The Solution Distribution of Influence Maximization

A High-level Experimental Study on Three Algorithmic Approaches

Naoto Ohsaka

NEC Corporation

naoto.ohsaka@gmail.com

ABSTRACT

Influence maximization is among the most fundamental algorithmic problems in social influence analysis. Over the last decade, a great effort has been devoted to developing efficient algorithms for influence maximization, so that identifying the “best” algorithm has become a demanding task. In SIGMOD’17, Arora, Galhotra, and Ranu reported benchmark results on eleven existing algorithms and demonstrated that there is no single state-of-the-art offering the best trade-off between computational efficiency and solution quality.

In this paper, we report a high-level experimental study on three well-established algorithmic approaches for influence maximization, referred to as Oneshot, Snapshot, and Reverse Influence Sampling (RIS). Different from Arora et al., our experimental methodology is so designed that we examine the *distribution* of random solutions, characterize the relation between the *sample number* and the actual solution quality, and avoid *implementation dependencies*. Our main findings are as follows: **1.** For a sufficiently large sample number, we obtain a unique solution regardless of algorithms. **2.** The average solution quality of Oneshot, Snapshot, and RIS improves at the same rate up to scaling of sample number. **3.** Oneshot requires more samples than Snapshot, and Snapshot requires fewer but larger samples than RIS. We discuss the time efficiency when *conditioning* Oneshot, Snapshot, and RIS to be of identical accuracy. Our conclusion is that Oneshot is suitable only if the size of available memory is limited, and RIS is more efficient than Snapshot for large networks; Snapshot is preferable for small, low-probability networks.

ACM Reference Format:

Naoto Ohsaka. 2020. The Solution Distribution of Influence Maximization: A High-level Experimental Study on Three Algorithmic Approaches. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data (SIGMOD’20)*, June 14–19, 2020, Portland, OR, USA. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3318464.3380564>

1 INTRODUCTION

Social influence among individuals plays an immense role in decision making and information acquisition, and the rise of online social networks has empowered it to spread out at a tremendous scale. Understanding, predicting, and controlling social influence and its diffusion have become a big field of research called computational social influence [11]. Among the most actively studied algorithmic problems in this field is the *influence maximization* problem [35, 36], initially motivated by viral marketing [21]. Conceptually, influence maximization involves identifying a small number of *seed* individuals in the network who can maximize the spread of influence. Kempe, Kleinberg, and Tardos [35] in 2003 formulated influence maximization as a combinatorial optimization problem on graphs, and their framework has been broadly accepted in the research community as well as database [17, 19, 23, 24, 30, 34, 60, 63, 65, 68–70].

One striking topic is the development of efficient algorithms. Under two well-established diffusion models called independent cascade [25] and linear threshold [28], Kempe et al. [35] proved that it is NP-hard to find the optimal solution, but the objective function referred to as the *influence spread* enjoys an excellent property called submodularity. A natural *Greedy algorithm* thus guarantees a $(1 - 1/e)$ -approximation [54]. However, the sheer size of today’s real-world networks and the stochastic nature of the diffusion process make it more challenging to execute the Greedy algorithm. Consequently, this topic has been an active research area for the past ten-odd years (see, e.g., [3, 12, 43]).

1.1 Quick Review of Existing Approaches

Our focus in this paper is on the empirical behavior of existing algorithms for influence maximization. Let us quickly review them (see Section 3 for details). The Greedy algorithm suffers from the intractability of evaluating the influence

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGMOD’20, June 14–19, 2020, Portland, OR, USA
© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6735-6/20/06...\$15.00

<https://doi.org/10.1145/3318464.3380564>

spread, which is defined as the expectation over exponentially many realizations. The requirement is thus an efficient scheme for approximating the influence spread.

The basic idea for this requirement is to construct an (unbiased) *estimator* of the influence spread. The representatives based on this idea can be classified into three approaches, namely, *Oneshot*, *Snapshot*, and *Reverse Influence Sampling (RIS)*. Each of them is parameterized by a single parameter called the *sample number* that trades between the computational complexity and the solution quality.

Up to now, there have been two different directions for realizing “efficient” algorithms. On the one hand, RIS-type algorithms aim to determine the least sample number required to guarantee the worst-case approximation factor theoretically. On the other hand, Oneshot- and Snapshot-type algorithms aim to run as fast as possible for a given sample number. Even though comparing such algorithms designed for distinct purposes is quite complicated, RIS-type algorithms have been regarded as the state-of-the-art supported by the near-linear time complexity “in theory” [30, 60, 68–70]. Note that there exist numerous heuristics that provide influence estimates quickly, but they often result in poorly influential solutions.

1.2 Benchmarking Study in 2017

Arora, Galhotra, and Ranu [2] published a paper titled “*Debunking the Myths of Influence Maximization: An In-Depth Benchmarking Study*” at SIGMOD 2017. The authors exposed a pitfall in comparing influence maximization algorithms by experiments through an exhaustive benchmarking study on eleven existing algorithms [15–17, 24, 26, 27, 31, 44, 62, 69, 70]. The results indicated that 1. the algorithmic efficiency is sensitive to the choice of problem instances, and 2. *no single* state-of-the-art achieves the best trade-off among computation time, memory consumption, and solution quality. Unfortunately, Arora et al.’s experimental methodology still contains several flaws as pointed out by Lu, Xiao, Goyal, Huang, and Lakshmanan [49]. For example, Arora et al. [2] used for each algorithm, a fixed parameter value determined based on preliminary experimental results, which does not match the research focus of RIS-type algorithms [49, Sect. 3.1]. This study complements previous studies from a different aspect.

1.3 Our Motivations

In this paper, we present a high-level experimental study on existing algorithmic approaches for influence maximization. Rather than attempting to choose the best one among them, we would like to clarify their potential applicability. Our key objective under this purpose is to elucidate

the empirical impact of the sample number on the solution distribution for each algorithmic approach

as driven by the following three facts.

1. **Existing algorithms are randomized.** Since an influence estimator is randomized, each algorithm run generates random solutions as well. Despite this nature, most of the previous studies conducted few-trials experiments only, e.g., the number of trials is 3 in [70], 5 in [69], 10 in [13, 15, 47], 20 in [30], 50 in [16], and not explicitly stated in [14, 17, 19, 24, 26, 27, 31, 38, 39, 56, 57, 60–62]; conclusions based on them would be questionable. In this paper, we analyze the empirical distribution of random solutions made from 1,000 trials to gain a deeper understanding of the stochastic behavior of randomized algorithms.
2. **Sample number controls the actual solution quality.** Selecting an appropriate sample number is crucial, but this has been actively studied only for RIS [7, 30, 56, 57, 60, 61, 68–70], and thus, we have no simple way to compare RIS-type algorithms with the other two types. Indeed, such *worst-case* lower bounds are too loose to explain the empirical success fully. In this paper, we run algorithm implementations for a wide range of sample numbers, from 1 to up to 16 million, to discover the relation between the sample number and the actual influence. We can, for example, find the minimum sample number required to obtain near-optimal solutions with high probability.
3. **There is a plethora of algorithm implementations.** We aim to evaluate the algorithmic efficiency, as well; however, a complete experimental comparison of existing implementations is hard for two reasons. First, one can neither completely understand nor modify complex source codes published by many different research groups. For example, on Arora et al. [2]’s setup, SimPath [27] got stuck in an infinite loop due to its different scheme for handling graph data [49, Sect. 3.2.2]. Second, we have many metrics for evaluating scalability, e.g., CPU time and RAM usage, which severely depend on implementations and machine configurations. In this paper, to avoid these implementation dependencies, we make use of simple implementations that capture the essence of each approach. We measure the number of vertices and edges traversed (*traversal cost*) and those stored in memory (*sample size*). Remark that the former (resp. latter) is proportional to running time (resp. memory usage), where the proportionality constant depends on testing environments.

1.4 Our Findings

Our empirical findings are summarized as follows (see Section 5 for more details).

- **Distribution of solutions:** We first reveal how solution distributions converge to what distribution. We find that Oneshot, Snapshot, and RIS return the unique solution for a sufficiently large sample number. We further find that *the Shannon entropy of solution distributions of Oneshot,*

Snapshot, and RIS drops at the same rate up to scaling of sample number.

- **Distribution of influence spread:** We then analyze the empirical distribution of influence spread. The minimum sample number required to obtain near-optimal solutions with probability 99% takes a wide range of values; e.g., from 64 to 8,192 for Oneshot. Those empirical numbers are far smaller than worst-case bounds that depend on the graph size and seed size. Comparing among the three approaches, *the mean value improves at the same rate up to scaling of sample size.* To achieve the same mean influence, *Oneshot requires up to 96 times as many samples as Snapshot requires, Snapshot requires “fewer” (say, 10^5 times fewer) but “larger” (say, 10^3 times larger) samples than RIS;* i.e., RIS is more space-saving than Snapshot.
- **Computational efficiency:** We finally report the traversal cost. We find that the presence of many high-probability edges causes expensive graph traversal due to the emergence of a *giant component* [5, 33, 64]. The *per-sample traversal-cost ratio among Oneshot, Snapshot, and RIS is that $1 : \frac{\tilde{m}}{m} : \frac{1}{n}$* , where n is the number of vertices, m is the number of edges, and $\tilde{m}$ is the sum of all edge probabilities representing the magnitude of influence, showing that RIS is the most per-sample time-efficient.

Section 6 further discusses the traversal cost when *conditioning Oneshot, Snapshot, and RIS to be of identical accuracy.* Our conclusion is that **1.** Oneshot is suitable only if the size of available memory is limited, and **2.** RIS is more efficient than Snapshot for large complex networks; Snapshot is preferable for small, low-probability networks.

Organization. Section 2 introduces formal definitions and known results of influence maximization. Section 3 is devoted to a systematic survey of existing algorithms. Section 4 designs our experimental methodology. Sections 5 and 6 report and discuss experimental results, respectively. Section 7 lists future directions of this study.

2 INFLUENCE MAXIMIZATION

2.1 Notations

For a positive integer ℓ , let $[\ell]$ denote the set $\{1, 2, \dots, \ell\}$. We deal with two types of graph. One is a (deterministic) graph $G = (V, E)$, where V is a set of vertices and E is a set of edges. The other is an *influence graph* $\mathcal{G} = (V, E, p)$, which captures the stochastic nature of network diffusion, where V is the vertex set, E is the edge set, and $p : E \rightarrow (0, 1]$ is an *influence probability* function representing the magnitude of influence between a pair of vertices. For a vertex v in G , we use $\Gamma_G^+(v)$ and $\Gamma_G^-(v)$ to denote the out-neighbors and the in-neighbors of v , respectively, and we use $d_G^+(v)$ and $d_G^-(v)$ to denote the out-degree and the in-degree of v , respectively.

For a vertex set S , let $r_G(S)$ denote the number of vertices reachable from S in G . We will omit the subscripts when G is clear from the context and use the same notations for influence graph $\mathcal{G}$. We conclude this paragraph by defining influence maximization [35], where the definition of the influence spread is deferred.

