

Analysis of a Learning Based Algorithm for Budget Pacing

Extended Abstract

MohammadTaghi Hajiaghayi
University of Maryland
College Park, USA
hajiagha@umd.edu

Max Springer
University of Maryland
College Park, USA
mss423@umd.edu

ABSTRACT

We analyze a natural learning algorithm for uniform pacing of advertising budgets, equipped to adapt to varying ad sale platform conditions. On the demand side, advertisers face a fundamental technical challenge in automating bidding in a way that spreads their allotted budget across a given campaign subject to hidden, and potentially dynamic, cost functions. This automation and calculation must be done in runtime, implying a necessarily low computational cost for the high frequency auction rate. Advertisers are additionally expected to exhaust nearly all of their sub-interval (by the hour or minute) budgets to maintain budgeting quotas in the long run. To resolve this challenge, our study analyzes a simple learning algorithm that adapts to the latent cost function of the market and learns the optimal *average* bidding value for a period of auctions in a small fraction of the total campaign time, allowing for smooth budget pacing in real-time. We prove our algorithm is robust to changes in the auction mechanism, and exhibits a fast convergence to a stable average bidding strategy.

KEYWORDS

budget pacing; auction systems; mechanism design; learning algorithm

ACM Reference Format:

MohammadTaghi Hajiaghayi and Max Springer. 2023. Analysis of a Learning Based Algorithm for Budget Pacing: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 3 pages.

1 INTRODUCTION

Ad impressions sold through real-time bidding (RTB) auctions are responsible for an ever-increasing portion of company expenditures as well as the revenue of large ad providers like Google, Facebook, Amazon, Microsoft, and Yahoo!. Google alone is responsible for upwards of 50 billion ad impressions on average per day [21], with a corresponding revenue on the order of \$100 million. While companies bidding in these advertising campaigns do not participate in every auction throughout the day, they are often required to participate enough to spend their allotted budget in this time. Thus, great interest is placed on adequately pacing budgets throughout the day so as to not spend too hastily at the beginning of a day and miss out on better impressions, or spend too frugally until the close of day.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

We here take the perspective of a demand-side platform (DSP), an intermediary serving as the interface connecting advertisers with ad-exchanges and offering customized bid strategies in online auctions. Our model of a DSP is given a users budget and target spent amount, and is asked to uniformly pace spending so as to exhaust this amount.

1.1 Problem Statement

We consider *online bid optimization* in the following framework: there are T auction periods ordered by an index $t \in \{1, \dots, T\}$. At the start of an auction period, an individual advertiser must make a decision as to how much to bid for desired ads within this period, denoted by b_t . Advertisers further have a total daily budget $B \in \mathbb{R}_{\geq 0}$ that limits the amount that can be spent within the day. Typically, advertisers would like to have smooth budget delivery [1, 3–5, 9, 15, 23], expressed as not buying more than a set fraction of the impressions before a given time in order to ensure that (1) budgets are not prematurely expended, thus resulting in missed opportunities later in the day, and (2) spending does not fluctuate substantially for ease of analysis.

The added complication of the smooth delivery problem we focus on in the present study is that oftentimes, when an advertiser submits a bid in an auction, this may not be the true value they pay for that impression. For example, by nature of the auction system implemented, a user may pay more or less than their desired bid a result of reserve prices or unexpected pricing fluctuations.

In response, the pacing of an advertiser’s budget relies on learning the actual cost of submitted bids to adequately scale bid values to meet desired spending goals in each period. We formally frame the budget pacing problems as:

$$\text{minimize } B - \sum_{t=1}^T c_t \quad (1)$$

$$\text{subject to } \left| \frac{B}{T} - c_t \right| \leq \epsilon \quad (2)$$

where c_t is the cost incurred in period t and for small ϵ . We note that here the optimization is framed using the information for all the auction periods, however, the problem itself is online. As a result, it is clear we need an algorithm that quickly learns the latent cost function, $f(b_t) = c_t$, in a small portion of the total number of auctions and subsequently uniformly paces their bidding for the duration of the campaign.

