Computing a Profit-Maximizing Sequence of Offers to Agents in a Social Network

Sayan Bhattacharya *, Dmytro Korzhyk **, and Vincent Conitzer **
Department of Computer Science, Duke University
Durham, NC 27708, USA
{dima, bsayan, conitzer}@cs.duke.edu

Abstract. Firms have ever-increasing amounts of information about possible customers available to them; furthermore, they are increasingly able to push offers to them rather than having to passively wait for a consumer to initiate contact. This opens up enormous new opportunities for intelligent marketing. In this paper, we consider the limit case in which the firm can predict consumers' preferences and relationships to each other perfectly, and has perfect control over when it makes offers to consumers. We focus on how to optimally introduce a new product into a social network of agents, when that product has significant externalities. We propose a general model to capture this problem, and prove that there is no polynomial-time approximation unless P=NP. However, in the special case where agents' relationships are symmetric and externalities are positive, we show that the problem can be solved in polynomial time.

1 Introduction

Often the utility that a person derives from a technology depends on whether her neighbors are using the same technology. Examples include various kinds of office software (calendar management, word processing, spreadsheets), mobile phones, etc. In such a context, the technology-provider may need to charge early adopters lower prices (or even give them compensations). Moreover, as firms obtain increasing amounts of data on consumers, they are able to individualize offers to them, in terms of both the timing of the offer and price quoted. This results in a challenging optimization problem for the provider: choose intelligently to which agents to make offers, and in which order.

We assume that a new provider is introducing a single new technology. There may be competing technologies in the market, but in any case the existing situation is static. This rules out possibilities such as existing providers modifying their own prices or otherwise acting in response to the new provider's actions. We also assume that the agents are myopically rational: when made an offer, an agent decides on the offer based on the technologies currently used by her neighbors. The agent does not attempt to predict whether her neighbors will later switch technologies themselves. Finally, we restrict ourselves to situations where the new provider can perfectly predict how much an agent is be willing to pay.

* Bhattacharya was supported by NSF awards CCF-0745761 and CCF-1008065.
** Korzhyk and Conitzer were supported by NSF Awards IIS-0812113, IIS-0953756, and CCF-1101659, as well as ARO MURI Grant W911NF-11-1-0332 and an Alfred P. Sloan fellowship.

We show that the general problem is hard to approximate unless $P = NP$ (Section 3). However, in an interesting special case where the agents have symmetric utilities and positive externalities, the problem can be formulated as an integer program whose constraint matrix is totally unimodular. Hence, we get a polynomial time algorithm (Section 4).

Previous Work. There is an extensive literature on marketing policies over a social network [7]. The generic setting is as follows. Initially, the firm convinces a certain subset of agents to use the new technology and those agents, in turn, influence their neighbors. The process continues, and more agents adopt the new technology due to a cascading effect. A standard objective [11] is to select an initial subset of at most k agents so as to maximize the *influence*, which is defined as the total number of agents who adopt the new technology at the end of the cascading process.

In contrast to the *influence* maximization, we optimize the *profit* over a social network [3, 9]. The two papers [10, 4] are particularly relevant to our setting. They consider a Bayesian model. Here, an agent's valuation for the new technology is private knowledge, but it is drawn from a publicly known distribution. This distribution depends on the subset of her neighbors who have already switched to the new technology. The new firm visits the agents one by one, and while visiting an agent, it offers her the new technology at some price. The agents behave myopically, and the objective is to maximize the expected sum of total payments collected from all the agents. The authors give simple *influence and exploit* policies that are constant factor approximations to optimal profit: In the first stage, a select subset of agents gets the new technology for free. In the next stage, the remaining agents are visited in a sequence chosen uniformly at random, and each of those agents is offered the new technology at the myopically optimal price. Our work is different from these results in three crucial aspects: 1) Unlike these previous papers, we consider a perfect-information (non-Bayesian) setting. 2) In our model, the firm incurs a nonnegative cost for producing each unit of the product, and the objective is to maximize the total payments made by the agents *minus* the total production cost. Hence, marketing policies that make offers to a large subset of agents at low prices can be extremely suboptimal. 3) We allow the agents to have positive utilities for being in the initial state, which captures settings where an existing technology is already in use, and our firm wants to enter the market and compete with an incumbent.

