



Efficient and Accurate SimRank-based Similarity Joins: Experiments, Analysis, and Improvement

Qian Ge[†]
Peking University
geqian@pku.edu.cn

Yu Liu^{†§}
Yinghao Zhao
Yuetian Sun
Beijing Jiaotong University
Beijing Key Laboratory of Traffic Data Analysis and
Mining, Beijing, China
{yul,yinghaozhao,yuetiansun}@bjtu.edu.cn

Lei Zou
Peking University
zoulei@pku.edu.cn

Yuxing Chen
Anqun Pan
Tencent Inc.
{axingguchen,aaronpan}@tencent.com

ABSTRACT

SimRank-based similarity joins, which mainly include threshold-based and top- k similarity joins, are important types of all-pair SimRank queries. Although a line of related algorithms have been proposed recently, they still fall short of providing approximation guarantee and suffer from scalability issues on medium and large graphs. Meanwhile, we also lack an extensive analysis of existing techniques in terms of accuracy and efficiency. Motivated by these challenges, we first conduct detailed analysis of state-of-the-art algorithms and provide additional theoretical results. Second, to address the limitations of existing techniques, we propose simple yet effective algorithm frameworks for both queries to theoretically guarantee the approximation bound, and present a more efficient all-pair algorithm inspired by randomized local push of Personalized PageRank. Next, we analyze the algorithmic complexity of threshold-based and top- k similarity joins by leveraging a reasonable assumption of SimRank distribution. Through extensive experiments, we find that our proposed methods far exceed existing ones with respect to query efficiency, approximation guarantee and practical accuracy, while our theoretical analysis nicely matches the empirical study.

PVLDB Reference Format:

Qian Ge, Yu Liu, Yinghao Zhao, Yuetian Sun, Lei Zou, Yuxing Chen, and Anqun Pan. Efficient and Accurate SimRank-based Similarity Joins: Experiments, Analysis, and Improvement. PVLDB, 17(4): 617-629, 2023. doi:10.14778/3636218.3636219

PVLDB Artifact Availability:

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment. Proceedings of the VLDB Endowment, Vol. 17, No. 4 ISSN 2150-8097. doi:10.14778/3636218.3636219

[†]These authors contributed equally to this work.

[§]Corresponding author.

The source code, data, and/or other artifacts have been made available at <https://github.com/xinghun0525/R2LP>.

1 INTRODUCTION

SimRank [12] is one of the most important measures for pairwise node similarity, which finds its applications in recommender systems [19], link prediction [27], and graph embeddings [34]. The SimRank similarity of node u and v is defined by the following recursive equation, with the intuition that “two nodes in a graph are similar if their in-neighbors are similar, and a node is most similar to itself”:

$$s(u, v) = \begin{cases} 1, & \text{if } u = v \\ \frac{c}{|I(u)||I(v)|} \sum_{x \in I(u)} \sum_{y \in I(v)} s(x, y), & \text{otherwise.} \end{cases} \quad (1)$$

Here, $I(u)$ denote the in-neighbors of node u , and $c \in (0, 1)$ is the decay factor, which is typically set to 0.6 [33, 36, 39] or 0.8 [12]. Due to its computational complexity, efficient computation of SimRank has been extensively studied in the past two decades [8, 16, 18, 22, 24, 26, 28, 32], and remains to be a hot topic in the very recent years [23, 31, 36, 39, 45, 48].

According to the query inputs and outputs, SimRank queries can be broadly categorized as *single-pair* queries, *single-source* queries, and *all-pair* queries with the following definition.

DEFINITION 1 (ALL-PAIR SIMRANK COMPUTATION). *Given a graph $G = (V, E)$ of n nodes and an error parameter ϵ , return the SimRank estimation $\hat{s}(u, v)$ for every node pair (u, v) with absolute error guarantee, so that $|\hat{s}(u, v) - s(u, v)| \leq \epsilon$ (with high probability).*

Among existing work [20, 26, 28, 29, 32, 39, 45–48] that tackles all-pair queries, state-of-the-art algorithms [39, 48] can hardly guarantee an additive error of $\epsilon = 0.01$ on million-node graphs. Besides, a recent work [36] indicates that it is “essentially hopeless” to compute all-pair queries on large graphs, as non-zero items can be as large as $O(n^2)$. Fortunately, two modified versions of all-pair queries have been studied, namely, *threshold-based* [29, 46, 47] and *top- k* [20, 32, 48] *similarity joins*, of which the outputs are small

subsets of all-pair SimRank computation. Both queries are more practically useful because we are only interested in the node pairs with non-negligible similarities [48].

DEFINITION 2 (THRESHOLD-BASED SIMILARITY JOIN). *Given a graph $G = (V, E)$ and a threshold $\theta > 0$, return the set of node pairs $\mathcal{R}(\theta)$ with $u \neq v$ and SimRank value $s(u, v) \geq \theta$.*

DEFINITION 3 (TOP- k SIMILARITY JOIN). *Given a graph $G = (V, E)$ and a parameter $k > 0$, return a k -sized set $\mathcal{R}(k)$ containing the node pairs with the top- k largest SimRank values among all pairs $\{(u, v) | u, v \in V, u \neq v\}$.*

Analogous to all-pair SimRank computation that an error parameter ϵ is allowed, we introduce the *approximation bound* ρ for similarity joins. As we shall see in Section 4, without ρ (i.e., setting $\rho = 1$), in worst cases no algorithm can be more efficient than computing the exact SimRank values¹.

DEFINITION 4 (APPROXIMATION BOUND). *Given an approximation parameter $\rho \in (0, 1)$, we say an algorithm \mathcal{A} has approximation bound for threshold-based (resp. top- k) similarity join of SimRank, if for any input graph G , the returned set of node pairs contains at least ρ fraction of the ground truth answer.*

1.1 Motivations

We note that existing algorithms for threshold-based and top- k similarity joins significantly fall short of both query efficiency and approximation guarantee. We identify the following issues.

Motivation 1. Lack of experimental analysis of efficiency and accuracy, especially on medium and large graphs. Even for state-of-the-art algorithms [39, 48] of general all-pair queries, the result accuracy is not evaluated except for small graphs due to lack of ground truth. Other representative algorithms [20, 46] even do not evaluate accuracy on small graphs.

Motivation 2. Lack of approximation guarantee for threshold-based and top- k similarity joins. Surprisingly, there does not exist an algorithm that guarantees the approximation bound for threshold-based or top- k similarity join. Algorithms with absolute error guarantee cannot be directly extended to achieve this goal.

Motivation 3. Lack of understanding of the computational complexity. As far as we know, it is unclear how the input parameters and graph properties impact the algorithmic complexity of both queries.

1.2 Contributions

For both types of all-pair SimRank queries, we analyze state-of-the-art algorithms in detail, provide improved algorithms with approximation guarantee and lower theoretical complexity, and discuss the problem complexity. Our contributions are summarized as follows.

Detailed analysis of state-of-the-art algorithms. We choose four representative algorithms, *UISim* [48], *FLP & Opt-LP* [39], *H-go SRJ* [46], and *KSimJoin* [20] for comparison. *UISim* and *FLP & Opt-LP* are state-of-the-art algorithms for general all-pair queries, which are directly extended to answer SimRank-based similarity joins [48]. *H-go SRJ* and *KSimJoin* are the state-of-the-art algorithms

¹We distinguish between the *exact* and the *ground truth* SimRank values, where the former is the exact solution to Equation 1 and the latter is a very accurate approximation.

for threshold-based and top- k similarity joins, respectively. Apart from their performance, we choose *UISim* [48] and *KSimJoin* [20] because they adopt the random walk interpretation, whereas *Opt-LP* [39] and *H-go SRJ* [46] are the best methods based on the SimRank matrix formation. For each algorithm, we discuss its basic idea, theoretical guarantee, complexity analysis and empirical study. We propose additional theoretical results (w.r.t. complexity or error guarantee) that are incorrect or missed in the original papers.

Improved algorithms for SimRank-based similarity joins. We first propose two simple yet effective algorithm frameworks for both queries respectively with approximation bound. The basic idea is to invoke an efficient algorithm for all-pair queries that estimates SimRank values within ϵ error, and to gradually shrink ϵ until the approximation bound is theoretically guaranteed. Next, we utilize the equivalence of all-pair SimRank on the input graph G and single-target Personalized PageRank (PPR) on the *SimRank graph* G^S [39], and devise Randomized Reverse Local Push (R^2LP) inspired by randomized backward push for PPR [35]. It improves the all-pair query complexity from $O(\frac{\sum_{u,v \in V} d_{in}(u)d_{in}(v)s(u,v)}{\epsilon})$ (the best known algorithm [39]) to $\tilde{O}(\frac{\sum_{u,v \in V} \sqrt{d_{in}(u)d_{in}(v)}s(u,v)}{\epsilon})$. We further propose a pruning strategy to dramatically enhance practical efficiency while retaining error guarantee.

Complexity analysis by utilizing SimRank distribution. We find that it is almost impossible to give a non-trivial complexity bound for both similarity join queries without any assumption of SimRank distribution. Our empirical analysis validates that for most real-world graphs, the ground truth SimRank values follow some *generalized* version of power-law distribution [2, 6], which is consistent with previous studies [23, 36]. We leverage this property to propose verifiable assumptions of SimRank distribution and present complexity analysis for our algorithms, which serves as the upper bound of the algorithmic complexity of SimRank-based similarity joins. Our theoretical results nicely match the empirical study in Section 6.

Extensive experiments on medium and large graphs. We conduct extensive experiments on small, medium, and large sized graphs and compare our algorithms with state-of-the-art methods. As far as we know, we first employ medium and large datasets to systematically evaluate threshold-based and top- k similarity joins in terms of running time, absolute error and query accuracy, and to demonstrate the necessities of holding approximation bound as well as devising more efficient algorithms. Experimental study demonstrates that our algorithms outperform existing ones in most cases, achieving better accuracy while being faster by up to an order of magnitude. More importantly, our empirical findings validate that it is the skewness of SimRank distribution that mainly determines the problem hardness. This results verifies the effectiveness of our theoretical analysis.

2 PRELIMINARIES

The definition in Equation 1 can be interpreted via the *node-pairs graph* G^2 [12]. Each node in G^2 represents a pair of nodes (u, v) in G . If there exist edges (u, v) and (x, y) in G , edge $((u, x), (v, y))$ is included in G^2 . Specifically, every node (v, v) is called the *singleton node*. To this end, SimRank can be thought of as propagating the similarity among nodes in G^2 , starting from all the singleton nodes.

Table 1: Table of notations.

Notation	Description
$G = (V, E)$	A graph with node set V and edge set E , where $n = V $ and $m = E $
G^2, G^s	The node-pairs graph [12] and the SimRank graph [39]
u, v	Nodes in G
$I(u), O(u)$	The incoming and outgoing neighbors of u , and we have $d_{in}(u) = I(u) , d_{out}(u) = O(u) $
$s(u, v)$	The SimRank value, i.e., the solution of Equation 1
$s^*(u, v)$	The ground truth SimRank value, with $ s^*(u, v) - s(u, v) \leq \epsilon_{min}$ (with high probability)
$\hat{s}(u, v)$	The estimated SimRank value
ϵ, δ	The additive error and the failure probability
θ	The input parameter of threshold-based similarity join
k	The input parameter of top- k similarity join
ρ	The approximation bound

Matrix formation of SimRank. The recursive definition can be represented in matrix formation [44], i.e., $S = cP^TSP \vee I$, where S is an $n \times n$ -dimensional similarity matrix with $S[u, v] = s(u, v)$, P is the column-normalized adjacency matrix of G , I is the identity matrix, and operator \vee denotes element-wise maximum. The all-pair ground truth of SimRank can be computed by the *Power Method* [12], which is essentially a fixed-point iteration process on G^2 . Recent literature [18, 26, 39] further derives various close-form solutions based on this formation.

