Theoretical Computer Science 535 (2014) 1–15

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Latency-bounded target set selection in social networks ☆

Ferdinando Cicalese [a], Gennaro Cordasco [b,*], Luisa Gargano [a], Martin Milanič [c], Ugo Vaccaro [a]

[a] Dept. of Computer Science, University of Salerno, Italy
[b] Dept. of Psychology, Second University of Naples, Italy
[c] University of Primorska, UP IAM and UP FAMNIT, SI 6000 Koper, Slovenia

ARTICLE INFO

Article history:
Received 23 October 2013
Accepted 20 February 2014
Communicated by P.G. Spirakis

Keywords:
Target set selection
Monotone irreversible dynamic monopolies (dynamo)

ABSTRACT

Motivated by applications in sociology, economy and medicine, we study variants of the Target Set Selection problem, first proposed by Kempe, Kleinberg and Tardos. In our scenario one is given a graph $G=(V,E)$, integer values $t(v)$ for each vertex v (*thresholds*), and the objective is to determine a small set of vertices (*target set*) that activates a given number (or a given subset) of vertices of G *within* a prescribed number of rounds. The activation process in G proceeds as follows: initially, at round 0, all vertices in the target set are activated; subsequently at each round $r \geqslant 1$ every vertex of G becomes activated if at least $t(v)$ of its neighbors are already active by round $r-1$. It is known that the problem of finding a minimum cardinality Target Set that eventually activates the whole graph G is hard to approximate to a factor better than $O(2^{\log^{1-\epsilon}|V|})$. In this paper we give *exact* polynomial time algorithms to find minimum cardinality Target Sets in graphs of bounded clique-width, and *exact* linear time algorithms for trees.

1. Introduction

Let $G=(V,E)$ be a graph, $S \subseteq V$, and let $t: V \to \mathbb{N} = \{1,2,\ldots\}$ be a function assigning integer thresholds to the vertices of G. An *activation process in G starting at S* is a sequence $\mathsf{Active}[S,0] \subseteq \mathsf{Active}[S,1] \subseteq \cdots \subseteq \mathsf{Active}[S,i] \subseteq \cdots \subseteq V$ of vertex subsets, with $\mathsf{Active}[S,0]=S$, and such that for all $i>0$,

$$\mathsf{Active}[S,i] = \mathsf{Active}[S,i-1] \cup \big\{u: \big|N(u) \cap \mathsf{Active}[S,i-1]\big| \geqslant t(u)\big\}$$

where $N(u)$ is the set of neighbors of u. In words, at each round i the set of active nodes is augmented by the set of nodes u that have a number of *already* activated neighbors greater or equal to u's threshold $t(u)$. The central problem we introduce and study in this paper is defined as follows:

(λ,β,α)-TARGET SET SELECTION ((λ,β,α)-TSS).
Instance: A graph $G=(V,E)$, thresholds $t: V \longrightarrow \mathbb{N}$, a latency bound $\lambda \in \mathbb{N}$, a budget $\beta \in \mathbb{N}$ and an activation requirement $\alpha \in \mathbb{N}$.

☆ An extended abstract of a preliminary version of this paper appeared in Proceedings of Computability in Europe 2013 (CiE 2013), The Nature of Computation: Logic, Algorithms, Applications, Lectures Notes in Computer Science, vol. 7921, P. Bonizzoni, V. Brattka, B. Loewe (Eds.), 2013, pp. 65–77.

* Corresponding author.
E-mail addresses: cicalese@dia.unisa.it (F. Cicalese), gennaro.cordasco@unina2.it (G. Cordasco), lg@dia.unisa.it (L. Gargano), martin.milanic@upr.si (M. Milanič), uv@dia.unisa.it (U. Vaccaro).

http://dx.doi.org/10.1016/j.tcs.2014.02.027
0304-3975/

Problem: Find $S \subseteq V$ s.t. $|S| \leqslant \beta$ and $|\text{Active}[S, \lambda]| \geqslant \alpha$ (or determine that no such a set exists).

We will be also interested in the case in which *a set of nodes that need to be activated* (within the given latency bound) is explicitly given as part of the input.

(λ, β, A)-Target Set Selection ((λ, β, A)-TSS).
Instance: A graph $G = (V, E)$, thresholds $t : V \longrightarrow \mathbb{N}$, a latency bound $\lambda \in \mathbb{N}$, a budget $\beta \in \mathbb{N}$ and a set to be activated $A \subseteq V$.
Problem: Find a set $S \subseteq V$ such that $|S| \leqslant \beta$ and $A \subseteq \text{Active}[S, \lambda]$ (or determine that such a set does not exist).

