Better Bounds on the Adaptivityf Gap of Influence Maximization under Full-adoption Feedback★

Gianlorenzo D'Angelo[a], Debashmita Poddar[a,], Cosimo Vinci[b,a,]

[a] *Gran Sasso Science Institute - L'Aquila, Italy*
[b] *University of Salento - Lecce, Italy*

Abstract

In the *influence maximization (IM)* problem, we are given a social network and a budget k, and we look for a set of k nodes in the network, called seeds, that maximize the expected number of nodes that are reached by an influence cascade generated by the seeds, according to some stochastic model for influence diffusion. Extensive studies have been done on the IM problem, since this definition by **?**]. However, most of the work focuses on the *non-adaptive* version of the problem where all the k seed nodes must be selected before the cascade starts. In this paper we study the *adaptive* IM, where the nodes are selected sequentially one by one, and the decision on the i-th seed can be based on the *observed* cascade produced by the first $i-1$ seeds. We focus on the *full-adoption feedback* in which we can observe the entire cascade of each previously selected seed under the independent cascade model where each edge is associated with an independent probability of diffusing influence.

Previous works showed that there are constant upper bounds on the adaptivity gap, which compares the performance of an adaptive algorithm against a non-adaptive one, but the analyses used to prove these bounds only works for specific graph classes such as in-arborescences, out-arborescences, and one-directional bipartite graphs. Our main result is the first sub-linear upper bound that holds for any graph. Specifically, we show that the adaptivity gap is upper-bounded by $\sqrt[3]{n}+1$, where n is the number of nodes in the graph. Moreover, we improve over the known upper bound for in-arborescences from $2e/(e-1) \approx 3.16$ to $2e^2/(e^2-1) \approx 2.31$. Then, we consider (β, γ)-bounded-activation graphs, where all nodes but β influence in expectation at most $\gamma \in [0, 1)$ neighbors each; for this class of influence graphs we show that the adaptivity gap is at most $\sqrt{\beta} + \frac{1}{1-\gamma}$. Finally, we study α-bounded-degree graphs, that is the class of undirected graphs in which the sum of node degrees higher than two is at most α, and show that the adaptivity gap is upper-bounded by $\sqrt{\alpha} + O(1)$; we also show that in 0-bounded-degree graphs, i.e. undirected graphs in which each connected component is a path or a cycle, the adaptivity gap is at most $3e^3/(e^3-1) \approx 3.16$.

To prove our bounds, we introduce new techniques to relate adaptive policies with

★A preliminary version of this work appeared in "Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21)" [**?**].

Email addresses: `gianlorenzo.dangelo@gssi.it` (Gianlorenzo D'Angelo), `debashmita.poddar@gssi.it` (Debashmita Poddar), `cosimo.vinci@unisalento.it` (Cosimo Vinci)

non-adaptive ones that might be of their own interest.

Keywords: influence maximization, adaptive optimization, adaptivity gap, graph algorithms

1. Introduction

In the Influence Maximization (IM) problem, we are given a social network, a stochastic model for diffusion of influence over the network, and a budget k. We are asked to find a set of k nodes, called *seeds*, that maximize their *spread of influence*, which is the expected number of nodes reached by a cascade of influence diffusion, generated by the seeds, according to the given diffusion model. One of the most studied model for influence diffusion is the Independent Cascade (IC), where each edge is associated with an independent probability of transmitting influence from the source node to the tail node. In the IC model, the spread of influence is a monotone submodular function of the seed set, therefore a greedy algorithm guarantees a $1-\frac{1}{e}$ approximation factor for the IM problem [**?**]. Since this definition [**? ?**] and formalization as an optimization problem [**? ?**], the IM problem and its variants have been extensively investigated, motivated by applications in viral marketing [**?**], adoption of technological innovations [**?**], and outbreak or failure detection [**?**]. See **? ?**] for surveys on the IM problem.

Recently, **?**] initiated the study of the IM problem under the framework of adaptive optimization, where, instead of selecting all the seeds at once in the beginning of the process, we can select one seed at a time and observe to some extent, the portion of the network reached by a new selected seed. The advantage is that the decision on the next seed to choose can be based on the observed spread of previously selected seeds, usually called *feedback*. Two main feedback models have been introduced: in the *full-adoption* feedback, the whole spread from each seed can be observed, while in the *myopic* feedback, one can only observe the direct neighbors of each seed.

An interesting application of adaptive strategies under both feedback models can be encountered in viral marketing [**?**]: Assume that a company is selling a new product and wants to advertise it to the largest number of possible buyers in order to increase its revenue. The company can use a limited budget to initially promote the product to some persons (i.e., the seeds) and, by exploiting the word-of-mouth effect, the product will be indirectly advertised to many other people in the underlying social network. In a non-adaptive strategy, the company selects the seeds at the beginning, without observing how the word-of-mouth effect contributes to the product advertising. A disadvantage of such a strategy is that some selected seed could be useless, e.g., it could have been already reached by the word-of-mouth induced by other seeds. To overcome this issue, the company could resort to an adaptive strategy to observe, for each selected seed, the set of people reached by the word-of-mouth and choose the next seeds focusing on the users that have not yet been reached. Under full-adoption feedback, the company is able to know exactly all people influenced by the advertising campaign, thus the company can avoid to advertise again the product to such people. Instead, under myopic feedback, the observation is restricted to the neighbors of the selected seeds, thus the useless seeding can only be partially avoided.

In this work, we focus on the model of full-adoption feedback. Golovin and Krause considered the Independent Cascade model and showed that, under full-adoption feedback, the objective function satisfies the property of *adaptive submodularity* (introduced in the same paper) and therefore a greedy algorithm achieves a $1 - \frac{1}{e}$ approximation for the adaptive IM problem. They also conjectured that there exists a constant factor approximation algorithm for the myopic feedback model, which has been found by **?**] who proposed a $\frac{1}{4}\left(1 - \frac{1}{e}\right)$-approximation algorithm.

However, the approximation ratio for the adaptive IM problem, which compares a given adaptive algorithm with an optimal adaptive one, does not measure the benefits of implementing adaptive policies over non-adaptive ones. In order to measure the superiority of adaptive policies over non-adaptive ones, **? ? ?**] analyzed the *adaptivity gap*, which is the supremum, over all possible inputs, of the ratio between the spread of an optimal adaptive policy and that of an optimal non-adaptive one. We remark that implementing a non-adaptive policy is simpler than using an adaptive one, as the former does not require to observe the feedback from the seed selection and all seeds are computed at the beginning. Thus, if the adaptivity gap is sufficiently low, we have that the seed set computed by an optimal non-adaptive policy (or an almost optimal one, such as the non-adaptive greedy algorithm) guarantees an expected influence spread that is close to that achievable in an optimal adaptive way (or that achieved by an almost optimal adaptive policy, such as the adaptive greedy algorithm).

In [**?**], Peng and Chen considered independent cascade model with myopic feedback and showed that the adaptivity gap is between $\frac{e}{e-1}$ and 4 for any graph. In [**?**], the same authors showed some upper and lower bounds on the adaptivity gap in the case of full-adoption feedback, still under independent cascade, for some particular graph classes. Specifically, they showed that the adaptivity gap is in the interval $\left[\frac{e}{e-1}, \frac{2e}{e-1}\right]$ for in-arborescences and it is in the interval $\left[\frac{e}{e-1}, 2\right]$ for out-arborescences. Moreover, it is equal to $\frac{e}{e-1}$ in one-directional bipartite graphs. In order to show these bounds, they followed an approach introduced by **?**], which consists in transforming an adaptive policy into a non-adaptive one by means of multilinear extensions, and constructing a Poisson process to relate the influence spread of the non-adaptive policy to that of the adaptive one. For general graphs and full-adoption feedback, the only known upper bounds on the adaptivity gap are linear in the size of the graph and can be trivially derived.

1.1. Our Contribution

In this paper, we consider the independent cascade model with full-adoption feedback, and show the first sub-linear upper bound on the adaptivity gap that holds for general graphs. In detail, we show that the adaptivity gap is at most $\lceil n^{1/3} \rceil$ (Theorem 3), where n is the number of nodes in the graph; as a corollary of Theorem 3 (Corollary 1), we also show that the non-adaptive greedy algorithm guarantees a $\Omega(1/n^{1/3})$ approximation to the adaptive optimum. Moreover, we tighten the upper bound on the adaptivity gap for in-arborescences by showing that it is at most $\frac{2e^2}{e^2-1} < \frac{2e}{e-1}$ (Theorem 1).

Using similar techniques we study the adaptivity gap of two classes of influence graphs introduced in this work: (β, γ)-*bounded-activation graphs*, where all nodes but β influence in expectation at most $\gamma \in [0, 1)$ neighbors each, and α-*bounded-degree* *graphs*, which is the class of influence graphs where the sum of node degrees higher than two is at most α. (β, γ)-bounded-activation graphs are well-motivated by social networks in which most nodes have a limited power of influence, and α-bounded-degree graphs can be encountered in several graph topologies (see, for instance, Example 1). We show that the adaptivity gap of (β, γ)-bounded-activation graphs and α-bounded-degree graphs is upper-bounded by $\sqrt{\beta} + \frac{1}{1-\gamma}$ and $\sqrt{\alpha} + O(1)$ (Theorems 4 and 5), respectively, and these values are smaller than that of $O(n^{1/3})$ for several influence graph classes. Furthermore, in 0-bounded-degree graphs, i.e. undirected graphs in which each connected component is a path or a cycle, the adaptivity gap is at most $\frac{3e^3}{e^3-1}$ (Theorem 6). In Table 1 we summarize the obtained results on the adaptivity gap of the considered graph classes, and we compare our bounds with the existing ones.

	Previous UB	Our UB
General Graphs	x	$\lceil n^{1/3} \rceil$
In-arborescences	$\frac{2e}{e-1}$ [**?**]	$\frac{2e^2}{e^2-1}$
(β, γ)-bounded-activation	//	$\sqrt{\beta} + \frac{1}{1-\gamma}$
α-bounded-degree	//	$\sqrt{\alpha} + O(1)$
Paths or Cycles	//	$\frac{3e^3}{e^3-1}$

Table 1: The table compares our upper bounds (UB) on the adaptivity gap with the existing ones, for the considered graph classes. Symbol "x" is used if no sub-linear upper bound was known, and symbol "//" denotes the cases that, prior to this work, had not been studied yet.

To prove our bounds, we introduce new techniques to connect adaptive policies with non-adaptive ones that might be of their own interest (further details are given in the paragraph "General outline of the proof technique" in Section 3). In particular, we resort to a simple randomized *hybrid non-adaptive policy*, that differs from the main approaches previously used in adaptive influence maximization and other adaptive optimization problems: (i) the *Poisson process* [**?**] combined with the *multi-linear extension of submodular set-functions* [**?**], which represent the main probabilistic technique adopted by **?**] and **?**], and (ii) the *random walk on the decision tree* [**? ?**], that is a tool applied by **? ?**] and **?**].

1.2. Related Work

Non-Adaptive Influence Maximization. Several studies based on general graphs [**? ? ? ?**] have been conducted since the seminal paper by **?**]. **?**] studied the influence maximization problem on undirected graphs and proved that it is APX-hard for both the independent cascade and the linear threshold problem. **?**] proposed an efficient algorithm that runs in quasilinear time and still guarantees an approximation factor of $1 - \frac{1}{e} - \epsilon$, for any $\epsilon > 0$. **?**] proposed an algorithm which is experimentally close to the optimal one under the independent cascade model. **?**] consider unknown graphs for the influence maximization problem and devised an algorithm which achieves a fraction between 0.6 and 0.9 of the influence spread with minimal knowledge of the graph topology. Extensive literature reviews on influence maximization and its machinery is provided by **?**] and **?**].

Adaptive Influence Maximization. **?**] considered the independent cascade model under the full-adoption feedback. By exploiting the adaptive submodularity property, they showed that the adaptive greedy algorithm guarantees a $1 - \frac{1}{e} - \epsilon$ approximation. **?**] analyzed the efficiency of the adaptive greedy under diffusion models that do not satisfy the adaptive submodularity, and showed a constant approximation factor for one-directional bipartite graphs under the triggering model.

