CSC304: Algorithmic Game Theory and Mechanism Design
Fall 2016

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

November 2, 2016

Lecture 15

- Announcements
  - Office hours: Tuesdays 3:30-4:30 SF 2303B; or schedule meeting; or drop by. But next week is the fall break so check first.
  - Term test will take place on Friday, November 4. Aids: One sheet of handwritten notes, both sides.
  - Nature and scope of term test:
    - ★ The test is long but many parts can be answered quickly.
    - ★ Any topic relating to questions that were asked in the first two assignments. Basic game theory, pure and mixed NE, network congestion and Braess paradox, Bayesian mechanism design; Myerson auction.
    - ★ Material covered in lectures 12,13, 14 and start of today on mechanisms VCG and GSP for sponsored search.
  - No class on Monday, November 7 (fall break)
- Todays agenda
  - Quick review of GSP for sponsored search
  - The public project problem continued.
  - Example of VCG for a combinatorial auction
  - General comments about truthful mechanism design (not on test)

Review: GSP equlibria with more and less revenue

For some instances of sponsored search, there can be a NE that achieves more revenue and another NE achieving less revenue than VCG.

Review: GSP equlibria with more and less revenue

For some instances of sponsored search, there can be a NE that achieves more revenue and another NE achieving less revenue than VCG.

Consider the instance $v_1 = 7, v_2 = 6$ and $v_3 = 1$, with 2 slots having CTRs $c_1 = 10, c_2 = 4$. (I am making the CTRs integers for convenience.)

Review: GSP equlibria with more and less revenue

For some instances of sponsored search, there can be a NE that achieves more revenue and another NE achieving less revenue than VCG.

Consider the instance $v_1 = 7, v_2 = 6$ and $v_3 = 1$, with 2 slots having CTRs $c_1 = 10, c_2 = 4$. (I am making the CTRs integers for convenience.)

1. With bidder 1 being present, the expected social welfare of the remaining two bidders is $c_2 \cdot 6 + 0 = 24$.

   If bidder 1 were not present, then the the expected social welfare for the two remaining bidders would be $c_1 \cdot 6 + c_2 \cdot 1 = 64$. So bidder 1 has caused them a loss of social welfare of $64 - 24 = 40$ and that is the expected revenue generated from bidder 1.
2. With bidder 2 present, the expected social welfare for bidders 1,3 is $c_1 \cdot 7 + 0 = 70$

   If bidder 2 were not present then the social welfare for the two remaining bidders would be $c_1 \cdot 7 + c_2 \cdot 1 = 74$. So bidder 2 has caused them an expected loss of social welfare of $74 - 70 = 4$.

Hence the total expected revenue generated by VCG is $40 + 4 = 44$

GSP with more and less that the expected VCG revenue

We have just seen that VCG generates expected revenue 44.
We now provide two other NE for this instance, one with more revenue and one with less.

- Bidding $(b_1, b_2, b_3) = (5, 4, 2)$ is a GSP equilibrium. It clearly achieves the same social optimal allocation but now with greater revenue $4 \cdot 10 + 2 \cdot 4 = 48$ than VCG.
- Bidding $(b_1, b_2, b_3) = (3, 5, 2)$ is a GSP equilibrium. It clearly does not achieve the same social optimal allocation and obtains less revenue $3 \cdot 10 + 2 \cdot 4 = 38$ than VCG.

GSP with more and less revenue at equilibria

What didn't we show last time?

GSP with more and less revenue at equilibria

What didn't we show last time?

While we computed the expected revenues, we still have to show that these bids $(5,4,2)$ and $(3,5,2)$ are equilibria. Lets just verify that $(3,5,2)$ is a NE. We will still refer to advertiser i as the agent with the highest valuation per click.

- Advertiser 1 (with $v_1 = 7$ per click) is bidding 3 and and will therefore, obtain expected utility $4(7-2) = 20$. Ignoring ties, nothing changes if the bid b_1 is modified so that $2 < b_1 < 5$. This advertiser could move back up to the first slot, say bidding 6 and then obtain utility $10(7-5) = 20$ so no advantage in doing that. The advertiser could bid below bidder 3, and then obtain no utility.
- Advertiser 2 (with $v_2 = 6$ per click) is bidding 5 and and will therefore, obtain expected utility $10(6-3) = 30$. This advertiser could move back to the second slot, say bidding 2.5 and then obtain utility $4(6-2) = 16$ so no advantage in doing that. The advertiser could bid below bidder 3, and then obtain no utility.

