CSC304: Algorithmic Game Theory and Mechanism Design
Fall 2016

Allan Borodin (instructor)
Tyrone Strangway and Young Wu (TAs)

October 5, 2016

Lecture 9

- Announcements
  - Tutorial this Friday.
  - Assignment due at start of this weeks turtorial. Submission link:. https://markus.teach.cs.toronto.edu/csc304-2016-09/en/main log in with your CDF credentials.
  - Slides for Lecture 7 have been posted. There are no slides for Lecture 8, the guest lecture by Tyrone Strangway. His lecture was based on a paper "Budgetary Effects on Pricing Equilibrium in Online Markets" that will be posted.
  - Next assignment will be due October 28. I plan to post the first couple of questions by this weekend.
  - Next week, Monday (October 10) is Thanksgiving and hence the University is closed and there is no lecture. I also am cancelling the Wednesday lecture. There will be a lecture or tutorial on Friday.
- Today's agenda

  Start new topic: Mechanism design (Part II in the KP text)
  - What is mechanism design and the various areas of mechanism design
  - Mechanism design with money; auctions

What is mechanism design and social choice theory?

In our first lecture, we made the following comments:
In game theory, we are given a "game" and the goal is to understand and analyze the strategies of agents and deduce the resulting outcome(s).

What is mechanism design and social choice theory?

In our first lecture, we made the following comments:
In game theory, we are given a "game" and the goal is to understand and analyze the strategies of agents and deduce the resulting outcome(s).

The raison d'etre of mechanism design is to design mechanisms (e.g. algorithms to allocate and price sets of items) so as to induce games with desireable outcomes. It is therefore sometimes called *inverse game theory*.

What is mechanism design and social choice theory?

In our first lecture, we made the following comments:
In game theory, we are given a "game" and the goal is to understand and analyze the strategies of agents and deduce the resulting outcome(s).

The raison d'etre of mechanism design is to design mechanisms (e.g. algorithms to allocate and price sets of items) so as to induce games with desireable outcomes. It is therefore sometimes called *inverse game theory*.

The agents in game theory and mechanism design traditionally have cardinal (e.g. monetary values) utilities wheres in **social choice theory** (e.g. voting, peer evaluation), agents usually have preferences. With that possible distinction in mind, we can view social choice theory as part of game theory/mechanism design.

More on the nature of mechanism design

In mechanism design, one designs games where preferences (or payoffs) are unknown (i.e. private information), so when players act rationally (in equilibrium) "desirable outcomes" emerge.

More on the nature of mechanism design

In mechanism design, one designs games where preferences (or payoffs) are unknown (i.e. private information), so when players act rationally (in equilibrium) "desirable outcomes" emerge.

Mechanism design is therefore a type of algorithm design where private inputs are coming from self interested strategic agents.

More on the nature of mechanism design

In mechanism design, one designs games where preferences (or payoffs) are unknown (i.e. private information), so when players act rationally (in equilibrium) "desirable outcomes" emerge.

Mechanism design is therefore a type of algorithm design where private inputs are coming from self interested strategic agents.

Traditionally, mechanism design involves money (to some extent) to incentivize agents to choose strategies leading to a desireable outcome.

More on the nature of mechanism design

In mechanism design, one designs games where preferences (or payoffs) are unknown (i.e. private information), so when players act rationally (in equilibrium) "desirable outcomes" emerge.

Mechanism design is therefore a type of algorithm design where private inputs are coming from self interested strategic agents.

Traditionally, mechanism design involves money (to some extent) to incentivize agents to choose strategies leading to a desireable outcome.

There is also an area of *mechanism design without money*. This includes stable matching (chapter 11), fair division (chapter 12) , and social choice theory (chapter 13).

More on the nature of mechanism design

In mechanism design, one designs games where preferences (or payoffs) are unknown (i.e. private information), so when players act rationally (in equilibrium) "desirable outcomes" emerge.

Mechanism design is therefore a type of algorithm design where private inputs are coming from self interested strategic agents.

Traditionally, mechanism design involves money (to some extent) to incentivize agents to choose strategies leading to a desireable outcome.