The random walk interpretation. Jeh and Widom explain the SimRank definition from random walk perspective in their original paper [12]. We adopt \sqrt{c} -walk, its revised version following [13, 22, 33]. Let $W_{\sqrt{c}}(u)$ denote a random walk starting from node u and following incoming edges at each step with walking probability \sqrt{c} . At each step, with probability \sqrt{c} , it randomly chooses an in-neighbor and continues the walk, otherwise the walk stops at the current node. [33] proves that the SimRank of u and v equals the probability that two \sqrt{c} -walks $W_{\sqrt{c}}(u)$ and $W_{\sqrt{c}}(v)$ meet.

Table 1 lists the frequently used notations throughout the paper.

3 ANALYSIS OF STATE OF THE ART

We present detailed comparison and analysis of state-of-the-art algorithms for threshold-based and top- k similarity joins (see Figure 1 and Table 2). We choose four methods including *UISim* [48], *FLP & Opt-LP* [39], *H-go SRJ* [46] and *KSimJoin* [20] for their state-of-the-art performance. We also correct the erroneous claims and make up the missing theoretical analysis in the original papers.

3.1 UISim

The algorithm depends on the random walk interpretation [12] which says that SimRank value $s(u, v)$ equals to the summation of *first-meeting probabilities* of all random walk pairs starting from u and v , respectively. To compute $s(u, v)$, *UISim* adopts the following idea: “Approximately prioritize all random walk pairs and first handle those with high meeting probabilities.” In practice, this is done by selecting a set of hub nodes H with high *in-degrees*. Note that a random walk passing many hubs will have small walking probability. Therefore, *UISim* prioritizes a random walk pair by the number of encountered hubs, which is an approximation of ordering walk pairs by their meeting probabilities.

For all-pair SimRank computation, *UISim* first computes the *prime out-subgraphs* of all nodes with at least two out-neighbors (denoted as $V_{\geq 2}^o$), and then assembles all generated random walk pairs (level-1 expansion in Figure 1(a)). The prime out-subgraph of node v contains all reachable nodes starting from v and following out-edges until a hub node is met. Next, we expand all encountered hubs with their prime out-subgraphs (level-2 expansion). The algorithm terminates with level- η expansion.

Theoretical Guarantee. *UISim* actually relaxes the first-meeting constraints of SimRank definition², and analyzes the expected error bound based on this assumption. Generally, it does not have an absolute error guarantee for the SimRank estimation $\hat{s}(u, v)$.

Complexity Analysis. According to [48], the time complexity of *UISim* for all-pair queries is bounded by $O(|V_{\geq 2}^o| |H|^\eta T)$ in expectation, where $T = O\left(\bar{d}^L (1 - \sum_{v \in H} d_{in}(v)/m)^L\right)$ stands for the expected number of random walks up to length L .

Empirical Analysis. *UISim* is evaluated for all-pair queries (by additive error) and for top- k similarity joins (by Precision@ k) but only on small graphs with tens of thousands nodes.

3.2 FLP & Opt-LP

Wang et al. [39] propose *SimRank graph*, a modified version of the node-pairs graph (see Figure 1(b)).

DEFINITION 5 (SIMRANK GRAPH). Given a graph G , its *SimRank graph* G^s is obtained from the node-pairs graph G^2 by removing all the in-edges of the singleton nodes. Besides, we slightly modify G^s by adding a virtual node (v_r, v_r) , and for each singleton node (v, v) , we add one in-edge from it to the virtual node.

Based on that, the idea is as follows: “*SimRank is reformalized as the solution of a new linear system, the computation of which can be approximated by Backward Push* [25].” In particular, the linear system is formulated as

$$[I - c(I - \text{Diag}(\overrightarrow{vec}(I)))(P^T \otimes P^T)]\overrightarrow{vec}(S) = \overrightarrow{vec}(I), \quad (2)$$

where \otimes denotes the Kronecker product, and $\overrightarrow{vec}(\cdot)$ (resp. $\text{Diag}(\cdot)$) transforms a matrix (resp. vector) to the corresponding vector (resp. diagonal matrix). The baseline algorithm, referred to as *Forward-LocalPush (FLP)*, is essentially *Backward Push* [25] for single-target PPR on G^s by setting the virtual node as target and following out-edges instead of in-edges (i.e., *reverse PPR*). An optimized algorithm named *OptimizeLocalPush (Opt-LP)* is proposed to avoid redundant computation and to efficiently handle self-loops. For implementation, the algorithm does not have to materialize G^s . We formalize the above claim as the following lemma³.

LEMMA 1. The SimRank similarity of u and v is exactly $\frac{1}{c(1-c)}$ times the reverse PPR value of node (u, v) w.r.t. the target node (v_r, v_r) on SimRank graph G^s and with stopping probability $1 - c$:

$$s(u, v) = \frac{1}{c(1-c)} \pi_{G^s}((u, v), (v_r, v_r)). \quad (3)$$

²Note that [48] provides a correction method to exclude *multi-hop double-meeting probabilities*, i.e., the probability of two random walks meeting exactly twice. It is non-trivial to extend the method to *multi-meeting probabilities*.

³For space constraint, we refer most of the proofs to the full version of the paper [1].

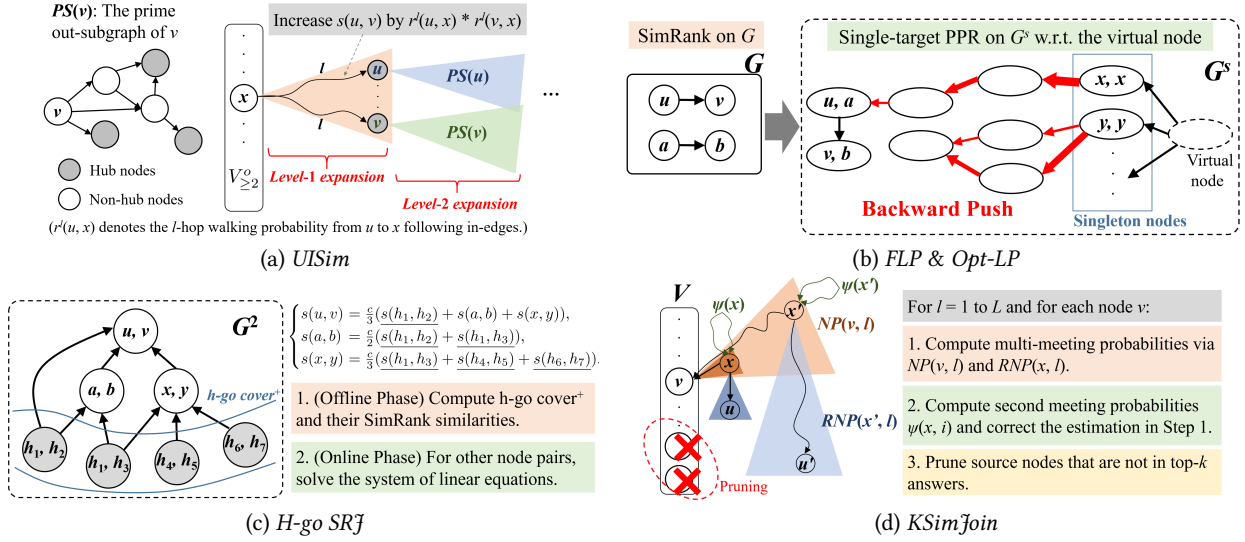


Figure 1: Overview of state-of-the-art algorithms for SimRank-based similarity joins.

Table 2: Comparison of state of the art. N/A means that the specific query is not studied in the original paper.

Algorithm	Theoretical Guarantee			Time Complexity	Empirical Study		
	Absolute	Threshold	Top-k		Absolute	Threshold	Top-k
UISim [48]	✗	N/A	✗	$O(V_{\geq 2}^o H ^T)$	✓	N/A	✓
FLP & Opt-LP [39]	✓	N/A	N/A	$O(\sum_{u,v \in V} d_{in}(u) d_{in}(v) s(u,v))$	✓	N/A	N/A
H-go SRJ [46]	✓	✗	N/A	$\tilde{O}(n\bar{d}^h + \frac{n H(G) }{\epsilon^2})$ (offline), $O(n\bar{d}^{h+3})$ (online)	N/A	Only query time	Only query time
KSimJoin [20]	✓	N/A	✗	$O(n\bar{d}^{\log \frac{1}{\epsilon}})$	N/A	N/A	Only query time

Theoretical Guarantee. FLP & Opt-LP guarantee absolute error following [25], i.e., $|\hat{s}(u, v) - s(u, v)| \leq \epsilon$ for each (u, v) .

Complexity Analysis. We observe that the proposed *average* time complexity $O(\frac{\bar{d}^2}{\epsilon})$ (Proposition 1 of [39]) is *incorrect* where \bar{d} denotes the average degree of G . This average complexity from [25] is achieved by considering *all possible target nodes* in the graph. However, for all-pair SimRank computation, the target node is *fixed* as the virtual node. We fix the issue using the following lemma.

LEMMA 2. To achieve an absolute error of ϵ for all-pair queries, following [25, 35], the time complexity of FLP and Opt-LP are bounded by $O(\frac{\sum_{u,v \in V} d_{in}(u) d_{in}(v) s(u,v)}{\epsilon})$.

Empirical Analysis. [39] conducts experiments for all-pair queries on both small and large graphs with respect to query time, but only evaluates estimation error on small graphs with thousands of nodes.

3.3 H-go SRJ

Inspired by the distance-based upper bound of SimRank which states that two nodes with a large shortest path distance cannot be very similar, the index-based algorithm H-go SRJ [46] answers threshold-based similarity joins with the following idea (Figure 1(c)): “First choose a proper set of node pairs (i.e., the h-go cover⁺) and precompute their SimRank estimations. For any other node pair, its SimRank estimation can be efficiently computed following recursive SimRank definition and with the help of the h-go cover⁺.”

In the offline phase, the algorithm first computes the h-go cover⁺ index with similar idea as the vertex cover, so that any h-hop path on the node-pairs graph intersects with the index. For each indexed node pair, its SimRank value can be estimated by a single-pair query. In the online phase, for every non-indexed node pair, H-go SRJ conducts graph traversal from it and builds a system of linear equations, which can be solved using Gaussian elimination.

Theoretical Guarantee. [46] does not discuss the error guarantee of H-go SRJ. We remedy this issue with the following analysis. Given an error parameter ϵ , for each indexed node pair we invoke the Monte Carlo algorithm [8] to sample $\tilde{O}(\frac{1}{\epsilon^2})$ pairs of random walks (the \tilde{O} notation ignores log factors). The following lemma bounds the estimation error of the online phase.

LEMMA 3. If all indexed SimRank similarities have at most ϵ additive error, then the ϵ error guarantee also holds for the similarity of any node pair computed in the online phase.

