Cascade Model with Contextual Externalities and Bounded User Memory for Sponsored Search Auctions*

(Extended Abstract)

Nicola Gatti
Politecnico di Milano
nicola.gatti@polimi.it

Marco Rocco
Politecnico di Milano
marco.rocco@polimi.it

Paolo Serafino
Teesside University
p.serafino@tees.ac.uk

Carmine Ventre
Teesside University
c.ventre@tees.ac.uk

ABSTRACT

In spite of the considerable research effort devoted to studying externalities in Sponsored Search Auctions (SSAs), even the basic question of modeling the problem has so far escaped a definitive answer. The popular cascade model appears too idealized to really describe the phenomenon yet it allows a good comprehension of the problem. Other models, instead, arguably describe the real setting more adequately but are too complex to permit a satisfactory theoretical analysis. In this work, we attempt to get the best of both approaches: (i) we generalize the cascade model along a number of directions in the attempt to have mathematical formulations that are close to SSAs in the real world and (ii) prove a host of results drawing a nearly complete picture about the computational complexity of the problem. We complement these approximability results with some considerations about mechanism design in our context.

Categories and Subject Descriptors

F.2.0 [**Analysis of Algorithms and Problem Complexity**]: General; J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Algorithms, Theory, Economics

Keywords

Algorithmic Mechanism Design; Sponsored Search Auctions

1. INTRODUCTION

The most adopted solution in practice for SSAs is the *generalized second price auction* (GSP), which is known to be *non–truthful*, i.e., advertisers can play the search engine in order to get a better outcome from their perspective. The GSP is a *social welfare* maximizing auction, i.e., maximizing the total "happiness" of the advertisers. The crucial point here is the definition of happiness: the more clicks their ads receive, the more content advertisers are. A measure to forecast clicks, named *click through rate* (CTR), must necessarily take into account the behavior of the user, exploiting information about the slot in which the ad is shown, the quality of the ad itself and cannot overlook the relative influences between ads, a.k.a. *externalities*.

The question "on which ads will users click?" is commonly tackled by designing a *user model* that describes how users "move" over the slots. In this respect, on one hand of the scale there is the simple, yet neat, *cascade* model [1, 4], where users are assumed to scan the ads *sequentially* from top to bottom and the probability with which a user clicks on the ad a_i shown in slot s_m is the product of an intrinsic quality q_i of the ad, the relevance λ_m of slot s_m (sometimes, called *slot-dependent externality*) and of *all* the ads allocated in slots s_1 through s_{m-1}. In its more general version, the optimization problem of social welfare maximization is conjectured to be NP-hard, shown to be in APX (i.e., a 1/4-approximation algorithm is given) and to admit a QPTAS (quasi-polynomial time approximation scheme) [4]. The cascade model has two main limitations to be considered a satisfactory model of externalities in SSAs. Firstly, it assumes that users have *unlimited memory* and, consequently, an ad in slot s_1 can exert externalities on an ad many slots below. This is quite unrealistic, as proven experimentally in [3]. Secondly, it assumes that the externality of an ad is the same no matter what ad is exerted on. Nevertheless, while, for instance, a BMW ad can have a strong externality on a Mercedes one (as both makers attract the high end of the market), the externality on makers in a different price bracket, e.g. KIA, is arguably much less strong.

On the other hand of the user model scale, we find models motivated by these limitations. The authors of [2] propose a model whereby users have a *limited memory* and externalities occur only in a *window* of c consecutive slots. They consider externalities of an ad to apply to ads c slots below (*forward* externalities) and ads c slots above (*backward* externalities). They also consider the possibility that externalities boost CTRs (*positive* externalities). Moreover, they notably introduce the concept of *contextual graph* to model the fact that externalities might have *ad-dependent* effect. The vertices of a contextual graph are ads and edge weights represent the externality between the endpoints. Their model turned out to be too rich to allow significant algorithmic results. Indeed, their main results apply to the arguably less interesting case of forward positive externalities. The more interesting case of negative externalities appears to be too hard to tackle in their model.