2 The Problem: OPTIMAL-OFFER-SEQUENCE

Consider a simple undirected graph $G = (V, E)$. Every node $i \in V$ denotes an agent, and there is an edge $\{i, j\} \in E$ iff $i \neq j$ and i and j are neighbors. Initially, every agent $i \in V$ is in state $\mathcal{A}$. A new firm (say $\mathcal{B}$) now wants to enter the market, and its objective is to maximize profit by exploiting the network structure. If some agent $i \in V$ decides to be a customer of firm $\mathcal{B}$, then we say that agent i *switches* (or *converts*) to state $\mathcal{B}$.

The vector $\mathbf{S}$ captures the states of all the agents at any particular instant. Component $i \in V$ of vector $\mathbf{S}$ is denoted by $\mathbf{S}_i$, and the notation $\mathbf{S}_{-i}$ denotes all the components *except* component i. Specifically, we set $\mathbf{S}_i = \mathcal{A}$ (resp. $\mathbf{S}_i = \mathcal{B}$) iff agent i is in state $\mathcal{A}$ (resp. state $\mathcal{B}$). Let $U_i(\mathbf{S})$ be the utility of agent $i \in V$. It is a function of the state vector, and can be expressed as the sum of two terms:

$$U_i(\mathbf{S}) = In_i(\mathbf{S}_i) + \mathit{\Gamma}_i(\mathbf{S}_i, \mathbf{S_{-i}}) \tag{1}$$

In the above equation, the term $In_i(\mathbf{S_i})$ denotes the *intrinsic* utility agent $i \in V$ derives from being in state $\mathbf{S_i}$; whereas her *extrinsic* utility is captured by the term $\mathit{\Gamma}_i(\mathbf{S_i}, \mathbf{S_{-i}})$ and it is determined in the following manner. Let $\mathit{\Phi}_{t,t'}(i,j)$ be the (non-negative) utility agent i derives from her friend j, when i is in state $t \in \{\mathcal{A}, \mathcal{B}\}$ and j is in state $t' \in \{\mathcal{A}, \mathcal{B}\}$. In general, these utilities may be *asymmetric*, that is, we may have $\mathit{\Phi}_{t,t'}(i,j) \neq \mathit{\Phi}_{t',t}(j,i)$. For all $t, t' \in \{\mathcal{A}, \mathcal{B}\}$ and $i, j \in V$, we set $\mathit{\Phi}_{t,t'}(i,j) = 0$ if the agents i, j are not friends with each other. Now:

$$\mathit{\Gamma}_i(\mathbf{S_i}, \mathbf{S_{-i}}) = \sum_{\mathbf{j \in V}} \mathbf{\Phi}_{\mathbf{S_i}, \mathbf{S_j}}(\mathbf{i}, \mathbf{j}) \tag{2}$$

Initially, every agent is in state $\mathcal{A}$. Next, firm $\mathcal{B}$ selects a subset $V^* \subseteq V$, and computes a *ranking* $\pi : V^* \rightarrow \{1, \ldots, |V^*|\}$ of the agents in V^*. The rank of agent $i \in V^*$ is given by $\pi(i)$. Firm $\mathcal{B}$ now *visits* the agents in V^* in increasing order of their ranks. While visiting an agent i, firm $\mathcal{B}$ offers her the new technology at a price p_i.

