

Efficient Delivery Policy to Minimize User Traffic Consumption in Guaranteed Advertising

Jia Zhang,^{1,2} Zheng Wang,³ Qian Li,^{1,2} Jialin Zhang,^{1,2}
Yanyan Lan,^{1,2} Qiang Li,^{1,2} Xiaoming Sun^{1,2}

¹CAS Key Lab of Network Data Science and Technology, Institute of Computing Technology, Chinese Academy of Sciences, China.

²University of Chinese Academy of Sciences, China.

¹{zhangjia, liqian, zhangjialin, lanyanyan, liqiang01, sunxiaoming}@ict.ac.cn

³Department of Computer Science, The University of Hong Kong, Hong Kong.

³zwang@cs.hku.hk

Abstract

In this work, we study the guaranteed delivery model which is widely used in online advertising. In the guaranteed delivery scenario, ad exposures (which are also called impressions in some works) to users are guaranteed by contracts signed in advance between advertisers and publishers. A crucial problem for the advertising platform is how to fully utilize the valuable user traffic to generate as much as possible revenue. Different from previous works which usually minimize the penalty of unsatisfied contracts and some other cost (e.g. representativeness), we propose the novel consumption minimization model, in which the primary objective is to minimize the user traffic consumed to satisfy all contracts. Under this model, we develop a near optimal method to deliver ads for users. The main advantage of our method lies in that it consumes nearly as least as possible user traffic to satisfy all contracts, therefore more contracts can be accepted to produce more revenue. It also enables the publishers to estimate how much user traffic is redundant or short so that they can sell or buy this part of traffic in bulk in the exchange market. Furthermore, it is robust with regard to priori knowledge of user type distribution. Finally, the simulation shows that our method outperforms the traditional state-of-the-art methods.

Introduction

Online advertising is now an industry which is worth tens of billions of dollars. According to the latest report released by Interactive Advertising Bureau (IAB 2016), the annual revenue of online advertising in the US reaches \$59.6 billion in 2015, which is \$10.1 billion higher than that in 2014. In online advertising, ad exposures (also called impressions) are sold in two major ways: either in spot through real-time bidding (or RTB) or in advanced by guarantee delivery (or GD) (Chen, Yuan, and Wang 2014). Although RTB is developing rapidly, only 13% of publishers' revenue is generated by RTB by the year 2012 (eMarketer 2012). Large amounts of ad exposures are still sold through GD.

GD is usually adopted by top tier advertisers (Apple, P&G, Coca-Cola etc.) and publishers (Yahoo!, Amazon

etc.). Generally, top tier publishers have frequently visited web pages and large amounts of first-hand user data. These resources are valuable, and undoubtedly, heavily attractive for advertisers, especially for those top tier advertisers. When promoting their products or brands, top tier advertisers usually prefer to sign a contract with publisher of high quality to ensure advertising effects. This contract contains targeted user types (what kinds of users the ads can only be displayed to) and the amount of exposures of corresponding ad to those targeted users. Publisher should comply with this contract and is paid when this contract is satisfied.

Previous works usually focus on meeting the requirements of all contracts. Specifically, they perform study on designing delivery policy to minimize some loss functions, which are related to unsatisfied contracts and other particular requirements such as representativeness and smoothness of delivery (Chen et al. 2012; Lefebvre, Contournet, and Crivelli 2012). In this work we propose a novel optimization model, in which the primary target is to minimize the user traffic (or page views of users) within the constraint of satisfying all exposure requirements. Compared to previous model, this model has more practical value:

- For most publishers, the *redundant user traffic*¹ will be sold in exchange market to avoid waste. However, compared with contract orders, selling user traffic in exchange market is usually less profitable. User traffic consumption for all contracts is minimized implies that more contracts can be accepted. Consequently, the proportion of user traffic sold to exchange market is decreased.
- Under the consumption minimization model, the amount of required user traffic can be estimated with proper delivery policy. This estimation can be viewed as an important indicator to predict how much user traffic is redundant or short. Then, publishers can sell or purchase user traffic in bulk, which is more profitable than doing that in scattered. This estimation is also useful for order management. When a particular contract comes, publisher can

¹The part of user traffic that is not targeted by any contract or corresponding contracts are already satisfied.

make decision on whether to accept it based on the change of user traffic consumption.

Contributions

In this work we propose a practical consumption minimization model for Guarantee Delivery. Under this model, we provide the theoretical analysis and design a near optimal delivery policy. Our contributions mainly consist of two parts which are listed as follows.

Theoretical analysis for optimal expected user traffic consumption. By employing some fundamental results in stopping theory, we prove that the optimal offline expected user traffic consumption is lower bounded by an expected flow (THEOREM 1). This is shown to be very useful for performance measurement of delivery policies. Besides, we construct an optimal online delivery policy based on a recursion. However, solving this recursion needs exponential time, which motivates us to design more efficient approximation delivery policy.

