Efficient Algorithms for Group Hitting Probability Queries on Large Graphs

Qintian Guo, Dandan Lin, Sibo Wang, Raymond Chi-Wing Wong, and Wenqing Lin

***Abstract*—Given a source node s and a target node t, the hitting probability tells us how likely an α-terminating random walk (which stops with probability α at each step) starting from s can hit t before it stops. This concept originates from the hitting time, a classic concept in random walks. In this paper, we focus on the group hitting probability (GHP) where the target is a set of nodes, measuring the node-to-group structural proximity. For this group version of the hitting probability, we present efficient algorithms for two types of GHP queries: the pairwise query which returns the GHP value of a target set T with respect to (w.r.t.) a source node s, and the top-k query which returns the top-k target sets with the largest GHP value w.r.t. a source node s. We first develop an efficient algorithm named SAMBA for the pairwise query, which is built on a group local push algorithm tailored for GHP, with rigorous analysis for correctness. Next, we show how to speed up SAMBA by combining the group local push algorithm with the Monte Carlo approach, where GHP brings new challenges as it might need to consider every hop of the random walk. We tackle this issue with a new formulation of the GHP and show how to provide approximation guarantees with a detailed theoretical analysis. With SAMBA as the backbone, we develop an iterative algorithm for top-k queries, which adaptively refines the bounds for the candidate target sets, and terminates as soon as it meets the stopping condition, thus saving unnecessary computational costs. We further present an optimization technique to accelerate the top-k query, improving its practical performance. Extensive experiments show that our solutions are orders of magnitude faster than their competitors.**

***Index Terms*—Data mining, graph algorithms, graphs and networks.**

I. Introduction

MEASURING node-to-node similarity or proximity in graphs is a fundamental tool for various graph-based mining applications [1], [2], [3], [4], [5], and has been recognized as an important research problem in the data mining community. One common approach to capture "similarity" in the literature is *random-walk-based approaches* due to its effectiveness and efficiency, resulting in a lot of research studies [6], [7], [8], [9], [10], [11] in the communities of database and data mining. Among them, *hitting time* [12], a classic concept in random walk, finds many real-world applications, such as link prediction [3], product recommendation [13], query suggestions [14] and collaborative filtering [15].

Given a graph G, hitting time tells us the expected number of steps a random walk takes to hit a target node t from a source node s *for the first time*. The *fewer* steps a random walk from s takes to hit t, the *more relevant* (similar) node t to s is from the perspective of s. In [16], [17], it has been proved that for an α-terminating random walk, the hitting time $h(s,t)$ from s to t can be obtained by exploiting the probability that an α-terminating random walk starting from node s hits t before terminating, denoted as $f(s,t)$. That is, if $f(s,t)$ is higher, a random walk from s takes fewer steps to hit t for the first time. Hence, solving hitting probability queries is equivalent to solving the hitting time problem.

In this paper, we consider a generalized case of the hitting probability problem where the target is a node-set T, called the *Group Hitting probability* (GHP). It is easy to realize that the original hitting probability $f(s,t)$ of a target node t is a special case of our GHP problem, where the target set T contains only one node t. The motivation for adopting a group target setting to measure node-to-group proximity finds validation within the existing literature. In [18], Grady presents a group-based algorithm to address the image segmentation problem. In Grady's solution, it treats images as graphs, with pixels as nodes and edges connecting adjacent pixels. Each edge is assigned a real-valued weight, indicating the likelihood of a random walker moving across that edge. Given a small subset of pixels with pre-defined labels, it computes the probability that a random walker, starting from each unlabeled pixel, will reach one of the labeled pixel groups. This information is then utilized to assign the pixel to the group label that corresponds to the highest calculated probability. This process facilitates the attainment of a superior image segmentation outcome. In [13], given a graph G with m nodes that share the same event Q, Guan et al. propose an approach that employs the average proximity between a node in these m nodes and the remaining $m-1$ nodes, to assess whether any structural correlation exists between the event Q and the graph G. The essence of this approach lies in accurately measuring the proximity between a node and a set or group of

Manuscript received 22 March 2023; revised 29 October 2023; accepted 23 December 2023. Date of publication 9 January 2024; date of current version 10 June 2024. This work was supported in part by NSFC under Grant U1936205, in part by Hong Kong RGC ECS under Grant 24203419, in part by RGC GRF under Grant 14217322, in part by Hong Kong RGC CRF under Grant C4158-20G, and in part by Hong Kong ITC ITF under Grant MRP/071/20X. Recommended for acceptance by G. Wang. *(Qintian Guo and Dandan Lin are co-first authors.) (Corresponding author: Dandan Lin.)*

Qintian Guo and Sibo Wang are with The Chinese University of Hong Kong, Hong Kong (e-mail: qtguo@se.cuhk.edu.hk; swang@se.cuhk.edu.hk).

Dandan Lin is with the Shenzhen Institute of Computing Sciences, Shenzhen, Guangdong 518172, China (e-mail: lindandan@sics.ac.cn).

Raymond Chi-Wing Wong is with The Hong Kong University of Science and Technology, Hong Kong (e-mail: raywong@cse.ust.hk).

Wenqing Lin is with Tencent Inc., Shenzhen, Guangdong 518054, China (e-mail: edwlin@tencent.com).

Digital Object Identifier 10.1109/TKDE.2023.3349164

1041-4347 © 2024 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See https://www.ieee.org/publications/rights/index.html for more information.

Authorized licensed use limited to: Hong Kong University of Science and Technology. Downloaded on July 08,2024 at 11:53:59 UTC from IEEE Xplore. Restrictions apply.

nodes. Guan et al. demonstrate that group hitting time serves as a suitable choice, effectively identifying events highly correlated with the graph structure.

Other applications of the node-to-group proximity measure include: (1) *Group Suggestions:* In practical scenarios, both Facebook and LinkedIn offer services to provide personalized recommendations, assisting users in discovering groups, as stated on their websites [19], [20]. By recommending pertinent groups to join, social network platforms can enhance user engagement, encourage community interactions, and subsequently increase the duration users spend on the platform. Using our GHP measure, a platform can pinpoint groups that are in close proximity to a user, based on the graph's connectivity information, and then make a recommendation. (2) *Collaborative Filtering:* Recommendation systems can estimate a user's structural proximity to groups of users with similar tastes. If a user is proximate to a group that has expressed a liking for a certain item, that item might be recommended to the user. (3) *Web Network Analysis:* By gauging the proximity of a web page (node) to groups of web pages with established categories, search engines can refine categorization and heighten the relevance of search results.

The first hitting constraint also makes the GHP more reasonable when considering the proximity between a node and a group, as it fundamentally views the group as a singular entity. This contrasts with other metrics like personalized PageRank (PPR) [6], [21], which only consider the probability that an α-terminating random walk stops at a specific node. To obtain the proximity of a group with respect to the source node, one might aggregate the PPR scores of nodes within that group. However, as the case study in Section V-C shows, the PPR-aggregation-based metric is still inferior to the GHP in proximity-based group recommendation. Thus, it motivates us to probe into the GHP-based query.

In this paper, we first study the *pairwise query*. This query takes as input a source node s, a target set T, and outputs the value $f(s, T)$. Such a pairwise query can be employed to measure the similarity between a node and a specified community or group. However, in certain scenarios, there may be a desire to pinpoint a limited number of group sets that are of utmost relevance to a particular node. Consider the application of group suggestions in social networks: instead of suggesting an extensive list of groups, social networking platforms might prefer to highlight a limited number of groups that possess the highest appeal to the user, thereby enhancing the recommendation success rate. Motivated by this, we delve into the second type of GHP-based query: the *top-k query*. This query takes a source node s as input and retrieves the k target sets with the top GHP values with respect to s.

Deriving an exact solution for these two types of queries causes prohibitive computational costs. Specifically, the pairwise query requires $O(n^{2.37})$ time since it needs to compute the inverse of an $n \times n$ matrix. For top-k queries, it is even more expensive, which requires computing the matrix inversion much more times, depending on the number of groups. Thus, we aim to present efficient approximate algorithms with performance guarantees.

Contributions: Our contributions are presented as follows:

1) We first develop an efficient algorithm called *SAMBA* to answer pairwise queries. Though SAMBA[1] is inspired by BiPPR [22], which is designed for the approximate pairwise PPR problem, as we will see in Section III, it is far from a facile adaption.
2) With SAMBA as the backbone, we further develop a framework called *PING* for the top-k GHP queries. PING runs in an iterative manner to gradually refine the lower and upper bounds for the estimated GHP values. The algorithm is designed to terminate as soon as the stopping condition is met, thus avoiding unnecessary computational costs. In addition, we present an optimization technique further to accelerate the computation of PING.
3) We run extensive experiments on real datasets to demonstrate the efficiency of our algorithms. Our experiments show that our algorithms SAMBA and PING take much less running time than their competitors on all datasets.

II. Preliminary

A. Problem Definition

Let $G = (V, E)$ be a directed graph with n nodes and m edges. For an undirected input graph, we convert it to a directed graph by treating each undirected edge as two directed edges in opposing directions [6], [8]. Given a source node $s \in V$ and a terminating probability α, an *α-terminating random walk* starts from s, and then, at each step, it chooses one of the following two options: either (i) terminates at the current node with α probability; or (ii), with $1 - \alpha$ probability, moves to an out-neighbor of the current node uniformly at random. From the definition, it immediately follows that the probability of such a random walk exceeding L steps is bounded by $(1 - \alpha)^L$. Consequently, the occurrence of a random walk spanning a large number of steps becomes exceedingly unlikely, making it negligible.

Definition 1 (Hitting probability): Given a graph G, let $(X_i)_{i=0}^{\tau}$ be an α-terminating random walk on G. Then for any two nodes s and t, we define the hitting probability as

$$f(s, t) = \Pr\left[\{t\} \cap \{X_i\}_{i=0}^{\tau} \neq \emptyset | X_0 = s\right]$$

where $\{X_i\}_{i=0}^{\tau}$ is the node set visited by random walk $(X_i)_{i=0}^{\tau}$.

To explain, the hitting probability $f(s, t)$ is the probability that an α-terminating random walk starting from s could hit t before it terminates. Obviously when $s = t$, $f(s, t) = 1$ since a node s always hits itself at the very beginning. In the remainder of this paper, we omit the term *α-terminating* if the context is clear.

Example 1: Consider the toy graph in Fig. 1(a). We simulate 7 random walks starting from s, shown in Fig. 1(b). Note that among them there are 4 (resp. 2) random walks that hit target node v_1 (resp. v_5). Thus, from the definition, the hitting probability $f(s, v_1)$) and $f(s, v_5)$ is estimated as $4/7$ and $2/7$, respectively.

Definition 2 (Group hitting probability): Given a graph G, let $(X_i)_{i=0}^{\tau}$ be an α-terminating random walk on G. Then for node

[1]Random Sampling with Backward Local Push

Authorized licensed use limited to: Hong Kong University of Science and Technology. Downloaded on July 08,2024 at 11:53:59 UTC from IEEE Xplore. Restrictions apply.

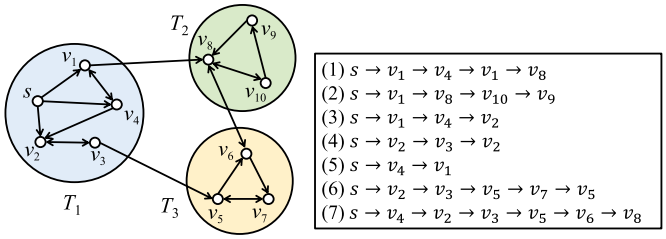

(a) Toy graph (b) 7 random walks from node s

Fig. 1. (a) Toy graph G whose nodes are partitioned into three target sets T_1, T_2, and T_3. (b) 7 random walks starting from s.

s and set of nodes $T \neq \emptyset$, the group hitting probability is

$$f(s,T) = \Pr\left[T \cap \{X_i\}_{i=0}^{\tau} \neq \emptyset | X_0 = s\right].$$

In this paper, we focus on the *group hitting probability* (GHP) $f(s,T)$. Formally, the GHP $f(s,T)$ is the probability that a random walk from the source node s hits any one of the nodes in T during the traversal. It is easy to see that the hitting probability $f(s,t)$ of a single target node t is a special case of the GHP $f(s,T)$ with $T = \{t\}$.