Without any loss of generality, we can assume that every agent $i \in V^*$ accepts her offer.[1] Let $\mathbf{S}$ be the state vector just before firm $\mathcal{B}$ makes an offer to agent i. Agent i behaves myopically and utilities are quasilinear. Hence, if she is to switch her state, then we must have: $In_i(\mathcal{B}) + \mathit{\Gamma}_i(\mathcal{B}, \mathbf{S}_{-i}) - p_i \geq In_i(\mathcal{A}) + \mathit{\Gamma}_i(\mathcal{A}, \mathbf{S_{-i}})$. Since firm $\mathcal{B}$ wants to maximize its profit, it sets p_i to the highest possible value. Thus, we have:

$$p_i = In_i(\mathcal{B}) + \mathit{\Gamma}_i(\mathcal{B}, \mathbf{S}_{-i}) - In_i(\mathcal{A}) - \mathit{\Gamma}_i(\mathcal{A}, \mathbf{S_{-i}}) \tag{3}$$

The price p_i can be negative, which implies a subsidy. The idea is that firm $\mathcal{B}$ may have to subsidize some agents in the beginning, when few agents are in state $\mathcal{B}$ and they may incur a loss for switching to the new technology. As more and more agents convert to state $\mathcal{B}$, the firm will be able to exploit the resulting positive externalities and generate a large profit, due to the customers who switch in later stages. Firm $\mathcal{B}$ also incurs a manufacturing cost of c per unit of the product. We want to maximize its net profit, given by the expression $\sum_{i \in V^*}(p_i - c)$. Throughout the rest of the paper, we refer to this optimization problem as OPTIMAL-OFFER-SEQUENCE.

Lemma 1. *Let* PROFIT(j) *be the profit from agent* j*. For all* $i \in V^*$*, let* $\pi_-(i)$ *be the set of agents switching to state* $\mathcal{B}$ before *agent* i*, i.e.,* $\pi_-(i) = \{j \in V^* : \pi(j) < \pi(i)\}$*.*

$$\textsc{Profit}(i) = \begin{cases} 0, & \text{if } i \in V \setminus V^*. \\ \big(In_i(\mathcal{B}) - In_i(\mathcal{A}) - c\big) + \sum_{j \in \pi_-(i)} \big(\mathit{\Phi}_{\mathcal{B},\mathcal{B}}(i,j) - \mathit{\Phi}_{\mathcal{A},\mathcal{B}}(i,j)\big) & \\ + \sum_{j \in V \setminus \pi_-(i)} \big(\mathit{\Phi}_{\mathcal{B},\mathcal{A}}(i,j) - \mathit{\Phi}_{\mathcal{A},\mathcal{A}}(i,j)\big), & \\ & \text{if } i \in V^*. \end{cases}$$

The total profit of firm $\mathcal{B}$ *is given by:* $\sum_{i \in V}$ PROFIT$(i) = \sum_{i \in V^*}$ PROFIT(i).

[1] Otherwise, we could delete agent i from the set V^*.

Proof. Fix any agent $i \in V^*$. Note that $\textsc{Profit}(i) = p_i - c$. Let $\mathbf{S}$ be the state vector just before i switches to state $\mathcal{B}$. By Equation 3, $\textsc{Profit}(i)$ is equal to:

$$In_i(\mathcal{B}) - In_i(\mathcal{A}) - c + \Gamma_i(\mathcal{B}, \mathbf{S}_{-i}) - \Gamma_i(\mathcal{A}, \mathbf{S}_{-i}) \tag{4}$$

Expanding the right hand side of Equation 2, we can show:

$$\Gamma_i(\mathcal{B}, \mathbf{S}_{-i}) = \sum_{j \in \pi_-(i)} \Phi_{\mathcal{B},\mathcal{B}}(i,j) + \sum_{j \in V \setminus \pi_-(i)} \Phi_{\mathcal{B},\mathcal{A}}(i,j)$$

$$\Gamma_i(\mathcal{A}, \mathbf{S}_{-i}) = \sum_{j \in \pi_-(i)} \Phi_{\mathcal{A},\mathcal{B}}(i,j) + \sum_{j \in V \setminus \pi_-(i)} \Phi_{\mathcal{A},\mathcal{A}}(i,j)$$

Finally, we substitute the above expressions back in Eq. 4. □

3 A Hardness Result

In this section, we show that (see Lemma 3) it is NP-hard to decide whether firm $\mathcal{B}$ can make positive profit, by a reduction from the Maximum Arc Set on Tournaments (MAST) problem. This rules out the existence of any polynomial-time approximation algorithm for OPTIMAL-OFFER-SEQUENCE, unless $P = NP$ (see Theorem 1).