Efficient near optimal delivery policy. Inspired by some insights derived from the proof of lower bound for optimal offline user traffic consumption, we design a flow based delivery policy (ALGORITHM 1) in the online setting. This delivery policy has several advantages:

- (i) It consumes nearly as least as possible user traffic to satisfy all exposure requirements. Let n stand for the number of user types. If the total amount of all ads exposures is large enough, which is a realistic assumption in practice, the delivery policy achieves a near optimal competitive ratio $1 + \frac{1}{\text{poly}(n)}$ (THEOREM 2).
- (ii) It is robust on the priori knowledge of user type distribution. As this distribution is learned from history data, it usually has a bias compared to the real one. It is a concern that whether this bias will influence the delivery efficiency significantly. For this issue, we provide a robustness result indicating that when the bias between the learned user type distribution and the ground truth is bounded, the loss on delivery efficiency is also bounded (THEOREM 3).
- (iii) Our policy can estimate user traffic consumption as well (THEOREM 2). As aforementioned, this estimation is very useful for order and traffic management.
- (iv) Our delivery policy is also flexible. It can be easily implemented to reach some other requirements, such as representativeness and smoothness of delivery (REMARK 1.).

According to simulation experiments, our delivery policy works well under different settings and outperforms other delivery policies including the well known HWM policy (Chen et al. 2012).

Related Works

Online advertising is a hot topic concerned by researchers in both economic and computer science communities. We discuss the line of works that are most related to our paper.

The GD problem has been well studied as a special version of online assignment in theoretical domain. The classical online assignment problem studies adversary (or worst case) setting, such as (Karp, Vazirani, and Vazirani 1990;

Mehta et al. 2007; Agrawal, Wang, and Ye 2014; Sun, Zhang, and Zhang 2016). Some generalizations take distribution information into account. In the work (Feldman et al. 2009), the authors study the online matching problem when types of users are drawn independently from a known distribution, which is similar to our setting. However, their target is to maximize the matching size when there are n users in total. They propose an algorithm with competitive ratio 0.6702. This result is improved by works such as (Bahmani and Kapralov 2010; Jaillet and Lu 2013). Correspondingly, authors in the work (Karande, Mehta, and Tripathi 2011) study the online matching problem when the distribution of user types is unknown. In the work (Vee, Vassilvitskii, and Shanmugasundaram 2010), the authors propose the online allocation problem with forecast and suppose there is a sample set of vertices that will come online. They provide a general framework to optimize several object functions which satisfy some constraints. According to (Chen et al. 2012), this framework is sensitive to the accuracy of the forecast and requires solving a non-linear convex problem. We refer to the survey (Mehta 2012) for a systemic view.

Some more practical works have also been done in this field. In the work (Chen et al. 2012), the authors propose a delivery policy which is called High Water Mark (or HWM). It is efficient, space saving and stateless so that can run in parallel on distributed machines. However, it is not designed to optimize the user traffic consumption. Numerical results in the experimental part of our work show that HWM is not good enough when considering user traffic consumption. In the work (Lefebvre, Contournet, and Crivelli 2012), the authors design a policy called SHALE, which can be treated as an improvement of HWM. Actually, SHALE achieves a better performance on delivery with a loss on efficiency. However, SHALE is a heuristic algorithm, and usually needs specific number of iterations, which is sensitive to input, to get good enough results. A common problem for HWM and SHALE is that expected user traffic consumption is hard to estimated. Both HWM and SHALE take representativeness into consideration, which means ads in one campaign should be delivered somehow *fairly* to users among all targeted user types, instead of only partial types. Besides, *user reach* (the number of different individuals who see the ad) and *delivery frequency* (the number of times a user should see the ad) are also concerned by some works (Turner et al. 2014; Shen, Li, and Chen 2014).

In this work, we assume the user type distribution \mathcal{D} is known in advance and the types of users are sampled from this distribution, while in other works, such as (Cetintas, Chen, and Si 2013; Cetintas et al. 2011), authors study how to estimate user traffic in a particular period. Another line of works concern the pricing mechanism and revenue maximization in GD (Bharadwaj et al. 2010; Radovanovic and Heavlin 2012). In the work (Bharadwaj et al. 2010), two algorithms are provided to calculate the price of each guaranteed contract based on the price of each targeted user.

As RTB is adopted by more and more publishers, some works study how to combine RTB and GD to make more profit. In the work (Ghosh et al. 2009), publishers are considered to act as bidders for guaranteed contracts. In

works (Araman and Popescu 2010; Balseiro et al. 2011; Chen, Yuan, and Wang 2014), a similar framework is considered that publishers can dynamically select which ad campaign contract to accept. In the work (Chen, Yuan, and Wang 2014), it is shown that RTB and GD can be seamlessly combined programmatically and the authors provide some results about what percentages of user traffic should be sold in GD and what the price is.

