CS364B: Frontiers in Mechanism Design Lecture #1: Ascending and Ex Post Incentive Compatible Mechanisms*

Tim Roughgarden†

January 8, 2014

1 Introduction

These twenty lectures cover advanced topics in mechanism design. They assume familiarity with some of the material covered in the instructor's CS364A course — specifically, lectures 2–4 and 7–9.

Recall that mechanism design is the "science of rule-making." The goal is to understand how to design systems with strategic participants — autonomous decision-makers whose objectives are generally different from the the designer's — that have good performance guarantees. For example, a mechanism designer might want to compute a socially efficient allocation of scarce resources or raise significant revenue, while a mechanism participant only cares about its own utility.

2 Course Outline

The plan is to cover the following five topics.

1. (6 lectures.) Welfare-maximization in combinatorial auctions: tractable special cases and ascending implementations.

2. (4 lectures.) Welfare-maximization in combinatorial auctions: dominant-strategy approximation mechanisms for NP-hard cases.

3. (3 lectures.) Welfare-maximization in combinatorial auctions: better guarantees for weaker solutions concepts (undominated strategies and Bayes-Nash equilibria).

†Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305. Email: tim@cs.stanford.edu.

4. (4 lectures.) The price of anarchy in simple auctions.

5. (3 lectures.) Revenue-maximization in multi-parameter settings.

3 Review: The k-Vickrey Auction

Let's recall a basic example from last quarter.

Scenario #1:

- k identical items (even $k = 1$ is interesting);
- each bidder is *unit demand*, meaning it only wants one item;
- each bidder has a private *valuation* v_i for a item.

The valuation of a bidder is its maximum willingness-to-pay for an item. It is private in the sense that the seller and other bidders have no idea what it is.

We already know an excellent — awesome, even — auction for Scenario #1, the *k-Vickrey auction*. This auction works as follows.

1. Each bidder submits a bid b_i.
2. The top k bidders are awarded an item.
3. All winners pay the $(k + 1)$th highest bid.

Recall the second and third steps are the *allocation rule* ("who wins?") and *payment rule* ("at what price?") of the auction. The auction is *direct-revelation*, in the sense that bidders are effectively asked to report all of their private information up front.

Why is this a good auction? Because of its incentive, performance, and tractability/simplicity guarantees.

1. **(Incentive guarantee.)** The k-Vickrey auction is *dominant-strategy incentive compatible (DSIC)*. First, truthful bidding is a dominant strategy, guaranteed to maximize a bidder i's quasilinear utility $v_i \cdot x_i(\mathbf{b}) - p_i(\mathbf{b})$ over all b_i for any fixed $\mathbf{b}_{-i}$. (Here, $x_i(\mathbf{b})$ is 1 if i wins and 0 if i loses.) We proved this property in Lecture #2 of CS364A and won't do it again here. Second, truthful bidding guarantees non-negative utility, since the k-Vickrey auction charges each loser 0 and each winner a price less than its bid.

2. **(Performance guarantee.)** For each valuation profile $\mathbf{v}$, assuming truthful bids, the k-Vickrey auction awards the k items to the bidders with the highest valuations and hence maximizes the welfare $\sum_{i=1}^{n} v_i \cdot x_i(\mathbf{b})$ over all feasible allocations.

3. **(Tractability guarantee.)** The k-Vickrey auction runs in polynomial time. It is "simple" in almost any reasonable sense of the word.

Looking ahead, we'll strive to get as close as possible to the "holy grail" of these three simultaneous guarantees, for as many different scenarios as possible. For the first several lectures we'll focus on special cases where we can get everything we want. Then we'll move on to more complex problems where we provably can't get everything want, and we'll explore relaxations of the above three guarantees and trade-offs between them.

4 Ascending Implementations

Let's illustrate an "ascending auction" by example. The following is more or less the "English auction", familiar from movies, art auction houses, and so on. The basic idea is to keep raising the proposed selling price until the demand equals the supply. The parameter $\epsilon > 0$ is an a priori fixed increment amount.

1. Initial the price p to 0.
2. The initial set S_0 of "active bidders" is all bidders.
3. For $t = 1, 2, \ldots, ::$
   (a) Ask each bidder of S_{t-1} if it wants an item at the price $p + \epsilon$. Let S_t be the set of bidders that say "yes."
   (b) If $|S_t| \leq k$, then halt. Sell an item to each bidder of S_t at the price p. If there are items leftover (i.e., $k - |S_t| > 0$), sell them to an arbitrary subset of the bidders of $S_{t-1} \setminus S_t$ at the price p.
   (c) Otherwise, increment p by ϵ.