The myopic feedback model (in which, one can only observe the nodes influenced by the seed nodes) has been introduced in [**?**] and further analyzed in [**? ? ?**]. **?**] showed that both non-adaptive and adaptive greedy algorithms achieve a fraction $\frac{1}{4}\left(1 - \frac{1}{e}\right) \approx 0.158$ of the optimal adaptive policy (up to an arbitrarily small addend $\epsilon > 0$), under the myopic model and for general graphs; this result was recently improved by **?**], who showed that the non-adaptive and adaptive greedy algorithms are respectively $\frac{1}{2}\left(1 - \frac{1}{e}\right) \approx 0.316$ and $1 - \frac{1}{\sqrt{e}} \approx 0.393$ approximate to the optimal adaptive policy (up to an arbitrarily small addend $\epsilon > 0$).

In Table 2 we provide the approximation factors guaranteed by the adaptive and the non-adaptive greedy algorithms in the context of adaptive influence maximization.

	Non-adaptive Greedy	Adaptive Greedy
Full-ad.	$\Omega(1/n^{1/3})$ [*]	$1 - \frac{1}{e} - \epsilon$ [**?**]
Myopic	$\frac{1}{4}\left(1 - \frac{1}{e}\right) - \epsilon$ [**?**]	$\frac{1}{4}\left(1 - \frac{1}{e}\right) - \epsilon$ [**?**]
	$\frac{1}{2}\left(1 - \frac{1}{e}\right) - \epsilon$ [**?**]	$\left(1 - \frac{1}{\sqrt{e}}\right) - \epsilon$ [**?**]

Table 2: The table represents the approximation ratios guaranteed by the adaptive and the non-adaptive greedy algorithms in the context of adaptive influence maximization, under both full-adoption and myopic feedback. We use symbol "$*$" to denote the results obtained in this work.

Several works on the adaptive influence maximization problem [**? ? ? ? ? ? ? ? ? ? ?**] evolved after the concept was introduced by **?**], and explore adaptive optimization under different settings.

?] capture the scenario in which, instead of considering one round, the diffusion process takes over T rounds, and a seed set of at most k nodes is selected at each round. The authors designed a greedy approximation algorithm that guarantees a constant approximation ratio. **?**] introduced a new version of the adaptive influence maximization problem by adding a time constraint. Other than the classic full-adoption and myopic feedback model, **?**], and **?**], have also introduced different feedback models that use different parameters to overcome the need of submodularity to guarantee a good approximation. **?**] and **?**] proposed a framework that generalizes the adaptive influence maximization problem in which, at each step, a batch of nodes is selected in a non-adaptive way (instead of a single node as in classic adaptive IM), and then the resulting spread is observed to select the subsequent batches; in particular, they showed an approximation factor of $1 - e^{\rho_b(\epsilon - 1)}$, where $\rho_b = 1 - (1 - 1/b)^b$ and b is the size of each batch. Another multi-batches variant of adaptive IM problems has been studied by **?**], who provided an approximation algorithm and a heuristic algorithm for the related computational problems.

Adaptivity Gaps. Adaptivity gaps for the problem of maximizing stochastic monotone submodular functions have been studied by **?**]. A series of work studied adaptivity gaps for a two-step adaptive influence maximization problem [**? ? ? ?**]. Gupta et al. [**? ?**] and **?**] worked on the adaptivity gaps for stochastic probing.

A recent line of studies [**? ? ?**] focuses on finding the adaptivity gaps on different graph classes using the classical feedback models. **?**] confirmed a conjecture of **?**], which states that the adaptivity gap of the independent cascade model with myopic feedback is constant. In particular, they showed that the adaptivity gap belongs to $[\frac{e}{e-1}, 4]$. The above upper bound has been recently lowered to $\frac{2e}{e-1} \approx 3.16$ by **?**]. **?**] introduced the greedy adaptivity gap, which compares the performance of the adaptive and the non-adaptive greedy algorithms. They showed that the infimum of the greedy adaptivity gap is $1 - \frac{1}{e}$ for every combination of diffusion and feedback models.

The most related work to our results is that of **?**], as they derived upper and lower bounds on the adaptivity gap under the independent cascade model with full-adoption feedback, when the considered graphs are in-arborescences, out-arborescences, or one-directional bipartite graphs. In particular, they showed that the adaptivity gaps of in-arborescences and out-arborescences are in the intervals $\left[\frac{e}{e-1}, \frac{2e}{e-1}\right]$ and $\left[\frac{e}{e-1}, 2\right]$, respectively, and they provided a tight bound of $\frac{e}{e-1}$ on the adaptivity gap of one-directional bipartite graphs. Under the more general triggering model, a constant upper bound of 2 on the adaptivity gap of one-directional bipartite graphs was provided by **?**].

In Table 3 we summarize the main results on the adaptivity gap of influence maximization under the independent cascade model.

		General	In-arb.	Out-arb.	Bipartite
Full-ad.	UB	$\lceil n^{1/3} \rceil$ [*]	$\frac{2e}{e-1}$ [**?**], $\frac{2e^2}{e^2-1}$ [*]	2 [**?**]	$\frac{e}{e-1}$ [**?**]
	LB	$\frac{e}{e-1}$ [**?**]	$\frac{e}{e-1}$ [**?**]	$\frac{e}{e-1}$ [**?**]	$\frac{e}{e-1}$ [**?**]
Myopic	UB	4 [**?**], $\frac{2e}{e-1}$ [**?**]	//	//	$\frac{e}{e-1}$ [**?**]
	LB	$\frac{e}{e-1}$ [**? ?**]	//	//	$\frac{e}{e-1}$ [**? ?**]

Table 3: The table represents the existing bounds on the adaptivity gap of influence maximization under the independent cascade model, for both full-adoption and myopic feedback, and different graph topologies (general graphs, in-arborescences, out-arborescences, one-directional bipartite graphs). The acronyms "UB" and "LB" refer to upper and lower bounds, respectively. We use symbol "$*$" to denote the results obtained in this work, and symbol "//" to denote cases which have not been studied yet (to the best of our knowledge).

1.3. Organization of the Paper

In Section 2 we give the preliminary definitions and notations which this work is based on. Sections 3–5 are devoted to the main technical contribution of the paper, i.e., adaptivity gaps of in-arborescences, general graphs and other influence graphs ((β, γ)-bounded-activation graphs and α-bounded-degree graphs). In Section 6, we highlight some future research directions.

2. Preliminaries

For two integers h and k, $h \leq k$, let $[k]_h := \{h, h+1, \dots, k\}$ and $[k] := [k]_1$.

Independent Cascade Model. In the *independent cascade model* (IC), we have an *influence graph* $G = (V = [n], E, (p_{uv})_{(u,v)\in E})$, where $p_{uv} \in [0, 1]$ is an *activation probability* associated to each edge $(u, v) \in E$. Given a set of *seed nodes* $S \subseteq V$ which are initially *active*, the diffusion process in the IC model is defined in $t \geq 0$ discrete steps as follows: (i) let A_t be the set of active nodes which are activated at each step $t \geq 0$; (ii) $A_0 := S$; (iii) given a step $t \geq 0$, for any edge (u, v) such that $u \in A_t$, node u can activate node v with probability p_{uv} independently from any other node, and, in case of success, v is included in A_{t+1}; (iv) the diffusion process ends at a step $r \geq 0$ such that $A_r = \emptyset$, i.e., no node can be activated at all, and $\bigcup_{t\leq r} A_t$ is the *influence spread*, i.e., the set of nodes activated/reached by the diffusion process.

The above diffusion process can be equivalently defined as follows. The *live-edge graph* $L = (V, L(E))$ of G is a random graph made from G, where $L(E) \subseteq E$ is a subset of edges such that each edge $(u, v) \in E$ is included in $L(E)$ with probability p_{uv}, independently from the other edges. Given a live-edge graph L, let $R(S, L) := \{v \in V :$ there exists a path from u to v in L for some $u \in S\}$, i.e., the set of nodes reached by S in the live-edge graph L. Informally, if S is the set of seed nodes, and L is a live-edge graph, $R(S, L)$ equivalently denotes the set of nodes which are reached/activated by the above diffusion process. Given a set of seed nodes S, the *expected influence spread* of S is defined as $\sigma(S) := \mathbb{E}_L[|R(S, L)|]$.

Non-adaptive Influence Maximization. The *non-adaptive influence maximization problem under the IC model* is a computational problem that is defined as follows: given an influence graph G and an integer $k \geq 1$, we are to find a set of seed nodes $S \subseteq V$ with $|S| = k$ such that $\sigma(S)$ is maximized.

Adaptive Influence Maximization. Differently from the non-adaptive setting, in which all the seed nodes are activated at the beginning and then the influence spread is observed, an *adaptive policy* activates the seeds sequentially in k steps, one seed node at each step, and the decision on the next seed node to select is based on the feedback resulting from the observed spread of the previously selected nodes. The feedback model considered in this work is *full-adoption*: when a node is selected, the adaptive policy observes its entire influence spread.

An adaptive policy under the full-adoption feedback model is formally defined as follows. Given a live-edge graph L, the *realization* $\phi_L : V \to 2^V$ associated to L assigns to each node $v \in V$ the value $\phi_L(v) := R(\{v\}, L)$, i.e., the set of nodes activated by v under a live-edge graph L. Given a set $S \subseteq V$, a *partial realization* $\psi : S \to 2^V$ is the restriction to S for some realization, i.e., there exists a live-edge graph L such that $\psi(v) = \phi_L(v)$ for any $v \in S$. Given a partial realization $\psi : S \to 2^V$, let $dom(\psi) := S$, i.e., $dom(\psi)$ is the domain of partial realization ψ, let $R(\psi) := \bigcup_{v\in S} \psi(v)$, i.e., $R(\psi)$ is the set of nodes reached/activated by the diffusion process when the set of seed nodes is S, and let $f(\psi) := |R(\psi)|$. A partial realization ψ' is a *sub-realization* of ψ (or, equivalently, $\psi' \subseteq \psi$), if $dom(\psi') \subseteq dom(\psi)$ and $\psi'(v) = \psi(v)$ for any $v \in dom(\psi')$. We observe that a partial realization ψ can be equivalently represented as $\{(v, R(\{v\}, L)) : v \in dom(\psi)\}$ for some live-edge graph L. Some of the above definitions are illustrated in Figure 2.

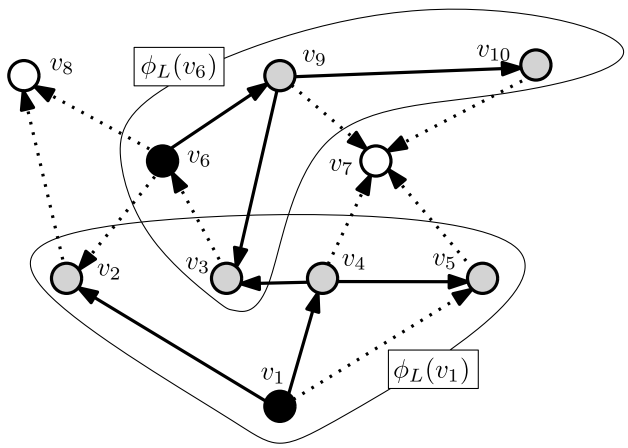

Figure 1: The figure represents an influence graph G having 10 nodes. Each arrow is an edge, and the set of non-dotted edges represents the live-edge graph L. The seed set is $S = \{v_1, v_6\}$ (i.e., the black nodes), the set of nodes reached by v_1 and v_6 is $R(S, L) = \{v_1, v_2, v_3, v_4, v_5, v_6, v_9, v_{10}\}$ (i.e., the grey nodes), and the nodes that remain non-active are v_7 and v_8 (i.e., the white ones). The partial realization $\psi : S \to 2^V$ observed by seeding the nodes in $dom(\psi) = S$ verifies $\psi(v_1) = \phi_L(v_1) = \{v_1, v_2, v_3, v_4, v_5\}$ and $\psi(v_6) = \phi_L(v_6) = \{v_3, v_6, v_9, v_{10}\}$, where ϕ_L denotes the realization associated with live-edge graph L.