Consumption Minimization Model

In guaranteed delivery (GD) scenario, users are categorized into a set of *atomic* types according to their traits, e.g., geographic, gender, age etc. *Atomic* means each user visiting the publish platform can only belong to a unique type. Each advertiser has one or more ad campaigns and each campaign corresponds to a contract with publisher. The contract specifies how many times the corresponding ad should be exposed and the set of targeted user types. See Figure 1 as an example. The first ad campaign reserves 200 thousand exposures and all users at the age of 20 to 25 are targets of this campaign. Thus when a user who satisfies this requirement visits the platform, corresponding ad can be displayed to this user.




Figure 1: An example of ads campaigns in GD.

The ad campaigns and user types relationship can be formally modeled as a demand-supply graph $G = (A, U, E)$ with $|A| = m$ and $|U| = n$. The nodes in A stand for m ad campaigns, while nodes in U stand for n user types. For each ad campaign $a_i \in A$ and user type $u_j \in U$, the edge $(a_i, u_j) \in E$ means that u_j is targeted by a_i . Conventionally, for any $v \in A \cup U$, we use $\Gamma(v)$ to stand for the neighbors of v in G . W_i stands for total amount of exposures that a_i reserves in the contract. We denote $W = \{W_1, W_2, \dots, W_m\}$ for convenience, and let $M \triangleq \sum_{i=1}^m W_i$. We can safely assume $M \gg n$.

We assume the users come one by one online. Once a user visits the publish platform (we call it *a user visit*), her type is *revealed* and the publisher should irrevocably deliver only one ad² (we call it *an exposure*). Usually, there are large amounts of users' visiting logs available on the publish platform, thus, we assume that each user type is drawn from a distribution $\mathcal{D} = \{p_j \mid 1 \leq j \leq n\}$ which can be learned from those log data. That is, for any unseen user, she belongs to type u_j with probability p_j , which satisfies $\sum_{j=1}^n p_j = 1$. We define p_{\min} (resp. p_{\max}) as the minimum (resp. maximum) arrival probability. \mathcal{D} is considered as a part of the input. Thus, the input instance I is a triple (G, W, \mathcal{D}) .

²Our results can be easily extended to meet multiple delivery requirements. See the full version (Zhang et al. 2016) for detail.

Our objective is to design an ads delivery policy to minimize the expectation of user traffic. Here the *user traffic consumed* means the number of users' visits when all contracts are fulfilled. It is a random variable due to the uncertainty of user traffic and the delivery policy. This is the *consumption minimization problem for user traffic* (abbreviated as CMPUT). For a delivery policy π , we employ the concept of competitive ratio to measure its performance. We say π has a *competitive ratio* c if and only if for any input instance, $\frac{E(\pi)}{E(OPT)} \leq c$, where $E(\pi)$ is the expected user traffic consumed by policy π and $E(OPT)$ is the expectation of optimal user traffic consumption when the publisher is allowed to modify previous delivery decisions. We also call $E(OPT)$ optimal offline expected user traffic consumption.

We provide several important definitions for further analysis. Given an input instance $I = (G, W, \mathcal{D})$ and a non-negative number T , we can construct an *expected network* $F_G(T)$ by adding two nodes s (source) and t (sink) into G and set the capacities as follows. For $a_i \in A$, we connect s with a_i and set the capacity of the edge (s, a_i) as W_i . For $u_j \in U$, we connect t with u_j and set the capacity of the edge (u_j, t) as Tp_j . For those edges in E , we just set their capacities as M . Let $\text{MAX-FLOW}(F_G(T))$ stand for the maximum flow that can pass through $F_G(T)$. Z_{flow} is defined as the minimum value of T that satisfies $\text{MAX-FLOW}(F_G(T)) = M$.

Theoretical Analysis

In this section we first prove a lower bound of the optimal offline expected user traffic consumption, and then construct an optimal online policy based on a recursion which needs exponential time to solve. To beat this hardness we develop an efficient and near optimal flow based policy.

Lower Bound for Optimal Offline User Traffic Consumption

The main tool we employ to develop our lower bound is *Wald's identity* (Wald 1944). Wald's identity is a foot stone in optimal stopping theory. Here we use a variant introduced in work (Blackwell 1946), and it can be described as follow.

Wald's Identity. Let X_1, X_2, \dots be i.i.d. random variables with $E(X_1) < \infty$, and let τ be a stopping rule which is independent with $X_{\tau+1}, X_{\tau+2}, \dots$ and satisfies $E(\tau) < \infty$, then there has $E(\sum_{i=1}^{\tau} X_i) = E(\tau)E(X_1)$.