The English auction stops as soon as the demand — the number of bidders who want an item at the current price — becomes at most the supply. It is an indirect auction, in that it never explicitly asks a bidder for its valuation, just a series of "yes/no" queries at different prices.

There are many variants of this ascending auction. For example, the auction above does not allow a bidder to re-enter after dropping out. Such rules are common in, for example, simultaneous ascending auctions for selling wireless spectrum licenses. Art auctions, however, typically do allow bidders to implicitly drop out and re-enter. These details don't matter much in Scenario #1, but some of them do matter in more complex settings.

Why bother with ascending implementations? The k-Vickrey auction is already a compelling solution to welfare-maximization in Scenario #1, why do we need another? There are a number of reasons.

1. Ascending auctions are easier for bidders. It is generally easier to answer simple queries than to report a valuation. This point will become especially relevant in more complex scenarios.

2. Less information leakage. The winner of an ascending auction does not reveal its valuation, just the fact that it is at least the second-highest bid.

3. Transparency. The cause of a high selling price is generally more obvious in open ascending auctions than in sealed-bid auctions.

4. Potentially more seller revenue. For example, ascending auctions encourage ‘bidding wars.” There is also some supporting theoretical work on this point [1].

5. When there are multiple items, the opportunity for “price discovery.” A bidder has the opportunity for mid-course corrections and to better coordinate with other bidders. See Lecture #8 from CS364A for a detailed discussion of this point in the context of spectrum auctions.

Given all of these potential benefits of ascending auctions, and their consequent popularity in practice, it is important to have a theory of incentive-compatible iterative mechanism that parallels the theory we've developed for direct-revelation mechanisms. This is the subject of the next several lectures.

5 Analysis of the English Auction

We next argue that, in Scenario #1, the English auction enjoys the same incentive, performance, and tractability guarantees as the k-Vickrey auction.

First, we define the analog of truthful revelation. In an iterative auction, *sincere bidding* means that a player answers all queries honestly. In the English auction above, sincere bidding simply means that a bidder i answers “yes” to a query if and only if its valuation v_i is at least the proposed price $p + \epsilon$.[1]

Our first proposition states that the incentive guarantee of the English auction is essentially as good as that of the k-Vickrey auction.

Proposition 5.1 *In Scenario #1, in an English auction, sincere bidding is a dominant strategy for every bidder (up to ϵ).*

Proposition 5.1 states that, no matter what the bidders other than i do, no strategy by bidder i can increase its payoff by more than ϵ over sincere bidding. Thus, as $\epsilon \to 0$, sincere bidding approaches a dominant strategy.

The proof of Proposition 5.1 parallels the proof that the k-Vickrey auction is DSIC, and we leave it as an exercise. The idea is that a deviation by a bidder either involves exiting early or remaining active for extra iterations. These deviations correspond to underbidding and overbidding, respectively, in a k-Vickrey auction. The former is never profitable and the latter never increase the bidder's payoff by more than ϵ.

[1] For clarity, we reserve the word “truthful” for direct-revelation mechanisms, and use “sincere” instead for iterative mechanisms.

Proposition 5.1 may seem trivial now, but we'll appreciate it more shortly. A word of caution about iterative auctions: the action set of a player can be *much* richer than in a direct-revelation mechanism. In particular, bidder behavior can be a function of what happened earlier in the auction.[2] Proposition 5.1 shows that, at least in Scenario #1, this rich action set doesn't affect incentives in the English auction.

Our second proposition states that the performance guarantee of the English auction is essentially as good as that of the k-Vickrey auction.

Proposition 5.2 *In Scenario #1, if all bidders bid sincerely in an English auction, the welfare of the outcome is within $k\epsilon$ of the maximum possible.*

We again leave the proof as an exercise. The idea is that, with sincere bidding, the auction correctly identifies the bidders with the k highest valuations, up to a discretization error of ϵ. Note that the auction terminates within $v_{\max}/\epsilon$ iterations, where $v_{\max}$ is the highest valuation.

6 Bidders with Additive Valuations

6.1 The Setting

We now consider *non-identical* items. These arise naturally in many applications. The original motivating application for combinatorial auctions was take-off and landing slots at airports, which obviously vary in time and location. In the spectrum license auctions discussed last quarter, licenses vary in geography and frequency. Assembling a vacation package requires purchasing airfare, hotel rooms, and tour tickets. The lost goes on.

