ before estimating D , we can obtain the value of $\|\vec{\pi}_i\|^2$ and scale R down by a factor of $\|\vec{\pi}_i\|^2$. This simple modification will reduce the running time to $O\left(\frac{\|\vec{\pi}_i\|^2 \log n}{\varepsilon^2}\right)$.

One small technical issue is that the maximum of the random variables $\frac{\sum_{\ell=0}^{\infty} \vec{\pi}_i^\ell(k) \cdot \vec{\pi}_j^\ell(k)}{\rho(k)} \cdot \hat{D}_r(k, k)$ may get too large as the fraction $\rho(k)$ gets too small. However, by the facts that $\rho(k) = \left\lceil \frac{R\vec{\pi}_i(k)^2}{\|\vec{\pi}_i\|^2} \right\rceil / R$ and $\hat{D}_r(k, k) \leq 1$, we have

$$\sum_{\ell=0}^{\infty} \vec{\pi}_i^\ell(k) \cdot \vec{\pi}_j^\ell(k) \cdot \hat{D}_r(k, k) \leq \frac{\vec{\pi}_i(k)}{\rho(k)} = R\vec{\pi}_i(k) / \left\lceil \frac{R\vec{\pi}_i(k)^2}{\|\vec{\pi}_i\|^2} \right\rceil.$$

If we view the right side of the above equality as a function of $\vec{\pi}_i(k)$, it takes maximum when $\frac{R\vec{\pi}_i(k)^2}{\|\vec{\pi}_i\|^2} = 1$, or equivalently $\vec{\pi}_i(k) = \sqrt{\frac{\|\vec{\pi}_i\|^2}{R}}$. Thus, the random variables in Equation (17) can be bounded by $R\sqrt{\frac{\|\vec{\pi}_i\|^2}{R}} =$

$\|\vec{\pi}_i\| \sqrt{R}$. Plugging $b = \|\vec{\pi}_i\| \sqrt{R}$ and $\text{Var}[\vec{s}^L(j)] \leq \frac{\|\vec{\pi}_i\|^2}{(1-\sqrt{c})^4 R}$ into Bernstein Inequality, and the lemma follows.

To demonstrate the effectiveness of sampling according to $\vec{\pi}_i(k)^2$, notice that in the worst case, $\|\vec{\pi}_i\|^2$ is as large as $\|\vec{\pi}_i\|_1^2 = 1$, so this optimization technique is essentially useless. However, it is known [7] that on scale-free networks, the Personalized PageRank vector $\vec{\pi}_i$ follows a power-law distribution: let $\vec{\pi}_i(k_j)$ denote the j -th largest entry of $\vec{\pi}_i$, we can assume $\vec{\pi}_i(k_j) \sim \frac{j^{-\beta}}{n^{1-\beta}}$ for some power-law exponent $\beta \in (0, 1)$. In this case, $\|\vec{\pi}_i\|^2$ can be bounded by $O\left(\sum_{j=1}^n \left(\frac{j^{-\beta}}{n^{1-\beta}}\right)^2\right) = O\left(\max\left\{\frac{\ln n}{n}, \frac{1}{n^{2-2\beta}}\right\}\right)$, and the $\|\vec{\pi}_i\|^2$ factor becomes significant for any power-law exponent $\beta < 1$.

Note that the expected length of every random walk is $\frac{1}{1-\sqrt{c}}$, which is a constant. Hence, by Lemma 5, the time cost of Algorithm 3 can be bounded by $O\left(\frac{\|\vec{\pi}_i\|^2 \log n}{\varepsilon^2}\right)$. Recall that after we derive the estimated matrix D ,

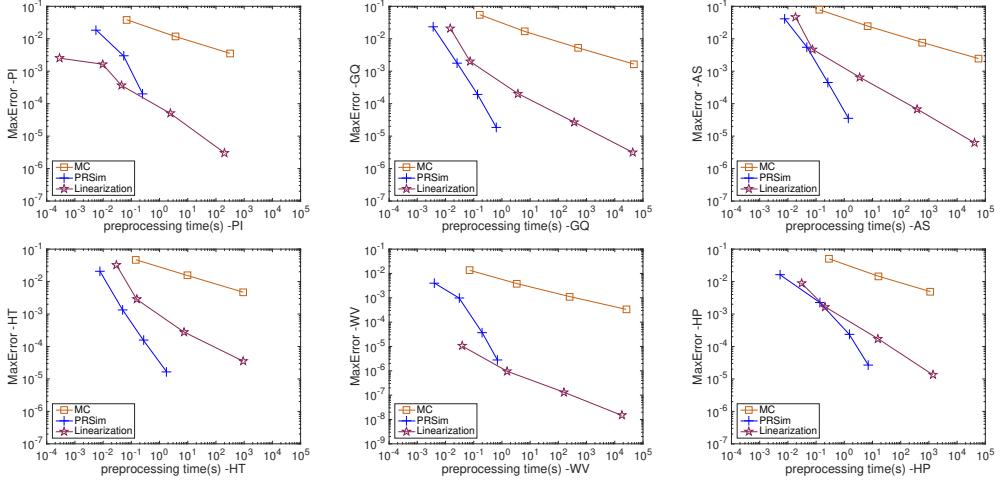


Fig. 3 MaxError v.s. Preprocessing time on small graphs

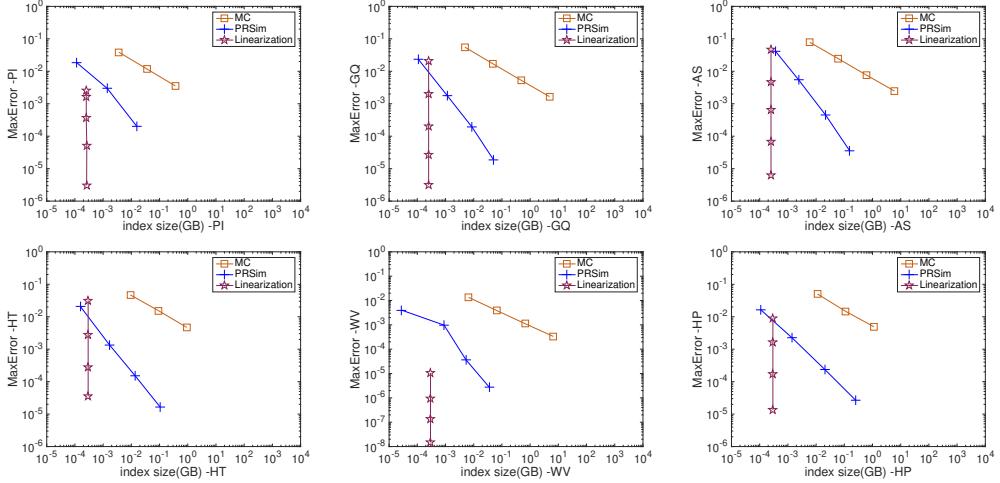


Fig. 4 MaxError v.s. Index size on small graphs

the linearized summation for \bar{s}^L takes $O(m \log \frac{1}{\varepsilon})$ time. Consequently, the total time cost of the optimized ExactSim is $O\left(\frac{\|\vec{n}_i\|^2 \log n}{\varepsilon^2} + m \log \frac{1}{\varepsilon}\right)$, which follows Theorem 2.

5 Experiments

In this section, we experimentally study ExactSim and the other single-source algorithms. We first evaluate ExactSim against four methods MC, ParSim, Linearization and PRSim to prove ExactSim's ability of exact computation (i.e., $\varepsilon_{min} = 10^{-7}$). Then we conduct an ablation study to demonstrate the effectiveness of the optimization techniques. Finally, based on the ground truths computed by ExactSim, we conduct a comprehensive empirical study on existing single-source SimRank algorithms and SimRank distributions.

Table 5 Datasets.