To use Wald's Identity, we need an upper bound of $E(OPT)$. Consider a simple delivery policy called *RANDOM*, in which each ad campaign a_i is expanded as W_i ads. These ads target the same types of users and each of them needs to be exposed only once. When a user that belongs to u_j comes, *RANDOM* uniformly chooses one ad from all valid ones waiting for exposure. For each u_j , denote $W(u_j) = \sum_{a_i \in \Gamma(u_j)} W_i$, we have the following result.

Proposition 1.

$$E(\text{RANDOM}) \leq \int_0^\infty (1 - \prod_{i=1}^m (1 - e^{-t \sum_{u_j \in \Gamma(a_i)} \frac{p_j}{W(u_j)} W_i}) dt.$$

Proof. In *RANDOM*, each expanded ad will be delivered with probability $\sum_{u_j \in \Gamma(a_i)} \frac{p_j}{W(u_j)}$ when the first

user comes. This probability will increase with more and more ads are delivered. If we fix the probability just as $\sum_{u_j \in \Gamma(a_i)} \frac{p_j}{W(u_j)}$, the GD problem is reduced to the *non-uniform coupon collection problem*, and according to the work (Flajolet, Gardy, and Thimonier 1992), the expected user traffic consumption can be calculated by an integral as shown in this proposition. Clearly, this integral will converge to a value, which is an upper bound of $E(\text{RANDOM})$. \square

According to above proposition, $E(\text{OPT}) \leq E(\text{RANDOM}) < \infty$.

Denote U_j as the number of users belong to type u_j when all contracts are fulfilled and let $[n]$ stand for the set $\{1, 2, \dots, n\}$. We can use Wald's Identity to show the following proposition.

Proposition 2. $E(U_j) = p_j E(\text{OPT})$ for any $j \in [n]$.

Proof. This is a direct application of Wald's Identity. Fix any $j \in [n]$, define

$$X_t = \begin{cases} 1, & \text{the } t\text{-th user belongs to type } u_j, \\ 0, & \text{otherwise.} \end{cases}$$

We can find that $\{X_t\}_{t \geq 1}$ is a sequence of i.i.d. random variables and $E(X_t) = p_j < \infty$. In addition, the stopping time of the policy, i.e. OPT is independent of $X_{\text{OPT}+1}, X_{\text{OPT}+2}, \dots$ and $E(\text{OPT}) < \infty$. Thus $E(U_j) = E\left(\sum_{t=1}^{\text{OPT}} X_t\right) = p_j E(\text{OPT})$. \square

Theorem 1. For any input instance, $E(\text{OPT}) \geq Z_{\text{flow}}$.

Proof. According to the definition of Z_{flow} and the property of maximum flow, there exists a subset $S \subseteq A$ such that $\sum_{a_i \in S} W_i \geq \sum_{u_j \in \Gamma(S)} p_j Z_{\text{flow}}$, where $\Gamma(S) \subseteq U$ is the set of user types targeted by ad campaigns in S . Otherwise, Z_{flow} can be decreased to a smaller value which still satisfies $\text{MAX-FLOW}(F_G(Z_{\text{flow}})) = M$. It is a contradiction to the definition of Z_{flow} .

On the other hand, the delivery procedure can stop only if the ad corresponding to campaign a_i has been delivered W_i times, thus $\sum_{u_j \in \Gamma(S)} U_j$ is always no less than $\sum_{a_i \in S} W_i$.

According to PROPOSITION 2, we have $E(U_j) = p_j E(\text{OPT})$ for any $i \in [n]$, thus

$$\sum_{u_j \in \Gamma(S)} p_j E(\text{OPT}) = E \sum_{u_j \in \Gamma(S)} U_j \geq \sum_{u_j \in \Gamma(S)} p_j Z_{\text{flow}},$$

which implies $E(\text{OPT}) \geq Z_{\text{flow}}$. \square

Optimal Online Delivery Policy

Suppose the first user comes with type u_j . If no ad campaign targets u_j , nothing will be delivered to this user; otherwise, the delivery policy should decide which ad in $\Gamma(u_j)$ will be delivered. It is easy to argue that if $\Gamma(u_j)$ is not empty, delivering nothing can not be better than delivering any arbitrary ad to the user. Suppose $a_i \in \Gamma(u_j)$ is delivered, then the problem is reduced to a sub problem in which ad a_i should be delivered $W_i - 1$ times. For convenience, we construct an input $I(i) = (G, W(i), \mathcal{D})$, where $W(i) =$

$\{W_1, \dots, W_{i-1}, \max\{0, W_i - 1\}, W_{i+1}, \dots, W_m\}$. Based on above analysis, in the optimal online delivery policy π^* , a_i must be chosen to satisfy that $a_i = \arg \min_{a_k \in \Gamma(u_j)} E(\pi^*(I(k)))$, where $E(\pi^*(I(k)))$ stands for expected user traffic consumed by π^* on input instance $I(k)$. According to the properties of conditional expectation, we have $E(\pi^*(I)) = 1 + \sum_{j=1}^n p_j \min_{a_i \in \Gamma(u_j)} E(\pi^*(I(i)))$. Thus, π^* can be constructed from bottom to top by dynamic programming (DP). However, it is easy to check that there are $\prod_{i=1}^m (W_i + 1)$ sub problems in this DP. Thus, exponential time is needed to construct π^* in this way. In fact, we believed that π^* is hard to construct in any way. This motivates us to design some delivery policies with polynomial running time.