Complexity Analysis. According to [46], the h-go cover⁺ selection takes $O(n\bar{d}^h)$ time where \bar{d} is the average degree of G . We notice that computing SimRank estimation for the indexed node pairs takes an extra $\tilde{O}(\frac{n|H(G)|}{\epsilon^2})$ time, where $H(G)$ denotes the h-go cover⁺ of G . Therefore, the offline phase needs $\tilde{O}(n\bar{d}^h + n\frac{|H(G)|}{\epsilon^2})$ time. For the online phase, given a node pair, building the linear system by graph traversal needs $O(\bar{d}^{2h})$ time, while solving it incurs

$O(\bar{d}^{2h+1})$ time with all optimizations [46]. Since the linear system contains $O(\bar{d}^{2(h-1)})$ unknown values on average, the amortized cost for computing each non-indexed node pair is $O(\bar{d}^3)$. Assume there exist $O(n\bar{d}^h)$ non-indexed node pairs after distance-based pruning, the online phase needs $O(n\bar{d}^h \cdot \bar{d}^3) = O(n\bar{d}^{h+3})$ time.

Empirical Analysis. [46] evaluates query efficiency on graphs with up to millions of nodes but ignores the evaluation of accuracy.

3.4 KSimJoin

Based on the random walk interpretation, *KSimJoin* [20] answers top- k similarity joins via decomposing the SimRank value by meeting steps, i.e., $s(u, v) = \sum_{l=0}^{\infty} s^{(l)}(u, v)$, where $s^{(l)}(u, v)$ is the *first meeting probability* of two \sqrt{c} -walks from u and v at exact l -th step. It employs the following key idea: “Compute the meeting probabilities in ascending order of meeting steps. Multi-meeting cases at each step are eliminated by carefully utilizing previous computation with fewer steps.”

Specifically, the following equation is adopted to compute $s^{(l)}(u, v)$:

$$s^{(l)}(u, v) = \sum_{x \in V} \Pr((u, v) \xrightarrow{l} (x, x)) - \sum_{i \in [1, l-1], y \in V} \Pr((u, v) \xrightarrow{i} (y, y))\psi(y, l-i), \quad (4)$$

where $\psi(y, l-i)$ is the probability of two \sqrt{c} -walks starting from y first meeting at exact $(l-i)$ -th step, which is referred to as the *second meeting probability*. The algorithm performs a breadth-first search (BFS) from every node v following in-edges (Figure 1(d)). All visited node at step l is referred to as $NP(v, l)$. Then, for each node $x \in NP(v, l)$, a BFS following out-edges finds all node u that is l -step from x , denoted as $RNP(x, l)$. The multi-meeting probabilities are computed from $NP(v, l)$ and $RNP(x, l)$ for all $x \in NP(v, l)$. By leveraging Equation 4, it excludes all multi-meeting cases, where the second meeting probabilities $\psi(x, i)$ can be computed following a similar BFS-based incremental approach. For top- k queries, after each iteration of l , *KSimJoin* uses a carefully designed upper bound to prune nodes that cannot be in the top- k answers, where the upper bound is computed by relaxing Equation 4.

Theoretical Guarantee. [20] does not analyze the error guarantee of *KSimJoin*. Our lemma states that *KSimJoin* (without upper bound pruning) is essentially equivalent to the *Power Method* [12].

LEMMA 4. *By considering the terms from $l = 0$ to L with $L = O(\log \frac{1}{\varepsilon})$, *KSimJoin* achieves $O(\varepsilon)$ absolute error guarantee.*

Complexity Analysis. Computing the multi-meeting probabilities and the upper bound incurs $O(|R|\bar{d}^{2L})$ time [20]. The cost for second meeting probabilities is asymptotically identical. Since $L = O(\log \frac{1}{\varepsilon})$ and $|R| = O(n)$, the overall time complexity is $O(n\bar{d}^{\log \frac{1}{\varepsilon}})$.

Empirical Analysis. Empirical study of [20] only focuses on the efficiency of top- k queries on small and medium graphs.

3.5 Comparison of State-of-the-art Algorithms

Apart from the specific interpretation of SimRank in the original papers, *KSimJoin*, *UISim* and *FLP & Opt-LP* can be viewed as different ways of prioritizing meeting probabilities. *KSimJoin* orders

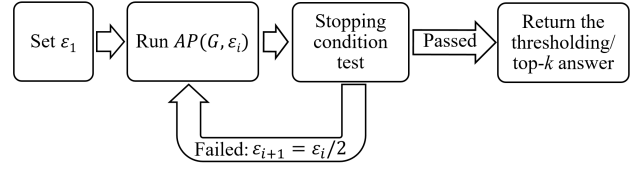


Figure 2: Illustration of the algorithm framework.

random walk pairs by walk length, *FLP & Opt-LP* prioritize them by the walking probabilities, and *UISim* can be considered as an approximation of the strategy in *FLP & Opt-LP*. However, existing experimental study is quite limited, including lack of evaluation on large graphs [20, 39, 48] and for estimation error [20, 46]. More importantly, none of the algorithms provides any approximation guarantee for both types of similarity joins.

4 OUR IMPROVED ALGORITHMS

4.1 Algorithm Framework

We adopt a simple yet effective idea [23, 37] to answer threshold-based and top- k queries with approximation bound (Figure 2). It takes an all-pair algorithm \mathcal{AP} as subprocedure and progressively decreases the input error parameter of \mathcal{AP} until the bound is satisfied. We present specific stopping conditions for our problems.

4.1.1 APThres for Threshold Queries. Recall that threshold-based similarity join finds (at least ρ fraction of) all node pairs with SimRank similarity above θ . Given an all-pair algorithm that guarantees an additive error of ε (with high probability), i.e., $|\hat{s}(u, v) - s(u, v)| \leq \varepsilon$ for each node pair (u, v) , if the estimated SimRank value $\hat{s}(u, v)$ satisfies that $\hat{s}(u, v) \geq \theta + \varepsilon$, then we have $s(u, v) \geq \theta$, that is, (u, v) is included in the result set. Similarly, for (u', v') with $\hat{s}(u', v') < \theta - \varepsilon$, we should exclude it from the result set. Intuitively, ε should be small enough to meet the approximation bound ρ . We iteratively decrease ε by half until the condition is satisfied.

Algorithm Description. The framework, denoted as *APThres*, is illustrated in Algorithm 1. The algorithm initializes the result set $\mathcal{R}(\theta)$ and the candidate set C as empty set, and set $\varepsilon_1 = O(\theta)$ (e.g. $\frac{\theta}{2}$) for the first iteration (Line 1). Then, at i -th iteration, *APThres* iteratively invokes \mathcal{AP} with error parameter ε_i (Lines 2-3). Particularly, we use $\hat{S}_i = \{\hat{s}_i(u, v) | u, v \in V, u \neq v\}$ to denote the all-pair SimRank estimations, where the subscript i specifies the iteration number. In practice, we do not store node pairs with zero estimated values. We use \hat{S}_i to update $\mathcal{R}(\theta)$ (Line 4) and C (Line 5). The algorithm terminates if $\frac{|\mathcal{R}(\theta)|}{|\mathcal{R}(\theta)| + |C|} \geq \rho$ holds, which implies that $\mathcal{R}(\theta)$ contains at least ρ fraction of the result (Lines 6-7). Otherwise, we halve the error parameter for the next iteration (Line 8). If our stopping condition is never satisfied, the algorithm will stop when the SimRank estimation of every node pair has no more than ε_{min} error. In both cases, $\mathcal{R}(\theta)$ is returned (Line 9).

Remark. The existence of approximation bound $\rho \in (0, 1)$ is a necessary condition for *APThres* to stop. Without ρ , which means setting $\rho = 1$, *APThres* stops only if $C = \emptyset$, which does not always hold. In worst cases, we claim that no such an algorithm exists with better efficiency than computing the exact SimRank values. Suppose that a graph has $|\mathcal{R}|$ node pairs with SimRank values being exactly θ , while other (non-singleton) node pairs have zero similarity. Even

for a ground truth algorithm that guarantees ε_{min} additive error, it still might return an answer that is arbitrary bad in terms of approximation bound, e.g., when all estimations are below θ .

Algorithm 1: The *APTThres* framework

Input: Directed graph $G = (V, E)$, threshold θ , approximation bound ρ , all-pair algorithm \mathcal{AP}
Output: $\mathcal{R}(\theta)$, which contains at least ρ fraction of all node pairs with similarities above θ

```

1  $\mathcal{R}(\theta) \leftarrow \emptyset, C \leftarrow \emptyset, \varepsilon_1 \leftarrow O(\theta)$ ;
2 for  $i = 1$  to  $\log_2 \frac{1}{\varepsilon_{min}}$  do
3    $\hat{S}_i \leftarrow \mathcal{AP}(G, \varepsilon_i)$ ;
4    $\mathcal{R}(\theta) \leftarrow \mathcal{R}(\theta) \cup \{(u, v), \hat{s}_i(u, v) \in \hat{S}_i | \hat{s}_i(u, v) \geq \theta + \varepsilon\}$ ;
5    $C \leftarrow (C \cup \{(u, v), \hat{s}_i(u, v) \in \hat{S}_i | \hat{s}_i(u, v) \in [\theta - \varepsilon, \theta + \varepsilon]\}) \setminus \mathcal{R}(\theta)$ ;
6   if  $\frac{|\mathcal{R}(\theta)|}{|\mathcal{R}(\theta)| + |C|} \geq \rho$  then
7     break;
8    $\varepsilon_{i+1} \leftarrow \varepsilon_i / 2$ ;
9 return  $\mathcal{R}(\theta)$ ;
```

Correctness. The following theorem proves that *APTThres* has approximation bound for threshold-based queries.

THEOREM 1. *APTThres* with an all-pair algorithm \mathcal{AP} of absolute error guarantee holds approximation bound ρ for threshold-based similarity joins.

PROOF (SKETCH). All node pairs included in $\mathcal{R}(\theta)$ belong to the actual result because $\hat{s}_i(u, v) \geq \theta + \varepsilon$ and $|\hat{s}_i(u, v) - s(u, v)| \leq \varepsilon$. Similarly, if some node pair in the real answer is not included, it must be in the candidate set C . Therefore, $\frac{|\mathcal{R}(\theta)|}{|\mathcal{R}(\theta)| + |C|}$ is a lower bound of the approximation guarantee. The above claims also hold for an algorithm \mathcal{AP} with a small failure probability. If the failure probability of *APTThres* is δ (e.g. $o(\frac{1}{n})$), it is sufficient to set the failure probability of \mathcal{AP} as $\delta / \log_2 \frac{1}{\varepsilon_{min}}$, so the theorem holds. \square

4.1.2 APTop-k for Top-k Queries. We provide a framework *APTtop-k* for top- k similarity joins with approximation bound ρ , which is analogous to *APTThres* but with different stopping condition.

Algorithm Description. We demonstrate the pseudocode in Algorithm 2. The algorithm sets the initial value of the error parameter as $O(c)$ (Line 1), and iteratively invokes the all-pair algorithm (Lines 2-3). At i -th iteration, the estimation values \hat{S}_i are sorted in descending order, denoted as $\{\hat{s}_{i,1}, \hat{s}_{i,2}, \dots, \hat{s}_{i,k}, \dots, \hat{s}_{i,n(n-1)}\}$ (Line 4). We set the stopping conditions as $\hat{s}_{i, \lceil \rho k \rceil} - \varepsilon_i \geq \hat{s}_{i, k+1} + \varepsilon_i$ (Lines 5-6), which says that the lower bound of the $\lceil \rho k \rceil$ -th largest estimation is no smaller than the upper bound of the $(k+1)$ -th one. The algorithm decreases the error parameter by half if it does not terminate this round (Line 7). Finally, *APTtop-k* returns the set of node pairs with the top- k largest estimations (Lines 8-9).

