European Journal of Operational Research 242 (2015) 960–974

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Decision Support

Cooperation through social influence ☆

Xavier Molinero [a,1], Fabián Riquelme [b,2,*], Maria Serna [b,3]

[a] Department of Applied Mathematics III, Universitat Politècnica de Catalunya, Manresa, Spain
[b] Department of Computer Science, Universitat Politècnica de Catalunya, Barcelona, Spain

ARTICLE INFO

Article history:
Received 18 January 2013
Accepted 3 November 2014
Available online 10 November 2014

Keywords:
Spread of influence
Simple games
Influence games
Computational complexity

ABSTRACT

We consider a simple and altruistic multiagent system in which the agents are eager to perform a collective task but where their real engagement depends on the willingness to perform the task of other influential agents. We model this scenario by an *influence game*, a cooperative simple game in which a team (or coalition) of players succeeds if it is able to convince enough agents to participate in the task (to vote in favor of a decision). We take the linear threshold model as the influence model. We show first the expressiveness of influence games showing that they capture the class of simple games. Then we characterize the computational complexity of various problems on influence games, including measures (length and width), values (Shapley–Shubik and Banzhaf) and properties (of teams and players). Finally, we analyze those problems for some particular extremal cases, with respect to the propagation of influence, showing tighter complexity characterizations.

1. Introduction

Cooperation towards task execution when tasks cannot be performed by a single agent is one of the fundamental problems in both social and multiagent systems. There has been a lot of research understanding collective tasks allocation under different models coming from cooperative game theory. Under such framework, in general, cooperation is achieved by splitting the agents into teams so that each team performs a particular task and the pay-off of the team is split among the team members. Thus, cooperative game theory provides the fundamental tools to analyze this context. Among the many references we point the reader to Wooldridge and Dunne (2004), Wooldridge and Dunne (2006), Chalkiadakis, Elkind, and Wooldridge (2011), Monroy and Fernández (2011), Bachrach, Parkes, and Rosenschein (2013), and Darmann, Nicosia, Pferschy, and Schauer (2014).

The ways in which people influence each other through their interactions in a social network has received a lot of attention in the last decade. Social networks have become a huge interdisciplinary research area with important links to sociology, economics, epidemiology, computer science, and mathematics (Apt & Markakis, 2011; Easley & Kleinberg, 2010; Hellmann & Staudigl, 2014; Jackson, 2008) (players face the choice of adopting a specific product or not; users choose among competing programs from providers of mobile telephones, having the option to adopt more than one product at an extra cost, etc.). A social network can be represented by a graph where each node is an agent and each edge represents the degree of influence of one agent over another one. Several "germs" (ideas, trends, fashions, ambitions, rules, etc.) can be initiated by one or more agents and eventually adopted by the system. The mechanism defining how these germs are propagated within the network, from the influence of a small set of initially *infected* nodes, is called a model for *influence spread*.

Motivated by viral marketing and other applications the problem that has been usually studied is the *influence maximization problem* initially introduced by Domingos and Richardson (2001) and Richardson and Domingos (2002) and further developed in Kempe, Kleinberg, and Tardos (2003) and Even-Dar and Shapira (2011). This problem addresses the question of finding a set with at most k players having maximum influence, and it is NP-hard (Domingos & Richardson, 2001), unless additional restrictions are considered, in which case some generality of the problem is lost (Richardson & Domingos, 2002). Two general models for spread of influence were defined in Kempe et al. (2003): the *linear threshold model*, suggested by Granovetter (1978) and Schelling (1978), and the *independent cascade model*, created in the context of marketing by Goldenberg, Libai, and Muller (2001a) and Goldenberg, Libai, and Muller (2001b). Models for influence spread in the presence of multiple competing products have

☆ This work is partially supported by grant 2009SGR1137 (ALBCOM) of Generalitat de Catalunya.

