

How to go Viral: Cheaply and Quickly

Ferdinando Cicalese¹, Gennaro Cordasco², Luisa Gargano¹,
M. Milanić³, Joseph G. Peters⁴, and Ugo Vaccaro¹

¹ Department of Computer Science, University of Salerno, Italy
{cicalese,lg,uv}@dia.unisa.it

² Department of Psychology, Second University of Naples, Italy
gennaro.cordasco@unina2.it

³ University of Primorska, UP IAM and UP FAMNIT, 6000 Koper, Slovenia, and
Institute of Mathematics, Physics and Mechanics, 1000 Ljubljana, Slovenia
martin.milanic@upr.si

⁴ School of Computing Science, Simon Fraser University, Canada
peters@cs.sfu.ca

Abstract. Given a social network represented by a graph G , we consider the problem of finding a bounded cardinality set of nodes S with the property that the influence spreading from S in G is as large as possible. The dynamics that govern the spread of influence is the following: initially only elements in S are influenced; subsequently at each round, the set of influenced elements is augmented by all nodes in the network that have a sufficiently large number of already influenced neighbors. While it is known that the general problem is hard to solve — even in the approximate sense — we present exact polynomial time algorithms for trees, paths, cycles, and complete graphs.

Keywords: Social Networks, Spread of Influence, Viral Marketing, Dynamic Monopolies, Exact Polynomial Time Algorithms.

1 The Motivations

Gaming giant FONY[®] is about to launch its brand new console PlayForFUN-7[®], and intends to maximize the adoption of the new product through a massive viral marketing campaign, exploiting the human tendency to conform [4].

This tendency occurs for three reasons: a) the basic human need to be liked and accepted by others [5]; b) the belief that others, especially a majority group, have more accurate and trustworthy information than the individual [29]; c) the “direct-benefit” effect, implying that an individual obtains an explicit benefit when he/she aligns his/her behavior with the behavior of others (e.g., [20], Ch. 17).

In the case in point, argument c) is supported by the fact that each player who buys the PlayForFUN-7 console will be able to play online with all of the people who already have bought the same console. Indeed, the (possible) success of an on-line gaming service comes from its large number of users; if this service had no members, there would be no point to anyone signing up for it. But as people begin using the service, the benefit for more people to sign up increases due to the increasing opportunities to play games with others online. This motivates more people to sign up for the service which further increases the benefit.

FONY is also aware that the much-feared competitor Nanosoft® will soon start to flood the market with a very similar product: FUNBox-14. For this reason, it is crucial to quickly spread the awareness of the new console PlayForFUN-7 to the whole market of potential customers.

The CEO of FONY enthusiastically embraced the idea of a viral marketing campaign¹, and instructed the FONY Marketing Division to plan a viral marketing campaign with the following requirements: 1) an initial set of influential people should be targeted and receive a complimentary personalized PlayForFUN-7 station (because of budget restrictions, this set is required to be *small*); 2) the group of influential people must be judiciously chosen so as to *maximize* the spread of influence within the set of potential PlayForFUN-7 buyers; 3) the spread of influence must happen *quickly*.

To comply with the CEO *desiderata*, FONY Marketing Division analyzed the behavior of players in the network during the past few years (i.e., when players bought the latest console, how many games they bought, how many links/friends they have in the network, and how long they play on average every week). On the basis of this analysis, an estimate of each player's tendency to conform was made, and the following mathematical model was put forward. The network of players is represented by a graph $G = (V, E)$, where V is the set of players, and there is an edge between two players if those two players are friends in the network. The individual's tendency to conform is quantified by a function $t : V \rightarrow \mathbb{N} = \{0, 1, 2, \dots\}$, with easy-to-convince players having "low" $t(\cdot)$ values, and hard-to-convince players having "high" $t(\cdot)$ values. If $S \subseteq V$ is any initial set of targeted people (*target set*), then an *influence spreading process in G* , starting at S , is a sequence of node subsets $\text{Influenced}[S, 0] \subseteq \text{Influenced}[S, 1] \subseteq \dots \subseteq \text{Influenced}[S, \rho] \subseteq \dots \subseteq V$, such that

$$\text{Influenced}[S, 0] = S$$

and for all $\rho > 0$,

$$\text{Influenced}[S, \rho] = \text{Influenced}[S, \rho-1] \cup \left\{ u : |N(u) \cap \text{Influenced}[S, \rho-1]| \geq t(u) \right\},$$

where $N(u)$ is the set of neighbors of u . In words, an individual v becomes influenced if the number of his influenced friends is at least its threshold $t(v)$. It will be said that v is influenced *within* round ρ if $v \in \text{Influenced}[S, \rho]$; v is influenced *at* round $\rho > 0$ if $v \in \text{Influenced}[S, \rho] \setminus \text{Influenced}[S, \rho-1]$.

Using this terminology and notation, we can formally state the original problem as:

(λ, β) -MAXIMALLY INFLUENCING SET ((λ, β)-MIS).