Correctness. We have the following theorem for *APTtop-k*.

THEOREM 2. *APTtop-k* with an all-pair algorithm \mathcal{AP} of absolute error guarantee holds approximation bound ρ for top- k queries.

Algorithm 2: The *APTtop-k* framework

Input: Directed graph $G = (V, E)$, parameter k , approximation bound ρ , all-pair algorithm \mathcal{AP}
Output: $\mathcal{R}(k)$, which contains at least ρ fraction of node pairs with top- k largest SimRank values

```

1  $\varepsilon_1 \leftarrow O(c)$ ;
2 for  $i = 1$  to  $\log_2 \frac{1}{\varepsilon_{min}}$  do
3    $\hat{S}_i \leftarrow \mathcal{AP}(G, \varepsilon_i)$ ;
4    $\{\hat{s}_{i,1}, \hat{s}_{i,2}, \dots, \hat{s}_{i,k}, \dots, \hat{s}_{i,n(n-1)}\} \leftarrow \text{sort}(\hat{S}_i)$  (in descending order);
5   if  $\hat{s}_{i, \lceil \rho k \rceil} - \varepsilon_i \geq \hat{s}_{i, k+1} + \varepsilon_i$  then
6     break;
7    $\varepsilon_{i+1} \leftarrow \varepsilon_i / 2$ ;
8 Let  $\mathcal{R}(k)$  be the node pairs with top- $k$  largest estimations;
9 return  $\mathcal{R}(k)$ ;
```

PROOF (SKETCH). We claim when the stopping condition is satisfied, for each $j \leq \lceil \rho k \rceil$ and each $j' \geq k+1$, the SimRank similarity of the j -th node pair is no smaller than that of the j' -th, as the lower bound of the former is no smaller than the upper bound of the latter. Consequently, it means that the top- $\lceil \rho k \rceil$ node pairs are always included in the exact result. Similar to Theorem 1, this claim holds for algorithm \mathcal{AP} with a small failure probability. \square

4.2 Randomized Reverse Local Push (R^2LP)

We present an all-pair algorithm following the random walk interpretation and leveraging the local push techniques [25, 35]. It is inspired by the advantage of *Randomized Backward Search (RBS)* [35] over *Backward Push* for PPR. As computing SimRank on the node-pairs graph with $O(n^2)$ nodes incurs much more computational and memory costs than computing PPR on the original graph, to guarantee the practical efficiency, we devise a pruning strategy with absolute error guarantee to dramatically enhance query time.

Algorithm Description. We refer to the algorithm as *Randomized Reverse Local Push (R^2LP)* (see Algorithm 3). The SimRank value $s(u, v)$ is split into $s^{(l)}(u, v)$ for $l \in [0, \infty)$, where $s^{(l)}(u, v)$ is the probability of two \sqrt{c} -walks $W_{\sqrt{c}}(u)$ and $W_{\sqrt{c}}(v)$ first meeting after exact l steps. It is obvious that $s(u, v) = \sum_{l=0}^{\infty} s^{(l)}(u, v)$. R^2LP estimates $s^{(l)}(u, v)$ for all node pairs in ascending order of l . Following previous work [24, 46], by only considering terms up to $s^{(L)}(u, v)$, the incurred error is bounded by c^{L+1} . Therefore, it is sufficient to set $L = O(\log \frac{1}{\varepsilon})$ to bound an $O(\varepsilon)$ error (Line 1). We initialize the algorithm in Lines 2-3, where $s^{(l)}(u, v)$ corresponds to the reserves of PPR [25, 35], and we omit the residues because it holds that $r^{(l)}(u, v) = \hat{s}^{(l)}(u, v) / (1 - c)$. Then, for $l = 0$ to $L - 1$, we compute the estimates $\hat{s}^{(l+1)}(u, v)$ based on $\hat{s}^{(l)}(u, v)$ (Lines 4-10). For each $\hat{s}^{(l)}(u, v)$ at level l , we check if its value is large enough (Line 5). For now, we set $f(\varepsilon) = 0$ to be the same as *RBS*. Similarly, our algorithm checks each out-neighbor (u', v') of (u, v) in the SimRank graph G^s (Lines 6&9). For node pair (u', v') with limited amount of in-neighbors in G^s , we conduct deterministic push as in *Backward Push* (Line 7). Otherwise, the randomized push is applied (Lines 8&10), which gives an unbiased estimation [35] but ignores many

node pairs in practice. The returned value $\hat{s}(u, v)$ summarizes the estimations of every step l (Lines 11-12).

Algorithm 3: Randomized Reverse Local Push (R^2LP)

Input: Directed graph $G = (V, E)$ with sorted adjacency lists, additive error parameter ε

Output: $\hat{s}(u, v)$ for all $u, v \in V$

- 1 $L \leftarrow \log_c(1-c)\varepsilon$;
- 2 $\hat{s}^{(l)}(u, v) \leftarrow 0$ for $l \in [0, L]$ and $u, v \in V$;
- 3 $\hat{s}^{(0)}(v, v) \leftarrow 1$ for all $v \in V$;
- 4 **for** $l = 0$ to $L - 1$ **do**
- 5 **for each** (u, v) with $\hat{s}^{(l)}(u, v) > f(\varepsilon)$ **do**
- 6 **for each** $u' \in O(u), v' \in O(v), u' \neq v'$ and
- 7 $\sqrt{d_{in}(u')d_{in}(v')} \leq \frac{c\hat{s}^{(l)}(u, v)}{(1-c)\varepsilon}$ **do**
- 8 $\hat{s}^{(l+1)}(u', v') \leftarrow \hat{s}^{(l+1)}(u', v') + \frac{c\hat{s}^{(l)}(u, v)}{d_{in}(u')d_{in}(v')}$;
- 9 $r \leftarrow \text{rand}(0, 1)$;
- 10 **for each** $u' \in O(u), v' \in O(v), u' \neq v'$ and
- 11 $\frac{c\hat{s}^{(l)}(u, v)}{(1-c)\varepsilon} < \sqrt{d_{in}(u')d_{in}(v')} \leq \frac{c\hat{s}^{(l)}(u, v)}{(1-c)\varepsilon r}$ **do**
- 12 $\hat{s}^{(l+1)}(u', v') \leftarrow \hat{s}^{(l+1)}(u', v') + \frac{(1-c)\varepsilon}{\sqrt{d_{in}(u')d_{in}(v')}}$;
- 13 $\hat{s}(u, v) \leftarrow \sum_{l=0}^L \hat{s}^{(l)}(u, v)$;
- 14 **return** non-zero $\hat{s}(u, v)$ for $u, v \in V, u \neq v$;

We discuss two implementation details. First, we do not need to materialize G^S . To get all out-neighbors of (u, v) in G^S , it is sufficient to traverse the out-neighbors of u and v in G simultaneously. Second, to guarantee the efficiency, the randomized push strategy [35] needs the out-neighbors of (u, v) to be sorted in *ascending* order of *in-degree*. This can be achieved by first preprocessing the adjacency lists of G before invoking the algorithm.

Pruning Strategy. Recall that the above algorithm (as well as *RBS*) combines deterministic and randomized push for algorithmic efficiency. Intuitively, it encourages most node pairs to be handled by the randomized process. When applied to SimRank computation, due to the excessive number of node pairs in G^S , the randomized part of the algorithm still needs to be improved to effectively prune the node pairs with negligible similarities. We additionally propose a deterministic pruning strategy to directly rule out estimations with small values at each step l . To be precise, we set $f(\varepsilon) = \tilde{O}(\varepsilon)$ (i.e., we omit the logarithmic factor of ε) in Line 5 of Algorithm 3.

For the deterministic process, it is almost unaffected as the algorithm requires $\hat{s}^{(l)}(u, v) > \frac{(1-c)\varepsilon}{c} \cdot \sqrt{d_{in}(u')d_{in}(v')}$ and the RHS usually exceeds ε . Hence, the deterministic part guarantees the large fraction of the similarity value is computed exactly. For the randomized process, the pruning strategy becomes effective when r is unfortunately very small, namely, $r \leq \frac{c}{(1-c)\sqrt{d_{in}(u')d_{in}(v')}}$. As we shall see in the experiments, the pruning power is significant. **Correctness.** The following lemma shows that our R^2LP algorithm guarantees additive error ε with high probability. Particularly, the above pruning strategy does not affect the error guarantee.

LEMMA 5. *By conducting the randomized push procedure of R^2LP for $O(\log \frac{n}{\delta})$ times and apply the Median-of-Mean trick [7], we have $|\hat{s}(u, v) - s(u, v)| \leq \varepsilon$ with at least $1 - \delta$ probability.*

PROOF (SKETCH). Note that without the pruning strategy, by Lemma 1, to guarantee ε absolute error it is sufficient to set the input error parameter to $c(1-c)\varepsilon$. Since without pruning the algorithm is equivalent to invoking *RBS* on G^S , the lemma follows by Theorem 4.2 of [35]. Next we show when conducting *RBS* on G^S , the caused error by pruning is also bounded.

First, we prove that $\mathbb{E}[\hat{\pi}'_{l+1}(u, t)] \geq \pi_{l+1}(u, t) - \frac{1-\alpha}{\alpha}\varepsilon$ for $\forall u \in V$ and $l \in [0, L]$. Here, $\hat{\pi}'$ and $\hat{\pi}$ denote the PPR estimation (on G^S) with and without pruning, respectively. We prove it by induction. For $l = 0$, $\hat{\pi}'_0(t, t) = \alpha$ and $\hat{\pi}'_0(v, t) = 0$ for any $v \neq t$ where t denotes the target node (v_r, v_r) of G^S (we assume $\alpha > \varepsilon$), and the claim holds. Suppose $\mathbb{E}[\hat{\pi}'_l(v, t)] \geq \pi_l(v, t) - \sum_{i=1}^{l-1} (1-\alpha)^i \varepsilon$ holds, then we have $\hat{\pi}'_{l+1}(u, t) = \sum_{v \in N_{out}(u)} X'_{l+1}(u, v)$ and

$$\begin{aligned} \mathbb{E}[\hat{\pi}'_{l+1}(u, t) | \{\hat{\pi}'_l\}] &= \sum_{v \in N_{out}(u)} \mathbb{E}[X'_{l+1}(u, v) | \{\hat{\pi}'_l\}] \\ &\geq \sum_{v \in N_{out}(u)} \frac{(1-\alpha)\hat{\pi}'_l(v, t)}{d_{out}(u)} - (1-\alpha)\varepsilon. \end{aligned} \quad (5)$$

Hence, we can derive that

$$\begin{aligned} \mathbb{E}[\hat{\pi}'_{l+1}(u, t)] &\geq \sum_{v \in N_{out}(u)} \frac{(1-\alpha)\mathbb{E}[\hat{\pi}'_l(v, t)]}{d_{out}(u)} - (1-\alpha)\varepsilon \\ &\geq \sum_{v \in N_{out}(u)} \frac{(1-\alpha)(\pi_l(v, t) - \sum_{i=1}^{l-1} (1-\alpha)^i \varepsilon)}{d_{out}(u)} - (1-\alpha)\varepsilon \\ &= \pi_{l+1}(u, t) - \sum_{i=1}^l (1-\alpha)^i \varepsilon \geq \pi_{l+1}(u, t) - \frac{1-\alpha}{\alpha}\varepsilon. \end{aligned} \quad (6)$$

Second, we show that with our pruning strategy, the variance of every $\hat{\pi}'_l(v, t)$ is still bounded by Lemma 4.4 of [35]. The key observation is that the variance comes from the randomized push and pruning reduces its estimation value. Specifically, we have $\text{Var}[\hat{\pi}'_{l+1}(u, t) | \{\hat{\pi}'_l\}] \leq \text{Var}[\hat{\pi}_{l+1}(u, t) | \{\hat{\pi}_l\}]$.