An adaptive policy π takes as input a partial realization ψ and, either returns a node $\pi(\psi) \in V$ and activates it as seed, or interrupts the activation of new seed nodes, e.g., by returning a string $\pi(\psi) := STOP$; furthermore, we assume w.l.o.g. that an adaptive policy π chooses each seed based only on the observation of $R(\psi)$ (i.e., the set of nodes activated by the previous seeds) and $|dom(\psi)|$ (i.e., the number of seeds previously selected), that is, $\pi(\psi) = \pi(R(\psi), |dom(\psi)|)$. An adaptive policy π can be run as in Algorithm 1.[1] The *expected influence spread* of an adaptive policy π is defined as $\sigma(\pi) := \mathbb{E}_L[f(\psi_{\pi,L})]$, i.e., it is the expected value (taken on all the possible live-edge graphs) of the number of nodes reached by the diffusion process at the end of Algorithm 1. We say that $|\pi| = k$ if policy π always return a partial realization $\psi_{\pi,L}$ with $|dom(\psi_{\pi,L})| = k$. The *adaptive influence maximization problem under the IC model* is the computational problem that, given an influence graph G and an integer $k \geq 1$, asks to find an adaptive policy π that maximizes the expected influence spread $\sigma(\pi)$ subject to a constraint $|\pi| = k$.

Adaptivity gap. Given an influence graph G and an integer $k \geq 1$, let $OPT_N(G, k)$ (resp. $OPT_A(G, k)$) denote the optimal value of the non-adaptive (resp. adaptive) influence maximization problem with input G and k. Given a class of influence graphs

[1]The assumption that π can only observe $R(\psi)$ and $dom(\psi)$ is done to make the adaptive policy more realistic. Indeed, assuming that a seed v_1 is activated before seed v_2, and a node u has been already activated by v_1, it might unfeasible to know if v_2 would have influenced u without the help of seed v_2. We also observe that the above assumption is done without loss of generality. Indeed, to the aim of designing an optimal adaptive policy under the full-feedback model, it is sufficient to observe which nodes have been influenced by the previous seed selections, without necessarily knowing if a new seed could have reached a node that was previously influenced by some previous seed.

Algorithm 1 Adaptive algorithm

Require: an influence graph G and an adaptive policy π;
Ensure: a partial realization;
let L be the live-edge graph;
let $\psi := \emptyset$ (i.e., ψ is the empty partial realization);
while $\pi(\psi) \neq STOP$ **do**
    $v := \pi(\psi)$;
    $\psi := \psi \cup \{(v, R(\{v\}, L))\}$;
end while
return $\psi_{\pi,L} := \psi$;

$\mathcal{G}$ and an integer $k \geq 1$, the *k-adaptivity gap* of $\mathcal{G}$ is defined as

$$AG(\mathcal{G}, k) := \sup_{G \in \mathcal{G}} \frac{OPT_A(G, k)}{OPT_N(G, k)},$$

and measures how much an adaptive policy outperforms a non-adaptive solution for the influence maximization problem applied to influence graphs in $\mathcal{G}$, when the maximum number of seed nodes is k. The *adaptivity gap* of $\mathcal{G}$ is defined as $AG(\mathcal{G}) := \sup_{k \geq 1} AG(\mathcal{G}, k)$. We observe that for $k = 1$ or $n \leq k$ the k-adaptivity gap is trivially equal to 1, thus we omit such cases in the following.

In Table 4 we summarize the notation in Influence Maximization defined in this section.

Symbol	Meaning
G	input graph
$V = [n]$	set of n nodes
E	set of edges
$p_{u,v}$	diffusion probability of edge (u, v)
L	live-edge graph
$R(S, L)$	set of nodes reached by S in graph L
$\sigma(S)$	expected influence spread $\mathbb{E}_L[\lvert R(S, L)\rvert]$ generated by S
$\phi_L : V \to 2^V$	realization associated with L
	(i.e., $\phi_L(v) := R(\{v\}, L)$ for any $v \in V$)
ψ	partial realization of ϕ_L
	(i.e., the realization ϕ_L restricted to the selected seed set S)
$dom(\psi)$	the domain S of ψ
$R(\psi)$	$\bigcup_{v \in dom(\psi)} \psi(v)$
$f(\psi)$	$\lvert R(\psi)\rvert$
$\psi' \subseteq \psi$	ψ' is a sub-realization of ψ
π	generic adaptive policy
$OPT_A(G, k)$	adaptive optimal value of a policy selecting k seeds in graph G
$OPT_N(G, k)$	non-adaptive optimal value for k seeds in graph G

Table 4: Notation on Influence Maximization

3. Adaptivity Gap for In-arborescences

An *in-arborescence* is a graph $G = (V, E)$ that can be constructed from a rooted tree $T = (V, F)$, by adding in E an edge (v, u) if u is a father of v in tree T. In-arborescences have been widely studied in the context of influence maximization (see, for instance, [**?** **?** **?**]), and can be successfully used to represent diffusion networks possessing a hierarchical structure, where each node can influence its father, and indirectly, some of its ancestors. An upper bound of $\frac{2e}{e-1} \approx 3.16$ on the adaptivity gap of in-arborescences has been provided in [**?**].

In this section we provide an improved upper bound for such graphs.

Theorem 1. If $\mathcal{G}$ is the class of all the in-arborescences, then

$$AG(\mathcal{G}, k) \leq \frac{2}{1-(1-2/k)^k} \leq \frac{2e^2}{e^2-1} \approx 2.31, \ \forall k \geq 2.$$

Let $G = (V = [n], E, (p_{uv})_{(u,v)\in E})$ be an in-arborescence, where $n > k$ is the number of nodes. To show the claim of Theorem 1, we need some preliminary notations and lemmas. Given a partial realization ψ, and a node $v \in V$, let

$$\Delta(v|\psi) := \mathbb{E}_L[f(\psi \cup \{(v, R(\{v\}, L))\}) - f(\psi)|\psi \subseteq \phi_L],$$

i.e., $\Delta(v|\psi)$ is the expected increment of the influence spread due to node v when the observed partial realization is ψ. We have the following claim (from [**?**]), holding even for general graphs, whose proof is trivial.

Claim 1 (Adaptive Submodularity [**?**]). Let G be an arbitrary influence graph. For any partial realizations ψ, ψ' of G such that $\psi \subseteq \psi'$, and any node $v \notin R(\psi')$, we have that $\Delta(v|\psi') \leq \Delta(v|\psi)$.

An adaptive policy π is called *randomized* if, for any partial realization ψ, node $\pi(\psi)$ is not selected deterministically (in general), but randomly (according to a probability distribution p_ψ depending on ψ). Given a vector $\boldsymbol{y} = (y_1, \ldots, y_n)$ such that $y_v \in [0, 1]$ for any $v \in V$, we say that $\mathbb{P}(\pi) = \boldsymbol{y}$ if the probability that each node v belongs to $dom(\psi_{\pi,L})$ is y_v, where $\psi_{\pi,L}$ is the partial realization returned by Algorithm 1 with policy π. Let $OPT_A(G, \boldsymbol{y})$ be the optimal expected influence spread $\sigma(\pi)$ over all the randomized adaptive policies π such that $\mathbb{P}(\pi) = \boldsymbol{y}$.[2]

Let π^* be an optimal adaptive policy for the adaptive influence maximization problem (with $|\pi^*| = k$), and let $\boldsymbol{x} = (x_1, \ldots, x_n)$ be the vector such that $\mathbb{P}(\pi^*) = \boldsymbol{x}$. As $|\pi^*| = k$, we have that $\sum_{v\in V} x_v = k$.

For any $t \in [k]_0$, let S_t be the optimal set of t seed nodes in the non-adaptive influence maximization problem, i.e., such that $OPT_N(G, t) = \mathbb{E}_L(|R(S_t, L)|)$. Let $\psi_{t,L}$ be the random variable denoting the sub-realization of ϕ_L such that $dom(\psi_{t,L}) = S_t$. Let ρ be the random variable equal to node $v \in V$ with probability x_v/k. Observe that the above random variable is well-defined, as $\sum_{v\in V}(x_v/k) = k/k = 1$. For any $t \in [k]$, let $\psi_{\rho,t,L}$ be the random variable denoting the sub-realization of ϕ_L such that $dom(\psi_{\rho,t,L}) = S_{t-1} \cup \{\rho\}$.

[2] We observe that, if $\boldsymbol{y}$ is arbitrary, a deterministic policy π verifying $\mathbb{P}(\pi) = \boldsymbol{y}$ might not exists, and the introduction of randomization solves this issue.

General outline of the proof technique. We observe that $\psi_{\rho,t,L}$ is the partial realization coming from the following *hybrid non-adaptive policy*: initially, we activate the first $t-1$ seed nodes as in the optimal non-adaptive solution guaranteeing an expected influence spread of $OPT_N(G, t-1)$; then, we randomly choose a node v according to random variable ρ and we select v as t-th seed node (if not already selected as seed). We use this hybrid non-adaptive policy as a main tool to obtain an improved upper bound on the adaptivity gap for in-arborescences. In Lemma 1, holding even for general graphs, we relate the hybrid non-adaptive policy and the optimal non-adaptive solution, with the optimal adaptive policy. Lemma 1, together with Lemma 2 (that is similar to Lemma 3.8 in [**?**]), is used in the main proof of the theorem to relate $OPT_N(G, t)$ with $OPT_A(G, k)$ for any $t \in [k]$, and this leads to our upper bound.

The proof structure of Lemma 1 exhibits some similarities with Lemma 6 of **?**] and Lemma 3.3 of **?**], but in their approach, they relate non-adaptive policies based on the Poisson process and multi-linear extensions, with the optimal adaptive policy. One disadvantage of the Poisson process adopted in [**?**] is that the number X of seed nodes selected by the corresponding non-adaptive policy is equal to k under expectation (i.e., $\mathbb{E}(X) = k$), and determining the expected influence spread w.r.t. random variable X has implied a further loss in the final upper bound (see Lemma 3.9 and inequality (21) of Theorem 3.1 in [**?**]). Instead, by using the hybrid-non-adaptive policy, we guarantee that the number of selected seed nodes at each step $t \in [k]$ is exactly equal to t, independently from the considered random execution. This property allow us to avoid the expectations w.r.t. the number of selected seed nodes, and this leads to a further improvement of the resulting upper bound on the adaptivity gap.

Lemma 1. Let G be an arbitrary influence graph. For any $t \in [k]$, and any fixed partial realization ψ of G such that $\mathbb{P}[\psi_{t-1,L} = \psi] > 0$, we have

$$OPT_A(G,k) \leq \sigma(R(\psi)) + k \cdot \mathbb{E}_{L,\rho}\left[f(\psi_{\rho,t,L}) - f(\psi_{t-1,L}) | \psi_{t-1,L} = \psi\right].$$

Proof. We have

$$\begin{aligned}
&k \cdot \mathbb{E}_{L,\rho}\left[f(\psi_{\rho,t,L}) - f(\psi_{t-1,L}) | \psi_{t-1,L} = \psi\right] \\
=&k \cdot \sum_{v \in V} \mathbb{P}[\rho = v] \cdot \Delta(v|\psi) \\
=&k \cdot \sum_{v \in V \setminus R(\psi)} \frac{x_v}{k} \cdot \Delta(v|\psi) \qquad (1)\\
=&\sum_{v \in V \setminus R(\psi)} x_v \cdot \Delta(v|\psi), \qquad (2)
\end{aligned}$$

where (1) holds since $\Delta(v|\psi) = 0$ for any $v \in R(\psi)$.