Instance: A graph $G = (V, E)$, thresholds $t : V \rightarrow \mathbb{N}$, a latency bound $\lambda \in \mathbb{N}$ and a budget $\beta \in \mathbb{N}$.

Question: Find a set $S \subseteq V$ such that $|S| \leq \beta$ and $|\text{Influenced}[S, \lambda]|$ is as large as possible.

¹ "If politicians can sell their stuff through a viral marketing campaign [9, 25, 30], then why not us?", an unconfirmed source claims the CEO said.

2 The Context

It did not spoil the fun(!) of FONY Marketing Division to learn that (variants of) the (λ, β) -MIS problem have already been studied in the scientific literature. We shall limit ourselves here to discussing the work that is most directly related to ours, and refer the reader to the monographs [13, 20] for an excellent overview of the area. We just mention that our results also seem to be relevant to other areas, like dynamic monopolies [21, 27] for instance.

The first authors to study the spread of influence in networks from an algorithmic point of view were Kempe *et al.* [23, 24]. However, they were mostly interested in networks with randomly chosen thresholds. Chen [11] studied the following minimization problem: given a graph G and fixed thresholds $t(v)$, find a set of minimum size that eventually influences all (or a fixed fraction of) nodes of G . He proved a strong inapproximability result that makes unlikely the existence of an algorithm with approximation factor better than $O(2^{\log^{1-\epsilon} |V|})$. Chen's result stimulated a series of papers [1, 6, 7, 10, 14–17, 19, 22, 28, 31], that isolated interesting cases in which the problem (and variants thereof) becomes tractable.

None of these papers considered the *number of rounds* necessary for the spread of influence in the network. However, this is a relevant question for viral marketing in which it is quite important to spread information quickly. Indeed, research in Behavioural Economics shows that humans make decisions mostly on the basis of very recent events, even though they might remember much more [2, 12]. The only paper known to us that has studied the spread of influence in the same diffusion model that we consider here, and with constraints on the number of rounds in which the process must be completed, is [18]. How our results are related to [18] will be elucidated in the next section. Finally, we point out that Chen's [11] inapproximability result still holds for general graphs if the diffusion process must end in a bounded number of rounds.

3 The Results

Our main results are polynomial time algorithms to solve the (λ, β) -MIS problem on Trees, Paths, Cycles, and Complete graphs, improving and extending some results from [18]. In particular, the paper [18] put forward an algorithmic framework to solve the (λ, β) -MIS problem (and related ones), in graphs of bounded clique-width. When instantiated on trees, the approach of [18] would give algorithms for the (λ, β) -MIS problem with complexity that is *exponential* in the parameter λ , whereas our algorithm has complexity polynomial in all the relevant parameters (cf., Theorem 1). We should also remark that, in the very special case $\lambda = 1$ and thresholds $t(v) = 1$, for each $v \in V$, problems of influence diffusion reduce to well known domination problems in graphs (and variants thereof). In particular, when $\lambda = 1$ and $t(v) = 1$, for each $v \in V$, our (λ, β) -MAXIMALLY INFLUENCING SET problem reduces to the MAXIMUM COVERAGE problem considered in [8]. Therefore, our results can also be seen as far-reaching generalizations of [8].

4 (λ, β) -Maximally Influencing Set on Trees

In this section, we give an algorithm for the (λ, β) -MAXIMALLY INFLUENCING SET problem on trees. Let $T = (V, E)$ be a tree, rooted at some node r . Once such a rooting is fixed, for any node v , we denote by $T(v)$ the subtree rooted at v . We will develop a dynamic programming algorithm that will prove the following theorem.

Theorem 1. *The (λ, β) -MAXIMALLY INFLUENCING SET problem can be solved in time $O(\min\{n\Delta^2\lambda^2\beta^3, n^2\lambda^2\beta^3\})$ on a tree with n nodes and maximum degree Δ .*

The rest of this section is devoted to the description and analysis of the algorithm that proves Theorem 1. The algorithm traverses the input tree T bottom up, in such a way that each node is considered after all its children have been processed. For each node v , the algorithm solves all possible (λ, b) -MIS problems on the subtree $T(v)$, for $b = 0, 1, \dots, \beta$. Moreover, in order to compute these values we will have to consider not only the original threshold $t(v)$ of v , but also the decreased value $t(v) - 1$ which we call the *residual threshold*. In the following, we assume without loss of generality that $0 \leq t(u) \leq d(u) + 1$ (where $d(u)$ denotes the degree of u) holds for all nodes $u \in V$ (otherwise, we can set $t(u) = d(u) + 1$ for every node u with threshold exceeding its degree plus one without changing the problem).

Definition 1. *For each node $v \in V$, integers $b \geq 0$, $t \in \{t(v) - 1, t(v)\}$, and $\rho \in \{0, 1, \dots, \lambda\} \cup \{\infty\}$, let us denote by $MIS[v, b, \rho, t]$ the maximum number of nodes that can be influenced in $T(v)$, within round λ , assuming that*

- at most b nodes among those in $T(v)$ belong to the target set;
- the threshold of v is t ;
- the parameter ρ is such that

1) if $\rho = 0$ then v must belong to the target set, (1)

2) if $1 \leq \rho \leq \lambda$ then v is not in the target set and at least t of its children are active within round $\rho - 1$, (2)

3) if $\rho = \infty$ then v is not influenced within round λ . (3)

We define $MIS[v, b, \rho, t] = -\infty$ when any of the above constraints is not satisfiable. For instance, if $b = \rho = 0$ we have² $MIS[v, 0, 0, t] = -\infty$.