By the above discussion, we have $\mathbb{E}[\hat{\pi}'(u, t)] \geq \pi(u, t) - L\frac{1-\alpha}{\alpha}\varepsilon$, and the variance of $\hat{\pi}'(u, t)$ is also bounded by $L\alpha\varepsilon^2$ [35]. By setting $L = O(\log \frac{1}{\varepsilon})$ to bound the truncation error by $O(\varepsilon)$, the total estimation error can be bounded by ε , and the lemma follows. \square

Complexity Analysis. We use the following theorem to guarantee the efficiency of R^2LP .

THEOREM 3. *The expected time complexity of R^2LP is bounded by $\tilde{O}(\frac{\sum_{u \in V} \sum_{v \in V} \sqrt{d_{in}(u)d_{in}(v)}s(u, v)}{\varepsilon})$.*

PROOF (SKETCH). It follows directly from Lemma 5 and Theorem 4.2 of [35]. \square

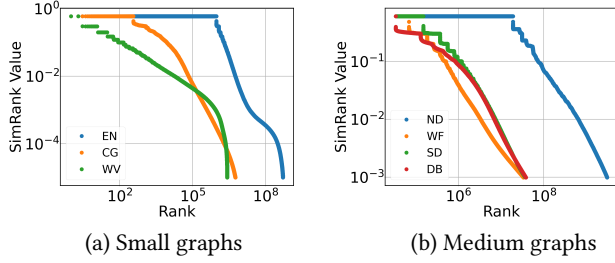


Figure 3: Distribution of all-pair SimRank values.

5 COMPUTATIONAL COMPLEXITY ANALYSIS

5.1 Modeling SimRank Distribution

We demonstrate in Figure 3 the all-pair SimRank distributions of three small graphs and four medium graphs used in our experiments (Section 6). We omit the base cases with $s(v, v) = 1$. Generally the distributions show different degrees of skewness with a fraction of the top ranking node pairs having a similarity of c . Inspired by [2, 6], we propose the following assumption of SimRank distribution.

DEFINITION 6 (POWER LAW BOUNDED DISTRIBUTION OF SIMRANK). Assume we have nnz node pairs with non-zero SimRank values for a given graph G . Let s_j denote the j -th largest SimRank value, for $j \in [1, nnz]$. We say the SimRank distribution is power law bounded (PLB) if for any $x = 0, \dots, \lfloor \log_2 nnz \rfloor$, there exist parameters $\beta > 0$ and $b_2 > b_1 > 0$ such that

$$\sum_{j=2^{x+1}}^{2^{x+1}-1} b_1 \cdot r(j)^{-\beta} \leq \sum_{j=2^x}^{2^{x+1}-1} s_j \leq \sum_{j=2^{x+1}}^{2^{x+1}-1} b_2 \cdot r(j)^{-\beta}, \quad (7)$$

where $r(j) = \max(1, j - t)$, and we set the shift $t = |\{(u, v) | s(u, v) = c, u, v \in G\}|$. Specifically, β is referred to as the PLB exponent of SimRank distribution, while b_1 and b_2 are small numbers that can be taken as constants.

Please refer to the full version of our paper [1] for empirical validation of the assumption.

5.2 Complexity of Threshold-based Queries

With the above assumption, we derive the *upper bound* of the complexity of threshold-based queries. We do this by bounding the complexity of our algorithm framework which integrates an all-pair algorithm with absolute error guarantee (e.g., *Opt-LP* or *R²LP*).

LEMMA 6. Given a graph G and an all-pair algorithm \mathcal{AP} that guarantees absolute error, assume that *APT* terminates when the error parameter reaches ϵ_t , and we have $\text{cost}(\mathcal{AP}(G, \epsilon_{i+1})) \geq p \cdot \text{cost}(\mathcal{AP}(G, \epsilon_i))$ for $i \in [1, t]$ with $p > 1$. The time complexity of the algorithm is then bounded by $O(\text{cost}(\mathcal{AP}(G, \epsilon_t)))$.

Next, we bound the error parameter ϵ_t for threshold-based queries with the help of Definition 6.

LEMMA 7. We are given the problem input as in Lemma 6. Moreover, the graph G has PLB SimRank distribution. For threshold-based queries with approximation bound ρ , we have

$$\epsilon_t = O(\max(\frac{(b_1 - \rho^\beta b_2)\theta}{b_1 + \rho^\beta b_2}, \epsilon_{min})). \quad (8)$$

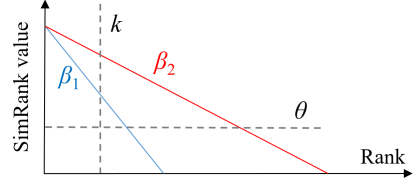


Figure 4: Illustration of the complexity of threshold-based and top- k similarity joins.

The following theorem bounds the computational complexity of threshold-based queries.

THEOREM 4. Given a graph G with PLB SimRank distribution, an approximation bound ρ with $b_1 - \rho^\beta b_2 > 0$, and an all-pair algorithm \mathcal{AP} that guarantees absolute error, let $\text{cost}(\mathcal{AP}(G, \epsilon))$ denote the complexity of \mathcal{AP} given error parameter ϵ . The *APT* framework with \mathcal{AP} answers ρ -approximation threshold-based queries in $O(\text{cost}(\mathcal{AP}(G, \frac{(b_1 - \rho^\beta b_2)\theta}{b_1 + \rho^\beta b_2})))$ time.

5.3 Complexity of Top- k Queries

For top- k queries, the following theorem is provided, with the intuition that the complexity is mainly determined by the gap between the $\lceil \rho k \rceil$ -th and the $(k + 1)$ -th largest values.

THEOREM 5. We are given the problem input as in Lemma 7. For top- k queries with approximation bound ρ , with the PLB assumption of SimRank distribution, let $x_1 = 2^{\lceil \log_2 \lceil \rho k \rceil \rceil} - 1$ and $x_2 = 2^{\lceil \log_2 (k+1) \rceil}$. We have

$$\epsilon'_t = O(\max(b_1(x_1 - t)^{-\beta} - b_2(x_2 - t)^{-\beta}, \epsilon_{min})). \quad (9)$$

Furthermore, with the power law assumption of all-pair SimRank distribution, i.e., $s_j = \Theta((j - t)^{-\beta})$, we have

$$\epsilon'_t = O(\max((b_1(\lceil \rho k \rceil - t)^{-\beta} - b_2(k + 1 - t)^{-\beta}), \epsilon_{min})). \quad (10)$$

The *APT* top- k framework with \mathcal{AP} answers ρ -approximation top- k queries in $O(\text{cost}(\mathcal{AP}(G, \epsilon'_t)))$ time.

Lastly, we present a qualitative analysis of the complexities w.r.t. the graph property β , which has different impacts on threshold-based and top- k queries. An illustrative demonstration is referred to Figure 4. For threshold-based queries, a smaller β (e.g. β_2) leads to a smaller ϵ_t in Lemma 7, and the complexity upper bound increases. Intuitively, more node pairs are included in the answer, while given an error parameter ϵ , the size of node pairs to be carefully handled (e.g., with similarity values between $\theta - \epsilon$ and $\theta + \epsilon$) also increases. In contrast, as Theorem 5 shows, a larger β (e.g. β_1) makes a smaller gap between the $\lceil \rho k \rceil$ -th and the $(k + 1)$ -th SimRank values, so the algorithm needs more effort to separate them. This is in consistent with Figure 4, where the k -th SimRank value becomes smaller (also note that the figure is in log-log scale).

6 EXPERIMENTAL EVALUATION

6.1 Experimental Setup

Datasets. We use three small graphs (CG, WV, and EN), four medium graphs (SD, DB, WF, and ND), and three large graphs (CP, LJ, and WZ) for evaluation. All graphs are obtained from [4, 5, 15, 17]

with basic statistics listed in Table 3, where \bar{d} is the average degree and β is the fitted PLB exponent of all-pair SimRank distribution[1].

Table 3: Datasets and their statistics.

Dataset	Type	n	m	\bar{d}	β
ca-GrQc (CG)	U	5.2K	14.5K	2.77	1.063
Wiki-Vote (WV)	D	7.1K	103.7K	14.57	0.495
email-Enron (EN)	U	36.7K	183.8K	5.01	1.303
Slashdot0922 (SD)	D	82.2K	948.5K	11.54	1.098
DBLP (DB)	U	317.1K	1.05M	3.31	0.866
Wikilinks-fy (WF)	D	65.6K	1.07M	16.35	1.098
Notre Dame (ND)	D	325.7K	1.5M	4.6	0.964
cit-Patents (CP)	D	3.77M	16.52M	4.38	-
LiveJournal (LJ)	D	4.85M	68.99M	14.23	-
Wikilinks-zh (WZ)	D	1.79M	72.61M	40.65	-

Ground truths. For small and medium graphs, we compute the all-pair ground truth using *Power Method* [12] with an absolute error about 10^{-7} . For large graphs, since *Power Method* is computationally intractable, we randomly choose 100 nodes in each graph and use *ExactSim* [36] to compute their single-source ground truths with $\epsilon_{min} = 10^{-7}$. The partial-pair ground truth is used for evaluation. **Baselines.** We use *FLP* & *Opt-LP* [39], *UISim* [48], and *KSimJoin* [20] as our baselines. The implementation of [39, 48] is obtained from the authors. We implement other baselines and our improved algorithms in C++. Specifically, we do not consider *H-go SRJ* [46] for efficiency issues, which shows significant inferior performance compared to other baselines.

We conduct all experiments on an Ubuntu server with an Hygon C86 7151 processor and 1TB memory, repeat each experimental settings for ten times and report the average measures. Throughout the paper, all reported query times are in seconds.

6.2 Evaluation of Additive Error

Experimental settings. For *FLP*, *Opt-LP*, *R²LP* and *KSimJoin* with absolute error guarantee, we vary the error parameter $\epsilon = \{0.05, 0.01, 0.005, 0.001\}$ for small and medium graphs, and set $\epsilon = \{0.1, 0.05, 0.01\}$ for large graphs. In particular, we slightly modify *KSimJoin* to answer all-pair queries by excluding the upper bound computation and the iterative pruning framework, and set $L = 2$. For *UISim*, we fix the number of iterations $\eta = 2$ and vary the number of hub nodes $|H| = \{0.2n, 0.5n\}$. In its implementation, the authors adopt another parameter named stopping reachability (*stopRea*), where the graph expansion stops either when a hub node is encountered or the walking probability falls below *stopRea*. We vary it in $\{0.005, 0.001\}$ for small graphs and $\{0.05, 0.01\}$ for medium graphs, and set it as 0.01 for large graphs.