Data Set	Type	<i>n</i>	<i>m</i>
PPI (PI)	undirected	3,890	38739
ca-GrQc (GQ)	undirected	5,242	28,968
AS-2000(AS)	undirected	6,474	25,144
CA-HepTh(HT)	undirected	9,877	51,946
Wikivote (WV)	directed	7,115	103,689
CA-HepPh (HP)	undirected	12,008	236978
DBLP-Author (DB)	undirected	5,425,963	17,298,032
LiveJournal (LJ)	directed	4,847,571	68,475,391
IndoChina (IC)	directed	7,414,768	191,606,827
Orkut-Links (OL)	undirected	3,072,441	234,369,798
It-2004 (IT)	directed	41,290,682	1,135,718,909
Twitter (TW)	directed	41,652,230	1,468,364,884
Wiki-Pl (WP)	dynamic	1,033,050	25,026,208
Wiki-De (WD)	dynamic	2,166,669	86,337,879

Datasets and Environment. We use six small datasets, six large datasets, and two dynamic datasets obtained from [1–3]. The details of these datasets can be found in Table 5. All experiments are conducted on a machine

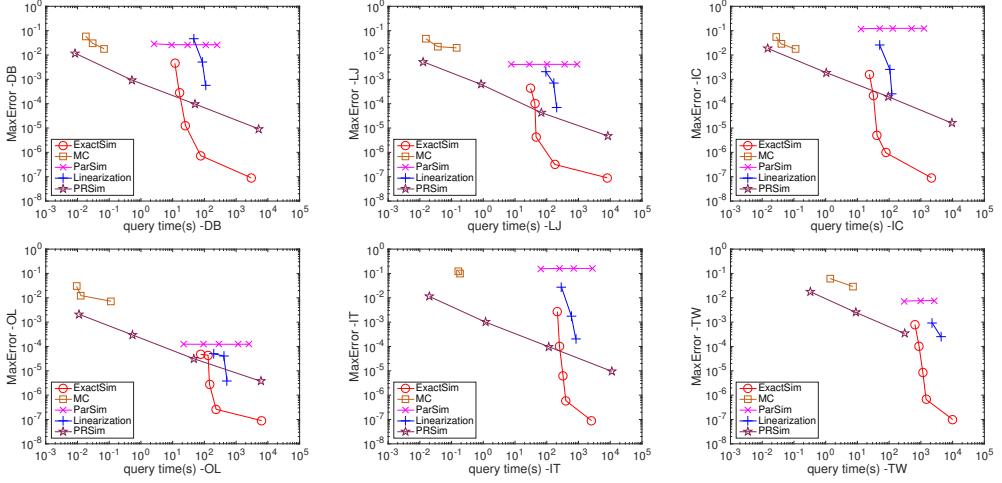


Fig. 5 MaxError v.s. Query time on large graphs

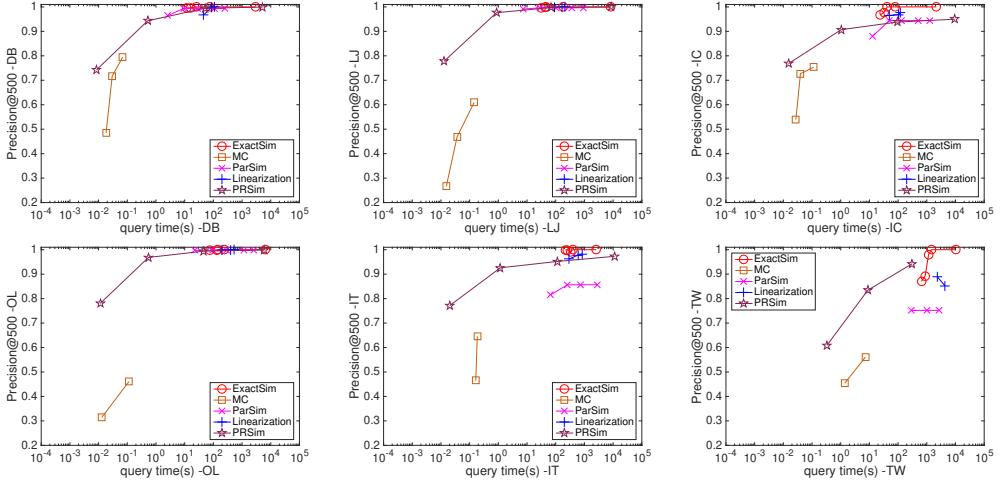


Fig. 6 Precision@500 v.s. Query time on large graphs

with an Intel(R) Xeon(R) E7-4809 @2.10GHz CPU and 196GB memory.

5.1 Evaluation towards ExactSim

Methods and Parameters. We evaluate ExactSim with the four state-of-the-art methods, including one Monte Carlo method: MC [9], two iterative methods: Linearization [29] and ParSim [45], and one Local push/sampling methods: PRSim [39]. For a fair comparison, we run each algorithm in the single thread mode on static graphs.

MC has two parameters: the length of each random walk L and the number of random walks per node r . We vary (L, r) from $(5, 50)$ to $(5000, 50000)$ on small graphs and from $(5, 50)$ to $(50, 500)$ on large graphs. ParSim has one parameter L , the number of iterations. We vary it from 50 to 5×10^5 on small graphs and from

5 to 500 on large graphs. Finally, Linearization, PRSim, and ExactSim share the same error parameter ε , and we vary ε from 10^{-1} to 10^{-7} (if possible) on both small and large graphs. We evaluate the optimized ExactSim unless otherwise stated. In all experiments, we set the decay factor c of SimRank as 0.6.

Metrics. We use *MaxError* and *Precision@k* to evaluate the quality of the single-source and top- k results. Given a source node v_i and an approximate single-source result with n similarities $\hat{S}(i, j), j = 1, \dots, n$, *MaxError* is defined to be the maximum error over n similarities: $\text{MaxError} = \max_{j=1}^n |\hat{S}(i, j) - S(i, j)|$. Given a source node v_i and an approximate top- k result $V_k = \{v_1, \dots, v_k\}$, *Precision@k* is defined to be the percentage of nodes in V_k that coincides with the actual top- k results. In our experiments, we set k to be 500. Note that this is the first time that top- k queries with $k > 100$ are evaluated on large graphs. On each

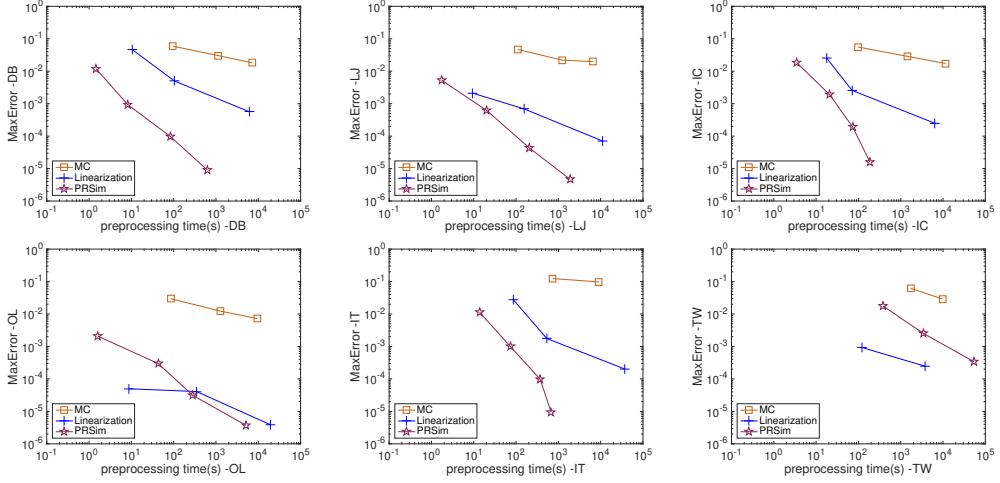


Fig. 7 MaxError v.s. Preprocessing time on large graphs

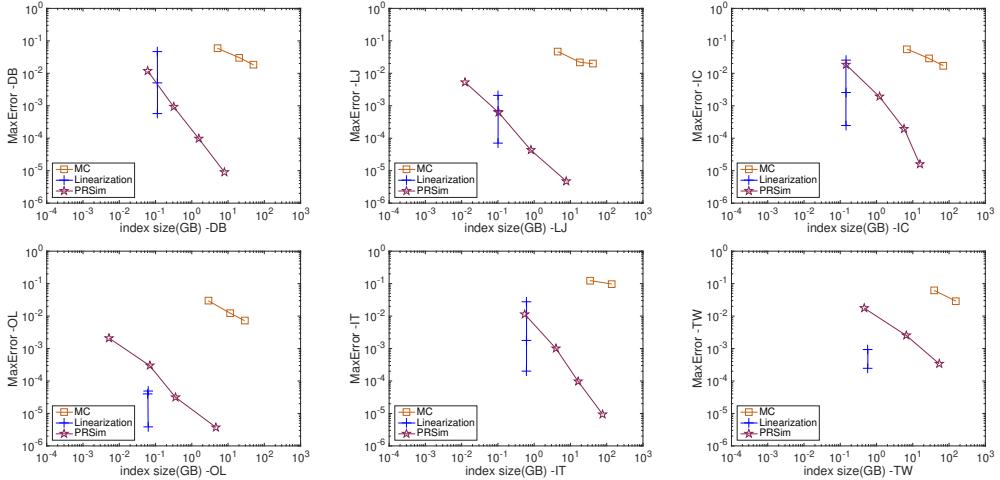


Fig. 8 MaxError v.s. Index size on large graphs

dataset, we generate 50 query nodes for each dataset. For each set of parameters and each method, we issue one query from each query node and report the average *MaxError* and *Precision@500* among the 50 query nodes.

Experiments on small graphs. We first evaluate ExactSim against other single-source algorithms on six small graphs. We compute the ground truths of the single-source and top- k queries using Power Method [13]. We omit a method if its query or preprocessing time exceeds 24 hours.