* Corresponding author. Tel.: +34 625932189.
E-mail addresses: xavier.molinero@upc.edu (X. Molinero), farisori@lsi.upc.edu (F. Riquelme), mjserna@lsi.upc.edu (M. Serna).

1 X.M. is partially funded by grant MTM2012-34426 of the Ministry of Economy and Competitiveness.

2 F.R. is supported by grant BecasChile of the "National Commission for Scientific and Technological Research of Chile" (CONICYT).

3 M.S. is supported by grant TIN2013-46181-C2-1-R (COMMAS) of the Ministerio de Economía y Competitividad.

http://dx.doi.org/10.1016/j.ejor.2014.11.006

also been proposed and analyzed (Apt & Markakis, 2011; Bharathi, Kempe, & Salek, 2007; Borodin, Filmus, & Oren, 2010). In this setting there is also work done towards analyzing the problem from the point of view of non-cooperative game theory. Non-cooperative *influence games* were defined in 2011 by Irfan and Ortiz (2011). Those games, however, analyze the strategic behavior of two firms competing on the social network and differ from our proposal.

We propose to analyze cooperation in multiagent systems based on a model for influence among the agents in their established network of trust and influence. Social influence is relevant to determine the global behavior of a social network and thus it can be used to enforce cooperation by targeting an adequate initial set of agents. From this point of view we consider a simple and altruistic multi-agent system in which the agents are eager to perform a collective task but where their real engagement depends on the perception of the willingness to perform the task of other influential agents. We model the scenario by an *influence game*, a cooperative simple game in which a team of players (or coalition) succeeds if it is able to convince sufficiently many agents to participate in the task. We take the deterministic linear threshold model (Apt & Markakis, 2011; Chen, 2009) as the mechanism for influence spread in the associated social network.

In the considered scenario we adopt the natural point of view of decision or voting systems, mathematically modeled as *simple games* (von Neumann & Morgenstern, 1944). Simple games were first introduced by von Neumann and Morgenstern (1944) as a fundamental model for social choice. This point of view brings into the analysis several parameters and properties that are relevant in the study of simple games and thus in the analysis of the proposed scenario. Among those we consider the *length* and the *width*, two fundamental parameters that are indicators of efficiency for making a decision (Ramamurthy, 1990), or the Shapley–Shubik value (SSval) and the Banzhaf value (Bval), which provide a measure of individual influence. The properties defining *proper*, *strong* and *decisive* games have been considered in the context of simple game theory from its origins (Taylor & Zwicker, 1999) and they are also studied. Besides those properties we also consider *equivalence* and *isomorphism*. Together with properties of the games there are several properties associated to players that are of interest. Among others we consider the critical players which were used at least since 1965 by Banzhaf (1965). We refer the reader to Taylor and Zwicker (1999) for a more detailed motivation of the viewpoint of simple games and to Aziz (2009) and Chalkiadakis et al. (2011) for computational aspects of simple games and in general of cooperative game theory.

To define an influence game we take the spread of influence, in the linear threshold model, as the value that measures the power of a team. An *influence game* is described by an influence graph, modeling a social network, and a quota, indicating the required minimum number of agents that have to cooperate to perform the task successfully. Therefore, a team will be successful, or winning, if it can influence at least as many individuals as the quota requires. Such an approach reveals the importance of the influence between some players over others in order to form successful teams. In this first analysis, we draw upon the deterministic version of the linear threshold model, in which node thresholds are fixed, as our model for influence spread following (Apt & Markakis, 2011; Chen, 2009). It will be of interest to analyze influence games under other spreading models, in particular in the linear threshold model with random thresholds.

Our first result concerns the expressiveness of the family of influence games. We show that unweighted influence games capture the complete family of simple games. Although the construction can be computed in polynomial time when the simple game is given in extensive winning or minimal winning form, the number of winning or minimal winning coalitions is, in general, exponential in the number of players. Interestingly enough the formalization as weighted influence games allows a polynomial time implementation of the operations of intersection and union of weighted simple games, thus showing that, in several cases, simple games that do not admit a succinct representation as weighted games can be represented succinctly as influence games, because their (co)dimension is small.