Evaluation metric. We use *AvgErr* for evaluation, which is defined as $AvgErr = \frac{\sum_{(u,v) \in S(\theta)} |s^*(u,v) - \hat{s}(u,v)|}{|S(\theta)|}$, where $s^*(u,v)$ and $\hat{s}(u,v)$ are the ground truth and estimated value respectively, and $S(\theta) = \{(u,v) | s^*(u,v) \geq \theta, u,v \in V, u \neq v\}$. We set $\theta = 10^{-3}$ in the following experiments.

The results are demonstrated in Figure 5. For *UISim*, the unfilled and filled markers correspond to *stopRea*'s default values (0.001 and 0.01) and enlarged values (0.005 and 0.05), respectively. Surprisingly,

although *UISim* does not guarantee the absolute error, it outperforms *FLP* and *Opt-LP* on 6 of 10 graphs. However, we note that varying the number of hubs has limited influence, that is, the two unfilled markers (as well as the two filled markers) are very close. The reason is that most of graph expansions terminate because of stopping reachability rather than encountering hub nodes.

For algorithms with absolute error guarantees, *Opt-LP* outperforms *FLP* on most graphs thanks to the optimizations, except that it fails on the largest graph WZ (not terminated for over 10 hours). *KSimJoin* only manages to answer all-pair queries on the two smallest graphs and has similar performance as *UISim*, while the tremendous memory consumption makes it infeasible on larger graphs. For example, on WV and with $L = 2$, the *NP* structure for all nodes has about 1.8M items. With the randomized push strategy, *R²LP* achieves significantly better performance compared to all other baselines. On most graphs, it is faster than *FLP*, *Opt-LP*, and *KSimJoin* by up to an order of magnitude with similar additive errors. The only exception is ND, where *R²LP* and *Opt-LP* show comparable results. Since ND has numerous node pairs of very large similarity values (see Figure 3(b)), the cost of deterministic push dominates.

Evaluation of *R²LP*'s pruning strategy. Figure 6 demonstrates the performance of *R²LP* with and without the pruning strategy. For each dataset, we tune both settings to ensure comparable additive errors and compare the query time. *R²LP* w/ pruning clearly achieves better effectiveness-efficiency tradeoff. On EN, WF, LJ, and WZ, the speedup is about one order of magnitude.

To better illustrate the pruning power, we use Table 4 to show the number of edge traversals for deterministic and randomized pushes of *R²LP* under both settings. For deterministic push, the costs are very close, while our pruning strategy manages to eliminate up to 47%-96% of the costs for randomized push. The pruning strategy also effectively reduces memory cost, as without pruning the algorithm causes out-of-memory (OOM) error on LJ.

6.3 Evaluation of Threshold-based Queries

Experimental settings and metrics. We adopt *Precision*, *Recall* and *F1-score* for evaluation. Given a threshold-based query with parameter $\theta \in (0, c]$, let $R^*(\theta)$ denote the set of node pairs with ground truth SimRank values above or equal to θ , and we assume $|R^*(\theta)| > 0$. Let $R_{\mathcal{A}}(\theta)$ denote the returned set by algorithm \mathcal{A} . We have $Precision = |R_{\mathcal{A}}(\theta) \cap R^*(\theta)| / |R_{\mathcal{A}}(\theta)|$, $Recall = |R_{\mathcal{A}}(\theta) \cap R^*(\theta)| / |R^*(\theta)|$, and $F1-score = 2 / (\frac{1}{Precision} + \frac{1}{Recall})$.

We first compare the three best all-pair algorithms *R²LP*, *Opt-LP*, and *UISim* with $\theta = \{0.1, 0.01, 0.001\}$. Note that none of these algorithms provides approximation bound for threshold-based queries, but we can still compare their empirical performance. The results are demonstrated in Figure 7⁴. First, the *F1-scores* vary significantly for different graphs and with different θ . It is not recommended to answer the threshold-based query with a general all-pair algorithm by simply setting $\epsilon = O(\theta)$. For example, on CG, EN, DB, and WF, the *F1-scores* with $\epsilon = \theta = 0.001$ are below 0.9. Second, *R²LP* generally achieves the best tradeoff between efficiency and accuracy except on ND. Third, for *UISim*, we choose the best one among the four configurations in Section 6.2. On a majority of graphs, it

⁴For the results of *Precision* and *Recall*, please refer to the full version of our paper [1].

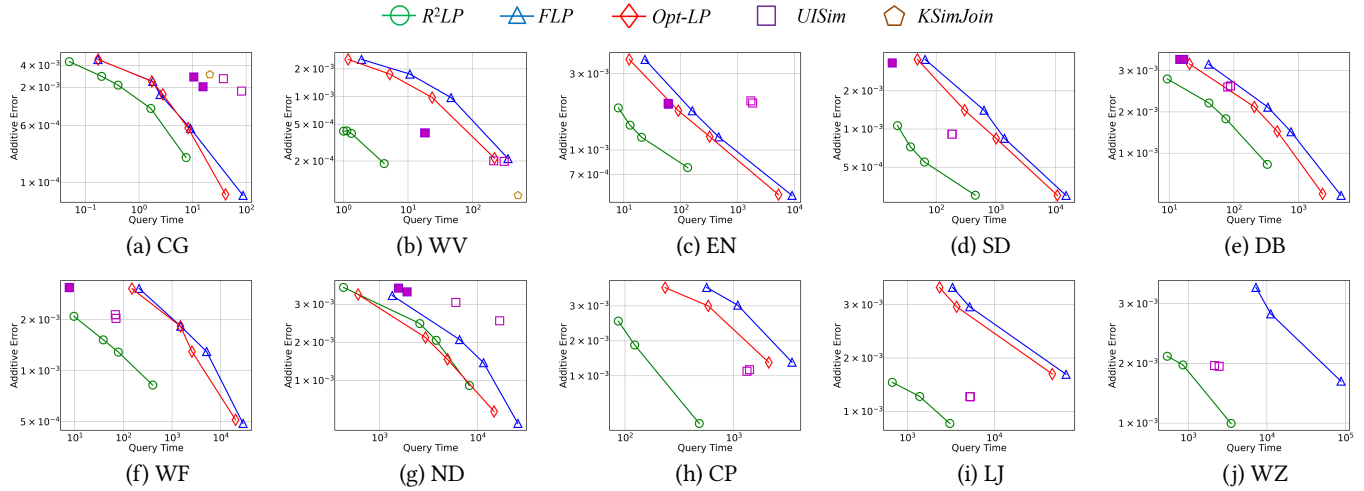


Figure 5: Additive error of SimRank estimation.

Table 4: R^2LP 's pruning power for deterministic and randomized push in terms of edge traversal ($\epsilon = 0.01$).

		CG	WV	EN	SD	DB	WF	ND	CP	LJ	WZ
Deterministic	w/ pruning	4.29e5	6.82e6	2.56e7	5.69e7	3.49e7	3.98e7	1.49e9	4.81e8	2.65e9	2.03e9
	w/o pruning	4.32e5	6.82e6	2.57e7	5.72e7	3.51e7	3.99e7	1.7e9	4.82e8	OOM	2.03e9
Randomized	w/ pruning	8.83e5	3.22e5	1.18e7	2.29e7	5.85e7	7.98e7	5.74e8	4.68e8	1.43e9	4.5e9
	w/o pruning	3.11e6	5.33e6	2.13e8	6.18e8	3.85e8	4.18e8	4.36e9	8.91e8	OOM	1.85e10

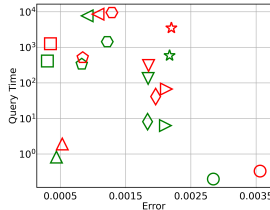


Figure 6: R^2LP with and without the pruning strategy.

surpasses $Opt-LP$ mainly due to the fast query time. However, as we decrease $stopRea$ or increase η , the algorithm fails to compute an output in several hours, which implies the hub-based strategy is less flexible than the probability-guided ones.

Next, we apply the $APThres$ framework and integrate it with R^2LP and $Opt-LP$, the best algorithms with absolute error guarantee. By fixing $\theta = 0.01$, we test their actual performance with different approximation bounds, which is shown in Table 5. Consistent with our theoretical guarantee, both $APThres+Opt-LP$ and $APThres+R^2LP$ hold the approximation bound on all graphs and all settings of ρ . $APThres+R^2LP$ achieves similar $F1$ -score with much less query time except on ND, of which the reasons are explained in Section 6.2.

Compared to Figure 7, both algorithms incur significantly more time, and fail on LJ for different reasons⁵ while $APThres+Opt-LP$ also fails on WZ. Recall that R^2LP is able to handle LJ with acceptable accuracy if we do not force an approximation bound. Meanwhile,

⁵We say an algorithm is time out if it takes over an hour in one iteration.

Table 5: Threshold-based queries with $\theta = 0.01$. We vary $(\rho_1, \rho_2) = (0.9, 0.99), (0.9, 0.95), (0.7, 0.9)$ for small, medium and large graphs respectively.

	$Opt-LP(\rho_1)$		$Opt-LP(\rho_2)$		$R^2LP(\rho_1)$		$R^2LP(\rho_2)$	
	Time	F1	Time	F1	Time	F1	Time	F1
CG	5.55	0.991	11.99	0.997	2.85	0.978	8.28	0.993
WV	152.4	0.998	159.7	0.998	31.2	0.998	32.0	0.999
EN	2944	0.995	10625	0.999	201.8	0.989	1040	0.996
SD	9451	0.998	12011	0.998	611.6	0.997	1457	0.998
DB	1994	0.994	3480	0.997	501.6	0.985	1119	0.992
WF	13336	0.994	8938	0.994	649.9	0.984	1370	0.99
ND	10691	0.991	13129	0.996	6607	0.936	10269	0.958
CP	5487	0.999	8504	0.999	1903	0.999	4350	0.999
LJ	Time Out				Out of Memory (OOM)			
WZ	Time Out				3743	0.985	7411	0.99

the actual $F1$ -scores shown in Table 5 far exceed the bound ρ . We believe this phenomenon is attributed to the pessimistic estimation of the absolute error. In practice, we have $|\hat{s}(u, v) - s(u, v)| \ll \epsilon$. Hence, it remains an open problem to answer threshold-based queries more efficiently and with theoretical approximation bound.