Scenario #2:

- A set U of m non-identical items.
- Each bidder i has a private valuation v_{ij} for each item j.
- Each bidder i has an *additive valuation*, meaning its value for a bundle $S \subseteq U$ of items is
$$v_i(S) := \sum_{j \in S} v_{ij}.$$

Scenario #2 is incomparable to scenario #1. Note that with additive valuations, bidders want as many items as possible. The additivity assumption means that a bidder's value for an item is independent of what other items it receives. Thus, there are no "substitutes" — redundancies between items, or "complements" — synergies between items. Additive valuations are seldom realistic, but they are a good segue into more complex settings.

[2]To specify the action precisely, we need to define exactly what information is available to bidders — the current price, the number of currently active bidders, the identities of the currently active bidders, etc. We'll be more formal about this as needed in forthcoming lectures.

The direct-revelation DSIC solution to scenario #2 is straightforward: just run a separate Vickrey auction for each of the m items. Additivity of valuations and payments implies that the incentive and performance guarantees of the Vickrey auction carry over. This mechanism is simple and runs in polynomial time.

6.2 Parallel English Auctions Are Not DSIC

The analogous ascending implementation is to run a separate English auction for each of the items — in parallel, say. That is, the auction maintains a set S_{tj} of active bidders for each item j in iteration t. Each iteration, for each item j on which it is active, a bidder is asked whether it would still want item j at a price ϵ higher than before. The jth constituent auction halts when there is only one active bidder remaining — if the final two or more bidders exit simultaneously, an arbitrary one of them is awarded the item at the last accepted price.

Warning: Sincere bidding is *not* a dominant strategy with parallel English auctions. That is, sincere bidding is not a best response with respect every set of actions that the other bidders might take. With parallel English auctions, the rich action space of iterative auctions rears its ugly head.

Example 6.1 To see the issue, consider two bidders and two items, with $v_{11} = 3$, $v_{12} = 2$, $v_{21} = 2$, and $v_{22} = 1$. If both bidders bid sincerely, the first bidder will win both items at prices of 2 and 1, respectively. Consider the following alternative action by the second bidder:

- If bidder 1 bids on item 1 in the first iteration, then keep bidding on both items forever (or up to a price of 3, say).[3]
- Otherwise, bid sincerely until the auction terminates.

If bidder 1 bids sincerely, then it certainly bids on item 1 in the first iteration (at a price of ϵ). This triggers the second bidder's threat of bidding forever, which causes bidder 1 to lose both items and receive utility 0. On the other hand, if bidder 1 abandons item 1 immediately, it at least wins the second item at a price of 1, for a utility of 1. Thus, for this action by bidder 2, sincere bidding is not a best response for bidder 1.

In almost all iterative auctions we'll discuss, similar examples show that sincere bidding is not a dominant strategy — Proposition 5.1 is quite unusual in this regard. Obviously, this does not mean that iterative auction are bad, or that sincere bidding is implausible. The example above merely illustrates that the rich auction space of iterative auctions necessitates a different incentive guarantee — a suitable analog of DSIC for iterative auctions. Informally, we want to formalize the following guarantee: *sincere bidding is always a best response provided other participants also bid sincerely.* This guarantee is attractive because when a bidder reasons about whether to bid sincerely or not, it does not need to other other bidders'

[3] Even if bidder 2 cannot directly observe whether or not bidder 1 bids for item 1 in the first iteration, it can infer this from what happens in the auction in the second iteration.

valuations — just that, whatever those valuations might be, the other bidders are bidding sincerely with respect to them. The example above does not rule out such a guarantee, since the second player's behavior does not correspond to sincere bidding with respect to any valuation.

6.3 EPIC Mechanisms

We next define a new equilibrium concept, weaker than a dominant-strategy equilibrium. The standard Nash equilibrium concept is not appropriate because it is defined for full-information games — games where all players' preferences are common knowledge. In most auction settings, a player has to reason about its action with no or incomplete information about other players' preferences (i.e., valuations).

Formally, consider n bidders with sets of possible private valuations $V_1, \ldots, V_n$. Let $A_1, \ldots, A_n$ be the sets of possible actions — again, in iterative auctions, actions can be history-dependent and so these sets are quite rich. A *strategy* s_i is a function from V_i to A_i — so a strategy specifies what a bidder does as a function of what it wants. For example, sincere bidding in an iterative auction is a strategy — not an action, but a function from valuations to actions — since a bidder's honest answer to a query (e.g., "do you want an item at price p?") depends on its valuation. The range of a sincere bidding strategy comprises only particularly simple, history-independent actions.

