

On Optimal Multidimensional Mechanism Design

YANG CAI, CONSTANTINOS DASKALAKIS and S. MATTHEW WEINBERG

Massachusetts Institute of Technology

We solve the *optimal multi-dimensional mechanism design problem* when either the number of bidders is a constant or the number of items is a constant. In the first setting, we need that the values of each bidder for the items are i.i.d., but allow different distributions for each bidder. In the second setting, we allow the values of each bidder for the items to be arbitrarily correlated, but assume that the bidders are i.i.d. For all $\epsilon > 0$, we obtain an efficient additive ϵ -approximation, when the value distributions are bounded, or a multiplicative $(1-\epsilon)$ -approximation when the value distributions are unbounded, but satisfy the Monotone Hazard Rate condition. When there is a single bidder, we generalize these results to independent but not necessarily identically distributed value distributions, and to independent regular distributions.

Categories and Subject Descriptors: F.2.2 [**Theory of Computation**]: Algorithm Analysis and Problem Complexity

General Terms: Algorithms, Economics, Mechanism Design

Additional Key Words and Phrases: Optimal Multidimensional Mechanism Design

Introduction. In his seminal paper, Myerson [1981] studies the following auction design problem. A seller has a single item to sell to a set of interested bidders and looks to leverage stochastic knowledge about the bidders' values for the item to design an auction maximizing her revenue, in expectation. Myerson provides an elegant, closed-form solution to this problem in a single sweep that encompasses many other auction settings, referred to in Economics as *single-dimensional*. These can be very general auction settings, such as selling multiple items to multiple bidders with various feasibility constraints on the allocations of items to bidders that are feasible, but need to satisfy that each bidder derives the same (unknown) value for receiving service from the auctioneer independently of the type of service, e.g. what bundle of items she ends up winning in the auction.

Following Myerson's work, a large body of research in both Economics and Engineering has been devoted to extending this result to *multi-dimensional settings*, where the above single-dimensionality assumption is violated. A natural and simple example is the following: The auctioneer has n (heterogeneous) items to sell to m bidders who may have different values for each item, additive valuations for bundles of items and potentially budget or demand constraints.¹ The auctioneer still looks to design an auction that maximizes her expected revenue, leveraging stochastic knowledge about the bidders' values for the items. An even simpler setting is when there is a single unit-demand bidder and the auctioneer only wants to compute prices for the items that maximize her revenue.² Despite the simplicity

¹We refer to this setting as the *Auction Setting* in this paper.

²We refer to this setting as the *Pricing Setting* in this paper.

Authors' addresses: ycai@csail.mit.edu, costis@csail.mit.edu, smw79@mit.edu

of these settings, there is no efficient, revenue-optimal mechanism known to date, and indeed it looks that we are far from obtaining such mechanism; see the survey by Manelli and Vincent [2007] for an overview of what is known about this problem in Economics.

On the algorithmic front, previous work by Chawla et al. [2007; 2010] has provided constant factor approximations to the unit-demand multi-dimensional pricing problem (Pricing Setting above.) They propose an elegant reduction of this problem to a related single-dimensional problem, such that Myerson's solution of the latter provides a price vector that gives a factor of 2 of the optimal revenue for the former. Moreover, their solution can be computed efficiently if the values of the items come from *regular distributions*. (This is a class of distributions commonly studied in mechanism design.) Chawla et al. [2010] and Bhattacharya et al. [2010] extend this result, obtaining constant factor approximations in more general multi-dimensional mechanism design problems, which include the multiple-item multiple-bidder problem (Auction Setting above.) The mechanism of Chawla et al. is a sequential posted price mechanism obtained via a direct reduction to Myerson's setting and applies to unit-demand bidders with no budget constraints, while the mechanism of Bhattacharya et al. is obtained via linear programming relaxations of the problem, and can also accommodate budget constraints and non-unit demand bidders.

And, while our algorithmic understanding of the optimal mechanism design problem is solid as far as constant factor approximations go, there has been virtually no result in designing efficient revenue-optimal mechanisms for multi-dimensional settings. In particular, one can argue that the previous approaches [Bhattacharya et al. 2010; Chawla et al. 2007; Chawla et al. 2010] are inherently limited to constant factor approximations, as ultimately the revenue of these mechanisms is compared against the optimal revenue in a related single-dimensional setting [Chawla et al. 2007; Chawla et al. 2010], or a linear programming relaxation of the problem [Bhattacharya et al. 2010]. Our focus in this work is to fill this important gap in the algorithmic mechanism design literature, i.e. to *obtain efficient near-optimal multi-dimensional mechanisms*, achieving a $(1 - \epsilon)$ -approximation to the optimal revenue in polynomial time, for any desired accuracy $\epsilon > 0$.