Lastly, it is important to note that the cost function, f , is *subject to change* at different times.

Algorithm 1 Budget Smoothing

```

1: Input:  $B, T, t, b_t, (c_0, \dots, c_t)$ 
2: Output:  $b_{t+1}$ 
3:  $B_r^t = B - \sum_{i=0}^t c_i$  {remaining budget after time  $t$ }
4:  $c_{\text{opt}} = \frac{B_r^t}{T-t}$  {optimal spend amount}
5:  $c_{\text{act}} = c_t$  {actual spend amount}
6:  $\alpha = \frac{c_{\text{opt}}}{c_{\text{act}}}$ 
7: return  $\alpha \cdot b_t$ 

```

2 ALGORITHM

Our approach to smoothly pacing an advertisers’s budget over a fixed length advertising campaign invokes an iterative control feedback mechanism to estimate the proper *average* bid to submit over each time period, which is later manipulated by a platform (potentially hidden to the bidder) to compute an “actual spent amount” for that period so that the total budget is approximately spent in a uniform fashion throughout the campaign. The algorithm relies on simplistic scaling of bids in response to learning of this latent mechanism in a naturalistic way, and is currently in implementation at Amazon and Overstock, two major companies with which the authors were previously affiliated.

The algorithm is formulated on the assumption the the advertiser has two pieces of information: its budget, B , and the number of auction periods in the campaign, T (ie. auction periods per day). Intuitively, the advertiser who is trying to uniformly pace their budget will initially bid its average budget for the corresponding number of auctions, $b_0 = \frac{B}{nT}$ (where n is the number of auctions per period). However, in general, the actual cost is much larger than the input bid and the convergence time remains low regardless of this initial bid selection, so any choice suffices. Once an initial average bid is set, it is utilized for the first period ($t = 0$), and the agent incurs cost c_0 throughout this time.

Following the initial bid, an advertiser now has the information of how much was spent in the first period and can assess the discrepancy between the *desired* amount to be spent and the *actual* amount. Let B_r^t denote the remaining budget after period t , and c_t be the incurred cost in this period, then we can define the scaling factor $\frac{B_r^t}{T-t} \cdot \frac{1}{c_t}$ as the ratio of the amount the advertiser wants to spend on average in each remaining auction period to how much it spent on the previous. Using this ratio as a budget pacing factor, we have the following iterative scheme:

$$b_{t+1} = \frac{B_r^t}{c_t(T-t)} \cdot b_t$$

Our main theoretical result bounds the number of periods until convergence to a stable bidding strategy and, thus, uniform budget pacing. We henceforth assume that the latent cost function takes the form $f_t(b_t) = C \cdot b_t^k$ for nonnegative parameters C and k . Most crucially, this convergence occurs within a small fraction of the total auction time and is dependent upon these problem parameters.

THEOREM 2.1. For $|1 - k| < 1$, Algorithm 1 has a bounded distance from the stable bid value at time t defined by:

$$\epsilon := |b_t - b^*| \leq \gamma^{-1/k} \cdot \frac{|1 - k|^{t-1 + \frac{1}{k}}}{1 - |1 - k|}$$

where $\gamma = \frac{CT}{B}$. Subsequently, we have a convergence time, t^* , to the stable bid value bounded as:

$$t^* \leq \frac{k-1}{k} + \frac{\ln \left| \epsilon \gamma^{1/k} (1 - |1 - k|) \right|}{\ln |1 - k|} \quad (3.4.1)$$

We can see that the convergence time is dependent upon the parameters of our RTB campaign, namely the budget, length, and degree of the latent cost function.

For the special case of *linear* spent amount functions ($k = 1$), we exhibit convergence in exactly one iteration. This simplistic cost function corresponds to a mere rescaling of input bids and naturally captures a large class of online auction systems.

THEOREM 2.2. For a linear cost function, $f(b_t) = C \cdot b_t$ where $C > 0$, and initial bid b_0 , Algorithm 1 converges to a fixed point bid value in exactly one iteration.