Fig. 1 shows the tradeoffs between *MaxError* and the query time of each algorithm. The first observation is that ExactSim is the only algorithm that consistently achieves an error of 10^{-7} within 10^4 seconds. Linearization is able to achieve a faster query time when the error parameter ε is large. However, as we set $\varepsilon \leq 10^{-5}$, Linearization is troubled by its $O\left(\frac{n \log n}{\varepsilon^2}\right)$ preprocess-

ing time and is unable to finish the computation of the diagonal matrix D in 24 hours.

Fig. 2 presents the tradeoffs between *Precision@500* and query time of each algorithm. We observe that ExactSim with $\varepsilon = 10^{-7}$ is able to achieve a precision of 1 on all six graphs. This confirms the exactness of ExactSim. We also note that ParSim is able to achieve high precisions on most of graphs despite its large *MaxError* in Fig. 1. This observation demonstrates the effectiveness of the $D \sim (1 - c)I$ approximation on small datasets. Finally, for the index-based methods MC, PRSim, and Linearization, we also plot the tradeoffs between *MaxError* and preprocessing time/index size in Fig. 3 and Fig. 4. The index sizes of Linearization form a vertical line, as the algorithm only recomputes and stores a diagonal matrix D . PRSim generally achieves the smallest error given a fixed amount of preprocessing time and index size.

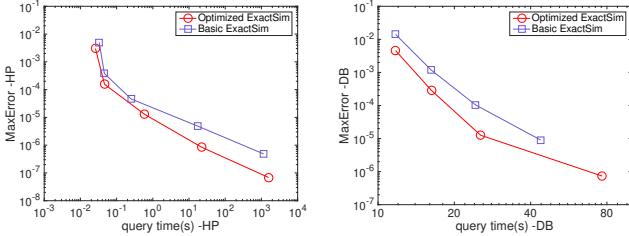


Fig. 9 Basic ExactSim v.s. Optimized ExactSim

Table 6 Memory overhead on large graphs.

Memory overhead (GB)	DB	IC	IT	TW
Basic ExactSim	2.49	3.40	18.95	19.12
Optimized ExactSim	0.47	0.58	3.26	3.54
Graph size (GB)	0.48	1.88	10.94	13.30

Experiments on large graphs. For now, we have both theoretical and experimental evidence that ExactSim is able to obtain the exact single-source and top- k SimRank results. In this section, we will treat the results computed by ExactSim with $\varepsilon = 10^{-7}$ as the ground truths to evaluate the performance of ExactSim with larger ε on large graphs.

Fig. 5 and Fig. 6 show the trade-offs between the query time and *MaxError/Precision@500* of each algorithm. Fig. 7 and Fig. 8 display the *MaxError* and preprocessing time/index size plots of the index-based algorithms. For ExactSim with $\varepsilon = 10^{-7}$, we set its *MaxError* as 10^{-7} and *Precision@500* as 1. We observe from Fig. 6 that ExactSim with $\varepsilon = 10^{-6}$ also achieves a precision of 1 on all four graphs. This suggests that the top-500 results of ExactSim with $\varepsilon = 10^{-6}$ are the same as that of ExactSim with $\varepsilon = 10^{-7}$. In other words, the top-500 results of ExactSim actually *converge* after $\varepsilon = 10^{-6}$. This is another strong evidence of the exact nature of ExactSim. From Fig. 5, we also observe that ExactSim is the only algorithm that achieves an error of less than 10^{-6} on all six large graphs. In particular, on the TW dataset, no existing algorithm can achieve an error of less than 10^{-4} , while ExactSim is able to achieve exactness within 10^4 seconds.

Ablation study. We now evaluate the effectiveness of the optimization techniques. Recall that we use sampling according to $\vec{\pi}_i(k)^2$ and local deterministic exploitation to reduce the query time, and sparse Linearization to reduce the space overhead. Fig. 9 shows the time/error tradeoffs of the basic ExactSim and the optimized ExactSim algorithms. Under similar actual error, we observe a speedup of 10 – 100 times. Table 6 shows the memory overhead of the basic ExactSim and the optimized ExactSim algorithms. We observe that the space overhead of the basic ExactSim algorithm is

usually larger than the graph size, while sparse Linearization reduces the memory usage by a factor of 3–5 times. This demonstrates the effectiveness of our optimizing techniques.

5.2 Benchmarking Approximate SimRank Algorithms

We have proved the effectiveness of ExactSim on both small and large graphs against the state-of-the-art methods in each category. In the following, we will use the ground truths computed by ExactSim to evaluate the performances of existing single-source SimRank algorithms. To the best of our knowledge, this is the first experimental study on the accuracy/cost tradeoffs of SimRank algorithms on large graphs.

Methods. Recall that in Section 2, we present a detailed analysis about all existing single-source SimRank algorithms which can support large graphs. Because Uniwalk only supports undirected graphs, we omit it methods in our evaluation and consider the other nine single-source algorithms, including three Monte Carlo methods: MC [9], READS [14] and TSF [32], two iterative methods: Linearization [29] and ParSim [45], and four Local push/sampling methods: ProbeSim [24], PRSim [39], SLING [34] and TopSim [17]. Among them, ProbeSim and ParSim are index-free methods, and the others are index-based methods; READS, TSF, ProbeSim, TopSim and ParSim can handle *dynamic* graphs, and the other methods can only handle *static* graphs. For the fairness of evaluation, we conduct each method in the single thread mode.

Experiments on Real-World Graphs. We first evaluate the performance of each method on real-world graphs. The parameters of MC, ParSim, Linearization and PRSim are the same as that in Section 5.1. Besides, READS has two parameters: the length of each random walk L and the number of random walks per node r . To cope with its better optimization, we vary (L, r) in larger ranges, from $(10^2, 10^3)$ to $(10^6, 10^7)$ on small graphs and from $(10, 100)$ to $(500, 5000)$ on large graphs. TSF has three parameters R_g , R_q and T , where R_g is the number of one-way graphs, R_q is the number of samples at query time and T is the number of iterations/steps. We vary (R_g, R_q, T) from $(100, 20, 10)$ to $(10000, 2000, 1000)$ on small graphs and from $(100, 20, 10)$ to $(4000, 800, 400)$ on large graphs. TopSim has four parameters T, h, η , and H , which correspond to the maximum length of a random walk, the lower bound of the degree to identify a high degree node, the probability threshold to eliminate a path, and the size of priority pool, respectively. As advised in paper [17], we fix $1/h = 100$ and $\eta = 0.001$ and vary (T, H) from $(3, 100)$

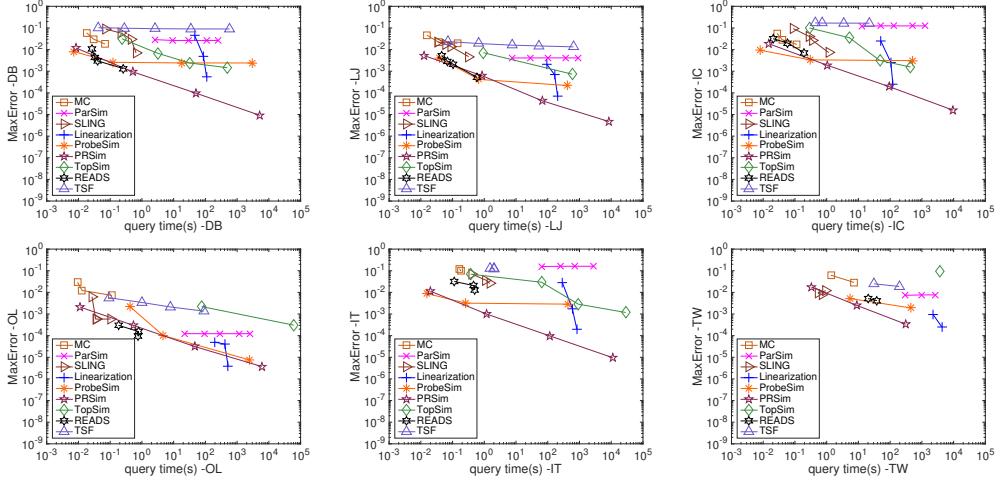


Fig. 10 Trade-offs: MaxError v.s. Query time on large graphs

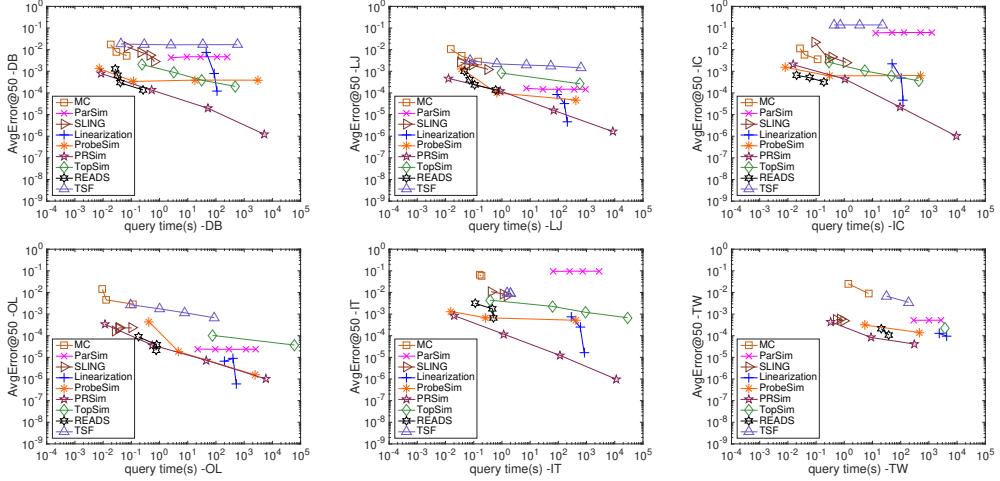


Fig. 11 Trade-offs: AvgError@50 v.s. Query time on large graphs

to $(20, 10^9)$ on small graphs and from $(3, 100)$ to $(7, 10^6)$ on large graphs. ProbeSim and SLING share the same error parameter ϵ , and we vary ϵ from 10^{-1} to 10^{-7} (if possible) on both small and large graphs.