A complexity comment on the combinatorial public projects problem (CPPP) problem

A problem that has been formulated with the AGT community is the CPPP optimization problem. (Initially studied by Papadimiriou, Schapira, Singer [2008]).
As mentioned last lecture, in the CPPP, the government wants to undertake a set of k projects from a relatively large set P of potential projects. Each individual i ($1 \leq i \leq n$) has a valuation function $v_i : 2^P \to \mathbb{R}$ for each possible subset $S \subset P$ of projects. The mechanism needs to choose a subset S so as to maximize $\sum_i^n v_i(S)$.

Without any structure to the space of valuations, this would be difficult to optimize even if every agent was truthful.

One reasonable assumption is that the valuation function v_i of every agent is a *monotone submodular set function* and this is the CPPP problem that has been studied. (Tyrone spoke about submodular fucntions in his lecture but we will repeat the defintions.)

CPPP with submodular valuations

Monotone submodular set functions

Let $f : 2^P \to \mathbb{R}$ be a set function.

- f is monotone if $f(S) \leq f(T)$ for all $S \subseteq T$.
- f is submodular if $f(S) + f(T) \geq f(S \cup T) + f(S \cap T)$ for all subsets S and T.

Equivalently, f is submodular if it satisifies the following diminishing marginal gains property:

$f(S \cup \{x\}) - f(S) \geq f(T \cup \{x\}) - f(T)$ for all $S \subset T$.

Submodular functions and especially monotone submodular functions play an important role in many areas of optimization. In particular, for auctions, a monotone submodular function represents in some sense that some items are (partial) substitutes for eah other. For example, an iPAD would probably be worth more to me if I just have a cell phone than it would be worth to me if I have a cell phone and a lab top.

Submodular CPPP continued

The sum of submodular set functions is a submodular function. The underlying allocation problem of the submodular CPPP is to maximize a monotone submodular function subject to a cardinality constraint. This problem is hard to approximate to a factor better than $1 - \frac{1}{e}$. This approximation can be achieved by a simple greedy algorithm. What is the algorithm?

This hardness is in terms of needing exponentially many value oracle calls to beat this approximation ratio or to beat this approximation under a standard complexity assumption. (Of course, some specific submodular functions (e,g. modular or linear function) are easy to optimize.

Moreover, as a game theory auction problem, agents may not be truthful.

Submodular CPPP continued

The sum of submodular set functions is a submodular function. The underlying allocation problem of the submodular CPPP is to maximize a monotone submodular function subject to a cardinality constraint. This problem is hard to approximate to a factor better than $1 - \frac{1}{e}$. This approximation can be achieved by a simple greedy algorithm. What is the algorithm?

This hardness is in terms of needing exponentially many value oracle calls to beat this approximation ratio or to beat this approximation under a standard complexity assumption. (Of course, some specific submodular functions (e,g. modular or linear function) are easy to optimize.

Moreover, as a game theory auction problem, agents may not be truthful.

An important AGT result is that **no truthful mechanism can achieve an approximation ratio better than $\frac{1}{n^{\frac{1}{2}-\epsilon}}$ for any $\epsilon > 0$.**

Combinatorial auctions and VCG

Lets return to VCG as applied to CAs. Lets recall the general defnition of a CA. A general CA consists of :

- A set N of n agents, and a set M of m items. We can view the items as being distinct. (A multi-init CA allows for having multiple copies of some items.)
- A feasible allocation is one in which no item is give to more than one person.
- Each agent has a valuation function $v_i : 2^M \rightarrow \mathbb{R}\geq 0$. The utility of an agent is $v_i(S) - p_i(S)$ where $p_i(S)$ is the price that agent i pays if allocated set S
- We usually assume free disposal $v_i(S) \leq v_i(T)$ for $S \subseteq T$ and insure individual rationality (IR) by assuming $p_i(\emptyset) = 0$.

A mechanism consist of a feasible allocation algorithm and a pricing algorithm. Mechanisms can be deterministic or randomized.

An example of VCG applied to a simple CA

Suppose we have three agents going to Tim Horton's shop at the end of the day. They have left one cup of coffee c and one donut d. Three customers arrive with single minded declarations:
$v_1(c,d) = \$11, v_2(d) = \$10, v_3(c) = \$2$. (Obviously a very good donut.)

The salesperson wants to award the items so as to achieve maximum social welfare but of course dosn't know the valuations. Fortunately, the cash register is a VCG mechaanism cash register that takes bids for the items and then sets prices. Who gets what items and at what prices?