Our second set of results settles the complexity of problems related to parameters and properties. Hardness results are obtained for unweighted influence games in which the number of agents in the network is polynomial in the number of players, while polynomial time algorithms are devised for general influence games. The new results are summarized in Table 1 as well as the known ones.

We refer the reader to Sections 2 and 4 for a formal definition of all the representations mentioned in the first row and the problems in the first column of Table 1. There P (polynomial time solvable), #PC (#P-complete), NPH (NP-hard), coNPH (coNP-hard), coNPC (coNP-complete), QP (quasi-polynomial time solvable) and GIso (the class of problems reducible to graph isomorphism) are known computational complexity classes (Garey & Johnson, 1979; Papadimitriou, 1994). The isomorphism problems for simple games, given either by $(N, \mathcal{W})$ or $(N, \mathcal{W}^m)$, are easily shown to be polynomially reducible to the graph isomorphism problem. For games given by $(N, \mathcal{W}^m)$, the Iso problem and the graph isomorphism problem are equivalent using arguments from Luks (1999).

Finally, we consider two extreme cases of influence spread in social networks for undirected and unweighted influence games. In a *maximum influence requirement*, agents adopt a behavior only when all their peers have already adopted it. This is opposed to a *minimum influence requirement* in which an agent gets convinced when at least one of its peers does. We show that, in both cases, the problems IsProper, IsStrong and IsDecisive, as well as computing Width, have polynomial time algorithms. Computing Length is NP-hard for maximum influence and polynomial time solvable for minimum influence. For the case of maximum influence and maximum spread, or minimum influence we can show that the problems IsDummy and AreSymmetric belong to P.

2. Definitions and preliminaries

Before introducing formally the family of influence games we need to define a family of labeled graphs and a process of spread of influence based on the *linear threshold model* (Granovetter, 1978; Schelling, 1978). We use standard graph notation following Bollobás (1998). As usual, given a finite set U, $\mathcal{P}(U)$ denotes its power set, and $|U|$ its cardinality. For any $0 \le k \le |U|$, $\mathcal{P}_k(U)$ denotes the subsets of U with exactly k elements. For a given graph $G = (V, E)$ we assume that $n = |V|$ and $m = |E|$. Also $G[S]$ denotes the subgraph induced by $S \subseteq V$ and, for a vertex $u \in V$, $N(u) = \{v \in V \mid (u, v) \in E\}$.

Definition 1. An *influence graph* is a tuple (G, w, f), where $G = (V, E)$ is a weighted, labeled and directed graph (without loops). As usual V is the set of vertices or agents, E is the set of edges and $w : E \to \mathbb{N}$ is a *weight function*. Finally, $f : V \to \mathbb{N}$ is a labeling function that quantifies how influenceable each agent is. An agent $i \in V$ has *influence* over another agent $j \in V$ if and only if $(i, j) \in E$. We also consider the family of *unweighted influence graphs* (G, f) in which every edge has weight 1.

Given an influence graph (G, w, f) and an initial activation set $X \subseteq V$, the *spread of influence* of X is the set $F(X) \subseteq V$ which is formed by the agents activated through an iterative process. We use $F_k(X)$ to denote the set of nodes activated at step k. Initially, at step 0, only the vertices in X are activated, that is $F_0(X) = X$. The set of vertices activated at step $i > 0$ consists of all vertices for which the total weight of the edges connecting them to nodes in $F_{i-1}(X)$ meets or exceeds their labels, i.e.,

$$F_i(X) = F_{i-1}(X) \cup \left\{ v \in V \mid \sum_{\{u \in F_{i-1}(X) | (u,v) \in E\}} w((u, v)) \ge f(v) \right\}.$$

The process stops when no additional activation occurs and the final set of activated nodes is denoted by $F(X)$.