Fig. 10 to Fig. 15 present the benchmarking studies of existing single-source algorithms against the ground truths. Specifically, Fig. 10 plots the trade-offs between query time and *MaxError*. Fig. 11 shows the trade-off lines between query time and *AvgError@50*, where

$$\text{AvgError}@k = \frac{1}{k} \sum_{v_j \in V_k} |\hat{S}(i, j) - S(i, j)|,$$

where V_k denotes the set of approximate top-k nodes. Fig. 12 draws the trade-off plots between query time and *Precision@500*. Fig. 13 shows the relations between memory cost and *MaxError*. Besides, as for those index-based methods, Fig. 14 and Fig. 15 plot the trade-offs between preprocessing time/index size and *MaxError*, respectively.

From these experimental results, we can derive the following observations. First of all, PRSim generally provides the best overall performance among the index-based methods in terms of query-time/error tradeoffs. This suggests that the local push/sampling approach is more suitable for large graphs. Secondly, the two recent dynamic methods, ProbeSim and READS, achieve similar accuracy on large graphs for the typical query time range (< 10 seconds) of the approximate algorithms. However, ProbeSim is an index-free algorithm and thus has better scalability. In particular, READS runs out of memory on the TW dataset with the number of samples per node $r > 1000$. Thirdly, ParSim is unable to achieve the same high precisions as it does on small graphs, which suggests that the $D \sim (1 - c)I$ approximation is not as effective on large graphs. SLING and Linearization also quickly become unbearable on large graphs due to their $O\left(\frac{n \log n}{\epsilon^2}\right)$ preprocessing time. This

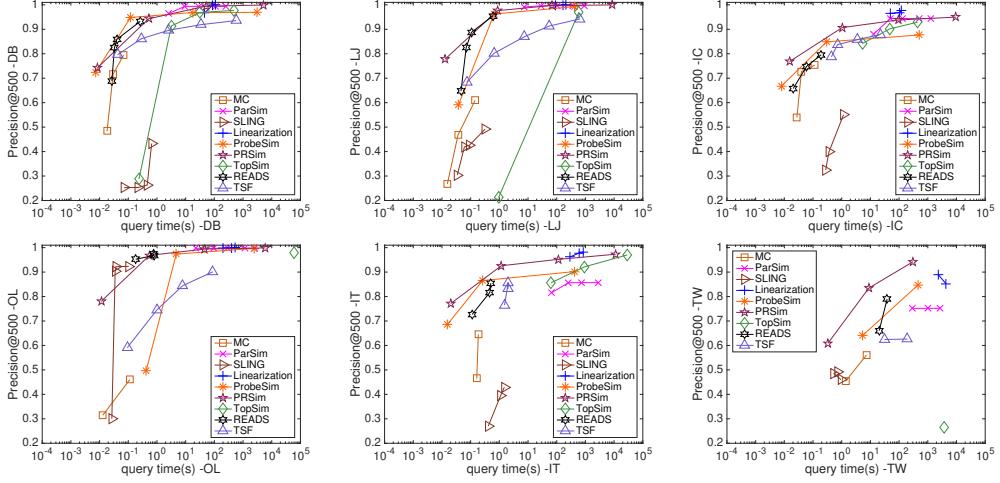


Fig. 12 Trade-offs: Precision@500 v.s. Query time on large graphs

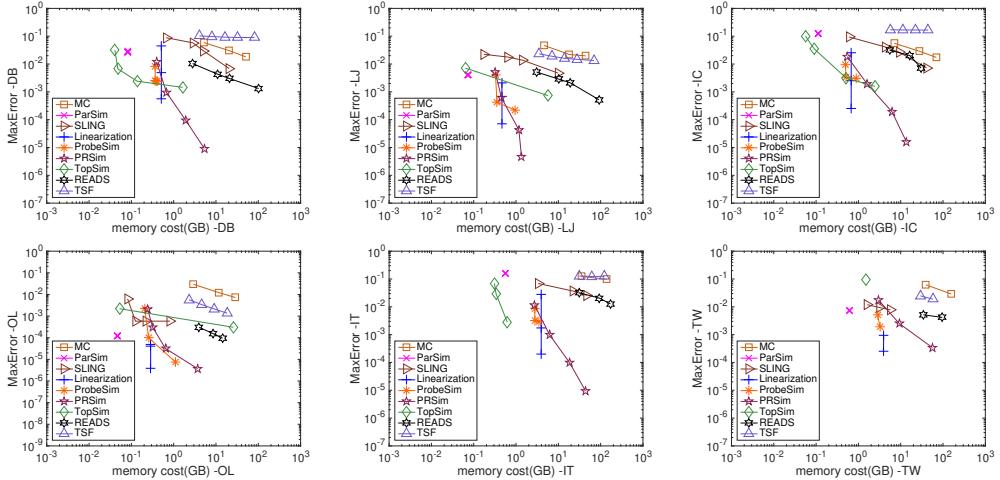


Fig. 13 Trade-offs: MaxError v.s. Memory Cost on large graphs

shows the necessity of evaluating the accuracy on large graphs. Finally, Fig. 13 shows iterative methods (ParSim and Linearization) perform the best in terms of space overhead.

Besides, we evaluate sensitivity of each method to the choice of k as for the $Precision@k$. Fig. 16 shows the precision plots with varying k from 10 to 1000 on DB and TW datasets. For each method, we only pick one group of parameters to view the change of $Precision@k$. For fairness, we try to keep each method staying in the same level of precision by appropriate parameter settings. In detail, we set $L = 20, r = 200$ for MC; $L = 5$ for ParSim; $L = 100, r = 10$ for READS; $R_g = 100, R_q = 20, T = 10$ for TSF; $T = 4$ and $H = 1000$ for TopSim; $\epsilon = 0.1$ for Linearization, PRSim, ProbeSim, and SLING. We observe that larger k always leads to low precisions. The only exception is ParSim on TW, which shows a slightly increment with larger k . This

reflects that ParSim can maintain the relative order of top- k nodes well.

Experiments on Synthetic Datasets. We also analyze the trade-off of each method with fixed parameters on synthetic datasets to vary network structures and sizes. For fairness, we choose the parameters to guarantee the accuracy of each method remains in the same level. In particular, we set $L = 50$ and $r = 500$ for MC; $L = 500$ for ParSim; $L = 10$ and $r = 100$ for READS; $R_g = 100, R_q = 20$ and $T = 10$ for TSF; $T = 3$ and $H = 100$ for TopSim; $\epsilon = 0.1$ for Linearization, PRSim, ProbeSim and SLING. On each dataset, we also generate 50 query nodes for each dataset. For each set of parameters and each method, we issue one query from each query node and report the average $MaxError$ and $Precision@500$ among the 50 query nodes.