Let $\boldsymbol{x}' = (x'_1, \dots x'_n)$ be the vector such that $x'_v = 1$ if $v \in R(\psi)$, and $x'_v = x_v$ otherwise. As $x'_v \geq x_v$ for any $v \in V$ we have

$$OPT_A(G,k) \leq OPT_A(G, \boldsymbol{x}) \leq OPT_A(G, \boldsymbol{x}'). \qquad (3)$$

Let π' be the optimal randomized adaptive policy such that $\mathbb{P}(\pi') = \boldsymbol{x}'$. Policy π' selects each node in $R(\psi)$ with probability 1, thus we can assume that such seed nodes are selected at the beginning and that the adaptive policy starts by observing the resulting partial realization. Furthermore, we can assume w.l.o.g. that each seed v selected by π' does not belong to $R(\psi')$, where ψ' denotes the partial realization observed before selecting the new seed v (that is, $R(\psi')$ is the set of nodes reached by the previous influence spread). Indeed, if it was $v \in R(\psi')$, we would have $R(\{v\}, L) \subseteq R(\psi')$, thus the selection of v would not influence any further node. Given $j \in [n]$, let $\Delta'(j)$ denote the expected increment of the influence spread when π' selects the j-th seed node (in order of selection, and without considering in the count the initial seeds of $R(\psi)$); analogously, let $\Delta'(j|v)$ denote the expected increment of the influence spread when π' selects the j-th seed node, conditioned by the fact that the j-th seed is node v.[3] We get

$$
\begin{aligned}
&OPT_A(G, \boldsymbol{x}') \\
=&\sigma(R(\psi)) + \sum_j \Delta'(j) \\
=&\sigma(R(\psi)) + \sum_j \sum_{v \in V \setminus R(\psi)} \mathbb{P}[\text{the } j\text{-th seed node is } v] \cdot \Delta'(j|v) \\
=&\sigma(R(\psi)) + \sum_{v \in V \setminus R(\psi)} \sum_j \mathbb{P}[\text{the } j\text{-th seed node is } v] \cdot \Delta'(j|v) \\
=&\sigma(R(\psi)) + \sum_{v \in V \setminus R(\psi)} \sum_j \mathbb{P}[\text{the } j\text{-th seed node is } v] \cdot \\
&\cdot \mathbb{E}_{\pi'}[\Delta(v|\psi')|v = \pi'(\psi') \text{ for some } \psi' \supseteq \psi \text{ observed at step } j] \\
\leq&\sigma(R(\psi)) + \sum_{v \in V \setminus R(\psi)} \sum_j \mathbb{P}[\text{the } j\text{-th seed node is } v] \cdot \Delta(v|\psi) && (4) \\
=&\sigma(R(\psi)) + \sum_{v \in V \setminus R(\psi)} \mathbb{P}[v \text{ is selected as seed}] \cdot \Delta(v|\psi) \\
=&\sigma(R(\psi)) + \sum_{v \in V \setminus R(\psi)} x'_v \cdot \Delta(v|\psi) \\
=&\sigma(R(\psi)) + \sum_{v \in V \setminus R(\psi)} x_v \cdot \Delta(v|\psi), && (5)
\end{aligned}
$$

where (4) holds since $\Delta(v|\psi') \leq \Delta(v|\psi)$ for any partial realization $\psi' \supseteq \psi$ by adaptive submodularity (Claim 1). By putting together (2), (3), and (5), we get

$$
\begin{aligned}
&\sigma(R(\psi)) + k \cdot \mathbb{E}_{L,\rho}\left[f(\psi_{\rho,t,L}) - f(\psi_{t-1,L})|\psi_{t-1,L} = \psi\right] \\
=&\sigma(R(\psi)) + \sum_{v \in V \setminus R(\psi)} x_v \cdot \Delta(v|\psi) \\
\geq& OPT_A(G, \boldsymbol{x}') \\
\geq& OPT_A(G, k),
\end{aligned}
$$

[3] If an execution of π' requires less than j steps, we assume that the increase of the influence spread at step j (that contributes to the expected values $\Delta'(j)$ and $\Delta'(j|v)$) is null.

thus showing the claim. $\square$

Lemma 2. If the input influence graph G is an in-arborescence, then

$$\sigma(R(\psi_{t-1,L})) \leq f(\psi_{t-1,L}) + OPT_N(G, t-1)$$

for any live-edge graph L and $t \in [k]$.

Proof. Given a subset $U \subseteq V$, let $\partial U := \{u \in U : \exists (u, v) \in E, v \notin U\}$. We have that $\sigma(R(\psi)) \leq |R(\psi)| + \sigma(\partial R(\psi)) = f(\psi) + \sigma(\partial R(\psi))$ for any partial realization ψ. Thus, to show the claim, it suffices to show that $\sigma(\partial R(\psi_{t-1,L})) \leq OPT_N(G, t-1)$. For in-arborescences, we have that $|\partial R(\psi_{t-1,L})| \leq |dom(\psi_{t-1,L})| = t-1$, thus $\sigma(\partial R(\psi_{t-1,L})) \leq OPT_N(G, t-1)$. $\square$

Armed with the above lemmas, we can now prove Theorem 1.

Proof of Theorem 1. For any $t \in [k]$, we have

$$\begin{aligned}
&k \cdot (OPT_N(G,t) - OPT_N(G,t-1)) \\
=&k \cdot (\sigma(S_t) - \sigma(S_{t-1})) \\
=&k \cdot (\mathbb{E}_L[f(\psi_{t,L})] - \mathbb{E}_L[f(\psi_{t-1,L})]) \\
\geq&k \cdot (\mathbb{E}_{L,\rho}[f(\psi_{\rho,t,L})] - \mathbb{E}_L[f(\psi_{t-1,L})]) && (6)\\
=&k \cdot (\mathbb{E}_{L,\rho}[f(\psi_{\rho,t,L})] - \mathbb{E}_{L,\rho}[f(\psi_{t-1,L})]) \\
=&k \cdot \mathbb{E}_{L,\rho}[f(\psi_{\rho,t,L}) - f(\psi_{t-1,L})] \\
=&\mathbb{E}_{\psi_{t-1,L}}\left[k \cdot \mathbb{E}_{L,\rho}[f(\psi_{\rho,t,L}) - f(\psi_{t-1,L}) | \psi_{t-1,L}]\right] \\
\geq&\mathbb{E}_{\psi_{t-1,L}}[OPT_A(G,k) - \sigma(R(\psi_{t-1,L}))] && (7)\\
\geq&\mathbb{E}_{\psi_{t-1,L}}[OPT_A(G,k) - f(\psi_{t-1,L}) - OPT_N(G,t-1)] && (8)\\
=&\mathbb{E}_{\psi_{t-1,L}}[OPT_A(G,k)] - \mathbb{E}_{\psi_{t-1,L}}[f(\psi_{t-1,L})] - \mathbb{E}_{\psi_{t-1,L}}[OPT_N(G,t-1)] \\
=&OPT_A(G,k) - \sigma(S_{t-1}) - OPT_N(G,t-1) \\
=&OPT_A(G,k) - 2 \cdot OPT_N(G,t-1), && (9)
\end{aligned}$$

where (6) holds since $dom(\psi_{t,L})$ is the optimal set of t seed nodes for the non-adaptive influence maximization problem, (7) comes from Lemma 1, and (8) comes from Lemma 2. Thus, by (9), we get $k \cdot (OPT_N(G,t) - OPT_N(G,t-1)) \geq OPT_A(G,k) - 2 \cdot OPT_N(G,t-1)$, that after some manipulations leads to the following recursive relation:

$$OPT_N(G,t) \geq \frac{1}{k} \cdot OPT_A(G,k) + \left(1 - \frac{2}{k}\right) \cdot OPT_N(G,t-1), \ \forall t \in [k]. \tag{10}$$

By applying iteratively (10), we get

$$OPT_N(G,k) \geq \frac{1}{k} \cdot \sum_{t=0}^{k-1} \left(1 - \frac{2}{k}\right)^t \cdot OPT_A(G,k) = \frac{1 - (1 - 2/k)^k}{2} \cdot OPT_A(G,k),$$

that leads to

$$\frac{OPT_A(G,k)}{OPT_N(G,k)} \leq \frac{2}{1 - (1 - 2/k)^k} \leq \frac{2}{1 - e^{-2}} = \frac{2e^2}{e^2 - 1},$$

and this shows the claim. $\square$

4. Adaptivity Gap for General Influence Graphs

In this section, we exhibit upper bounds on the k-adaptivity gap of general graphs. In the following theorem, we first give an upper bound that is linear in the number of seed nodes.

Theorem 2. Given an arbitrary class of influence graphs $\mathcal{G}$ and $k \geq 2$, we get $AG(\mathcal{G}, k) \leq k$.

Proof. Let $G = (V = [n], E, (p_{uv})_{(u,v)\in E})$ be an arbitrary influence graph. Let π^* be an optimal adaptive policy subject to $|\pi^*| = k$, and let $\psi_{t,\pi^*,L}$ be the partial realization observed when the t-th seed node has been selected by Algorithm 1 with policy π^*. Fix $t \in [k]$, a partial realization ψ such that $\mathbb{P}[\psi_{t-1,\pi^*,L} = \psi] > 0$, and let $v = \pi^*(\psi)$ be the node selected by policy π^* under partial realization ψ. We have that

$$\begin{aligned}
&\mathbb{E}_L[f(\psi_{t,\pi^*,L}) - f(\psi_{t-1,\pi^*,L})|\psi_{t-1,\pi^*,L} = \psi] \\
=&\Delta(v|\psi) \\
\leq&\Delta(v|\emptyset) \qquad (11)\\
=&\sigma(\{v\}) \\
\leq&OPT_N(G,1), \qquad (12)
\end{aligned}$$

where (11) holds by adaptive submodularity (Claim 1). Thus, we get

$$\begin{aligned}
OPT_A(G,k) =&\mathbb{E}_L[f(\psi_{k,\pi^*,L})] \\
=&\sum_{t=1}^{k} \mathbb{E}_L[f(\psi_{t,\pi^*,L}) - f(\psi_{t-1,\pi^*,L})] \\
=&\sum_{t=1}^{k} \mathbb{E}_{\psi_{t-1,\pi^*,L}}[\mathbb{E}_L[f(\psi_{t,\pi^*,L}) - f(\psi_{t-1,\pi^*,L})|\psi_{t-1,\pi^*,L}]] \\
\leq&k \cdot \mathbb{E}_{\psi_{t-1,\pi^*,L}}[OPT_N(G,1)] \qquad (13)\\
=&k \cdot OPT_N(G,1) \\
\leq&k \cdot OPT_N(G,k), \qquad (14)
\end{aligned}$$

where (13) comes from (12), and the claim follows. □

In the next theorem we give an upper bound on the adaptivity gap that is sublinear in the number of nodes of the considered graph.

Theorem 3. If $\mathcal{G}$ is the class of influence graphs having at most n nodes, we get $AG(\mathcal{G}) \leq \lceil n^{1/3} \rceil$.

Let $G = (V, E, (p_{uv})_{(u,v)\in E})$ be the input influence graph. To show Theorem 3, we recall the preliminary notations considered for the proof of Theorem 1, and we give a further preliminary lemma.

Lemma 3. Given a set $U \subseteq V$ of cardinality $h \geq k$, we have $\sigma(U) \leq \frac{h}{k} \cdot OPT_N(G, k)$.