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a directed tournament graph; that is, for any two distinct nodes $i, j \in \mathcal{V}$, we have $|\mathcal{E} \cap \{(i,j),(j,i)\}| = 1$. Let $\pi : \mathcal{V} \to \{1, \ldots, |\mathcal{V}|\}$ be a ranking of the set of nodes $\mathcal{V}$, where $\pi(i)$ denotes the rank of node $i \in \mathcal{V}$, and $\pi(i) \neq \pi(j)$ if $i \neq j$. We say that an edge $(i,j) \in \mathcal{E}$ is a *forward edge* (resp. *backward edge*) w.r.t. ranking π if $\pi(i) < \pi(j)$ (resp. $\pi(i) > \pi(j)$).

Maximum Acyclic Subgraph on Tournaments (MAST): An instance $\mathcal{F}$ of the problem consists of an ordered pair $(\mathcal{G}, \theta)$, where $\theta \geq 1$ is a positive integer, and $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a directed tournament graph. The objective is to decide if there exists a ranking of $\mathcal{V}$ where the number of backward edges is at least θ. This problem is NP-hard [6, 2, 5, 1].

The Reduction. Given an instance $\mathcal{F}$ of the MAST problem $(\mathcal{G} = (\mathcal{V}, \mathcal{E}), \theta)$, we construct the following instance $\mathcal{I}_{\mathcal{F}}$ of OPTIMAL-OFFER-SEQUENCE. It is easy to see that the reduction can be implemented in polynomial time.

- $G = (V, E)$ is a complete undirected graph, defined on the same node set as that of $\mathcal{G}$; that is, $V = \mathcal{V}$ and $E = \{\{i,j\} : i, j \in V, i \neq j\}$.
- For all $i, j \in V$: if $(i,j) \in \mathcal{E}$ then $\Phi_{\mathcal{B},\mathcal{B}}(i,j) = 1$, else $\Phi_{\mathcal{B},\mathcal{B}}(i,j) = 0$.
- For all $i, j \in V$: we have $\Phi_{\mathcal{A},\mathcal{B}}(i,j) = \Phi_{\mathcal{B},\mathcal{A}}(i,j) = \Phi_{\mathcal{A},\mathcal{A}}(i,j) = 0$.
- For all $i \in V$: we set $In_i(\mathcal{A}) = In_i(\mathcal{B}) = 0$.
- The cost per unit c is set in such a way that

$$-c \times |V| + \theta = 1 \tag{5}$$

According to the above reduction, the profit (Lemma 1) from the instance $\mathcal{I}_{\mathcal{F}}$ equals:

$$\textsc{Profit} = -c|V^*| + \sum_{i \in V^*} \sum_{j \in \pi_-(i)} \Phi_{\mathcal{B},\mathcal{B}}(i,j) \tag{6}$$

Let $\mathcal{G}[V^*] = (V^*, \mathcal{E}^*)$ be the subgraph of $\mathcal{G}$ induced by the node set $V^* \subseteq V$, so that:

$$\mathcal{E}^* = \{(i,j) \in \mathcal{E} : i,j \in V^*, i \neq j\} \tag{7}$$

Let $\mathcal{E}^*_\pi$ be the set of backward edges in $\mathcal{G}[V^*]$ w.r.t. π. Since $\Phi_{\mathcal{B},\mathcal{B}}(i,j) = 1$ when $(i,j) \in \mathcal{E}$, and $\Phi_{\mathcal{B},\mathcal{B}}(i,j) = 0$ when $(i,j) \notin \mathcal{E}$, Equation 6 implies that

$$\text{PROFIT} = -c|V^*| + \sum_{i \in V^*} \sum_{j \in \pi_-(i)} \Phi_{\mathcal{B},\mathcal{B}}(i,j) = -c|V^*| + |\mathcal{E}^*_\pi| \tag{8}$$

Lemma 2. *In the instance $\mathcal{I}_\mathcal{F}$ of* OPTIMAL-OFFER-SEQUENCE*, the profit-maximizing solution either converts* all *the agents to state $\mathcal{B}$, or it does not convert* any *agent to state $\mathcal{B}$; that is, it sets either $V^* = \emptyset$ or $V^* = V$.*