There is also an area of *mechanism design without money*. This includes stable matching (chapter 11), fair division (chapter 12) , and social choice theory (chapter 13). While the text does these chapters first, I think it is more natural to transition to auctions, arguably the prime example of *mechanism design with money*.

Auctions

The text devotes 4 chapters (14-17) to various types of auctions and the relevant issues. A very general setting (and one well beyond what we will consider) consists of the following ingrediants:

- A set (or multiset) M of items to be sold.
- A set U of buyers having valuations for various sets (or multisets) of items and possibly a budget.
- A number of sellers having costs for producing items.
- The outcome of an auction mechanism is a "feasible allocation" of the items to the buyers so as to achieve certain desired goals. That is, there are contraints on what allocations are feasible.

Before we consider specific types of mechanisms, the text uses a more abstract and general formulation where we do not mention items but rather have the concept of *feasible outcomes* where agents have valuations for outcomes. (**Our formulation precludes** *externalities.*)

Auctions and mechanism design

If all information is publicly known, then this is a pure combinatorial problem that requires an algorithm to compute an optimal or near optimal desireable allocation. There are no strategic considerations.

Auctions and mechanism design

If all information is publicly known, then this is a pure combinatorial problem that requires an algorithm to compute an optimal or near optimal desireable allocation. There are no strategic considerations.

But in mechanism design we assume that some information (e.g. the valuations and/or budgets of the buyers, the costs of the sellers) is private.

Auctions and mechanism design

If all information is publicly known, then this is a pure combinatorial problem that requires an algorithm to compute an optimal or near optimal desireable allocation. There are no strategic considerations.

But in mechanism design we assume that some information (e.g. the valuations and/or budgets of the buyers, the costs of the sellers) is private.

Because agents (e.g. buyers or sellers) are strategic they may not be truthful in reporting their private information. A mechanism needs to incentivize strategic behaviour to achieve a desireable outcome given only partial or no prior knowledge of the private information.

Auctions and mechanism design

If all information is publicly known, then this is a pure combinatorial problem that requires an algorithm to compute an optimal or near optimal desireable allocation. There are no strategic considerations.

But in mechanism design we assume that some information (e.g. the valuations and/or budgets of the buyers, the costs of the sellers) is private.

Because agents (e.g. buyers or sellers) are strategic they may not be truthful in reporting their private information. A mechanism needs to incentivize strategic behaviour to achieve a desireable outcome given only partial or no prior knowledge of the private information.

In Bayesian mechanism design, the mechanism (and perhaps the agents) have prior common distributional knowledge about the private information.

Auctions and mechanism design

If all information is publicly known, then this is a pure combinatorial problem that requires an algorithm to compute an optimal or near optimal desireable allocation. There are no strategic considerations.

But in mechanism design we assume that some information (e.g. the valuations and/or budgets of the buyers, the costs of the sellers) is private.

Because agents (e.g. buyers or sellers) are strategic they may not be truthful in reporting their private information. A mechanism needs to incentivize strategic behaviour to achieve a desireable outcome given only partial or no prior knowledge of the private information.

In Bayesian mechanism design, the mechanism (and perhaps the agents) have prior common distributional knowledge about the private information.

The objective of a mechanism may be to try to optimize social welfare (also called social surplus as in the text) or the mechanism may itself be viewed as an agent trying to maximize its revenue.

Combinatorial auctions: an example of a mechanism design problem

As we mentioned, the suggested framework is very general and one usually considers more restrictive settings. Lets consider one reasonably general setting that is perhaps the most studied in theoretical computer science.

The combinatorial auction (CA) problem

In the CA problem, there is a set M of m items, a set U of n buyers, and one seller which we can also view as the Mechanism. Each agent i has a private valuation function $v_i : 2^M \rightarrow \mathbb{R}^{\geq 0}$. Each agent will submit bids $b_i(S) \geq 0$ for the subsets S it desires. The mechanism will allocate a desired subset S_i (possibly the empty subset) to each agent and will charge a price $p_i(S_i) \leq b_i(S_i)$ for the set allocated. The *quasi linear utility* of agent i for this allocation is $u_i(S_i) = v_i(S_i) - p_i(S_i)$. A feasible allocation is a collection of disjoint subsets S_i.