PROBLEM 2.1 (INFLUENCE MAXIMIZATION). *Given an influence graph $\mathcal{G} = (V, E, p)$ and a seed size k , the influence maximization problem is to find a seed set $S \subseteq V$ of size k that maximizes the influence spread of S .*

2.2 Diffusion Model

Network diffusion models describe the process by which influence (e.g., information, contamination, and virus) triggered by a set of seed vertices spreads over the network. In this paper, we adopt one of the well-studied models in the literature of influence maximization. The *independent cascade (IC)* model introduced by [25] mimics the spread of infectious diseases. In the IC model, each vertex takes either of the two states: *active* or *inactive*. An inactive vertex may become active but not vice versa. Let us define A_t as the set of newly activated vertices at discrete time step t , and let $A_{\leq t} = \bigcup_{t' \leq t} A_{t'}$. For a *seed* set $S \subseteq V$, the vertices in S are initially activated at time step $t = 0$ and the others are inactive, i.e., $A_0 = S$. Given A_t at each time step t , we construct A_{t+1} as follows. Each newly activated vertex $u \in A_t$ is given a single chance to influence its inactive out-neighbors $v \notin A_{\leq t}$, which succeeds with probability $p(u, v)$. If this is the case, v becomes active at time step $t + 1$, i.e., v is added into A_{t+1} . This repetition terminates within a finite time step (at most n), and $A_{\leq n}$ is the set including all activated vertices.

The *influence spread* of a seed set S in $\mathcal{G}$, denoted $\text{Inf}_{\mathcal{G}}(S)$, is defined as the expected number of activated vertices by initially activating seed vertices in S , i.e., $\text{Inf}_{\mathcal{G}}(S) = \mathbb{E}[|A_{\leq n}|]$. Since $\text{Inf}_{\mathcal{G}}(\cdot)$ can be viewed as a function on a subset of vertices, we call $\text{Inf}_{\mathcal{G}} : 2^V \rightarrow \mathbb{R}_{\geq 0}$ an *influence function*.

Here, we describe the *random-graph interpretation* [35] that characterizes the IC model. For an influence graph $\mathcal{G} = (V, E, p)$, consider the distribution over deterministic graphs (V, E') , where E' is obtained from E by maintaining each edge e with probability $p(e)$. We use $G \sim \mathcal{G}$ to mean that G is a *random graph* sampled from this distribution. Then, the influence spread of a seed set S in $\mathcal{G}$ is equal to the expected number of vertices reachable from S in $G \sim \mathcal{G}$, i.e., $\text{Inf}_{\mathcal{G}}(S) = \mathbb{E}_{G \sim \mathcal{G}}[r_G(S)]$. This fact tells us that we do not have to consider the chronological order of activation trials, but we need to consider reachability on random graphs.

2.3 Intractability and Approximability

Let us recall the intractability. Formally, it has been proven that for the IC model, influence maximization is NP-hard

Table 1: Three popular algorithmic approaches for influence maximization.

approach	sample number (typical value)	sample size = (# vertices) + (# edges)	exp. traversal cost (at $k = 1$) vertex	edge	time complexity (naive impl.)
Oneshot	# simulations β (10^4)	–	$\beta \sum_v \text{Inf}_{\mathcal{G}}(v)$	$\leq \beta m \max_v \text{Inf}_{\mathcal{G}^T}(v)$	$O(\beta knm)$
Snapshot	# random graphs τ (10^2)	$\tau \tilde{m}$ in exp. ($\tilde{m} := \sum_e p(e)$)	$\tau \sum_v \text{Inf}_{\mathcal{G}}(v)$	$\leq \tau m \max_v \text{Inf}_{\mathcal{G}^T}(v)$	$O(\tau knm)$
RIS	# RR sets θ (n/a)	θ EPT in exp. (EPT := $\frac{1}{n} \sum_v \text{Inf}_{\mathcal{G}}(v)$)	$\theta \frac{1}{n} \sum_v \text{Inf}_{\mathcal{G}}(v)$	$\leq \theta \frac{m}{n} \max_v \text{Inf}_{\mathcal{G}}(v)$	$O(\theta \frac{m}{n} \max_v \text{Inf}_{\mathcal{G}}(v))$

Table 2: Representative existing algorithms.

approach	representatives
Oneshot	CELF [44], CELF++ [26], UBLF [73, 74], SIEA [58, 59].
Snapshot	Bond Percolation [39, 40], NewGreedy [14], MixedGreedy [14], StaticGreedy [17], PMC [62], SKIM [19].
RIS	RIS [7] TIM+ [70], IMM [69], LISA [20, 56], BCT [55, 61], SSA [60], SSA-Fix [30], WebGraph framework [65], OPIM [68].

to solve exactly [35, Theorem 2.4], and it is even #P-hard to compute the influence spread exactly [13, Theorem 1].

Albeit the negative results, we can obtain approximation. The striking result of Kempe et al. [35, Theorem 2.2] is that the influence function is monotone and submodular. Here, a set function $f : 2^V \rightarrow \mathbb{R}$ is said to be *monotone* if it holds that $f(S) \leq f(T)$ whenever $S \subseteq T \subseteq V$ and *submodular* if it holds that $f(S + v) - f(S) \geq f(T + v) - f(T)$ whenever $S \subseteq T \subseteq V, v \in V \setminus T$. The classical result on submodular functions by Nemhauser, Wolsey, and Fisher [54] then tells us that the simple Greedy, which iteratively adds an element that makes the maximum marginal increase in function value, produces a $(1 - 1/e)$ -approximate solution, i.e., it holds that $f(S) \geq (1 - 1/e)f(S^*)$, where S is the k -seed Greedy solution and S^* is the optimal solution of size k .

Putting it all together, we have that Greedy achieves a constant-factor approximation in polynomial time. The influence spread can be approximated within a factor of $(1 \pm \epsilon)$ for any $\epsilon > 0$ by running simulations $\Omega(\epsilon^{-2}n^2)$ times [37].

3 ALGORITHMS REVIEW

In this section, we review existing algorithms for influence maximization systematically. Based on the mechanism of influence estimation, we partition existing algorithms into three approaches shown in Table 1, namely, *Oneshot* (Section 3.3), *Snapshot* (Section 3.4), *Reverse Influence Sampling (RIS)* (Section 3.5), and others (Section 3.6).

3.1 Kempe et al.’s Greedy Algorithm

Before taking up each approach, we explain the Greedy algorithm of [35]. It begins with an empty solution and iteratively adds an element that makes the largest marginal increase in influence into the solution until k elements have been added. Remark that the influence function is evaluated for at most nk seed sets. Due to monotonicity and submodularity of the influence function, the resulting solution is a $(1 - 1/e)$ -approximation [54]; this factor is the best possible [22]. By contrast, it often provides near-optimal (e.g., 0.95 times the optimum) solutions empirically [41, 44, 67].

Algorithm 3.1 Simple greedy framework.

```

1: call Build( $\mathcal{G} = (V, E, p)$ , sample number) and  $S_0 \leftarrow \emptyset$ .
2: randomly shuffle the order of vertices in  $V$ .
3: for all  $\ell = 1, 2, \dots, k$  do
4:   call Estimate( $S_{\ell-1}, v$ ) for all  $v \in V$ .
5:    $v_\ell \leftarrow$  last vertex with maximum (marginal) influence.
6:   call Update( $v_\ell$ ) and  $S_\ell \leftarrow S_{\ell-1} + v_\ell$ .
7: return  $S_k$ .

```

The limitation of Kempe et al.’s Greedy algorithm is that we are unable to evaluate $\text{Inf}(\cdot)$ exactly. Fortunately, even if we are given only an approximate value oracle, which approximates the actual value within a factor of $(1 \pm \epsilon)$, running Greedy on this achieves a $(1 - 1/e - O(k\epsilon))$ -approximation [29]. Monte-Carlo simulations can be employed for this purpose; however, running them for nk seed sets is computationally prohibitive, e.g., it takes a few days on graphs with ten thousand vertices [38]. Subsequent studies thus seek efficient and accurate methods for estimating the influence spread.

3.2 A Simple Greedy Framework

Here, we introduce a simple greedy framework to describe existing algorithms in a unified manner. Most of the existing algorithms fall into this framework, and Table 2 lists up representative algorithms that belong to either of the three approaches. Our greedy framework shown in Algorithm 3.1 requires the following three procedures. Hereafter, we denote the seed vertex chosen at the ℓ -th iteration by v_ℓ .

- *Build*($\mathcal{G}$, *sample number*) builds an estimator for the influence function given influence graph $\mathcal{G}$ with an approach-specific parameter, called “*sample number*” in this paper.
- *Estimate*($S_{\ell-1}, v$) returns an estimate for the marginal influence of v with respect to $S_{\ell-1}$ (i.e., $\text{Inf}(S_{\ell-1} + v) - \text{Inf}(S_{\ell-1})$), or the influence spread of $S_{\ell-1} + v$. The results will be the same regardless.
- *Update*(v_ℓ) updates the current estimator given the next seed v_ℓ to be able to estimate the marginal influence with respect to the latest seed set $S_\ell := S_{\ell-1} + v_\ell$.

In what follows, we clarify the concept and features of each algorithmic approach and explain how to implement *Build*, *Estimate*, and *Update* in a simple manner. We then analyze the computational efficiency. Besides the asymptotic complexity, we measure the computational effort in terms of how many times vertices and edges were touched. Concretely speaking, the *vertex (resp. edge) traversal cost* is defined as the number of vertices (resp. edges) examined (possibly more than once), and the *sample size* is defined as the number of

Algorithm 3.2 Naive implementation of Oneshot.

```

1: procedure Build( $\mathcal{G} = (V, E, p), \beta$ )                                ▶ do nothing.
2: procedure Estimate( $S_{\ell-1}, v$ )
3:   for  $\beta$  times do
4:     simulate the diffusion model from  $S_{\ell-1} + v$  in Sect. 3.3.
5:     count # activated vertices.
6:   return total # activated vertices divided by  $\beta$ .
7: procedure Update( $v_\ell$ )                                              ▶ do nothing.

```

vertices and edges stored as approach-specific samples in memory. Traversal cost and sample size are implementation independent and approximately proportional to the running time and memory usage, respectively. Further, we present a brief survey on existing “efficient” implementations.

There are two types of “efficient.” One means that the sample number is *as small as possible* while satisfying a specified precision requirement, which demands a good criterion for sample number determination. The other means that for a given sample number (which may be determined based on the above-mentioned criterion), the implementation runs *as fast as possible*, which demands the complexity reduction. We note that Oneshot- and Snapshot-type algorithms have focused on the latter kind of efficiency while RIS-type algorithms have focused on the former kind of efficiency as will be clarified. This discrepancy would be the reason why previous experimental comparisons led to questionable claims.

3.3 Oneshot Algorithms

3.3.1 Concept. *Oneshot-type algorithms* (a.k.a. simulation-based [48]) execute Monte-Carlo simulations of the diffusion process many times on the spot when estimates are needed. Algorithm 3.2 shows pseudocode of Oneshot, of which sample number β specifies the *number of simulations* to be performed. In *Estimate*, given a seed set, we simulate the diffusion process β times and return the average number of activated vertices. This estimate is unbiased. Nothing is done in *Build* and *Update*. One concern is that neither submodularity nor monotonicity is guaranteed [17] because *Estimate*’s return values are independent of each other.

3.3.2 Computational Complexity. We analyze the computational complexity of Oneshot. Each call of *Estimate* for a seed set S examines $\text{Inf}(S)$ vertices in expectation and their outgoing edges. In particular, the vertex and edge traversal cost at $k = 1$ is $\beta \sum_v \text{Inf}(v)$ and $\beta m \max_v \text{Inf}_{\mathcal{G}^T}(v)$, respectively, whose proof is deferred to Appendix. Note that the entire computational complexity is bounded by $O(\beta kmn)$. In *Estimate*, we store $|A_{\leq n}| \leq n$ vertices, which is negligible.