We first evaluate the performance of each method on power-law graphs. Using the hyperbolic graph gen-

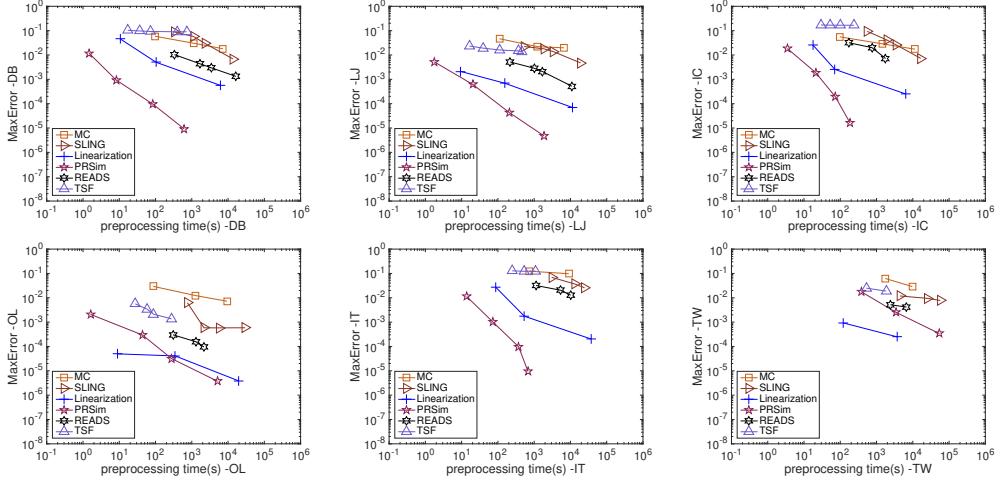


Fig. 14 Trade-offs: MaxError v.s. Preprocessing time on large graphs

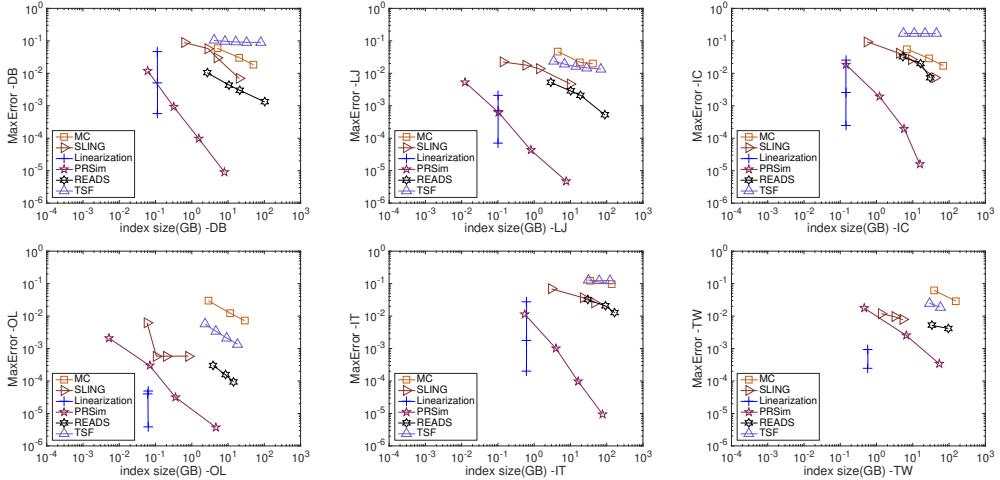


Fig. 15 Trade-offs: MaxError v.s. Index size on large graphs

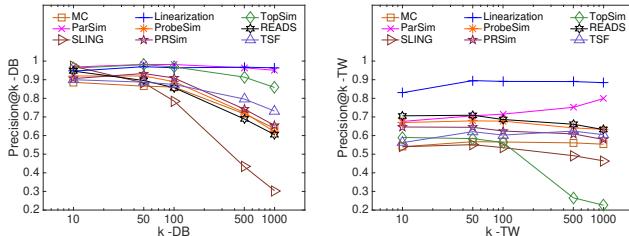


Fig. 16 Precision@k v.s. k

erator given in [4, 15], we generate a set of graphs with various power-law exponent γ , graph size n and average degree \bar{d} . We fix the graph size $n = 100,000$, the average degree $\bar{d} = 10$ and vary γ from 2.0 to 3.0. Fig. 17(a) reports the query time of each γ . From Fig. 17(a), we observe that the query time of most of methods increase with $1/\gamma$ except for Linearization, ParSim and SLING. For Linearization and ParSim, in the query phase, the two iterative methods repeat to do matrix multiplications with fixed times, leading to the unchanged query

time. As for SLING, it heavily relies on the index and its query time with large ε is too short to be impacted by γ . In Fig. 17(b), we fix $\gamma = 3$ and $\bar{d} = 10$, and vary n from 10^4 to 10^7 to evaluate the trade-offs between query time and the graph size n . We observe that local push/sampling methods' scalabilities outperform other methods in general. This is because these methods mainly focus on local information and are less influenced by the graph size. For Fig. 17(c), we try to explore the performance of each method on the power-law graphs with different average degrees. Specifically, we fix $\gamma = 3$ and $n = 100,000$, and vary \bar{d} from 5 to 1,000. We observe that the query time of PRSim increases at the slowest speed among these methods. This reveals the ability of PRSim to support dense graphs. On the contrary, TopSim shows a rapidly growing query time as the average degree increases.

Besides, we use Erdős and Rényi (ER) model to generate non-power-law graphs for evaluations. According

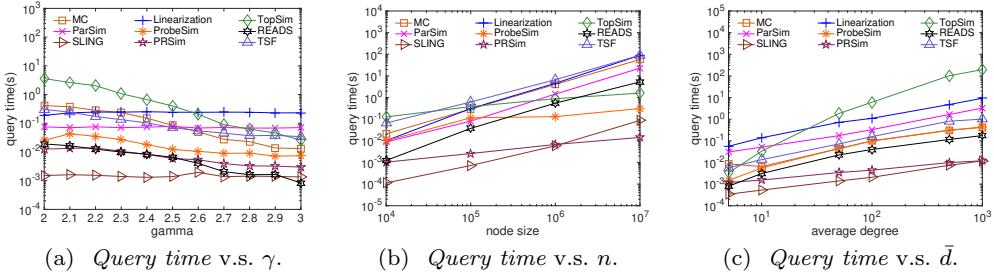


Fig. 17 Results on power-law graphs

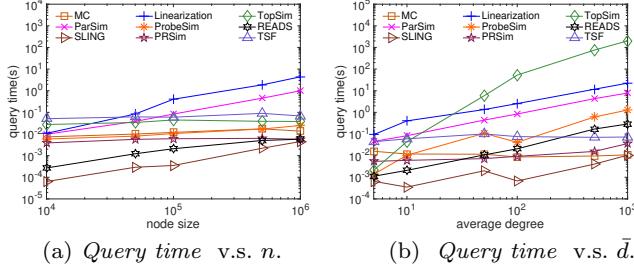


Fig. 18 Results on non-power-law graphs

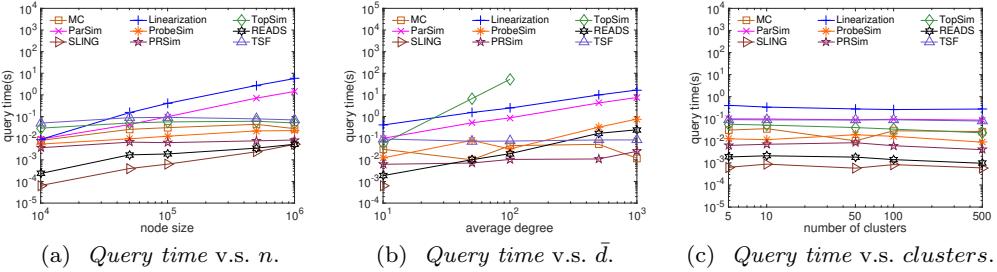


Fig. 19 Results on stochastic block graphs

to ER model, any pair of node will be assigned an edge with a specified probability p . In Fig. 18(a), we vary the graph size n from 10^4 to 10^6 . We adjust the probability p to fix the average degree $\bar{d} = 10$. In Fig. 18(b), we vary \bar{d} from 5 to 10^3 with fixed $n = 100,000$. Because by fixing the average degree, the structures of ER graphs nearly remain unchanged with the increment of n . As shown in Fig. 18(a), the query time of MC-based methods (MC,READS and TSF) does not increase with n on the ER graphs. However, we observe that the query time of the three methods show obvious increments on power-law graphs. We attribute this difference to the existence of the hub nodes on power-law graphs.

Finally, we generate graphs using the stochastic block model with four parameters, including the graph size n , the number of clusters c , the probability p to assign an edge for any pair of node belonging to the same cluster and the probability q to assign an edge for any two nodes belonging to different clusters. In Fig. 19(a), we modulate the values of p and q to keep the average degree $\bar{d} = 10$, the number of clusters $c = 5$, and vary the graph size n from 10^4 to 10^6 . In Fig. 19(b), we fix

$n = 10^5$ and $c = 5$, adjust p and q to vary the average degree \bar{d} from 10 to 1000. In Fig. 19(c), we vary the number of clusters c from 5 to 500 and fix $n = 10^5$, $\bar{d} = 10$. We observe that the result of each method is similar with that on ER graphs, which reflects that stochastic block model is a generalized version of ER model. Fig. 19(c) shows that the number of clusters does not have a significant effect on the query time of these methods.