Proof. In the profit-maximizing solution, suppose that the agents in V^* switch to state $\mathcal{B}$ according to the ranking $\pi : V^* \to \{1, \ldots, |V^*|\}$. For the sake of contradiction, suppose that the lemma is false, and the profit-maximizing solution sets $\emptyset \subset V^* \subset V$. Since the profit is nonnegative, Equation 8 implies that $-c|V^*| + |\mathcal{E}^*_\pi| \geq 0$, or equivalently, $c \leq |\mathcal{E}^*_\pi|/|V^*|$. Since $|\mathcal{E}^*_\pi| \leq \binom{|V^*|}{2}$, we derive $c < |V^*|/2$.

Fix any $k \in V \setminus V^*$. Let $\delta^+(k, V^*)$ (resp. $\delta^-(k, V^*)$) be the number of outgoing (resp. incoming) edges of k whose other endpoints lie in V^*. Since the graph $\mathcal{G}$ is a tournament, either $\delta^-(k, V^*) \geq |V^*|/2$ or $\delta^+(k, V^*) \geq |V^*|/2$.

Case 1. $\delta^-(k, V^*) \geq |V^*|/2$.
In this case, we construct a new solution that converts all the nodes in $V^* \cup \{k\}$ to state $\mathcal{B}$ in the following order: First, it converts node k. Next, it converts the nodes in V^* according to ranking π. Let the new profit be P′. Clearly, we have:

$$\text{P}' = -c(|V^*| + 1) + \delta^-(k, V^*) + |\mathcal{E}^*_\pi| > -c|V^*| + |\mathcal{E}^*_\pi|$$

The inequality holds since $c < |V^*|/2$ and $\delta^-(k, V^*) \geq |V^*|/2$. Thus, the new profit is strictly greater than the maximum profit, which is a contradiction.

Case 2. $\delta^+(k, V^*) \geq |V^*|/2$.
In this case, we construct another solution that converts all the nodes in $V^* \cup \{k\}$ to state $\mathcal{B}$ in the following order: First, it converts the nodes in V^* according to ranking π. Next, it converts node k. Applying an argument similar to Case 1, we show that the new profit is strictly greater than the maximum profit, which is a contradiction. □

Lemma 3. *Firm $\mathcal{B}$ can get positive profit from the instance $\mathcal{I}_\mathcal{F}$ of the* OPTIMAL-OFFER-SEQUENCE *problem* if and only if *the instance $\mathcal{F}$ of the MAST problem admits a ranking where the number of backward edges is at least θ.*

Proof. Suppose that the optimal solution to the instance $\mathcal{I}_\mathcal{F}$ converts the agents in $V^* \subseteq V$ to state $\mathcal{B}$ according to the ranking π. Lemma 2 implies that it is possible to get positive profit from the instance $\mathcal{I}_\mathcal{F}$ iff $V^* = V$, and in that case, applying Equation 8:

$$\text{PROFIT} = -c|V| + |\mathcal{E}^*_\pi| = 1 - \theta + |\mathcal{E}^*_\pi| > 0.$$

The second equality holds because of Equation 5. Since θ is an integer, $1 - \theta + |\mathcal{E}^*_\pi| > 0$ iff $|\mathcal{E}^*_\pi| \geq \theta$. Since π is also a ranking for the MAST instance $\mathcal{F}$, the lemma follows.

Lemma 3 implies Theorem 1.

Theorem 1. *The* OPTIMAL-OFFER-SEQUENCE *problem does not admit any polynomial-time approximation algorithm, unless $P = NP$.*

Next, we describe a family of instances that admit a 2-approximation in poly-time. Theorem 2 follows from a result by Guruswami et al. [8].