3.3.3 Efficient Implementations.

Estimate Call Pruning. We then review existing implementation techniques. Since simulating network diffusion is computationally expensive, reduction of *Estimate* calls is a promising direction. A convenient technique uses *upper bounds* for

Algorithm 3.3 Naive implementation of Snapshot.

```

1: procedure Build( $\mathcal{G} = (V, E, p), \tau$ )
2:   generate  $\tau$  random graphs  $G^{(1)}, \dots, G^{(\tau)}$  from  $\mathcal{G}$ .
3: procedure Estimate( $S_{\ell-1}, v$ )
4:   return  $\frac{1}{\tau} \sum_{i \in [\tau]} r_{G^{(i)}}(S_{\ell-1} + v) - r_{G^{(i)}}(S_{\ell-1})$ .
5: procedure Update( $v_\ell$ )                                              ▶ do nothing.

```

the marginal influence to identify vertices that would never be selected as a seed. Such upper bounds are quickly derived without simulations by utilizing submodularity [26, 44] and linear systems [73, 74].

Sample Number Determination. Another direction is to determine an appropriate sample number β given a precision requirement. Theoretically, setting $\beta = \Omega(\epsilon^{-2} k^2 n (\log \delta^{-1} + \log k) / \text{OPT}_k)$ for any $\epsilon > 0$, where $\text{OPT}_k = \max_{S: |S|=k} \text{Inf}(S)$, achieves a $(1 - 1/e - \epsilon)$ -approximation with probability $1 - \delta$ [70, Lemma 10]; however, none of the existing algorithms adopt such bounds due to the inefficiency of estimating OPT_k . In contrast, tens of thousands of simulations are sufficient to obtain reasonable solutions in practice [17, 35].

3.4 Snapshot Algorithms

3.4.1 Concept. *Snapshot-type algorithms* generate random graphs from an influence graph in advance and share them over the entire greedy seed selection, which makes room for a bunch of speed-up techniques. Algorithm 3.3 shows pseudocode of Snapshot algorithms, of which sample number τ specifies the *number of random graphs* to be generated. In *Build*, we independently sample τ random graphs, denoted $G^{(1)}, G^{(2)}, \dots, G^{(\tau)}$, from $\mathcal{G}$. An unbiased estimator of the influence spread of seed set $S \subseteq V$ is then defined as $\frac{1}{\tau} \sum_{i \in [\tau]} r_{G^{(i)}}(S)$. Hence, in *Estimate*, we take the size of reachable sets and return the average. *Update* does nothing by default. Unlike Oneshot, this estimator enjoys both monotonicity and submodularity since $G^{(i)}$ ’s are fixed [17].

3.4.2 Computational Complexity. *Build* touches each edge only τ times, which does not dominate the whole time complexity. In *Estimate*, for each random graph, we scan vertices reachable from the seed set and their outgoing edges (e.g., by a breadth-first search). Hence, the expected vertex (resp. edge) traversal cost at $k = 1$ is up to $\tau \sum_v \text{Inf}(v)$ (resp. $\tau m \max_v \text{Inf}_{\mathcal{G}^T}(v)$), whose proof is similar to that of Oneshot. Remark that the entire running time is bounded by $O(\tau kmn)$. The sample size is expected to be $\tau \tilde{m}$, where $\tilde{m} = \sum_e p(e)$ is the expected number of edges in $G \sim \mathcal{G}$.

3.4.3 Efficient Implementations.

Fast Reachability Computation. Besides applying the reduction technique for Oneshot, *quick reachability computation* in *Estimate* seems helpful. Here, we explain a graph reduction technique in [39, 40, 62]. Consider the situation

Algorithm 3.4 Naive implementation of RIS.

```

1: procedure Build( $\mathcal{G} = (V, E, p)$ ,  $\theta$ )
2:    $\mathcal{R} \leftarrow \emptyset$ .
3:   for  $\theta$  times do
4:      $R \leftarrow$  an RR set under  $\mathcal{M}$  and  $\mathcal{R} \leftarrow \mathcal{R} + R$ .
5:   procedure Estimate( $S_{\ell-1}, v$ )
6:     return  $n \cdot F_{\mathcal{R}}(v)$ .
7:   procedure Update( $v_{\ell}$ )
8:     remove RR sets including  $v_{\ell}$  from  $\mathcal{R}$ .

```

where we are about to choose the second seed. In $Update(v_1)$, we construct the subgraph $H^{(i)}$ obtained from $G^{(i)}$ by deleting vertices reachable from the first seed v_1 and their incident edges. Then, it follows that $r_{G^{(i)}}(\{v_1, v\}) - r_{G^{(i)}}(v_1) = r_{H^{(i)}}(v)$ for any vertex v . Therefore, it suffices to traverse on smaller $H^{(i)}$. Repeatedly applying this reduces the traversal cost in the subsequent iterations without disturbing estimates.

Meanwhile, finding the most influential seed (i.e., the iteration at $\ell = 1$) is still laborious. In the first iteration, we need to compute the number of vertices reachable from every vertex. This computation problem is precisely the *descendant counting problem* [18] and unsolvable in truly-subquadratic time, assuming the strong exponential time hypothesis [6]. Previous work used fast approximation or heuristics, e.g., reachability sketches [14, 18], bottom- k min-hash sketches [19], and pruned breadth-first searches [62].

Sample Number Reduction. Since the sample number τ governs not only the computation time but also the memory consumption unlike Oneshot, it is more crucial to determine the value of τ appropriately. Generating $\tau = \Omega(\frac{n^2}{\epsilon^2}(k \log n + \log \delta^{-1}))$ random graphs suffices to ensure that the seed set obtained has influence at least $(1 - 1/e) \text{OPT}_k - \epsilon$ with probability at least $1 - \delta$ [32, Prop. 3]. Perhaps surprisingly, it was empirically observed [17, 39, 40] that Snapshot requires fewer random graphs (e.g., $\tau = 100$) than the number of simulations required by Oneshot (e.g., $\beta = 10,000$), which might be unpredictable from their principles. One reason is that Snapshot’s estimator preserves monotonicity and submodularity while Oneshot’s estimator does not so. However, to the best of our knowledge, no clear theoretical justification for such empirical reduction in sample number has been made.

3.5 Reverse Influence Sampling

3.5.1 Concept. Reverse Influence Sampling (RIS) (a.k.a. sketch-based [48]) pioneered by Borgs, Brautbar, Chayes, and Lucier [7], is the first near-linear time algorithm (for fixed k) for influence maximization. The elegant insight is the reduction of influence maximization to stochastic maximum coverage, which is efficiently approximable. The critical notion of serving as a bridge between them is *reverse reachable sets*.

DEFINITION 3.1 (REVERSE REACHABLE SET [7]). For a target $z \in V$, a reverse reachable (RR) set for z under IC is defined

as a set of vertices that can reach z in random $G \sim \mathcal{G}$. An RR set for a random target, denoted R , is referred to as an RR set.

Then, we have that $\Pr_R[R \cap S \neq \emptyset] = \frac{\text{Inf}(S)}{n}$ for any vertex set $S \subseteq V$ [7, Observation 3.2], which intuitively means that influential vertices frequently appear in RR sets.

Influence maximization is therefore equivalent to a maximum coverage problem on exponentially many RR sets. To solve it approximately, RIS polls a few RR sets and runs Greedy on them to return an approximate solution. Let $\mathcal{R}$ be a collection of RR sets, and we define $F_{\mathcal{R}}(S)$ as the fraction of RR sets in $\mathcal{R}$ intersecting S , i.e., $F_{\mathcal{R}}(S) = \frac{|\{R \in \mathcal{R} | R \cap S \neq \emptyset\}|}{|\mathcal{R}|}$. We then have $n \cdot F_{\mathcal{R}}(S)$ an unbiased estimate for $\text{Inf}(S)$.

Algorithm 3.4 shows pseudocode of RIS, which is characterized by the number θ of RR sets to be drawn. *Build* constructs a collection $\mathcal{R}$ of θ RR sets. For IC, an RR set can be found efficiently by a reverse simulation [7, 69]. *Estimate* computes the marginal coverage of vertex v as $F_{\mathcal{R}}(v)$. *Update* given the new seed v_{ℓ} removes RR sets including v_{ℓ} from $\mathcal{R}$ so that $F_{\mathcal{R}}(v)$ is the marginal coverage with respect to S_{ℓ} .

3.5.2 Computational Complexity. Remark that the number of steps taken by generating an RR set is the sum of in-degrees of vertices in the set [7]. The *weight* of an RR set R , denoted $w(R)$, is defined as $\sum_{v \in R} d^-(v)$. By definition, an RR set contains vertex v with probability $\text{Inf}(v)/n$, and it turns out that the expected size of an RR set is $\text{EPT} := \mathbb{E}[|R|] = \sum_v \text{Inf}(v)/n$ (corresp. to the vertex traversal cost). The expected number of steps taken by generating an RR set is $\mathbb{E}[w(R)]$, which is bounded from above by $\frac{m}{n} \max_v \text{Inf}(v)$ (corresp. to the edge traversal cost). Accordingly, *Build* finishes in $O(\theta \frac{m}{n} \max_v \text{Inf}(v))$ expected time, and so does the entire algorithm,¹ and the sample size is bounded by θEPT . Here, it holds that $\text{EPT} \leq 1 + \bar{m}$, i.e., the sample size of RIS is less than that of Snapshot if $\theta = \tau$ up to addition by one, of which proof is deferred to Appendix.

3.5.3 Efficient Implementations.

Sample Number Determination. Unlike the case of Oneshot and Snapshot, most of the research on RIS focus on a proper selection of sample number θ , or equivalently, a stopping condition for RR-set generation. The standard requirement is to draw as few RR sets as possible that yield a “theoretical worst-case guarantee” on a $(1 - 1/e - \epsilon)$ -approximation with probability $1 - \delta$ [7, 30, 56, 57, 60, 61, 68–70].

Borgs et al. [7, 8] first showed that terminating RR-set generation when the total weight exceeds $O(\epsilon^{-2} k(m + n) \log n)$, which implies that $\theta \approx O(\frac{\epsilon^{-2} k n \log n}{\text{OPT}_k})$ (this factor is k times smaller than Oneshot’s upper bound [70]), we have the desired guarantee. The whole process thus can be done in almost linear time, which is runtime-optimal up to a

¹See [7, Theorem 3.1] for fast implementations of *Estimate* and *Update*.

logarithmic factor. However, a hidden constant sorely limits the usability against large graphs, and subsequent research has aimed at designing tight stopping criteria such as based on optimal influence bounds [69, 70], degree-based thresholding [56, 61], and search-and-verify [30, 47, 60]. It is worth mentioning that these algorithms “have roughly the same running time when their empirical accuracies are the same” [49, Sect. 3.1] (wavy underline was added by the author). This is not the case for Oneshot and Snapshot.

Complexity Reduction. A few reduction techniques exist, e.g., graph compression [65] and importance sampling [57].

3.6 Other Approaches

We review two approaches not tested in this paper.

Binary Decision Diagrams for Exact Computation. Exact computation of the influence spread has been studied though it consumes exponential time. The current fastest exact algorithm is the one designed by Maehara, Suzuki, and Ishihata [51]. The idea is to make use of the *binary decision diagram* [10] for computing the two-terminal network reliability, which is $\#P$ -hard [71]. Maehara et al.’s algorithm was successful with computing the exact influence on graphs with up to a hundred edges.