3 BEYOND UNIFORM PACING

We supplement our main result by further generalizing our algorithm to handle multiple different objectives throughout an advertising campaign. Concretely, we demonstrate that our algorithm can further capture non-uniform pacing and mitigate subthreshold budgets.

3.1 Non-Uniform Pacing

While our main analysis demonstrates the stability of average bidding for the desired goal of uniform pacing, the algorithm can be simply adapted to meet changing target spend amounts throughout the campaign. For instance, if an advertiser wants to decrease spending during the morning since the bulk of their target market is not online, they may decrease their target spend amount until later in the day. A variant of this form is easily achieved by adjusting the submitted budget for any given time period(s). This is achieved by adding another input parameter $\kappa_t \in (0, 1]$ which scales the budget by the multiplier at any given time point.

3.2 Subthreshold Budgets

In the instance where users submit a budget to the DSP that are too small for meaningful uniform pacing, the provider can implement a variant on our algorithm for forcibly exhausting a budget early in the campaign period, thus resulting in the user needing to update with a larger budget for the next campaign.¹ This is achieved by implementing a “virtual budget” where the DSP simply scales up the input budget and thus the algorithm establishes a higher average bid than is feasible for the user’s actual budget, leading to early exit from the campaign. While this outcome is counter to the presented practicality of our algorithm, it is an effective means by which a provider can ensure that advertisers are submitting a high enough budget to be competitive in the market space.

The full version of our paper is available on arXiv and contains proofs of our theoretical claims as well as experimental validation of the algorithm on real-world datasets.

¹We note that this is common practice at certain companies, in an effort to force users to meet a minimum budget threshold.

ACKNOWLEDGEMENTS

MohammadTaghi Hajiaghayi was supported in part by NSF CCF grants 2114269 and 2218678. Max Springer was supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 1840340. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. The work is partially support by DARPA QuICC, NSF AF:Small # 2218678, and NSF AF:Small # 2114269.