Theorem 2. *Consider a family of instances of the* OPTIMAL-OFFER-SEQUENCE *problem where $c = 0$, $In_i(\mathcal{A}) = In_i(\mathcal{B}) = 0$ for all $i \in V$, and $\Phi_{\mathcal{A},\mathcal{B}}(i,j) = \Phi_{\mathcal{B},\mathcal{A}}(i,j) = \Phi_{\mathcal{A},\mathcal{A}}(i,j) = 0$ and $\Phi_{\mathcal{B},\mathcal{B}}(i,j) \geq 0$ for all $i, j \in V$. Under such settings, there exists a poly-time 2 approximation algorithm for the* OPTIMAL-OFFER-SEQUENCE *problem, and it is Unique Games hard to get better than 2 approximation.*

4 Symmetric Utility Functions: Polynomial Time Algorithm

In this section, for all $\{i,j\} \in E$, we require that $\Phi_{\mathcal{A},\mathcal{A}}(i,j) = \Phi_{\mathcal{A},\mathcal{A}}(j,i)$, $\Phi_{\mathcal{B},\mathcal{B}}(i,j) = \Phi_{\mathcal{B},\mathcal{B}}(j,i)$, and $\Phi_{\mathcal{A},\mathcal{B}}(i,j) = \Phi_{\mathcal{B},\mathcal{A}}(j,i) = 0$. Such utility functions are *symmetric*, and we write $\Phi_{\mathcal{A},\mathcal{A}}(\{i,j\})$ and $\Phi_{\mathcal{B},\mathcal{B}}(\{i,j\})$ instead of $\Phi_{\mathcal{A},\mathcal{A}}(i,j)$ and $\Phi_{\mathcal{B},\mathcal{B}}(i,j)$. Under symmetric utilities, the problem can be solved in polynomial time (see Theorem 3).

Lemma 4. *If the utility functions are symmetric, then the profit of firm $\mathcal{B}$ is given by:*

$$\sum_{i \in V^*} \big(In_i(\mathcal{B}) - In_i(\mathcal{A}) - c\big) + \sum_{\{i,j\} \subseteq V^*} \Phi_{\mathcal{B},\mathcal{B}}(\{i,j\}) - \sum_{\{i,j\} \cap V^* \neq \emptyset} \Phi_{\mathcal{A},\mathcal{A}}(\{i,j\})$$

Proof. Since the utility functions are symmetric, we have:

$$\sum_{i \in V^*} \sum_{j \in \pi_-(i)} \big(\Phi_{\mathcal{B},\mathcal{B}}(i,j) - \Phi_{\mathcal{A},\mathcal{B}}(i,j)\big) = \sum_{\{i,j\} \subseteq V^*} \Phi_{\mathcal{B},\mathcal{B}}(\{i,j\}) \tag{9}$$

$$\sum_{i \in V^*} \sum_{j \in V \setminus \pi_-(i)} \big(\Phi_{\mathcal{B},\mathcal{A}}(i,j) - \Phi_{\mathcal{A},\mathcal{A}}(i,j)\big) = - \sum_{\{i,j\} \cap V^* \neq \emptyset} \Phi_{\mathcal{A},\mathcal{A}}(\{i,j\}) \tag{10}$$

The lemma follows from Equations 9, 10 and Lemma 1. □

Lemma 4 implies that the profit of firm $\mathcal{B}$, under symmetric utility functions, is uniquely determined by the set of agents who switch to state $\mathcal{B}$, and is independent of the order in which those agents are offered the new technology. We now give an integer programming formulation (IP-1) for our problem. Note that in IP-1, the variables $\gamma_{\{i,j\}}, \lambda_{\{i,j\}}$ are defined over *unordered* pairs of nodes $\{i,j\} \in E$.