Heuristics for Quick Guesses. To avoid expensive sampling procedures, there have been developed cheap heuristics at the price of estimation accuracy. One popular approach is to make a strong assumption on the process of network diffusion under which the influence spread can be computed quickly, e.g., influence follows along with shortest paths [38] or trees [13], and the zone of influence is bounded by neighbors [14]. Others employ linear systems [16, 31] and graph reduction [52, 64, 66]. Such heuristics are faster than the three approaches, but resulting seed sets have less influence.

4 DESIGN

In this section, we design our experimental methodology. We recapitulate two features of the algorithms reviewed so far. First, Oneshot, Snapshot, and RIS run Greedy on the respective influence estimator to yield random solutions. The issue is that examining the results of a few trials may not tell much about mean, tail, and diversity accurately. Our strategy is thus to construct the *distribution* of seed sets and discover its structural properties. Second, whereas the sample number governs the influence spread, we have no way of choosing a sample number of Oneshot and Snapshot according to a given precision requirement, and worst-case approximation factors are less than 0.64, which is too loose to explain the empirical success adequately. Our strategy is thus to discover the relation between the sample size and the *actual* influence spread rather than worst-case approximation factors.

To realize these, we run algorithm alg (e.g., Oneshot) with sample size s (e.g., β) T times, and we then record each obtained seed set and its influence spread to empirically construct the seed set distribution $S^{(s)}$ and the influence distribution $I^{(s)}$. We can, for example, understand the limit behavior by verifying the convergence of $S^{(s)}$ for a large s , understand the tail behavior by analyzing the first percentile of $I^{(s)}$, and compare Oneshot, Snapshot, and RIS, assuming that the sample size was ideally chosen.

In the following, we contemplate the selection of 1. algorithm implementations, 2. network data, and 3. edge probabilities to grasp the behavioral trends of each approach.

4.1 Algorithm Implementations

Since Oneshot, Snapshot, and RIS use random numbers, we here clarify where to invoke what type of pseudorandom number generator (PRNG). For each algorithm run, we use different seed values to initialize the state of a PRNG to obtain randomized solutions. The implementations are written in C++ and compiled using g++ v4.8.2 with the -O2 option. We used the Mersenne Twister [53], a general-purpose PRNG, to draw random numbers.

- Oneshot: In *Estimate*, we invoke a PRNG to generate a random $x \in [0, 1]$ for each examined edge e and think of e as alive if $x < p(e)$.
- Snapshot: In *Build*, we invoke a PRNG to generate a random $x \in [0, 1]$ for each edge e and each $i \in [r]$, and we keep e in the i -th random graph $G^{(i)}$ if $x < p(e)$.
- RIS: We use two kinds of PRNG, where one chooses a vertex in V randomly, and the other generates a random real number between $[0, 1]$. During a single RR-set generation in *Build*, we first invoke the first kind of PRNG to generate a random target $v \in V$, and we then invoke the second kind of PRNG to generate a random $x \in [0, 1]$ for each examined edge e and think of e as alive if $x < p(e)$.

During greedy iterations in Algorithm 3.1, two or more “tie” vertices may have the same maximum estimate. To avoid dependency on the initially given order of vertices, we shuffle the order randomly beforehand so that simply running through vertices in this order breaks any ties randomly.

4.2 Network Data

Table 3 summarizes the basic statistics of network data used in this paper. Our selection includes not only real-world networks but also synthetic networks, some of which are simple and small unlike those adopted in previous studies so that we can run algorithms with a huge sample number.

4.2.1 Real-world Networks. We use six social networks. They are called *complex networks* sharing common structural features that occur in neither random graphs nor grid graphs.

Table 3: Network statistics. Δ^+ and Δ^- denote maximum out- and in-degree.

network	n	m	type	Δ^+	Δ^-	clus. coef.	avg. dis.
<i>real-world networks</i>							
Karate [42]	34	156	social	17	17	0.26	2.41
Physicians [42]	241	1,098	social	9	26	0.25	–
ca-GrQc [45]	5,242	28,968	collab.	81	81	0.63	–
Wiki-Vote [45]	7,115	103,689	voting	893	457	0.13	–
com-Youtube	1,134,889	5,975,248	social	28,754	28,754	–	–
soc-Pokec	1,632,802	30,622,564	social	8,763	13,733	–	–
<i>synthetic networks</i>							
BA_s	1,000	999	BA	20	23	0.00	7.22
BA_d	1,000	10,879	BA	100	107	0.06	2.50

Scale-free property [4]: The distribution of vertex degrees follows power-law, i.e., the fraction of vertices having k neighbors is $\propto k^{-\gamma}$, where $\gamma \in [2, 3]$ typically.

Small-world property [72]: The expected distance between two vertices chosen randomly is small, typically $O(\log n)$.

Cluster property [72]: Complex networks have a high clustering coefficient, which is defined as three times the number of triangles over the number of connected triplets.

Core-whisker structure [9, 46, 50]: Complex networks can be decomposed generally into two parts; the expander-like dense “core” part and the tree-like “whisker” part.

4.2.2 Synthetic Networks. There exist plenty of generative models of random graphs aimed at resembling real networks’ properties. We adopt the Barabási-Albert model [1], which generates scale-free undirected graphs by a preferential attachment process, in which every time a new vertex has been added, it is randomly connected to M existing vertices. In this paper, we generate a sparse version BA_s with $n = 1,000$, $M = 1$ and a dense version BA_d with $n = 1,000$, $M = 11$, and we assigned random directions for each edge.

4.3 Edge Probabilities

We finally designate edge probability settings. Since publicly-available network data do not usually include influence probabilities, we assign them artificially. Our policy consists of the following well-established four strategies.

- **Uniform cascade** (denoted $uc_{0.1}$ (resp. $uc_{0.01}$)): Each edge probability is a constant 0.1 (resp. 0.01).
- **In-degree weighted cascade** (denoted irc): The influence probability of edge (u, v) is set to $1/d^-(v)$. It turns out that $\sum_{u \in \Gamma^-(v)} p(u, v) = 1$.
- **Out-degree weighted cascade** (denoted owc): The influence probability of edge (u, v) is set to $1/d^+(u)$. It turns out that $\sum_{v \in \Gamma^+(u)} p(u, v) = 1$.

Intuitively, each vertex owns equalized influence on its neighbors on owc while higher-degree vertices have more chances to activate their neighbors on uc .

5 RESULTS

In this section, we present the results of the experiments designed previously. We conducted experiments on a Linux server with an Intel Xeon E5-2670 2.60GHz CPU and 500GB memory. We constructed the empirical distribution of seed sets from $T = 1,000$ trials for small instances and from $T = 20$ trials for large instances, which are marked with “★”. Unless otherwise specified, the seed size k was set to 1, 4, 16, 64, or 1,024, and the sample number of Oneshot (β), Snapshot (τ), and RIS (θ) was set to a power of two up to $2^{16} = 65,536$, $2^{16} = 65,536$, and $2^{24} = 16,777,216$, respectively,² which results in the convergence of seed set distributions as shown in Section 5.1. Our research focus is on the following.

- **Empirical distribution of seed sets** (Section 5.1): Since an influence estimator approaches to the exact influence function as the sample number increases, the seed set distribution is expected to settle down eventually. We will investigate how fast or slow the seed set distribution loses the diversity and verify its convergence for a variety of instances. We will also compare the diversity decay speed among Oneshot, Snapshot, and RIS.
- **Empirical distribution of influence spread** (Section 5.2): Similarly to the case of seed set distributions, it is expected that influence distributions converge eventually. We will observe how influence distributions concentrate, and assess the minimum sample number required to obtain near-optimal seed sets. We will further compare influence distributions among Oneshot, Snapshot, and RIS. For this purpose, we will discuss how to evaluate the *quality* of influence distributions.
- **Empirical traversal cost** (Section 5.3): Rather than merely measuring running time, which is dependent on the machine configuration, we employ the traversal cost. We will identify what kind of instance incurs high traversal cost. From Table 1, the traversal-cost ratio of Oneshot, Snapshot, and RIS when $\beta = \tau = \theta$ is that $1 : 1 : \frac{1}{n}$ in theory; we will verify whether this is the case empirically.

5.1 Distribution of Seed Sets

We first examine the distribution of seed sets. Our task is to discover the speed at which the seed set distribution settles down. To measure the diversity of a given distribution over sets, we use the *Shannon entropy*, $H = -\sum_S p_S \log_2 p_S$, where p_S is the probability mass of vertex set S . If the distribution takes only a single set, it is called *degenerate* and has entropy 0. Since the empirical distribution is constructed from 10^3 trials, its entropy never exceeds $\log_2 10^3 \approx 9.97$.

5.1.1 Overall Tendencies. We found that basically, the entropy in the early stages is nearly maximum, and it then

²Some of the results are missing when it took over weeks.

Figure 1: Entropy of the seed set distributions on Karate ($UC_{0,1}$). The entropy converged to 0 when $k = 1, 4$.

Figure 2: Two instances for which the entropy hits a plateau, which means that there exist two almost-the-same-influence seed sets.

monotonically decreases. Figure 1 shows the change in entropy as the sample number increases on Karate ($UC_{0,1}$, $k = 1, 4, 16$). In Figure 1a, where the maximum possible entropy is $\log_2 \binom{34}{1} \approx 5.09$, Oneshot reduced the entropy from 4.74 to 0, Snapshot reduced the entropy from 4.92 to 0, and RIS reduced the entropy from 5.06 to 0. On instances including Karate ($UC_{0,1}$, $UC_{0,01}$, IWC, owc, $k = 1$), Physicians ($UC_{0,01}$, owc, $k = 1$), Wiki-Vote ($UC_{0,01}$, IWC, $k = 1$), BA_s ($UC_{0,1}$, $UC_{0,01}$, IWC, owc, $k = 1$), and BA_d ($UC_{0,01}$, IWC, $k = 1$), BA_s (IWC, $k = 16$), ★ com-Youtube (IWC, $k = 1, 16$), and ★ soc-Pokec (IWC, $k = 1, 4$), the entropy for some algorithms eventually converged to 0. We verified in those cases that the resulting seed sets are *unique* regardless of the choice of algorithm and sample number. Therefore, the three algorithms have the *same limit behavior*. Figures 1a and 1b indicate that for a fixed instance, the entropy of algorithm alg at sample number s can be represented by $h(\alpha_{\text{alg}} \cdot s)$ for some decreasing function h , where α_{alg} is a scaling parameter. That is, *the entropy of Oneshot, Snapshot, and RIS drops at the same rate up to scaling*.

5.1.2 When is the Convergence Slow (or Fast)? We explain several cases where the entropy slowly decreased or did not even converge to 0 (as far as we observed). The main reason is that there are several ties so that we would need more

Figure 3: Difference of the entropy decay speed by edge probability settings.

Table 4: Top three influence spread of a single vertex. Highlighted digits match to those of $\lnf(v^{1st})$.