Experiments on Dynamic Datasets. In this section, we evaluate the performances of the methods which can support dynamic graphs. Recall that ParSim [45], ProbeSim [24], and TopSim [17] are index-free methods and can support dynamic graphs naturally. READS [14] and TSF [32] are two index-based methods which can support dynamic graphs by modifying index structures. Since the vertex modification can be treated as several edge modifications, we use the two dynamic graphs WD and WP which only contains edge modifications for ease of readability. The parameters of each method are the same with that in Section 5.2. For the four

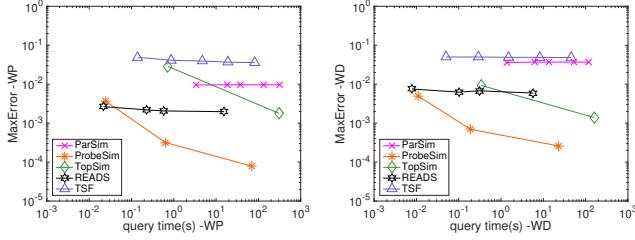


Fig. 20 MaxError v.s. Query time on dynamic graphs

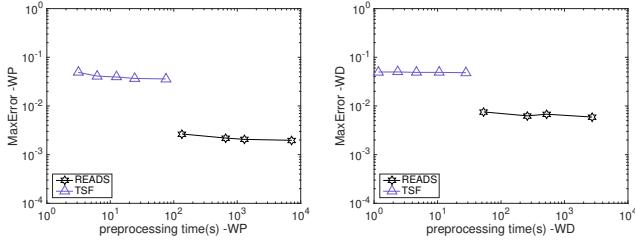


Fig. 22 MaxError v.s. Preprocessing time on dynamic graphs

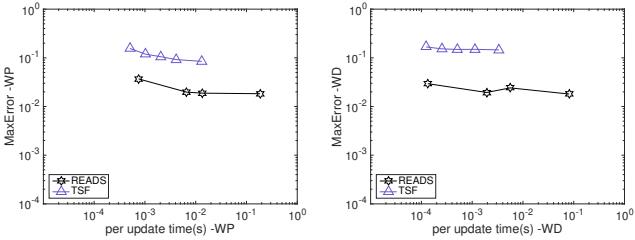


Fig. 24 MaxError v.s. Update Time per edge insertion/deletion on dynamic graphs

index-free methods, we run them on the final graphs of WP and WD. For READS and TSF, we first load the initial graph without the last 10,000 edge modifications and construct the index. Then we run the two methods on the dynamic graphs with 10,000 edge modifications. After the updating process, we compare the computational quality of the six methods and plot their trade-offs between the query time and MaxError/Precision@500 in Fig. 20 and Fig. 21, respectively. In Fig. 20, we observe that each method's performance is similar with that on static graphs. ProbeSim achieves the highest approximation quality within the same query time. We observe that the performances of index-free methods are similar with that on static graphs. ProbeSim still shows the best performance among these methods. However, the *MaxError* of READS is hard to be reduced with increasing query time. This is very different from what we have observed on static graphs, where READS and ProbeSim achieve similar accuracy. In Fig. 22, 23, and 24, we plot the trade-offs between MaxError and preprocessing time/index size/updating time of the two index-based methods TSF and READS. Note that the updating time is the average time per

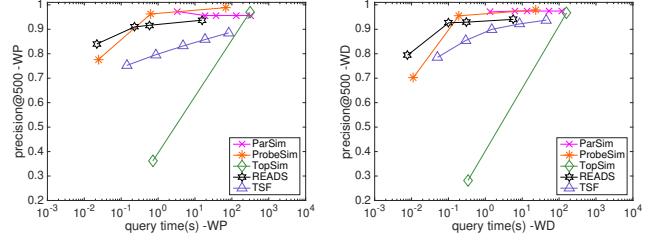


Fig. 21 Precision@500 v.s. Query time on dynamic graphs

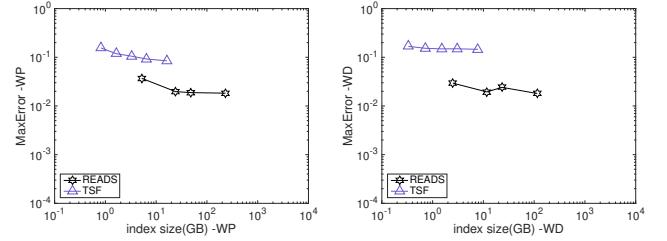


Fig. 23 MaxError v.s. Index size on dynamic graphs

edge insertion/deletion in the updating process. We observe that the two methods both incur large maximum error. READS shows a better performance than TSF.

5.3 SimRank Distribution

We now design experiments to seek the answers for two open questions regarding the distribution of SimRank:

- Does the single-source SimRank result follow the power-law distribution on real-world graphs?
- What is the density of SimRank values on real-world graphs?

We use ExactSim to compute the ground truths of 50 random single-source queries on each of the six large graphs. Then we compute the average frequency of SimRank values in every range of length 10^{-5} , and plot these frequencies against the SimRank values in Fig. 25. Besides, we plot the frequency distribution of Personalized PageRank(PPR) computed by its Power Method [31] with teleport probability $\alpha = 0.2$, which has been proved following the power law [23]. The results suggest SimRank values indeed exhibit a power-law shaped distribution on real-world graphs as PPR does. In particular, the power-law exponent (slope) on TW appears to be significantly more skewed than that on IT, which explains why TW is a harder dataset for computing single-source SimRank queries. For sake of completeness, we also plot the degree distributions of the six graphs in Fig. 26. We compute the average frequency of each degree in every range of length 10. We observe that the largest degree can achieve 10^6 on TW, which is apparently larger than other datasets. This

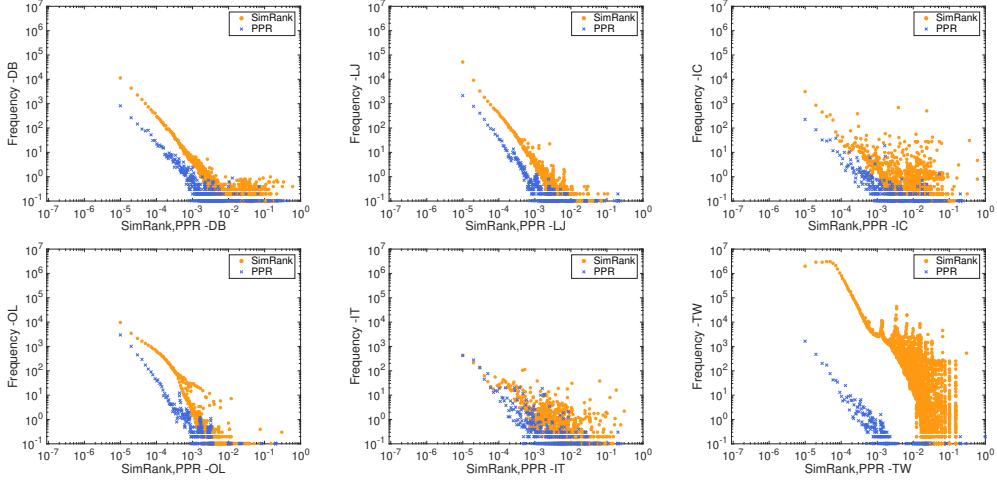


Fig. 25 The distribution of SimRank and PPR on real-world graphs

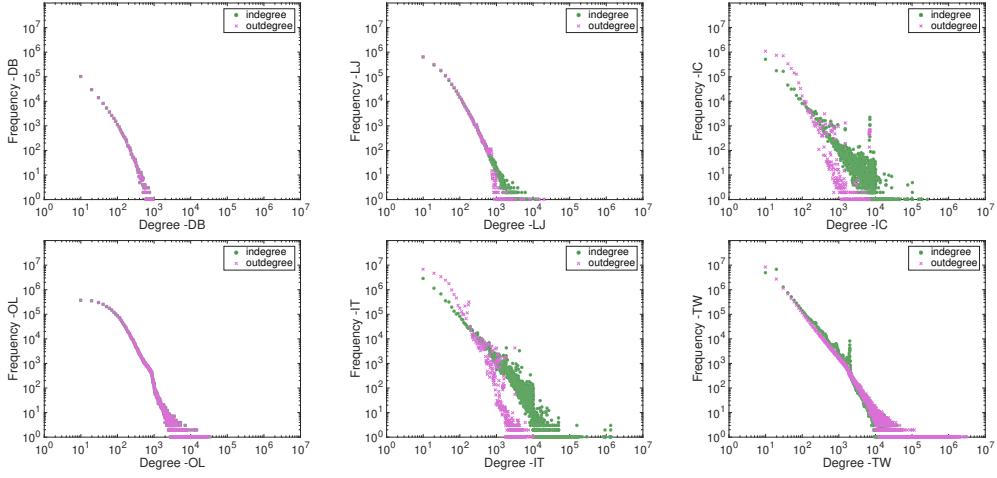


Fig. 26 Degree distribution of real-world graphs

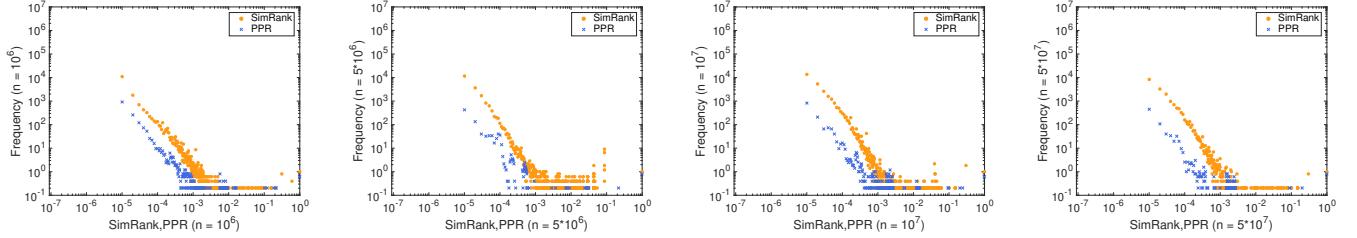


Fig. 27 SimRank and PPR distribution of the Kronecker graphs (varying n from 10^6 to 5×10^7)

also demonstrates the hardness to compute SimRank on TW.