IP-1

$$\text{Max.} \sum_{i \in V} \big(In_i(\mathcal{B}) - In_i(\mathcal{A}) - c\big)x_i + \sum_{\{i,j\}} \big(\Phi_{\mathcal{B}\mathcal{B}}(\{i,j\})\gamma_{\{i,j\}} - \Phi_{\mathcal{A}\mathcal{A}}(\{i,j\})\lambda_{\{i,j\}}\big)$$

$$
\begin{aligned}
\text{s.t.} \qquad \gamma_{\{i,j\}} - x_i &\le 0 \quad && \forall i \in V, \{i,j\} \in E && (11)\\
x_i - \lambda_{\{i,j\}} &\le 0 && \forall i \in V, \{i,j\} \in E && (12)\\
x_i &\in \{0,1\} && \forall i \in V && (13)\\
\gamma_{\{i,j\}}, \lambda_{\{i,j\}} &\in \{0,1\} && \forall \{i,j\} \in E && (14)
\end{aligned}
$$

Lemma 5. *The constraints of IP-1 ensure that in an optimal solution:*

- *The variable $x_i = 1$ iff node $i \in V$ switches to state $\mathcal{B}$, that is, when $i \in V^*$.*
- *The variable $\gamma_{\{i,j\}} = 1$ iff both the endpoints of edge $\{i,j\}$ switch to state $\mathcal{B}$.*
- *The variable $\lambda_{\{i,j\}} = 1$ iff at least one endpoint of edge $\{i,j\}$ switches to state $\mathcal{B}$.*

Hence, Lemma 4 implies that IP-1 gives an integer programming formulation of the OPTIMAL-OFFER-SEQUENCE *problem in the special case of symmetric utilities.*

Proof. We show that the interpretation of $\gamma_{\{i,j\}}$ is consistent with the interpretation of x_i. Each $\gamma_{\{i,j\}}$ has a nonnegative coefficient in the objective. Hence, in an optimal solution, $\gamma_{\{i,j\}}$ is set to the largest possible value. Constraint 11 establishes an upper bound of $\min(x_i, x_j)$ on the variable $\gamma_{\{i,j\}}$. It follows that $\gamma_{\{i,j\}} = 1$ iff $x_i = x_j = 1$.

Each $\lambda_{\{i,j\}}$ has a nonpositive coefficient in the objective. Thus, in an optimal solution, $\lambda_{\{i,j\}}$ is set to the smallest possible value. Constraint 12 establishes a lower bound of $\max(x_i, x_j)$ on the variable $\gamma_{\{i,j\}}$. Hence, $\lambda_{\{i,j\}}$ is set to 0 iff $x_i = x_j = 0$. $\square$

Theorem 3. *The constraint matrix of IP-1 is totally unimodular. Hence, we can find an optimal solution of IP-1 in polynomial time. Thus, the* OPTIMAL-OFFER-SEQUENCE *problem can be solved efficiently when the utility functions are symmetric.*

References

[1] N. Ailon, M. Charikar, and A. Newman. Aggregating inconsistent information: ranking and clustering. In *STOC*, pages 684–693, 2005.
[2] N. Alon. Ranking tournaments. *SIAM J. Discrete Math.*, 20(1):137–142, 2006.
[3] D. Arthur, R. Motwani, A. Sharma, and Y. Xu. Pricing strategies for viral marketing on social networks. In *WINE*, pages 101–112, 2009.
[4] M. Babaei, B. Mirzasoleiman, M. Safari, and M. Jalili. Revenue maximization in social networks through discounting. *Working Paper*, 2012.
[5] P. Charbit, S. Thomassé, and A. Yeo. The minimum feedback arc set problem is NP-hard for tournaments. *Combinatorics, Probability & Computing*, 16(1):1–4, 2007.
[6] V. Conitzer. Computing Slater rankings using similarities among candidates. In *AAAI*, 2006.
[7] P. Domingos and M. Richardson. Mining the network value of customers. In *KDD*, pages 57–66, 2001.
[8] V. Guruswami, R. Manokaran, and P. Raghavendra. Beating the random ordering is hard: Inapproximability of maximum acyclic subgraph. In *FOCS*, pages 573–582, 2008.
[9] N. Haghpanah, N. Immorlica, V. S. Mirrokni, and K. Munagala. Optimal auctions with positive network externalities. In *ACM EC*, pages 11–20, 2011.
[10] J. D. Hartline, V. S. Mirrokni, and M. Sundararajan. Optimal marketing strategies over social networks. In *WWW*, pages 189–198, 2008.
[11] D. Kempe, J. M. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *KDD*, pages 137–146, 2003.