The CA problem continued

What is the goal or utility of the seller/mechanism?

The CA problem continued

What is the goal or utility of the seller/mechanism?

One possible goal is social welfare/surplus; namely, to obtain a feasible allocation that maximizes the values of the allocated sets.
Note: This is equivalent to maximizing the utilties of the allocated sets + the revenue collected by the seller/mechanism

The CA problem continued

What is the goal or utility of the seller/mechanism?

One possible goal is social welfare/surplus; namely, to obtain a feasible allocation that maximizes the values of the allocated sets.
Note: This is equivalent to maximizing the utilties of the allocated sets + the revenue collected by the seller/mechanism

A specific example of a CA is the spectrum auction that is discussed in Chapter 16 of the KP text. Here we envision the government (i.e. the mechanism) is allocating licenses for various collections of spectrum frequencies and it is not unreasonable to assume that the goal of the mechanism is social welfare.

CA problem continued

In its generality, each agent has a value for each possible subset of items. This requires an exponential (in m) input representation. To algorithmically accomodate (possibly exponential) size set system problems, one can assume some sort of "oracle" such as a value oracle that given S will return a value, say $b_i(S)$.

CA problem continued

In its generality, each agent has a value for each possible subset of items. This requires an exponential (in m) input representation. To algorithmically accomodate (possibly exponential) size set system problems, one can assume some sort of "oracle" such as a value oracle that given S will return a value, say $b_i(S)$.

In practice, many CA problems can be represented explicitly and succinctly. For example, if each agent is interested in only one set, then this just requires specifying the set and the bid associated with that set. This is referred to as a *single-minded CA*. As long, as each agent is only interested in a few sets, the CA problem can be represented explicitly and succinctly.

CA problem continued

In its generality, each agent has a value for each possible subset of items. This requires an exponential (in m) input representation. To algorithmically accomodate (possibly exponential) size set system problems, one can assume some sort of "oracle" such as a value oracle that given S will return a value, say $b_i(S)$.

In practice, many CA problems can be represented explicitly and succinctly. For example, if each agent is interested in only one set, then this just requires specifying the set and the bid associated with that set. This is referred to as a *single-minded CA*. As long, as each agent is only interested in a few sets, the CA problem can be represented explicitly and succinctly.

Another explictly represented CA is the sCA problem where every desired set has cardinality at most s (for some small constant s).

CA problem continued

In its generality, each agent has a value for each possible subset of items. This requires an exponential (in m) input representation. To algorithmically accomodate (possibly exponential) size set system problems, one can assume some sort of "oracle" such as a value oracle that given S will return a value, say $b_i(S)$.

In practice, many CA problems can be represented explicitly and succinctly. For example, if each agent is interested in only one set, then this just requires specifying the set and the bid associated with that set. This is referred to as a *single-minded CA*. As long, as each agent is only interested in a few sets, the CA problem can be represented explicitly and succinctly.

Another explictly represented CA is the sCA problem where every desired set has cardinality at most s (for some small constant s).

Note: We assume that $v_i(\varnothing) = 0$ and that $v_i(S)$ is monotone for every i (i.e. *free disposal*). Thus, sets that are not explicitly given values inherit their value from a desired subset of largest value.

The CA problem continued and what is *algorithmic* mechanism design

What is the algorithmic challenge of this problem?

The CA problem continued and what is *algorithmic* mechanism design

What is the algorithmic challenge of this problem?

The underlying allocation problem is the *set packing problem*; namely, given a collection of sets $\mathcal{S} = \{S_1, S_2, \ldots, S_t\}$, where each set S_j has a value v_j, choose a subcollection $\mathcal{S}'$ of disjoint sets so as to maximize $\sum_{\{j:S_j \in \mathcal{S}'\}} v_j$.

The CA problem continued and what is *algorithmic* mechanism design

What is the algorithmic challenge of this problem?

The underlying allocation problem is the *set packing problem*; namely, given a collection of sets $\mathcal{S} = \{S_1, S_2, \ldots, S_t\}$, where each set S_j has a value v_j, choose a subcollection $\mathcal{S}'$ of disjoint sets so as to maximize $\sum_{\{j : S_j \in \mathcal{S}'\}} v_j$.