BA_s (Figure 3a)	$UC_{0,1}$	$UC_{0,01}$	IWC	owc
$\lnf(v^{1st})$	3.2089	1.1901	21.4167	5.1365
$\lnf(v^{2nd})$	2.9921	1.1739	20.5095	5.0286
$\lnf(v^{3rd})$	2.9779	1.1604	20.3244	5.0272
BA_d (Figure 3b)	$UC_{0,1}$	$UC_{0,01}$	IWC	owc
$\lnf(v^{1st})$	377.0686	2.1710	101.7954	15.5098
$\lnf(v^{2nd})$	377.0483	2.1162	100.1006	15.5031
$\lnf(v^{3rd})$	375.7611	2.1066	96.1872	15.4971

samples to distinguish them from the limit. On Karate (IWC, $k = 4$) and Physicians (IWC, $k = 1$), shown in Figure 2, the entropy stays around 1 for a long time, though eventually escapes from there. This means that there are two almost-the-same-influence seed sets, and the tie-breaking rule of Algorithm 3.1 selects either of them with almost equal probability. In fact, the influence estimate for the two ties is 21.444 and 21.446 on Karate (IWC, $k = 4$) and 12.403 and 12.412 on Physicians (IWC, $k = 1$), which are nearly indistinguishable. We also observed in Figure 1 that the larger the seed size k , the higher the entropy, which is consistent with that the solution space is of size $\binom{n}{k}$.

Here, we investigate the impact of edge probability settings on the entropy decay speed. Figure 3 shows the result for the two Barabási-Albert networks when $k = 1$. IWC shows the lowest entropy at any sample number on both networks while $UC_{0,01}$ and owc show the highest entropy on BA_s and

Figure 4: Influence distribution in notched box plot on Physicians ($UC_{0.1}$, $k = 16$).

BA_d, respectively. Table 4 reports the top three influence spread of a single vertex for each instance. One can see that *the large difference between the maximum influence and the second maximum is critical for quick convergence of entropy*. Such differences in convergence speed tell us that experimental evaluation using a single setting of edge probabilities is not sufficient (e.g., *only* rwc was tested in [30, 56, 60, 61, 68–70]), even though real data is unavailable.

5.2 Distribution of Influence Spread

Let us now go into the distribution of influence spread. Our task is to reveal how fast the influence distribution concentrates, and when the three algorithms yield the-same-quality influence distributions. Since the exact computation of the influence spread is intractable, we used the following approximation scheme. For each influence graph $\mathcal{G}$, we generated 10^7 RR sets, denoted $\mathcal{R}_{\mathcal{G}}$. The unbiased estimator was then defined as $n \cdot F_{\mathcal{R}_{\mathcal{G}}}(\cdot)$. This estimator was reused over different runs of different algorithms to ensure that multiple identical seed sets have a unique estimate. Note that the event that “an RR set intersects S ” is a Bernoulli trial with success probability $\frac{\text{Inf}(S)}{n}$, and thus the 99% confidence interval for the actual influence spread is given by $n \cdot F_{\mathcal{R}_{\mathcal{G}}}(\cdot) \pm \frac{1.29}{\sqrt{10^7}} n$.

5.2.1 Overall Tendencies. We visualize influence distributions. Figure 4 shows notched box plots of influence distributions on Physicians ($UC_{0.1}$, $k = 16$).

See the right figure for explanation of notched box plots. Both mean and median are consistently increasing and (are expected to) eventually converge to the unique influence, which would be identical to Greedy on the exact influence.

We then assess the sample number needed to obtain “near-optimal” seed sets “almost certainly.” Specifically, we regard the seed set uniquely obtained when $H = 0$ as *Exact Greedy* and a seed set of influence at least 0.95 times Exact Greedy as *near-optimal*. Table 5 reports the least sample number (denoted β^* , τ^* , θ^*) and the respective entropy (denoted H^*)

Table 5: Least sample number and corresponding entropy required to obtain near-optimal seed sets with probability 99% for each algorithm. Blues and reds are for minimum and maximum, respectively.

network	prob.	k	Oneshot		Snapshot		RIS	
			$\log_2 \beta^*$	H^*	$\log_2 \tau^*$	H^*	$\log_2 \theta^*$	H^*
Karate	$UC_{0.1}$	1	8	0.98	7	1.03	12	0.92
Karate	$UC_{0.1}$	4	9	2.91	7	2.35	12	2.94
Karate	$UC_{0.01}$	1	8	1.14	8	1.13	16	1.14
Karate	$UC_{0.01}$	4	7	7.10	7	3.43	15	5.46
Karate	rwc	1	10	0.06	10	0.04	14	0.02
Karate	owc	1	11	0.35	11	0.24	14	0.46
Karate	owc	4	10	4.49	8	5.57	12	5.67
Physicians	$UC_{0.01}$	1	7	5.95	7	5.90	20	5.15
Physicians	rwc	4	10	2.15	8	1.80	13	2.67
Physicians	owc	1	13	0.23	13	0.14	17	0.45
Wiki-Vote	$UC_{0.01}$	1	7	0.92	7	0.99	18	0.83
Wiki-Vote	$UC_{0.01}$	4	7	1.79	6	1.26	17	1.26
Wiki-Vote	rwc	1	7	0.25	7	0.22	17	0.09
Wiki-Vote	rwc	4	7	0.87	5	0.77	15	0.56
Wiki-Vote	$UC_{0.01}$	1	8	0.62	8	0.61	19	0.51
Wiki-Vote	$UC_{0.01}$	4	8	0.78	6	1.26	17	1.26
Wiki-Vote	rwc	1	9	0.01	8	0.07	18	0.01
Wiki-Vote	rwc	4	8	0.37	6	0.23	15	0.56
BA_s	$UC_{0.1}$	1	10	0.03	10	0.02	20	0.02
BA_s	$UC_{0.01}$	1	8	1.32	9	0.97	> 20	–
BA_s	rwc	1	9	0.08	10	0.01	18	0.14
BA_s	rwc	16	6	8.18	4	5.10	13	7.52
BA_s	owc	1	9	0.07	9	0.08	19	1.51
BA_d	$UC_{0.01}$	1	8	1.35	8	1.42	18	1.47
BA_d	rwc	1	11	0.78	11	0.74	14	0.83

for which each algorithm was able to obtain a near-optimal seed set with probability at least 99% over 10^3 trials. The required sample number heavily depends on the problem instance, e.g., β^* varies from 64 to 8,192 and τ^* varies from 16 to 8,192, and it is therefore mandatory to appropriately select the sample number of Oneshot and Snapshot although the previous studies use a fixed sample number. Remark also that the entropy does not have to be low, e.g., $H^* > 4$ on Karate (owc, $k = 4$). We here show a large gap between worst-case upper bounds and empirical least sample numbers: On Wiki-Vote ($UC_{0.01}$, $k = 4$) (resp. BA_s (rwc, $k = 16$)), the bound for Oneshot [70] with $\epsilon = 0.05$, $\delta = 0.01$ is $1.0 \cdot 10^8$ (resp. $5.8 \cdot 10^7$) while the empirical number is 256 (resp. 64), and the bound for RIS [70] with $\epsilon = 0.05$, $\delta = 0.01$ is $1.6 \cdot 10^7$ (resp. $1.4 \cdot 10^6$) and the empirical number is 131,072 (resp. 8,192).

Figure 5: Influence distribution in notched box plot for RIS on ca-GrQc ($k = 1$).

5.2.2 Different Tendencies by Instances. We discuss the difference by problem instances. Reconsider Table 5 to examine the relation between the seed size k and the required sample number. In a theoretical analysis on stochastic submodular maximization [8, 32], the sample number linearly scales in k , but this is not the case for our instances. In particular, on Karate ($UC_{0.01}$, owc) and BA_s (IWC), the required sample number decreases. One possible reason is that there are numerous near-optimal seed sets; a seed set that we extracted from them is reasonably influential even when the entropy is *high*. In fact, an RIS-type algorithm required fewer RR sets when $k > 1$ than when $k = 1$, e.g., see [69, Figure 4d].

We then depict contrasting results from ca-GrQc ($UC_{0.1}$, $k = 1$) and ca-GrQc (owc, $k = 1$) in Figure 5. In the former case, the mean influence starts with less than 20% of the maximum and then quickly improves. The 1st and 99th percentiles at $\theta = 2^{20}$ are 208.797 and 208.900, respectively. In the latter case, it slowly increases although the starting point is better than half of the maximum, and the 1st and 99th percentiles at $\theta = 2^{20}$ are 6.795 and 7.429, respectively. The core-whisker structure can best explain this. In the random graph counterpart of ca-GrQc ($UC_{0.1}$), a certain portion of the densely-connected core part forms a *giant component* [5, 33, 64], and the tree-like whisker part becomes entirely disconnected. Hence, there exist the two extremes of very influential vertices and poorly influential vertices, and it becomes easier to identify some core vertices as soon as the sample number grows. On the other hand, in the random graph counterpart of ca-GrQc (owc), all vertices have exactly one outgoing edge in expectation, and so are similarly influential; however, it is difficult to distinguish the best influential vertex from a number of *slightly less* influential ones.

5.2.3 Comparison among Algorithms. We present the comparison among Oneshot, Snapshot, and RIS. We first discuss how to measure the *quality* of influence distributions. One might think that we need to evaluate several statistics such as mean, standard deviation, and percentiles simultaneously. Our important finding is that for each instance, the *mean*

Figure 6: Relation between mean value and other statistics of influence distributions, where solid lines are for Physicians (owc, $k = 4$) and dashed lines are for Physicians ($UC_{0.1}$, $k = 16$).

can be a dominant factor. In Figure 6a (resp. 6b), we demonstrate that the relation between mean and standard deviation (resp. first percentile) is almost independent of the choice among Oneshot, Snapshot, and RIS. We thus declare that influence distribution I_1 is *better than* influence distribution I_2 if the mean of I_1 is greater than that of I_2 .

We then determine what sample number and sample size make two distributions comparable with each other as follows. Fix an instance. For each $i \in \{1, 2\}$, let alg_i be one of Oneshot, Snapshot, and RIS, s_i denote the sample number of alg_i , and $I_i^{(s_i)}$ be the distribution of influence spread obtained by running alg_i with s_i . Then, s_2 is said to be *comparable* to s_1 if “ s_2 is the least sample number such that $I_2^{(s_2)}$ is better than $I_1^{(s_1)}$ ”, s_2/s_1 is called the *comparable number ratio* of alg_2 to alg_1 , and $(\text{sample size of } alg_2 \text{ at } s_2)/(\text{sample size of } alg_1 \text{ at } s_1)$ is called the *comparable size ratio* of alg_2 to alg_1 . A large “number” ratio means that alg_2 requires more samples than alg_1 while a large “size” ratio means that alg_2 requires larger samples in size than alg_1 . Note that s_2/s_1 takes a power of two since so do s_1 and s_2 .

Oneshot versus Snapshot. We first investigate the comparable number ratio of Oneshot to Snapshot as shown in Figure 7.³ The comparable ratio almost always lies within the range from 1 to 32; i.e., Snapshot requires fewer samples than Oneshot. Besides, the ratio is stable with Snapshot’s sample number τ , that is, *the mean influence of Oneshot and Snapshot improves at the same rate up to scaling of sample number*. We thus report the median of the comparable number ratio in Table 6, which increases to up to 96 as the seed size k increases and further study might be needed.

Snapshot versus RIS. We next investigate both the comparable size ratio and comparable number ratio of RIS to Snapshot shown in Figure 8 and Table 7. *The comparable size ratio is stable with Snapshot’s sample size $\tau \cdot \bar{m}$, and this is*

³Since the sample size of Oneshot is 0, the size ratio is not analyzed.

Figure 7: Comparable number ratio of Oneshot to Snapshot.