*Partially supported by EPSRC grant EP/M018113/1.

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May, 4–8, 2015, Istanbul, Turkey.*

	c<K		c=K	
	nr	**r**	**nr**	**r**
LB	APX-hard	APX-comp	poly-APX-comp	APX-comp
UB	$\frac{log(N)}{2\min\{N,K\}}$			
SP	$\alpha\gamma_{min}^c$	1/2	$1/K$	1/2

Table 1: Summary of our results

2. NEW MODELS FOR EXTERNALITIES

In an instance of SSA, we have N advertisers, each having a single ad, and K slots to fill. For $i \in [N]$, we let a_i denote advertiser i's ad; $q_i \in [0,1]$ be its *quality* (i.e., the probability a user will click ad a_i when observed); and, $v_i \in \mathbb{R}^+$ be the *valuation* of advertiser i for a click on ad a_i. As in the cascade model, users behave in a Markovian fashion observing the slots from the top to the bottom one. W.r.t. the cascade model, the novelty in our model is constituted by parameters $\gamma_{i,j} \in [0,1]$, for $i, j \in [N]$, which model the *continuation probability* of the pair (a_i, a_j), i.e., the probability the user will move, in a Markovian fashion, from ad a_j to ad a_i that is allocated immediately after a_j (besides other forms of externalities, see below). (Note that in the original cascade model, $\gamma_{i,j} = \gamma_i$.) We encode the γ's in a *contextual graph* $G = ([N], \mathcal{E})$ where the direct edges represent the way an ad influences the others. In particular, the graph is weighted and the weight of the edge (i,j) is $\gamma_{i,j} \in (0,1]$; $(i,j) \notin \mathcal{E}$ means the probability the user will move from ad a_i to a_j is 0, i.e., $\gamma_{i,j} = 0$. For $m \in [K]$, we let s_m denote the m-th slot and $\lambda_m \in [0,1]$ denote the *factorized prominence* of s_m. We define the *prominence* $\Lambda_m = \prod_{i \le m} \lambda_i$. It is assumed w.l.o.g. that $\Lambda_1 = \lambda_1 = 1$. For notational convenience, we add an additional *fictitious ad* $a_\perp$, with $q_\perp = v_\perp = 0$, and an additional *fictitious slot* $s_\perp$, with $\Lambda_\perp = \lambda_\perp = 0$. Feasible allocations of ads to slots are those in which each ad a_i, $i \in [N]$, is allocated to either a unique slot s_m, $m \in [K]$, or $s_\perp$ (when not displayed). The fictitious ad $a_\perp$ can be assigned to more than one s_m, meaning that those slots are left empty. We also assume that each slot can be allocated to at most a single ad a_i, $i \in [N]$. We let Θ denote the set of feasible solutions and θ a specific allocation $\theta = (a_1, \dots, a_K)$, where the ads are ordered from the top slot to the bottom one. Given $\theta \in \Theta$, with abuse of notation, we denote as $\theta(m)$ the ad allocated in the m-th slot and with $\theta(a_j)$ the slot ad a_j is allocated to, i.e., $\theta(a_j) = \{m \in [K] \cup \{\perp\} | \theta(m) = a_j\}$. Given an allocation $\theta \in \Theta$, we can now also define the click through rate of an ad a_i, $CTR_i(\theta)$, as the probability the ad is clicked by the user taking externalities into consideration. We define the optimal allocation θ^* as the one maximizing the *social welfare*, namely: $\theta^* \in \arg\max_{\theta \in \Theta} SW(\theta)$, where $SW(\theta) = \sum_{i \in [N]} CTR_i(\theta) v_i$. An α-approximate solution θ satisfies $SW(\theta) \ge \alpha SW(\theta)$.