This problem is NP-hard and NP-hard to approximatte to within a factor $m^{\frac{1}{2}-\epsilon}$, even when all $v_j = 1$ (i.e. the unweighted case).

The CA problem continued and what is *algorithmic* mechanism design

What is the algorithmic challenge of this problem?

The underlying allocation problem is the *set packing problem*; namely, given a collection of sets $\mathcal{S} = \{S_1, S_2, \ldots, S_t\}$, where each set S_j has a value v_j, choose a subcollection $\mathcal{S}'$ of disjoint sets so as to maximize $\sum_{\{j : S_j \in \mathcal{S}'\}} v_j$.

This problem is NP-hard and NP-hard to approximatte to within a factor $m^{\frac{1}{2}-\epsilon}$, even when all $v_j = 1$ (i.e. the unweighted case).

Furthermore, the set packing problem is NP hard even when all sets have cartdinality at most s (i.e. the underlying allocation problem for the sCA problem) for $s \geq 3$ and hard to approximate to a factor $\Omega(\frac{s}{\ln s})$.

Strategic behaviour meets computational complexity

We will soon discuss *truthfulness* where bidding truthfully is a dominant strategy. A truthful mechanism (also referred to as an *incentive compatible IC or dominant strategy incentive compatible DSIC mechanism*) is one that results in truthful bidding being a dominant strategy. And now here is the issue that began algorithmic game theory. (See the seminal papers by Nisan and Ronen, and by Lehmann, O'Callahan and Shoham.)

If we could compute an optimal allocation (i.e. one optimizing social welfare), we could then rely upon Vickrey-Clarke-Groves (VCG) pricing (to be dsicussed) to obtain a truthful mechanism.

Strategic behaviour meets computational complexity

We will soon discuss *truthfulness* where bidding truthfully is a dominant strategy. A truthful mechanism (also referred to as an *incentive compatible IC or dominant strategy incentive compatible DSIC mechanism*) is one that results in truthful bidding being a dominant strategy. And now here is the issue that began algorithmic game theory. (See the seminal papers by Nisan and Ronen, and by Lehmann, O'Callahan and Shoham.)

If we could compute an optimal allocation (i.e. one optimizing social welfare), we could then rely upon Vickrey-Clarke-Groves (VCG) pricing (to be dsicussed) to obtain a truthful mechanism.

But the NP-hardness of this problem precludes an optimal allocation (at least in the worst case, if we assume P $\neq$ NP) and hence we would have to rely on some approximation algorithm.

Strategic behaviour meets computational complexity

We will soon discuss *truthfulness* where bidding truthfully is a dominant strategy. A truthful mechanism (also referred to as an *incentive compatible IC or dominant strategy incentive compatible DSIC mechanism*) is one that results in truthful bidding being a dominant strategy. And now here is the issue that began algorithmic game theory. (See the seminal papers by Nisan and Ronen, and by Lehmann, O'Callahan and Shoham.)

If we could compute an optimal allocation (i.e. one optimizing social welfare), we could then rely upon Vickrey-Clarke-Groves (VCG) pricing (to be dsicussed) to obtain a truthful mechanism.

But the NP-hardness of this problem precludes an optimal allocation (at least in the worst case, if we assume P $\neq$ NP) and hence we would have to rely on some approximation algorithm.

But VCG pricing does not always result in a truthful mechanism for approximation algorithms!
VCG mechanism = optimal allocation + VCG pricing.

The discussion on auctions in the KP text

Various aspects of auctions are discussed in the KP text and we will be presenting material from each of the relevant chapters (although not in the same order). Here is a brief summary of the relevant chapters.

- Chapter 14 introduces auctions beginning with the special (but still very interesting) case of one item and one seller where the buyer's valuation is drawn from a known distribution (i.e. a Bayesian setting). The goal here is maximize the expected revenue of the seller/mechanism. The Vickrey second price auction is shown to be truthful. Myerson's optimal auction for one buyer is presented.
- Chapter 15 introduces the VCG mechanism with a focus on *single parameter settings* where the agents provide a single bid. The main application in Chapter 15 is a sponsored search auction.
- An application of the VCG mechanism for a multi-parameter problem is given in Chapter 16. The chapter also discusses scoring rules.
- Finally, Chapter 17 discusses matching markets where we have buyers and sellers and each buyer is a *unit demand buyer* who only one item (from some set of possibilities) and sellers who have one item.