Table 6: Median of comparable number ratio of Oneshot to Snapshot.

network	k	$uc_{0.1}$	$uc_{0.01}$	rwc	owc
Karate	1	2	1	2	2
Karate	4	8	6	8	8
Karate	16	16	8	8	12
Physicians	1	2	1	2	2
Physicians	4	4	3	6	4
Physicians	16	8	8	16	16
★ Physicians	64	32	16	32	32
ca-GrQc	1	1	2	2	1
ca-GrQc	4	1	4	8	4
★ ca-GrQc	64	1	64	32	96
Wiki-Vote	1	–	2	2	1
Wiki-Vote	4	–	4	8	2
BA_s	1	2	2	1	2
BA_s	4	4	4	4	4
BA_s	16	8	8	8	8
BA_d	1	2	2	2	4
BA_d	4	2	4	8	16
BA_d	16	–	8	24	–

the case for the comparable number ratio. The comparable number ratio varies from $2^2 = 4$ to up to $2^{19} = 524,288$, and it is less dependent on seed size k unlike the comparison between Oneshot and Snapshot.

Comparable number ratios over 4,096 are observed on Physicians ($uc_{0.01}$), ca-GrQc ($uc_{0.01}$), com-Youtube ($uc_{0.01}$, rwc, owc), soc-Pokec ($uc_{0.01}$, owc), and BA_s ($uc_{0.01}$). The primary reason is that the influence spread on these instances is tiny and so is an RR set; the upper bound on RIS's sample number θ [8] is inversely proportional to the maximum influence (divided by n). We emphasize that *just because Snapshot requires much fewer samples than RIS does not mean that Snapshot is more efficient than RIS*. Comparable size ratios have a quite different trend, which are significantly less than 1, e.g., 0.00033 on com-Youtube (rwc, $k = 1$) and 0.016 on soc-Pokec (rwc, $k = 1, 4$). We thus conclude that when RIS is comparable to Snapshot, *Snapshot requires fewer but larger samples than RIS*, or equivalently, *RIS is more space-saving than Snapshot* for large networks.

5.3 Graph Traversal Cost per Sample

We finally examine the per-sample traversal cost of naive implementations. Table 8 reports the average traversal cost corresponding to the case of $k = 1$ and $\beta = \tau = \theta = 1$.

Figure 8: Comparable size ratio of RIS to Snapshot.

Table 7: Median of comparable ratio of RIS to Snapshot. **Number ratios** > 4,096 and **size ratios** < 0.1 are highlighted.

network	k	number ratio θ/τ				size ratio $(\theta \text{EPT})/(\tau \tilde{m})$			
		$uc_{0.1}$	$uc_{0.01}$	rwc	owc	$uc_{0.1}$	$uc_{0.01}$	rwc	owc
Karate	1	32	384	16	16	4	258	2	2
Karate	4	32	512	16	32	4	344	2	3
Karate	16	128	1,024	64	64	16	689	7	7
Physicians	1	256	8,192	32	64	4	782	0.62	1
Physicians	4	256	8,192	64	64	4	782	1	1
Physicians	16	256	8,192	128	128	4	782	2	2
★ Physicians	64	512	8,192	256	256	8	782	5	5
ca-GrQc	1	32	8,192	256	768	0.13	30	0.19	0.57
ca-GrQc	4	128	12,288	512	1,024	0.53	45	0.34	0.76
★ ca-GrQc	64	2,048	65,536	1,024	2,048	9	241	0.76	2
★ ca-GrQc	1,024	16,384	98,304	4,096	8,192	68	362	3	6
Wiki-Vote	1	–	2,048	1,024	512	–	2	0.75	0.31
Wiki-Vote	4	–	2,048	1,024	768	–	2	0.75	0.46
★ Wiki-Vote	64	16,384	8,192	4,096	2,048	607	10	3	1
★ Wiki-Vote	1,024	49,152	32,768	32,768	32,768	1,820	40	24	20
★ com-Youtube	1	–	128	64	65,536	–	0.13	0.00033	0.34
★ com-Youtube	4	–	512	1,536	65,536	–	0.53	0.0080	0.34
★ com-Youtube	16	–	32,768	2,048	65,536	–	34	0.011	0.34
★ com-Youtube	64	–	65,536	4,096	65,536	–	68	0.021	0.34
★ com-Youtube	1,024	–	393,216	32,768	524,288	–	409	0.17	3
★ soc-Pokec	1	–	32,768	1,536	2,048	–	0.16	0.016	0.021
★ soc-Pokec	4	–	24,576	1,536	1,536	–	0.12	0.016	0.016
★ soc-Pokec	16	–	32,768	3,072	4,096	–	0.16	0.033	0.042
★ soc-Pokec	64	–	81,920	4,096	8,192	–	0.41	0.043	0.085
★ soc-Pokec	1,024	–	262,144	16,384	49,152	–	1	0.17	0.51
BA_s	1	1,536	32,768	512	1,024	17	3,314	2	4
BA_s	4	2,048	32,768	512	2,048	23	3,314	2	8
BA_s	16	4,096	65,536	1,024	2,048	46	6,627	4	8
BA_d	1	4	2,048	16	24	0.54	21	0.24	0.36
BA_d	4	8	2,048	32	48	1	21	0.48	0.72
BA_d	16	512	4,096	128	256	69	43	2	4

5.3.1 Overall Tendencies. The edge traversal cost is higher than the vertex traversal cost excepting Snapshot on several instances. Generally, the larger the influence graph, the higher the traversal cost. The setting of edge probabilities also has a significant impact; $uc_{0.1}$ often incurs the most expensive cost, e.g., BA_d ($uc_{0.1}$) and ca-GrQc ($uc_{0.1}$). This can be explained by the core-whisker structure: A giant component is likely to appear in the random graph counterpart of BA_d ($uc_{0.1}$) but not of BA_d ($uc_{0.01}$, rwc, and owc), whose expected number of edges $\tilde{m}$ is less than n . Such an influence graph that a giant component is likely to appear will be referred to as a *giant-component influence graph*.

5.3.2 Comparison among Algorithms. Firstly, we compare Oneshot with Snapshot. The vertex traversal cost is the

Table 8: Traversal cost at $k = 1$ and sample number 1 for each algorithm.

network	algorithm	$UC_{0.1}$		$UC_{0.01}$		IWC		OWC	
		vertex	edge	vertex	edge	vertex	edge	vertex	edge
Karate $n = 34$ $m = 156$	Oneshot	66.6	375.3	35.7	168.8	126.2	560.6	126.2	858.9
	Snapshot	66.6	37.5	35.7	1.7	126.2	119.2	126.2	126.2
	RIS	2.0	11.0	1.1	5.0	3.7	25.3	3.7	16.5
Physicians $n = 241$ $m = 1,098$	Oneshot	429.9	2,008.8	252.5	1,153.2	1,020.8	4,636.5	986.6	4,700.4
	Snapshot	429.9	200.9	252.5	11.5	1,020.9	904.3	986.5	900.4
	RIS	1.8	8.3	1.1	4.8	4.2	19.3	4.1	18.4
ca-GrQc $n = 5,242$ $m = 28,968$	Oneshot	63,248.7	1,247,121.3	5,595.5	35,844.5	20,373.7	129,789.9	20,373.2	220,480.0
	Snapshot	63,210.5	124,630.7	5,595.4	358.4	20,374.3	19,625.0	20,377.5	20,377.4
	RIS	12.1	237.9	1.1	6.8	3.9	42.1	3.9	24.8
Wiki-Vote $n = 7,114$ $m = 103,689$	Oneshot	–	–	8,959.0	184,956.1	12,449.6	233,880.9	26,242.1	924,225.6
	Snapshot	–	–	8,959.0	1,849.6	12,449.7	5,365.7	26,243.6	19,385.0
	RIS	–	–	1.3	26.0	1.8	46.3	3.7	46.8
★ com-YouTube $n = 1.1M, m = 6.0M$	Snapshot	–	–	70,630,278.9	77,780,873.5	6,713,554.8	6,133,043.3	6,712,576.5	6,712,316.1
	RIS	–	–	62.2	6,851.5	5.9	2,360.7	5.9	35.0
★ soc-Pokec $n = 1.6M, m = 31M$	Snapshot	–	–	2,481,201.0	855,397.3	26,287,270.5	25,490,537.8	24,272,264.1	23,489,755.3
	RIS	–	–	1.5	52.2	16.1	890.8	14.9	325.1
BA_s $n = 1,000$ $m = 999$	Oneshot	1,131.8	1,318.0	1,010.2	1,026.8	2,276.0	1,844.1	2,233.1	3,987.3
	Snapshot	1,131.8	131.8	1,010.3	10.3	2,276.2	1,276.2	2,233.4	1,233.4
	RIS	1.1	1.3	1.0	1.0	2.3	4.1	2.2	1.8
BA_d $n = 1,000$ $m = 10,879$	Oneshot	146,555.6	2,054,009.7	1,134.1	13,416.0	14,992.8	162,203.9	15,078.3	263,505.4
	Snapshot	146,577.9	205,464.5	1,134.0	134.1	15,014.4	14,924.5	15,055.5	15,053.0
	RIS	146.5	2,055.1	1.1	13.4	15.0	263.4	15.1	163.2

same as expected from Table 1, but the edge traversal cost is different. Specifically, the edge traversal cost of Snapshot is approximately $\tilde{m}/m$ times that of Oneshot, where $\tilde{m}$ is $0.1m$ on $UC_{0.1}$, $0.01m$ on $UC_{0.01}$, and n on IWC and OWC . This factor comes from that Snapshot scans only live edges in *Estimate*.

We then compare Oneshot with RIS. The vertex traversal cost of Oneshot is n times that of RIS as expected from Table 1. The edge traversal cost indicates a similar tendency, e.g., the cost of Oneshot divided by that of RIS is within [240, 255] on Physicians with 241 vertices and [616, 1,615] on BA_d with 1,000 vertices, which is not implied by Table 1. It seems that there is not much difference between an influence graph and its transposed counterpart with regard to the edge traversal cost though $\text{Inf}_G \neq \text{Inf}_{G^T}$. To sum up, *there is a simple edge-traversal-cost relation: Oneshot $\approx m/\tilde{m} \cdot$ Snapshot $\approx n \cdot$ RIS, and RIS is the most per-sample time-efficient.*

5.4 Recap of Findings

So far, we have tested naive implementations of Oneshot, Snapshot, and RIS on eight graph data under four edge probability settings. Our empirical findings are summarized below, and the next section is devoted to further discussions.

5.4.1 Distribution of Seed Sets. Seed set distributions approach to degenerate as the sample number grows, i.e., for a sufficiently large sample number, we obtain a single seed set. This seed set is unique regardless of the choice of algorithms, which means that Oneshot, Snapshot, and RIS have the same limit behavior. The entropy of Oneshot, Snapshot, and RIS drops at the same rate up to scaling of sample number. However, the entropy slowly decreases if several

ties have almost-the-same influence. In particular, the large difference between the maximum influence and the second maximum is critical for quick convergence.

5.4.2 Distribution of Influence Spread. Following the convergence nature of seed set distributions, influence distributions converge to the unique influence. In particular, the mean influence monotonically increases as the sample number increases. Concerning the tail behavior, the minimum sample number required to obtain near-optimal seed sets takes a wide range of values, i.e., $\beta^* \in [64, 8,192]$, $\tau^* \in [16, 8,192]$, and $\theta^* \in [4,096, 1,048,576]$, depending on edge probability settings, seed size, and graph size. Hence, it is mandatory for Oneshot- and Snapshot-type algorithms to appropriately select the sample number. Comparing among the three approaches, the mean influence of Oneshot, Snapshot, and RIS improves at the same rate up to scaling of sample number and sample size. Oneshot requires up to 96 times as many samples as Snapshot requires, Snapshot requires fewer (say, 10^5 times fewer) but larger (say, 10^3 times larger) samples in size than RIS; i.e., RIS is more space-saving than Snapshot.