Our Results. In recent work [Daskalakis and Weinberg 2011; Cai and Daskalakis 2011], we provide efficient algorithms for the multi-dimensional auction and pricing problems (Auction and Pricing Settings defined above,) with arbitrarily good approximations under certain restrictions on the value distributions. We solve the multi-dimensional auction problem when either the number of bidders is a constant or the number of items is a constant. In the first setting, we need that the values of each bidder for the items are i.i.d., but allow different distributions for each bidder. In the second setting, we allow the values of each bidder for the items to be arbitrarily correlated, but assume that the bidders are i.i.d. For the pricing problem, we can relax our assumptions, only requiring the value distributions to be independent but not identical. As far as the quality of our solution goes, for all $\epsilon > 0$, we efficiently obtain an additive ϵ -approximation to the optimal revenue when the value distributions are bounded. Furthermore, we can strengthen all our results to efficiently obtain a $(1 - \epsilon)$ -fraction of the optimal revenue when the value distributions are unbounded, but satisfy the Monotone Hazard Rate condition (this is a subclass

of regular distributions commonly studied in mechanism design.) Especially for the pricing problem, we can also obtain a $(1 - \epsilon)$ -fraction of the optimal revenue when the value distributions are independent and regular in quasi-polynomial time.

Formally, we show the following theorems (as well as their extensions, described above, to MHR and regular distributions.)

THEOREM 1.1 AUCTION SETTING: CONSTANT NUMBER OF BIDDERS. *For all k and $\epsilon > 0$, there is an efficient algorithm for computing an explicit mechanism whose revenue is within an additive ϵ of the optimal revenue when the values of each of the k bidders for each of the n items are drawn i.i.d. from an arbitrary (possibly different for each bidder) distribution on $[0, 1]$.*

THEOREM 1.2 AUCTION SETTING: CONSTANT NUMBER OF ITEMS. *For all k and $\epsilon > 0$, there is an efficient algorithm for computing an explicit mechanism whose revenue is within an additive ϵ of the optimal revenue when the valuation vectors for each of the m bidders for the k items are drawn independently from an arbitrary (possibly correlated) distribution over $[0, 1]^k$.*

THEOREM 1.3 PRICING SETTING: NON-I.I.D. DISTRIBUTIONS. *For all $\epsilon > 0$, there is an efficient algorithm for computing a price vector whose revenue is within an additive ϵ of the optimal revenue when the values of the bidder are independent (but not necessarily identical) and distributed in $[0, 1]$.*

Techniques. Our solutions to both problems come in two parts, probabilistic and algorithmic. On the probabilistic side, we obtain structural theorems for the optimal mechanism as a function of the structure in the value distributions of the bidders. First, we prove a very general symmetrization lemma, which is reminiscent of Nash’s theorem on symmetric Nash equilibria [Nash 1951] and applies to auction settings that are much broader than the ones considered here. Our lemma states that, if the distribution of the bidders’ values for the items (viewed as a distribution over $\mathbb{R}^{m \times n}$) satisfies any symmetries, then any given revenue-optimal incentive compatible mechanism can be turned into one that simultaneously satisfies all these symmetries, in a formal technical sense of the word “satisfies.” Second, we obtain extreme value theorems—discussed in further detail below, which allow us to restrict the prices used by a near-optimal mechanism to a finite interval.

And, while our structural theorems bear witness to the existence of a simple solution, it is non-trivial to efficiently find one. Our second contribution is algorithmic: we show how to efficiently find the simple solutions guaranteed by our structural theorems. In the pricing problem, instead of searching over all possible price vectors (which are exponentially many even if there is just two possible prices for each item,) we shift our focus to the set of all possible revenue distributions that may arise from price vectors, computing a sparse probabilistic cover of this space under the total variation distance between distributions via dynamic programming. In the auction problem, we prove additional structural theorems that allow us to write a polynomial-size linear program for computing an optimal ϵ -truthful mechanism. We then complete our solution using a reduction similar to that of [Hartline et al. 2011] to turn our ϵ -truthful mechanism into one that is exactly truthful.

