Mechanism for Optimally Trading Off Revenue and Efficiency in Multi-unit Auctions*

Anton Likhodedov
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
likh@cs.cmu.edu

Tuomas Sandholm
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
sandholm@cs.cmu.edu

ABSTRACT

We study auctioning multiple units of the same good to potential buyers with single unit demand (i.e. every bidder wants only one unit of the good). Depending on the objective of the seller, different selling mechanisms are desirable. The Vickrey auction with a truthful reserve price is optimal when the objective is efficiency - allocating the units to the parties who values them the most. The Myerson auction is optimal when the objective is the seller's expected utility. These two objectives are generally in conflict, and cannot be maximized with one mechanism. In many real-world settings—such as privatization and competing electronic marketplaces—it is not clear that the objective should be either efficiency or seller's expected utility. Typically, one of these objectives should weigh more than the other, but both are important. We account for both objectives by designing a new *deterministic* dominant strategy auction mechanism that maximizes expected social welfare subject to a minimum constraint on the seller's expected utility. This way the seller can maximize social welfare subject to doing well enough for himself.

Categories and Subject Descriptors

J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms

Algorithms, Economics

Keywords

economics, efficiency, tradeoff, revenue, auctions

1. INTRODUCTION

Electronic commerce has spawned the use of increasingly sophisticated auction mechanisms. In many ecommerce settings the bidders are automated agents that are programmed to act rationally even in complex situations. Also human users of ecommerce systems are typically quite savvy and able to recognize attractive properties of sophisticated mechanisms. It is in the ecommerce setting that we believe that unintuitive auction mechanisms are palatable, and it is justified to introduce auctions that meet more complicated objectives.

One example is the Vickrey auction [2]. In this auction with q_0 units of the same good on sale, q_0 highest bidders win, but only pay the price of the first *unsuccessful* bid. The Vickrey auction maximizes economic *efficiency*, aka. *social welfare* (assuming the reserve price is set to equal the seller's valuation for a unit of good being sold), that is, the units end in the hands of the party who values it the most.

Another example is the Myerson auction [1], which maximizes the *seller's expected utility* (expected *revenue* in case the seller does not value the good).The unintuitive aspect of the Myerson auction is that it sometimes allocates the goods to bidders other than the q_0 highest bidders.

Expected social welfare and the seller's expected utility cannot be maximized with the same auction mechanism in general because these objectives conflict. Furthermore, in many real-world settings it is not clear that the objective should be either of the two.

For example, most privatization auctions are motivated by the belief that private companies can make more efficient use of an asset than the government can. It seems thus reasonable to allocate the asset to the party who can make the most effective use of it, that is, to use efficiency as the auction objective. At the same time, the government would like to raise as much money from the sale as possible (maximize the seller's expected utility) because the asset is owned by the tax payers who prefer to pay for government expenditures out of the auction revenue rather than taxes.

We account for both objectives by designing a new deterministic auction mechanism that maximizes expected social welfare subject to a minimum constraint on the seller's expected utility. This way the seller can maximize social welfare subject to subject to doing well enough for himself. We prove that this auction mechanism belongs to a family of mechanisms, maximizing a linear combination of the seller's expected utility and expect. We then present an algorithm for determining the optimal value for this parameter and constructing an auction with desired characteristics. This approach is different from the traditional mechanism design - the mechanism is not completely specified upfront - rather we compute the particular mechanism for every instance of the problem (that is, given the constraint on the revenue and the distributions of buyers' valuations).

We also present a family of much simpler randomized mechanisms that *under some assumptions*, which we state later in the paper achieve the same expected revenue and efficiency. However, we argue, that randomization is inappropriate for the settings, which motivate the present work.

2. FRAMEWORK AND NOTATION

We focus on settings with one seller, multiple buyers with *single*

*This material is based upon work supported by the National Science Foundation under CAREER Award IRI-9703122, Grant IIS-9800994, ITR IIS-0081246, and ITR IIS-0121678.

EC'04, May 17–20, 2004, New York, New York, USA.
ACM 1-58113-711-0/04/0005.

unit demand, and multiple units of the same good on sale. For convenience buyers are indexed with numbers from 1 to n and the set of all buyers is denoted by $N = \{1 \dots n\}$. Index 0 always refers to the seller. We analyze the problem under the (potentially asymmetric) independent private values model: the valuation v_i of bidder i is drawn from the p.d.f. $f_i : [a_i, b_i] \rightarrow \mathbb{R}$ (the corresponding c.d.f. is denoted by F_i).The set of all possible combinations of valuations of bidders is $V = \times_{j \in N}[a_i, b_i]$. The vector of valuations is denoted by $v = (v_1, \dots, v_n)$.

By the *revelation principle* we may restrict attention to mechanisms where each bidder *truthfully* bids for a unit of good in the following sealed-bid format: each bidder i submits a bid v_i for the unit. The seller then computes the *allocation* $(p_1(v), \dots p_n(v))$ and the *payment* $(t_1(v), \dots, t_n(v))$, where $p_i(v)$ is the probability that bidder i receives the unit, and $t_i(v)$ is the amount bidder i pays.