Proof. For any $t \in [h]_0$, let $U_t := \emptyset$ if $t = 0$, and $U_t := U_{t-1} \cup \{v_t\}$, where $v_t \in \arg\max_{v \in U \setminus U_{t-1}} (\sigma(U_{t-1} \cup \{v\}) - \sigma(U_{t-1}))$. We have that $\Delta_t := \sigma(U_t) - \sigma(U_{t-1})$ is non-increasing in $t \in [h]$. Indeed, given $t \in [k-1]$, we have that

$$\begin{aligned}
\Delta_{t+1} =& \sigma(U_{t+1}) - \sigma(U_t) \\
=& \sigma(U_t \cup \{v_{t+1}\}) - \sigma(U_t) \\
\leq& \sigma(U_{t-1} \cup \{v_{t+1}\}) - \sigma(U_{t-1}) && (15) \\
\leq& \max_{v \in U \setminus U_{t-1}} (\sigma(U_{t-1} \cup \{v\}) - \sigma(U_{t-1})) \\
=& \sigma(U_{t-1} \cup \{v_t\}) - \sigma(U_{t-1}) \\
=& \Delta_t, && (16)
\end{aligned}$$

where (15) holds since σ is a submodular set-function (see [**?**]). Thus, we necessarily have

$$\begin{aligned}
\frac{\sigma(U)}{h} =& \frac{\sum_{t=1}^{h} \Delta_t}{h} \\
\leq& \frac{\sum_{t=1}^{k} \Delta_t \ (h/k)}{h} && (17) \\
=& \frac{\sum_{t=1}^{k} \Delta_t}{k} \\
=& \frac{\sigma(U_k)}{k} \\
\leq& \frac{OPT_N(G,k)}{k}, && (18)
\end{aligned}$$

where (17) comes from (16). By (18), the claim follows. □

We are now ready to show Theorem 3. In particular, we will first invoke Theorem 2 to show the claim if k is sufficiently small. Then, to cope with the remaining cases, we will use Lemma 1 as in the proof of Theorem 1, and Lemma 3 will play a similar role as Lemma 2 in Theorem 1.

Proof of Theorem 3. We assume w.l.o.g. that $k > \lceil n^{1/3} \rceil$ and that $OPT_N(G,k) < (\lceil n^{1/3} \rceil)^2$. Indeed, if $k \leq \lceil n^{1/3} \rceil$, by Theorem 2 the claim holds, and if $OPT_N(G,k) \geq (\lceil n^{1/3} \rceil)^2$, then $\frac{OPT_A(G,k)}{OPT_N(G,k)} \leq \frac{|V|}{OPT_N(G,k)} \leq \frac{n}{(\lceil n^{1/3} \rceil)^2} \leq \lceil n^{1/3} \rceil$, and the claim holds as well. For any $t \in [k]$, we have

$$\begin{aligned}
& k \cdot (OPT_N(G,t) - OPT_N(G,t-1)) \\
\geq& k \cdot (\mathbb{E}_{L,\rho}[f(\psi_{\rho,t,L})] - \mathbb{E}_{L,\rho}[f(\psi_{t-1,L})]) && (19) \\
=& \mathbb{E}_{\psi_{t-1,L}} \left[k \cdot \mathbb{E}_{L,\rho}[f(\psi_{\rho,t,L}) - f(\psi_{t-1,L}) | \psi_{t-1,L}] \right] && (20) \\
\geq& \mathbb{E}_{\psi_{t-1,L}} [OPT_A(G,k) - \sigma(R(\psi_{t-1,L}))] && (21) \\
=& \mathbb{E}_{\psi_{t-1,L}} [OPT_A(G,k)] - \mathbb{E}_{\psi_{t-1,L}} [\sigma(R(\psi_{t-1,L}))] \\
\geq& \mathbb{E}_{\psi_{t-1,L}} [OPT_A(G,k)] - \mathbb{E}_{\psi_{k,L}} [\sigma(R(\psi_{k,L}))]
\end{aligned}$$

$$\geq \mathbb{E}_{\psi_{t-1,L}}[OPT_A(G,k)] - \mathbb{E}_{\psi_{k,L}}\left[\frac{|R(\psi_{k,L})|}{k} \cdot OPT_N(G,k)\right] \tag{22}$$
$$= OPT_A(G,k) - \frac{\mathbb{E}_{\psi_{k,L}}[|R(\psi_{k,L})|]}{k} \cdot OPT_N(G,k)$$
$$\geq OPT_A(G,k) - \frac{\mathbb{E}_{\psi_{k,L}}[|R(\psi_{k,L})|]}{\lceil n^{1/3}\rceil + 1} \cdot ((\lceil n^{1/3}\rceil)^2 - 1) \tag{23}$$
$$= OPT_A(G,k) - (\lceil n^{1/3}\rceil - 1) \cdot \mathbb{E}_{\psi_{k,L}}[|R(\psi_{k,L})|]$$
$$= OPT_A(G,k) - (\lceil n^{1/3}\rceil - 1) \cdot OPT_N(G,k), \tag{24}$$

where (19) and (20) are obtained similarly as in Theorem 1, (21) comes from Lemma 1, (22) comes from Lemma 3, and (23) comes from the hypothesis $k > \lceil n^{1/3}\rceil$ and $OPT_N(G,k) < (\lceil n^{1/3}\rceil)^2$. By (24), we get $OPT_N(G,t) - OPT_N(G,t-1) \geq (OPT_A(G,k) - (\lceil n^{1/3}\rceil - 1) \cdot OPT_N(G,k))/k$ for any $t \in [k]$, and by summing such inequality over all $t \in [k]$, we get

$$OPT_N(G,k)$$
$$= \sum_{t=1}^{k}(OPT_N(G,t) - OPT_N(G,t-1))$$
$$\geq \sum_{t=1}^{k} \frac{OPT_A(G,k) - (\lceil n^{1/3}\rceil - 1) \cdot OPT_N(G,k)}{k}$$
$$= OPT_A(G,k) - (\lceil n^{1/3}\rceil - 1) \cdot OPT_N(G,k). \tag{25}$$

Finally, (25) implies that $OPT_A(G,k) \leq \lceil n^{1/3}\rceil \cdot OPT_N(G,k)$, and this shows the claim. □

By definition of adaptivity gap, we have that any α-approximation to the non-adaptive optimal solution is also an $\alpha/AG(\mathcal{G})$ approximation to the adaptive optimum, where $\mathcal{G}$ denotes the class of influence graphs with at most n nodes and $AG(\mathcal{G})$ denotes its adaptivity gap. The *non-adaptive greedy algorithm* for IM selects k seeds in k steps, and at each step it selects, as new seed, a node that approximately maximizes the expected increment of the influence spread without observing any feedback; **?**] showed that the non-adaptive greedy algorithm guarantees a $\left(1 - \frac{1}{e} - \epsilon\right)$-approximation to the non-adaptive optimum. Thus, the following corollary of Theorem 3 immediately holds:

Corollary 1. *The non-adaptive greedy algorithm guarantees an expected influence spread of at least* $\left(1 - \frac{1}{e} - \epsilon\right)/\lceil n^{1/3}\rceil = \Omega(1/n^{1/3})$ *the adaptive optimum.*

5. Adaptivity Gap for Other Influence Graphs

In this section, as further application of the approaches exploited in Theorems 1 and 3, we introduce the classes of (β,γ)-bounded-activation graphs and α-bounded-degree graphs, that include several interesting graph topologies, and we get upper bounds on the adaptivity gap for such graphs.

5.1. (β,γ)-bounded-activation graphs.

Let $N(u) := \{v \in V : (u,v) \in E\}$ denote the set of out-neighbors of node u. Given an integer $\beta \geq 0$ and a real value $\gamma \in [0,1)$, an influence graph $G = (V, E, (p_{uv})_{(u,v)\in E})$ is a *(β,γ)-bounded-activation graph* if there exists $\hat{S} \subseteq V$ with $|\hat{S}| \leq \beta$ such that $\sum_{v\in N(u)} p_{u,v} \leq \gamma$ for any $u \in V \setminus \hat{S}$. Informally, the class of (β,γ)-bounded-activation graphs coincides with all the influence graphs such that all nodes but β influence in expectation at most γ neighbors each. We observe that (β,γ)-bounded-activation graphs are well-motivated by social networks in which most nodes have a limited power of influence. By doing a parallelism with epidemic processes in social networks, the parameter γ is analogue to the *basic reproduction number* [**? ?**] (often denoted as R_0), that measures the transmission potential of a disease and is defined as the average number of people who can be directly infected by an already infected person. Indeed, in our influence maximization scenario, by setting $\gamma \in [0,1)$ we are assuming that the basic reproduction number is lower than one for all people, except at most β.

In the following theorem, whose proof is partially based on Theorem 2 and Lemma 3, we provide an upper bound on the adaptivity gap of (β,γ)-bounded-activation graphs, for any integer $\beta \geq 0$ and $\gamma \in [0,1)$.

Theorem 4. Given an integer $\beta \geq 0$ and $\gamma \in [0,1)$, let $\mathcal{G}$ be the class of (β,γ)-bounded-activation graphs. Then

$$\begin{aligned} AG(\mathcal{G},k) &\leq \min\left\{k, \frac{\max\{\beta,k\}\cdot\min\{1,\beta\}+\frac{k}{1-\gamma}}{k}\right\} \qquad (26)\\ &\leq \max\left\{\frac{\sqrt{\left(\frac{1}{1-\gamma}\right)^2+4\beta}+\frac{1}{1-\gamma}}{2}, \min\{1,\beta\}+\frac{1}{1-\gamma}\right\}\\ &\leq \sqrt{\beta}+\frac{1}{1-\gamma}. \end{aligned}$$

for any $k \geq 2$.

Let $G = (V = [n], E, (p_{uv})_{(u,v)\in E})$ be a (β,γ)-bounded-activation graph. To show the claim of Theorem 4, we recall the notations from Theorem 1 and we give some preliminary lemmas.

Let $\hat{S}$ be the set of nodes such that $|\hat{S}| \leq \beta$ and $\sum_{v\in N(u)} p_{u,v} \leq \gamma$ for any $u \in V \setminus \hat{S}$. Let $G \setminus \hat{S}$ be the graph obtained from G by removing the nodes in $\hat{S}$ and their adjacent edges, and let $OPT_N(G\setminus\hat{S}, 1)$ denote the optimal non-adaptive influence spread $OPT_N(G \setminus \hat{S}, 1)$ achieved by a unique seed in graph $G \setminus \hat{S}$. In the following lemma, we provide an upper bound for the optimal adaptive influence spread $OPT_A(G,k)$ in G.

Lemma 4. We have that $OPT_A(G,k) \leq \sigma(\hat{S}) + k \cdot OPT_N(G \setminus \hat{S}, 1)$.

Proof. Let $\hat{\pi}$ be an optimal adaptive policy that first selects the nodes in $\hat{S}$ and then adaptively selects k nodes. By construction, we have that

$$OPT_A(G,k) \leq \sigma(\hat{\pi}). \qquad (27)$$

Let $\psi_{0,\hat{\pi},L}$ denote the partial realization observed by $\hat{\pi}$ after the selection of $\hat{S}$, and let $\psi_{t,\hat{\pi},L}$ be the partial realization observed after the selection of the t-th seed node in $V \setminus \hat{S}$. By exploiting the adaptive submodularity (Claim 1) similarly as in the proof of Theorem 2, one can easily show that

$$\mathbb{E}_L[f(\psi_{t,\hat{\pi},L}) - f(\psi_{t-1,\hat{\pi},L})|\psi_{t-1,\hat{\pi},L} = \psi] \leq OPT_N(G \setminus \hat{S}, 1) \tag{28}$$

for any fixed partial realization ψ (with $\mathbb{P}[\psi_{t,\hat{\pi},L}) = \psi] > 0$) and any $t \in [k]$. Thus, we get

$$\begin{aligned}
\sigma(\hat{\pi}) =& \mathbb{E}_L[f(\psi_{k,\hat{\pi},L})] \\
=& \mathbb{E}_L[f(\psi_{0,\hat{\pi},L})] + \sum_{t=1}^{k} \mathbb{E}_L[f(\psi_{t,\hat{\pi},L}) - f(\psi_{t-1,\hat{\pi},L})] \\
=& \mathbb{E}_L[f(\psi_{0,\hat{\pi},L})] + \sum_{t=1}^{k} \mathbb{E}_{\psi_{t-1,\hat{\pi},L}}[\mathbb{E}_L[f(\psi_{t,\hat{\pi},L}) - f(\psi_{t-1,\hat{\pi},L})|\psi_{t-1,\hat{\pi},L}]] \\
\leq& \mathbb{E}_L[f(\psi_{0,\hat{\pi},L})] + k \cdot \mathbb{E}_{\psi_{t-1,\hat{\pi},L}}[OPT_N(G \setminus \hat{S}, 1)] \qquad (29) \\
=& \mathbb{E}_L[f(\psi_{0,\hat{\pi},L})] + k \cdot OPT_N(G \setminus \hat{S}, 1) \\
=& \sigma(\hat{S}) + k \cdot OPT_N(G \setminus \hat{S}, 1), \qquad (30)
\end{aligned}$$

where (29) comes from (28). By putting (27) and (30) together, the claim follows. □

In the following lemma, we provide an upper bound for $OPT_N(G \setminus \hat{S}, 1)$ in terms of parameter γ.