Denote by $S(v, b, \rho, t)$ any target set attaining the value $MIS[v, b, \rho, t]$.

We notice that in the above definition if $1 \leq \rho \leq \lambda$ then, the assumption that v has threshold t implies that v is influenced within round ρ and is able to influence its neighbors starting from round $\rho + 1$. The value $\rho = \infty$ means that no condition are imposed on v : It could be influenced after round λ or not influenced at all. In the sequel, $\rho = \infty$ will be used to ensure that v will not contribute to the influence any neighbor (within round λ).

² Since $\rho = 0$ then v should belong to the target set, but this is not possible because the budget is 0.

Remark 1. It is worthwhile mentioning that $MIS[v, b, \rho, t]$ is monotonically non-decreasing in b and non-increasing in t . However, $MIS[v, b, \rho, t]$ is not necessarily monotonic in ρ .

The maximum number of nodes in G that can be influenced within round λ with any (initial) target set of cardinality at most β can be then obtained by computing

$$\max_{\rho \in \{0, 1, \dots, \lambda, \infty\}} MIS[r, \beta, \rho, t(r)]. \quad (4)$$

In order to obtain the value in (4), we compute $MIS[v, b, \rho, t]$ for each $v \in V$, for each $b = 0, 1, \dots, \beta$, for each $\rho \in \{0, 1, \dots, \lambda, \infty\}$, and for $t \in \{t(v) - 1, t(v)\}$.

We proceed in a bottom-up fashion on the tree, so that the computation of the various values $MIS[v, b, \rho, t]$ for a node v is done after all the values for v 's children are known.

For each leaf node ℓ we have

$$MIS[\ell, b, \rho, t] = \begin{cases} 1 & \text{if } (\rho = 0 \text{ AND } b \geq 1) \text{ OR } (t = 0 \text{ AND } 1 \leq \rho \leq \lambda) \\ 0 & \text{if } \rho = \infty \\ -\infty & \text{otherwise.} \end{cases} \quad (5)$$

Indeed, a leaf ℓ gets influenced, in the single node subtree $T(\ell)$, only when either ℓ belongs to the target set ($\rho = 0$) and the budget is sufficiently large ($b \geq 1$) or the threshold is zero (either $t = t(\ell) = 0$ or $t = t(\ell) - 1 = 0$) independently of the number of rounds.

For an internal node v , we show how to compute each value $MIS[v, b, \rho, t]$ in time $O(d(v)^2 \lambda \beta^2)$.

We recall that when computing a value $MIS[v, b, \rho, t]$, we already have computed all the $MIS[v_i, *, *, *]$ values for each child v_i of v .

We distinguish three cases for the computation of $MIS[v, b, \rho, t]$ according to the value of ρ .

CASE 1: $\rho = 0$. In this case we assume that $b \geq 1$ (otherwise $MIS[v, 0, 0, t] = -\infty$). Moreover, we know that $v \in S(v, b, 0, t)$ hence the computation of $MIS[v, b, 0, t]$ must consider all the possible ways in which the remaining budget $b - 1$ can be partitioned among v 's children.

Lemma 1. *It is possible to compute $MIS[v, b, 0, t]$, where $b \geq 1$, in time $O(d\lambda b^2)$, where d is the number of children of v .*

Proof. Fix an ordering v_1, v_2, \dots, v_d of the children of node v .

For $i = 1, \dots, d$ and $j = 0, \dots, b - 1$, let $AMAX_v[i, j]$ be the maximum number of nodes that can be influenced, within λ rounds, in $T(v_1), T(v_2), \dots, T(v_i)$ assuming that the target set contains v and at most j nodes among those in $T(v_1), T(v_2), \dots, T(v_i)$.

By (1) we have

$$MIS[v, b, 0, t] = 1 + AMAX_v[d, b - 1]. \quad (6)$$

We now show how to compute $AMAX_v[d, b - 1]$ by recursively computing the values $AMAX_v[i, j]$, for each $i = 1, 2, \dots, d$ and $j = 0, 1, \dots, b - 1$.

For $i = 1$, we assign all of the budget to $T(v_1)$ and

$$AMAX_v[1, j] = \max_{\rho_1, t_1} \{MIS[v_1, j, \rho_1, t_1]\},$$

where $\rho_1 \in \{0, \dots, \lambda, \infty\}$, $t_1 \in \{t(v_1), t(v_1) - 1\}$, and if $t_1 = t(v_1) - 1$ then $\rho_1 \geq 1$.