A strategy profile $(s_1, \ldots, s_n)$ is an *ex post Nash equilibrium (EPNE)* if, for every bidder i and valuation $v_i \in V_i$, the action $s_i(v_i)$ is a best-response to every action profile $\mathbf{s}_{-i}(\mathbf{v}_{-i})$ with $\mathbf{v}_{-i} \in V_{-i}$. That is, bidder i is confident that $s_i(v_i)$ is a best response knowing *only* that the other bidders use the strategies $\mathbf{s}_{-i}$, and *without* knowing their actual valuations. In the context of iterative auctions, "all players bid sincerely" corresponds to a strategy profile, so one can discuss whether or not it is an EPNE.

Stronger than an EPNE is a *dominant-strategy equilibrium (DSE)*, where for every bidder i and valuation v_i, the action $s_i(v_i)$ is a best response to every action profile $\mathbf{a}_{-i}$ of A_{-i}, *whether of the form* $\mathbf{s}_{-i}(\mathbf{v}_{-i})$ *or not.* Thus, in a DSE, each bidder is confident that its strategy is a best response without knowing anything at all about others' actions. For example, Proposition 5.1 shows that, in scenario #1, sincere bidding by all bidders is a DSE (up to ϵ).

Example 6.1 shows that, in scenario #2, sincere bidding in parallel English auctions is not a DSE. It is, however, an EPNE.

Proposition 6.2 *In Scenario #2, in parallel English auctions, sincere bidding by all bidders is an ex post Nash equilibrium (up to $m\epsilon$).*

We leave the proof as an exercise. The idea is that, when bidders have additive valuations and bid sincerely, the different English auctions can be analyzed separately. Proposition 5.1 then carries over to each of the m single-item auctions.

We'll call a mechanism *ex post incentive compatible (EPIC)* if sincere bidding is an ex post Nash equilibrium in which all bidders always receive nonnegative utility. Proposition 6.2 implies that, in scenario #2, parallel English auctions is an EPIC mechanism.

7 Where We're Headed

Our goal in the next several lectures is to identify scenarios where we can achieve best-possible incentive, performance, and tractability guarantees via ascending auctions:

1. **(Incentive guarantee.)** EPIC. Again, this means that sincere bidding is an ex post Nash equilibrium that guarantees all bidders nonnegative utility.

2. **(Performance guarantee.)** If all bidders bid sincerely, then the outcome of the auction maximizes the welfare.

3. **(Tractability guarantee.)** The auction should be "simple." At the very least, with sincere bidding, it should terminate in a reasonable number of iterations.[4]

After we identify tractable special cases where we can achieve all three of these goals, we'll study more complex settings where compromises are required, and will study the trade-offs between them.

8 Necessary Conditions for EPIC Welfare-Maximization

This section outlines some simple necessary conditions for an ascending auction to satisfy the incentive and performance guarantees of Section 7. These conditions will guide our auction designs in all of the scenarios to follow.

8.1 EPIC vs. DSIC Implementations

Because every DSE is an EPNE, every DSIC mechanism is EPIC. We've now seen a natural ascending auction that is EPIC but not DSIC. In a direct-revelation mechanism, however, the two concepts coincide. The reason is that, in such a mechanism, every available action is consistent with the truthful revelation of a possible private valuation — every bid might well be the bidder's actual valuation. (Cf., Example 6.1.) It follows that truthful revelation in a direct-revelation mechanism is a DSE if and only if it is EPNE.

We emphasize that the EPIC guarantee, while technically weaker than the DSIC guarantee we're used to, is still very strong. It asserts that sincere bidding a universal (over $\mathbf{v}_{-i}$) best response for bidder i, assuming only that other bidders bid sincerely. Lots of auctions don't have this property — in a first-price auction, for example, a bidder is not content knowing only other bidders' strategies (e.g., knowing that its opponent always shades its bid by 20%) — also knowing others' valuations would be extremely useful in formulating a bid. We'll adopt as the EPIC condition as the strongest incentive guarantee that we can generally hope for in an iterative auction.

[4] At the very least, the running time should be pseudopolynomial, meaning polynomial when all of the numbers provided in the input are polynomially bounded. For example, an English auction for a single item, as we have described it, requires a pseudopolynomial number of iterations to terminate: $v^{(2)}/\epsilon$ iterations, where $v^{(2)}$ is the second-highest valuation.

Designing EPIC iterative mechanisms is only harder than designing DSIC direct-revelation mechanisms, in the following sense. Given an EPIC iterative mechanism M, we can apply the Revelation Principle (see Lecture #4 of CS364A) to it. For example, applying the Revelation Principle to English auction in scenario #1 yields the k-Vickrey auction. This produces an equivalent mechanism — meaning the outcome of direct revelation in M' is the outcome of sincere bidding in M — that is EPIC (see the Exercises). Since M' is a direct-revelation mechanism, the observation above implies that it is in fact DSIC. Thus, a DSIC direct-revelation mechanism with good properties (like high welfare) is a logical prerequisite to an EPIC iterative mechanism with the same good properties. Designing the former is a useful "sanity check" before trying to design the latter.

