Reverse Ranking by Graph Structure: Model and Scalable Algorithms

Eliav Buchnik^{*}
School of Computer Science
Tel Aviv University, Israel
eliavbuh@gmail.com

Edith Cohen
Google Research, CA, USA
Tel Aviv University, Israel
edith@cohenwang.com

ABSTRACT

Distances in a network capture relations between nodes and are the basis of centrality, similarity, and influence measures. Often, however, the relevance of a node u to a node v is more precisely measured not by the magnitude of the distance, but by the number of nodes that are closer to v than u . That is, by the *rank* of u in an ordering of nodes by increasing distance from v .

We identify and address fundamental challenges in rank-based graph mining. We first consider single-source computation of reverse-ranks and design a “Dijkstra-like” algorithm which computes nodes in order of increasing approximate reverse rank while only traversing edges adjacent to returned nodes. We then define *reverse-rank influence*, which naturally extends reverse nearest neighbors influence [Korn and Muthukrishnan 2000] and builds on a well studied distance-based influence. We present near-linear algorithms for greedy approximate reverse-rank influence maximization. The design relies on our single-source algorithm. Our algorithms utilize near-linear preprocessing of the network to compute all-distance sketches. As a contribution of independent interest, we present a novel algorithm for computing these sketches, which have many other applications, on multi-core architectures.

We complement our algorithms by establishing the hardness of computing *exact* reverse-ranks for a single source and *exact* reverse-rank influence. This implies that when using near-linear algorithms, the small relative errors we obtain are the best we can currently hope for.

Finally, we conduct an experimental evaluation on graphs with tens of millions of edges, demonstrating both scalability and accuracy.

1. INTRODUCTION

Shortest-paths distances in a network are a classic measure of the relation between nodes and are the basis of similarity [29, 9], centrality [3, 33, 19, 4, 12, 31, 7, 10], and influence [21, 2, 15, 11] measures. Often, however, the relation of a node j to i is more correctly modeled not by the magnitude of the distance d_{ji} from j to i , but by i ’s position π_{ji} in an ordering of nodes according to increasing distance from j [14, 23, 37, 20]. A classic use of rank as an indicator of relevance in metric spaces is the k nearest neighbors (kNN) classifier, which classifies points based on the k closest labeled examples [14, 23]. In terms of popularity, kNN outweighs the respective distance-based classifiers, which instead use all examples within a certain distance.

^{*}Part of the research was done during an internship at Google.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGMETRICS ’16 June 14-18, 2016, Antibes Juan-Les-Pins, France

© 2016 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4266-7/16/06.

DOI: <http://dx.doi.org/10.1145/2896377.2901458>

Figure 1: Example undirected social network (edge lengths are proportional to drawn lengths).

More formally, we view a node j as *ranking* other nodes according to their distance order from j . The *rank* π_{ji} is the position of i in increasing order from j . Accordingly, from the perspective of node j , we can refer to π_{ij} as a *reverse rank*.

An advantage of using rank is that it provides a different signal than distance by “factoring out” the effects of uneven density. This is illustrated in the toy social network in Figure 1: We expect node A to be more important to node C than it is to node B , even though, A is closer to B than to C ($d_{CA} > d_{BA}$). This is because B has a dense neighborhood of closer nodes than A , but C has only two nodes closer to it than A . In terms of ranks, we have $\pi_{CA} = 3$ and $\pi_{BA} = 6$ and thus $\pi_{CA} < \pi_{BA}$, which reflects this intuition.

The rank relation is asymmetric: In the example network in Figure 1, we have $\pi_{BA} = 6$, since there are 5 nodes closer to B than A , and $\pi_{AB} = 1$, since B is the closest node to A . Therefore, $\pi_{AB} \neq \pi_{BA}$ even though the distance is symmetric ($d_{AB} = d_{BA}$). The asymmetry $\pi_{BA} > \pi_{AB}$ reflects our intuition that the higher degree node (B) has more influence on its neighbor (A) than the reverse.

In particular, with tie breaking on distances, a node v has exactly one nearest neighbor, but can have 0 to many *reverse nearest neighbors*, which are nodes u which satisfy $\pi_{uv} = 1$. In our example, node A has no reverse nearest neighbors. The number of reverse nearest neighbors of a point v is a well studied notion of v ’s influence, proposed by Korn and Muthukrishnan [26], and considered in metric spaces and in graphs [27].

In the basic model all nodes both rank and get ranked. A natural extension [26] allows only a subset of the nodes to get ranked (*rankees*) and also permits nodes that provide ranks (*rankers*) to have different weights. In this model, π_{ij} relates a ranker i to rankee j . This distinction is useful when nodes have two or more types of entities, for example, users (rankers) and content (rankees). It also allows us to specify a special small set of certified rankees such that we can characterize properties of other nodes by this smaller set of ranks. Importance weights $\beta(i) \geq 0$ assigned to rankers can correspond to properties like purchase power or trust level. The ranks assigned by ranker i are then weighted by $\beta(i)$. This weighting is

useful when we aggregate the scores of multiple rankers to obtain centrality/influence scores of rankees.

1.1 Contributions and Overview

Rank-based measures provide a natural alternative to distance-based ones, but algorithmically pose different challenges. We identify and motivate fundamental challenges and present scalable algorithmic tools which facilitating rank-based graph mining.

Reverse-rank Single-source computation

An important tool in working with distances is an efficient single-source computation. Dijkstra’s algorithm from a source s computes the distances d_{si} for all nodes i in near-linear time. A powerful property of Dijkstra’s algorithm is *sorted access*: Nodes are revealed in order of increasing distance from s . Thus, for any k , the k closest nodes to s are computed while traversing only edges adjacent to these k nodes. Therefore, if we are only interested in a prefix of the closest nodes, we can terminate the execution after they are computed, performing a fraction of the computation of a full execution. When we work with ranks, we would be instead interested in a prefix of the highest ranks. Such sorted access to rankings is also important for efficient aggregating rankings [16].

The *reverse-rank single-source problem* is, for a node $i \in U$, to compute the reverse ranks π_{ji} with respect to all nodes j . Moreover, we aim for an efficient algorithm that provides sorted-access: Listing nodes in increasing π_{ji} order with an algorithm that only traverses edges adjacent to listed nodes.

A naive solution for exact reverse-rank single source computation from i is to run Dijkstra’s algorithm from each node j , until node i is processed. For the average node, this is equivalent to performing n runs of Dijkstra until revealing on average $n/2$ nodes. Note that even on sparse networks, this scales quadratically in the number of nodes, which is prohibitive even on mid-size networks. This is in sharp contrast to the shortest-path single-source computation which takes (near) linear time. Previous work [27] proposed ways to scalably identify the set of reverse nearest neighbors of nodes, but did not address higher reverse ranks. We are able here (Section 6) to provide an explanation, establishing that the naive solution is in a sense the best we can do for the exact problem: We leverage the theory of subcubic equivalence [36] and construct a reduction from graph radius computation to reverse-rank single source computation. The former is known to have subcubic equivalence to all-pairs shortest-paths computation (APSP) [1].

This hardness result, fortunately, applies to the *exact* problem. An important contribution we make here (see Section 4) is devising a novel, scalable, Dijkstra-like (sorted access), *approximate* reverse-rank single-source algorithm, which provides estimates $\hat{\pi}_{ji}$ with a small relative error. Since ranks are intrinsically slightly noisy measures of the actual relations, estimates with a small relative errors are often as good as the exact values.

An essential component of our design is a preprocessing step where we compute All-Distances Sketches (ADS) for all nodes [6, 13, 7]. The sketches provide us with a fast oracle which estimates $\hat{\pi}_{ji}$ from the distance d_{ji} . In Section 3 we review the sketches and estimators as applied in our context. We note that we can apply any ADS algorithm, and existing designs are suitable for sequential, shared-memory, and node-centric message-passing computations [6, 32, 13, 5, 7]. A stand-alone contribution we make here is engineering an ADS algorithm for multicore architectures which provides provable tunable tradeoff between overhead and concurrency. Our algorithm can be used for many other applications of the sketches which include estimating distances, closeness similar-

ity, the distance distribution, and timed-influence [6, 32, 13, 5, 9, 15, 11, 7].

The estimation of single-source reverse ranks can be done by first running Dijkstra’s algorithm to compute the single-source distances, and then apply the oracle we obtained in the preprocessing step to the computed distances. This method, however, will not provide us sorted access. Our sorted access algorithm, similarly to Dijkstra, also traverses a shortest-path tree rooted at the source, but critically, instead of doing so in distance order, which would violate rank-based sorted access, does so in order of increasing estimated reverse ranks. The correctness of our design relies on key insights on properties of reverse ranks.

Reverse-rank influence

Distance or reachability-based notions of centrality and influence of a set S of seed nodes are fundamental measures in network analysis. In the general form [11], distance-based influence is defined with respect to a non-increasing decay function $\alpha(x) \geq 0$ (smoothing kernel) and node weights $\beta(i) \geq 0$. The contribution of each node j to the influence of S is proportional to its weight $\beta(j)$ and decays with the distance of j from S , $d_{Sj} = \min_{i \in S} d_{ij}$:

$$\text{Inf}^{(d)}(S) = \sum_j \beta(j) \alpha(d_{Sj}). \quad (1)$$

Well studied special cases include *Closeness centrality* [3, 33, 19, 4, 12, 31], where S contains a single node i , and the celebrated *reachability-based* influence model of [25], obtained when $\alpha(x) = 1$ for finite x and 0 otherwise. Distance-based influence with threshold function $\alpha(x) = 1$ for $x \leq T$ and $\alpha(x) = 0$ otherwise was studied in [21, 2, 15] (With distance interpreted as elapsed time).

Here we define *reverse-rank influence*

$$\text{Inf}^{(\pi^{-1})}(S) = \sum_j \beta(j) \alpha(\pi_{jS}), \quad (2)$$

where $\pi_{jS} = \min_{i \in S} \pi_{ji}$. The special case of $\text{Inf}^{(\pi^{-1})}(v)$, the influence of a single node, with $\alpha(1) = 1$ and $\alpha(x) = 0$ otherwise is the number of reverse nearest neighbors of v , is the influence measure proposed in [26, 27]. Our more flexible definition (2) is able to account for the contribution of nodes with higher reverse-rank to the influence of our node. For example, by setting $\alpha(x) = 1/x$ we achieve the effect that a reverse rank of x contributes $1/x$ to the total influence of v ; A node for which v is the 5th closest neighbor contributes to its influence 20% of what it would have contributed as a reverse nearest neighbor. With α being a T -threshold function, rankers u that rank v in their top T contribute $\beta(u)$ to v ’s influence.

