Theoretical Computer Science 539 (2014) 68–86

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

On the efficiency of Influence-and-Exploit strategies for revenue maximization under positive externalities [☆]

Dimitris Fotakis [a,*], Paris Siminelakis [b,1]

[a] School of Electrical and Computer Engineering, National Technical University of Athens, 15780 Athens, Greece
[b] Department of Electrical Engineering, Stanford University, Stanford, CA 94304, United States

ARTICLE INFO

Article history:
Received 11 September 2013
Received in revised form 11 February 2014
Accepted 22 April 2014
Communicated by X. Deng

Keywords:
Pricing
Externalities
Social network monetization
Influence and Exploit strategies

ABSTRACT

The mitigated effectiveness of traditional forms of advertising along with winner-take-all phenomena caused by globalization and the Internet necessitates a new approach in marketing. Hartline et al. (2008) [16] introduced a marketing model for social networks, where a seller is trying to exploit positive externalities between the buyers and to maximize his revenue by designing an intelligent series of individualized offers. Under this setting, we study the problem of revenue maximization and mostly focus on Influence-and-Exploit (IE) marketing strategies. We show that in undirected social networks, revenue maximization is **NP**-hard not only when we search for an optimal marketing strategy, but also when we search for the best IE strategy. Rather surprisingly, we observe that allowing IE strategies to offer prices smaller than the myopic price in the exploit step leads to a significant improvement on their performance. Thus, we show that the best IE strategy approximates the maximum revenue within a factor of 0.911 for undirected and of roughly 0.553 for directed social networks. Utilizing a connection between good IE strategies and large cuts in the underlying social network, we obtain polynomial-time algorithms that approximate the revenue of the best IE strategy within a factor of roughly 0.9. Hence, we significantly improve on the best known approximation ratio for revenue maximization to 0.8229 for undirected and to 0.5011 for directed networks (from 2/3 and 1/3, respectively).

1. Introduction

Understanding the flow of information, influence and epidemics through the social fabric has become increasingly important due to the high interconnectedness brought about by technological advances. The digitization of communications (cell phones, emails, text messages) and social interaction (Facebook, Twitter, Second Life) has provided researchers not only with a strong empirical footing upon which to base their theories and test their predictions, but also has opened

☆ This research was supported by the project AlgoNow, co-financed by the European Union (European Social Fund – ESF) and Greek national funds, through the Operational Program "Education and Lifelong Learning" of the National Strategic Reference Framework (NSRF) – Research Funding Program: THALES, investing in knowledge society through the European Social Fund. Part of this work was done while P. Siminelakis was with the School of Electrical and Computer Engineering, National Technical University of Athens, Greece. An extended abstract of this work appeared in the Proceedings of the 8th International Workshop on Internet and Network Economics (WINE 2012), Paul W. Goldberg (Editor), Lecture Notes in Computer Science 7695, pp. 270–283, Springer, 2012.
* Corresponding author.
E-mail addresses: fotakis@cs.ntua.gr (D. Fotakis), psimin@stanford.edu (P. Siminelakis).
1 This author is partially supported by a scholarship from the Onassis Foundation.

http://dx.doi.org/10.1016/j.tcs.2014.04.026
0304-3975/

the frontier of direct algorithmic applications. Particularly, there has been a shift from aggregate descriptive theories, in the spirit of *Diffusion of Innovations*, to models incorporating the micro-structure of social networks, culminating with the algorithmic paradigm of *Influence Maximization*.

Firms operating in such a reticular environment, where information about products and services diffuses rapidly between individuals, have acknowledged the importance of revisiting their approach. The availability of information about users and the mitigated effectiveness of traditional forms of marketing occasion the need for more intelligent marketing strategies. Towards realizing that goal, there are three significant challenges: mining individual preferences, quantifying the influence that buyers exert upon each other, and fusing these information along a promotion process. The ideal is an algorithm that intelligently adjusts itself (prices, individuals to approach) based on the current state of the network aiming to maximize the sellers revenue.