5.4.3 Traversal Cost. High traversal costs are incurred on a giant-component influence graph such as ca-GrQc ($UC_{0.1}$) and BA_d ($UC_{0.1}$) having both high-degree vertices and high-probability edges. The empirically observed relation among Oneshot, Snapshot, and RIS is that $1 : 1 : \frac{1}{n}$ for vertex traversal cost and $1 : \frac{\tilde{m}}{m} : \frac{1}{n}$ for edge traversal cost, demonstrating that RIS offers the best per-sample time-efficiency and Snapshot offers the second best.

6 DISCUSSIONS ON TRAVERSAL COST

We finally discuss the graph traversal cost *conditioning* that the sample number is carefully specified so that Oneshot, Snapshot, and RIS yield seed set distributions of identical accuracy. Let cr_1 and cr_2 be the comparable number ratios of Oneshot and RIS to Snapshot, respectively. Setting $\beta = cr_1 \cdot \gamma$, $\tau = \gamma$, and $\theta = cr_2 \cdot \gamma$ for any γ ensures that the three algorithms yield the same-mean influence distributions. Referring to Tables 6, 7, and 8, we can estimate the traversal cost when this is the case, which are reported in Table 9.

Oneshot is almost always the least time-efficient. Comparing between Snapshot and RIS, we observe that RIS significantly surpasses Snapshot on the two largest networks, com-Youtube and soc-Pokec. On the other hand, Snapshot runs more than twice as fast as RIS on ca-GrQc ($uc_{0.1}$) and BA_s ($uc_{0.1}$, $uc_{0.01}$). On such instances, edge probabilities are too low, or the underlying graph is quite sparse and small, resulting in large comparable ratios (see Table 7).

Summarizing the above, we conclude that *as long as* we use naive implementations, 1. either Snapshot or RIS is preferable; Oneshot can be promising only if the size of available memory is limited, and 2. RIS is more time-efficient than Snapshot for large complex networks; conversely, Snapshot is preferable for small, low-influence-probability networks.

Remarks. Our discussion would need to be adapted for existing efficient implementations, and we should note the following two facts. First, there is a variety of boosted implementations designed for Oneshot and Snapshot. Some of them *significantly reduce the traversal cost without sacrificing approximation guarantee*, e.g., SKIM proposed by Cohen, Delling, Pajor, and Werneck [19] is Snapshot-type and guaranteed to run in near-linear time $O(\tau(n + m) + m\epsilon^{-2} \log^2 n)$ (for some ϵ independent of sample number). Second, none of the existing Oneshot- and Snapshot-type algorithms adopted bounds on the sample number in Sections 3.3 and 3.4. This situation motivates applying (tight) bounds on the sample number to Oneshot and Snapshot.

7 CONCLUDING REMARKS

In this paper, we established an experimental study on three algorithmic approaches for influence maximization. Possible directions for developing of new algorithms are listed below.

- **Sample number selection for Oneshot and Snapshot:** If the sample number is specified so that the three algorithms are of identical accuracy, then Oneshot can be the most space-saving and Snapshot can have less expensive traversal costs than RIS for small graphs. Can we *efficiently* calculate a tight bound on the sample number for Oneshot and Snapshot (in Sections 3.3 and 3.4), or apply RIS's sample number determination to Oneshot and Snapshot?

Table 9: Traversal cost at $k = 1$ conditioning that Oneshot, Snapshot, and RIS are of identical accuracy.

network	algorithm	$uc_{0.1}$	$uc_{0.01}$	IWC	owc
ca-GrQc	Oneshot	1,310,453 γ	82,882 γ	300,318 γ	240,829 γ
ca-GrQc	Snapshot	187,970 γ	5,954γ	39,997 γ	40,744 γ
ca-GrQc	RIS	8,000γ	64,757 γ	11,762γ	22,000γ
Wiki-Vote	Oneshot	–	387,830 γ	492,661 γ	950,468 γ
Wiki-Vote	Snapshot	–	10,809γ	17,815γ	45,629 γ
Wiki-Vote	RIS	–	55,830 γ	49,202 γ	25,859γ
★ com-Youtube	Snapshot	–	148,411,152 γ	12,846,598 γ	13,424,893 γ
★ com-Youtube	RIS	–	884,953γ	151,464γ	2,679,453γ
★ soc-Pokec	Snapshot	–	3,336,598 γ	51,777,808 γ	47,762,019 γ
★ soc-Pokec	RIS	–	1,760,893γ	1,393,038γ	696,292γ
BA_s	Oneshot	4,899 γ	4,074 γ	4,120 γ	12,441 γ
BA_s	Snapshot	1,264γ	1,021γ	3,552 γ	3,466γ
BA_s	RIS	3,762 γ	66,751 γ	3,282γ	4,154 γ
BA_d	Oneshot	4,400,487 γ	29,100 γ	354,374 γ	1,114,474 γ
BA_d	Snapshot	351,923 γ	1,268γ	29,902 γ	30,176 γ
BA_d	RIS	8,806γ	29,798 γ	4,454γ	4,278γ

- **Space reduction for Snapshot and RIS:** Both Snapshot and RIS may consume much memory despite their time efficiency. Can we cut down the memory usage of Snapshot and RIS, e.g., by compressing reverse-reachable sets?

We also give two aspects of experimental evaluation. Only carrying out either of them would not be enough. One is to test the trade-off between scalability and actual influence with *varying sample numbers* as in this paper. One can identify the fastest algorithm on the assumption that the sample number is ideally chosen. Note that RIS-type algorithms would show nearly identical trade-off [49, Sect. 3.1]. The other is to test the scalability of algorithms for a *given precision requirement* (e.g., a $(1 - 1/e - \epsilon)$ -approximation with probability $1 - \delta$). One can identify the most efficient algorithm taking into account the sample number selection. Currently, existing Oneshot- and Snapshot-type algorithms would not be applicable for this test due to the lack of a mechanism for sample number selection.

ACKNOWLEDGMENTS

The author would like to thank anonymous reviewers for their constructive comments and suggestions and was supported by JST ERATO Grant Number JPMJER1201, Japan.

APPENDIX

Traversal cost for Oneshot at $k = 1$. Since a vertex is scanned if it is activated, the vertex traversal cost is equal to $\sum_v \text{Inf}_{\mathcal{G}}(v)$. For each scanned vertex, we touch its out-going edges; we then have that the edge traversal cost is

$$\sum_v \mathbb{E}_G[\sum_{w:v \text{ can reach } w} d_G^+(w)] = \sum_w d^+(w) \cdot \mathbb{E}_G[r_{G^\top}(w)] = \sum_w d^+(w) \cdot \text{Inf}_{\mathcal{G}^\top}(w) \leq m \cdot \max_v \text{Inf}_{\mathcal{G}^\top}(v).$$

Sample size of Snapshot and RIS. By definition, it follows that $\sum_e p(e) = \mathbb{E}_G[|E(G)|]$. Observing that $r_G(v) \leq 1 + |E(G)|$ for all v , we have that $\frac{1}{n} \sum_v \text{Inf}(v) = \frac{1}{n} \sum_v \mathbb{E}_G[r_G(v)] \leq \frac{1}{n} \sum_v \mathbb{E}_G[1 + |E(G)|] = 1 + \sum_e p(e)$.