Lemma 5. We have that $OPT_N(G \setminus \hat{S}, 1) \leq \frac{1}{1-\gamma}$.

Proof. Let v_0 be the node that maximizes the expected influence spread in $G \setminus \hat{S}$ when selected as unique seed. For any live-edge graph L and $j \in [n-1]_0$, let $A(j, L)$ denote the set of nodes activated at the j-th round of diffusion when v_0 is the initial seed node of $G \setminus \hat{S}$, i.e., $A(0, L) = \{v_0\}$ and $A(j, L)$ is the set of neighbors of $A(j-1, L)$ activated by $A(j-1, L)$. We can easily observe that

$$\begin{aligned}
&\mathbb{E}_L[|A(j, L)|] \\
=& \sum_{v \in G \setminus \hat{S}} \mathbb{P}[v \in A(j, L)|] \\
\leq& \sum_{v \in G \setminus \hat{S}} \sum_{\substack{P = (v_0, v_1, \ldots, v_j := v): \\ P \text{ is a path of } G \setminus \hat{S} \text{ from } v_0 \text{ to } v \text{ having } j \text{ edges}}} p_{v_0,v_1} \cdot p_{v_1,v_2} \cdots p_{v_{j-2},v_{j-1}} \cdot p_{v_{j-1},v_j} \\
=& \sum_{\substack{P=(v_0,v_1,\ldots,v_j): \\ P \text{ is a path of } G \setminus \hat{S} \text{ from } v_0 \text{ having } j \text{ edges}}} p_{v_0,v_1} \cdot p_{v_1,v_2} \cdots p_{v_{j-2},v_{j-1}} \cdot p_{v_{j-1},v_j} \\
=& \sum_{v_1 \in N(v_0) \cap \hat{S}} p_{v_0,v_1} \sum_{v_2 \in N(v_1) \cap \hat{S}} p_{v_1,v_2} \sum_{v_3 \in N(v_2) \cap \hat{S}} \cdots p_{v_{j-2},v_{j-1}} \sum_{v_j \in N(v_{j-1}) \cap \hat{S}} p_{v_{j-1},v_j}
\end{aligned}$$

for any $j \in [n-1]$. Thus

$$
\begin{aligned}
& OPT_N(G \setminus \hat{S}, 1) \\
=& \mathbb{E}_L \left[\sum_{j=0}^{n-1} |A(j, L)| \right] \\
=& \mathbb{E}_L [|A(0, L)|] + \sum_{j=1}^{n-1} \mathbb{E}_L \left[|A(j, L)| \right] \\
\leq & 1 + \sum_{j=1}^{n-1} \underbrace{\sum_{v_1 \in N(v_0) \cap \hat{S}} p_{v_0, v_1} \underbrace{\sum_{v_2 \in N(v_1) \cap \hat{S}} \cdots \overbrace{\sum_{v_{j-1} \in N(v_{j-2}) \cap \hat{S}} p_{v_{j-2}, v_{j-1}} \overbrace{\sum_{v_j \in N(v_{j-1}) \cap \hat{S}} p_{v_{j-1}, v_j}}^{\leq \gamma}}^{\leq \gamma^2}}_{\leq \gamma^{j-1}}}_{\leq \gamma^j} \\
\leq & 1 + \sum_{j=1}^{n-1} \gamma^j \leq \sum_{j=0}^{\infty} \gamma^j = \frac{1}{1-\gamma},
\end{aligned}
$$

and this shows the claim. $\square$

We are ready to show Theorem 4.

Proof of Theorem 4. We first show the upper bound in (26). By Theorem 2, we have that k is an upper bound on the k-adaptivity gap. Thus, it is sufficient to show that $\frac{\max\{\beta, k\} \cdot \min\{1, \beta\} + \frac{k}{1-\gamma}}{k}$ is an upper bound on the k-adaptivity gap. We have that

$$
\begin{aligned}
\frac{OPT_A(G, k)}{OPT_N(G, k)} &\leq \frac{\sigma(\hat{S}) + OPT_N(G \setminus \hat{S}, 1) \cdot k}{OPT_N(G, k)} && (31) \\
&\leq \frac{\sigma(\hat{S}) + \frac{k}{1-\gamma}}{OPT_N(G, k)} && (32) \\
&\leq \frac{OPT_N(G, \beta) + \frac{k}{1-\gamma}}{OPT_N(G, k)}, && (33)
\end{aligned}
$$

where (31) and (32) follows from Lemmas 4 and 5, respectively.

If $\beta = 0$, by using $OPT_N(G, \beta) = 0$ and $OPT_N(G, k) \geq k$ in (33), we get that (33) is at most $\frac{\frac{k}{1-\gamma}}{k} = \frac{\max\{\beta, k\} \cdot \min\{1, \beta\} + \frac{k}{1-\gamma}}{k}$, and this shows (26).

If $1 \leq \beta \leq k$, by continuing from (33), we get

$$
\frac{OPT_A(G, k)}{OPT_N(G, k)} \leq \frac{OPT_N(G, \beta) + \frac{k}{1-\gamma}}{OPT_N(G, k)}
$$

$$
\begin{aligned}
&\le \frac{OPT_N(G,k) + \frac{k}{1-\gamma}}{OPT_N(G,k)} \\
&\le \frac{k + \frac{k}{1-\gamma}}{k}, \qquad (34)
\end{aligned}
$$

where (34) holds since $OPT_N(G,k) \ge k$. As (34) is equal to $\frac{\max\{\beta,k\}\cdot\min\{1,\beta\}+\frac{k}{1-\gamma}}{k}$, inequality (26) holds if $1 \le \beta \le k$.

Finally, if $\beta > k$, by continuing from (33), we get

$$
\begin{aligned}
\frac{OPT_A(G,k)}{OPT_N(G,k)} &\le \frac{OPT_N(G,\beta) + \frac{k}{1-\gamma}}{OPT_N(G,k)} \\
&\le \frac{\frac{\beta}{k}\cdot OPT_N(G,k) + \frac{k}{1-\gamma}}{OPT_N(G,k)} \qquad (35) \\
&\le \frac{\frac{\beta}{k}\cdot k + \frac{k}{1-\gamma}}{k}, \qquad (36)
\end{aligned}
$$

where (35) follows from Lemma 3 and (36) holds since $OPT_N(G,k) \ge k$. As (36) is equal to $\frac{\max\{\beta,k\}\cdot\min\{1,\beta\}+\frac{k}{1-\gamma}}{k}$, inequality (26) holds if $\beta > k$.

We conclude that $\min\left\{k, \frac{\max\{\beta,k\}\cdot\min\{1,\beta\}+\frac{k}{1-\gamma}}{k}\right\}$ is an upper bound on the k-adaptivity gap. Furthermore, as k and $\frac{\max\{\beta,k\}\cdot\min\{1,\beta\}+\frac{k}{1-\gamma}}{k}$ are respectively increasing and non-increasing in k (for any fixed integer $\beta \ge 0$), the real value k such that the two quantities are equal is a further upper bound on the adaptivity gap, and such value is

$$
\begin{aligned}
k &= \max\left\{\frac{\sqrt{\left(\frac{1}{1-\gamma}\right)^2 + 4\beta} + \frac{1}{1-\gamma}}{2}, \min\{1,\beta\} + \frac{1}{1-\gamma}\right\} \\
&\le \max\left\{\frac{\sqrt{\left(\frac{1}{1-\gamma}\right)^2} + \sqrt{4\beta} + \frac{1}{1-\gamma}}{2}, \min\{1,\beta\} + \frac{1}{1-\gamma}\right\} \\
&= \max\left\{\sqrt{\beta} + \frac{1}{1-\gamma}, \min\{1,\beta\} + \frac{1}{1-\gamma}\right\} \\
&= \sqrt{\beta} + \frac{1}{1-\gamma}.
\end{aligned}
$$

Thus, $\max\left\{\frac{\sqrt{\left(\frac{1}{1-\gamma}\right)^2+4\beta}+\frac{1}{1-\gamma}}{2}, \min\{1,\beta\} + \frac{1}{1-\gamma}\right\}$ and $\sqrt{\beta} + \frac{1}{1-\gamma}$ are further upper bounds on the adaptivity gap.

□

5.2. α-bounded-degree graphs.

In the following, when we refer to undirected influence graphs, we assume that, for any undirected edge $\{u, v\}$, there are two directed edges (u, v) and (v, u) having respectively two (possibly) distinct probabilities p_{uv} and p_{vu}.

Given an undirected graph $G = (V, E)$ and a node $v \in V$, let $deg_v(G)$ denote the degree of node v in graph G. Given an integer $\alpha \geq 0$, an influence graph $G = (V, E, (p_{uv})_{(u,v)\in E})$ is an *α-bounded-degree graph* if it is undirected and $\sum_{v\in V:deg_v(G)>2} deg_v(G) \leq \alpha$, i.e., the sum all the node degrees higher than 2 is at most α; we observe that the definition of α-bounded degree graphs does not depends on the influence probabilities, but on the graph topology only.

Example 1. Given an undirected graph G, a simple subpath P (resp. cycle C) of G is *standard* if all the nodes of P but the first and the last one (resp. all the nodes of C) have degree 2. The *standard contraction* of G is the multigraph G' obtained by replacing each standard simple subpath $P = (v_1, \ldots, v_t)$ of G with an edge connecting v_1 and v_t, and by deleting all the standard cycles. There are several interesting classes of α-bounded-degree graphs characterized by the topological structure of their standard contraction:

- The set of 0-bounded-degree graphs is made of all the graphs G such that each connected component of G is either an undirected path or an undirected cycle; equivalently, the set of 0-bounded-degree graphs is made of all the graphs G whose standard contraction is the (possibly empty) union of several disconnected edges.
- If the standard contraction of a graph G is homeomorphic to a star with h edges, then G is a h-bounded-degree graph.
- If the standard contraction of a graph G is homeomorphic to a parallel-link multigraph with h edges (that is, a multigraph with h edges connecting two nodes), then G is a $2h$-bounded-degree graph.
- If the standard contraction of a graph G is homeomorphic to a cycle with h chords, then G is a $6h$-bounded-degree graph.
- If the standard contraction of a graph G is homeomorphic to a clique with h nodes, then G is a $h(h-1)$-bounded-degree graph.

Also in light of the above examples, we have that the study of α-bounded-degree-graphs is mainly of theoretical interest, as most of the networks belonging to this class are interesting from a geometric point of view (e.g., cycles, paths, standard contractions of parallel-link graphs) but they are unlikely to appear as real-world social networks. Anyway, such graphs constitute an interesting application of our techniques when the node degrees are somewhat constrained, and could probably constitute a further starting point to study more realistic network topologies having similar node degrees constraints (e.g., power-law networks [**?**]).

In the following theorem, we provide an upper bound on the adaptivity gap of α-bounded-degree graphs for any $\alpha \geq 0$ (the proof is deferred to the appendix).

Theorem 5. Given $\alpha \geq 0$, let $\mathcal{G}$ be the class of α-bounded-degree graphs. Then

$$AG(\mathcal{G}, k) \leq \min\left\{k, \frac{\alpha}{k} + 2 + \frac{1}{1-(1-1/k)^k}\right\} \leq \frac{\sqrt{4(e-1)^2\alpha + (3e-2)^2} + 3e - 2}{2(e-1)}$$

for any $k \geq 2$, i.e., $AG(\mathcal{G}) \leq \sqrt{\alpha} + O(1)$.