For $i > 1$, we consider all possible ways of partitioning the budget j into two values a and $j - a$, for each $0 \leq a \leq j$. The budget a is assigned to the first $i - 1$ subtrees, while the budget $j - a$ is assigned to $T(v_i)$. Hence,

$$AMAX_v[i, j] = \max_{0 \leq a \leq j} \left\{ AMAX_v[i - 1, a] + \max_{\rho_i, t_i} \{MIS[v_i, j - a, \rho_i, t_i]\} \right\}$$

where $\rho_i \in \{0, \dots, \lambda, \infty\}$, $t_i \in \{t(v_i), t(v_i) - 1\}$, and if $t_i = t(v_i) - 1$ then $\rho_i \geq 1$.

The computation of $AMAX_v$ comprises $O(db)$ values and each one is computed recursively in time $O(\lambda b)$. Hence we are able to compute it, and by (6), also $MIS[v, b, 0, t]$, in time $O(d\lambda b^2)$.

CASE 2: $1 \leq \rho \leq \lambda$. In this case v is not in the target set and at round $\rho - 1$ at least t of its children must be influenced. The computation of a value $MIS[v, b, \rho, t]$ must consider all the possible ways in which the budget b can be partitioned among v 's children in such a way that at least t of them are influenced within round $\rho - 1$.

Lemma 2. *For each $\rho = 1, \dots, \lambda$, it is possible to compute $MIS[v, b, \rho, t]$ recursively in time $O(d^2 \lambda b^2)$, where d is the number of children of v .*

Proof. Fix any ordering v_1, v_2, \dots, v_d of the children of the node v .

We first define the values $BMAX_{v, \rho}[i, j, k]$, for $i = 1, \dots, d$, $j = 0, \dots, b$, and $k = 0, \dots, t$.

If $i \geq k$, we define $BMAX_{v, \rho}[i, j, k]$ to be the maximum number of nodes that can be influenced, within λ rounds, in the subtrees $T(v_1), T(v_2), \dots, T(v_i)$ assuming that

- v is influenced within round ρ ;
- at most j nodes among those in $T(v_1), T(v_2), \dots, T(v_i)$ belong to the target set;
- at least k among v_1, v_2, \dots, v_i , will be influenced within round $\rho - 1$.

We define $BMAX_{v, \rho}[i, j, k] = -\infty$ when the above constraints are not satisfiable. For instance, if $i < k$ we have $BMAX_{v, \rho}[i, j, k] = -\infty$.

By (2) and by the definition of $BMAX$, we have

$$MIS[v, b, \rho, t] = 1 + BMAX_{v, \rho}[d, b, t]. \quad (7)$$

We can compute $BMAX_{v, \rho}[d, b, t]$ by recursively computing the values of $BMAX_{v, \rho}[i, j, k]$ for each $i = 1, 2, \dots, d$, for each $j = 0, 1, \dots, b$, and for each $k = 0, 1, \dots, t$, as follows.

For $i = 1$, we have to assign all the budget j to the first subtree of v . Moreover, if $k = 1$, then by definition v_1 has to be influenced before round ρ and consequently we

can not use threshold $t(v_1) - 1$ (which assumes that v contributes to the influence of v_i). Hence, we have

$$BMAX_{v,\rho}[1, j, k] = \begin{cases} \max_{\rho_1, t_1} \{MIS[v_1, j, \rho_1, t_1]\}, & \text{if } k = 0 \\ \max_{\delta} \{MIS[v_1, j, \delta, t(v_1)]\}, & \text{if } k = 1 \\ -\infty, & \text{otherwise,} \end{cases} \quad (8)$$

where

- $\rho_1 \in \{0, \dots, \lambda, \infty\}$
- $t_1 \in \{t(v_1), t(v_1) - 1\}$
- if $t_1 = t(v_1) - 1$ then $\rho_1 \geq \rho + 1$
- $\delta \in \{0, \dots, \rho - 1\}$.

The third constraint ensures that we can use a reduced threshold on v_1 only after the father v has been influenced.

To show the correctness of equation (8), one can (easily) check that, for $k < 2$, any target set solution S that maximizes the value on the left side of the equation is also a feasible solution for the value on the right, and vice versa.

For $i > 1$, as in the preceding lemma, we consider all possible ways of partitioning the budget j into two values a and $j - a$. The budget a is assigned to the first $i - 1$ subtrees, while the remaining budget $j - a$ is assigned to $T(v_i)$. Moreover, in order to ensure that at least k children of v , among children v_1, v_2, \dots, v_i , will be influenced before round ρ , there are two cases to consider: a) the k children that are influenced before round ρ are among the first $i - 1$ children of v . In this case v_i can be influenced at any round and can use a reduced threshold; b) only $k - 1$ children among nodes v_1, v_2, \dots, v_{i-1} are influenced before round ρ and consequently v_i has to be influenced before round ρ and cannot use a reduced threshold. Formally, we prove that

$$BMAX_{v,\rho}[i, j, k] = \max \left\{ \max_{\substack{0 \leq a \leq j \\ \rho_i, t_i}} (BMAX_{v,\rho}[i-1, a, k] + MIS[v_i, j-a, \rho_i, t_i]), \right. \\ \left. \max_{\substack{0 \leq a \leq j \\ \delta}} (BMAX_{v,\rho}[i-1, a, k-1] + MIS[v_i, j-a, \delta, t(v_i)]) \right\} \quad (9)$$

where

- $\rho_i \in \{0, \dots, \lambda, \infty\}$
- $t_i \in \{t(v_i), t(v_i) - 1\}$
- if $t_i = t(v_i) - 1$ then $\rho_i \geq \rho + 1$
- $\delta \in \{0, \dots, \rho - 1\}$.