Delivery Policy Design for CMPUT

According to previous analysis, Z_{flow} is highly related to $E(\text{OPT})$. It somehow reminds us to use expected flow to guide ads delivery. Inspired by this insight, we design a flow based delivery policy.

Flow Based Delivery Policy As shown in ALGORITHM 1, the flow based delivery policy can be divided into two parts: the offline part and the online part. In the offline part, we first construct an expected network with integer capacities on all edges, that's $\hat{F}_G(\hat{Z})$ used in ALGORITHM 1. The difference between $F_G(\hat{Z})$ and $\hat{F}_G(\hat{Z})$ lies in that the capacity of edge (u_j, t) is set to be $\lceil \hat{Z} p_j \rceil$ for $j \in [n]$ in the latter. Then we find the minimum value of \hat{Z} while ensuring the max flow on $\hat{F}_G(\hat{Z})$ is no less than M . To find such a value of \hat{Z} , we extend the idea of binary search, that's to say, we enumerate a value \hat{Z} , and calculate $\text{MAX-FLOW}(\hat{F}_G(\hat{Z}))$, then adjust the value of \hat{Z} with a similar idea of binary search. According to the properties of maximum flow, the traffic on each edge in maximum flow is a non-negative integer when the capacities of all edges are integers. This is why we use $\hat{F}_G(\hat{Z})$ instead of $F_G(\hat{Z})$. In the online part, when a user comes, we deliver a proper ad according to the maximum flow of $\hat{F}_G(\hat{Z})$.

Note that we haven't specified how to deliver an ad according to the maximum flow of $\hat{F}_G(\hat{Z})$ in the line REM. Actually, FLOW-BASED-DELIVERY-RULE stands for a class of delivery rules which satisfy that when there are no less than $\lceil \hat{Z} p_j \rceil$ users of type u_j come to the publish platform, either all ads of a_i are delivered or no less than $C_{i,j}$ ads of a_i are delivered to users of type u_j . We can implement a FLOW-BASED-DELIVERY-RULE according to particular requirements, such as representativeness and smoothness of delivery. Here we give a simple implementation in ALGORITHM 2.

Next, we will prove that the performance of ALGORITHM 1 is quite good by showing that it is close to the lower bound provided in THEOREM 1 when $M p_{\min} = \Omega(n^{-\varepsilon})$ for $\varepsilon > 0$. The main tool we employ is the Chernoff Bound, which is an important tool in probability theory. Besides the Chernoff Bound, the full proof of THEOREM 2 mainly depends on delicate analysis of expectation. We leave this proof in supplemental material of the full version (Zhang et al. 2016).

Algorithm 1: Expected Flow Based Delivery Policy

INPUT: $I = (G, W, \mathcal{D})$.
OFFLINE PART:
Set $\hat{Z} \leftarrow M/2$.
repeat
 Set $\hat{Z} \leftarrow 2\hat{Z}$.
 Construct an expected network $\hat{F}_G(\hat{Z})$.
until $\text{MAX-FLOW}(\hat{F}_G(\hat{Z})) \geq M$;
Set $Begin \leftarrow \hat{Z}/2, End \leftarrow \hat{Z}$.
repeat
 Set $Mid \leftarrow (Begin + End)/2$.
 if $\text{MAX-FLOW}(\hat{F}_G(\hat{N})) < M$ **then**
 Set $Begin \leftarrow Mid + 1$.
 else
 Set $End \leftarrow Mid$.
until $Begin \geq End$;
Set $\hat{Z} \leftarrow Begin$.
for each $(a_i, u_j) \in E$ **do**
 Set $C_{i,j}$ as the traffic passing through (a_i, u_j) in
 the maximum flow of $\hat{F}_G(\hat{Z})$.
ONLINE PART:
for each user comes to the publish platform do
 if $M = 0$ **then** **exit**.
 Suppose u_j is the type of this user.
 Set $i \leftarrow \text{FLOW-BASED-DELIVERY-RULE}(u_j)$.
 if $i = NULL$ **then**
 Taking this user for other purpose.
 Deliver an ad of a_i to current user.
 Set $W_i \leftarrow W_i - 1$ and $M \leftarrow M - 1$.

Algorithm 2: GREEDY-DELIVERY-RULE

GREEDY-DELIVERY-RULE(u_j)
Set $i' \leftarrow \arg \max_{a_i \in \Gamma(u_j), W_j > 0} \{C_{i,j}\}$.
if we can't find such an i' **then**
 return $NULL$.
Set $C_{i',j} \leftarrow C_{i',j} - 1$.
return i' .