Both results rely on extreme value theorems, which might be of independent interest. For example, we establish the following theorem for independent Mono-

tone Hazard Rate distributions: For any set of independent, but not necessarily identically distributed, MHR random variables $\{X_i\}_{i \in [n]}$, there exists an anchoring point β such that the contribution to the expectation of $\max_i X_i$ from values above $\frac{1}{\epsilon} \log \frac{1}{\epsilon} \beta$ is only $O(\epsilon\beta)$, while with constant probability $\max_i X_i \geq \beta$. This implies that the optimal revenue is $\Omega(\beta)$, if the bidders' value distributions are independent and MHR. Moreover, we can ignore “extreme” bidders, because even if we could perfectly price to extract full surplus from every such extreme bidder, we would still only make an extra $O(\epsilon\beta)$ of revenue, i.e. at most an extra fraction of $O(\epsilon)$ of the optimal revenue.

Conclusion and Extensions. In conclusion, this paper provides the first near-optimal efficient algorithms for the multi-dimensional pricing problem, for a unit-demand bidder whose values are independent (but not necessarily identically distributed.) In addition, we provide the first near-optimal efficient algorithms for the multi-dimensional multi-bidder multi-item auction problem, for the settings described in Theorems 1.2 and 1.1. Our results provide algorithmic, structural and probabilistic insights into the properties of the optimal solution for the case of MHR, regular, and more general distributions.

In the pricing frontier, it would be interesting to extend our results (algorithmic and/or structural) to mechanisms that price lotteries over items [Thanassoulis 2004; Briest et al. 2010], to bundle-pricing [Manelli and Vincent 2006] and to a non unit-demand bidder. We can certainly obtain such extensions, albeit when sizes of lotteries, bundles, demand etc. are an absolute constant. We believe that our extreme value theorems, and our probabilistic view of the problem in terms of revenue distributions will be helpful in obtaining more general results. We also leave the complexity of the exact problem for independent distributions as an open question, and conjecture that it is *NP*-hard, referring the reader to [Briest 2008] for hardness results in the case of correlated distributions.

In the auction frontier, it would be interesting to extend our results to cover all the cases we can solve for the pricing problem, specifically by replacing the i.i.d. assumption with just independence in Theorems 1.2 and 1.1. One would need to combine our techniques for the auction problem with our probabilistic covering theorem for the pricing problem, but we suspect that significant technical work will be required. In contrast, our work on the auction problem already prices lotteries and bundles, and does not require a unit-demand assumption. Our work on the auction problem is very general in the types of input we can consider, but so far restricted in the class of value distributions we can solve.

REFERENCES

- BHATTACHARYA, S., GOEL, G., GOLLAPUDI, S., AND MUNAGALA, K. 2010. Budget constrained auctions with heterogeneous items. In *STOC*. 379–388.
- BRIEST, P. 2008. Uniform budgets and the envy-free pricing problem. In *ICALP (1)*. 808–819.
- BRIEST, P., CHAWLA, S., KLEINBERG, R., AND WEINBERG, S. M. 2010. Pricing randomized allocations. In *SODA*. 585–597.
- CAI, Y. AND DASKALAKIS, C. 2011. Extreme-value theorems for optimal multidimensional pricing. *Manuscript*.
- CHAWLA, S., HARTLINE, J. D., AND KLEINBERG, R. D. 2007. Algorithmic pricing via virtual valuations. In *ACM Conference on Electronic Commerce*. 243–251.

- CHAWLA, S., HARTLINE, J. D., MALEC, D. L., AND SIVAN, B. 2010. Multi-parameter mechanism design and sequential posted pricing. In *STOC*. 311–320.
- DASKALAKIS, C. AND WEINBERG, S. M. 2011. On optimal multi-dimensional mechanism design. *Manuscript*.
- HARTLINE, J., KLEINBERG, R., AND MALEKIAN, A. 2011. Bayesian incentive compatibility via matchings. In *SODA*. 734–747.
- MANELLI, A. M. AND VINCENT, D. R. 2006. Bundling as an optimal selling mechanism for a multiple-good monopolist. *Journal of Economic Theory* 127, 1, 1–35.
- MANELLI, A. M. AND VINCENT, D. R. 2007. Multidimensional mechanism design: revenue maximization and the multiply-good monopoly. *Journal of Economic Theory* 137, 1, 153–185.
- MYERSON, R. B. 1981. Optimal auction design. *Mathematics of Operation Research* 6, 58–73.
- NASH, J. 1951. Non-Cooperative Games. *The Annals of Mathematics* 54, 2 (Sept.), 286–295.
- THANASSOULIS, J. 2004. Haggling over substitutes. *Journal of Economic Theory* 117, 2, 217–245.