6.4 Evaluation of Top- k Queries

Experimental settings and metrics. For top- k similarity joins, $Precision@k$ is used for evaluation. Let $R^*(k)$ and $R_{\mathcal{A}}(k)$ denote the ground truth and the results of algorithm \mathcal{A} , respectively. We have $Precision@k = |R_{\mathcal{A}}(k) \cap R^*(k)|/k$. Because the numbers of node pairs with large similarities vary significantly for different graphs, instead of setting different k for our datasets, for medium and large graphs we conduct *all-pair queries* but only evaluate the

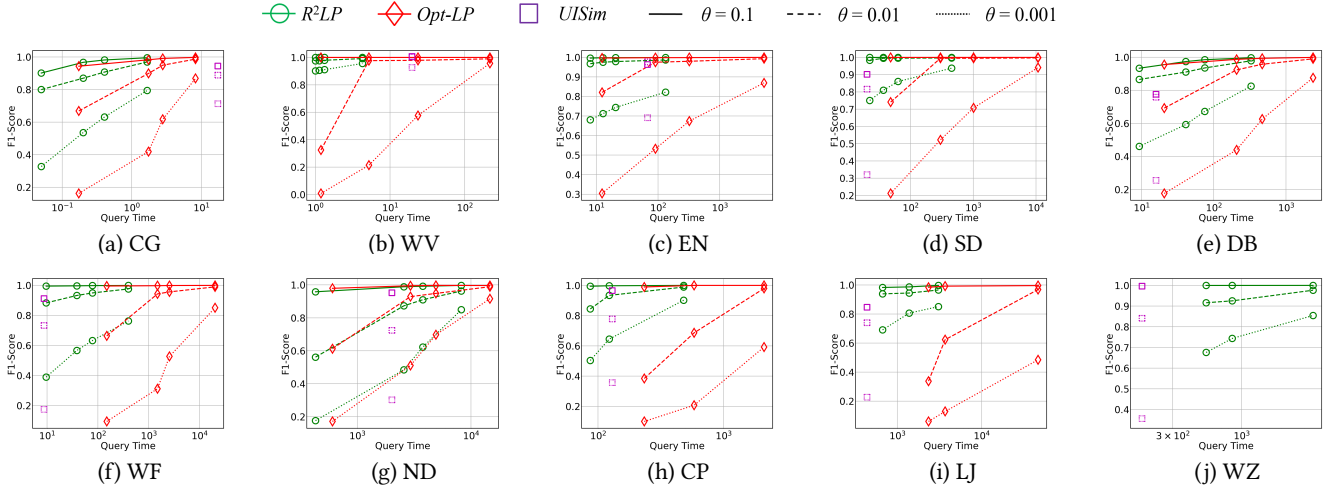


Figure 7: F1-score for threshold-based queries. No approximation bound is provided.

result accuracy on a pair of query node sets A and B . Specifically, we set $A = B = V$ for CG and WV, while for EN, SD, WF, and ND we set $|A| = |B| = 10^3$. Since the all-pair ground truth of DB is very sparse, we set $|A| = |B| = 10^4$. For three large graphs, we fix A as the 100 query nodes in generating ground truth, and set $|B| = 10^3$. The query nodes are sampled at random.

Table 6: Precision@5000 and query time ($\rho = 0.9$). The practical accuracy is underlined if it falls below ρ .

	UISim		Opt-LP		KSimJoin		R ² LP	
	Time	P@k	Time	P@k	Time	P@k	Time	P@k
CG	82.5	0.935	2.16	0.999	6.3	0.926	1.42	0.998
WV	308.2	0.992	110.3	0.999	405.9	0.994	24.6	0.998
EN	1189	0.993	443.7	0.997	OOM		88.0	0.996
SD	185.4	0.922	7682	0.996	OOM		2424	0.997
DB	84.0	<u>0.777</u>	1082	0.998	OOM		333.7	0.995
WF	56.0	<u>0.776</u>	10501	0.99	OOM		3488	0.995
ND	13448	0.959	4964	0.997	OOM		11749	0.993
CP	1416	0.95	2037	0.992	OOM		1190	0.994
LJ	5604	0.917	8586	0.964	OOM		8824	0.97
WZ	2375	0.904	Time Out		OOM		7171	1

We first compare the state-of-the-art algorithms by fixing $k = 5000$ and $\rho = 0.9$ with the results listed in Table 6. For *UISim*, we fix $\eta = 2$ and set *stopRea* as 10^{-4} for CG and WV, 10^{-3} for EN, and 10^{-2} for all other datasets. Although being fast, *UISim* mainly suffers from accuracy issues. On the other hand, *KSimJoin* has severe scalability problem on medium and large graphs in spite of its pruning strategy. We integrate *Opt-LP* and *R²LP* with *APTop-k* to answer top- k queries. For both of them, the query accuracy is guaranteed. *APTop-k+R²LP* is the fastest on a majority of graphs.

Next, we evaluate the query efficiency w.r.t. k (Figure 8). In most cases, *APTop-k+R²LP* outperforms *APTop-k+Opt-LP* by a significant margin. On three large graphs CP, LJ, and WZ, both algorithms lose the theoretical guarantee due to early termination. Fortunately, the practical accuracy of *APTop-k+R²LP* is always above ρ , whereas

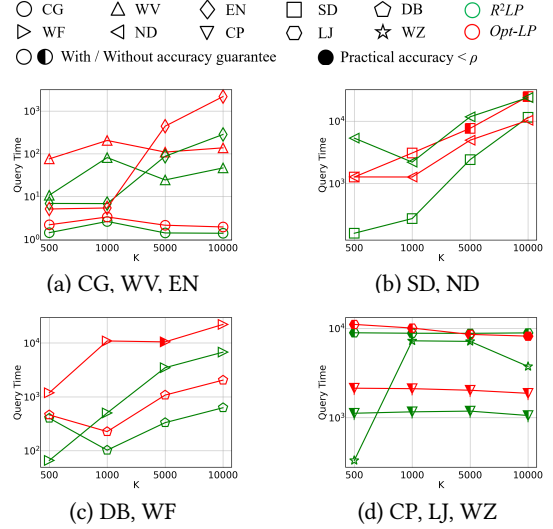


Figure 8: Query time vs. k , with $\rho = 0.9$.

APTop-k+Opt-LP has accuracy below ρ on LJ with $k = 10000$ and even fails on WZ.

6.5 Empirical Study of Computational Complexity

We conduct two sets of experiments to validate our theoretical analysis w.r.t. graph properties. We conduct threshold queries with $\theta = 0.01$ and top- k queries with $k = 1000$ following the settings in the above subsections. To illustrate the intrinsic problem complexities w.r.t. β , the PLB exponent of SimRank distribution, we first make sure the *practical accuracy* of *R²LP* slightly surpasses $\rho = 0.9$ by adjusting the error parameter ϵ . We then normalize the query time by approximating the numerator of *R²LP*'s complexity (see Theorem 3) and rescale the normalized values so that the value of CG equals 1. The normalized query time (Figure 9) only reflects the impact of ϵ , which is mainly affected by both ρ (fixed in our experiment) and graph properties, i.e., β , b_1 and b_2 . Generally

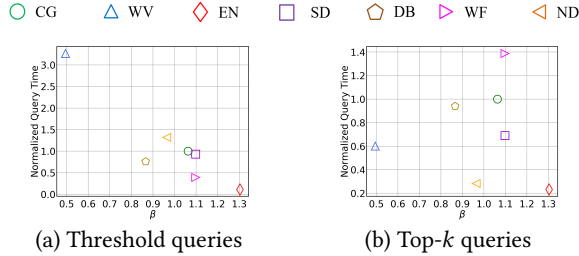


Figure 9: Query time vs. β .

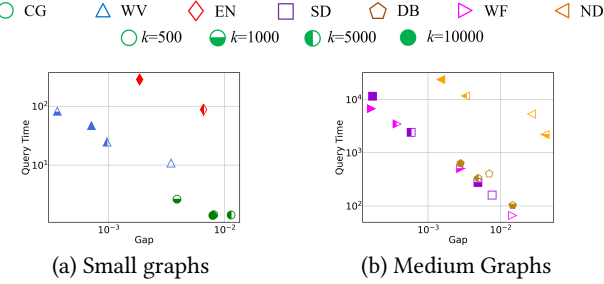


Figure 10: Query time vs. $s_{\lceil \rho k \rceil} - s_{k+1}$ for top- k queries.

speaking, the time is in negatively (resp. positively) correlated with β for threshold-based (resp. top- k) queries, which coincides with the derived computational complexity (Theorem 4 & 5). For top- k queries, the data points are more scattered because the impact of β is more complicated as shown by Theorem 5. Therefore, an additional analysis is conducted.

For top- k queries, our theoretical analysis suggests that the time complexity heavily relies on the gap between the $\lceil \rho k \rceil$ -th and the $(k + 1)$ -th largest SimRank values, as shown in Figure 10. Note that for a majority of studied graphs including CG, WV, DB and ND, the gap does not necessarily decrease as we increase k . This explains the phenomenon in Figure 8 that query time does not always increase with k .

6.6 Case Study

We adopt a very recent study [21] of graph neural networks (GNNs) to verify the effectiveness of our methods, where SimRank is treated as a new interpretation of GNNs. Instead of iteratively conducting neighborhood aggregation [14] which can be simplified as computing the Personalized PageRank matrix [3, 10], SIMGA [21] precomputes the SimRank matrix and applies it for aggregation. Specifically, we have $Z = \text{softmax}(S \cdot f_{\theta}(F, A))$, where $Z \in \mathbb{R}^{n \times d}$ is the embedding matrix, S is the SimRank matrix, and f_{θ} is a mapping function (e.g., an MLP) to transform node features F and adjacency matrix A into embeddings. SIMGA uses *Opt-LP* to compute S , and generally outperforms sophisticated GNNs on both homophily and heterophily graphs in terms of classification accuracy. We substitute *Opt-LP* with *R²LP* and demonstrate in Figure 11 the accuracy of SIMGA and the cost of computing S on a small dataset Cora ($n = 2.7K$, $m = 5.3K$) and a large dataset Pokec ($n = 1.6M$, $m = 30.6M$). It turns out that the accuracy is comparable to state-of-the-art GNNs [21], e.g., 82.30 ± 0.12 on Pokec. Meanwhile, our method achieves several times of speedup.

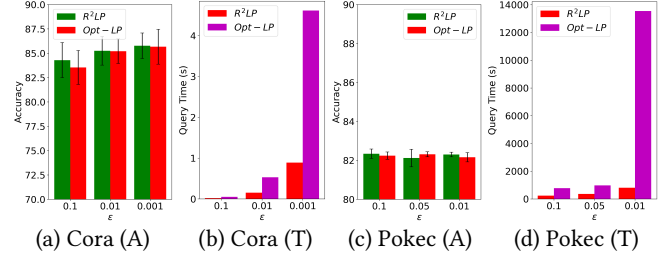


Figure 11: A case study of our approach on GNN. “A” stands for accuracy and “T” stands for query time.

7 OTHER RELATED WORK

The *Power Method* [12] is used to compute the ground truth SimRank values on small graphs (e.g., with tens of thousands of nodes), which guarantees a very small absolute error. To accelerate it on larger graphs, [24, 44] improves the algorithmic complexity by pruning redundant computation. Another line of work [11, 18, 41] studies all-pair SimRank computation on dynamic graphs. Several all-pair algorithms [9, 40–43] rely on a modified definition of SimRank, i.e. $S = cP^TSP + (1 - c)I$, which results in estimations without absolute error guarantee. A very recent study [45] employs divide-and-conquer to enhance state of the art [39], but is only able to improve the practical efficiency. Their contribution is orthogonal to ours and the implementation can be integrated naturally.

Most of the existing work for SimRank (e.g., [13, 16, 22, 23, 30, 31, 33, 36]) focuses on single-source settings, including *single-source SimRank computation* [13, 30, 31, 33, 36], *single-source top- k queries* [16, 22, 23], and *single-source thresholding queries* [23]. Various algorithms have been proposed based on random walk sampling [13, 30], graph traversal [16, 22, 23, 33], and local push [31]. Besides, a single-pair query for two nodes u and v computes an estimation of $s(u, v)$ with additive error up to ϵ . Recent algorithms [38] can answer single-pair queries within milliseconds for $\epsilon \geq 10^{-4}$.

8 CONCLUSION

This paper presents detailed empirical studies, new algorithms with better efficiency and approximation bounds, and novel theoretical analysis for two types of all-pair SimRank queries, namely, threshold-based and top- k similarity joins. We conduct a detailed comparison of state-of-the-art algorithms followed by the introduction of algorithm frameworks that theoretically guarantee the result accuracy for both queries, along with a more efficient all-pair algorithm inspired by the randomized local push. To gain a deep understanding of algorithm performance and the computational complexities of these problems, we propose theoretical analysis by leveraging the SimRank distribution and conduct extensive experimental studies. Our experiments show that the proposed algorithms significantly improve both query efficiency and accuracy, while our theoretical results nicely match the empirical analysis.