8.2 Uniqueness of Payments and the VCG Mechanism

Recall the *VCG mechanism* from last quarter. This is a direct revelation mechanism that can be defined in very general mechanism design settings. The allocation rule is to choose the outcome that maximizes welfare with respect to the reported bids. The payment of a bidder i is its "externality" — the welfare other to others caused by i's participation in the mechanism.

The point of this section is that *a necessary condition for an ascending auction to be EPIC and welfare-maximizing is that sincere bidding always yields the VCG allocation and payments.* This fact follows from two observations. The first is that, as above, applying the Revelation Principle to an EPIC welfare-maximizing mechanism yields a direct-revelation DSIC welfare-maximizing mechanism, with the same mapping from valuations to allocations and payments. The second, proved below, is that the VCG mechanism is the unique DSIC welfare-maximizing mechanism (up to an additive constant).

More generally, consider an allocation rule $\mathbf{x}$ that has a finite range A.[5] We regard a valuation v_i as a vector indexed by A. We assume that each set $V_i \subseteq \mathcal{R}^A$ of possible valuations for bidder i is connected.[6] We prove that, for each i and $\mathbf{v}_{-i}$, all payment functions $p_i(v_i, \mathbf{v}_{-i})$ that satisfy the DSIC condition are the same, up to an additive constant. This implies that the VCG mechanism is the only DSIC welfare-maximizing mechanism in which $p_i(0, \mathbf{v}_{-i})$ is always 0.

Fix i and $\mathbf{v}_{-i}$. Write $x(\cdot)$ and $p(\cdot)$ for $x_i(\cdot, \mathbf{v}_{-i})$ and $p_i(\cdot, \mathbf{v}_{-i})$, respectively. Let A denote the (finite) range of x and suppose x and p satisfy the DSIC condition. We immediately have that $p(v) = p(v')$ whenever $x(v) = x(v')$ — if $p(v) < p(v')$, for example, a bidder with true valuation v' would have an incentive to misreport v. For $a \in A$, it is therefore well defined to write p_a for the payment $p(v)$ made under every declaration v with $x(v) = a$. We can assume that $|A| \geq 2$.

Call two outcomes $a, b \in A$ *close* if for every $\epsilon > 0$ there are valuations v_a, v_b with $\|v_a - v_b\|_\infty < \epsilon$, $x(v_a) = a$, and $x(v_b) = b$. We can use the DSIC condition to pin down

[5] More general results are possible, but this version is good enough for our purposes.

[6] This assumption holds in our applications, and is mathematically necessary. For example, in a single-item auction with only integer valuations, the second-price rule is not unique: marking up the sale price by $\frac{1}{2}$ whenever there is a unique highest bid gives an alternative DSIC auction.

$p_b - p_a$ as follows:

$$\underbrace{v_a(a) - p_a}_{\text{truthful report } v_a} \geq \underbrace{v_a(b) - p_b}_{\text{false report } v_b}$$

and

$$\underbrace{v_b(b) - p_b}_{\text{truthful report } v_b} \geq \underbrace{v_b(a) - p_a}_{\text{false report } v_a},$$

implying that

$$p_b - p_a \in \underbrace{[v_i(b) - v_i(a), v_i'(b) - v_i'(a)]}_{\leq v_i(b) - v_i(a) + 2\epsilon}.$$

Since a and b are close, we can take $\epsilon \downarrow 0$, implying that there is at most one possible value for $p_b - p_a$.

For an arbitrary pair a, b of outcomes, connectedness of V_i implies that we can find a sequence $a = a_0, a_1, a_2, \ldots, a_{k-1}, a_k = b$ such that a_j, a_{j+1} are close for every j.[7] By transitivity, there is at most one possible value for $p_b - p_a$ for each $a, b \in A$. Thus, all payment functions $p(\cdot)$ that satisfy the DSIC condition with $x(\cdot)$ differ only by an additive constant.

References

[1] P. R. Milgrom and R. J. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50(5):1089–1122, 1982.

[7]Formally, let A' be the outcomes reachable from a by sequences of close pairs. Let X and Y denotes the valuations vV_i such that $x(v) \in A'$ and $x(v) \notin A'$, respectively. If $A' \neq A$, then X and Y are a partition of V_i into two sets with strictly positive Hausdorff distance, contradicting the connectedness of V_i.