

How to Approximate Optimal Auctions

NIMA HAGHPANAH

Northwestern University

and

NICOLE IMMORLICA

Northwestern University

and

VAHAB MIRROKNI

Google Research, New York

and

KAMESH MUNAGALA

Duke University

Bayesian auction design investigates how to sell scarce resources to agents with private values drawn according to known distributions. A natural objective in this setting is revenue maximization. The seminal work of Roger Myerson presents a recipe for designing revenue-maximizing auctions in single-parameter settings. Myerson defines a function that maps each agent's value to a new number called the *virtual value*. The optimal auction then chooses the solution with maximum virtual surplus subject to the feasibility constraints. Often this underlying optimization problem is NP-hard and is complicated by the fact that virtual values may be negative. Prior work approximates the optimal auction by approximating this optimization problem. Such a solution can be interpreted as approximating the optimal auction point-wise, i.e., for each realization of values the revenue of the derived auction is close to that of the Myerson auction. In this letter, we suggest a different approach:

*Approximate the maximum virtual surplus on average
with respect to the induced virtual value distributions.*

Our approach has the advantage that sometimes the point-wise optimization problem is not approximable to within any reasonable factor. However, by leveraging the fact that any distribution of virtual values has non-negative expectation, it is sometimes possible to get good average-case approximations. Furthermore, the optimal auction is itself an average-case guarantee: it maximizes the revenue on average with respect to the distributions, but it may lose substantial revenue on certain value profiles. Thus our average-case guarantee is without loss of generality.

We showcase our approach using the problem of selling goods with positive network externalities. For this problem, the underlying optimization problem is inapproximable, yet with our approach we are able to prove that a natural greedy auction is a constant-factor approximation.

Categories and Subject Descriptors: J.4 [Social and Behavior Sciences]: Economics

General Terms: Algorithms, Economics, Theory

Authors' addresses: nima.haghpanah@gmail.com, nicimm@gmail.com, mirrokni@gmail.com,
kamesh@cs.duke.edu.

1. INTRODUCTION

Auctions are often used to allocate resources to agents. In the most simple auction, a single item, say a painting, is sold to a single agent. The seller is unsure of the exact value of the agent, but knows the distribution of the agent's value. Based on this distribution, the seller sets a price that maximizes his expected revenue with respect to the value distribution. For example, if the agent's value is uniformly distributed in the interval $[0, 1]$, then for a fixed price $p \in [0, 1]$, the expected revenue is $p(1 - p)$. A price of $p = 1/2$ achieves the maximum expected revenue of $1/4$ for this setting. For any distribution, the solution to this simple setting is a straight-forward maximization program.

For more complex settings with multiple agents and potentially complicated constraints on the set of feasible outcomes, the seminal paper by Myerson [Myerson 1981] fully characterizes auctions that maximize revenue in expectation over the value distributions. By this characterization, the expected revenue of any auction is equal to the total expected *virtual value* $\phi(v) = v - \frac{1-F(v)}{f(v)}$ of the allocated agents, where $F(\cdot)$ is the cumulative distribution function and $f(\cdot)$ is the corresponding density function of the value distribution.¹ Thus the auction that maximizes expected virtual value is optimal. For the example above, the virtual value function is $\phi(v) = 2v - 1$, and so the auction that allocates to the agent whenever his virtual value is positive has revenue of $(1/2)E[2v - 1 | v > 1/2] = 1/4$, matching the above pricing scheme as expected.

In many applications, computing the maximum expected virtual value is computationally hard due to the feasibility constraints on the set of winners. In such settings, the best known computationally tractable auctions first design an α -approximation algorithm for the underlying virtual value optimization problem, and then use this solution to achieve an α fraction of the optimal revenue. Unfortunately, for many settings, there is provably no good worst-case approximation algorithm for the relevant optimization problem. In particular, for feasibility constraints that are not downward-closed², the fact that virtual values may be negative often implies hardness of approximation results. However, an approximately optimal auction need not provide worst-case guarantees: an average-case guarantee is sufficient. This is because the optimal auction anyway provides an optimal average-case guarantee, and there is no mechanism with high revenue for every instantiation of values.