Example 2: Consider the toy graph shown in Fig. 1(a) and the random walks in Fig. 1(b), where the nodes are partitioned into 3 target sets, namely $T_1 = \{s, v_1, v_2, v_3, v_4\}$, $T_2 = \{v_8, v_9, v_{10}\}$ and $T_3 = \{v_5, v_6, v_7\}$. First, we could know that $f(s,T_1)$ is always 1 since all the walks hit the target set T_1 at the very beginning. Note that the target set T_2 is hit by the walks with index 1, 2, and 7, thus yielding an estimated GHP $\hat{f}(s,T_2) = 3/7$.

To the best of our knowledge, we are the first to consider the approximate query for the group hitting probability with performance guarantee and aim to provide inspirational insights for future research. In this paper, we study the approximate pairwise GHP query and the approximate top-k GHP query, which are defined as follows.

Definition 3 (Approximate Pairwise GHP Query): Given a source node s, a target set T, a relative error ϵ, a threshold δ, and a failure probability p_f, the approximate pairwise GHP query returns the estimated GHP $\hat{f}(s,T)$ such that

$$|f(s,T) - \hat{f}(s,T)| \leq \epsilon \cdot f(s,T), \text{ if } f(s,T) \geq \delta, \quad (1)$$

holds with at least $1 - p_f$ probability.

Definition 4 (Approximate Top-k GHP query): Given a graph $G(V,E)$, a set $\mathcal{T}$ of target sets such that $\mathcal{T} = \{T_1, T_2, \ldots, T_{|\mathcal{T}|}\}$ where $|\mathcal{T}|$ is the number of target sets, a source node s, a positive integer k, a relative error ϵ, and a threshold δ, the approximate top-k GHP query returns a set $\mathcal{R} = \langle T'_1, \ldots, T'_k \rangle$ such that for any $i \in \{1, \ldots, k\}$ with $f(s,T^*_i) \geq \delta$:

$$\hat{f}(s,T'_i) \geq (1-\epsilon) \cdot f(s,T'_i) \quad (2)$$

$$f(s,T'_i) \geq (1-\epsilon) \cdot f(s,T^*_i) \quad (3)$$

hold with at least $1 - p_f$ probability, where T^*_i is the set whose exact GHP w.r.t. node s is the i-th largest.

Note that (2) guarantees the accuracy of the estimated GHP values of returned top-k sets, while (3) ensures that the returned i-th set T'_i should have an exact GHP *close* to the exact i-th largest GHP $f(s,T^*_i)$.

In this paper, we make utilization of the Chernoff bounds [23] to provide concentration guarantee.

Lemma 1 (Chernoff Bounds [23]): Suppose that X_i are independent random variables satisfying $X_i \in [0,1]$ for $1 \leq i \leq n$. Let $X = \sum_{i=1}^{n} X_i$. Then, we have

$$\Pr\left[X - E[X] \leq -\lambda\right] \leq \exp\left(-\frac{\lambda^2}{2 \cdot E[X]}\right),$$

$$\Pr\left[X - E[X] \geq \lambda\right] \leq \exp\left(-\frac{\lambda^2}{2 \cdot E[X] + \frac{2}{3}\lambda}\right).$$

B. Existing Solutions

We first introduce the Monte Carlo (MC) approach for computing the GHP, and next, we present the existing local-update techniques (i.e., Forward Push and Backward Push) that are widely used for other random walk based queries [6], [22], [24].

MC Method: To estimate the GHP value $f(s,T)$, we may simulate a sufficient number of random walks starting from s, and then estimate $f(s,T)$ as the fraction of random walks that hit at least one node in T, denoted by $\hat{f}(s,T)$. In order to obtain an estimation with a guarantee of at most ϵ relative error and probability of at least $1 - 1/n$, the number of random walks required, denoted as ω, is $O(\frac{\log n}{\delta \cdot \epsilon^2})$. This conclusion can be derived by concentration bounds (e.g., Chernoff bounds). When $\delta = O(1/n)$, the running time is $O(\frac{n \log n}{\epsilon^2})$. When it comes to the top-k GHP queries, for each target set $T \in \mathcal{T}$, the MC method computes an approximate GHP $\hat{f}(s,T)$ and then returns the top-k sets with the highest estimated GHP as the results. The time complexity is still $O(\frac{n \log n}{\epsilon^2})$. However, on large graphs, the MC method is not efficient enough. In particular, answering the top-k GHP query on the Orkut dataset, which comprises 3 million nodes, takes several hours to complete.

Forward Push [25]: Forward Push is usually used for accelerating the single source PPR query (which approximates the PPR values of each node $t \in V$ w.r.t. a source s) and top-k PPR query [6], [7], [26]. In particular, for each $t \in V$, it maintains a reserve $\pi^f(s,t)$ (i.e., the amount of PPR values that have been propagated to node t from s) and a residue $r^f(s,v)$ (i.e., the amount of PPR values that are "temporarily" held by node v and waiting for further propagation from node v to other nodes). Initially, Forward Push assigns unit PPR value to the residue of source node s, and then, continually updates the reserves and residues by performing forward push operations on each node v with non-zero residue (which distributes the residue of v evenly to its out-neighbours) [27]. When the algorithm finishes, the final reserve of each $t \in V$ could be considered as the approximate PPR value.

However, this technique does not work for the GHP. For ease of explanation, we assume the target set T includes only one node t. The Forward Push algorithm might perform forward push operations on node t multiple times during the propagation process, which puts computations of GHP in trouble. In particular, after performing a forward push operation on node t, the residue of node t is distributed to out-neighbours of node t, and

Authorized licensed use limited to: Hong Kong University of Science and Technology. Downloaded on July 08,2024 at 11:53:59 UTC from IEEE Xplore. Restrictions apply.

then a few push iterations later, it might happen that a portion of the residue mass returns to t, resulting in double counting of the returning probability mass. Therefore, to represent this feature, we say the GHP is ***target-orient***. We still take Fig. 1(a) as an example. Consider node v_1 is the target node. The residue of v_1 comes from source s and v_4, and then, this residue will be distributed to its out-neighbours v_4 and v_8. Since the residue of v_4 is non-zero, it will be pushed to v_1 again and consequently, this probability mass is counted more than once. One might consider a simple solution to correct the above error: for a target node t, Forward Push does not distribute any residue from node t. From the definition of the GHP, it is easy to understand that such an adaption is correct. Therefore, we will take it as a competitor in our experimental evaluation and make a comparison with our proposed solution.

Backward Push [8]: Unlike Forward Push, Backward Push propagates the PPR value of a target node t via the reverse direction of each edge in the graph. The underlying rationale of Backward Push lies in a recursive definition where the PPR $\pi(s,t)$ equals a weighted sum of the PPR values $\pi(s,v)$, where v's are the in-neighbors of node t. The invariant is presented as follows: $\pi(s,t) = \alpha \cdot \mathbb{I}_s(t) + \sum_{v \in \mathcal{N}_T^-(t)} \frac{(1-\alpha)}{d^+(v)} \cdot \pi(s,v)$, where $\mathbb{I}_s(t)$ is an indicator function that is 1 if $t = s$ and otherwise 0. Backward push technique, however, cannot be directly applied to the GHP problem, as there exists no such a recursion for GHP. To explain, consider the simple case where the target set T contains only one target node $T = \{t\}$. Intuitively, the hitting probability $f(s,t)$ of target node t is related to the number of arrivals at t's in-neighbors (say v) in the traversal of a random walk, since for each arrival of in-neighbor node v, the walk could hit t in the next step with a certain probability. However, by the GHP definition, $f(s,v)$ is the probability that a random walk could hit node v before terminating, rather than how many times it hits v.

In the case where target set T contains more than one node, it is not correct to sum up $f(s,v_i)$ for all $v_i \in T$ to get the GHP $f(s,T)$. To explain, consider the second random walk in Fig. 1(b). This random walk first hits node $v_8 \in T_2$ and then node $v_{10} \in T_2$ and finally node $v_9 \in T_2$. If we simply sum up $f(s,v_8)$, $f(s,v_9)$ and $f(s,v_{10})$ to compute the GHP $f(s,T_2)$, the second random walk will be counted triple times, yielding a wrong answer. This is in contrast to the group setting in the PPR problem, which could be computed by summing up all the PPRs $\pi(s,v_i)$ for each node $v_i \in T$. It is because the PPR just takes into account the terminating node of a random walk, and thus when simulating a random walk, the events of stopping at node v_8 and stopping at node v_9 are disjoint.

C. Related Work

In graph analysis, measuring structural proximity between two graph objects based on graph topology is a fundamental problem that has been studied for decades. A significant category in this realm is the random-walk-based proximity measures. Personalized PageRank and SimRank are two notable representatives of this category. However, both focus on node-to-node proximity, which is distinct from the node-to-group case we explored in this paper.

Personalized PageRank (PPR) is first introduced by [28] to measure personalized views of importance between node pairs. The PPR score $\pi(s,t)$ between a source node s and a target node t represents the probability that an α-terminating random walk, starting from node s, concludes at node t. Given its efficacy in various graph mining applications, numerous studies aim to devise efficient algorithms for PPR computation. Andersen et al. [8] introduce the Forward Push algorithm for single-source PPR queries, though it lacks an accuracy guarantee. The Backward Push algorithm is presented in [25] to cater to single-target PPR queries. Lofgren et al. [22] combine the Backward Push approach with the Monte Carlo method, resulting in BiPPR, which addresses pairwise PPR queries. In [6], Wang et al. formulate an index-free algorithm named FORA, offering solutions for both single-source and top-k PPR queries with assured accuracy. Wu et al. [29] blend Power Iteration with Forward Push to create the SpeedPPR algorithm, tailored for single-source PPR queries. Nonetheless, all these methods do not apply to our GHP query.

The SimRank score $s(u,v)$ for nodes u and v, as proposed by Jeh et al. in [30], is recursively defined as:

$$s(u,v) = \begin{cases} 1, & \text{if } u = v \\ \frac{c}{|\mathcal{I}(u)| \cdot |\mathcal{I}(v)|} \sum_{u' \in \mathcal{I}(u)} \sum_{v' \in \mathcal{I}(v)} s(u',v'), & \text{otherwise} \end{cases}$$

where $\mathcal{I}(u)$ denotes the in-neighbors of node u, and c is a constant between 0 and 1. Tian et al. [31] interpret the SimRank concept $s(u,v)$ as the probability that two $(1-\sqrt{c})$-terminating random walks, starting from nodes u and v respectively, intersect with each other. They introduce an index-based algorithm, SLING, to address the single-source and top-k queries. In [32], Zhe et al. harness the reverse PageRank distribution of the input graph to formulate the index-based algorithm, PRSim, and demonstrate that it attains sub-linear time complexity for power-law graphs. ExactSim [33] offers high-precision single-source and top-k SimRank results for large graphs. Analogous to PPR, these methods cannot be applied to our GHP query.

Besides random-walk-based measures, other categories of proximity measures in the literature include common neighbors, shortest paths, and effective conductance. Let $N(u)$ represent the set of neighbors of node u. The Jaccard similarity, defined as $\frac{|N(u) \cap N(v)|}{|N(u) \cup N(v)|}$, proves useful in graph structural clustering [34]. The Katz measure [35] postulates that the more paths exist from node s to node t, and the shorter these paths, the greater the similarity of t to s. Koren et al. [36] employ cycle-free effective conductance to evaluate node proximity within a graph. These measures are distinct from our GHP problem.