Besides, we plot the SimRank distribution on synthetic power-law graphs in Fig. 27 using the Kronecker graph model [18], which can generate large graphs of million nodes. We fix the probability seed matrix as $(0.9, 0.5; 0.5, 0.1)$ and vary the graph size n from 10^6 to 5×10^7 . On the four synthetic graphs, SimRank values still exhibit a power-law shaped distribution. We also plot the degree distribution of the four synthetic

power-law graphs in Fig. 28. The degree distribution of the four synthetic graphs are all power-law shaped.

In comparison, we generate non-power-law graphs using the Erdős and Rényi(ER) model, and show the SimRank distributions on the synthetic non-power-law graphs. According to the settings of ER-model, an edge is attached to each node with a user-defined probability p . We vary the number of nodes n from 10^4 to 5×10^5 , and tune p to guarantee the average degree $d = 10$. Fig. 29 plots the SimRank and PPR distributions.

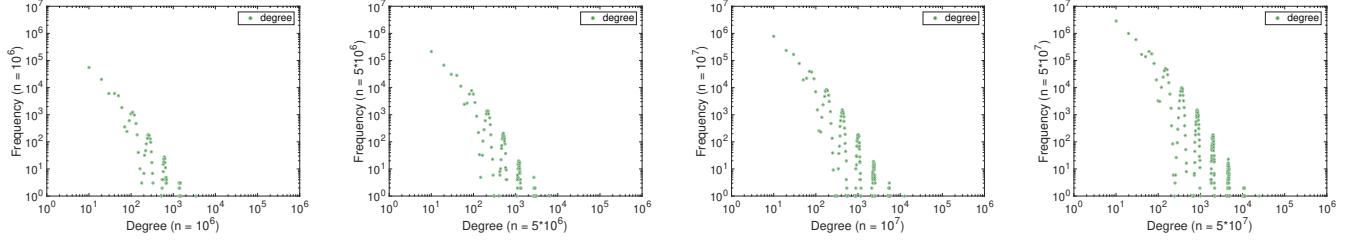


Fig. 28 Degree distribution of the Kronecker graphs (varying n from 10^6 to 5×10^7)

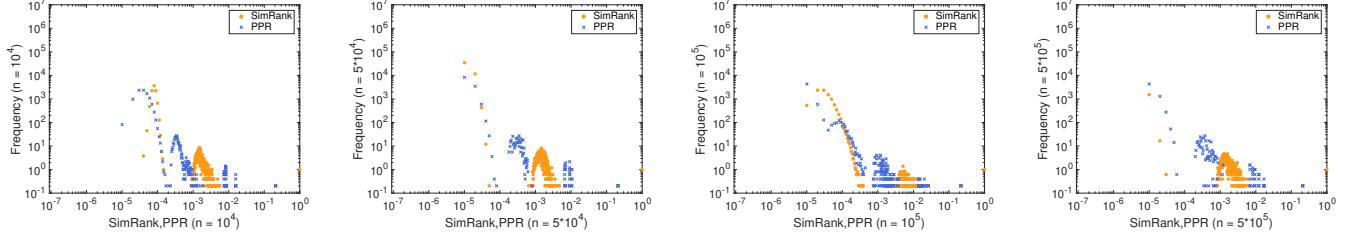


Fig. 29 The distribution of SimRank and PPR on E-R graphs (varying n from 10^4 to 5×10^5)

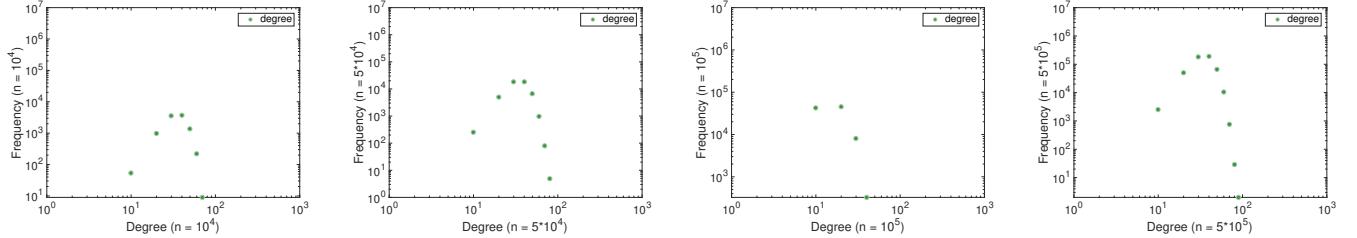


Fig. 30 Degree distribution on E-R graphs (varying n from 10^4 to 5×10^5)

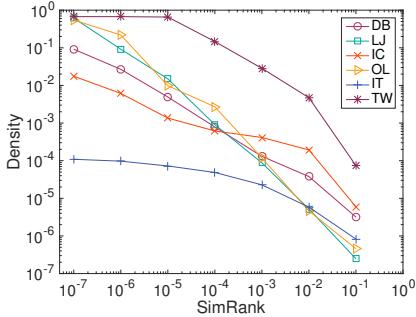


Fig. 31 SimRank density on large graphs

Fig. 30 displays the degree distributions on these ER graphs. We observe that the distributions of SimRank and PPR both show non-power-law shaped curves on ER graphs.

Next, we analyze the density of single-source SimRank queries. The density of SimRank is the percentage of SimRank values that are larger than some threshold ε . Fig. 31 shows the average density of 50 queries on six large datasets, with ε varying from 0.1 to 10^{-7} . The result shows that the densities can vary widely on different datasets. For example, on the TW dataset, the density of SimRank values quickly reaches close to 1 for $\varepsilon < 10^{-4}$. On the other hand, the density on the

IT dataset seems to converge on 10^{-4} . This suggests that density-sensitive methods such as [38] can achieve satisfying results on IT and may run out of memory on dense graphs such as TW. This result also implies that it is essentially hopeless to design an exact algorithm for all-pair queries on large real-world graphs, as the number of non-zero entries in the SimRank matrix can be as large as $O(n^2)$.

6 Conclusions

This paper presents ExactSim, an algorithm that produces the ground truths for single-source and top- k SimRank queries with precision up to 7 decimal places on large graphs. Using the ground truths computed by ExactSim, we present the first experimental study of the accuracy/cost tradeoffs of existing SimRank algorithms on large graphs. We also exploit various properties of the distributions of SimRank on large real-world graphs. For future work, we note that the $O\left(\frac{\log n}{\varepsilon^2}\right)$ complexity of ExactSim prevents it from achieving a precision of 10^{-14} (i.e., the precision of the double type). To achieve such extreme precision, we need an algorithm with $O\left(\frac{\log n}{\varepsilon}\right)$ complexity, which remains a major open problem in SimRank study.

7 ACKNOWLEDGEMENTS

Zhewei Wei was supported by National Natural Science Foundation of China (NSFC) No. 61972401 and No. 61932001, by the Fundamental Research Funds for the Central Universities and the Research Funds of Renmin University of China under Grant 18XNLG21, and by Alibaba Group through Alibaba Innovative Research Program. The work is partially done at Beijing Key Laboratory of Big Data Management and Analysis Methods, MOE Key Lab DEKE, Renmin University of China, and Pazhou Lab, Guangzhou, 510330, China. Hanzhi Wang was supported by the Outstanding Innovative Talents Cultivation Funded Programs 2020 of Renmin University of China. Ye Yuan was supported by NSFC No. 61932004 and No. 61622202, and by FRFCU No. N181605012. Ji-Rong Wen was supported by NSFC No. 61832017, and by Beijing Outstanding Young Scientist Program NO. BJJWZYJH012019100020098. Xiaoyong Du was supported by NSFC No. U1711261.