Eliminating any one of the parameters λ and β, one obtains two natural minimization problems. For instance, eliminating β, one obtains the following problem:

(λ, A)-Target Set Selection ((λ, A)-TSS).
Instance: A graph $G = (V, E)$, thresholds $t : V \longrightarrow \mathbb{N}$, a latency bound $\lambda \in \mathbb{N}$ and a set $A \subseteq V$.
Problem: Find a set $S \subseteq V$ of minimum size such that $A \subseteq \text{Active}[S, \lambda]$.

Notice that in the above problems we may assume without loss of generality that $0 \leqslant t(u) \leqslant d(u) + 1$ 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).

The above algorithmic problems have roots in the general study of the *spread of influence* in Social Networks (see [22] and references quoted therein). For instance, in the area of viral marketing [21,20] companies wanting to promote products or behaviors might try initially to target and convince a few individuals which, by word-of-mouth effects, can trigger a cascade of influence in the network, leading to an adoption of the products by a much larger number of individuals. Recently, viral marketing has also become an important tool in the communication strategies of politicians [33,38].

It is clear that the (λ, β, α)-TSS problem represents an abstraction of that scenario, once one makes the reasonable assumption that an individual decides to adopt the products if a certain number of his/her friends have adopted said products. Analogously, the (λ, β, α)-TSS problem can describe other diffusion problems arising in sociological, economical and biological networks, again see [22]. Therefore, it comes as no surprise that special cases of our problem (or variants thereof) have recently attracted much attention by the algorithmic community. In this version of the paper we shall limit ourselves to discuss the work which is strictly related to the present paper (we just mention that our results are also relevant to other areas, like dynamic monopolies [24,36,5], for instance). The first authors to study problems of spread of influence in networks from an algorithmic point of view were Kempe et al. [30,31]. However, they were mostly interested in networks with randomly chosen thresholds. Chen [8] studied the following minimization problem: Given a graph G and fixed thresholds $t(v)$, find a target set of minimum size that eventually activates all (or a fixed fraction of) vertices 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 the work [1,3,11]. In particular, in [3], Ben-Zwi et al. proved that the $(|V|, \beta, \alpha)$-TSS problem can be solved in time $O(t^w|V|)$ where t is the maximum threshold and w is the treewidth of the graph, thus showing that this variant of the problem is fixed-parameter tractable if parametrized w.r.t. both treewidth and the maximum degree of the graph. Paper [11] isolated other interesting cases in which the problems become tractable. The Target Set Selection (and variants thereof) has been studied also in the works [6,7,10,12,13,16,25,37,39], among others.

All the above mentioned papers did not consider the issue of the *number of rounds* necessary for the activation of the required number of vertices. However, this is a relevant question: In viral marketing, for instance, it is quite important to spread information quickly. Indeed, research in Behavioural Economics shows that humans take decisions mostly on the basis of very recent events, even though they might hold much more in their memory [2,9]. It is equally important, before embarking on a possible onerous investment, to try estimating the maximum amount of influence spread that can be guaranteed within a certain amount of time (i.e., for some λ fixed in advance), rather than simply knowing that eventually (but maybe too late) the whole market might be covered. These considerations motivate our first generalization of the problem, parametrized by the number of rounds λ. The practical relevance of parameterizing the problem also with bounds on the initial budget or the final requirement should be equally evident.

For general graphs, Chen's [8] inapproximability result still holds if one demands that the activation process ends in a bounded number of rounds. We show that the general (λ, β, α)-TSS problem is polynomially solvable in graphs of bounded clique-width and constant latency bound λ (see Theorem 1 in Section 2). Since graphs of bounded treewidth are also of bounded clique-width [17], this result implies a polynomial solution of the (λ, β, α)-TSS problem with constant λ also for graphs of bounded treewidth, complementing the result of [3] showing that for bounded-treewidth graphs, the TSS problem without the latency bound (equivalently, with $\lambda = |V| - 1$) is polynomially solvable. Moreover, the result settles the status of the computational complexity of the Vector Domination problem for graphs of bounded tree- or clique-width, which was posed as an open question in [14].

We also consider the case when G is a tree. For this special case we give an *exact linear time* algorithm for the (λ, A)-TSS problem, *for any* λ and $A \subseteq V$. When $\lambda = |V| - 1$ and $A = V$ our result is equivalent to the (optimal) linear time algorithm for the classical TSS problem (i.e., without the latency bound) on trees proposed in [8].