In the following we show the correctness of equation (9). First we show that

$$BMAX_{v,\rho}[i, j, k] \leq \max \left\{ \max_{\substack{0 \leq a \leq j \\ \rho_i, t_i}} (BMAX_{v,\rho}[i-1, a, k] + MIS[v_i, j-a, \rho_i, t_i]), \right. \\ \left. \max_{\substack{0 \leq a \leq j \\ \delta}} (BMAX_{v,\rho}[i-1, a, k-1] + MIS[v_i, j-a, \delta, t(v_i)]) \right\}$$

Let $S \subseteq \bigcup_{z=1}^i T(v_z)$ be a feasible target set solution that maximizes the number of nodes that can be influenced, within λ rounds, in the subtrees $T(v_1), T(v_2), \dots, T(v_i)$

and satisfies the constraints defined in the definition of $BMAX_{v,\rho}[i, j, k]$. Hence $|S| \leq j$. We can partition S into two sets S_a , where $|S_a| \leq a$, and S_b ($|S_b| \leq j - a$) in such a way that $S_a \subseteq \bigcup_{z=1}^{i-1} T(v_z)$ while $S_b \subseteq T(v_i)$. Since S satisfies the constraints defined in the definition of $BMAX_{v,\rho}[i, j, k]$, we have that, starting with S , at least k children of v , among children v_1, v_2, \dots, v_i , will be influenced before round ρ . Hence, starting with S_a , at least $k - 1$ children of v , among children v_1, v_2, \dots, v_{i-1} , will be influenced before round ρ . We distinguish two cases:

- If S_a influences $k - 1$ children of v , among children v_1, v_2, \dots, v_{i-1} , before round ρ , then we have that S_b must also influence v_i before round ρ . Hence S_a is a feasible solution for $BMAX_{v,\rho}[i-1, a, k-1]$ and S_b is a feasible solution for $\max_{\delta} \{MIS[v_i, j-a, \delta, t(v_i)]\}$.
- On the other hand when S_a influences at least k children of v , among children v_1, v_2, \dots, v_{i-1} , before round ρ then S_a is a feasible solution for $BMAX_{v,\rho}[i-1, a, k]$ and S_b is a feasible solution for $\max_{\rho_i, t_i} \{MIS[v_i, j-a, \rho_i, t_i]\}$.

In either case we have that the solution S is also a solution for the right side of the equation. Perfectly similar reasoning can be used to show that

$$BMAX_{v,\rho}[i, j, k] \geq \max \left\{ \max_{\substack{0 \leq a \leq j \\ \rho_i, t_i}} (BMAX_{v,\rho}[i-1, a, k] + MIS[v_i, j-a, \rho_i, t_i]), \right. \\ \left. \max_{\substack{0 \leq a \leq j \\ \delta}} (BMAX_{v,\rho}[i-1, a, k-1] + MIS[v_i, j-a, \delta, t(v_i)]) \right\}$$

and hence equation (9) is proved.

The computation of $BMAX_{v,\rho}$ comprises $O(d^2b)$ values (recall that $t \leq d + 2$) and each one is computed recursively in time $O(\lambda b)$. Hence we are able to compute it, and by (7), also $MIS[v, b, \rho, t]$, in time $O(d^2 \lambda b^2)$.

CASE 3: $\rho = \infty$. In this case we only have to consider the original threshold $t(v_i)$ for each child v_i of v . Moreover, we must consider all the possible ways in which the budget b can be partitioned among v 's children.

Lemma 3. *It is possible to compute $MIS[v, b, \infty, t]$ in time $O(d\lambda b^2)$, where d is the number of children of v .*

Proof. Fix any ordering v_1, v_2, \dots, v_d of the children of the node v .

For $i = 1, \dots, d$ and $j = 0, \dots, b$, let $CMAX_v[i, j]$ be the maximum number of nodes that can be influenced, within λ rounds, in $T(v_1), T(v_2), \dots, T(v_i)$ assuming that

- v will not be influenced within λ rounds and
- at most j nodes, among nodes in $T(v_1), T(v_2), \dots, T(v_i)$, belong to the target set.

By (3) and by the definition of $CMAX$, we have

$$MIS[v, b, \infty, t] = CMAX_v[d, b]. \quad (10)$$

We can compute $CMAX_v[d, b]$ by recursively computing the values $CMAX_v[i, j]$ for each $i = 1, 2, \dots, d$ and for each $j = 0, 1, \dots, b$, as follows.

