AUCTIONS FOR ONLINE ADVERTISING WITH CONSTRAINTS

ABSTRACT

Auctions are central in online advertising and are deployed billions of times a day. Modern real-world auctions are increasingly constrained by contracts, budgets, privacy and other concerns. Classic work on revenue-maximizing auctions turns out to be insufficient as such constraints may not be satisfied automatically, and alternate solutions largely resort to heuristic techniques. We prove that optimal constrained auctions in many natural settings turn out to be a variant of Myerson's auction when the virtual values are *shifted* and/or *scaled*, give efficient algorithms for finding them, and use real-world data to validate the benefit of our approach vs the state-of-the-art.

INTRODUCTION

Many Internet companies generate revenue by selling the advertisement space on their webpages, and revenue from online advertising is now over $50 Billion per year in the US alone. Ads are often sold via an *ad exchange* in which publishers post advertisement slots, advertisers post bids, and in order to match the ad to a user an auction is run every time a user loads the page [3]. The bids are allowed to vary depending on the properties of the user who loads the page, as some users may be have more value than others to the advertiser. These values can be thought of as being drawn (independently) from known distributions. In this setting, the classical result of Myerson [4] can be used to obtain a revenue-optimal auction, which is then deployed repeatedly. However, the world of online advertising has gained significant complexity in the last few years and auctions are often required to satisfy important additional constraints, e.g., related to contracts, budgets, privacy or financial risk. For example, ad exchanges often engage in outside *contracts* [2], which specify a price and number of guaranteed ad slots to an advertiser; this is done in order to attract large advertisers to their exchange. In practice, such constraints can make it impossible to repeatedly deploy Myerson's auction as-is because it may not satisfy the contracts; this leads to significant algorithmic challenges. In [1], these challenges are explored in detail and it is noted that "the central problem for the advertisement server is to decide which ad(s) to serve each user while satisfying all the contracts", moreover, *"the key question here is how to maximize auction revenue while ensuring that all reservation contract buyers are satisfied."* In this work we answer this question by providing the optimal auction that maximizes revenue while satisfying contract constraints. More generally, we consider how to tackle a variety of constraints that arise in online auctions, giving a characterization and efficient algorithms for finding the optimal auctions.

Our Contributions

Often, constraints as outlined above can be interpreted as constraints *in expectation*. The reason is that billions of auctions are run in one day and each auction is effectively independent of the other; thus, the number of items won by an agent is highly concentrated around its expectation. Hence, constraints including contracts, budgets, privacy or risk aversion, can often be formulated as either constraints on the expected probabilities of winning or on the expected payments of the agents. Thus, conceptually, we study the problem of designing auctions that optimize the revenue where the range of allowable expected probabilities or expected payments are restricted. This steps beyond the classical Myerson setting of Bayesian optimal auction design by imposing constraints on the expected outcomes; a revenue-optimal auction for this setting must hence allocate items in a clever manner in order to maximize revenue while still satisfying the constraints.

Characterization of Optimal Constrained Auctions. Our first result is a mathematical characterization of optimal auctions in the constrained settings.

Theorem (Informal)**.** *The optimal probability-constrained auction is defined by a unique vector $\boldsymbol{\alpha}$; it applies Myerson's allocation and pricing rule to the virtual values of each agent* shifted *by $\boldsymbol{\alpha}$. Similarly, the optimal payment-constrained auction is defined by a unique vector $\boldsymbol{\beta}$; it applies Myerson's allocation and pricing rule to the virtual values* scaled *by $\boldsymbol{\beta}$.*

We prove this by setting up the appropriate optimization problem for which we can then apply Lagrangian relaxation to attain our results. This characterization is rather general and can be extended to other single-parameter environments.

Efficient Algorithm for Finding Optimal Constrained Auctions. Unfortunately, the characterization above is not constructive, and it is easy to show that arbitrary constraints may be NP-hard to satisfy. However, the type of natural constraints mentioned above simply require restricting the expected probability (or expected payment) of each agent to fall within a fixed range. For such *interval-constraints* we provide an efficient algorithms for finding the optimal $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$.

Theorem (Informal)**.** *There exists an efficient algorithm for finding the optimal $\boldsymbol{\alpha}$ (respectively $\boldsymbol{\beta}$) for single-item auctions when there are interval constraints on the expected probabilities (respectively expected payments) under natural assumptions on the distributions.*

To prove this, we first show that we can construct an efficient oracle that gives the optimal $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$ when agents have exact constraints – i.e., the probabilities *exactly* match pre-defined quantities – under natural conditions on the distributions. We then show that we can find an optimal expected probability (or payment) vector that satisfies the interval constraints by recursively adding one tight constraint at a time in a particular order. This allows us to find the optimal expected probability (or payment) vector after linearly many iterations, and then, using our oracle, the optimal $\boldsymbol{\alpha}$ or $\boldsymbol{\beta}$.

Applications and Empirical Results. We show how to translate contract and other constraints mentioned above into constraints on the expected probabilities or payments and demonstrate our auction's effectiveness using historical data from the Yahoo! online auction dataset [5], a platform for display advertising transactions. Since the auction format was a second-price auction and hence truthful, we can interpret bids as values and thus simulate how bidders would counterfactually behave under other auctions. In this manner, we conduct experiments in order to compare the revenue of our auctions against state-of-the-art approaches for various constrained settings. For example, our experiments show that for contracts, our auction attains more revenue than state-of-the-art auctions, roughly by a margin of 5% over the best alternative; a significant amount in such billion-dollar markets. As another example, when there are budget constraints, Myerson's auction is neither optimal nor truthful. In this setting, our algorithm is not only truthful, but it outperforms Myerson's auction with respect to revenue; e.g., even if only a single agent runs out of budget halfway through all of the auctions, a $\boldsymbol{\beta}$-scaled auction can attain approximately 0.5% more revenue than Myerson's. Finally, our work should be of particular interest to the online advertising community because it allows for the ease of running a single one-step auction repeatedly – an important requirement in practice – while simultaneously allowing the imposition of practical constraints.

References

[1] Nitish Korula, Vahab Mirrokni, and Hamid Nazerzadeh. Optimizing display advertising markets: Challenges and directions. *IEEE Internet Computing*, 20(1):28–35, 2016.

[2] Vahab Mirrokni and Hamid Nazerzadeh. Deals or no deals: Contract design for online advertising. *Available at SSRN 2693037*, 2015.

[3] S Muthukrishnan. Ad exchanges: Research issues. In *International Workshop on Internet and Network Economics*, pages 1–12. Springer, 2009.

[4] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58—73, 1981.

[5] Yahoo! A1 - search marketing advertiser bidding data, version 1.0. `https://webscope.sandbox.yahoo.com/catalog.php?datatype=a`.