One potential application of our GHP proximity measure is community or group recommendations. Most existing works on group recommendations utilize the Collaborative Filtering (CF) technique, which typically requires user history records or profiles. In [37], Chen et al. introduce the CCF model, considering both community-user co-occurrences and community-description co-occurrences, and derive the joint probability distribution over community, user, and description. With extensive profile data, Sharma [38] adopts a content-based approach and proposes a probabilistic latent preference model (Pairwise PLSI) for group recommendations on LinkedIn. Wang et al. [39] leverage production and consumption engagement

Authorized licensed use limited to: Hong Kong University of Science and Technology. Downloaded on July 08,2024 at 11:53:59 UTC from IEEE Xplore. Restrictions apply.

records to determine the user-group affinity and introduce a time-dependent matrix factorization model for recommending groups to users. These methods stand in contrast to our application of group recommendation based on structural proximity. In fact, as highlighted in our experimental evaluation, our GHP solution exhibits high scalability. Resultantly, the research community can immediately reap benefits: utilizing our solution to pre-prune groups with minimal structural relevance to the user. This reduces the pool of candidate groups for other existing algorithms, thereby enhancing their scalability.

III. Pairwise GHP Query

In this section, we present *SAMBA*, an efficient algorithm for pairwise GHP queries, which will also serve as the backbone of our top-k algorithm in Section IV. As mentioned earlier, the target-oriented nature of GHP makes it challenging to formulate a correct local push algorithm. To tackle these issues, we first bring in a new concept *T-absorbed reachability* and then make a connection between GHP and the T-absorbed reachability. With the new tool, we derive a recursive definition for T-absorbed reachability from scratch and propose a group local push algorithm for T-absorbed reachability in Section III-B, which ensures the correctness for computing $f(s,T)$. Finally, we show how to combine MC with group local push to answer pairwise queries in Section III-C.

A. The T-Absorbed Reachability

Let $\Pr[s \rightsquigarrow v|L]$ be the probability that a random walk from s reaches v at the L-th hop. Given a target set T, we define the L hop *T-absorbed visiting probability* $\Pr[s \rightsquigarrow v|L,T]$ of node v w.r.t. node s as the probability that a random walk from s reaches v at the L-th hop (not necessary to stop at the L-th hop) but with a constraint that no node in T is visited unless a node in T is visited at the L-th hop (in this case $v \in T$). Now, the T-absorbed reachability $r_T(s,v)$ of node v w.r.t. s is defined as the sum of the T-absorbed visiting probability of node v from node s at all possible number of hops, namely $L \in [0,+\infty)$, such that

$$r_T(s,v) = \sum_{L=0}^{+\infty} \Pr[s \rightsquigarrow v|L,T].$$

For each node $v \notin T$, we can know that the value $r_T(s,v)$ is the expected number of times that a random walk visits v before reaching any node in the target set T. A special case is $v \in T$. In this case, v is the first node in T that is reached by this random walk (and then it terminates). Thus, if $v \in T$, $r_T(s,v) \leq 1$, which can be interpreted as the probability that node v is the first one among nodes in T visited by the random walk. Recall that the hitting time $h(s,t)$ is the expected number of steps for a random walk from source node s to hit node t for the first time, while the hitting probability $f(s,t)$ is the probability that the random walk from s can hit node t before it terminates.

Subsequently, we define a T-absorbed reachability $r_T(s,T)$ of a target set T w.r.t. s as the sum of $r_T(s,v)$ for each $v \in T$, i.e., $r_T(s,T) = \sum_{v \in T} r_T(s,v)$. From the definition of GHP,

Theorem 1: For a target set T, $f(s,T) = r_T(s,T)$.

The theorem allows us to build a bridge between GHP and T-absorbed reachability. We can estimate $f(s,T)$ by computing $r_T(s,T)$. To estimate the T-absorbed reachability, we define the T-Absorbed Random Walk (T-ARW) that runs as follows: it starts from node s, and at each step: *(i)* it terminates *if the current node is a node t in T*; *(ii)* or terminates *with α probability*; *(iii)* otherwise, it moves to an out-neighbour of the current node *with $1-\alpha$ probability*. Then we can compute the T-absorbed reachability of a node v by exploiting the Monto Carlo approach.

B. Group Local Push Algorithm

Next, we demonstrate how to carefully devise the group local push algorithm for T-absorbed reachability. First, we show that there exists a recursive relationship between $r_T(s,u)$ and $r_T(s,v)$ for $v \in \mathcal{N}_T^-(u)$ where $\mathcal{N}_T^-(u)$ is the set of in-neighbors of u except the node in T. Specifically, for a T-ARW that reaches u at the L-th hop, it must first reach one in-neighbour v of u at the $(L-1)$-th step; otherwise, it cannot reach u at the following step. Therefore, for the L-hop T-absorbed visiting probability $\Pr[s \rightsquigarrow u|L,T]$, we can derive the following equation:

$$\Pr[s \rightsquigarrow u|L,T] = \sum\nolimits_{v \in \mathcal{N}_T^-(u)} \Pr[s \rightsquigarrow v|L-1,T] \cdot \frac{1-\alpha}{d^+(v)},$$

where $d^+(v)$ is the out-degree of node v. To explain, for a random walk that visits v at the $(L-1)$-th hop, with the remaining $1-\alpha$ probability, it randomly jumps to an out-neighbour of node v. Thus, it moves to node u from node v with $\frac{(1-\alpha)}{d^+(v)}$ probability. Adding all such in-neighbours v of node u together, the above equation is obtained. Now, by summing up over all possible hops L where $L \in [1,+\infty)$ since the value of L starts from 1 since this case must visit an in-neighbour of node u, we have that:

$$\sum_{L=1}^{+\infty} \Pr[s \rightsquigarrow u|L,T] = \sum_{v \in \mathcal{N}_T^-(u)} \left(\sum_{L=0}^{+\infty} \Pr[s \rightsquigarrow v|L,T] \right) \cdot \frac{1-\alpha}{d^+(v)},$$

which is equal to $\sum_{v \in \mathcal{N}_T^-(u)} r_T(s,v) \cdot \frac{1-\alpha}{d^+(v)}$ from the definition of $r_T(s,v)$. Since $r_T(s,u) = \sum_{L=0}^{+\infty} \Pr[s \rightsquigarrow u|L,T]$, we can derive that:

$$r_T(s,u) = \Pr[s \rightsquigarrow u|L=0,T] + \sum\nolimits_{v \in \mathcal{N}_T^-(u)} r_T(s,v) \cdot \frac{1-\alpha}{d^+(v)}.$$

Obviously, the value of $\Pr[s \rightsquigarrow u|L=0,T]$ is 1 if $s=u$ since node s always visits itself at the first step, and otherwise is 0. Thus, the final recursive equation is obtained as follows:

$$r_T(s,u) = \mathbb{I}_s(u) + \sum\nolimits_{v \in \mathcal{N}_T^-(u)} r_T(s,v) \cdot \frac{(1-\alpha)}{d^+(v)}, \quad (4)$$

where $\mathbb{I}_s(u)$ is an indicator function that is 1 if $u=s$ and otherwise 0. This recursive definition plays an important role in our local push algorithm. It is because it indicates that the T-absorbed reachability $r_T(s,u)$ could be obtained by summing up the T-absorbed reachability $r_T(s,v)$ of all in-neighbours v of u.

Rationale of Group Local Push: Now, we present the rationale behind our group local push algorithm to compute the T-absorbed reachability $r_T(s,T)$, by exploiting (4). In the R.H.S

Authorized licensed use limited to: Hong Kong University of Science and Technology. Downloaded on July 08,2024 at 11:53:59 UTC from IEEE Xplore. Restrictions apply.

Algorithm 1: GroupLocalPush$(G, s, T, \alpha, R_{max})$.

1 $R(v,T) \leftarrow 0$ for $v \in V \backslash T$;
2 $R(t,T) \leftarrow 1$ for each $t \in T$;
3 $z(s,T) \leftarrow 0$;
4 **for** *each* $t \in T$ **do**
5 **for** *each* v *such that* $v \in \mathcal{N}_T^-(t)$ **do**
6 $R(v,T) \leftarrow R(v,T) + \frac{1-\alpha}{d^+(v)} \cdot R(t,T)$;
7 $R(t,T) = 0$;
8 **while** $\exists v \in V \backslash T$ *such that* $R(v,T) > R_{max}$ **do**
9 **if** v *is* s **then**
10 $z(v,T) \leftarrow z(v,T) + R(v,T)$;
11 **for** *each* u *such that* $u \in \mathcal{N}_T^-(v)$ **do**
12 $R(u,T) \leftarrow R(u,T) + \frac{1-\alpha}{d^+(u)} \cdot R(v,T)$;
13 $R(v,T) = 0$;
14 **return** $z(s,T)$, and $R(v,T)$ of each node $v \in V \backslash T$;

of the above equation, the first term is a constant whose value is 1 or 0 depending on whether $t = s$. We call this constant a *reserve* of node t w.r.t. the source node s, denoted as $z(s,t)$, which is the amount of reachability value $r_T(s,t)$ that has been distributed to node s from t via the incoming edges in the graph. Besides, The second term reveals a probability that a T-ARW reaches node t from an in-neighbour v of node t, namely $\frac{(1-\alpha)}{d^+(v)}$. We call this probability a *residue* of node v w.r.t. node t, denoted as $R(v,t)$, which is waiting for further backward propagation via v's in-neighbours. We thus rewrite the reachability $r_T(s,t)$ as follows:

$$r_T(s,t) = z(s,t) + \sum\nolimits_{v \in \mathcal{N}_T^-(t)} r_T(s,v) \cdot R(v,t).$$

Similarly, for each in-neighbour $v \in \mathcal{N}_T^-(t)$, the reachability $r_T(s,v)$ could be represented as: $r_T(s,v) = z(s,v) + \sum_{w \in \mathcal{N}_T^-(v)} r_T(s,w) \cdot R(w,v)$. As a result, we could find a relationship between the reachability of t and the reachability of t's 2-hop in-neighbours, 3-hop in-neighbors, and so on. Thus, the reachability $r_T(s,t)$ could be represented by the reachability value $r_T(s,v)$ for each node $v \in V$:

$$r_T(s,t) = z(s,t) + \sum\nolimits_{v \in V \backslash T} r_T(s,v) \cdot R(v,t). \quad (5)$$

As analyzed in Section III-A, the reachability value $r_T(s,T)$ is the sum of $r_T(s,t)$ of each node $t \in T$, and we can derive the following invariant:

$$r_T(s,T) = z(s,T) + \sum\nolimits_{v \in V \backslash T} r_T(s,v) \cdot R(v,T). \quad (6)$$

Hence, we could use $z(s,T)$ to approximate $r_T(s,T)$ by performing push operations at each v with non-zero $R(v,T)$. A detailed derivation of (6) is deferred to our technical report [40].

We develop our group local push algorithm, as shown in Algorithm 1, where $R_{\max}$ is the user-specific threshold. We defer the illustration of Algorithm 1 to our technical report for the interest of space. To better understand our GroupLocalPush algorithm, we give a simple example where the target set T has only one node t.

Example 3: Consider the example in Fig. 2. Assume $\alpha = 0.2$ and the target set $T = \{t\}$. We set $R_{\max} = 0.4$. Initially, the residue value of the target node t is set to $R(t,T) = 1$, while other nodes have zero residue values, as depicted in Fig. 2(a). Next, a backward push operation is performed at node t as shown in Fig. 2(b). This operation proceeds as follows: (i) for each in-neighbour v of node t, the residue of node v increases by $\frac{1-\alpha}{d^+(v)} \cdot R(t,T)$ (see Algorithm 1 Line 12). Therefore, for the sole in-neighbour v_1 of t, $R(v_1,T)$ becomes $\frac{1-0.2}{1} = 0.8$; and (ii) it sets $R(t,T) = 0$ (see Algorithm 1 Line 13). Subsequently, since $R(v_1,T) = 0.8 > R_{\max}$, a backward push operation is performed at node v_1, as depicted in Fig. 2(c). Here: (i) for the in-neighbour s of v_1, $R(s,T)$ increases by $\frac{1-\alpha}{d^+(s)} \cdot R(v_1,T)$, becoming 0.32; and (ii) it sets $R(v_1,T)$ to 0. Currently, as no node has a residue value exceeding $R_{\max}$, the algorithm terminates. By (6), we have $r_T(s,T) = 0 + R(s,T) \times r_T(s,s) = 0.32 \times r_T(s,s)$, given that $z(s,T) = 0$.