Reverse-rank Influence Computation: We show (Section 6) that the computation of exact reverse-rank influence, even for a single node, and even when α is a threshold function, has subcubic equivalence to APSP. We therefore consider approximate influence $\widehat{\text{Inf}}$ computed using approximate ranks $\hat{\pi}$. Clearly $\widehat{\text{Inf}}(S)$ can be computed using $|S|$ single-source approximate reverse-rank computations (and using $\hat{\pi}_{jS} = \min_{i \in S} \hat{\pi}_{ji}$). Surprisingly, however, we show in Section 5, that even with large $|S|$, one single-source computation suffices.

Reverse-rank Influence Maximization: An important coverage problem which is extensively explored for reachability and distance-based influence, is *influence maximization (IM)* [25]: For a given $s \geq 1$, identify a set of s *seed* nodes with maximum influence. Intuitively, such a set provides the best “coverage” for its size with respect to the influence measure at hand. Here we consider IM with respect to our reverse-rank influence function $\text{Inf}^{(\pi^{-1})}$. The reverse-rank IM problem with α being a threshold function with parameter T on our example user and movies data set is to find a

set of s movies which maximizes the number of users for which there is at least one movie from S in their top T choices.

Similar to the distance-based influence function $\text{Inf}^{(d)}$, $\text{Inf}^{(\pi^{-1})}$ is monotone and submodular, and even for simple threshold α , when s is a parameter, the IM problem is NP hard. The most common and hugely successful algorithm for such coverage problems is the greedy algorithm [30], which iteratively builds a seed set by selecting in each step a node with maximum marginal contribution. For submodular and monotone functions, greedy has the property that each prefix of the sequence of size s has influence that is at least $1 - (1 - 1/s)^s \geq 1 - 1/e$ of the influence of the optimal seed set of that size [30]. Exact greedy, however, does not scale well for very large graphs. For reverse-rank influence, an exact greedy sequence can be computed in cubic time in the number of nodes. When all nodes are both rankers and rankees, the graph is sparse, and we work with a threshold function α with parameter T , the computation reduces to $O(nT)$, by performing a single-source search from all nodes to find the T nearest neighbors and computing a greedy cover. But even this special case does not scale well for large graphs for larger values of T .

Approximate greedy and heuristics had been extensively studied for reachability-based [28, 22, 35, 10] and distance-based [15, 11] influence. In particular, the SKIM algorithm [10, 11] computes in near-linear time a full greedy permutation so that each prefix of size s has approximation ratio of $1 - (1 - 1/s)^s - \epsilon$.

In Section 5 we present a near-linear algorithm for an approximate greedy sequence with respect to the *approximate* reverse-rank influence objective with threshold function: $\text{Inf}(S) = \{|z \in Z \mid \hat{\pi}_{zS} \leq T|\}$. We use the SKIM framework with our sorted-access approximate reverse-rank single-source computations as a critical component.

Experiments

Our experimental evaluation, detailed in Section 7 was focused on scalability and solution quality, using publicly available anonymized social graph data sets. Our ADS implementation runs on graphs with tens of millions of edges in tens of minutes on a single core, providing estimates with NRMSE (normalized mean square errors) of 6%-13%. Our multithreaded design achieved speedup factors of 3 to 4 on a machine with two CPUs and multiple cores.

With the preprocessing in place, our approximate reverse-rank single-source computations have similar running time to Dijkstra's algorithm (which computes single-source distances). In particular, a reverse-rank single-source computation was performed in less than 15 seconds on a single core on a graph with 4×10^6 nodes and 35×10^6 edges. For comparison, this should be contrasted with the running time of an exact reverse-rank single source computation, which would have taken an estimated 6000 hours on the same instance.

Using our implementation, we are able to visualize the reverse-rank distributions of some nodes in a large network, demonstrating how the distribution reveals information on the relative importance of a node in its locality. Prior to our work, it was not possible to scalably compute these distributions on large graphs.

Our approximate greedy IM implementation computes the full sequence on graphs with tens of millions of edges in minutes. We also observe that for small graphs or small values of T , where we could compute an exact greedy sequence, the solution quality of our approximate sequence was very close to the exact one.

2. PRELIMINARIES

We introduce some necessary notation. For a numeric function

$r : X$ over a set X , the function $k_r^{\text{th}}(X)$ returns the k^{th} smallest value in the range of r on X . If $|X| < k$, we define $k_r^{\text{th}}(X)$ as the supremum in the range of r . If r is not specified, we return the k^{th} smallest value in X .

We work with networks modeled as directed or undirected graphs $G = (V, E)$ with nodes $V = [n] = \{1, \dots, n\}$ and edges E with lengths $w(e) > 0$. We use $m = |E|$ for the number of edges. A subset or all nodes $U \subset V$ are specified as *rankee* nodes. We use the notation G^T for the *transpose graph*, which is the graph with edges reversed.

For nodes i, j , let d_{ij} be the shortest-paths distance from i to j . For $y \geq 0$, the *rankee y -neighborhood* of i is the set of rankee nodes within distance y from i . We denote the neighborhood by

$$N_i(y) = \{j \in U \mid d_{ij} \leq y\}$$

and its cardinality by $n_i(y) = |N_i(y)|$. The respective strict neighborhood is denoted by $\underline{N}_i(y) = \{j \in U \mid d_{ij} < y\}$ and its cardinality by $\underline{n}_i(y)$. For $i \in V$ and $j \in U$, π_{ij} denoted the *rank* of j with respect to i . When distances are unique, we have $\pi_{ij} = n_i(d_{ij})$, that is, equal to the number of rankee nodes that are at least as closer to i as j . When distances are not unique, we consider the range $(\underline{\pi}_{ij}, \bar{\pi}_{ij}]$, where

$$\begin{aligned} \underline{\pi}_{ij} &= \underline{n}_i(d_{ij}), \\ \bar{\pi}_{ij} &= n_i(d_{ij}). \end{aligned} \quad (3)$$

According to what we want to capture, we can define the rank π_{ij} as either, $\bar{\pi}_{ij}, \underline{\pi}_{ij} + 1$, a uniform at random choice from the range, or as the midpoint of this range: $\pi_{ij} \equiv \frac{\underline{\pi}_{ij} + 1 + \bar{\pi}_{ij}}{2}$. Our algorithms and implementation can be adapted to support all these choices.

An important ingredient of our design is the computation of a data structure which allows us to efficiently estimate the number of rankees in a neighborhood. That is, for a query specified by a node i and $d \geq 0$, return $\hat{n}_i(d)$. The data structure can be viewed as a set of lists $L(i)$, one for each node $i \in V$. Each list $L(i)$ consists of pairs (d, y) where d is a distance value and $y = \hat{n}_i(d) > 0$ is an estimate on $n_i(d)$. The lists are sorted and increasing in both d and y . To estimate for $n_i(x)$ and $\underline{n}_i(x)$, from the list $L(i)$, we use

$$\hat{n}_i(d) = y \text{ such that } (x, y) = \arg \max_{(x, y) \in L(i) \mid x \leq d} x \quad (4)$$

$$\underline{\hat{n}}_i(d) = y \text{ such that } (x, y) = \arg \max_{(x, y) \in L(i) \mid x < d} x. \quad (5)$$

That is, we look at the pair $(x, y) \in L(i)$ such that $x \leq d$ (or $x < d$) is maximum and return y . From the relations (3), we can obtain estimates $\hat{\pi}_{ij} = \hat{n}_i(d_{ij})$ and $\underline{\hat{\pi}}_{ij} = \underline{\hat{n}}_i(d_{ij})$ from $L(i)$ if we know d_{ij} .

The lists $L(i)$ are computed from All-Distances Sketches $\text{ADS}(i)$, which are the subject of the next section.

3. ALL-DISTANCES SKETCHES

We preprocess the graph to compute All-Distances Sketches (ADS) [6, 7] for the nodes in the graph. The sketches are defined with respect to a parameter k and a random permutation of rankee nodes. We find it convenient to work with $r(i) \in [0, 1]$ which is the permutation position of i divided by $|U|$. Alternatively, it is sometimes convenient to work instead with random hash based $r(i) \sim U[0, 1]$. The sketch $\text{ADS}(i)$ of a node $i \in V$ consists of a set of entries of the form (j, d_{ij}) , consisting of a node $j \in U$ and the distance d_{ij} . We assume that $r(j)$ is either included in the entry or can be easily retrieved from j . The set of rankee nodes included in $\text{ADS}(i)$ is a random variable which depends on the assignment r :

$$j \in \text{ADS}(i) \iff r(j) \leq k_r^{\text{th}}\{h \in U \mid d_{ih} \leq d_{ij}\}. \quad (6)$$

This ADS definition (6) applies with unique and non-unique distances. A technical point is that for estimation with non-unique distances, we also maintain with $\text{ADS}(i)$, as auxiliary, entries (j, d_{ij}) that satisfy for some $z \in \text{ADS}(i)$ $r(j) = k_r^{\text{th}}\{h \in U \setminus \{z\} \mid d_{ih} \leq d_{iz}\}$ when these entries are not already included in $\text{ADS}(i)$ [8] (When distances are unique, all these entries are already in $\text{ADS}(i)$). With unique distances, the expected size of the sketches is exactly $\sum_{i=1}^{|U|} \min\{1, k/i\} \leq k \ln |U|$ with good concentration, but the sketch can be much smaller when distances are not unique, while providing the same statistical guarantees on estimate quality, which is why we separately treat non-unique distances rather than tie break.

Our implementation of ADS computation is based on PRUNED DIJKSTRA'S [6, 13, 7]. The pseudocode for the basic sequential version is provided as Algorithm 1 and uses $O(km \ln n)$ edge traversals. When applied with non-unique distances, the algorithm also includes the auxiliary entries.