For $i = 1$, we can assign all of the budget to the first subtree of v and we have

$$CMAX_v[1, j] = \max_{\rho_1} \{MIS[v_1, j, \rho_1, t(v_1)]\}$$

where $\rho_1 \in \{0, \dots, \lambda, \infty\}$.

For $i > 1$, we consider all possible ways of partitioning the budget j into two values a and $j - a$, for each $0 \leq a \leq j$. The budget a is assigned to the first $i - 1$ subtrees, while the remaining budget $j - a$ is assigned to $T(v_i)$. Hence, the following holds:

$$CMAX_v[i, j] = \max_{0 \leq a \leq j} \left\{ CMAX_v[i - 1, a] + \max_{\rho_i} \{MIS[v_i, j - a, \rho_i, t(v_i)]\} \right\}$$

where $\rho_i \in \{0, \dots, \lambda, \infty\}$.

The computation of $CMAX_v$ comprises $O(db)$ values and each one is computed recursively in time $O(\lambda b)$. Hence, by (10), we are able to compute $MIS[v, b, \infty, t]$ in time $O(d\lambda b^2)$.

Thanks to the three lemmas above we have that for each node $v \in V$, for each $b = 0, 1, \dots, \beta$, for each $\rho = 0, 1, \dots, \lambda, \infty$, and for $t \in \{t(v) - 1, t(v)\}$, $MIS[v, b, \rho, t]$ can be computed recursively in time $O(d(v)^2 \lambda \beta^2)$. Hence, the value

$$\max_{\rho \in \{0, 1, \dots, \lambda, \infty\}} MIS[r, \beta, \rho, t(r)]$$

can be computed in time

$$\sum_{v \in V} O(d(v)^2 \lambda \beta^2) \times O(\lambda \beta) = O(\lambda^2 \beta^3) \times \sum_{v \in V} O(d(v)^2) = O(\min\{n\Delta^2 \lambda^2 \beta^3, n^2 \lambda^2 \beta^3\}),$$

where Δ is the maximum node degree. Standard backtracking techniques can be used to compute a target set of cardinality at most β that influences this maximum number of nodes in the same $O(\min\{n\Delta^2 \lambda^2 \beta^3, n^2 \lambda^2 \beta^3\})$ time. This proves Theorem 1.

5 (λ, β) -Maximally Influencing Set on Paths, Cycles, and Complete Graphs

The results of Section 4 obviously include paths. However, we are able to significantly improve on the computation time for paths.

Let $P_n = (V, E)$ be a path on n nodes v_1, v_2, \dots, v_n , and edges (v_i, v_{i+1}) , for $i = 1, \dots, n - 1$. Moreover, we denote by C_n the cycle on n nodes that consists of the path P_n augmented with the edge (v_1, v_n) . In the following, we assume that $1 \leq t(i) \leq 3$, for $i = 1, \dots, n$. Indeed, paths with 0-threshold nodes can be dealt with by removing up to λ 1-threshold nodes on the two sides of each 0-threshold node. In case we remove strictly less than λ nodes, we can reduce by 1 the threshold of the first node that is not removed (which must have threshold greater than 1). The path gets split into several subpaths, but the construction we provide below still works (up to taking care of boundary conditions).

Theorem 2. *The (λ, β) -MAXIMALLY INFLUENCING SET problem can be solved in time $O(n\beta\lambda)$ on a path P_n .*

Proof. (Sketch.) For $i = 1, 2, \dots, n$, let $r(i)$ be the number of consecutive nodes having threshold 1 on the right of node v_i , that is, $r(i)$ is the largest integer such that $i + r(i) \leq n$ and $t(v_{i+1}) = t(v_{i+2}) = \dots = t(v_{i+r(i)}) = 1$. Analogously we define $l(i)$ as the largest integer such that $i - l(i) \geq 1$ and $t(v_{i-1}) = t(v_{i-2}) = \dots = t(v_{i-l(i)}) = 1$.

We use $P(i, r, t)$ to denote the subpath of P induced by nodes v_1, v_2, \dots, v_{i+r} , where the threshold of each node v_j with $j \neq i$ is $t(v_j)$, while the threshold of v_i is set to $t \in \{t(v_i) - 1, t(v_i)\}$.

We define $MIS[i, b, r, t]$ to be the maximum number of nodes that can be influenced in $P(i, r, t)$ assuming that at most b nodes among v_1, v_2, \dots, v_i belong to the target set while v_{i+1}, \dots, v_{i+r} do not.

Noticing that $P(n, 0, t(v_n)) = P$ and we require that $|S| \leq \beta$, the desired value is $MIS[n, \beta, 0, t(v_n)]$.

In order to get $MIS[n, \beta, 0, t(v_n)]$, we compute $MIS[i, b, r, t]$ for each $i = 0, 1, \dots, n$, for each $b = 0, 1, \dots, \beta$, for each $r = 0, 1, \dots, \min\{\lambda, r(i)\}$, and for $t \in \{t(v_i) - 1, t(v_i)\}$.