Theorem 2. Given $\varepsilon > 0$, when $Mp_{\min} = \Omega(n^\varepsilon)$, the expected user traffic consumption of ALGORITHM 1 is at most $(1 + O(n^{-\varepsilon/2})) \cdot Z_{flow}$, which implies that it achieves a competitive ratio $1 + O(n^{-\varepsilon/2})$.

Remark 1. It should be pointed out that the flow based delivery policy can be generalized to solve the representativeness, smoothness and multiple delivery issues. The details can be found in the supplemental material of (Zhang et al. 2016).

Remark 2. The flow based delivery policy provides a threshold that how many users of each type are needed to satisfy all contracts, thus user traffic beyond this threshold can be safely utilized for other purpose, such as sell it in an exchange market. It is shown in Line IF of ALGORITHM 1.

Robustness In our model, we assume the user type distribution is accurate. However, there usually has a bias between estimated distribution and the real one. The following analysis provides a robustness guarantee for flow based delivery policy.

Given a delivery policy π , denote $E_{\mathcal{D}}(\pi(\hat{\mathcal{D}}))$ as the expected user traffic consumption when the distribution input to π is $\hat{\mathcal{D}}$ while the real distribution is \mathcal{D} . We also define \hat{p}_j as the estimated probability for p_j in $\hat{\mathcal{D}}$. Now we have the following theorem for the flow based delivery policy. The full proof is moved to supplemental material of (Zhang et al. 2016).

Theorem 3. Denote the flow based delivery policy as π_f , for any $\delta > 0$, if $(1 - \delta)p_j \leq \hat{p}_j \leq (1 + \delta)p_j$ for each user type a_j , we have $E_{\mathcal{D}}(\pi_f(\hat{\mathcal{D}})) \leq \frac{(1+O(n^{-\varepsilon/2}))(1+\delta)}{1-\delta} E_{\mathcal{D}}(\pi_f(\mathcal{D}))$.

Experiments

In this section, we perform some simulations of the flow based delivery policy and several other policies on different random graphs. All experiment arguments are selected to approximate the real world system. The results show the superiority of the flow based delivery policy.

Experimental Setup

Hardware A desktop with an Intel i5-4570 CPU @ 3.2GHZ and a 4GB DDR3 memory.

Settings Restricted by the business model of GD, the size of demand-supply graph in real system is usually not large. Thus we conduct experiments on a random demand and supply graph with 500 ad campaigns and 1000 user types. Edges are generated uniformly and randomly while the average degree of user types are fixed. To generate the user types distribution \mathcal{D} , we choose a random value $r_j > 0$ for each u_j , and set $p_j = \frac{r_j}{\sum_{l=1}^{1000} r_l}$. r_j is chosen in two ways: r_j is drawn uniformly from $[0, 1]$ (we call this method Random-Normalization) or r_j equals to $\frac{1}{1000} + \varepsilon$, where $\varepsilon \sim \mathcal{N}(0, 1/6000^2)$ (we call this method Gauss-Perturbation).

Delivery Policies We compare the flow based delivery policy (FB) with other four delivery policies, including RANDOM, HWM, and two intuitive greedy policies: PROBABILITY-GREEDY (PG) and DEGREE-GREEDY (DG).

Policies DG and PG are based on a straightforward idea that ad campaigns that are *difficult* to be fulfilled should have a higher priority. Specifically, when a user comes, the ad with minimum degree on demand-supply graph is chosen in DG, while in PG, it calculates a *delivery value* $r_i = \sum_{u_j \in \Gamma(a_i)} \frac{p_j}{W(u_j)}$ for each ad campaign a_i , and delivers an ad with minimum delivery value. Recall that $W(u_j) = \sum_{a_i \in \Gamma(u_j)} W_i$. For these policies, we compare their competitive ratios and the robustness on different random graphs and different user visiting sequences. On each generated random graph, we simultaneously run five policies with 100 sampled user sequences. The competitive ratio for each policy is estimated as dividing the average user traffic consumed on these sampled sequences by the average optimal user traffic needed.

Experiments under Different Average Degrees

The first experiment is conducted under different average degrees of user types. With a specific average degree, there

are 100 input instances generated independently and randomly. Half of them use Random-Normalization to generate \mathcal{D} , the others use Normalization-Perturbation. W_i is drawn uniformly from [50, 100]. In the second experiment we will show that the amount of exposures has little influence on performance, thus a small value for W_i is enough.

Experiment results for average degrees 5, 10, 15, 20 and 25 are shown in Figure 2. The height of the bar stands for the average competitive ratios on 50 input instances with corresponding setting while the vertical lines indicate the fluctuation ranges of those competitive ratios. According to this figure, the flow based delivery policy has a much better competitive ratio under all average degrees. These fluctuation ranges of competitive ratios also show that the flow based delivery policy is more robust on different graphs.