Algorithm 1 ADS set for G via PRUNED DIJKSTRA'S

```

for rankee node  $u \in U$  by increasing  $r(u)$  do
  Run Dijkstra's algorithm from  $u$  on  $G^T$ 
  foreach scanned node  $v$  do
    if  $d_{vu} > k^{\text{th}}\{y \mid (x, y) \in \text{ADS}(v)\}$  or
     $|\{x \in \text{ADS}(v) \mid d_{vx} \leq d_{vu}\}| > k$  then
      prune Dijkstra at  $v$ 
    else
       $\text{ADS}(v) \leftarrow \text{ADS}(v) \cup \{(r(u), d_{vu})\}$ 

```

The term *scanned node* in the pseudocode refers to the event where the node $v \in V$ is popped from the Dijkstra priority queue. Each node can be scanned at most once in each (pruned) Dijkstra search. The scanned nodes are always a prefix of the nodes when sorted by increasing distance from u in G^T . When a node v is scanned, either u is inserted to $\text{ADS}(v)$ or the search is pruned at v . Therefore, the number of node scans is equal to the ADS size.

The algorithm builds the ADS of all nodes by considering one node $u \in U$ at a time and adding it as an entry in $\text{ADS}(v)$ for all relevant v . To do so efficiently, we maintain the entries in $\text{ADS}(v)$ as an array sorted by decreasing distances. The insertion condition then amounts to testing if $|\text{ADS}(v)| < k$ or if the entry (x, y) in the $|\text{ADS}(v)| - k$ position in the array (the k th smallest distance) has $y > d_{vu}$, or if it has $y = d_{vu}$ but either $|\text{ADS}(v)| = k$ or the entry (x, z) in the $|\text{ADS}(v)| - k - 1$ position has $z > d_{vu}$.

We refer to the k th smallest distance in $\text{ADS}(v)$ as the *threshold distance* and denote it by $\Delta(v)$. We also use the notation $\ast(v)$ for the bit indicating if the $k + 1$ th smallest is equal to the k th smallest distance. The prune condition can then be written as

$$d_{vu} > \Delta(v) \text{ or } d_{vu} = \Delta(v) \text{ and } \ast(v). \quad (7)$$

Observe that insertions can only affect the k last entries in the current ADS. Therefore, it suffices to keep only that “tail” part in active memory. When k is small we can keep it as an array and implement insertions by shifting. When k is larger we can use a data structure that supports efficient insertions.

3.1 Multithreading

PRUNED DIJKSTRA'S, as stated, sequentially performs possibly dependent searches from all rankee nodes. We propose here a design which allows us to control in a principled way the tradeoff between overhead and concurrency. We partition the $|U|$ pruned Dijkstra searches to batches, where each batch is a consecutive set of nodes when ordered by increasing r . All the searches in the same

batch are made independent so that they can be executed concurrently. Each search computes a set of *proposed entries* to sketches $\text{ADS}(v)$, as contributions to a set $PE(v)$. A proposed entry is created when a node v is visited and the pruning condition (7) is not satisfied. The pruning, however, and hence the proposed entries, are computed with respect to the set of threshold distances and bits $(\Delta(v), \ast(v))$ for $v \in V$, as it was at the beginning of the batch. Pseudocode for an independent search thread is provided as Algorithm 2. Each such Dijkstra search may generate a proposed ADS entry for multiple nodes.

At the end of a batch, for each node v , the proposed entries $PE(v)$ from all the searches in the batch are merged with (the k -tail of) $\text{ADS}(v)$ (as it was in the beginning of the batch) to compute an updated $\text{ADS}(v)$ with respect to the end of the batch. The merge is performed by scanning the entries (u, d_{vu}) in $PE(v)$ in order of increasing r and applying the insertion procedure to $\text{ADS}(v)$ as used in Algorithm 1: If the pruning condition 7 does not hold, we insert u and update $\text{ADS}(v)$ (note that this updates $\Delta(v)$ and $\ast(v)$). Note that not all proposed entries are incorporated, since the insertion rule is not satisfied with respect to the updated $(\Delta(v), \ast(v))$ after processing previous $PE(v)$ entries.

Algorithm 2 A PRUNED DIJKSTRA thread (search from u)

```

Run Dijkstra's algorithm from  $u$  on  $G^T$ 
foreach scanned node  $v$  do
  if  $d_{vu} > \Delta(v)$  or  $d_{vu} = \Delta(v)$  and  $\ast(v)$  then
    prune Dijkstra at  $v$ 
  else
     $PE(v) \leftarrow PE(v) \cup \{(u, d_{vu})\}$ 

```

3.2 Concurrency/Overhead tradeoff analysis

The sequential algorithm has the property that all generated entries constitute final ADS entries. The multithreading algorithm computes proposed entries that may be eventually discarded. These discarded entries are the overhead of the multithreading algorithm.

More precisely, we define the overhead as the ratio of the expected number of discarded entries to the expected number of ADS entries per node. The overhead depends on how we partition the searches to batches. Placing each search in a separate batch would result in no overhead, but also no concurrency. Putting all searches in the same batch would have a very large overhead, as none of the searches would be pruned.

Note that the overhead of discarded entries corresponds to an overhead on edge traversals, which are the main cost of the algorithm. In particular, we can bound the total work performed by the multithreaded algorithm by multiplying the sequential bound of $km \ln |U|$ by $(1 + h)$, where h is a bound on the overhead.

We next propose batch partitions which allow us to bound the overhead. We first observe that the search is never pruned for the k nodes with lowest r values. Our first batch would contain these nodes, and we can perform those searches independently without overhead. At the end of this first batch, all generated proposed entries $PE(i)$ would be sorted by distance to form $\text{ADS}(i)$ with respect to those k nodes. As for subsequent batches, we propose exponentially increasing batch sizes and show the following:

LEMMA 3.1. *For a parameter $\mu > 0$, consider a partition to batches so that the j th batch ends at node in position $\lceil (1 + \mu)^{j-1} k \rceil$ in the sorted order by increasing $r(v)$. Then the expected overhead is at most $h \leq \mu / \ln(1 + \mu) - 1$.*

PROOF. Consider processing a batch that starts at position $b_0 + 1$. The probability of a node in the batch to be included in $PE(i)$

is $\min\{1, \frac{k}{b_0+1}\}$. Note that to generate a proposed entry, the node needs to be with distance smaller than $\Delta(i)$, that is be among the k smallest distances among all the nodes processed up to the previous batch and itself. This probability is exactly that of being in one of the first k positions in a random permutation of $b_0 + 1$ nodes, which is $\min\{1, \frac{k}{b_0+1}\}$. Now we can consider the probability that a node in the batch is a final member of $\text{ADS}(i)$. If the node is in position $b_0 + j$, the probability is $\min\{1, \frac{k}{b_0+j}\}$.

We now consider a batch that has nodes in permutation positions $b_0 + 1$ to b_t , such that $b_0 > k$. The ratio of good work to total work is

$$\frac{\sum_{j=1}^{b_t-b_0} 1/(b_0+j)}{(b_t-b_0)/(b_0+1)} = \frac{b_0+1}{b_t-b_0} \sum_{j=1}^{b_t-b_0} \frac{1}{b_0+j}$$

$$\approx \frac{b_0+1}{b_t-b_0} \ln\left(\frac{b_t}{b_0+1}\right)$$

Thus, if we choose $b_t = (1 + \mu)b_0$ for some $\mu > 0$, we obtain the ratio $\ln(1 + \mu)/\mu$. The overhead, by definition, is the inverse of this ratio minus 1. $\square$

In particular, we can see that $\mu = 0.5$ results in overhead of 20% more edge traversals than the sequential algorithm. Using $\mu = 0.1$, has overhead of about 5%. The total number of batches is $O(\log_{1+\mu}(|U|/k)) \approx O(\frac{\log(|U|/k)}{\mu})$, which is logarithmic in the number of rankee nodes.

3.3 Cardinality estimation

We now discuss the computation of a list $L(i)$ from $\text{ADS}(i)$. Recall that $L(i)$ is a list of pairs in the form $(d, \hat{n}_i(d))$. There is one pair for each unique distances d in $\text{ADS}(i)$, and we assume $L(i)$ is sorted by increasing d .

We review two estimators $\hat{n}_i(d)$: The bottom- k estimator and the HIP estimator. Both are unbiased, nonnegative, and have a small relative error, with good concentration which depend on the ADS parameter k . The HIP estimate is tighter: Estimates are at least as good as bottom- k , and with unique distances, has half the variance of the bottom- k estimator. The bottom- k estimator, however, is useful to us because it has some monotonicity property. We can compute the lists $L(i)$ using either estimator or both.

The bottom- k estimator: This inverse probability estimator [24] has coefficient of variation (CV) at most $1/\sqrt{k-2}$ [6, 7]. To estimate $n_i(x)$, we take the k th smallest r value among nodes in $N_i(x)$, which we denote by τ . If there are fewer than k nodes in $N_i(x)$, we return the number of entries as our estimate. Otherwise, we compute the probability p that an r -value is below τ . When $r(v) \sim U[0, 1]$, $p = \tau$. We then use the estimate $\hat{n}_i(x) = (k-1)/p$.

The HIP estimator: This estimator has CV at most $1/\sqrt{k-2}$ and with unique distances is most $1/\sqrt{2k-2}$ [7] (see [8] for extension to non-unique distances). The estimates are obtained as follows: For each (non auxiliary) entry j in $\text{ADS}(i)$, we compute the threshold value

$$\tau_{ij} = k_r^{\text{th}}\{h \in \text{ADS}(i) \mid d_{ih} < d_{ij}\}. \quad (8)$$

We then compute p_{ij} as the probability of $r(j) < \tau_{ij}$. If there are fewer than k entries lower than d_{ij} then $p_{ij} = 1$. Otherwise, when $r(j) \sim U[0, 1]$, we have $p_{ij} = \tau_{ij}$. We then take $a_{ij} = 1/p_{ij}$. Finally, the HIP estimate (summed over non-auxiliary entries) is

$$\hat{n}_i(x) = \sum_{j \in \text{ADS}(i) \mid d_{ij} \leq x} a_{ij}.$$

Computing the estimates The estimation list $L(i)$ for both the bottom- k and the HIP estimators can be computed by processing the entries of $\text{ADS}(i)$ in increasing distance order, maintaining the k th smallest values in the prefix processed so far and accordingly the k th smallest value τ , and computing the estimates $\hat{n}_i(d)$ when entries of distance d are processed.