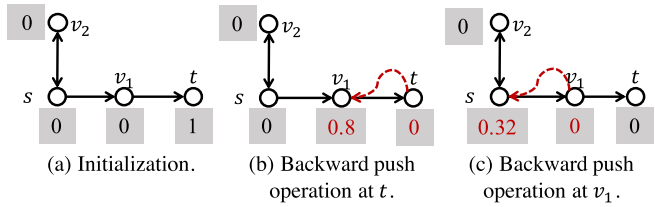

(a) Initialization. (b) Backward push operation at t. (c) Backward push operation at v_1.

Fig. 2. Running example of the GroupLocalPush algorithm with target set $T = \{t\}$, where the residue value $R(v,T)$ for each v is shown inside grey box and the newly-updated ones are highlighted in red.

Theorem 2: Given a target set T, the amortized cost for Algorithm 1 is bounded by $O\left(\frac{|T| \cdot m}{n \cdot \alpha \cdot R_{\max}}\right)$.

All the omitted proofs are deferred to our technical report [40].

C. Combining Group Local Push With MC

The main idea of our SAMBA can be understood as a combination of the MC method and group local push: it first performs group local push with an elaborately designed threshold $R_{\max}$, and then simulates a collection of T-ARWs. Since SAMBA exploits information obtained from the group local push, it can significantly cut down the number of required T-ARWs, saving computational overhead while still satisfying the desired quality guarantee.

Challenge: According to Section III-B, after performing the phase of group local push, the invariant (i.e., (6)) still holds. Thus, in order to compute $r_T(s,T)$, what we need is to estimate the GHP $r_T(s,v)$ for each v. Intuitively, we can sample a sufficiently large number of T-ARWs and get an unbiased estimation $\hat{r}_T(s,T)$, and then apply a concentration bound like Chernoff bound to bound the error of $\hat{r}_T(s,T)$. Unfortunately, all existing concentration bounds require that the range of the random variable (or its variance) is bounded. In our case, both the range and the variance of $r_T(s,v)$ are unbounded since the length of a T-ARW from s might be infinity. This makes it infeasible to guarantee the accuracy of the estimated $\hat{r}_T(s,T)$.

Hop-Wise Method: Recall that $r_T(s,v) = \sum_{L=0}^{+\infty} \Pr[s \rightsquigarrow v|L,T]$, where $\Pr[s \rightsquigarrow v|L,T]$ is bounded for each L. Thus we may approximate an $r_T(s,v)$ by separately estimating $\Pr[s \rightsquigarrow v|L,T]$ $(L = 1, 2, \ldots)$, and then combining them together. From

Authorized licensed use limited to: Hong Kong University of Science and Technology. Downloaded on July 08,2024 at 11:53:59 UTC from IEEE Xplore. Restrictions apply.

Algorithm 2: SAMBA$(G, s, T, \alpha, \epsilon, p_f, \delta, R_{\max})$.

1 **if** $s \in T$ **then**
2 $\quad \hat{f}(s,T) = 1$;
3 **else**
4 $\quad [z(s,T), \mathbf{R}(T)] \leftarrow$ GroupLocalPush(s, T, α, R_{max});
5 $\quad Y \leftarrow R(s,T)$, $L_{max} \leftarrow \log_{1-\alpha}(\frac{\alpha}{1-\alpha} \cdot \frac{\epsilon}{2} \cdot \delta)$;
6 $\quad \omega \leftarrow \frac{3 \cdot R_{max} \cdot \log(2/p_f)}{(1 - R_{max}/2) \cdot \delta \cdot \epsilon^2}$;
7 $\quad$ **for** $L = 1$ *to* L_{max} **do**
8 $\quad\quad Y_L \leftarrow 0$;
9 $\quad\quad \omega_L \leftarrow \lceil \omega \cdot (1-\alpha)^L \rceil$;
10 $\quad\quad a_L = \frac{(1-\alpha)^L \cdot \omega}{\omega_L}$;
11 $\quad\quad$ **for** $i = 1$ *to* ω_L **do**
12 $\quad\quad\quad W_i \leftarrow$ RandomWalk(s, L);
13 $\quad\quad\quad$ **if** $W_i \cap T = \emptyset$ **then**
14 $\quad\quad\quad\quad v \leftarrow$ the visited node at the L-th hop of W_i;
15 $\quad\quad\quad\quad Y_L \leftarrow Y_L + R(v,T) \cdot a_L$;
16 $\quad\quad Y \leftarrow Y + Y_L$;
17 $\quad \hat{f}(s,T) \leftarrow z(s,T) + R(s,T) + Y/\omega$;
18 **return** $\hat{f}(s,T)$;

this idea, we develop our hop-wise method with approximation guarantee for $\sum_{v \in V \setminus T} r_T(s,v) \cdot R(v,T)$.

For ease of clarification, we use $H_L(s,v)$ to denote the L hop T-absorbed visiting probability $\Pr[s \rightsquigarrow v | L, T]$. In order to estimate $\sum_{v \in V \setminus T} r_T(s,v) \cdot R(v,T)$ by $H_L(s,v)$, we need to answer these two questions: (i) how to handle the infinite series of $H_L(s,v)$ $(L = 0, 1, \ldots, \infty)$; (ii) how many L-hop random walks are required to estimate $H_L(s,v)$ respectively such that the subsequent aggregation of all $H_L(s,v)$ $(v \in V \setminus T)$ has accuracy guarantee to approximate $\sum_{v \in V \setminus T} r_T(s,v) \cdot R(v,T)$. We answer the first question with the truncation technique. Recall that an α-terminating random walk stops at each step with α probability. It can be directly derived that the probability that a random walk has L hops (or more) is $(1-\alpha)^L$. Therefore, for any $H_L(s,v)$, it is bounded by $(1-\alpha)^L$. That is, when L is sufficiently large, $H_L(s,v)$ is insignificant, and we can omit it. We give the following lemma about the truncation.

Lemma 2: Let $L_{\max} = \log_{1-\alpha}\left(\frac{\alpha}{1-\alpha} \cdot \frac{\epsilon\delta}{2}\right)$. If $r_T(s,v) > \delta$, then the relative error of $\sum_{L=0}^{L_{\max}} H_L(s,v)$ and $r_T(s,v)$ is at most $\frac{\epsilon}{2}$.

Before answering the second question about the random walk number issues, we first present the pseudo-code of SAMBA, as shown in Algorithm 2. Specifically, for the first case (where $s \in T$), SAMBA sets $\hat{f}(s,T)$ as 1 (Lines 1–2). Otherwise, it starts the estimation by integrating the group local push with our random walk sampling (Lines 3–14): it first invokes our group local push algorithm with $R_{\max}$ which returns the reserve $z(s,T)$ and residue $R(v,T)$ for each node $v \in V \setminus T$ (Line 4). Next, it generates sufficient random walks (Lines 7–16), where the procedure **RandomWalk**(s, L) returns a random walk with L hops starting from s. Given a specific L, the number of L-hop random walks is $\omega_L = \lceil \omega \cdot (1-\alpha)^L \rceil$. Let $W_i[L]$ denote the node visited by the i-th random walk W_i at its L-th hop (Line 12). Let X_i be the random variable that takes value $R(v,T) \cdot a_L$ if $W_i[L] = v$ and $W_i \cap T = \emptyset$ (i.e., W_i does not visit any node in T), otherwise take value 0. Then the expectation of X_i is:

$$E[X_i] = \sum_{v \in V \setminus T} \frac{H_L(s,v)}{(1-\alpha)^L} \cdot R(v,T) \cdot a_L.$$

Then the expectation of Y_L is as follows:

$$E[Y_L] = \omega_L \cdot E[X_i] = \omega \cdot \sum_{v \in V \setminus T} H_L(s,v) \cdot R(v,T).$$

Therefore, we have:

$$E[Y/\omega] = E\left[\frac{1}{\omega} \cdot \sum_{L=1}^{L_{\max}} Y_L\right] = \frac{1}{\omega} \cdot \sum_{L=1}^{L_{\max}} E[Y_L]$$

$$= \sum_{v \in V \setminus T} \sum_{L=1}^{L_{\max}} H_L(s,v) \cdot R(v,T). \quad (7)$$

As $H_0(s,s) = 1$ and $H_0(s,v) = 0$ for $v \neq s$, we have:

$$R(s,T) + E[Y/\omega] = \sum_{v \in V \setminus T} \sum_{L=0}^{L_{\max}} H_L(s,v) \cdot R(v,T).$$

We claim that by setting $\omega = \frac{3 \cdot R_{\max} \cdot \log(2/p_f)}{(1 - R_{\max}/2) \cdot \delta \cdot \epsilon^2}$, SAMBA returns a estimated GHP $\hat{f}(s,T)$ which is close to $f(s,T)$. Besides, to minimize the total time complexity, we can choose an $R_{\max}$ such that the cost of group local push and the MC are balanced.

Theorem 3: SAMBA (Algorithm 2) guarantees that (1) holds with at least $1 - p_f$ probability. By setting $R_{\max} = \epsilon \cdot \sqrt{\frac{\alpha \cdot |T| \cdot m \cdot \delta}{3n \cdot \log(2/p_f)}}$, the time complexity of SAMBA is $O\left(\frac{1}{\epsilon \cdot \alpha}\sqrt{\frac{|T| \cdot m \log(1/p_f)}{n \cdot \delta \cdot \alpha}}\right)$.

IV. Top-K GHP Query

In this section, we study the approximate top-k GHP problem (see Definition 4) and present an efficient algorithm PING[2] to answer the top-k queries with desirable approximation guarantee. Before that, we first derive upper and lower bounds, namely $f^u(s,T)$ and $f^l(s,T)$ for the GHP $f(s,T)$ if we invoke SAMBA. Note that these two bounds hold for any fixed δ and $R_{\max}$. The derivation is based on Chernoff bounds.

Lemma 3: Given a source node s, a target set T and a failure probability p_f, by running SAMBA, we have that with $(1 - p_f)$ probability, $f(s,T)$ is bounded:

$$f(s,T) \leq f^u(s,T) \triangleq z(s,T) + R(s,T) + \Delta^u,$$

$$f(s,T) \geq f^l(s,T) \triangleq z(s,T) + R(s,T) + \Delta^l.$$

Let $\beta = \log(2/p_f)$. The notations Δ^u and Δ^l are defined as:

$$\Delta^u = \left(\sqrt{\frac{Y}{R_{\max}} + \frac{\beta}{2}} + \sqrt{\frac{\beta}{2}}\right)^2 \cdot \frac{R_{\max}}{\omega} + \frac{\epsilon}{n}$$

[2]Top-k GHP with Refining Bounds

Authorized licensed use limited to: Hong Kong University of Science and Technology. Downloaded on July 08,2024 at 11:53:59 UTC from IEEE Xplore. Restrictions apply.

Algorithm 3: PING$(G, \mathcal{T}, s, k, \epsilon)$.