References

1. <http://snap.stanford.edu/data>.
2. <http://law.di.unimi.it/datasets.php>.
3. <http://konect.uni-koblenz.de/networks>.
4. Rodrigo Aldecoa, Chiara Orsini, and Dmitri Krioukov. Hyperbolic graph generator. *Computer Physics Communications*, 196:492–496, 2015.
5. Reid Andersen, Fan R. K. Chung, and Kevin J. Lang. Local graph partitioning using pagerank vectors. In *FOCS*, pages 475–486, 2006.
6. Ioannis Antonellis, Hector Garcia Molina, and Chi Chao Chang. Simrank++: query rewriting through link analysis of the click graph. *VLDB*, 1(1):408–421, 2008.
7. Bahman Bahmani, Abdur Chowdhury, and Ashish Goel. Fast incremental and personalized pagerank. *VLDB*, 4(3):173–184, 2010.
8. Fan R. K. Chung and Lincoln Lu. Concentration inequalities and martingale inequalities: A survey. *Internet Mathematics*, 3(1):79–127, 2006.
9. Daniel Fogaras and Balazs Racz. Scaling link-based similarity search. In *WWW*, pages 641–650, 2005.
10. Dániel Fogaras, Balázs Rácz, Károly Csalogány, and Tamás Sarlós. Towards scaling fully personalized pagerank: Algorithms, lower bounds, and experiments. *Internet Mathematics*, 2(3):333–358, 2005.
11. Yuichiro Fujiwara, Makoto Nakatsuji, Hiroaki Shiokawa, and Makoto Onizuka. Efficient search algorithm for simrank. In *ICDE*, pages 589–600, 2013.
12. Guoming He, Haijun Feng, Cuiping Li, and Hong Chen. Parallel simrank computation on large graphs with iterative aggregation. In *KDD*, pages 543–552, 2010.
13. Glen Jeh and Jennifer Widom. Simrank: a measure of structural-context similarity. In *SIGKDD*, pages 538–543, 2002.
14. Minhao Jiang, Ada Wai-Chee Fu, and Raymond Chi-Wing Wong. Reads: a random walk approach for efficient and accurate dynamic simrank. *PPVLDB*, 10(9):937–948, 2017.
15. Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguná. Hyperbolic geometry of complex networks. *Physical Review E*, 82(3):036106, 2010.
16. Mitsuuru Kusumoto, Takanori Maehara, and Ken-ichi Kawarabayashi. Scalable similarity search for simrank. In *SIGMOD*, pages 325–336, 2014.
17. Pei Lee, Laks V. S. Lakshmanan, and Jeffrey Xu Yu. On top-k structural similarity search. In *ICDE*, pages 774–785, 2012.
18. Jure Leskovec, Deepayan Chakrabarti, Jon Kleinberg, Christos Faloutsos, and Zoubin Ghahramani. Kronecker graphs: an approach to modeling networks. *Journal of Machine Learning Research*, 11(2), 2010.
19. Cuiping Li, Jiawei Han, Guoming He, Xin Jin, Yizhou Sun, Yintao Yu, and Tianyi Wu. Fast computation of simrank for static and dynamic information networks. In *EDBT*, pages 465–476, 2010.
20. Lina Li, Cuiping Li, Hong Chen, and Xiaoyong Du. Mapreduce-based simrank computation and its application in social recommender system. In *2013 IEEE international congress on big data*, pages 133–140. IEEE, 2013.
21. Zhenguo Li, Yixiang Fang, Qin Liu, Jiefeng Cheng, Reynold Cheng, and John Lui. Walking in the cloud: Parallel simrank at scale. *VLDB*, 9(1):24–35, 2015.
22. Zhenjiang Lin, Michael R Lyu, and Irwin King. Matchsim: a novel similarity measure based on maximum neighborhood matching. *KAIS*, 32(1):141–166, 2012.
23. Nelly Litvak, Werner RW Scheinhardt, and Yana Volkovich. In-degree and pagerank: why do they follow similar power laws? *Internet mathematics*, 4(2-3):175–198, 2007.
24. Yu Liu, Bolong Zheng, Xiaodong He, Zhewei Wei, Xiaokui Xiao, Kai Zheng, and Jiaheng Lu. Probesim: scalable single-source and top-k simrank computations on dynamic graphs. *VLDB*, 11(1):14–26, 2017.
25. Dmitry Lizorkin, Pavel Velikhov, Maxim Grinev, and Denis Turdakov. Accuracy estimate and optimization techniques for simrank computation. *The VLDB Journal*, 19(1):45–66, 2010.
26. Dmitry Lizorkin, Pavel Velikhov, Maxim N. Grinev, and Denis Turdakov. Accuracy estimate and optimization techniques for simrank computation. *VLDB J.*, 19(1):45–66, 2010.
27. Linyuan Lü and Tao Zhou. Link prediction in complex networks: A survey. *Physica A: statistical mechanics and its applications*, 390(6):1150–1170, 2011.
28. X. Luo, J. Gao, C. Zhou, and J. X. Yu. Uniwalk: Unidirectional random walk based scalable simrank computation over large graph. In *2017 IEEE 33rd International Conference on Data Engineering (ICDE)*, pages 325–336, April 2017.
29. Takanori Maehara, Mitsuuru Kusumoto, and Ken-ichi Kawarabayashi. Efficient simrank computation via linearization. *CoRR*, abs/1411.7228, 2014.
30. Takanori Maehara, Mitsuuru Kusumoto, and Ken-ichi Kawarabayashi. Scalable simrank join algorithm. In *ICDE*, pages 603–614, 2015.
31. Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: bringing order to the web. 1999.
32. Yingxia Shao, Bin Cui, Lei Chen, Mingming Liu, and Xing Xie. An efficient similarity search framework for simrank over large dynamic graphs. *VLDB*, 8(8):838–849, 2015.

33. Wenbo Tao, Minghe Yu, and Guoliang Li. Efficient top-k simrank-based similarity join. *PVLDB*, 8(3):317–328, 2014.
34. Boyu Tian and Xiaokui Xiao. SLING: A near-optimal index structure for simrank. In *SIGMOD*, pages 1859–1874, 2016.
35. Anton Tsitsulin, Davide Mottin, Panagiotis Karras, and Emmanuel Müller. Verse: Versatile graph embeddings from similarity measures. In *WWW*, pages 539–548. International World Wide Web Conferences Steering Committee, 2018.
36. Hanzhi Wang, Zhewei Wei, Ye Yuan, Xiaoyong Du, and Ji-Rong Wen. Exact single-source simrank computation on large graphs. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, pages 653–663, 2020.
37. Yue Wang, Yulin Che, Xiang Lian, Lei Chen, and Qiong Luo. Fast and accurate simrank computation via forward local push and its parallelization. *IEEE Transactions on Knowledge and Data Engineering*, 2020.
38. Yue Wang, Lei Chen, Yulin Che, and Qiong Luo. Accelerating pairwise simrank estimation over static and dynamic graphs. *The VLDB Journal*, 28(1):99–122, February 2019.
39. Zhewei Wei, Xiaodong He, Xiaokui Xiao, Sibo Wang, Yu Liu, Xiaoyong Du, and Ji-Rong Wen. Prsim: Sublinear time simrank computation on large power-law graphs. In *SIGMOD*, pages 1042–1059. ACM, 2019.
40. Wensi Xi, Edward A Fox, Weiguo Fan, Benyu Zhang, Zheng Chen, Jun Yan, and Dong Zhuang. Simfusion: measuring similarity using unified relationship matrix. In *SIGIR*, pages 130–137. ACM, 2005.
41. Weiren Yu, Xuemin Lin, and Wenjie Zhang. Fast incremental simrank on link-evolving graphs. In *ICDE*, pages 304–315, 2014.
42. Weiren Yu, Xuemin Lin, Wenjie Zhang, Lijun Chang, and Jian Pei. More is simpler: Effectively and efficiently assessing node-pair similarities based on hyperlinks. *PVLDB*, 7(1):13–24, 2013.
43. Weiren Yu and Julie McCann. Gauging correct relative rankings for similarity search. In *CIKM*, pages 1791–1794, 2015.
44. Weiren Yu and Julie A. McCann. Efficient partial-pairs simrank search for large networks. *PVLDB*, 8(5):569–580, 2015.
45. Weiren Yu and Julie A McCann. Efficient partial-pairs simrank search on large networks. *Proceedings of the VLDB Endowment*, 8(5):569–580, 2015.
46. Weiren Yu and Julie Ann McCann. High quality graph-based similarity search. In *SIGIR*, pages 83–92, 2015.
47. Weiren Yu, Wenjie Zhang, Xuemin Lin, Qing Zhang, and Jiajin Le. A space and time efficient algorithm for simrank computation. *World Wide Web*, 15(3):327–353, 2012.
48. Jing Zhang, Jie Tang, Cong Ma, Hanghang Tong, Yu Jing, and Juanzi Li. Panther: Fast top-k similarity search on large networks. In *SIGKDD*, pages 1445–1454. ACM, 2015.
49. Peixiang Zhao, Jiawei Han, and Yizhou Sun. P-rank: a comprehensive structural similarity measure over information networks. In *CIKM*, pages 553–562. ACM, 2009.
50. Peixiang Zhao, Jiawei Han, and Yizhou Sun. P-rank: a comprehensive structural similarity measure over information networks. In *CIKM*, pages 553–562, 2009.
51. Weiguo Zheng, Lei Zou, Yansong Feng, Lei Chen, and Dongyan Zhao. Efficient simrank-based similarity join over large graphs. *PVLDB*, 6(7):493–504, 2013.