An easy to verify property that is useful to us is that the neighborhood size estimates for each node are non-decreasing with distance from the node and therefore can only increase when the distance does:

LEMMA 3.2. *When $L(i)$ is computed using either bottom- k and HIP estimates, the estimates (4) and (5) satisfy*

$$d_1 \leq d_2 \implies \hat{n}_i(d_1) \leq \hat{n}_i(d_2)$$

$$d_1 \leq d_2 \implies \underline{\hat{n}}_i(d_1) \leq \underline{\hat{n}}_i(d_2).$$

4. REVERSE-RANK SINGLE-SOURCE

As noted in the introduction, if we are interested in computing reverse-ranks from a source i to all nodes, we can compute the distances d_{ji} by applying Dijkstra's algorithm from i on G^T , and return estimated reverse-ranks from the distances using (4) and (5). The nodes, however, are processed in order of increasing distance, which does not necessarily corresponds to the order by increasing reverse ranks (recall the example in the introduction). Therefore, if we are only interested in correctly identifying nodes with highest reverse ranks and we apply this algorithm, we can not prune the computation and we will scan a much larger portion of the graph than needed.

In this section we present an approximate reverse-rank single-source algorithm that provides sorted-access: Computing nodes in order of increasing (approximate) reverse-rank. We start by establishing a basic monotonicity property of reverse-ranks that is essential for the correctness of our design.

4.1 SP monotonicity of reverse-ranks

We noted in the introduction that reverse-rank order does not necessarily correspond to distance order. For nodes on a shortest path, however, we can show that the reverse-ranks, and the respective bottom- k estimates, are monotone:

LEMMA 4.1. *Consider a shortest path $i_t, \dots, i_0$ in G . Then $\bar{\pi}_{i_j i_0}$, $\underline{\pi}_{i_j i_0}$, and the bottom- k estimates $\hat{\pi}_{i_j i_0}$ and $\bar{\pi}_{i_j i_0}$, are all non-decreasing with j .*

PROOF. Consider $j < h$, then $d_{i_j i_0} < d_{i_h i_0}$. The neighborhoods relations is $N_{i_j}(d_{i_j i_0}) \subseteq N_{i_h}(d_{i_h i_0})$. Therefore, $\bar{\pi}_{i_j i_0} \leq \bar{\pi}_{i_h i_0}$. Similarly, $\underline{\pi}_{i_j i_0} \leq \underline{\pi}_{i_h i_0}$ and thus the cardinalities satisfy $\hat{\pi}_{i_j i_0} \leq \hat{\pi}_{i_h i_0}$.

In the case of bottom- k estimates, the claim follows again from containment of neighborhoods. Let τ_1 be the k th smallest r value in the contained set and let τ_2 be the k th smallest r value in the containing set. Then clearly, $\tau_1 \geq \tau_2$. Recall that the bottom- k cardinality estimate is $(k-1)/\tau$. We have $(k-1)/\tau_1 \leq (k-1)/\tau_2$. $\square$

4.2 Algorithm and analysis

The pseudocode for our reverse-rank single source algorithms is presented as Algorithm 1. The algorithm has the same structure as Dijkstra's algorithm from source s in the transposed graph G^T . The algorithm maintains each unprocessed node j that is adjacent to an already processed node in a min priority queue. The entry contains $(j, \bar{d}_{js})$, where $\bar{d}_{js}$ is the upper bound on the distance from j to s . This upper bound serves as the priority in Dijkstra's algorithm and

is the minimum over processed nodes h of $d_{hs} + w_{jh}$. Our reverse-rank single source algorithm uses instead a priority as follows. We first look at $\hat{n}_{\bar{d}_{js}}(j)$, which is an upper bound on the estimated reverse-rank, computed according to the best current upper bound $\bar{d}_{js}$ on the distance. From Lemma 3.2, $\hat{n}_{\bar{d}_{js}}(j) \geq \hat{n}_{d_{js}}(j)$. Therefore, when the upper bounds on the distances are tightened, the priority can only decrease. Now, two nodes in the priority queue can have the same estimate $\hat{n}_{\bar{d}_{js}}(j)$. In this case, we break ties according to the distance upper bounds $\bar{d}_{js}$, always preferring the node with lower $\bar{d}_{js}$. If both $\hat{n}_{\bar{d}_{js}}(j)$ and $\bar{d}_{js}$ are the same, the tie can be broken arbitrarily.

The next node h that is selected from the queue is the one with minimum priority according to lexicographic order on $(\hat{n}_{\bar{d}_{js}}(j), \bar{d}_{js})$.

For this node h we set $d_{hs} \leftarrow \bar{d}_{hs}$ (a correctness proof that indeed $\bar{d}_{hs}$ is the distance is provided below). We then scan all in-coming edges (j, h) . If j is not already in the priority queue, we insert it with $\bar{d}_{ji} = d_{hs} + w(j, h)$ and the respective priority. If j is already in the queue we compare $x \leftarrow d_{hs} + w(j, h)$ to the current $\bar{d}_{js}$. If $x < \bar{d}_{js}$, we update $\bar{d}_{js} \leftarrow x$ and update the priority to $\hat{n}_{\bar{d}_{js}}(j)$.

Note that the algorithm applies for both directed and undirected graphs. When applied to directed graphs, the algorithm returns reverse ranks only for nodes that can reach s . For completeness, we explain how to extend this, if needed, also for nodes v that can not reach s , that is, $d_{vs} = \infty$. We first need to precisely define the rank π_{vs} in this case. All rankee nodes that can not be reached from s can be viewed as having rank range $(|R_v|, |U|]$, where R_v is the set of rankee nodes reachable from v . Now note that we can estimate $|R_v|$ by the cardinality estimate associated with the maximum-distance entry in $L(v)$.

Algorithm 1: Approximate Reverse-rank Single-Source

Input: Source node s
Output: A sequence $(v, d_{sv}, \hat{\pi}_{vs})$ in increasing $\hat{\pi}_{vs}$
 // Node object v : $v.dist, v.rDr$ are upper bounds
 on d_{sv} and $\hat{\pi}_{vs}$. Initialization $v.Init(d, r)$:
 $v.dist \leftarrow d$; $v.rDr \leftarrow r$
 $Q \leftarrow \emptyset$ // Initialize an empty min priority queue
 Q of node objects prioritized by lex order of
 $(v.rDr, v.dist)$
 $s.Init(1, 0)$ // Initialize source node object
 $Q.add(s)$ // put source node in queue
while Q is not empty **do**
 $v \leftarrow Q.extract_min()$
 Output (v) // output and scan v
foreach $u \mid (u, v) \in E$ and u not scanned **do**
 $d \leftarrow v.dist + w_{vu}$
if $u \notin Q$ **then**
 $u.Init(\hat{n}_u(d), d)$
 $Q.add(u)$
else
if $d < u.dist$ **then**
 $Q.decrease_key(u, (\hat{n}_u(d), d))$ // Update
 priority of u in Q : $u.dist \leftarrow d$;
 $u.rDr \leftarrow \hat{n}_u(d)$

We now establish correctness, that is, when a node v is popped from the priority queue, its priority is its correct (approximate) reverse rank and distance. We also show that the algorithm traverses a shortest path tree (but not necessarily in increasing distance order).

THEOREM 4.1. *When Algorithm 1 is applied with exact cardinalities $n_i(d)$ or with bottom- k estimates $\hat{n}_i(d)$, it correctly computes reverse ranks (respectively, approximate reverse ranks) and*

distances. Moreover, the traversal forms a shortest-paths tree from s .

PROOF. Let s be the source node. Consider a node i and its shortest path $s = x_0, \dots, x_t = i$ from s . We show correctness by induction on the number of nodes t in the shortest path. Recall that the algorithm maintains a minimum heap sorted lexicographically by a priority that bound the (approximate) reverse rank and distance. We denote the distance upper bound by $\delta(s, i)$. Since it always corresponds to some actual path, $\delta(s, i) \geq d_{si}$. The reverse rank component of the priority is computed from $\delta(s, i)$ and therefore upper bound the reverse rank.

In each iteration we scan (traverse) the top of the heap. The base case where $t = 0$ is trivial: The source node s is inserted to the heap with priority equal to its reverse rank of 1 and distance of 0, and is then scanned. By the induction hypothesis, the priority is correct, that is $\delta(s, x_{t-1}) = d_{sx_{t-1}}$, for node x_{t-1} when it is scanned by the algorithm. At that point $\delta(s, x_t) \geq d_{sx_t} > d_{sx_{t-1}}$. Therefore from the monotonicity property (Lemma 4.1), the priority of x_t is lower than that of x_{t-1} , and therefore x_t was not yet scanned. When scanning x_{t-1} , the algorithm relaxes all outgoing edges including the edge (x_{t-1}, x_t) . The relaxation procedure assigns $\delta(x_t, s)$ to the correct value and thus its heap priority is now its correct (approximate) reverse rank and distance. When it will be scanned, it will be with the correct priority. $\square$

5. REVERSE-RANK INFLUENCE

In this section, we consider the computation and maximization of reverse-rank influence. Consider a graph with a set of rankee nodes $U \subset V$ and ranks π_{ji} defined for rankees $i \in U$ and $j \in V$. Let $\beta(j) \geq 0$ be the ranker weights of $j \in V$. For a set $S \subset U$ of seed nodes, the reverse-rank influence is $\text{Inf}(S) = \sum_{j \in Z} \beta(j) \alpha(\pi_{js})$, where $Z \subset V$ is the set of ranker nodes (those with $\beta(z) > 0$). From Corollary 6.2, the exact computation of $\text{Inf}(S)$ has subcubic equivalence to APSP, even when restricted to threshold functions α and a single seed.

We therefore focus on scalably computing the approximate influence

$$\widehat{\text{Inf}}(S) = \sum_{j \in Z} \beta(j) \alpha(\hat{\pi}_{js}),$$

where $\hat{\pi}_{js} = \min_{i \in S} \hat{\pi}_{ji}$.