1 $\hat{f}(s,T) \leftarrow 0$ for each $T \in \mathcal{T}$;
2 $\mathcal{C} \leftarrow \mathcal{T}$, $p_f \leftarrow \frac{1}{n}$;
3 $p'_f \leftarrow \frac{p_f}{|\mathcal{T}|\cdot\log_2(n/k)}$;
4 **for** $\delta = \frac{1}{k}, \frac{1}{2k}, \frac{1}{4k}, \ldots, \frac{1}{n}$ **do**
5 $\quad R_{max} \leftarrow \frac{\epsilon\cdot\sqrt{|\mathcal{C}|\cdot|\mathcal{T}|\cdot\alpha\cdot m\cdot\delta}}{\sqrt{3n\cdot\log(2/p'_f)}}$;
6 $\quad \omega \leftarrow \frac{3\cdot R_{max}\cdot\log(2/p'_f)}{(1-R_{max}/2)\cdot\delta\cdot\epsilon^2}$;
7 $\quad$ Invoke SAMBA, and compute $f^l(s,T)$ and $f^u(s,T)$ for each $T \in \mathcal{C}$ according to Lemma 3;
8 $\quad \mathcal{R} \leftarrow$ top-k target sets from C according to $f^u(s,T)$;
9 $\quad$ **if** $\frac{f^l(s,T)}{f^u(s,T)} \geq (1-\epsilon)$ *for each* $T \in \mathcal{R}$ **then**
10 $\quad\quad$ **return** $\mathcal{R}$;
11 $\quad f_k^l \leftarrow$ the k-largest $f^l(s,T)$ for $T \in \mathcal{C}$;
12 $\quad$ **for** *each* $T \in \mathcal{C}$ **do**
13 $\quad\quad$ **if** $f^u(s,T) < f_k^l$ **then**
14 $\quad\quad\quad$ Prune T from the candidate set $\mathcal{C}$;
15 **Return** $\mathcal{R}$;

$$\Delta^l = \left(\left(\sqrt{\frac{Y}{R_{\max}} + \frac{2}{9}\beta} - \sqrt{\frac{\beta}{2}}\right)^2 - \frac{\beta}{18}\right)\cdot\frac{R_{\max}}{\omega} - \frac{\epsilon}{n}.$$

A. PING Algorithm

In order to answer a top-k GHP query, one may consider invoking SAMBA for each target set $T \in \mathcal{T}$, and return the k target sets with top-k largest estimated GHPs. However, if we want to satisfy the accuracy requirement described in Definition–4, we have to set up the parameters according to the exact k-th largest GHP $f(s,T_k^*)$, which is unknown in advance. To tackle this problem, a naive solution is to conservatively set $\delta = \frac{1}{n}$. It yields huge unnecessary computational costs, as the value of $f(s,T_k^*)$ is usually much larger than $\frac{1}{n}$.

Motivated by the aforementioned deficiency, we propose an iterative algorithm, whose main idea is that we adaptively diminish the value of δ, obtain tighter and tighter bounds for each GHP in the candidate set, and then examine whether the terminating condition is met. If it is met, it indicates that the approximate top-k GHPs satisfying Definition 4 have been found, and we stop the algorithm immediately without wasting computational resources. If not, we halve δ to refine the lower and upper bounds and repeat the process until the approximate top-k results are found. Furthermore, in each iteration, we prune those target sets that are impossible to be the top-k results, such that in subsequent iterations we do not need to compute their GHP, further saving computational overhead.

Algorithm 3 shows the pseudo-code of PING. Initially, we initialize the candidate set $\mathcal{C}$ to include all the target sets. Then the algorithm goes into the iteration loops. In each iteration, it calculates the parameters $R_{\max}$ and ω according to the current value of δ. Then it invokes SAMBA to estimate the GHP for each $T \in \mathcal{C}$, and meanwhile derive its bounds $f^l(s,T)$ and $f^u(s,T)$ according to Lemma 3. Next, it constructs a set $\mathcal{R}$ of target sets whose upper bound $f^u(s,T)$ are top k largest among $\mathcal{C}$ (Line 8). If each target set T in $\mathcal{R}$ satisfies that the ratio $\frac{f^l(s,T)}{f^u(s,T)}$ is greater than $1-\epsilon$, the algorithm terminates. If this stopping condition is not met, we have to halve the value of δ such that in the next iteration it can provide tighter bounds for each target set in $\mathcal{C}$. Before beginning with the next iteration, we further prune those target sets with upper bound $f^u(s,T)$ being smaller than the line f_k^l, which must not be the top-k answers, as there exist at least k target sets with larger GHP.

Next, we analyze the correctness of Algorithm 3, that is, it can provide approximate top-k results satisfying Definition 4.

Theorem 4: Algorithm 3 returns the top-k result $\mathcal{R}$ which satisfies Definition 4 with at least $1 - 1/n$ probability.

Theorem 5: The expected time complexity of PING (Algorithm 3) is $O\left(\frac{1}{\epsilon\cdot\alpha}\sqrt{\frac{m\cdot\log(1/p_f)}{\alpha\cdot\delta^*}}\right)$, where δ^* is the value of δ when the algorithm terminates.

Lemma 4: When $\delta = O(f(s,T_k^*))$, Algorithm 4 meets the stopping condition with high probability, where $f(s,T_k^*)$ is the exact k-th largest GHP.

B. Optimization

Recall that in the algorithm PING, it has to compute $\hat{f}(s,T)$ for each candidate set $T \in \mathcal{C}$ by invoking SAMBA, where $|\mathcal{C}|$ is the size of the candidate sets. For the sake of saving computational cost, we can generate only $\sum_{L=1}^{L_{\max}}\omega_L$ random walks (Line 12), and reuse them to estimate the GHP values for each $T \in \mathcal{C}$. However, due to the target-orient property, we have to estimate the term $\sum_{v\in V\setminus T} r_T(s,v)R(v,T)$ (Lines 13-15) for each $T \in \mathcal{C}$ one by one. As a result, we need to scan these random walks $|\mathcal{C}|$ times in total. Such a solution is expensive and significantly hampers its practical performance.

Optimization: Our optimization is based on an observation that most candidate target sets are not frequently hit by random walks. Thus we can first assume that all random walks do not visit any node in any target set $T \in \mathcal{C}$. Then, we can scan random walks only once and estimate $r_T(s,v)$ for all $v \in V\setminus T$. Surely, the estimated $\hat{f}(s,T)$ under such an assumption is incorrect. Subsequently, a rectifying step is needed to correct the outcome.

Motivated by this idea, we propose our optimization technique whose pseudocode is given in Algorithm 4. Algorithm 4 consists of two phases: the *sampling* phase (Lines 6–16) and the *correction* phase (Lines 17–23). In particular, in the sampling phase, we sample $\sum_{L=1}^{L_{\max}}\omega_L$ random walks. For each random walk W_i whose last node is v (Line 11), we directly increase the counter $c(v)$ by a_L since we have assumed W_i does not visit any target set. Then if any node of a target set T has been visited by W_i, we record the walk id in $\mathcal{H}(T)$ (Lines 14–16). Finding out all the target sets visited by W_i can be solved efficiently, if we maintain a data structure for all v indicating which target sets include v. In the second phase, for each $T \in \mathcal{C}$ it first estimates $\hat{f}(s,T)$ (Lines 18–20) by summing up $c(v)\cdot R(v,T)$ for each node v whose $R(v,T) \neq 0$. However, as the counter $c(v)$ has counted in those random walks that visit certain nodes in T, $\hat{f}(s,T)$ is over-estimated. Thus we go through the records in $\mathcal{H}(T)$, and remove the over-estimating amount from $\hat{f}(s,T)$ (Lines 21–23).

Example 4: Consider the graph presented in Fig. 1. We set $L_{\max} = 3$ and the candidate set $\mathcal{C}$ to consist solely of T_2. For simplicity, we omit the random walks of lengths

Authorized licensed use limited to: Hong Kong University of Science and Technology. Downloaded on July 08,2024 at 11:53:59 UTC from IEEE Xplore. Restrictions apply.

Algorithm 4: Optimization $(G, s, \alpha, \omega, R_{\max}, \mathcal{C})$.

1 **for** *each node* $T \in \mathcal{C}$ **do**
2 $[z(s,T), \mathbf{R}(T)] \leftarrow$ Invoke Algorithm 1 with G, s, T, α, R_{max};
3 $\mathcal{H}(T) \leftarrow \emptyset$ for each $T \in \mathcal{C}$;
4 $i \leftarrow 0$;
5 $c(v) \leftarrow 0$ for each $v \in V$;
6 **for** $L = 1$ *to* L_{max} **do**
7 $i \leftarrow i + 1$;
8 $\omega_L \leftarrow \lceil (1-\alpha)^L \omega \rceil$;
9 $a_L = (1-\alpha)^L \omega / \omega_L$;
10 **for** $j = 1$ *to* ω_L **do**
11 $W_i \leftarrow$ RandomWalk(s, L);
12 $v \leftarrow$ the last node of W_i ;
13 $c(v) \leftarrow c(v) + a_L$;
14 **for** *each node* v *in* W_i **do**
15 **if** v *is a member node of any candidate set* T **then**
16 Append walk id i to $\mathcal{H}(T)$;
17 **for** *each node* $T \in \mathcal{C}$ **do**
18 $\hat{f}(s,T) \leftarrow z(s,T) + R(s,T)$;
19 **for** *each node* v *whose residue* $R(v,T) \neq 0$ **do**
20 $\hat{f}(s,T) \leftarrow \hat{f}(s,T) + c(v) \cdot R(v,T)/\omega$;
21 **for** *each walk* i *in* $\mathcal{H}(T)$ **do**
22 $v \leftarrow$ the last node of W_i;
23 $f(s,T) \leftarrow \hat{f}(s,T) - a_L \cdot R(v,T)/\omega$;
24 **Return** $\hat{f}(s,T)$ for each node $T \in \mathcal{C}$;

$L = 1$ and $L = 2$. Suppose we sample $\omega_3 = 4$ random walks starting from the source node s. These walks are as follows: $\{W_1$: $s \to v_1 \to v_8 \to v_6$, W_2: $s \to v_4 \to v_2 \to v_3$, W_3: $s \to v_2 \to v_3 \to v_5$, W_4: $s \to v_4 \to v_2 \to v_3\}$. The terminal nodes of W_1, W_2, W_3, and W_4 are v_6, v_3, v_5, and v_3, respectively. This results in the counts $c(v_6) = a_3$, $c(v_3) = 2 \cdot a_3$, and $c(v_5) = a_3$ (as described in Algorithm 4, Line 9 and Line 13). As random walk W_1 visits node v_8 from T_2, we deduce $\mathcal{H}(T_2) = 1$, where the value 1 represents W_1's id. Next, we describe the calculation of $\hat{f}(s, T_2)$ using the sampled random walks, W_1 through W_4. Let $z(s, T_2) + R(s, T_2) = z_0$. We initially estimate $f(s, T_2)$ as $\hat{f}(s, T_2) = z_0 + (a_3 \cdot R(v_6, T_2) + 2 \cdot a_3 \cdot R(v_3, T_2) + a_3 \cdot R(v_5, T_2))/\omega$ (Line 20). However, this estimate is flawed as it incorrectly takes into account W_1, which has traversed the target set T_2 before terminating. To correct this, we adjust for the impact of W_1 (Line 23), resulting in $\hat{f}(s, T_2) = z_0 + (2 \cdot a_3 \cdot R(v_3, T_2) + a_3 \cdot R(v_5, T_2))/\omega$.

V. Experiment

In this section, we experimentally evaluate our solutions SAMBA and PING against alternatives on real datasets. All experiments are conducted on a Linux server clocked at 2.3 GHz, and all codes are implemented in C++ and compiled with -O3.

Dataset: We use eight real datasets that can be obtained from public sources SNAP [41] and Konect [42]. These datasets are frequently used in previous work of graph-based algorithms [6], [24]. Their statistics are reported in Table I. Among them, the largest dataset Friendster has 1.8 billion edges and is used to examine the scalability of our algorithms.