Different auctions are usually evaluated either according to the expected utility of the seller:

$$U_0(p,t) = E_v\big[(q_0 - \sum_{i=1}^{n} p_i(v_i)) \cdot v_0 + \sum_{i=1}^{n} t_i(v)\big]$$

or expected social welfare:

$$SW(p) = E_v\left[\sum_{i=1}^{n} p_i(v_i) \cdot v_i + \big(q_0 - \sum_{i=1}^{n} p_i(v_i)\big) \cdot v_0\right]$$

In this paper we account for the importance of both objectives by setting up a constrained optimization problem: optimize one of the objectives (efficiency), subject to the following constraint on the other (seller's expected utility):

$$U_0(p,t) \geq R_0$$

and to the standard incentive compatibility (IC), individual rationality (IR) and probability normalization (PN) constraints. R_0 is the minimal expected utility that satisfies the seller.

3. OUTLINE OF THE DERIVATION

We first show that the optimal payment rule x is the same as in the Myerson auction: We also demonstrate that the problem of designing the optimal mechanism (p,t) can be reduced to an optimization problem in p only: the optimal allocation rule is computed by maximizing $SW(p)$ subject to the constraint

$$U_0(p) = \sum_{i=1}^{n} E_v\left[(v_i - v_0 - \frac{1-F_i(v_i)}{f_i(v_i)})p_i(v_i)\right] + q_0 v_0 \geq R_0 \qquad (3.1)$$

where $p(v)$ also satisfies the IC, IR, and PN constraints.

1. If Constraint 3.1 is inactive, the optimal auction is the standard Vickrey auction (yielding unconstrained global maximum of SW) with the reserve price set to equal the seller's valuation for a unit of the good on sale.

2. To solve the case where Constraint 3.1 is active, we show that the optimal allocation rule $p^{opt}(v)$ can be found as the maximum of

   $$\hat{L}(p,\lambda) = SW(p) + \lambda \cdot (\hat{U}_0(p) - R_0)$$

   with respect to (p,λ) on the convex set of feasible allocation rules. Then we derive the allocation rule $p^{\lambda}(v)$ that, for given λ, maximizes $\hat{L}(p,\lambda)$:

   $$p_i^{\lambda}(v) = \begin{cases} 1, & \text{if } \hat{c}_i^{\lambda}(v_i) > \hat{c}_j^{\lambda}(v_j), \text{ for all } j \neq i \\ & \text{and } \hat{c}_i^{\lambda}(v_i) > (1+\lambda)v_0 \\ 0, & \text{otherwise.} \end{cases} \qquad (3.2)$$

   Here $\hat{c}_i^{\lambda}(v_i)$ is a closest increasing continuous approximation to the *virtual valuation* $c_i^{\lambda}(v_i)$ of the bidder i with type v_i. Our virtual valuations are defined as follows:

   $$c_i^{\lambda}(v_i) = (1+\lambda)\left(v_i - \frac{\lambda}{1+\lambda} \cdot \frac{1-F_i(v_i)}{f_i(v)}\right)$$

   Our mechanism is similar to Myerson auction, but with different virtual valuations.

   The problem reduces to finding a value λ so that p^{λ} maximizes the objective $SW(p)$ subject to Constraint 3.1. Since Constraint 3.1 is active, the maximum can be found by solving the following integral equation:

   $$\hat{U}_0(p^{\lambda}) = R_0$$

   We prove that $\hat{U}_0(p^{\lambda})$ is increasing and continuous in λ. This allows us to find the optimal λ numerically, with a simple binary search:
   - set λ_{min} to 0 and λ_{max} to some value so that Constraint 3.1 is satisfied. and evaluate $\hat{U}_0(p^{\lambda})$ numerically in both cases.
   - Repeat the following step until $\lambda_{max} - \lambda_{min}$ converges to zero: Set $\lambda_{new} = \frac{\lambda_{min}+\lambda_{max}}{2}$. Construct the allocation rule $p^{\lambda_{new}}$ according to mechanism, given by 3.2 and check Constraint 3.1. If it is satisfied, set $\lambda_{max} = \lambda_{new}$, otherwise set $\lambda_{max} = \lambda_{new}$ and repeat. ($\hat{U}_0(p^{\lambda})$ is an n-dimensional integral, but it can be estimated with a Monte-Carlo method, which converges to the true value of the integral. Sampling v is easy because the valuations v_i can be drawn independently.)

4. RANDOMIZED MECHANISMS

When virtual valuations of the buyers c_i^{λ} are non-decreasing,[1] there exist simple randomized mechanisms where the Vickrey rules are used w.p. x and the Myerson rules w.p. $1-x$ that yield the same expected social welfare and seller's expected revenue as our mechanism.

Rather than computing the optimal λ using the algorithm above, the seller can use the distributions f_i to evaluate the expected revenue of the Myerson auction and the Vickrey auction in advance, and use these revenue values to analytically determine x.

However, randomization is often undesirable. In many settings an auction is only run once. For instance, each privatization auction usually has different participants and/or a different objects (company) for sale. Similarly, in Internet auctions, the set of buyers generally differs for every auction, as may the objects for sale. Now, say that in a given setting, the auction designer is unsatisfied with the seller's expected utility in the Vickrey auction, and with the expected social welfare of the Myerson auction. Still, the designer can be satisfied with the seller's expected utility and expected social welfare in our deterministic mechanism. So, the deterministic mechanism is satisfactory, but the randomized mechanism would run an unsatisfactory auction for sure. (For randomization to really make sense, the designer would have to be able to repeat the random drawing multiple times, i.e., to repeat the same auction in the same setting.)

5. REFERENCES

[1] Myerson, R. *Optimal auction design*, In Mathematics of operational research. Vol. 6, pp. 58-73, 1981.

[2] Vickrey W. *Counterspeculation, auctions and competitive sealed tenders*. Journal of Finance, 16 pp.8-37, 1961.

[1]this is called a *regular* case in economics literature