Note that to compute $\widehat{\text{Inf}}(S)$ it suffices to compute $\hat{\pi}_{js}$ for all ranker nodes $j \in Z$. Moreover, when α is a threshold function for some $T \ll n$, or more generally, any function with $\alpha(x) = 0$ for all $x > T$, it suffices to compute $\hat{\pi}_{js}$ only for nodes with $\hat{\pi}_{js} \leq T$. A naive way to compute these values is to perform, from each seed $i \in S$, a single-source reverse ranks search from i , using Algorithm 1, and terminate the search when we scan a node with $\hat{\pi}_{ji} > T$. We can then combine the results of the different searches, computing the minimum $\hat{\pi}_{ji}$ of each node j that is scanned in at least one of the searches, to obtain the values $\hat{\pi}_{js}$. This naive computation requires $|S||E| \log n$ operations (assuming the lists $L(j)$ are provided) when T is large. But even with smaller T , a node j can be scanned multiple times, once for each seed $i \in S$ with $\hat{\pi}_{ji} \leq T$. We now show how to remove the dependence on the number of seeds $|S|$.

THEOREM 5.1. *For a set of seeds S , we can compute the values $\hat{\pi}_{js}$ for all $j \in V$ using $O(|S| + |E| \log n)$ operations. When $\alpha(x) = 0$ for $x \geq T$, $|E|$ is replaced by the number of incoming edges to nodes j that satisfy $\hat{\pi}_{js} \leq T$. These bounds assume that the lists $L(j)$ are provided for $j \in V$.*

PROOF. We slightly modify Algorithm 1 by initializing the priority queue with entries with priorities $(i.\text{dist}, i.\text{rDr}) = (0, 1)$ for each node $i \in S$. The algorithm execution then proceeds as with a single source node. For correctness, we can show that nodes are scanned (popped from the queue) in increasing lex order of $(\hat{n}_j(d_{jS}), d_{jS})$, and at the point they are scanned we have $j.\text{dist} = d_{jS}$ and thus $j.\text{rDr} = \hat{n}_j(d_{jS})$.

To see that, note that d_{jS} suffices to obtain $\hat{\pi}_{jS}$. This is because, using Lemma 3.2, $\hat{n}_j(d_{jS}) \equiv \min_{i \in S} \hat{n}_j(d_{ji}) = \hat{n}_j(\min_{i \in S} d_{ji})$.

For correctness, we need to show that the monotonicity property (Lemma 4.1) to sets S : Consider a shortest path $i_t, \dots, i_0$ from i_t to S (that is, i_0 is the closest S node to i_t). Note that this implies that i_0 is the closest S node to all i_j .

It follows from 4.1 that $\bar{\pi}_{i_j i_0}$, $\underline{\pi}_{i_j i_0}$, and the bottom- k estimates $\hat{\pi}_{i_j i_0}$ and $\hat{\pi}_{i_j i_0}$ are non-decreasing with j . We observe that $\bar{\pi}_{i_j i_0} = \bar{\pi}_{i_j S}$ and similarly $\underline{\pi}_{i_j i_0} = \underline{\pi}_{i_j S}$ and this holds for the bottom- k estimates obtained using $d_{i_j S}$. Therefore, monotonicity holds also when i_0 is substituted with S . $\square$

5.1 Influence Maximization

Here we consider uniform ranker weights $\beta(z) = 1$ and α that is a threshold function for some T . The influence of a set $S \subset U$ of rankee nodes is then the number of rankers that have at least one node from S among their top T rankees:

$$\text{Inf}(S) = |\{z \in Z \mid \pi_{zS} \leq T\}|. \quad (9)$$

The goal of the IM problem is to find a set S of rankee nodes of a certain size which maximizes $\text{Inf}(S)$. A common approach to such coverage problems is the greedy algorithm. Greedy repeatedly selects a rankee node which has maximum marginal influence. For each $s \geq 1$, the set S of the first s selected seeds is guaranteed to have influence that is at least $1 - (1 - 1/s)^s$ of the maximum possible by s seeds. Algorithm 2 is an exact greedy algorithm for our reverse-rank IM problem with influence function (9). The computation of the algorithm is dominated by Dijkstra computations from each ranker node that are stopped once T rankees are popped from the priority queue. Recall that when distance are not unique, we can work with multiple definition of the rank (see Section 2), but with all of them, we can determine the ranks once at most $T + 1$ rankees are popped. example, when we use $\bar{\pi}$, then if the $T + 1$ rankee has the same distance d as the T rankee, then all rankees of distance d are excluded (they have rank larger than T). Even when all nodes are rankees, and thus at most T nodes are popped in total in each Dijkstra run, the required computation is $\Omega(T|E| \log |V|)$, which does not scale well for large values of the threshold T .

5.2 Approximate Greedy IM

We next obtain a near-linear algorithm using two relaxations. First, the greedy selection, and thus the statistical guarantees we obtain, are with respect to the relaxed influence function (9) where $\hat{\pi}$ replaces π :

$$\widehat{\text{Inf}}(S) = |\{z \in Z \mid \hat{\pi}_{zS} \leq T\}|. \quad (10)$$

Second, we do not compute an exact greedy sequence for $\widehat{\text{Inf}}(S)$ but instead use an approximate greedy algorithm: At each step, selects a node with marginal influence that is approximately (within a small relative error) the maximum.

Our design adapts the influence maximization algorithms SKIM and T-SKIM[10, 11] which are designed for reachability-based [25] and distance-based [21, 15, 11] influence with threshold functions. We quickly review SKIM, which remarkably, when all nodes are both rankers and rankees, computes a full approximate

Algorithm 2: Exact greedy reverse-rank IM

Input: Directed graph $G = (V, E)$, ranker nodes $Z \subset V$, rankee nodes $U \subset Z$, threshold T

Output: Exact greedy sequence seedlist

seedlist $\leftarrow \perp$ // Output list of (seed, marginal influence)

forall the rankee nodes $u \in U$ **do** coverage[u] $\leftarrow \emptyset$

forall the ranker nodes $z \in Z$ **do**

coverers[z] $\leftarrow \emptyset$

Run Dijkstra from z in G , until (we can determine that) $\pi_{zu} > T$

foreach rankee $u \in U$ with $\pi_{zu} \leq T$ **do**

coverers[z] $\leftarrow$ coverers[z] $\cup \{u\}$

coverage[u] $\leftarrow$ coverage[u] $\cup \{z\}$

while There are rankees u with |coverage[u] | > 0 **do**

$v \leftarrow \arg \max_{u \in U \setminus S} |\text{coverage}[u]|$

Append $(u, |\text{coverage}[u]|)$ to seedlist

foreach $z \in \text{coverage}[u]$ **do**

foreach $v \in \text{coverers}[z]$ **do**

Remove z from coverage[v]

Delete coverage[u]

return seedlist

greedy permutation in near linear time. To do so efficiently, SKIM samples nodes not covered by previously selected seeds, and maintains for each candidate seed node the number of sampled nodes it covers. Reachability-based SKIM performs a pruned reverse graph searches from the node to determine the nodes that cover it. The distance-based SKIM performs backward pruned Dijkstra searches. The node that first reaches some number K of samples has approximately maximum marginal influence and is selected as a seed. The *sample-size parameter* K determines a tradeoff between computation and accuracy. SKIM then updates the samples so that they are with respect to the updated marginal influences with the coverage of the new seed node removed. SKIM also updates the representation of the residual problem. The updates are performed using a respective forward (Dijkstra) search from the new seed to reveal all nodes that it covers. When a previously sampled node becomes covered, the samples of the nodes covering it are adjusted to reflect their reduced marginal coverage. Sampling is then resumed until another node reaches a sample size of K . We repeat the process of sampling, selecting a seed, and updating the residual problem until a desired number of seeds is selected, a desired coverage is achieved, or all nodes are covered.

Our Algorithm 3, reverse-rank SKIM (RR-SKIM), follows the SKIM design, of iterating the selection of a new seed node (rankee) via sample building and updates. The reverse-rank problem, however, requires some critical adaptations.

When sample building, we repeatedly select random uncovered ranker nodes z . We then run Dijkstra's algorithm from z but stop the search when the approximate rank $\hat{\pi}_{zu}$ exceeds T . For each visited rankee node u , we increment the sample size sample_size[u] and also add u to a list inverted_sample[z] (the list of nodes where z is included in the sample). This process stops when the first rankee u reaches sample_size[u] = K . The node u then becomes the next seed node. We then apply our sorted-access reverse-rank single source computation from u , up to rank T , to determine the coverage of the new seed u . We mark all uncovered visited nodes as covered. For each newly covered ranker z , we scan inverted_sample[z] and decrement sample_size[v] for each $v \in \text{inverted_sample}[z]$. We then delete inverted_sample[z]. For each covered ranker z , we maintain the (approximate) rank of the best seed best_seed[z].rank = $\min_{v \in S} \hat{\pi}_{vz}$ and the corresponding minimum distance best_seed[z].dist = $\min_{v \in S} d_{vz}$

(note that the node with minimum distance must have minimum estimated rank). The purpose of maintaining `best_seed` is to enable pruning of reverse searches. Pruning is critical for the near-linear computation bound of the algorithm (without it, we can construct examples where the bulk of covered nodes is revisited with each new seed, resulting in $\Omega(|\text{seedlist}|m)$ computation).

A search from the new seed u is always pruned at z when $\hat{\pi}_{uz} > T$, but is also pruned when

$$\begin{aligned} \hat{\pi}_{uz} &> \text{best_seed}[z].\text{rank} \text{ or} \\ \hat{\pi}_{uz} &= \text{best_seed}[z].\text{rank} \text{ and } d_{uz} \geq \text{best_seed}[z].\text{dist} . \end{aligned} \quad (11)$$

We now need to show that also with this pruning, the algorithm maintains the following invariant

LEMMA 5.1. *After the processing of a new seed node, all nodes z with $\min_{v \in S} \hat{\pi}_{vz} \leq T$ have $\text{best_seed}[z].\text{rank} = \min_{v \in S} \hat{\pi}_{zv}$ and $\text{best_seed}[z].\text{dist} = \min_{v \in S} d_{zv}$ and all other nodes have $\text{best_seed}[z].\text{rank} = +\infty$.*

PROOF. This property clearly holds when pruning only when $\hat{\pi}_{uz} > T$, since after inserting a new seed node u our reverse-rank search from u visits all nodes with $\hat{\pi}_{vz} \leq T$.