TABLE I
Datasets ($K = 10^3, M = 10^6, B = 10^9$)

Name	Abbr.	n	m	Type
DBLP	*DB*	317K	1.0M	undirected
YouTube	*YT*	1.1M	3.0M	undirected
Pokec	*PK*	1.6M	30.6M	directed
Wiki	*WK*	1.8M	28.5M	directed
Flickr	*FL*	2.3M	33.1M	directed
LiveJournal	*LJ*	4.0M	34.7M	undirected
Orkut	*OK*	3.1M	117.2M	undirected
Friendster	*FS*	65.6M	1.8B	undirected

Algorithms: For the pairwise GHP query, we compare our solution SAMBA with the MC method and Ad-FORA [6]. For the MC method, we simulate $O(\frac{\log n}{\delta \cdot \epsilon^2})$ random walks such that it could provide ϵ-approximation guarantee. Ad-FORA is an adapted version of the state-of-the-art PPR method FORA with Forward Push. It is also designed to provide ϵ-approximation guarantee.

For the top-k GHP query, we evaluate our solution PING against the MC method, which generates sufficient random walks to provide an approximation guarantee presented in Definition 4. We also conduct an adapted version of PING which runs MC only without group local push (namely, $R_{\max}$ is always set as 1), denoted as *MC-PING*. In this way, we examine the effectiveness of the termination condition. We omit Ad-FORA in the top-k query as from the experimental results of pairwise queries, it is significantly worse than SAMBA.

We set the relative error parameter ϵ as 0.1 for both the pairwise query and for the top-k query. The values of both δ and p_f are set as $\frac{1}{n}$, following existing work [6], [7], [24]. The value of k ranges from 20 to 100.

Target sets: For the pairwise query, we vary the size of a target set T with $\{1, 10, 20, 50, 100\}$. We regard a target set as a certain community where nodes are usually intimate with each other. Under this assumption, we use a neighbour-based method to sample a target set. Specifically, we randomly sample a node as the seed of a target set and then add its neighbours to the target set. For the top-k GHP query, we set the size of a target set T as $|T| = 20$. The total number of all target sets is $|\mathcal{T}| = 10000$.

Accuracy Metrics: For the pairwise GHP query, obtaining the exact GHP value $f(s, T)$ requires prohibitive computational cost. By setting $\epsilon = 0.005$, we employ the SAMBA algorithm to obtain an estimated result denoted as $\tilde{f}(s, T)$. Our SAMBA guarantees that if $f(s, T) \geq \delta$, then

$$(1 - 0.005) \cdot f(s,T) \leq \tilde{f}(s,T) \leq (1 + 0.005) \cdot f(s,T).$$

Given this guarantee, we treat $\tilde{f}(s, T)$ as the exact GHP value of node s with respect to target set T. Suppose there's another estimated result $\hat{f}(s, T)$ for the GHP $f(s, T)$. We define the relative error of $\hat{f}(s, T)$ as

$$\text{Relative error} = \frac{|\hat{f}(s,T) - \tilde{f}(s,T)|}{\max\left(\tilde{f}(s,T), \delta\right)}.$$

Authorized licensed use limited to: Hong Kong University of Science and Technology. Downloaded on July 08,2024 at 11:53:59 UTC from IEEE Xplore. Restrictions apply.

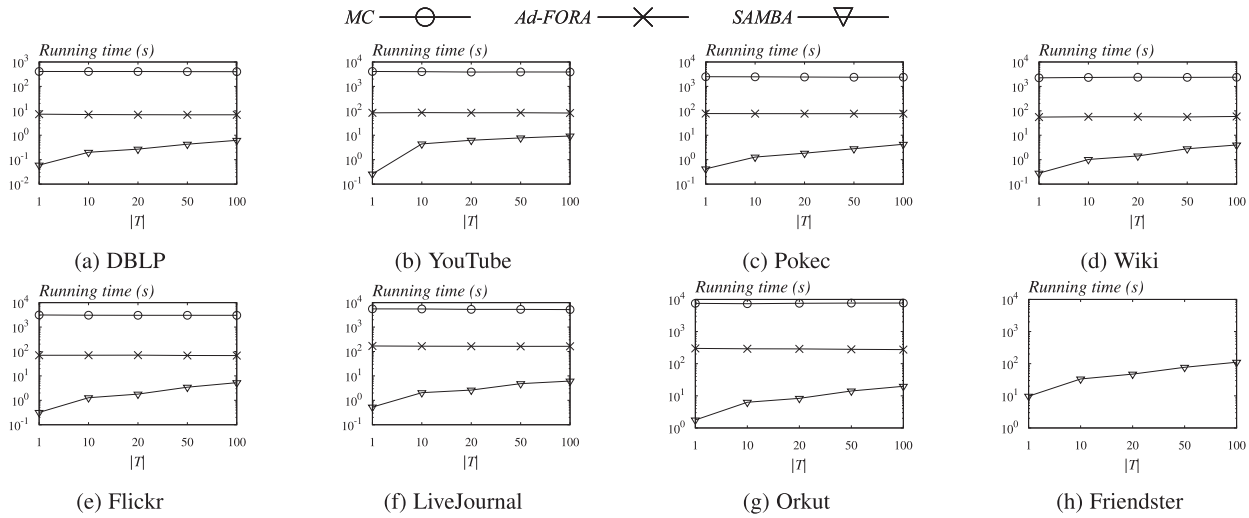

Fig. 3. Varying $|T|$: Running time of each method for the pairwise GHP query.

For the top-k GHP query, we consider the results obtained by MC as the exact top-k results. To gauge the accuracy performance of each method, we use the precision metrics, defined as

$$\text{Precision} = \frac{|\mathcal{R} \cap \mathcal{R}'|}{k},$$

where $\mathcal{R}$ and $\mathcal{R}'$ are the exact top-k result and the approximate top-k result returned by other algorithms, respectively.

A. Pairwise GHP Query

In the first set of experiments, we evaluate the performance of our SAMBA as the target set size $|T|$ varies. As the target set size $|T|$ ranges from 1 to 100, we generate 50 random pairwise queries. We then employ SAMBA, along with baseline methods Ad-FORA and MC, to answer these queries. Subsequently, we report the average running time and the practical maximum relative error for each algorithm, to measure efficiency and accuracy performance.

Fig. 3 reports the average query time of each method for the pairwise query with varying target set size $|T|$. The results demonstrate that our algorithm SAMBA always runs the fastest, followed by Ad-FORA, with the MC method being the slowest. We omit the curves of Ad-FORA and MC on the largest dataset Friendster, since they take more than 4 hours and 72 hours, respectively, to complete a single pairwise query. Specifically, compared to the second fastest method Ad-FORA, our SAMBA is two orders of magnitude faster when $|T| = 1$ and one order of magnitude faster when $|T| = 100$, on all eight datasets. For example, on dataset Orkut, SAMBA is 170x (resp. 14x) faster than Ad-FORA when $T = 1$ (resp. $T = 100$). Furthermore, in comparison to MC, our SAMBA consistently takes advantage by at least two orders of magnitude across all eight datasets, even when $T = 100$. To illustrate, SAMBA is 400x (resp. 560x) faster than MC on dataset Orkut (resp. Pokec) with $T = 100$. The high efficiency of our SAMBA can be explained in that it integrates group local push algorithm to significantly reduce the number of random walks required to guarantee accuracy. We also observe from Fig. 3 that the relationship between the running time of SAMBA (with fixed parameters δ, ϵ, and α) and $\sqrt{|T|}$ is roughly linear. That is, for all eight datasets, SAMBA takes a tenfold increase in running time when $|T|$ varies from 1 to 100. These experimental results conform with our analysis about the time complexity of SAMBA $O\left(\frac{\sqrt{|T|}}{\epsilon \cdot \alpha} \cdot \sqrt{\frac{m \log(1/p_f)}{n \cdot \delta \cdot \alpha}}\right)$, which is presented in Theorem 3.

Then, as shown in Fig. 4, we report the practical maximum relative errors of SAMBA, Ad-FORA, and MC to present their accuracy performance. From these figures, it's observed that the maximum relative errors of all algorithms across all datasets are less than $\epsilon = 0.1$. This confirms the algorithms' correctness, given their design to offer ϵ-approximation guarantees. Notably, our SAMBA algorithm exhibits smaller maximum relative errors than both Ad-FORA and MC on every dataset tested. This indicates that SAMBA outperforms Ad-FORA and MC in accuracy. This improved performance can be attributed to our utilization of the group local push technique, which significantly cuts down the number of required random walks. This reduction also aids in minimizing randomness during the estimation of the GHP value. We also note that in most cases MC suffers from higher maximum relative errors than Ad-FORA.

In the second set of experiments, we evaluate the impact of the relative error parameter ϵ on SAMBA's efficiency and accuracy, setting $|T| = 20$. We execute SAMBA to answer 50 random pairwise queries and subsequently report the average running time and the maximum relative error. The value of ϵ is set to $\{0.5, 0.1, 0.05, 0.01\}$. As depicted in Fig. 5, the results indicate that as ϵ diminishes, the maximum relative error correspondingly decreases, but this improvement comes with a longer running time. A smaller maximum relative error signifies better accuracy performance. Notably, the maximum relative error consistently remains below the specified ϵ value across all eight datasets. For instance, in the LiveJournal dataset with ϵ values of 0.5, 0.1, 0.05, and 0.01, the maximum relative errors are 6.33×10^{-2},

Authorized licensed use limited to: Hong Kong University of Science and Technology. Downloaded on July 08,2024 at 11:53:59 UTC from IEEE Xplore. Restrictions apply.

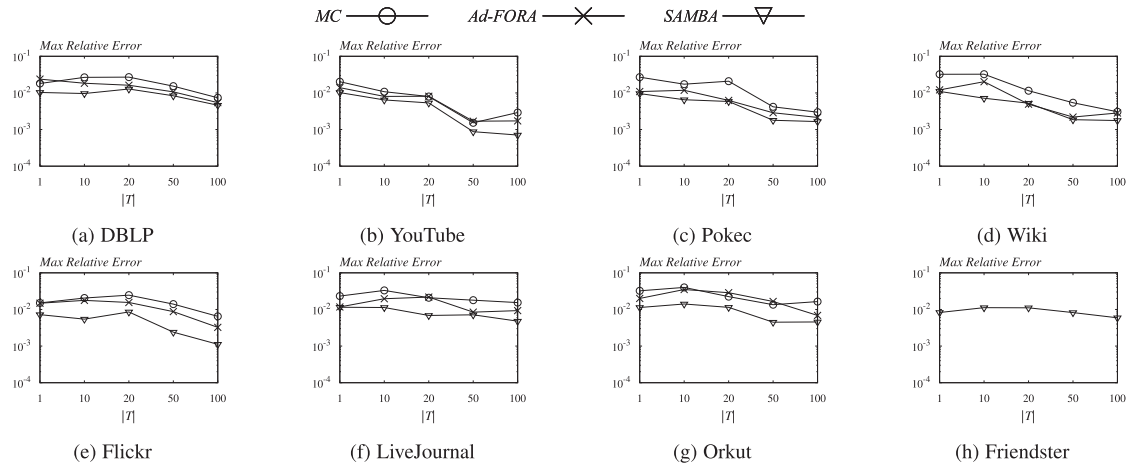

Fig. 4. Varying $|T|$: Maximum relative error of each method for the pairwise GHP query.

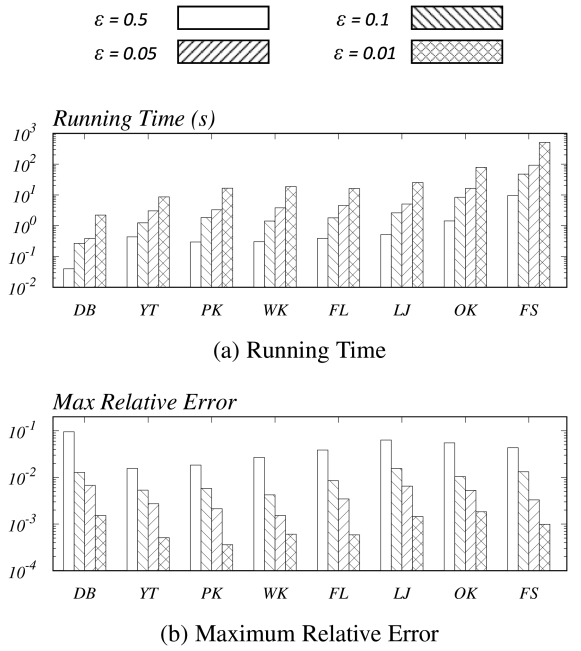

Fig. 5. Performance of SAMBA with varying parameter ϵ.