Surprisingly, DG is the second best policy. When the average degree is 5, its average competitive ratio is very close to that of the flow based delivery policy with a gap about 0.015 for Random-Normalization and 0.035 for Gauss-Perturbation. However, these gaps will be enlarged as the average degree increases to 10, 15 and 20. From Figure 2 we can also find the fact that the performances of all delivery policies will be improved continuously with average degree increasing. In fact, when the average degree is large enough, all polices are near optimal³. However, for the purpose of accurately targeting, the demand-supply graph is usually very sparse in real world.

In Figure 3, we show more detail experiment results of five delivery polices on 50 input instances with Random-Normalization and average degree 10. Figure 3a shows the results of competitive ratios. The horizontal axis stands for 50 input instances. As we can see, the flow based delivery policy achieves a much better competitive ratio on each input instance and it works more steadily. The worst case performance among all sampled user visiting sequences on 50 input instances is shown in Figure 3b. Clearly, the flow based delivery policy is more robust on the worst sampled sequence too.

Experiments under Different Exposure

In this part, we compare five polices under different exposure settings. The number of exposures for each ad campaign is uniformly drawn from 100 to 2500, 5000, 10000 and 20000. We only consider average degrees of 5 and 10. Similarly, 100 graphs are drawn independently, and for 50 of them, we construct distributions by Random-Normalization and for the rest, Gauss-Perturbation is used. Then input instances are constructed with different exposure settings.

Figure 4 shows the results for Random-Normalization. The flow based delivery policy outperforms other polices under all exposure settings. We point out that the amount of exposures seems to have little influence on the performance for each delivery policy. It is somehow means the user traffic consumption of each studied policy has a linear relationship with the total amount of exposures for a fixed demand-supply graph. Besides, comparing Figure 4a

³Considering an extreme case that when the demand-supply graph is complete, any reasonable delivery policy will be optimal.




Figure 2: Comparison under different average degree with Random-Normalization and Gauss-Perturbation. The height of the bar is the average of competitive ratios on 50 input instances and the vertical lines indicate the fluctuation ranges.

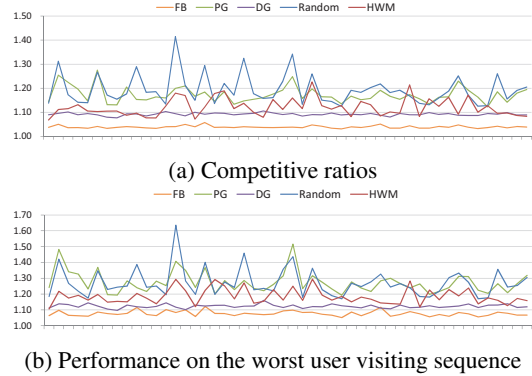


Figure 3: Results for 50 input instances with average degree 10 and Random-Normalization. The horizontal axis stands for those input instances.




Figure 4: Comparison under different exposures with Random-Normalization.

and 4b and combining previous experiments, we can find that the structure of demand-supply is much more crucial for the GD problem. Similar results are obtained for Gauss-Perturbation. These results can be found in supplemental material of (Zhang et al. 2016).

Conclusions and Future Works

In this work, we study the consumption minimization problem for user traffic, which lies in the heart of GD scenario. We provide a lower bound for the optimal offline user traffic consumption with an expected flow, which is shown to be useful for performance evaluation. We then propose a flow based delivery policy. This delivery policy consumes nearly as least as possible user traffic to satisfy all contracts under a realistic constraint. What's more important, it is robust with regard to estimation error of user type distribution and

shown to be very useful for order and traffic management.

There are some works remained to be done. For example, there may be a time window for each advertiser indicating the ad campaign's starting time and deadline. Another interesting problem is how to design delivery policy when advertisers also come online instead of being revealed in advance.

Acknowledgments.

This work was funded by 973 Program of China under Grants No. 2014CB340401, 2016YFB1000201, the National Natural Science Foundation of China (NSFC) under Grants 61222202, 61433014, 61502449, 61602440, the China National Program for support of Top-notch Young Professionals, and the Youth Innovation Promotion Association CAS under Grants No. 2016102.