We establish the claim using induction on added seeds. The base of the induction is when S is empty and $\text{best_seed}[z].\text{rank} = +\infty$. Assume now that our invariant holds and let s_2 be a newly selected seed node. Let u_1 be a node on which we prune the search from s_2 . From the condition (11) there exist a seed node s_1 such that $\hat{\pi}_{u_1, s_1} < \hat{\pi}_{u_1, s_2}$ or $\hat{\pi}_{u_1, s_1} = \hat{\pi}_{u_1, s_2}$ and $d_{u_1, s_1} \leq d_{u_1, s_2}$. From the definition of our estimators, $\hat{\pi}_{u_1, s_1} < \hat{\pi}_{u_1, s_2}$ implies $d_{u_1, s_1} < d_{u_1, s_2}$. Combining, we obtain that $d_{u_1, s_1} \leq d_{u_1, s_2}$.

Now assume to the contrary there is a node u_2 such that u_1 is on the shortest path from u_2 to s_2 and $\hat{\pi}_{u_2, s_2} < \hat{\pi}_{u_2, s_1}$ or $\hat{\pi}_{u_2, s_2} = \hat{\pi}_{u_2, s_1}$ and $d_{u_2, s_2} < d_{u_2, s_1}$. We show that this is not possible.

Using the above and triangle inequality we obtain $d_{u_2, s_1} \leq d_{u_2, u_1} + d_{u_1, s_1} \leq d_{u_2, u_1} + d_{u_1, s_2} = d_{u_2, s_2}$. A property of our estimates is that for any three nodes $d_{u_2, s_1} \leq d_{u_2, s_2}$ implies $\hat{\pi}_{u_2, s_1} \leq \hat{\pi}_{u_2, s_2}$. $\square$

The analysis of computation and approximation quality uses components from the analysis of T -SKIM[11]. An important critical component in the analysis is that we can “charge” edge traversals used for sample building to increases in sample sizes. When there are many non-rankees, we can construct worst-case graph where non-rankees are repeatedly traversed without incrementing sample counts. In realistic models, however, and when all nodes are rankers or rankees, we would expect such popular ranker hub nodes to be covered quickly by the first few selected seeds. Another component of the analysis that carries over from T -SKIM is bounding the number of updates to `best_seed`[z]. The argument there critically relies on the sample based approximate greedy selection. The approximation quality of the algorithm can only be guaranteed probabilistically and with respect to approximate ranks $\hat{\pi}_{uz}$. To summarize, when we run the algorithm with $K = O(\epsilon^{-2} \log n)$, and prune sampling searches using the approximate ranks $\hat{\pi}$, we obtain the following.

THEOREM 5.2. *With very high probability, for all $s \geq 1$, the influence $\widehat{\text{Inf}}$ of the first s seed nodes is at least $1 - (1 - 1/s)^s - \epsilon$ times the maximum possible $\widehat{\text{Inf}}$ with s seeds. When all nodes are both rankers and rankees, the algorithm uses $O(|E|\epsilon^{-1} \log^3 n + |E|\epsilon^{-2} \log n)$ operations.*

5.3 Approximability of the exact problem

Distance-based influence maximization is known to be at least as hard as max cover also in terms of inapproximability, by a seminal

result of Feige [17]. Thus, we know that in a sense Greedy is the best scalable algorithm. What we can say about reverse-rank influence maximization with a threshold kernel T is that it is at least as hard as max cover, when each element can be a member of at most T sets. The problem is NP-hard for $T \geq 2$ (by reduction to max vertex cover), but Feige’s inapproximability result does not apply. This leaves open the possibility that some polynomial-time algorithms have better approximation ratio than Greedy.

When $T = 1$, the influence function is simply the number of reverse nearest neighbors. In this case, the coverage sets of different nodes are disjoint and influence maximization is trivial: The greedy permutation which selects nodes in decreasing order according to number of reverse nearest neighbors is optimal.

When $T = 2$, each node can be covered by at most two other nodes, which is similar to max vertex cover, which is also NP hard, but has a polynomial approximation algorithms that achieves a slightly better approximation ratio than the greedy guarantee of $1 - (1 - 1/s)^s$ [18]. The Linear Programming based algorithm, however, does not scale for large inputs and also does not seem to apply for our general case of $T > 2$.

6. HARDNESS OF EXACT REVERSE-RANK SINGLE SOURCE COMPUTATION

Exact single source reverse-rank computation from a node u will return π_{iu} for all nodes i . Clearly, it can be solved using an APSP computation. We show the following

THEOREM 6.1. *The reverse-rank single source problem has sub-cubic equivalence to APSP.*

We give a reduction from the *Graph Radius Problem*: The radius of a graph G , is defined as the minimum over nodes u of the maximum distance from u to another node

$$R = \min_{u \in V} \max_{v \in V} d_{uv} . \quad (12)$$

The graph radius problem on undirected graphs is known to have subcubic equivalence to APSP [1].

Given a graph $G = (V, E)$, and a length parameter x , we construct a new graph $G_x = (V', E')$ by adding a new node $V' = V \cup \{u\}$ and adding edges from u to all $v \in V$ with length x .

LEMMA 6.1. *Let G be a graph with radius R and consider $x > 0$. If $R > x$ then in G_x , for all nodes $v \in V$, $\pi_{vu} < |V|$. If $R < x$ then there must exist a node $z \in V$ such that $\pi_{zu} = |V|$ in G_x .*

PROOF. Suppose that $R > x$ then by definition of radius, for all nodes z , $\max_{v \in V} d_{zv} \geq R > x$. Therefore node u will not be the farthest from z and we have $\pi_{zu} < |V|$.

Suppose now that $R < x$. let $z \in V$ be such that $R = \max_{v \in V} d_{zv}$. Then for all $v \in V$, $d_{zu} = x > R \geq d_{zv}$ and thus $\pi_{zu} = |V|$. $\square$

From the lemma, we can compute the graph radius R by performing a logarithmic number (in the representation of G) of exact reverse-rank single source computations on graphs the size of G . This concludes the proof of Theorem 6.1.

COROLLARY 6.2. *Exact computation of reverse-rank influence, even with a single seed node $|S| = 1$, uniform β , and a threshold function α , is sub-cubic equivalent to APSP.*

PROOF. We use the same construction and compute influence (centrality) of node u with α being a threshold function with $T = |V| - 1$. $\square$

Algorithm 3: reverse-rank SKIM

Input: Directed graph $G = (V, E)$, ranker nodes $Z \subset V$, rankee nodes $U \subset V$, threshold T , parameter K
Output: Approximate greedy sequence *seedlist*

// Initialization
forall the nodes $u \in V$ **do** $\text{best_seed}[u].\text{rank} \leftarrow \infty$
forall the rankee nodes $v \in U$ **do** $\text{sample_size}[v] \leftarrow 0$
 $\text{inverted_sample} \leftarrow \perp$ // Hash map of ranker nodes to sets of rankee nodes
 $\text{coverage} \leftarrow 0$ // Coverage of current *seedlist*
 $\text{seedlist} \leftarrow \perp$ // Output list
 $F \leftarrow$ random shuffle of the ranker nodes Z

while $|\text{coverage}| < |Z|$ **and** $(\exists \text{unscanned } u \in F \text{ or } \max_{u \in U} \text{sample_size}[u] > 0)$ **do** // select seed
 $x \leftarrow \perp$ // next seed node
while $\exists \text{unscanned } u \in F$ **do** // Build samples
 $u \leftarrow$ next node in shuffled sequence F
if $\text{best_seed}[u].\text{rank} < \infty$ **then** // Node u is covered, skip it
 $\perp$ **Continue**
// Find all rankees v with $\hat{\pi}_{uv} \leq T$
Run a Dijkstra search from u in G , during which
foreach scanned rankee node $v \in U$ **do**
if $\hat{\pi}_{uv} > T$ **then** terminate the search
 $\text{sample_size}[v] \leftarrow \text{sample_size}[v] + 1$
 $\text{inverted_sample}[u] \leftarrow \text{inverted_sample}[u] \cup \{v\}$
if $\text{sample_size}[v] = K$ **then**
 $x \leftarrow v$ // Next seed node
abort sample building loop
if $x = \perp$ **then**
 $x \leftarrow \arg \max_{u \in U} \text{sample_size}[u]$
if $\text{sample_size}[x] = 0$ **then** **abort** main loop

$I_x \leftarrow 0$ // Estimated coverage of x
// Compute I_x and update residual
run pruned reverse-rank single-source search from x in transposed graph G^T , during which
foreach scanned node $v \in V$ with $(\hat{\pi}_{vu}, d_{vu})$ **do**
if $\hat{\pi}_{vu} > T$ **or** $\text{best_seed}[v].\text{rank} < \hat{\pi}_{vu}$ **or**
 $\text{best_seed}[v].\text{rank} = \hat{\pi}_{vu}$ **and** $\text{best_seed}[v].\text{dist} \leq d_{vu}$ **then**
 $\perp$ **prune** at v
if $\text{best_seed}[v] = \infty$ **and** $v \in Z$ **then** // v is a newly covered ranker
 $I_x \leftarrow I_x + 1$; $\text{coverage} \leftarrow \text{coverage} + 1$
forall the nodes w in $\text{inverted_sample}[v]$ **do**
 $\text{sample_size}[w] \leftarrow \text{sample_size}[w] - 1$
 $\text{inverted_sample}[v] \leftarrow \perp$ // Delete
 $\text{best_seed}[v].\text{rank} \leftarrow \hat{\pi}_{vu}$
 $\text{best_seed}[v].\text{dist} \leftarrow d_{vu}$

$\text{seedlist.append}(x, I_x)$

return *seedlist*

7. EXPERIMENTAL EVALUATION

We implemented and evaluated our algorithms for computing ADS, approximate reverse-rank single-source, and influence maximization. Our implementations are in C++ and were compiled using gcc (g++) with full optimization. Our testing machine runs Centos 6.5 and uses Dell PowerEdge R720 server with two Intel Xeon E5-2640 CPUs. Each with 12 cores (2.50GHz, 12×32KiB L1, 6×256KiB L2, and 15MiB L3) and 264GiB of RAM. The disk capacity is 1T. We open sourced our implementation and it is available at https://github.com/eliavb/all_distance_sketch.