Web design issues. The definition of the γ's is enriched to handle the cases in which $a_\perp$ is allocated to some s_m, $m \le K$. We define two models. In the ***no–reset*** (nr) model, we have $\gamma_{i,\perp} = 0$ and $\gamma_{\perp,i} = 0$ $\forall i \in [N] \cup \{\perp\}$. This model captures the situation in which leaving a slot empty kills user's attention for the remainder of the slots. This implies that the web page should be designed so to have all the allocated ads together as a contiguous list. In the ***reset*** (r) model, instead, $\gamma_{i,\perp} = \gamma_{\perp,i} = 1$ $\forall i \in [N] \cup \{\perp\}$. This model captures the situation where slots can be distributed in the page in different positions (a.k.a., slates) and slots with no ads can be used to raise the user's attention, e.g., by allocating pictures that annul the externality between the ads allocated before and after the picture.

User Memory. We denote by c the number of ads displayed above a_i in θ, from $s_{\theta(a_i)-1}$ to $s_{\theta(a_i)-c}$, that affect $CTR_i(\theta)$. While in the cascade model $c = K$, here we assume that the memory of the user, also called *window*, might be limited and thus $c < K$ in general, and often $c = O(1)$.

CTRs. The probability that ad a_i is clicked is defined as $CTR_i(\theta) = \Lambda_{\theta(a_i)} \Gamma_i(\theta) q_i$, where

$$\Gamma_i(\theta) = \prod_{m=\max\{\theta(a_i)-c,1\}}^{\theta(a_i)-1} \gamma_{\theta(m),\theta(m+1)}$$

while $\Lambda_{\theta(a_i)}$ is defined as in the cascade model. When the window size equals the number of slots, i.e., $c = K$, the above reduces to $\Gamma_i(\theta) = \prod_{m=1}^{\theta(a_i)-1} \gamma_{\theta(m),\theta(m+1)}$. We denote as CFNE($c$)-r and CFNE($c$)-nr (Cascade Forward Negative Externalities) the problems of computing θ^* in the reset and no–reset model, respectively, when the window is $c \le K$.

Contributions. Table 1 summarizes our results (wherein the first row contains lower bounds, the second upper bounds, and the third strategyproof (SP) upper bounds). When the user memory is *bounded* ($c < K$), we prove that the problem of optimizing the social welfare is APX-hard for both the reset and no–reset models. Moreover, for CFNE(c)-nr we are not able to provide a constant approximation algorithm matching the APX-hardness lower bound. We provide instead a $\frac{log(N)}{2\min\{N,K\}}$ approximation algorithm based on the Color Coding technique. The algorithm proposed is MIR (Maximal In Range) so it can be used to design a SP mechanism. We also devise an approximate algorithm that turns any α-approximation for the Weighted 3-Set Problem into a $\alpha\gamma_{min}^c$ approximation for the subclass of instances of CFNE(c)-nr in which the minimum externality in the contextual graph is $\gamma_{min} > 0$, thus improving the upper bound for unconstrained instances wherein γ_{min} is constant.

As regards the models with unbounded user memory, we prove that CFNE(K)-nr is poly-APX-complete, although the polynomial approximation algorithm we use to prove the upper bound (i.e., membership in poly-APX) is not monotone, and hence cannot be used to design SP mechanisms. For CFNE(K)-nr, we identify a tractable subclass of instances having DAG contextual graphs, for which we provide an $O(KN^2)$ optimal (hence, MIR) algorithm, leading to a SP mechanism. For the unconstrained class of instances we are only able to provide a $1/K$-approximate MIR algorithm. Finally, we prove that CFNE(K)-r is APX-complete, providing a simple 1/2-approximate monotone algorithm.

REFERENCES

[1] G. Aggarwal, J. Feldman, S. Muthukrishnan, and M. Pál. Sponsored search auctions with markovian users. In *WINE*, 2008.

[2] D. Fotakis, P. Krysta, and O. Telelis. Externalities among advertisers in sponsored search. In *SAGT*, 2011.

[3] P. Jeziorski and I. Segal. What makes them click: Empirical analysis of consumer demand for search advertising. *AEJ: Microeconomics*, 2014.

[4] D. Kempe and M. Mahdian. A cascade model for externalities in sponsored search. In *WINE*, 2008.