1.55×10^{-2}, 6.51×10^{-3}, and 1.45×10^{-3}, respectively. This observation aligns with SAMBA's design, which guarantees approximation performance. Moreover, the average running time appears to increase roughly in proportion to $\frac{1}{\epsilon}$. For example, SAMBA's running time on the LiveJournal dataset with ϵ values of 0.5, 0.1, 0.05, and 0.01 is 0.52, 2.63, 5.08, and 25.34 seconds, respectively. This is consistent with the derived time complexity of SAMBA, given as $O\left(\frac{1}{\epsilon \cdot \alpha}\sqrt{\frac{|T| \cdot m \log(1/p_f)}{n \cdot \delta \cdot \alpha}}\right)$.

B. Top-K GHP Query

In the third set of experiments, we evaluate the effectiveness and accuracy of our PING/MC-PING in comparison to MC for answering top-k GHP queries. We set $|T| = 20$. By sampling 20 random sources, we execute PING, MC-PING, and MC to identify the target sets with the largest k GHP scores from the entire target set collection $\mathcal{T}$. Each algorithm answers each query five times. We then report the average running time and P@k value, assessing both efficiency and accuracy. We opt not to incorporate the terminating condition with the algorithm Ad-FORA, as the experimental results in Fig. 3 already demonstrate its inferior performance compared to our SAMBA.

Fig. 6 presents the average query time of MC, MC-PING, and PING for each dataset. The figures show that our PING and MC-PING algorithms consistently outpace the MC method significantly. Without prior knowledge of the exact top-k GHP scores, the MC method must conservatively provide approximation guarantees for all target sets with GHP scores greater than $\frac{1}{n}$. In contrast, our algorithms, PING and MC-PING, operate iteratively with a diminishing threshold δ, allowing them to terminate once they identify the top-k results that meet the user-specified approximation requirement. Consequently, our algorithms demonstrate greater efficiency than the MC method. Additionally, PING consistently outperforms MC-PING regarding query time across all tested datasets. For instance, with $k = 100$, PING processes a top-k query on the Orkut dataset in an average of 30 seconds, while MC-PING requires 250 seconds. This efficiency arises because PING integrates with our SAMBA and employs an optimization technique, enabling efficient estimation of the GHP value for each target set in $\mathcal{C}$. We omit the curve of running time of the MC method on the Friendster dataset, as it incurs prohibitive computational overhead.

Fig. 7 presents the precision results of PING and MC-PING. Notably, MC-PING utilizes the same terminating condition as described in Lines 9-10 of Algorithm 3. The precision scores for the MC method consistently equal 1 since we regard its output as the ground truth. It's noteworthy that the precision scores of PING and MC-PING closely align, and in most instances, they exceed 99%. This underscores the effectiveness of our terminating condition. Additionally, our PING achieves an optimal balance between efficiency and accuracy. In comparison to the MC method, we realize a significantly reduced running time (by up to three orders of magnitude) while maintaining nearly

Authorized licensed use limited to: Hong Kong University of Science and Technology. Downloaded on July 08,2024 at 11:53:59 UTC from IEEE Xplore. Restrictions apply.

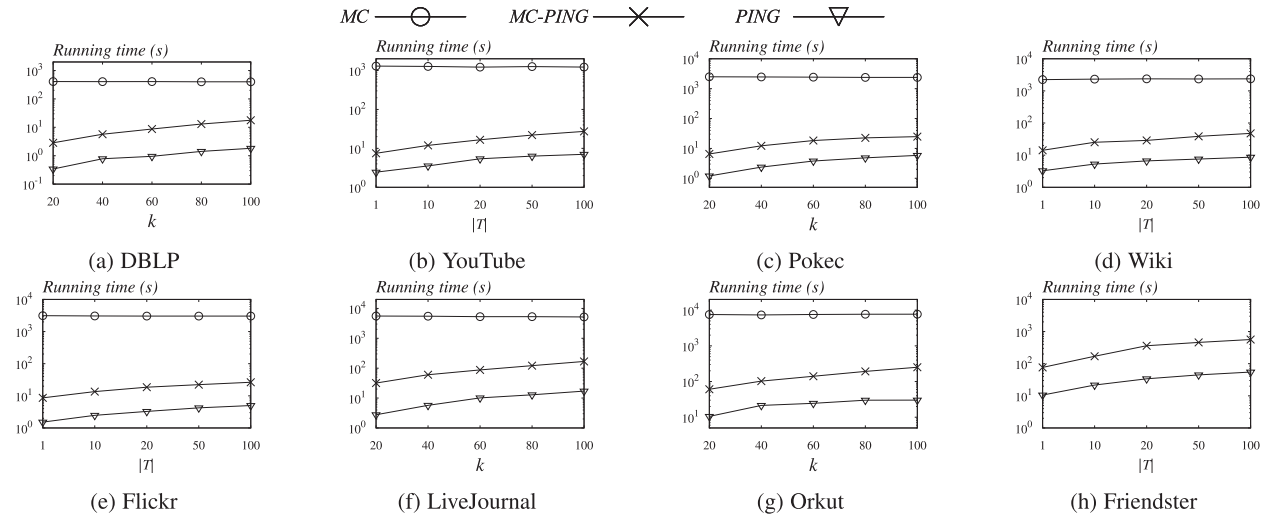

Fig. 6. Varying k: Running time of each method for the top-k GHP query where $|T| = 20$.

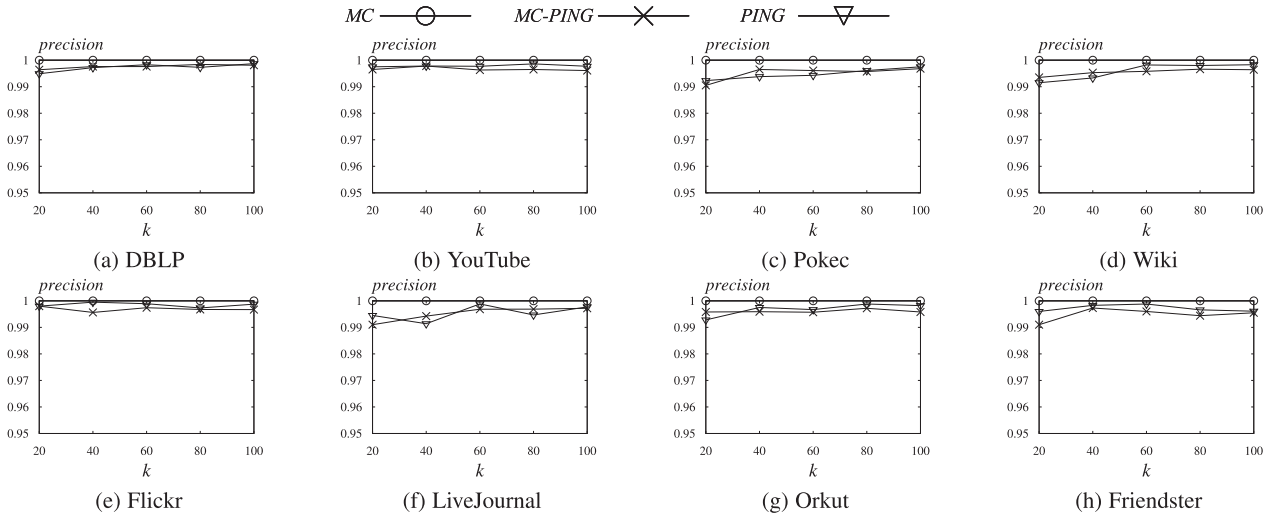

Fig. 7. Varying k: Accuracy performance of each method for the top-k GHP query where $|T| = 20$.

identical query quality, enhancing its appeal for real-world applications. For the Friendster dataset, the single-threaded version of the MC method takes an excessive amount of time to compute the top-k target sets. To address this, we implemented a multi-threaded version of MC. This allowed us to obtain the ground truth, facilitating accuracy evaluation for PING and MC-PING.

We next explore the performance of MC-PING/PING with respect to different values of relative error ϵ, set at $\{0.5, 0.1, 0.05, 0.01\}$. The value of k is set to 100. The experimental results for the DBLP dataset are depicted in Fig. 8(a). They indicate that PING consistently achieves better runtime than MC-PING. As the relative error ϵ decreases, the performance difference between PING and MC-PING grows, with PING being more than 100x faster when $\epsilon = 0.01$. Although similar trends appear in other datasets, we omit these results due to space limitations. To further showcase the effectiveness of our optimization strategy, as detailed in Section IV-B, we present the runtime of PING for top-k queries, where $k = 100$, in Fig. 8(b). This figure highlights that our optimization technique significantly reduces PING's execution time compared to

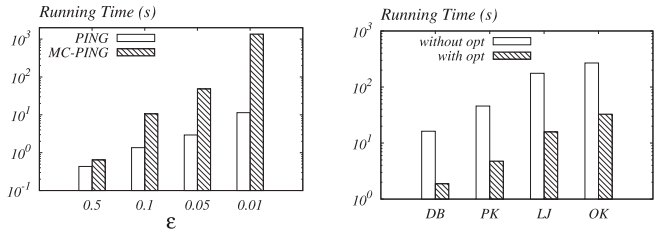

Fig. 8. (a) Time consumption of PING/MC-PING when varying relative error ϵ on dataset DBLP; (b) Time consumption of PING with/without optimization technique.

Authorized licensed use limited to: Hong Kong University of Science and Technology. Downloaded on July 08,2024 at 11:53:59 UTC from IEEE Xplore. Restrictions apply.

TABLE II
COMMUNITY PREDICTION PERFORMANCE (%)

Datasets		GHP	PPR-max	PPR-sum	PPR-avg
DBLP	NDCG@1	**78.1**	<u>71.2</u>	69.9	56.9
	NDCG@5	**86.7**	<u>83.4</u>	82.5	64.7
LJ	NDCG@1	**92.5**	83.4	<u>90.8</u>	81.4
	NDCG@5	**95.2**	90.5	<u>94.2</u>	88.2

The best and second-best performances are highlighted in bold and underline, respectively.

the non-optimized version. Specifically, using our optimization leads to up to an 11x improvement in efficiency. This boost is attributed to our optimization design, which combines a single scan of the random walks with a subsequent refinement step for adjusting the GHP score estimates.

C. Case Study

To examine the performance of our GHP in measuring node-to-group proximity, we conducted a case study on proximity-based group recommendation. We compared our GHP method against three PPR-aggregation-based methods. Given that PPR is a widely-used random-walk-based tool for measuring node-to-node similarity [43], [44], [45], we select it as the baseline. We aggregate the PPR scores of the nodes within a group using max, sum, and average. These are denoted as PPR-max, PPR-sum, and PPR-avg, respectively. The real datasets DBLP and LiveJournal (abbreviated as LJ) available on the SNAP website [41] are associated with ground-truth communities, facilitating the evaluation of prediction performance.

We define $\mathcal{S}$ as the set of nodes associated with a specific community in the dataset. The experimental setup is as follows:

1) The set $\mathcal{S}$ is randomly split into two subsets: the training set $\mathcal{S}_1$ and the test set $\mathcal{S}_2$. The sizes of $\mathcal{S}_1$ and $\mathcal{S}_2$ are $70\% \times |\mathcal{S}|$ and $30\% \times |\mathcal{S}|$, respectively.
2) Nodes in $\mathcal{S}_2$ are then removed from their respective communities, aiming at hiding their community information.
3) From $\mathcal{S}_2$, we sample 1000 nodes. For each sampled node, its GHP scores and PPR aggregated scores with respect to communities are computed based on the modified community structure. The top-k results are then used to predict the community to which the node belongs.