REFERENCES

- [1] Réka Albert and Albert-László Barabási. 2002. Statistical mechanics of complex networks. *Rev. Mod. Phys.* 74, 1 (2002), 47.
- [2] Akhil Arora, Sainyam Galhotra, and Sayan Ranu. 2017. Debunking the Myths of Influence Maximization: An In-Depth Benchmarking Study. In *SIGMOD*. 651–666.
- [3] Çigdem Aslay, Laks V. S. Lakshmanan, Wei Lu, and Xiaokui Xiao. 2018. Influence Maximization in Online Social Networks. In *WSDM*. 775–776.
- [4] Albert-László Barabási and Réka Albert. 1999. Emergence of scaling in random networks. *Science* 286, 5439 (1999), 509–512.
- [5] Béla Bollobás. 2001. *Random graphs*. Number 73. Cambridge University Press.
- [6] Michele Borassi. 2016. A note on the complexity of computing the number of reachable vertices in a digraph. *Inf. Process. Lett.* 116, 10 (2016), 628–630.
- [7] Christian Borgs, Michael Brautbar, Jennifer Chayes, and Brendan Lucier. 2014. Maximizing Social Influence in Nearly Optimal Time. In *SODA*. 946–957.
- [8] Christian Borgs, Michael Brautbar, Jennifer T. Chayes, and Brendan Lucier. 2016. Maximizing Social Influence in Nearly Optimal Time. *CoRR* abs/1212.0884v5 (2016).
- [9] Andrei Z. Broder, Ravi Kumar, Farzin Maghoul, Prabhakar Raghavan, Sridhar Rajagopalan, Raymie Stata, Andrew Tomkins, and Janet L. Wiener. 2000. Graph structure in the Web. *Comput. Netw.* 33, 1-6 (2000), 309–320.
- [10] Randal E. Bryant. 1986. Graph-Based Algorithms for Boolean Function Manipulation. *IEEE Trans. Comput.* 35, 8 (1986), 677–691.
- [11] Wei Chen. 2015. Computational Social Influence. In *SocInf*. 1–1.
- [12] Wei Chen, Laks V. S. Lakshmanan, and Carlos Castillo. 2013. *Information and Influence Propagation in Social Networks*. Morgan & Claypool Publishers.
- [13] Wei Chen, Chi Wang, and Yajun Wang. 2010. Scalable Influence Maximization for Prevalent Viral Marketing in Large-Scale Social Networks. In *KDD*. 1029–1038.
- [14] Wei Chen, Yajun Wang, and Siyu Yang. 2009. Efficient Influence Maximization in Social Networks. In *KDD*. 199–208.
- [15] Wei Chen, Yifei Yuan, and Li Zhang. 2010. Scalable Influence Maximization in Social Networks under the Linear Threshold Model. In *ICDM*. 88–97.
- [16] Suqi Cheng, Huawei Shen, Junming Huang, Wei Chen, and Xueqi Cheng. 2014. IMRank: Influence Maximization via Finding Self-consistent Ranking. In *SIGIR*. 475–484.
- [17] Suqi Cheng, Huawei Shen, Junming Huang, Guoqing Zhang, and Xueqi Cheng. 2013. StaticGreedy: Solving the Scalability-Accuracy Dilemma in Influence Maximization. In *CIKM*. 509–518.
- [18] Edith Cohen. 1997. Size-estimation Framework with Applications to Transitive Closure and Reachability. *J. Comput. Syst. Sci.* 55, 3 (1997), 441–453.
- [19] Edith Cohen, Daniel Delling, Thomas Pajor, and Renato F Werneck. 2014. Sketch-based Influence Maximization and Computation: Scaling up with Guarantees. In *CIKM*. 629–638.
- [20] Thang Dinh, Hung Nguyen, Preetam Ghosh, and Michael Mayo. 2015. Social Influence Spectrum with Guarantees: Computing More in Less Time. In *CSoNet*. 84–103.
- [21] Pedro Domingos and Matt Richardson. 2001. Mining the Network Value of Customers. In *KDD*. 57–66.
- [22] Uriel Feige. 1998. A Threshold of $\ln n$ for Approximating Set Cover. *J. ACM* 45, 4 (1998), 634–652.
- [23] Shanshan Feng, Gao Cong, Arijit Khan, Xiucheng Li, Yong Liu, and Yeow Meng Chee. 2018. Inf2vec: Latent Representation Model for Social Influence Embedding. In *ICDE*. 941–952.
- [24] Sainyam Galhotra, Akhil Arora, and Shourya Roy. 2016. Holistic Influence Maximization: Combining Scalability and Efficiency with Opinion-Aware Models. In *SIGMOD*. 743–758.
- [25] Jacob Goldenberg, Barak Libai, and Eitan Muller. 2001. Talk of the Network: A Complex Systems Look at the Underlying Process of Word-of-Mouth. *Mark. Lett.* 12, 3 (2001), 211–223.
- [26] Amit Goyal, Wei Lu, and Laks V. S. Lakshmanan. 2011. CELF++: Optimizing the Greedy Algorithm for Influence Maximization in Social Networks. In *WWW*. 47–48.
- [27] Amit Goyal, Wei Lu, and Laks V. S. Lakshmanan. 2011. SIMPATH: An Efficient Algorithm for Influence Maximization under the Linear Threshold Model. In *ICDM*. 211–220.
- [28] Mark Granovetter. 1978. Threshold Models of Collective Behavior. *Am. J. Sociol.* 83, 6 (1978), 1420–1443.
- [29] Thibaut Horel and Yaron Singer. 2016. Maximization of Approximately Submodular Functions. In *NIPS*. 3045–3053.
- [30] Keke Huang, Sibor Wang, Glenn S. Bevilacqua, Xiaokui Xiao, and Laks V. S. Lakshmanan. 2017. Revisiting the Stop-and-Stare Algorithms for Influence Maximization. *Proc. VLDB Endow.* 10, 9 (2017), 913–924.
- [31] Kyomin Jung, Wooram Heo, and Wei Chen. 2012. IRIE: Scalable and Robust Influence Maximization in Social Networks. In *ICDM*. 918–923.
- [32] Mohammad Reza Karimi, Mario Lucic, S. Hamed Hassani, and Andreas Krause. 2017. Stochastic Submodular Maximization: The Case of Coverage Functions. In *NIPS*. 6856–6866.
- [33] Richard M. Karp. 1990. The transitive closure of a random digraph. *Random Struct. Algor.* 1, 1 (1990), 73–93.
- [34] Xiangyu Ke, Arijit Khan, and Gao Cong. 2018. Finding Seeds and Relevant Tags Jointly: For Targeted Influence Maximization in Social Networks. In *SIGMOD*. 1097–1111.
- [35] David Kempe, Jon Kleinberg, and Éva Tardos. 2003. Maximizing the Spread of Influence through a Social Network. In *KDD*. 137–146.
- [36] David Kempe, Jon Kleinberg, and Éva Tardos. 2005. Influential Nodes in a Diffusion Model for Social Networks. In *ICALP*. 1127–1138.
- [37] David Kempe, Jon Kleinberg, and Éva Tardos. 2015. Maximizing the Spread of Influence through a Social Network. *Theory Comput.* 11 (2015), 105–147.
- [38] Masahiro Kimura and Kazumi Saito. 2006. Tractable Models for Information Diffusion in Social Networks. In *PKDD*. 259–271.
- [39] Masahiro Kimura, Kazumi Saito, and Ryohei Nakano. 2007. Extracting Influential Nodes for Information Diffusion on a Social Network. In *AAAI*. 1371–1376.
- [40] Masahiro Kimura, Kazumi Saito, Ryohei Nakano, and Hiroshi Motoda. 2010. Extracting influential nodes on a social network for information diffusion. *Data Min. Knowl. Discov.* 20, 1 (2010), 70–97.
- [41] Andreas Krause, Ajit Paul Singh, and Carlos Guestrin. 2008. Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies. *J. Mach. Learn. Res.* 9 (2008), 235–284.
- [42] Jérôme Kunegis. 2017. Konect Network Dataset – KONECT. <http://konect.uni-koblenz.de/networks/konect>.
- [43] Laks V. S. Lakshmanan, Panayiotis Tsaparas, and Yuichi Yoshida. 2018. *Influence Analytics in Graphs*. Springer International Publishing, 1–8.
- [44] Jure Leskovec, Andreas Krause, Carlos Guestrin, Christos Faloutsos, Jeanne VanBriesen, and Natalie Glance. 2007. Cost-effective Outbreak Detection in Networks. In *KDD*. 420–429.
- [45] Jure Leskovec and Andrej Krevl. 2014. SNAP Datasets: Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data>.
- [46] Jure Leskovec, Kevin J. Lang, Anirban Dasgupta, and Michael W. Mahoney. 2008. Statistical Properties of Community Structure in Large Social and Information Networks. In *WWW*. 695–704.
- [47] Xiang Li, J. David Smith, Thang N. Dinh, and My T. Thai. 2017. Why approximate when you can get the exact? Optimal Targeted Viral

- Marketing at Scale. In *INFOCOM*. 1–9.
- [48] Yuchen Li, Ju Fan, Yanhao Wang, and Kian-Lee Tan. 2018. Influence Maximization on Social Graphs: A Survey. *IEEE Trans. Knowl. Data Eng.* 30, 10 (2018), 1852–1872.
 - [49] Wei Lu, Xiaokui Xiao, Amit Goyal, Keke Huang, and Laks V. S. Lakshmanan. 2017. Refutations on "Debunking the Myths of Influence Maximization: An In-Depth Benchmarking Study". *CoRR abs/1705.05144* (2017).
 - [50] Takanori Maehara, Takuya Akiba, Yoichi Iwata, and Ken-ichi Kawarabayashi. 2014. Computing Personalized PageRank Quickly by Exploiting Graph Structures. *Proc. VLDB Endow.* 7, 12 (2014), 1023–1034.
 - [51] Takanori Maehara, Hirofumi Suzuki, and Masakazu Ishihata. 2017. Exact Computation of Influence Spread by Binary Decision Diagrams. In *WWW*. 947–956.
 - [52] Michael Mathioudakis, Francesco Bonchi, Carlos Castillo, Aristides Gionis, and Antti Ukkonen. 2011. Sparsification of Influence Networks. In *KDD*. 529–537.
 - [53] Makoto Matsumoto and Takuji Nishimura. 1998. Mersenne Twister: A 623-Dimensionally Equidistributed Uniform Pseudo-Random Number Generator. *ACM Trans. Model. Comput. Simul.* 8, 1 (1998), 3–30.
 - [54] George L. Nemhauser, Laurence A. Wolsey, and Marshall L. Fisher. 1978. An analysis of the approximations for maximizing submodular set functions. *Math. Program.* 14 (1978), 265–294.
 - [55] Hung T. Nguyen, Thang N. Dinh, and My T. Thai. 2016. Cost-aware Targeted Viral Marketing in Billion-scale Networks. In *INFOCOM*. 1–9.
 - [56] Hung T. Nguyen, Preetam Ghosh, Michael L. Mayo, and Thang N. Dinh. 2017. Social Influence Spectrum at Scale: Near-Optimal Solutions for Multiple Budgets at Once. *ACM Trans. Inf. Syst.* 36, 2 (2017), 14:1–14:26.
 - [57] Hung T. Nguyen, Tri P. Nguyen, Tam N. Vu, and Thang N. Dinh. 2017. Importance Sketching of Influence Dynamics in Billion-scale Networks. In *ICDM*. 337–346.
 - [58] Hung T. Nguyen, Tri P. Nguyen, Tam N. Vu, and Thang N. Dinh. 2017. Outward Influence and Cascade Size Estimation in Billion-scale Networks. In *SIGMETRICS*. 63.
 - [59] Hung T. Nguyen, Tri P. Nguyen, Tam N. Vu, and Thang N. Dinh. 2017. Outward Influence and Cascade Size Estimation in Billion-scale Networks. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 1, 1 (2017), 20:1–20:30.
 - [60] Hung T. Nguyen, My T. Thai, and Thang N. Dinh. 2016. Stop-and-Stare: Optimal Sampling Algorithms for Viral Marketing in Billion-scale Networks. In *SIGMOD*. 695–710.
 - [61] Hung T. Nguyen, My T. Thai, and Thang N. Dinh. 2017. A Billion-Scale Approximation Algorithm for Maximizing Benefit in Viral Marketing. *IEEE/ACM Trans. Netw.* 25, 4 (2017), 2419–2429.
 - [62] Naoto Ohsaka, Takuya Akiba, Yuichi Yoshida, and Ken-ichi Kawarabayashi. 2014. Fast and Accurate Influence Maximization on Large Networks with Pruned Monte-Carlo Simulations. In *AAAI*. 138–144.
 - [63] Naoto Ohsaka, Takuya Akiba, Yuichi Yoshida, and Ken-ichi Kawarabayashi. 2016. Dynamic Influence Analysis in Evolving Networks. *Proc. VLDB Endow.* 9, 12 (2016), 1077–1088.
 - [64] Naoto Ohsaka, Tomohiro Sonobe, Sumio Fujita, and Ken-ichi Kawarabayashi. 2017. Coarsening Massive Influence Networks for Scalable Diffusion Analysis. In *SIGMOD*. 635–650.
 - [65] Diana Popova, Akshay Khot, and Alex Thomo. 2018. Data Structures for Efficient Computation of Influence Maximization and Influence Estimation. In *EDBT*. 505–508.
 - [66] Manish Purohit, B. Aditya Prakash, Chanhyun Kang, Yao Zhang, and V.S. Subrahmanian. 2014. Fast Influence-based Coarsening for Large Networks. In *KDD*. 1296–1305.
 - [67] Dravyansh Sharma, Ashish Kapoor, and Amit Deshpande. 2015. On Greedy Maximization of Entropy. In *ICML*. 1330–1338.
 - [68] Jing Tang, Xueyan Tang, Xiaokui Xiao, and Junsong Yuan. 2018. Online Processing Algorithms for Influence Maximization. In *SIGMOD*. 991–1005.
 - [69] Youze Tang, Yanchen Shi, and Xiaokui Xiao. 2015. Influence Maximization in Near-Linear Time: A Martingale Approach. In *SIGMOD*. 1539–1554.
 - [70] Youze Tang, Xiaokui Xiao, and Yanchen Shi. 2014. Influence Maximization: Near-Optimal Time Complexity Meets Practical Efficiency. In *SIGMOD*. 75–86.
 - [71] Leslie G. Valiant. 1979. The Complexity of Enumeration and Reliability Problems. *SIAM J. Comput.* 8, 3 (1979), 410–421.
 - [72] Duncan J. Watts and Steven H. Strogatz. 1998. Collective dynamics of 'small-world' networks. *Nature* 393, 6684 (1998), 440–442.
 - [73] Chuan Zhou, Peng Zhang, Jing Guo, and Li Guo. 2014. An Upper Bound based Greedy Algorithm for Mining Top-k Influential Nodes in Social Networks. In *WWW*. 421–422.
 - [74] Chuan Zhou, Peng Zhang, Jing Guo, Xingquan Zhu, and Li Guo. 2013. UBLF: An Upper Bound Based Approach to Discover Influential Nodes in Social Networks. In *ICDM*. 907–916.