For the particular case of 0-bounded-degree graphs, the following theorem provides a better upper bound on the adaptivity gap (the proof is deferred to the appendix).

Theorem 6. Let $\mathcal{G}$ be the class of 0-bounded-degree graphs. Then

$$AG(\mathcal{G}, k) \leq \min\left\{k, \frac{3}{1-(\max\{0, 1-3/k\})^k}\right\} \leq \frac{3e^3}{e^3-1} \approx 3.16,$$

for any $k \geq 2$.

6. Future Works

The first problem that is left open by our results is the gap between the constant lower bound provided by **?**] and our upper bound on the adaptivity gap for general graphs. Besides trying to lower the upper bound, a possible direction could be that of increasing the lower bound by finding instances with a non constant adaptivity gap. Since the lower bound given in [**?**] holds even when the graph is a directed path, one direction could be to exploit different graph topologies.

Although in this work we have improved the upper bound on the adaptivity gap of in-arborescence, there is still a gap between upper and lower bound, thus another open problem is to close it. It would be also interesting to find better bounds on the adaptivity gap of other graph classes, like e.g. out-arborescences.

A further interesting research direction is to study the adaptivity gap of some graph classes modelling real-world networks, both theoretically and experimentally.

Finally, most of the work on adaptive IM has been done for the independent cascade model, and other diffusion models (e.g., the linear threshold and the triggering models) have been less investigated. We observe that in many diffusion models different from the independent cascade (e.g., the linear treshold and the triggering models) the objective function is not adaptive submodular under both myopic and full-adoption feedbacks and the standard analysis of the greedy approach does not guarantee an efficient approximation. For the general triggering model, **?**] overcome this problem by exploiting submodularity ratio, but constant bounds on both the adaptivity gap and the approximation ratio are guaranteed for bipartite graphs only, and the study of other graph topologies is still open. The techniques introduced in this paper to relate adaptive policies with non-adaptive ones might be useful to find better upper bounds on the adaptivity gaps in different diffusion/feedback models[4] or in different graph classes.

[4] An interesting diffusion mdoel that could be investigated in terms of adaptivity gap is the batch model [**? ?**] discussed in the related work, in which a batch of nodes is selected at each step, and then the influence spread is observed. We strongly believe that each batch of nodes could be treated as a single node, when applying our approach to relate a non-adaptive policy with the optimal adaptive one. Then, by applying the adaptive submodularity, we think that our approach would allow to show upper bounds on the adaptivity gap that are similar to those obtained in this work.

7. Acknowledgments

This work was partially supported by the Italian MIUR PRIN 2017 Project ALGADIMAR "Algorithms, Games, and Digital Markets", by GNCS-INdAM and by the PON R&I 2014-2020 Project TEBAKA "Sistema per acquisizione conoscenze di base del territorio".

Appendix A. Missing Proofs

Appendix A.1. Proof of Theorem 5

Let $G = (V = [n], E, (p_{uv})_{(u,v)\in E})$ be an α-bounded-degree graph, and we recall the preliminary notations from Theorem 1. The proof of Theorem 5 is a non-trivial generalization of Theorem 1. In particular, the proof resorts to Theorem 2 to get the upper bound of k and, by following the approach of Theorem 1, the following technical lemma is used in place of Lemma 2 to get the final upper bound.

Lemma 6. When the input influence graph G is an α-bounded-degree graph with $\alpha \geq 0$, we have that $\sigma(R(\psi_{t-1,L})) \leq f(\psi_{t-1,L}) + \left(\frac{\alpha}{k} + 2\right) \cdot OPT_N(G,k)$ for any $t \in [k]$ and live-edge graph L.

Proof. Given a subset $U \subseteq V$, let $\partial U := \{u \in U : \exists (u,v) \in E, v \notin U\}$. We have that $\sigma(R(\psi)) \leq |R(\psi)| + \sigma(\partial R(\psi)) = f(\psi) + \sigma(\partial R(\psi))$ for any partial realization ψ. Thus, to show the claim, it suffices to show that

$$\sigma(\partial R(\psi_{t-1,L})) \leq \left(\frac{\alpha}{k} + 2\right) \cdot OPT_N(G,k).$$

Let $U \subseteq V$ such that U has at most k connected components. Let A be the set of connected components containing at least one node of degree higher than 2, and let B be the set of the remaining components, i.e., containing nodes with degree in $[2]_0$ only. By definition of A and B, we necessarily have that $|\partial A| \leq \sum_{v\in V: deg_v(G)>2} deg_v(G) \leq \alpha$ and $|\partial B| \leq 2k$. Thus $|\partial U| \leq |\partial A| + |\partial B| \leq \alpha + 2k$, and the next claim follows.

Claim 2. Given a subset $U \subseteq V$ made of at most k connected components, then $|\partial U| \leq \alpha + 2k$.

Now, we have that

$$\begin{aligned}\sigma(\partial R(\psi_{t-1,L})) \leq & \sigma(\partial R(\psi_{k,L})) \\ \leq & \frac{|\partial R(\psi_{k,L})|}{k} \cdot OPT_N(G,k) \qquad \text{(A.1)} \\ \leq & \frac{\alpha + 2k}{k} \cdot OPT_N(G,k), \qquad \text{(A.2)}\end{aligned}$$

where (A.1) comes from Lemma 3, and (A.2) holds since $R(\psi_{k,L})$ contains at most k connected components and because of Claim 2. Thus, by (A.2), the claim of the lemma follows. □

We can now prove Theorem 5.

Proof of Theorem 5. For any $t \in [k]$, we have

$$
\begin{aligned}
& k \cdot (OPT_N(G,t) - OPT_N(G,t-1)) \\
\geq & k \cdot (\mathbb{E}_{L,\rho}[f(\psi_{\rho,t,L})] - \mathbb{E}_{L,\rho}[f(\psi_{t-1,L})]) && \text{(A.3)} \\
= & \mathbb{E}_{\psi_{t-1,L}}\left[k \cdot \mathbb{E}_{L,\rho}[f(\psi_{\rho,t,L}) - f(\psi_{t-1,L}) | \psi_{t-1,L}]\right] && \text{(A.4)} \\
\geq & \mathbb{E}_{\psi_{t-1,L}}[OPT_A(G,k) - \sigma(R(\psi_{t-1,L}))] && \text{(A.5)} \\
\geq & \mathbb{E}_{\psi_{t-1,L}}\left[OPT_A(G,k) - f(\psi_{t-1,L}) - \left(\frac{\alpha}{k} + 2\right) OPT_N(G,k)\right] && \text{(A.6)} \\
= & \mathbb{E}_{\psi_{t-1,L}}[OPT_A(G,k)] - \mathbb{E}_{\psi_{t-1,L}}[f(\psi_{t-1,L})] - \left(\frac{\alpha}{k} + 2\right) \cdot \mathbb{E}_{\psi_{t-1,L}}[OPT_N(G,k)] \\
= & OPT_A(G,k) - \sigma(S_{t-1}) - \left(\frac{\alpha}{k} + 2\right) \cdot OPT_N(G,k) \\
= & OPT_A(G,k) - \left(\frac{\alpha}{k} + 2\right) \cdot OPT_N(G,k) - OPT_N(G,t-1), && \text{(A.7)}
\end{aligned}
$$

where (A.3) and (A.4) are obtained similarly as in Theorem 1, (A.5) comes from Lemma 1 and (A.6) comes from Lemma 6. Thus, by (A.7), we get the following recursive relation:

$$
OPT_N(G,t) \geq \frac{1}{k}\left(OPT_A(G,k) - \left(\frac{\alpha}{k} + 2\right) OPT_N(G,k)\right) + \left(1 - \frac{1}{k}\right) OPT_N(G,t-1), \tag{A.8}
$$

for any $t \in [k]$. By applying iteratively (A.8), we get

$$
\begin{aligned}
OPT_N(G,k) \geq & \frac{1}{k} \cdot \left(OPT_A(G,k) - \left(\frac{\alpha}{k} + 2\right) \cdot OPT_N(G,k)\right) \cdot \sum_{t=0}^{k-1} \left(1 - \frac{1}{k}\right)^j \\
= & \left(OPT_A(G,k) - \left(\frac{\alpha}{k} + 2\right) \cdot OPT_N(G,k)\right) \cdot \left(1 - \left(1 - \frac{1}{k}\right)^k\right);
\end{aligned}
$$

then, by manipulating inequality

$$
OPT_N(G,k) \geq \left(OPT_A(G,k) - \left(\frac{\alpha}{k} + 2\right) \cdot OPT_N(G,k)\right) \cdot \left(1 - \left(1 - \frac{1}{k}\right)^k\right),
$$

we get

$$
\frac{OPT_A(G,k)}{OPT_N(G,k)} \leq \frac{\alpha}{k} + 2 + \frac{1}{1 - (1 - 1/k)^k} \leq \frac{\alpha}{k} + 2 + \frac{1}{1 - e^{-1}} \,. \tag{A.9}
$$

By Theorem 2, we have that $\frac{OPT_A(G,k)}{OPT_N(G,k)} \leq k$, thus, by (A.9), we get

$$
\begin{aligned}
\frac{OPT_A(G,k)}{OPT_N(G,k)} & \leq \min\left\{k, \frac{\alpha}{k} + 2 + \frac{1}{1 - (1 - 1/k)^k}\right\} \\
& \leq \min\left\{k, \frac{\alpha}{k} + 2 + \frac{1}{1 - e^{-1}}\right\} \\
& \leq \frac{\sqrt{4(e-1)^2\alpha + (3e-2)^2} + 3e - 2}{2(e-1)}, && \text{(A.10)}
\end{aligned}
$$

where (A.10) is equal to the real value of $k \geq 0$ such that $k = \frac{\alpha}{k} + 2 + \frac{1}{1-e^{-1}}$. By (A.10) the claim follows. $\square$

Appendix A.2. Proof of Theorem 6

The proof of Theorem 6 is similar to that of Theorem 1. Let $G = (V = [n], E, (p_{uv})_{(u,v)\in E})$ be a 0-bounded-degree graph. We recall the notation from Theorem 1 and we give the following preliminary lemma, whose proof is analogue to that of Lemma 2.

Lemma 7. When the input influence graph G is a 0-bounded-degree graph, we have

$$\sigma(R(\psi_{t-1,L})) \leq f(\psi_{t-1,L}) + 2 \cdot OPT_N(G, t-1), \tag{A.11}$$

for any $t \in [k]$ and live-edge graph L.

Proof. As in Lemma 2, we show that $\sigma(\partial R(\psi_{t-1,L})) \leq 2 \cdot OPT_N(G, t-1)$. First of all, we assume that $t \geq 2$, otherwise $\sigma(R(\psi_{t-1,L}))$ and the claim holds. By Lemma 3, we have that $\sigma(\partial R(\psi_{t-1,L})) \leq \frac{|\partial R(\psi_{t-1,L})|}{t-1} \cdot OPT_N(G, t-1)$. As G is a 0-bounded-degree graph, we have that $|\partial R(\psi_{t-1,L})| \leq 2(t-1)$. By considering the above inequalities, we get $\sigma(\partial R(\psi_{t-1,L})) \leq \frac{|\partial R(\psi_{t-1,L})|}{t-1} \cdot OPT_N(G, t-1) \leq \frac{2(t-1)}{t-1} \cdot OPT_N(G, t-1) = 2 \cdot OPT_N(G, t-1)$, and the claim follows. $\square$

We are ready to show Theorem 6.