References

- Agrawal, S.; Wang, Z.; and Ye, Y. 2014. A dynamic near-optimal algorithm for online linear programming. *Operations Research* 62(4):page. 876–890.
- Araman, V. F., and Popescu, I. 2010. Media revenue management with audience uncertainty: Balancing upfront and spot market sales. *Manufacturing & Service Operations Management* 12(2):190–212.
- Bahmani, B., and Kapralov, M. 2010. Improved bounds for on-line stochastic matching. In *18th Annual European Symposium Algorithms*, 170–181.
- Balseiro, S.; Feldman, J.; Mirrokni, V.; and Muthukrishnan, S. 2011. Yield optimization of display advertising with ad exchange. In *Proceedings of the 12th ACM conference on Electronic commerce*, 27–28.
- Bharadwaj, V.; Ma, W.; Schwarz, M.; Shanmugasundaram, J.; Vee, E.; Xie, J.; and Yang, J. 2010. Pricing guaranteed contracts in online display advertising. In *Proceedings of the 19th ACM international conference on Information and knowledge management*, 399–408.
- Blackwell, D. 1946. On an equation of wald. *Annals of Mathematical Statistics* 17(1):84–87.
- Cetintas, S.; Chen, D.; Si, L.; Shen, B.; and Datbayev, Z. 2011. Forecasting counts of user visits for online display advertising with probabilistic latent class models. In *Proceedings of the 34th International ACM Conference on Research and Development in Information Retrieval*, 1217–1218.
- Cetintas, S.; Chen, D.; and Si, L. 2013. Forecasting user visits for online display advertising. *Information Retrieval* 16(3):369–390.
- Chen, P.; Ma, W.; Mandalapu, S.; Nagarjan, C.; Shanmugasundaram, J.; Vassilvitskii, S.; Vee, E.; Yu, M.; and Zien, J. 2012. Ad serving using a compact allocation plan. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, 319–336.
- Chen, B.; Yuan, S.; and Wang, J. 2014. A dynamic pricing model for unifying programmatic guarantee and real-time bidding in display advertising. In *Proceedings of the 8th International Workshop on Data Mining for Online Advertising*, 1:1–1:9.
- eMarketer. 2012. Real-time bidding poised to make up quarter of all display spending. <http://goo.gl/04Ykx>. Accessed: 2016-05-02.
- Feldman, J.; Mehta, A.; Mirrokni, V.; and Muthukrishnan, S. 2009. Online stochastic matching: Beating $1-1/e$. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science*, 117–126.
- Flajolet, P.; Gardy, D.; and Thimonier, L. 1992. Birthday paradox, coupon collectors, caching algorithms and self-organizing search. *Discrete Applied Mathematics* 39(3):207–229.
- Ghosh, A.; McAfee, P.; Papineni, K.; and Vassilvitskii, S. 2009. Bidding for representative allocations for display advertising. In *Internet and Network Economics*, volume 5929. 208–219.
- IAB. 2016. 2015 internet advertising revenue full-year report.
- Jaillet, P., and Lu, X. 2013. Online stochastic matching: New algorithms with better bounds. *Mathematics of Operations Research*.
- Karande, C.; Mehta, A.; and Tripathi, P. 2011. Online bipartite matching with unknown distributions. In *Proceedings of the 43rd annual ACM symposium on Theory of computing*, 587–596.
- Karp, R. M.; Vazirani, U. V.; and Vazirani, V. V. 1990. An optimal algorithm for on-line bipartite matching. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing*, 352–358.
- Lefebvre, F.; Contournet, P.; and Crivelli, A. J. 2012. Shale: An efficient algorithm for allocation of guaranteed display advertising. In *the 18th ACM international conference on Knowledge discovery and data mining*, 1195–1203.
- Mehta, A.; Saberi, A.; Vazirani, U.; and Vazirani, V. 2007. Adwords and generalized online matching. *Journal of the ACM* 54(5):264–273.
- Mehta, A. 2012. Online matching and ad allocation. *Foundations & Trends in Theoretical Computer Science* 8(4).
- Radovanovic, A., and Heavlin, W. D. 2012. Risk-aware revenue maximization in display advertising. In *the 21st international conference on World Wide Web*, 91–100.
- Shen, H.; Li, Y.; and Chen, Y. 2014. Robust ad delivery plan for guaranteed display advertising. In *IEEE International Conference on Information and Automation*, 1125–1130.
- Sun, X.; Zhang, J.; and Zhang, J. 2016. Near optimal algorithms for online weighted bipartite matching in adversary model. *Journal of Combinatorial Optimization*, accepted.
- Turner, J.; Hojjat, A.; Cetintas, S.; and Yang, J. 2014. Delivering guaranteed display ads under reach and frequency requirements. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*.
- Vee, E.; Vassilvitskii, S.; and Shanmugasundaram, J. 2010. Optimal online assignment with forecasts. In *Proceedings of the 11th ACM conference on Electronic commerce*, 109–118.
- Wald, A. 1944. On cumulative sums of random variables. *The Annals of Mathematical Statistics* 15(3):283–296.
- Zhang, J.; Wang, Z.; Li, Q.; Zhang, J.; Lan, Y.; Li, Q.; and Sun, X. 2016. Efficient Delivery Policy to Minimize User Traffic Consumption in Guaranteed Advertising. *ArXiv e-prints: 1611.07599*.