ACKNOWLEDGMENTS

This work was supported by the Fundamental Research Funds for the Central Universities (No. 2022RC034), NSFC under grant 62202037 and 61932001, and CCF-Tencent Open Fund under grant RAGR20220131.

REFERENCES

- [1] [n.d.]. <https://github.com/xinghun0525/R2LP>.
- [2] Jonathan W Berry, Luke K Fostvedt, Daniel J Nordman, Cynthia A Phillips, C Seshadhri, and Alyson G Wilson. 2014. Why do simple algorithms for triangle enumeration work in the real world?. In *Proceedings of the 5th conference on Innovations in theoretical computer science*. 225–234.
- [3] Aleksandar Bojchevski, Johannes Gasteiger, Bryan Perozzi, Amol Kapoor, Martin Blais, Benedek Rózemerczki, Michal Lukasik, and Stephan Günnemann. 2020. Scaling graph neural networks with approximate pagerank. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2464–2473.
- [4] Paolo Boldi, Marco Rosa, Massimo Santini, and Sebastiano Vigna. 2011. Layered label propagation: A multiresolution coordinate-free ordering for compressing social networks. In *Proceedings of the 20th international conference on World Wide Web*. 587–596.
- [5] Paolo Boldi and Sebastiano Vigna. 2004. The WebGraph Framework I: Compression Techniques. In *Proc. of the Thirteenth International World Wide Web Conference (WWW 2004)*. ACM Press, Manhattan, USA, 595–601.
- [6] Paweł Brach, Marek Cygan, Jakub Łkacki, and Piotr Sankowski. 2016. Algorithmic complexity of power law networks. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 1306–1325.
- [7] Moses Charikar, Kevin Chen, and Martin Farach-Colton. 2002. Finding frequent items in data streams. In *International Colloquium on Automata, Languages, and Programming*. Springer, 693–703.
- [8] Dániel Fogaras and Balázs Rácz. 2005. Scaling link-based similarity search. In *Proceedings of the 14th international conference on World Wide Web*. 641–650.
- [9] Yasuhiro Fujiwara, Makoto Nakatsuji, Hiroaki Shiokawa, and Makoto Onizuka. 2013. Efficient search algorithm for SimRank. In *2013 IEEE 29th International Conference on Data Engineering (ICDE)*. IEEE, 589–600.
- [10] Johannes Gasteiger, Aleksandar Bojchevski, and Stephan Günnemann. 2018. Predict then propagate: Graph neural networks meet personalized pagerank. *arXiv preprint arXiv:1810.05997* (2018).
- [11] Guoming He, Haijun Feng, Cuiping Li, and Hong Chen. 2010. Parallel SimRank computation on large graphs with iterative aggregation. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*. 543–552.
- [12] Glen Jeh and Jennifer Widom. 2002. SimRank: a measure of structural-context similarity. In *SIGKDD*. 538–543.
- [13] Minhao Jiang, Ada Wai-Chee Fu, and Raymond Chi-Wing Wong. 2017. READS: a random walk approach for efficient and accurate dynamic SimRank. *PVLDB* 10, 9 (2017), 937–948.
- [14] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
- [15] Jérôme Kunegis. 2013. Konect: the koblenz network collection. In *Proceedings of the 22nd international conference on world wide web*. 1343–1350.
- [16] Pei Lee, Laks VS Lakshmanan, and Jeffrey Xu Yu. 2012. On top-k structural similarity search. In *2012 IEEE 28th International Conference on Data Engineering*. IEEE, 774–785.
- [17] Jure Leskovec and Andrej Krevl. 2014. SNAP Datasets: Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data>.
- [18] Cuiping Li, Jiawei Han, Guoming He, Xin Jin, Yizhou Sun, Yintao Yu, and Tianyi Wu. 2010. Fast computation of simrank for static and dynamic information networks. In *Proceedings of the 13th International Conference on Extending Database Technology*. 465–476.
- [19] Lina Li, Cuiping Li, Hong Chen, and Xiaoyong Du. 2013. Mapreduce-based SimRank computation and its application in social recommender system. In *2013 IEEE international congress on big data*. IEEE, 133–140.
- [20] Ruiqi Li, Xiang Zhao, Haichuan Shang, Yifan Chen, and Weidong Xiao. 2017. Fast top-k similarity join for SimRank. *Inf. Sci.* 381 (2017), 1–19. <https://doi.org/10.1016/j.ins.2016.10.042>
- [21] Haoyu Liu, Ningyi Liao, and Siqiang Luo. 2023. SIMGA: A Simple and Effective Heterophilous Graph Neural Network with Efficient Global Aggregation. *arXiv preprint arXiv:2305.09958* (2023).
- [22] Yu Liu, Bolong Zheng, Xiaodong He, Zhewei Wei, Xiaokui Xiao, Kai Zheng, and Jiaheng Lu. 2017. Probesim: scalable single-source and top-k simrank computations on dynamic graphs. *arXiv preprint arXiv:1709.06955* (2017).
- [23] Yu Liu, Lei Zou, Qian Ge, and Zhewei Wei. 2020. SimTab: accuracy-guaranteed simrank queries through tighter confidence bounds and multi-armed bandits. *Proceedings of the VLDB Endowment* 13, 12 (2020), 2202–2214.
- [24] Dmitry Lizorkin, Pavel Velikhov, Maxim Grinev, and Denis Turdakov. 2008. Accuracy Estimate and Optimization Techniques for SimRank Computation. *Proc. VLDB Endow.* 1, 1 (aug 2008), 422–433. <https://doi.org/10.14778/1453856.1453904>
- [25] Peter Lofgren and Ashish Goel. 2013. Personalized pagerank to a target node. *arXiv preprint arXiv:1304.4658* (2013).
- [26] Juan Lu, Zhiguo Gong, and Xuemin Lin. 2016. A novel and fast simrank algorithm. *IEEE transactions on knowledge and data engineering* 29, 3 (2016), 572–585.
- [27] Linyuan Lü and Tao Zhou. 2011. Link prediction in complex networks: A survey. *Physica A: statistical mechanics and its applications* 390, 6 (2011), 1150–1170.
- [28] Takanori Maehara, Mitsuru Kusumoto, and Ken-ichi Kawarabayashi. 2014. Efficient simrank computation via linearization. *arXiv preprint arXiv:1411.7228* (2014).
- [29] Takanori Maehara, Mitsuru Kusumoto, and Ken-ichi Kawarabayashi. 2015. Scalable simrank join algorithm. In *2015 IEEE 31st International Conference on Data Engineering*. IEEE, 603–614.
- [30] Yingxia Shao, Bin Cui, Lei Chen, Mingming Liu, and Xing Xie. 2015. An efficient similarity search framework for SimRank over large dynamic graphs. *PVLDB* 8, 8 (2015), 838–849.
- [31] Jieming Shi, Tianyuan Jin, Renchi Yang, Xiaokui Xiao, and Yin Yang. 2020. Real-time Index-Free Single Source SimRank Processing on Web-Scale Graphs. *Proc. VLDB Endow.* 13, 7 (2020), 966–978. <https://doi.org/10.14778/3384345.3384347>
- [32] Wenbo Tao, Minghe Yu, and Guoliang Li. 2014. Efficient Top-K SimRank-based Similarity Join. *Proc. VLDB Endow.* 8, 3 (2014), 317–328. <https://doi.org/10.14778/2735508.2735520>
- [33] Boyu Tian and Xiaokui Xiao. 2016. SLING: A Near-Optimal Index Structure for SimRank. In *SIGMOD*. 1859–1874.
- [34] Anton Tsitsulin, Davide Mottin, Panagiotis Karras, and Emmanuel Müller. 2018. Verse: Versatile graph embeddings from similarity measures. In *Proceedings of the 2018 world wide web conference*. 539–548.
- [35] Hanzhi Wang, Zhewei Wei, Junhao Gan, Sibowang, and Zengfeng Huang. 2020. Personalized pagerank to a target node, revisited. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 657–667.
- [36] Hanzhi Wang, Zhewei Wei, Yu Liu, Ye Yuan, Xiaoyong Du, and Ji-Rong Wen. 2021. ExactSim: benchmarking single-source SimRank algorithms with high-precision ground truths. *The VLDB Journal* 30, 6 (2021), 989–1015.
- [37] Sibowang, Renchi Yang, Xiaokui Xiao, Zhewei Wei, and Yin Yang. 2017. FORA: simple and effective approximate single-source personalized pagerank. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 505–514.
- [38] Yue Wang, Lei Chen, Yulin Che, and Qiong Luo. 2019. Accelerating pairwise simrank estimation over static and dynamic graphs. *The VLDB Journal* 28, 1 (2019), 99–122.
- [39] Yue Wang, Xiang Lian, and Lei Chen. 2018. Efficient simrank tracking in dynamic graphs. In *2018 IEEE 34th international conference on data engineering (ICDE)*. IEEE, 545–556.
- [40] Weiren Yu, Xuemin Lin, and Wenjie Zhang. 2013. Towards efficient SimRank computation on large networks. In *2013 IEEE 29th international conference on data engineering (ICDE)*. IEEE, 601–612.
- [41] Weiren Yu, Xuemin Lin, and Wenjie Zhang. 2014. Fast incremental simrank on link-evolving graphs. In *2014 IEEE 30th international conference on data engineering*. IEEE, 304–315.
- [42] Weiren Yu, Xuemin Lin, Wenjie Zhang, Lijun Chang, and Jian Pei. 2013. More is simpler: Effectively and efficiently assessing node-pair similarities based on hyperlinks. *Proceedings of the VLDB Endowment* 7, 1 (2013), 13–24.
- [43] Weiren Yu and Julie A McCann. 2015. Efficient partial-pairs simrank search on large networks. *Proceedings of the VLDB Endowment* 8, 5 (2015), 569–580.
- [44] Weiren Yu, Wenjie Zhang, Xuemin Lin, Qing Zhang, and Jiajin Le. 2012. A space and time efficient algorithm for SimRank computation. *World Wide Web* 15, 3 (2012), 327–353.
- [45] Liangfu Zhang, Cuiping Li, Xue Zhang, and Hong Chen. 2023. Hierarchical All-Pairs SimRank Calculation. In *International Conference on Database Systems for Advanced Applications*. Springer, 252–268.
- [46] Weiguo Zheng, Lei Zou, Lei Chen, and Dongyan Zhao. 2017. Efficient SimRank-Based Similarity Join. *ACM Trans. Database Syst.* 42, 3 (2017), 16:1–16:37. <https://doi.org/10.1145/3083899>
- [47] Weiguo Zheng, Lei Zou, Yansong Feng, Lei Chen, and Dongyan Zhao. 2013. Efficient SimRank-based Similarity Join Over Large Graphs. *Proc. VLDB Endow.* 6, 7 (2013), 493–504. <https://doi.org/10.14778/2536349.2536350>
- [48] Fanwei Zhu, Yuan Fang, Kai Zhang, Kevin Chang, Hongtai Cao, Zhen Jiang, and Minghui Wu. 2021. Unified and Incremental SimRank: Index-free Approximation with Scheduled Principle. *IEEE Transactions on Knowledge and Data Engineering* (2021).