This observation seems useless at first glance: a worst-case instance can be simulated by an average-case one where the distribution is a point mass. However, virtual value distributions are not arbitrary. They are guaranteed to have non-negative expectation, suggesting that non-downward-closed problems that are hard in the worst-case may be approximable on average. This fact enables us to design algorithms that perform well on average with respect to *any* distribution whose expectation is non-negative. Such algorithms imply approximately optimal auctions.

¹Note the prices are not explicit in this characterization; they are implied by the incentive constraints. To gain intuition for this, observe, as the instantiation of values are known only to agents, prices can not always extract full surplus or otherwise agents would misreport their values.

²Feasibility constraints are downward-closed if any subset of feasible winners is also feasible.

2. SHOWCASE PROBLEM

In prior work [Haghpanah et al. 2011], we use our technique to design auctions in social networks for goods that exhibit *single-parameter submodular network externalities* in which a bidder’s value for an outcome is a fixed private type times a known submodular function of the allocation of his friends. Our main result considers *step-function externalities* in which a bidder’s value for an outcome is either zero, or equal to his private type if at least one friend has the good. In this setting, the virtual value optimization problem equates to finding, in a vertex-weighted graph with possibly negative vertex weights, a maximum-weight subset of vertices whose induced subgraph has no singleton components. We observe via reduction to set-buying that approximating this optimization problem within even a linear factor on *every* sampling of the values is NP-hard. Even on average, we prove that our problem remains APX-hard. However, we are able to design constant approximations for several versions of the problem.

We first note that there is a simple $(1/2)$ -approximation for our problem. The algorithm divides the graph into two subsets of vertices, such that each vertex in each set has a neighbor in the other. This can be done, for example, by constructing a spanning tree of the graph and then taking a bipartite partitioning of it. The allocation strategy is to then pick the set with better expected revenue (computed with samples from the value distributions and without looking at the instantiation of values), extract revenue from that set (by allocating to each agent in the set with positive virtual value), and allocate to agents in the other set in order to maintain feasibility (this does not decrease the total expected virtual value on average as the expected virtual value of these agents is non-negative). This very simple algorithm does not use the structure of the social network in any deep way, and is therefore unable to give better approximations in even very simple social networks consisting of a single edge. In order to leverage knowledge of the network structure, we consider a greedy algorithm that iteratively allocates to influential vertices and their neighbors. Our main result shows that this can be used to obtain an $\frac{e}{e+1} \approx 0.73$ -approximation to the optimal revenue for any distribution of values.

3. CONCLUSION

This letter suggests a new approach to designing optimal auctions. The approach can be summarized as follows:

- (1) Design an algorithm A for the virtual value optimization problem. Let $E_{x \sim \mathcal{D}}[A(x)]$ be the expected value of A when inputs are drawn according to distribution \mathcal{D} . Let $E_{x \sim \mathcal{D}}[\text{OPT}(x)]$ be an algorithm that maximizes this expectation.
- (2) Calculate the average-case approximation $\alpha = \min_{\mathcal{D}}[E[A(x)]/E[\text{OPT}(x)]]$ where the minimization is over distributions that have non-negative expectation.

Then the resulting auction is an α -approximation to the optimal auction.

We demonstrated an application of this technique to selling a good with network externalities. In general, we believe our technique is of great use in non-downward closed settings.

REFERENCES

- HAGHPANAH, N., IMMORLICA, N., MIRROKNI, V., AND MUNAGALA, K. 2011. Optimal auctions with positive network externalities. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC)*. 11–20.
- MYERSON, R. 1981. Optimal auction design. *Mathematics of Operations Research*, 58–73.