Denote by $S(i, b, r, t)$ any target set attaining the value $MIS[i, b, r, t]$.

If $i = 0$ OR $b = 0$ we set $MIS[i, b, r, t] = 0$.

If $i > 0$ AND $b > 0$. Consider the following quantities

$$\begin{aligned} \ell &= \min\{\lambda, l(i)\} \\ M_0 &= \begin{cases} MIS[i-\ell-1, b-1, 0, t(v_{i-\ell-1}) - 1] + r + \ell + 1 & \text{if } \ell < \lambda \\ MIS[i-\ell-1, b-1, 0, t(v_{i-\ell-1})] + r + \ell + 1 & \text{otherwise} \end{cases} \\ M_1 &= \begin{cases} MIS[i-1, b, 0, t(v_{i-1})] & \text{if } t > 1 \\ MIS[i-1, b, \min\{\lambda, r+1\}, t(v_{i-1})] & \text{otherwise.} \end{cases} \end{aligned}$$

By distinguishing whether v_i belongs to the target set $S(i, b, r, t)$ or not we are able to prove that

$$MIS[i, b, r, t] = \max\{M_0, M_1\}$$

and $v_i \in S(i, b, r, t)$ if and only if $MIS[i, b, r, t] = M_0$.

For cycles, the problem can be solved by simply solving two different problems on a path and taking the minimum. Indeed, starting with a cycle we can consider any node v such that $t(v) \geq 2$ (if there is no such node, then the problem is trivial). If node v belongs to the target set, we can consider the path obtained by removing all the nodes influenced only by v and then solve the problem on this path with a budget $\beta - 1$. On the other hand, if we assume that v does not belong to the target set, then we simply consider the path obtained by eliminating v . Therefore, we obtain the following result.

Theorem 3. *The (λ, β) -MAXIMALLY INFLUENCING SET problem can be solved in time $O(n\beta\lambda)$ on a cycle C_n .*

Since complete graphs are of clique-width at most 2, results from [18] imply that the (λ, β) -MIS problem is solvable in polynomial time on complete graphs if λ is constant.

Indeed, one can see that for complete graphs the (λ, β) -MAXIMALLY INFLUENCING SET can be solved in linear time, independently of the value of λ , by using ideas of [26].

If G is a complete graph, we have that for any $S \subseteq V$, and any round $\rho \geq 1$, it holds that

$$\text{Influenced}[S, \rho] = \text{Influenced}[S, \rho - 1] \cup \{v : t(v) \leq |\text{Influenced}[S, \rho - 1]|\}.$$

Since $\text{Influenced}[S, \rho - 1] \subseteq \text{Influenced}[S, \rho]$, we have

$$\text{Influenced}[S, \rho] = S \cup \{v : t(v) \leq |\text{Influenced}[S, \rho - 1]|\}. \quad (11)$$

From (11), and by using a standard exchanging argument, one immediately sees that a set S with largest influence is the one containing the nodes with highest thresholds. Since $t(v) \in \{0, 1, \dots, n\}$, the selection of the β nodes with highest threshold can be done in linear time. Summarizing, we have the following result.

Theorem 4. *There exists an optimal solution S to the (λ, β) -MAXIMALLY INFLUENCING SET problem on a complete graph $G = (V, E)$, consisting of the β nodes of V with highest thresholds, and it can be computed in linear time.*

6 Concluding Remarks

We considered the problems of selecting a *bounded* cardinality subset of people in (classes of) networks, such that the influence they spread, in a *fixed* number of rounds, is the *highest* among all subsets of same bounded cardinality. It is not difficult to see that our techniques can also solve closely related problems, in the same classes of graphs considered in this paper. For instance, one could fix a requirement α and ask for the *minimum* cardinality target set such that after λ rounds the number of influenced people in the network is at least α . Or, one could fix a budget β and a requirement α , and ask about the *minimum* number λ such that there exists a target set of cardinality at most β that influences at least α people in the network within λ rounds (such a minimum λ could be equal to ∞). Therefore, it is likely that the FONY[®] Marketing Division will have additional fun in solving these problems (and similar ones) as well.

References

1. Ackerman, E., Ben-Zwi, O., Wolfvitz, G.: Combinatorial model and bounds for target set selection. *Theoretical Computer Science* 411, 4017–4022 (2010)
2. Alba, J., Hutchinson, J.W., Lynch, J.: Memory and Decision Making. In: Robertson, T.S., Kassarjian, H. (eds.) *Handbook of Consumer Behavior* (1991)
3. Aral, S., Walker, D.: Identifying Influential and Susceptible Members of Social Networks. *Science* 337(6092), 337–341 (2012)
4. Asch, S.E.: Studies of independence and conformity: A minority of one against a unanimous majority. *Psychological Monographs* 70 (1956)
5. Baumeister, R.F., et al.: The need to belong: Desire for interpersonal attachments as a fundamental human motivation. *Psychological Bulletin* 117(3), 497–529 (1995)