Table 1 shows the social graphs used for our evaluation, all taken from the SNAP project [34]. For each graph we list the number of nodes and edges and whether edges are directed. These data sets did not distinguish between edges, so we used uniform lengths of 1. Our implementations, however, are designed to work with general positive edge lengths.

7.1 Sequential ADS computation

Table 1 also lists, for each instance, performance figures (time and memory usage) of our optimized sequential implementation of *Pruned Dijkstras* (Algorithm 1). We list performance for ADS parameter values $k = 16, 64, 128$ (higher k implies larger sketch size and processing and higher estimation quality). The listed times are broken into *load time* – loading the graph into memory data structures, *ADS time* – computing the sketches, and *ests time* – process the ADS sketches to compute the distance to cardinality estimation lists $L(i)$. We can see that ADS computation was the dominant component. Overall, the preprocessing time is of the order of few hours, even on our largest data set. The table also lists the virtual memory usage of the different runs.

7.2 Multithreaded ADS algorithm

We next evaluate our implementation of the multithreaded ADS algorithm (Section 3.1). The evaluation was done by generating 1 to 14 concurrent threads. We used batch size parameter $\mu = 0.1$. The parameter $\mu = 0.1$ was selected since it had the best performance on a test of sweeping μ between 0.05 and 1 and considering 1–14 threads on the slashdot graph. We note that the amount of concurrency provided in the algorithmic design is much larger, but the architecture of our machine, mainly number of cores and shared caches, limited the benefit of using more threads. We show the results for executions with ADS parameter $k = 16$. The time to load the graph into memory and the total virtual memory used did not vary much for the same instance and different numbers of threads. Table 2 lists the load-time and virtual memory numbers for 7 threads. The table also shows the run time on a single thread. Note that it can be slightly larger than our optimized sequential implementation. Figure 2 shows the running times, as a fraction of the running time on a single-thread, as a function of the number of concurrent threads. We note that the number of threads listed is the concurrency generated by our program scheduler – the actual number of cores allocated by the OS was sometimes smaller and we had no control over it. We observe significant benefit of the multithreaded design, in particular for the larger graphs where we obtain up to a factor of 3 speedup, also with respect to the optimized sequential implementation. We note that most of the speedup is obtained using only 2 – 6 threads.

Table 2: Multithreaded ADS computation

instance	load [Sec]	Memory [GiB]	1-thread [Sec]
Facebook	0.13	0.56	0.69
Slashdot	0.57	1.1	10.0
Twitter	23	2.2	120
YouTube	3.9	3.3	157
LiveJournal	36	11	1541

7.3 Reverse-rank single-source computation

Table 3 shows running times of our approximate reverse-rank single source computations, averaged over 1000 source nodes selected uniformly at random. The times listed are net per computation after loading the graph and pre-computed sketches $L(i)$ into

Table 1: Test Instances and preprocessing time (single thread)

instance (un/directed)	#nodes	#edges	load [Sec]	Preprocess ($k = 16$)			Preprocess ($k = 64$)			Preprocess ($k = 128$)		
				ADS [Sec]	mem [GB]	ests [Sec]	ADS [Sec]	mem [GB]	ests [Sec]	ADS [Sec]	mem [GB]	ests [Sec]
Facebook (u)	3,959	84,243	0.3	0.5	0.03	0.043	0.7	0.03	0.140	1.4	0.04	0.341
Slashdot (d)	77,360	905,468	0.5	5.7	0.1	0.276	23	0.22	1.91	42	0.4	5.35
Twitter (d)	456,626	14,855,842	22	88	0.9	2.00	338	1.4	13.5	523	2	39
YouTube (u)	1,134,890	2,987,624	3.3	116	1.4	6.69	404	3.3	41.6	770	5.3	118
LiveJournal (u)	3,997,962	34,681,189	32	1481	6.4	29.7	4,901	18	209	9,555	31.5	642

Figure 2: Multithreaded computation. Speedup (ratio of time to single-thread time) as a function of the number of threads.

memory. We show running times for the different ADS parameter values $k = 16, 64, 128$. For reference, we also show running time for Dijkstra’s algorithm (single-source distances) averaged over the same 1000 source nodes. We can observe that the running times of our reverse-rank single-source computation do not depend on the sketch parameter k and are similar to Dijkstra computations. The table also shows extrapolated running time for APSP computation. The extrapolation was obtained by multiplying the time for a single Dijkstra run by the number of nodes. This is listed for reference, since exact reverse-rank single-source computation is equivalent to an all-pairs computation.

The table also displays the average relative errors (ARE) for each sketch parameter. Since it was not possible to scalably compute the exact reverse-rank values even for a single source, we computed instead the estimation errors on the ranks using the Dijkstra runs: The errors were therefore averaged over all the ranks provided by 1000 different rankers instead of all the ranks “received” by 1000 different rankers. We can see that the ARE, as expected, decreases with the ADS parameter k and are within the theoretical bounds. Note that a fixed set of sketches was computed during preprocessing and used in all subsequent computations. Therefore the estimates on reverse-ranks of different source-destination pairs can be highly dependent.

7.4 Reverse-rank distributions

Our implementation allowed us, for the first time, to view the reverse rank distributions of nodes in a large network. Figure 3 (left) shows the cumulative reverse rank distributions $\bar{\pi}_{js}$ for 4 selected source nodes in the YouTube network. For each node s , we sort the (estimated) values $\bar{\pi}_{js}$ for all nodes j in increasing order. The cumulative distribution plot then shows the value y at each position x . The figure also includes a reference line where for any i there is a node with rank i . The reference line is in a sense

Figure 3: Cumulative distributions on YouTube graph. Left: Reverse-rank Right: Distance.

corresponds to an “average” source node, which gives and receives the same influence.

We can get information on the relative importance of a node in its “locality,” for varying locality ranges, from its reverse-rank distribution. Nodes that are important in their locality would have distributions that at least initially lie well below the reference line. This means that for some i , there are many more than i nodes that rank them below i . Node #2711 and #480 are example influential nodes that remain important across neighborhood scales. Node #368749 has low influence with distribution above the reference line across ranges. Node #3394 has low influence on most ranges except for its immediate neighborhood, where it has average influence, and on the longest scale, when looking at its 7×10^5 and above highest rankers (which is 35% of total nodes), which indicates that it lies closer to the “core” of the network. Note that we plot $\bar{\pi}$, meaning ties are broken “upwards,” which biases towards being above the reference line.

Figure 3 (right) provides, for comparison, the cumulative distance distributions for the same nodes: For each number of hops y , we see the number x of nodes within y hops. The distance distribution captures the expansion rate, but does not quantify well the relative status of a node within its locality: A less influential member of a dense community would have higher expansion than an influential member of a sparser community. As a simplified example think of two nodes A and B with the same degree Δ such that all neighbors of A have degree $\ll \Delta$ and all neighbors of B have degree $\gg \Delta$. In this case we may view A as being influential in its neighborhood whereas B will not be. The reverse-rank distribution will correctly make this distinction whereas the distance-distribution will not.

7.5 Influence maximization

We next evaluate the performance of reverse-rank SKIM, in terms of both the quality of the coverage and the running time. The evaluation used the social graphs listed in Table 1, with uniform edge lengths and all nodes being both rankers and rankees. We used the $\pi \equiv \bar{\pi}$ (3) interpretation of rank. We study dependence on three parameters, T , k , and K : The threshold value T specifies the coverage rate. The ADS sketch parameter k determines the quality of $\hat{\pi}$ as estimates of the true ranks π , and thus, the relation between $\hat{\text{Inf}}$ (10), which we optimize for, and the true Inf. Finally, the sample size parameter K determines the quality of the coverage in terms

Table 3: Reverse-Rank Single Source computation, averaged over 1000 source nodes selected uniformly at random.

instance	reverse-rank single source						Dijkstra	APSP
	$k = 16$		$k = 64$		$k = 128$			
	[Sec]	ARE	[Sec]	ARE	[Sec]	ARE	[Sec]	$\approx$ [hours]
Facebook	0.006	0.11	0.006	0.072	0.006	0.067	0.002	0.002
Slashdot	0.071	0.34	0.077	0.12	0.074	0.076	0.06	1.3
Twitter	0.96	0.29	1.00	0.050	0.88	0.076	0.54	68
YouTube	1.44	0.24	1.51	0.041	1.34	0.023	0.98	309
LiveJournal	13	0.11	16	0.13	16	0.053	5.8	6440

Figure 4: Fractional coverage $\widehat{\text{Inf}}(S)/|V|$ as a function of seed set size $|S|$ for rR -SKIM, for varying T .

of the approximate influence $\widehat{\text{Inf}}(10)$. Larger K mean that we are more likely to select seeds with marginal $\widehat{\text{Inf}}$ influence that is closer to the maximum.

Recall that the computation of an exact greedy sequence, with respect to either $\widehat{\text{Inf}}$ or Inf (Algorithm 2), is $\tilde{O}(T|E|)$, where the $\tilde{O}$ notation suppresses logarithmic factors. The computation of reverse-rank SKIM uses ADS computation of $\tilde{O}(k|E|)$ and additional computation with worst-case bound of $\tilde{O}(K|E|)$. Moreover, note that in actuality, the time is $\tilde{O}(K|E|\rho)$, where ρ is the ratio between the average and maximum (marginal) influence of a node, and with the typical skew we have $\rho \ll 1$. Therefore, the scalability advantage of SKIM to become more significant for larger T .