Proof of Theorem 6. By Theorem 2, we have that k is an upper bound on the k-adaptivity gap, thus it is sufficient showing that $\frac{3}{1-(\max\{0,1-3/k\})^k}$ is a further upper bound. If $k \leq 3$ the claim trivially holds, since k is an upper bound on the k-adaptivity gap. Then, we assume that $k > 3$, and it is sufficient showing that $\frac{3}{1-(1-3/k)^k}$ is an upper bound. For any $t \in [k]$, we have

$$\begin{aligned}
& k \cdot (OPT_N(G, t) - OPT_N(G, t-1)) \\
\geq & k \cdot (\mathbb{E}_{L,\rho}[f(\psi_{\rho,t,L})] - \mathbb{E}_L[f(\psi_{t-1,L})]) && \text{(A.12)} \\
= & \mathbb{E}_{\psi_{t-1,L}}\left[k \cdot \mathbb{E}_{L,\rho}[f(\psi_{\rho,t,L}) - f(\psi_{t-1,L}) | \psi_{t-1,L}]\right] && \text{(A.13)} \\
\geq & \mathbb{E}_{\psi_{t-1,L}}[OPT_A(G, k) - \sigma(R(\psi_{t-1,L}))] && \text{(A.14)} \\
\geq & \mathbb{E}_{\psi_{t-1,L}}[OPT_A(G, k) - f(\psi_{t-1,L}) - 2 \cdot OPT_N(G, t-1)] && \text{(A.15)} \\
= & \mathbb{E}_{\psi_{t-1,L}}[OPT_A(G, k)] - \mathbb{E}_{\psi_{t-1,L}}[f(\psi_{t-1,L})] - 2 \cdot \mathbb{E}_{\psi_{t-1,L}}[OPT_N(G, t-1)] \\
= & OPT_A(G, k) - \sigma(S_{t-1}) - 2 \cdot OPT_N(G, t-1) \\
= & OPT_A(G, k) - 3 \cdot OPT_N(G, t-1), && \text{(A.16)}
\end{aligned}$$

where (A.12) and (A.13) are obtained similarly as in Theorem 1, (A.14) comes from Lemma 1 and (A.15) comes from Lemma 7. Thus, by (A.16), we get $k \cdot (OPT_N(G, t) - OPT_N(G, t-1)) \geq OPT_A(G, k) - 3 \cdot OPT_N(G, t-1)$, that after some manipulations leads to the following recursive relation:

$$OPT_N(G, t) \geq \frac{1}{k} \cdot OPT_A(G, k) + \left(1 - \frac{3}{k}\right) \cdot OPT_N(G, t-1), \ \forall t \in [k]. \tag{A.17}$$

By applying iteratively (A.17), we get

$$OPT_N(G,k) \geq \frac{1}{k} \cdot \sum_{t=0}^{k-1} \left(1 - \frac{3}{k}\right)^t \cdot OPT_A(G,k) = \frac{1-(1-3/k)^k}{3} \cdot OPT_A(G,k),$$

that leads to

$$\frac{OPT_A(G,k)}{OPT_N(G,k)} \leq \frac{3}{1-(1-3/k)^k} \leq \frac{3}{1-e^{-3}}, \tag{A.18}$$

and this shows the claim. □

References

[] Adamczyk, M., Grandoni, F., Mukherjee, J., 2015. Improved approximation algorithms for stochastic matching, in: Proceedings of the 23rd European Symposium on Algorithms (ESA), pp. 1–12.

[] Anderson, R., May, R., 1992. Infectious Diseases of Humans: Dynamics and Control. Infectious Diseases of Humans: Dynamics and Control, OUP Oxford.

[] Asadpour, A., Nazerzadeh, H., 2016. Maximizing stochastic monotone submodular functions. Management Science 62, 2374–2391.

[] Badanidiyuru, A., Papadimitriou, C., Rubinstein, A., Seeman, L., Singer, Y., 2016. Locally adaptive optimization: Adaptive seeding for monotone submodular functions, in: Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms (SODA), pp. 414–429.

[] Barabási, A.L., Albert, R., 1999. Emergence of scaling in random networks. Science 286, 509–512.

[] Borgs, C., Brautbar, M., Chayes, J.T., Lucier, B., 2014. Maximizing social influence in nearly optimal time, in: Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 946–957.

[] Bradac, D., Singla, S., Zuzic, G., 2019. (near) optimal adaptivity gaps for stochastic multi-value probing, in: Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM), pp. 49:1–49:21.

[] Călinescu, G., Chekuri, C., Pál, M., Vondrák, J., 2011. Maximizing a monotone submodular function subject to a matroid constraint. SIAM J. Comput. 40, 1740–1766.

[] Cautis, B., Maniu, S., Tziortziotis, N., 2019. Adaptive influence maximization, in: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD), ACM. pp. 3185–3186.

[] Chen, N., Immorlica, N., Karlin, A.R., Mahdian, M., Rudra, A., 2009. Approximating matches made in heaven, in: Proceedings of the 36th International on Automata, Languages and Programming (ICALP), pp. 266–278.

[] Chen, W., Collins, A., Cummings, R., Ke, T., Liu, Z., Rincón, D., Sun, X., Wang, Y., Wei, W., Yuan, Y., 2011. Influence maximization in social networks when negative opinions may emerge and propagate, in: Proceedings of the Eleventh SIAM International Conference on Data Mining, SDM 2011, pp. 379–390.

[] Chen, W., Lakshmanan, L.V.S., Castillo, C., 2013. Information and Influence Propagation in Social Networks. Synthesis Lectures on Data Management, Morgan & Claypool Publishers.

[] Chen, W., Peng, B., 2019. On adaptivity gaps of influence maximization under the independent cascade model with full-adoption feedback, in: Proceedings of the 30th International Symposium on Algorithms and Computation (ISAAC), pp. 24:1–24:19.

[] Chen, W., Peng, B., Schoenebeck, G., Tao, B., 2022. Adaptive greedy versus non-adaptive greedy for influence maximization. J. Artif. Intell. Res. 74, 303–351.

[] Chen, W., Wang, C., Wang, Y., 2010. Scalable influence maximization for prevalent viral marketing in large-scale social networks, in: Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1029–1038.

[] D'Angelo, G., Poddar, D., Vinci, C., 2021a. Better bounds on the adaptivity gap of influence maximization under full-adoption feedback, in: Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI), pp. 12069–12077.

[] D'Angelo, G., Poddar, D., Vinci, C., 2021b. Improved approximation factor for adaptive influence maximization via simple greedy strategies, in: Proceedings of the 48th International Colloquium on Automata, Languages, and Programming (ICALP), pp. 59:1–59:19.

[] Domingos, P., Richardson, M., 2001. Mining the network value of customers. Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining , 57–66.

[] Fujii, K., Sakaue, S., 2019. Beyond adaptive submodularity: Approximation guarantees of greedy policy with adaptive submodularity ratio, in: Proceedings of the 36th International Conference on Machine Learning (ICML), pp. 2042–2051.

[] Goldberg, S., Liu, Z., 2013. The diffusion of networking technologies, in: Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 1577–1594.

[] Golovin, D., Krause, A., 2011. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. Journal of Artificial Intelligence Research 42, 427–486.

[] Gupta, A., Nagarajan, V., Singla, S., 2016. Algorithms and Adaptivity Gaps for Stochastic Probing. Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms (SODA) , 1731–1747.

[] Gupta, A., Nagarajan, V., Singla, S., 2017. Adaptivity gaps for stochastic probing: Submodular and XOS functions, in: Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 1688–1702.

[] Han, K., Huang, K., Xiao, X., Tang, J., Sun, A., Tang, X., 2018. Efficient algorithms for adaptive influence maximization. Proceedings of the VLDB Endowment 11, 1029–1040.

[] Huang, K., Tang, J., Han, K., Xiao, X., Chen, W., Sun, A., Tang, X., Lim, A., 2020. Efficient approximation algorithms for adaptive influence maximization. VLDB J. 29, 1385–1406.

[] Kempe, D., Kleinberg, J., Tardos, É., 2003. Maximizing the spread of influence through a social network, in: Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 137–146.

[] Kempe, D., Kleinberg, J., Tardos, É., 2015. Maximizing the spread of influence through a social network. Theory of Computing 11, 105–147.

[] Leskovec, J., Krause, A., Guestrin, C., Faloutsos, C., VanBriesen, J.M., Glance, N.S., 2007. Cost-effective outbreak detection in networks, in: Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 420–429.

[] Li, Y., Fan, J., Wang, Y., Tan, K., 2018. Influence maximization on social graphs: A survey. IEEE Transactions on Knowledge and Data Engineering 30, 1852–1872.

[] Liang, G., Yao, X., Gu, Y., Huang, H., Gu, C., 2022. Multi-batches revenue maximization for competitive products over online social network. Journal of Network and Computer Applications 201, 103357.

[] Lowalekar, M., Varakantham, P., Kumar, A., 2016. Robust influence Maximization, in: Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), pp. 1395–1396.

[] Lu, Z., Zhang, Z., Wu, W., 2017. Solution of bharathi-kempe-salek conjecture for influence maximization on arborescence. J. Comb. Optim. 33, 803–808.

[] Mihara, S., Tsugawa, S., Ohsaki, H., 2015. Influence maximization problem for unknown social networks, in: 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), pp. 1539–1546.

[] Norris, J.R., 1997. Markov Chains. Cambridge University Press.

[] Pastor-Satorras, R., Castellano, C., Van Mieghem, P., Vespignani, A., 2015. Epidemic processes in complex networks. Rev. Mod. Phys. 87, 925–979.

[] Peng, B., Chen, W., 2019. Adaptive influence maximization with myopic feedback, in: Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019 (NeurIPS), pp. 5575–5584.

[] Richardson, M., Domingos, P.M., 2002. Mining knowledge-sharing sites for viral marketing, in: Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM. pp. 61–70.

[] Rubinstein, A., Seeman, L., Singer, Y., 2015. Approximability of adaptive seeding under knapsack constraints, in: Proceedings of the 2015 ACM Conference on Economics and Computation (EC), pp. 797–814.

[] Salha, G., Tziortziotis, N., Vazirgiannis, M., 2018. Adaptive submodular influence maximization with myopic feedback, in: Proceedings of the 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), pp. 455–462.

[] Schoenebeck, G., Tao, B., 2019. Influence maximization on undirected graphs: Towards closing the (1-1/e) gap, in: Proceedings of the 2019 ACM Conference on Economics and Computation (EC), ACM. pp. 423–453.

[] Seeman, L., Singer, Y., 2013. Adaptive seeding in social networks, in: Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS), pp. 459–468.

[] Singer, Y., 2016. Influence maximization through adaptive seeding. ACM SIGecom Exchanges 15, 32–59.

[] Sun, L., Huang, W., Yu, P.S., Chen, W., 2018. Multi-round influence maximization, in: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 2249–2258.

[] Tang, J., Lakshmanan, L.V., Huang, K., Tang, X., Sun, A., Xiao, X., Lim, A., 2019. Efficient approximation algorithms for adaptive seed minimization. Proceedings of the ACM SIGMOD International Conference on Management of Data , 1096–1113.

[] Tang, Y., Xiao, X., Shi, Y., 2014. Influence Maximization : Near-Optimal Time Complexity Meets Practical Efficiency. Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data , 75–86.

[] Tong, G., Wang, R., 5555. On adaptive influence maximization under general feedback models. IEEE Transactions on Emerging Topics in Computing , 1–1.

[] Tong, G., Wang, R., Dong, Z., Li, X., 2021. Time-constrained adaptive influence maximization. IEEE Trans. Comput. Soc. Syst. 8, 33–44.

[] Tong, G., Wu, W., Tang, S., Du, D.Z., 2017. Adaptive Influence Maximization in Dynamic Social Networks. IEEE/ACM Transactions on Networking 25, 112–125.

[] Vaswani, S., Lakshmanan, L.V.S., 2016. Adaptive influence maximization in social networks: Why commit when you can adapt? CoRR abs/1604.08171.

[] Yuan, J., Tang, S., 2017. No time to observe: Adaptive influence maximization with partial feedback. Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI) , 3908–3914.

[] Zhang, Y., Chen, S., Xu, W., Zhang, Z., 2022. Adaptive influence maximization under fixed observation time-step. Theor. Comput. Sci. 928, 104–114.