6. Ben-Zwi, O., Hermelin, D., Lokshtanov, D., Newman, I.: Treewidth governs the complexity of target set selection. *Discrete Optimization* 8, 87–96 (2011)
7. Bazgan, C., Chopin, M., Nichterlein, A., Sikora, F.: Parameterized approximability of maximizing the spread of influence in networks. In: Du, D.-Z., Zhang, G. (eds.) *COCOON 2013*. LNCS, vol. 7936, pp. 543–554. Springer, Heidelberg (2013)
8. Blair, J.R.S., Goddard, W., Hedetniemi, S.T., Horton, S., Jones, P., Kubicki, G.: On domination and reinforcement numbers in trees. *Discrete Mathematics* 308(7), 1165–1175 (2008)
9. Bond, R.M., et al.: A 61-million-person experiment in social influence and political mobilization. *Nature* 489, 295–298 (2012)
10. Centeno, C.C., Dourado, M.C., Draque Penso, L., Rautenbach, D., Szwarcfiter, J.L.: Irreversible conversion of graphs. *Theoretical Computer Science* 412(29), 3693–3700 (2011)
11. Chen, N.: On the approximability of influence in social networks. *SIAM J. Discrete Math.* 23, 1400–1415 (2009)
12. Chen, J., Iyer, G., Pazgal, A.: Limited Memory, Categorization and Competition. *Marketing Science* 29, 650–670 (2010)
13. Chen, W., Lakshmanan, L.V.S., Castillo, C.: *Information and Influence Propagation in Social Networks*. Morgan & Claypool (2013)
14. Chiang, C.Y., et al.: The Target Set Selection Problem on Cycle Permutation Graphs, Generalized Petersen Graphs and Torus Cordalis. arXiv:1112.1313 (2011)
15. Chopin, M., Nichterlein, A., Niedermeier, R., Weller, M.: Constant thresholds can make target set selection tractable. In: Even, G., Rawitz, D. (eds.) *MedAlg 2012*. LNCS, vol. 7659, pp. 120–133. Springer, Heidelberg (2012)
16. Chiang, C.-Y., Huang, L.-H., Li, B.-J., Wu, J., Yeh, H.-G.: Some results on the target set selection problem. *Journal of Combinatorial Optimization* 25(4), 702–715 (2013)
17. Chiang, C.-Y., Huang, L.-H., Yeh, H.-G.: Target Set Selection Problem for Honeycomb Networks. *SIAM J. Discrete Math.* 27(1), 310–328 (2013)
18. Cicalese, F., Cordasco, G., Gargano, L., Milanić, M., Vaccaro, U.: Latency-Bounded Target Set Selection in Social Networks. In: Bonizzoni, P., Brattka, V., Löwe, B. (eds.) *CiE 2013*. LNCS, vol. 7921, pp. 65–77. Springer, Heidelberg (2013)
19. Coja-Oghlan, A., Feige, U., Krivelevich, M., Reichman, D.: Contagious sets in expanders. arXiv:1306.2465
20. Easley, D., Kleinberg, J.: *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press (2010)
21. Flocchini, P., Králóvic, R., Ruzicka, P., Roncato, A., Santoro, N.: On time versus size for monotone dynamic monopolies in regular topologies. *J. Discrete Algorithms* 1, 129–150 (2003)
22. Gargano, L., Hell, P., Peters, J., Vaccaro, U.: Influence diffusion in social networks under time window constraints. In: Moscibroda, T., Rescigno, A.A. (eds.) *SIROCCO 2013*. LNCS, vol. 8179, pp. 141–152. Springer, Heidelberg (2013)
23. Kempe, D., Kleinberg, J.M., Tardos, E.: Maximizing the spread of influence through a social network. In: *Proc. of the Ninth ACM SIGKDD*, pp. 137–146 (2003)
24. Kempe, D., Kleinberg, J.M., Tardos, É.: Influential Nodes in a Diffusion Model for Social Networks. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) *ICALP 2005*. LNCS, vol. 3580, pp. 1127–1138. Springer, Heidelberg (2005)
25. Leppaniemi, M., Karjaluoto, H., Lehto, H., Goman, A.: Targeting Young Voters in a Political Campaign: Empirical Insights into an Interactive Digital Marketing Campaign in the 2007 Finnish General Election. *Journal of Nonprofit & Public Sector Marketing* 22, 14–37 (2007)
26. Nichterlein, A., Niedermeier, R., Uhlmann, J., Weller, M.: On tractable cases of target set selection. *Social Network Analysis and Mining* (2012)
27. Peleg, D.: Local majorities, coalitions and monopolies in graphs: a review. *Theoretical Computer Science* 282, 231–257 (2002)

28. Reddy, T.V.T., Rangan, C.P.: Variants of spreading messages. *J. Graph Algorithms Appl.* 15(5), 683–699 (2011)
29. Surowiecki, J.: *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*. Doubleday (2004)
30. Tumulty, K.: Obama's Viral Marketing Campaign. *TIME Magazine* (July 5, 2007)
31. Zaker, M.: On dynamic monopolies of graphs with general thresholds. *Discrete Mathematics* 312(6), 1136–1143 (2012)