REFERENCES

- [1] Deepak Agarwal, Souvik Ghosh, Kai Wei, and Siyu You. 2014. Budget Pacing for Targeted Online Advertisements at LinkedIn. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (New York, New York, USA) (*KDD '14*). Association for Computing Machinery, New York, NY, USA, 1613–1619. <https://doi.org/10.1145/2623330.2623366>
- [2] James Dugundji Andrzej Grana. 2003. *Fixed Point Theory*. Springer, New York.
- [3] Vashist Avadhanula, Riccardo Colini-Baldeschi, S. Leonardi, Karthik Abinav Sankararaman, and Okke Schrijvers. 2021. Stochastic bandits for multi-platform budget optimization in online advertising. *Proceedings of the Web Conference 2021* (2021).
- [4] S. Balseiro and Y. Gur. 2017. Learning in Repeated Auctions with Budgets: Regret Minimization and Equilibrium. *Proceedings of the 2017 ACM Conference on Economics and Computation* (2017).
- [5] S. Balseiro, A. Kim, Mohammad Mahdian, and V. Mirrokni. 2021. Budget-Management Strategies in Repeated Auctions. *Oper. Res.* 69 (2021), 859–876.
- [6] C. Borgs, J. Chayes, Nicole Immorlica, K. Jain, O. Etesami, and Mohammad Mahdian. 2007. Dynamics of bid optimization in online advertisement auctions. In *WWW '07*.
- [7] Han Cai, Kan Ren, Weinan Zhang, Kleanthis Malialis, Jun Wang, Yong Yu, and Defeng Guo. 2017. Real-Time Bidding by Reinforcement Learning in Display Advertising. *CoRR abs/1701.02490* (2017). [arXiv:1701.02490](http://arxiv.org/abs/1701.02490) <http://arxiv.org/abs/1701.02490>
- [8] A. Celli, Riccardo Colini-Baldeschi, Christian Kroer, and Eric Sodomka. 2021. The Parity Ray Regularizer for Pacing in Auction Markets. *ArXiv abs/2106.09503* (2021).
- [9] Vincent Conitzer, Christian Kroer, Eric Sodomka, and N. Stier-Moses. 2018. Multiplicative Pacing Equilibria in Auction Markets. In *WINE*.
- [10] Zhe Feng, Chara Podimata, and Vasilis Syrgkanis. 2018. Learning to Bid Without Knowing your Value. *Proceedings of the 2018 ACM Conference on Economics and Computation* (2018).
- [11] Joaquin Fernandez-Tapia. 2015. An analytical solution to the budget-pacing problem in programmatic advertising. *Journal of Information and Optimization Sciences* 40 (November 2015). <https://doi.org/10.1080/02522667.2017.1303946>
- [12] Sahin Cem Geyik, Abhishek Saxena, and Ali Dasdan. 2015. Multi-Touch Attribution Based Budget Allocation in Online Advertising. [arXiv:1502.06657](https://arxiv.org/abs/1502.06657) [cs.AI]
- [13] Grégoire Jauvion and Nicolas Grislain. 2020. Optimal Allocation of Real-Time-Bidding and Direct Campaigns. *CoRR abs/2006.07070* (2020). [arXiv:2006.07070](https://arxiv.org/abs/2006.07070)
- [14] Orcun Karaca, Pier Giuseppe Sessa, A. Leidi, and M. Kamgarpour. 2019. No-Regret Learning from Partially Observed Data in Repeated Auctions. *ArXiv abs/1912.09905* (2019).
- [15] Kuang-Chih Lee, Ali Jalali, and Ali Dasdan. 2013. Real Time Bid Optimization with Smooth Budget Delivery in Online Advertising. *CoRR abs/1305.3011* (2013). <http://arxiv.org/abs/1305.3011>
- [16] Alfonso Lobos, Paul Grigas, Zheng Wen, and Kuang chih Lee. 2018. Optimal Bidding, Allocation and Budget Spending for a Demand Side Platform Under Many Auction Types. [arXiv:1805.11645](https://arxiv.org/abs/1805.11645) [math.OC]
- [17] Thomas Nedelec, Nouredine El Karoui, and Vianney Perchet. 2019. Learning to bid in revenue-maximizing auctions. *ArXiv abs/1902.10427* (2019).
- [18] Claudia Perlich, Brian Dalessandro, Rod Hook, Ori Stitelman, Troy Raeder, and Foster Provost. 2012. Bid Optimizing and Inventory Scoring in Targeted Online Advertising. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Beijing, China) (*KDD '12*). Association for Computing Machinery, New York, NY, USA, 804–812. <https://doi.org/10.1145/2339530.2339655>
- [19] Walter Rudin. 1976. *Principles of mathematical analysis / Walter Rudin* (3d ed. ed.). McGraw-Hill New York. <http://www.loc.gov/catdir/toc/mh031/75017903.html>
- [20] Steven H. Strogatz. 2000. *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry and Engineering*. Westview Press.
- [21] Srinivas Tunuguntla and Paul R. Hoban. 2021. A Near-Optimal Bidding Strategy for Real-Time Display Advertising Auctions. *Journal of Marketing Research* 58 (2021), p. 1–21. <https://doi.org/10.1177/0022243720968547>
- [22] J. Weed, Vianney Perchet, and P. Rigollet. 2016. Online learning in repeated auctions. In *COLT*.
- [23] Jian Xu, Kuang-chih Lee, Wentong Li, Hang Qi, and Quan Lu. 2015. Smart Pacing for Effective Online Ad Campaign Optimization. *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Aug 2015). <https://doi.org/10.1145/2783258.2788615>
- [24] Weinan Zhang, Shuai Yuan, and Jun Wang. 2014. Optimal Real-Time Bidding for Display Advertising. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (New York, New York, USA) (*KDD '14*). Association for Computing Machinery, New York, NY, USA, 1077–1086. <https://doi.org/10.1145/2623330.2623633>