To illustrate, consider a community $T = \{v_1, v_2, v_3\}$ in the dataset. After Step 1, we may have $v_1 \in \mathcal{S}_1, v_2 \in \mathcal{S}_1$, and $v_3 \in \mathcal{S}_2$. In Step 2, we remove node v_3 from T, thus hiding the fact that v_3 is a part of T. Consequently, T only contains the nodes v_1 and v_2. In Step 3, if node v_3 is among the sampled nodes, we compute its GHP and PPR aggregated scores regarding community sets. The top-k results are chosen as the predicted communities, with the ground truth for v_3 being community T.

Prediction performance is evaluated using the Normalized Discounted Cumulative Gain (NDCG), a prevalent metric in the literature for ranking quality [46], [47]. A higher NDCG value indicates better prediction performance. Table II presents the experimental results. The highest score is highlighted in bold, and the second highest is underlined. GHP significantly surpasses the PPR-aggregation-based methods. Notably, the NDCG@1 score of GHP leads by 6.9% (and 1.7%) over the second-best score in DBLP (and LJ). Furthermore, we observe that PPR-max performs best among the PPR-aggregation methods in DBLP, whereas PPR-sum leads in LJ, suggesting the sensitivity of PPR-aggregation methods to the input graph structure and its potential limitations in real-world scenarios.

VI. CONCLUSION

In this paper, we introduce the concept of group hitting probability (GHP) to measure the node-to-group structural proximity on graphs. We investigate two basic types of GHP queries: the pairwise query and the top-k query. To answer the pairwise query, we present the SAMBA algorithm which combines the MC and group local push algorithm to achieve high efficiency while still providing performance guarantee. Further, to answer the top-k GHP query with approximation guarantee, we present the PING algorithm which could stop immediately as soon as the desired top-k results are found. We also propose an optimization technique to improve efficiency. Experimental results demonstrate that our solutions are more efficient than competitors.

REFERENCES

[1] Y. Zhou, H. Cheng, and J. X. Yu, "Graph clustering based on structural/attribute similarities," in *Proc. VLDB Endowment*, vol. 2, pp. 718–729, 2009.
[2] R. Mihalcea and D. Radev, *Graph-Based Natural Language Processing and Information Retrieval*. Cambridge, U.K.: Cambridge Univ. Press, 2011.
[3] D. Liben-Nowell and J. Kleinberg, "The link-prediction problem for social networks," *J. Assoc. Inf. Sci. Technol.*, vol. 58, pp. 1019–1031, 2007.
[4] A. A. Benczúr, K. Csalogány, and T. Sarlós, "Link-based similarity search to fight web spam," in *Proc. Int. Workshop Adversarial Inf. Retrieval Web*, 2006, pp. 9–16.
[5] D. Goldberg, D. Nichols, B. M. Oki, and D. Terry, "Using collaborative filtering to weave an information tapestry," *Commun. ACM*, vol. 35, pp. 61–70, 1992.
[6] S. Wang, R. Yang, X. Xiao, Z. Wei, and Y. Yang, "FORA: Simple and effective approximate single-source personalized PageRank," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2017, pp. 505–514.
[7] Z. Wei, X. He, X. Xiao, S. Wang, S. Shang, and J.-R. Wen, "TopPPR: Top-k personalized PageRank queries with precision guarantees on large graphs," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2018, pp. 441–456.
[8] R. Andersen, F. Chung, and K. Lang, "Local graph partitioning using PageRank vectors," in *Proc. 47th Annu. IEEE Symp. Found. Comput. Sci.*, 2006, pp. 475–486.
[9] B. Bahmani, A. Chowdhury, and A. Goel, "Fast incremental and personalized PageRank," in *Proc. VLDB Endowment*, vol. 4, pp. 173–184, 2010.
[10] P. Berkhin, "Bookmark-coloring algorithm for personalized PageRank computing," *Internet Math.*, vol. 3, pp. 41–62, 2006.
[11] Y. Fujiwara, M. Nakatsuji, T. Yamamuro, H. Shiokawa, and M. Onizuka, "Efficient personalized PageRank with accuracy assurance," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2012, pp. 15–23.
[12] L. Lovász, "Random walks on graphs," *Combinatorics, Paul Erdos Eighty*, vol. 2, no. 1–46, p. 4, 1993.
[13] Z. Guan, J. Wu, Q. Zhang, A. Singh, and X. Yan, "Assessing and ranking structural correlations in graphs," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2011, pp. 937–948.
[14] Q. Mei, D. Zhou, and K. Church, "Query suggestion using hitting time," in *Proc. Conf. Inf. Knowl. Manage.*, 2008, pp. 469–478.
[15] M. Brand, "A random walks perspective on maximizing satisfaction and profit," in *Proc. SIAM Int. Conf. Data Mining*, 2005, pp. 12–19.
[16] B. Liu, "Better than PageRank: Hitting time as a reputation mechanism," Doctoral Dissertation, Univ. Harvard, 2014.
[17] P. Sarkar and A. W. Moore, "Fast nearest-neighbor search in disk-resident graphs," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2010, pp. 513–522.

Authorized licensed use limited to: Hong Kong University of Science and Technology. Downloaded on July 08,2024 at 11:53:59 UTC from IEEE Xplore. Restrictions apply.

[18] L. Grady, "Random walks for image segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 11, pp. 1768–1783, Nov. 2006.
[19] Meta, "Find facebook groups to join," 2023. [Online]. Available: https://www.facebook.com/help/197647530805955
[20] Linkedin, "Find and join a linkedin group," 2021. [Online]. Available: https://www.linkedin.com/help/linkedin/answer/a544795
[21] P. Lofgren, S. Banerjee, A. Goel, and C. Seshadhri, "Fast-PPR: Scaling personalized PageRank estimation for large graphs," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2014, pp. 1436–1445.
[22] P. Lofgren, S. Banerjee, and A. Goel, "Personalized PageRank estimation and search: A bidirectional approach," in *Proc. SIAM Int. Conf. Data Mining*, 2016, pp. 163–172.
[23] F. Chung and L. Lu, "Concentration inequalities and martingale inequalities: A survey," *Internet Math.*, vol. 3, pp. 79–127, 2006.
[24] S. Wang, Y. Tang, X. Xiao, Y. Yang, and Z. Li, "HubPPR: Effective indexing for approximate personalized PageRank," in *Proc. VLDB Endowment*, vol. 10, pp. 205–216, 2016.
[25] R. Andersen, C. Borgs, J. Chayes, J. Hopcraft, V. S. Mirrokni, and S.-H. Teng, "Local computation of PageRank contributions," in *Proc. Int. Workshop Algorithms Models Web-Graph*, 2007, pp. 150–165.
[26] S. Wang et al., "Efficient algorithms for approximate single-source personalized PageRank queries," *ACM Trans. Database Syst.*, vol. 44, pp. 1–37, 2019.
[27] D. Lin, R. C.-W. Wong, M. Xie, and V. J. Wei, "Index-free approach with theoretical guarantee for efficient random walk with restart query," in *Proc. IEEE 36th Int. Conf. Data Eng.*, 2020, pp. 913–924.
[28] G. Jeh and J. Widom, "Scaling personalized web search," in *Proc. Int. Conf. World Wide Web*, 2003, pp. 271–279.
[29] H. Wu, J. Gan, Z. Wei, and R. Zhang, "Unifying the global and local approaches: An efficient power iteration with forward push," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2021, pp. 1996–2008.
[30] G. Jeh and J. Widom, "SimRank: A measure of structural-context similarity," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2002, pp. 538–543.
[31] B. Tian and X. Xiao, "Sling: A near-optimal index structure for SimRank," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2016, pp. 1859–1874.
[32] Z. Wei et al., "PRSim: Sublinear time SimRank computation on large power-law graphs," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2019, pp. 1042–1059.
[33] H. Wang, Z. Wei, Y. Yuan, X. Du, and J. Wen, "Exact single-source SimRank computation on large graphs," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2020, pp. 653–663.
[34] F. Zhang and S. Wang, "Effective indexing for dynamic structural graph clustering," in *Proc. VLDB Endow.*, vol. 15, no. 11, pp. 2908–2920, 2022.
[35] L. Katz, "A new status index derived from sociometric analysis," *Psychometrika*, vol. 18, no. 1, pp. 39–43, 1953.
[36] Y. Koren, S. C. North, and C. Volinsky, "Measuring and extracting proximity in networks," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2006, pp. 245–255.
[37] W. Chen, D. Zhang, and E. Y. Chang, "Combinational collaborative filtering for personalized community recommendation," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2008, pp. 115–123.
[38] A. Sharma and B. Yan, "Pairwise learning in recommendation: Experiments with community recommendation on linkedin," in *Proc. ACM Conf. Recommender Syst.*, 2013, pp. 193–200.
[39] X. Wang, R. Donaldson, C. Nell, P. Gorniak, M. Ester, and J. Bu, "Recommending groups to users using user-group engagement and time-dependent matrix factorization," in *Proc. Conf. Assoc. Advance. Artif. Intell.*, 2016, pp. 1331–1337.
[40] Q. Guo, "Technical report," 2023. [Online]. Available: https://github.com/qtguo/group-hitting-probability.git
[41] SNAP Datasets, 2014. [Online]. Available: http://snap.stanford.edu/data
[42] KONECT, 2013. [Online]. Available: http://konect.cc/networks/
[43] X. Zhang, K. Xie, S. Wang, and Z. Huang, "Learning based proximity matrix factorization for node embedding," in *Proc. Int. Conf. Knowl. Discov. Data Mining*, 2021, pp. 2243–2253.
[44] R. Yang, J. Shi, X. Xiao, Y. Yang, and S. S. Bhowmick, "Homogeneous network embedding for massive graphs via reweighted personalized pagerank," in *Proc. VLDB Endow.*, vol. 13, no. 5, pp. 670–683, 2020.
[45] L. Backstrom and J. Leskovec, "Supervised random walks: Predicting and recommending links in social networks," in *Proc. Int. Conf. Web Search Data Mining*, 2011, pp. 635–644.
[46] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T. Chua, "Neural collaborative filtering," in *Proc. Int. Conf. World Wide Web*, 2017, pp. 173–182.
[47] G. Zheng et al., "DRN: A deep reinforcement learning framework for news recommendation," in *Proc. Int. Conf. World Wide Web*, 2018, pp. 167–176.

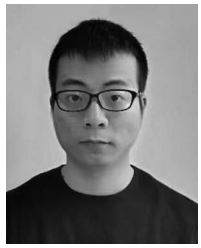

Qintian Guo received the BE degree from the University of Science and Technology of China and the MSc degree from the University of Alberta. He is currently working toward the PhD degree with the Department of Systems Engineering and Engineering Management, the Chinese University of Hong Kong. His research interests include graph data management and graph mining.

Dandan Lin received the PhD degree from The Hong Kong University of Science and Technology in 2021. Currently, she is a research fellow with the Shenzhen Institute of Computing Sciences. Her research interests include large-scale graph analysis, graph neural networks, and graph mining.

Sibo Wang received the BE degree in software engineering from Fudan University in 2011 and the PhD degree in computer science from Nanyang Technological University in 2016. He is an assistant professor with the Department of Systems Engineering and Engineering Management, the Chinese University of Hong Kong. He is currently interested in graph data management, Big Data analysis.

Raymond Chi-Wing Wong received the BSc, MPhil, and PhD degrees in computer science and engineering from the Chinese University of Hong Kong (CUHK), Hong Kong, in 2002, 2004, and 2008, respectively. He is a professor with the Department of Computer Science and Engineering, Hong Kong University of Science and Technology. His research interests include database and data mining.

Wenqing Lin received the PhD degree in computer science from Nanyang Technological University, in 2015. Currently, he is a principal researcher with Tencent in Shenzhen, China. His research interests include graph databases and data mining.

Authorized licensed use limited to: Hong Kong University of Science and Technology. Downloaded on July 08,2024 at 11:53:59 UTC from IEEE Xplore. Restrictions apply.