In this work, we are interested in the latter goal of designing a marketing strategy that exploits the influence between buyers to maximize the seller's revenue. We focus on the setting where the utility of the product depends inherently on the scale of the product's adoption (e.g., operating systems, VCR format) and particularly on the case where buyers' valuations depend on the specific set of their friends using the product (e.g., cell phones, on-line gaming). For instance, if 4 of 5 closest friends of an individual share the same mobile phone company, she might be willing to pay a higher monthly rate to join that company instead of choosing a cheaper one, due to the special intracompany discounts (free calls or text messages). But, if it happens that the one friend left out is her partner, then again it might be wise to join the same company. This example illustrates the complexity of and the extent that network externalities can influence the final decision.

Modeling network externalities is inarguably a non-trivial matter. The networked nature of the process along with the corresponding combinatorial explosion necessitate a trade-off between realism and tractability. The purpose of any mathematical model is to capture the few details that really matter and abstract the rest. Towards that goal, we believe that the crucial features are:

(i) *Individuation*: buyers should be allowed to have different valuations for the same product and be influenced in different amounts by their social contacts.
(ii) *Uncertainty*: social interaction is itself stochastic in nature and the seller in practice can only estimate the model parameters(buyer's valuations).
(iii) *Succinct representation*: buyers do not have a hard time deciding how much they are willing to pay for the product, which means that their internal representation of the product's value should have a simple form.

1.1. The marketing model and revenue maximization

In this work, we adopt the model of Hartline, Mirrokni, and Sundararajan [16], where a digital product is sold to a set of potential buyers under positive externalities. We assume an unlimited supply of the product and that there is no production cost for it. A (possibly directed) weighted social network $G(V, E, w)$ on the set V of potential buyers models how buyers' value of the product is affected by other buyers who already own the product. Specifically, an edge $(j, i) \in E$ denotes that the event that j owns the product has a positive influence on i's value of the product. The strength of this influence is quantified by a non-negative weight w_{ji} associated with edge (j, i) (we assume that $w_{ji} = 0$ if $(j, i) \notin E$). Also, buyer i may have an intrinsic value of the product,[2] quantified by a non-negative weight w_{ii}. The model incorporates the aforementioned features as follows:

(i) *Individuation*: The product's value to each buyer i is given by a non-decreasing function $v_i : 2^{N_i} \mapsto \mathbb{R}_+$, which depends on w_{ii} and on the set $S \subseteq N_i$ of i's neighbors who already own the product, where $N_i = \{j \in V \setminus \{i\} : (j, i) \in E\}$.
(ii) *Uncertainty*: The exact values $v_i(S)$ are unknown and are treated as random variables of which only the distributions $F_{i,S}$ are known to the seller. In particular, we assume that for each buyer i and each set $S \subseteq N_i$, the seller only knows the probability distribution $F_{i,S}(x) = \mathbb{P}\mathrm{r}[v_i(S) < x]$ that buyer i rejects an offer of price x for the product.
(iii) *Succinct representation*: Regarding the distribution of $v_i(S)$'s, the most interesting cases outlined in [16] are the *Concave Graph Model*, where the weights w_{ji} are random variables, and each $v_i(S)$ is a concave function of the total influence $M_{i,S} = \sum_{j \in S \cup \{i\}} w_{ji}$ perceived by buyer i from the set S of her neighbors owning the product, and the *Uniform Additive Model*, where weights w_{ji} are deterministic, and each $v_i(S)$ is uniformly distributed in $[0, M_{i,S}]$. Hence, in the Uniform Additive Model the probability that buyer i rejects an offer of price x is $F_{i,S}(x) = x/M_{i,S}$.

Marketing strategies. We assume that the seller approaches each potential buyer once and makes an individualized offer to him. Thus, a *marketing strategy* $(\vec{\pi}, \vec{x})$ consists of a permutation $\vec{\pi}$ of the buyers and a pricing vector $\vec{x} = (x_1, \ldots, x_n)$, where $\vec{\pi}$ determines the order in which the buyers are approached and $\vec{x}$ the prices offered to them. Given the set S of i's neighbors who own the product when the seller approaches her, buyer i accepts the offer with probability $1 - F_{i,S}(x_i)$, in which case she pays the price x_i, or rejects it, with probability $F_{i,S}(x_i)$, in which case she pays nothing and never receives

[2] For simplicity, we ignore w_{ii}'s for directed social networks. This is without loss of generality, since we can replace each w_{ii} by an edge (i', i) of weight w_{ii} from a new node i' with a single outgoing edge (i', i) and no incoming edges.