When evaluating the effect of the sample size K , we fixed the ADS parameter to be $k = 128$ for the four smaller graphs and $k = 64$ for the larger one and used several values of $T \leq 10^4$. We computed the exact greedy selection with respect to $\widehat{\text{Inf}}$, which is obtained by selecting a node with maximum marginal $\widehat{\text{Inf}}$ in each step. This was done on the four smaller graphs. On these graphs, sample size $K = 100$ was almost always within a fraction of a percent of the exact greedy coverage. The one exception was on Slashdot and $T = 10^4$ where the first seed had coverage that is 6% lower than the optimal one and the gap closed up for the first

three seeds. With $K = 500$, the approximate greedy coverage almost exactly matched the exact greedy coverage. For LiveJournal, we only evaluated the coverage for sample size up to $K = 500$, but the performance with $K = 100$ already matched that. We note that these observed errors are much lower than the worst-case guarantees provided in our analysis. The explanation is the skew of the influence distribution is skewed, where the node of maximum marginal influence is well separated from the second maximum, and with very few nodes having influence that is more than a fraction of the maximum.

Our implementation allows us to examine the coverage to seed set size tradeoffs as a function of the threshold T . These tradeoffs provide structural insights on the networks and results are shown in Figure 4. Higher values of T as expected have higher coverage with fewer seeds. We can also see a highly skewed and asymmetric distribution of importance. For example, the LiveJournal graph with nearly four million nodes, there is a single node that 4×10^4 other nodes rank within their top $T = 100$. The first 11 nodes have 1.6×10^5 nodes ranking at least one of them in their top 100. For $T = 1000$, the top seed covers 3×10^5 nodes and the top 12 cover 7.5×10^5 (a quarter of all nodes).

Table 4 lists selected single thread running times for reverse-rank SKIM. Listed times do not include ADS computation (see Table 1), but this preprocessing time was only a fraction of the total. We note that the running time did not significantly depend on ADS size (the parameter k). The parameter k can impact running time only because it can generate longer neighborhood estimate lists $L(i)$. The size of these lists, even with very large k , is below the effective diameter of the graph, which was small in our data sets. The listed times in the table use $k = 128$ for the four smaller graphs and $k = 64$ for the largest one. They correspond to computing the full sequence (until all nodes are covered). Note that the running time can be significantly reduced if we stop when a desired coverage or number of seeds are reached. We can also observe that the running time grows linearly with the sample size K . An interesting observation is that for the largest graphs, the computation is faster for larger values of T – This is because SKIM works with the residual problem, and its size decreases more rapidly for higher values of T . This is in contrast to an exact greedy computation (Algorithm 2), where the running time increases rapidly with T . Selected running times, and the slowdown factor with respect to reverse-rank SKIM (including ADS computation), are provided in Table 5 (For technical reasons, we had to perform this comparison on a different machine). We can see that exact computation is feasible for small values of T , but the running time increases rapidly with T .

Finally, we evaluate the quality of our approximate greedy sequences which were optimized for $\widehat{\text{Inf}}$, in terms of the exact influence objective Inf . To do so, we used a variation of Algorithm 2 to compute the exact influence of the sequence of seeds returned by reverse-rank SKIM. We observed that even for ADS parameter $k = 64$ and $k = 128$, the Inf coverage of the approximate greedy

Table 4: Reverse-rank SKIM running times [Sec]

Instance T:	K = 100			K = 500		
	10 ²	10 ³	10 ⁴	10 ²	10 ³	10 ⁴
Slashdot	10.22	16.35	-	17.37	32.81	-
Twitter	95.4	68.6	53.4	114	98	82
YouTube	151.3	138.6	228.9	289.3	256.6	346.9
LiveJournal	6533	3517	2040	8261	6123	4313

Table 5: Exact greedy running times and speedup factor of reverse-rank SKIM (including sketch computation)

Instance	T	[hours]	≈ factor
YouTube	10 ²	2.81	×13
YouTube	10 ³	9.50	×51
YouTube	10 ⁴	33.96	×135

sequence for $\widehat{\text{Inf}}$ generally almost matched, and never less than 8% of the exact greedy sequence for Inf.

8. CONCLUSION

Rank-based measures were used for decades as an alternative to distance-based measures. Here, we defined and motivated rank-based measures of centrality and influence, which we believe will become important tools in network mining and analysis. We then presented novel highly scalable algorithms for fundamental rank-based computations on graphs, including a Dijkstra-like approximate reverse rank single-source algorithm which facilitates reverse-rank influence computation and reverse-rank greedy influence maximization. We complement our work with hardness results that indicate that exact computation inherently scales poorly.

A contribution of broader interest is our novel multithreaded design for computing all-distance sketches (ADS) which provided the fastest implementation for computing these sketches on multi-core architectures. This design is relevant to many other applications of distance sketches.

Going forward, we plan to extend our reverse-rank IM computation to general decay functions and design a multithreaded implementation. We also hope to use our newly available tools to explore the relation between distance-based and rank-based influence.

Acknowledgements: The authors would like to thank Shiri Chechik and Amos Fiat for discussions and pointers.

9. REFERENCES

- [1] A. Abboud, F. Grandoni, and V. Vassilevska Williams. Subcubic equivalences between graph centrality problems, APSP and diameter. In *SODA*. ACM-SIAM, 2015.
- [2] B. D. Abrahamo, F. Chierichetti, R. Kleinberg, and A. Panconesi. Trace complexity of network inference. In *KDD*, 2013.
- [3] A. Bavelas. A mathematical model for small group structures. *Human Organization*, 7:16–30, 1948.
- [4] F. Bloch and M. O. Jackson. The formation of networks with transfers among players. *Journal of Economic Theory*, 133(1):83–110, 2007.
- [5] P. Boldi, M. Rosa, and S. Vigna. HyperANF: Approximating the neighbourhood function of very large graphs on a budget. In *WWW*, 2011.
- [6] E. Cohen. Size-estimation framework with applications to transitive closure and reachability. *J. Comput. System Sci.*, 55:441–453, 1997.
- [7] E. Cohen. All-distances sketches, revisited: HIP estimators for massive graphs analysis. *TKDE*, 2015.
- [8] E. Cohen. Multi-objective weighted sampling. In *HotWeb*. IEEE, 2015. full version: <http://arxiv.org/abs/1509.07445>.
- [9] E. Cohen, D. Delling, F. Fuchs, A. Goldberg, M. Goldszmidt, and R. Werneck. Scalable similarity estimation in social networks: Closeness, node labels, and random edge lengths. In *COSN*. ACM, 2013.
- [10] E. Cohen, D. Delling, T. Pajor, and R. F. Werneck. Sketch-based influence maximization and computation: Scaling up with guarantees. In *CIKM*. ACM, 2014.
- [11] E. Cohen, D. Delling, T. Pajor, and R. F. Werneck. Timed influence: Computation and maximization. Technical Report cs.SI/1410.6976, arXiv, 2014.
- [12] E. Cohen and H. Kaplan. Spatially-decaying aggregation over a network: Model and algorithms. *J. Comput. System Sci.*, 73:265–288, 2007. Full version of a SIGMOD 2004 paper.
- [13] E. Cohen and H. Kaplan. Summarizing data using bottom-k sketches. In *ACM PODC*, 2007.
- [14] T. M. Cover and P. E. Hart. Nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, 13(1):21–27, 1967.
- [15] N. Du, L. Song, M. Gomez-Rodriguez, and H. Zha. Scalable influence estimation in continuous-time diffusion networks. In *NIPS*. Curran Associates, Inc., 2013.
- [16] R. Fagin, A. Lotem, and M. Naor. Optimal aggregation algorithms for middleware. *J. Comput. System Sci.*, 66, 2003.
- [17] U. Feige. A threshold of $\ln n$ for approximating set cover. *J. Assoc. Comput. Mach.*, 45:634–652, 1998.
- [18] U. Feige and M. Langberg. Approximation algorithms for maximization problems arising in graph partitioning. *J. Algorithms*, 41, 2001.
- [19] L. C. Freeman. Centrality in social networks: Conceptual clarification. *Social Networks*, 1, 1979.
- [20] J. Gao, H. V. Jagadish, B. C. Ooi, and S. Wang. Selective hashing: Closing the gap between radius search and k-nn search. In *ACM KDD*, 2015.
- [21] M. Gomez-Rodriguez, D. Balduzzi, and B. Schölkopf. Uncovering the temporal dynamics of diffusion networks. In *ICML*, 2011.
- [22] A. Goyal, W. Lu, and L.V.S. Lakshmanan. Celf++: Optimizing the greedy algorithm for influence maximization in social networks. In *WWW*. ACM, 2011.
- [23] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning*. Springer New York Inc., 2001.
- [24] D. G. Horvitz and D. J. Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47(260):663–685, 1952.
- [25] D. Kempe, J. M. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *KDD*. ACM, 2003.
- [26] F. Korn and S. Muthukrishnan. Influence sets based on reverse nearest neighbor queries. In *SIGMOD*, 2000.
- [27] Yiu. M. L., D. Papadias, N. Mamoulis, and Y. Tao. Reverse nearest neighbors in large graphs. *IEEE TKDE*, 18, 2006.
- [28] J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, J. VanBriesen, and Glance N. Cost-effective outbreak detection in networks. In *KDD*. ACM, 2007.
- [29] D. Liben-Nowell and J. Kleinberg. The link prediction problem for social networks. In *CIKM*. ACM, 2003.
- [30] G. Nemhauser, L. Wolsey, and M. Fisher. An analysis of the approximations of maximizing submodular set functions. *Mathematical Programming*, 14, 1978.
- [31] T. Opsahl, F. Agneessens, and J. Skvoretz. Node centrality in weighted networks: Generalizing degree and shortest paths. *Social Networks*, 32, 2010. <http://toreopsahl.com/2010/03/20/>.
- [32] C. R. Palmer, P. B. Gibbons, and C. Faloutsos. ANF: A fast and scalable tool for data mining in massive graphs. In *KDD*, 2004.
- [33] G. Sabidussi. The centrality index of a graph. *Psychometrika*, 31(4):581–603, 1966.
- [34] Stanford network analysis. <http://snap.stanford.edu>.
- [35] Y. Tang, X. Xiao, and Y. Shi. Influence maximization: Near-optimal time complexity meets practical efficiency. In *SIGMOD*, 2014.
- [36] V. Vassilevska Williams and R. Williams. Subcubic equivalences between path, matrix and triangle problems. In *FOCS*. IEEE, 2010.
- [37] J. Weston, S. Bengio, and N. Usunier. WSABIE: scaling up to large vocabulary image annotation. In *IJCAI*, 2011.