The Vickrey auction for a single item

In a sealed auction for a single item, the auctioneer (the mechanism, the seller) receives bids $(b_1, \ldots, b_n)$ from (say) n bidders who have private values $(v_1, \ldots, v_n)$ for the item. Notably, the bids may not be equal to the values and the mechanism may not know anything about the true values (or it might know a prior distribution on these values).

The Vickrey auction for a single item

In a sealed auction for a single item, the auctioneer (the mechanism, the seller) receives bids $(b_1, \ldots, b_n)$ from (say) n bidders who have private values $(v_1, \ldots, v_n)$ for the item. Notably, the bids may not be equal to the values and the mechanism may not know anything about the true values (or it might know a prior distribution on these values).

Based on these bids, the mechanism allocates the item (i.e. determines the winner) and sets a price for the winner to pay. If the mechanism only wishes to maximize social welfare, what should it do?

The Vickrey auction for a single item

In a sealed auction for a single item, the auctioneer (the mechanism, the seller) receives bids $(b_1, \ldots, b_n)$ from (say) n bidders who have private values $(v_1, \ldots, v_n)$ for the item. Notably, the bids may not be equal to the values and the mechanism may not know anything about the true values (or it might know a prior distribution on these values).

Based on these bids, the mechanism allocates the item (i.e. determines the winner) and sets a price for the winner to pay. If the mechanism only wishes to maximize social welfare, what should it do?

For example, suppose I am interested in selling my legally unlocked 8GB iphone 4. I am assuming you know what it is worth to you! I will announce my allocation and pricing algorithm and you will then bid. My goal is to make sure that the person who values it the most will be the winner of my auction. (I may also believe that this person will bid reasonably and I will get some revenue.) What should I do?

The Vickrey auction continued: selling my ipone 4

1. I will allocate the phone to the person with the highest bid (ignore tying bids) and charge that person their bid. This is the *first price auction*.

The Vickrey auction continued: selling my ipone 4

1. I will allocate the phone to the person with the highest bid (ignore tying bids) and charge that person their bid. This is the *first price auction*.

2. I will allocate the phone to the person with the highest bid (ignore tying bids) and charge that person the second highest bid. This is the Vickrey second price auction.

The Vickrey auction continued: selling my ipone 4

1. I will allocate the phone to the person with the highest bid (ignore tying bids) and charge that person their bid. This is the *first price auction*.

2. I will allocate the phone to the person with the highest bid (ignore tying bids) and charge that person the second highest bid. This is the Vickrey second price auction.

How many people bid their true value for the first mechanism? If not, what fraction of the true value did you bid?

The Vickrey auction continued: selling my ipone 4

1. I will allocate the phone to the person with the highest bid (ignore tying bids) and charge that person their bid. This is the *first price auction*.

2. I will allocate the phone to the person with the highest bid (ignore tying bids) and charge that person the second highest bid. This is the Vickrey second price auction.

How many people bid their true value for the first mechanism? If not, what fraction of the true value did you bid?

How many people bid their true value for the second mechanism? If not what fraction of the true value did you bid?

The second price auction is truthful

The Vickrey auction is truthful

Bidding truthfully is a dominant strategy (no matter how the other buyers bid) and reasonably assuming everyone bids truthfully, social welfare is maximized.

There is an obvious issue with both the first and second price auctions in that the auctioneer/seller may value the item more than any of the buyers.

The second price auction is truthful

The Vickrey auction is truthful

Bidding truthfully is a dominant strategy (no matter how the other buyers bid) and reasonably assuming everyone bids truthfully, social welfare is maximized.

There is an obvious issue with both the first and second price auctions in that the auctioneer/seller may value the item more than any of the buyers.

This is dealt with by the mechanism announcing a *reserve price* so that no bid will be accepted that is under the reserve price.