An example of VCG applied to a simple CA

Suppose we have three agents going to Tim Horton's shop at the end of the day. They have left one cup of coffee c and one donut d. Three customers arrive with single minded declarations:
$v_1(c, d) = \$11, v_2(d) = \$10, v_3(c) = \$2$. (Obviously a very good donut.)

The salesperson wants to award the items so as to achieve maximum social welfare but of course dosn't know the valuations. Fortunately, the cash register is a VCG mechaanism cash register that takes bids for the items and then sets prices. Who gets what items and at what prices?

The customers know that VCG is truthful and hence bid truthfully. VCG then optimizes the allocation by giving the donut to the second customer and the coffee to the third customer and the first customer will go away hungry and thirsty.

An example of VCG applied to a simple CA

Suppose we have three agents going to Tim Horton's shop at the end of the day. They have left one cup of coffee c and one donut d. Three customers arrive with single minded declarations:
$v_1(c, d) = \$11, v_2(d) = \$10, v_3(c) = \$2$. (Obviously a very good donut.)

The salesperson wants to award the items so as to achieve maximum social welfare but of course dosn't know the valuations. Fortunately, the cash register is a VCG mechaanism cash register that takes bids for the items and then sets prices. Who gets what items and at what prices?

The customers know that VCG is truthful and hence bid truthfully. VCG then optimizes the allocation by giving the donut to the second customer and the coffee to the third customer and the first customer will go away hungry and thirsty.

The VCG prices are $v_2 = \$8$ and $v_3 = \$0$. (There can be free coffee unless the coffee shop sets a reserve price.) What happens if the coffee shop has a minimum purchase policy of \$3?

Comments about truthful CAs

As we remarked when we began discussing auctions, the CA problem is an interesting AGT problem because in its generality the underlying allocation problem (i.e. set packing) is NP-hard to $\frac{1}{m^{\frac{1}{2}-\epsilon}}$ approximate. So *in the worst case setting*) we cannot use VCG.

For a single-minded auction we can use a greedy algorithm (say, sorting bids by highest bid first) to approximate any sCA when s is small.

Comments about truthful CAs

As we remarked when we began discussing auctions, the CA problem is an interesting AGT problem because in its generality the underlying allocation problem (i.e. set packing) is NP-hard to $\frac{1}{m^{\frac{1}{2}-\epsilon}}$ approximate. So *in the worst case setting*) we cannot use VCG.

For a single-minded auction we can use a greedy algorithm (say, sorting bids by highest bid first) to approximate any *sCA* when *s* is small.

But as we have indicated before, we cannot use such a simple greedy allocation when bidders are not single-minded (i.e. cannot obtain truthfulness in the worst case setting).

Comments about truthful CAs

As we remarked when we began discussing auctions, the CA problem is an interesting AGT problem because in its generality the underlying allocation problem (i.e. set packing) is NP-hard to $\frac{1}{m^{\frac{1}{2}-\epsilon}}$ approximate. So *in the worst case setting*) we cannot use VCG.

For a single-minded auction we can use a greedy algorithm (say, sorting bids by highest bid first) to approximate any *sCA* when *s* is small.

But as we have indicated before, we cannot use such a simple greedy allocation when bidders are not single-minded (i.e. cannot obtain truthfulness in the worst case setting).

Recently, it was shown that for submodular CAs, (where greedy can provide an approximation within a factor of 2 for the allocation problem), we cannot have a truthful mechanism that achieves better than a factor $\min(n, \sqrt{m})$ approximation. Dobzinski et al [2006] show how to achieve this approximation ratio for all CAs by a *universally truthful* randomized mechanism. It is not known if there is any deterministic truthful mechanism that achieves say an $\frac{1}{m^{1-\epsilon}}$ approximation for any $\epsilon > 0$.

More general comments about IC mechanism design

What do we know about computationally efficient, truthful (i.e. incentive compatible IC) mechanisms that can essentially achieve the best known approximations for the underlying allocation problem?. Here we are asking this question for auctions in general with quasi-linear utilities.

This question can be asked for the objectives of social welfare or revenue maximization. For now, we are still considering social welfare.

More general comments about IC mechanism design

What do we know about computationally efficient, truthful (i.e. incentive compatible IC) mechanisms that can essentially achieve the best known approximations for the underlying allocation problem?. Here we are asking this question for auctions in general with quasi-linear utilities.

This question can be asked for the objectives of social welfare or revenue maximization. For now, we are still considering social welfare.

There are many interesting *special cases of auctions* where one can achieve good approximations by IC auctions.

There are also some more general approaches that consider *special classes of algorithms* and show how they can lead to IC mechanisms with good approximations.

IC mechanism design comments continued

These results consider restricted classes of allocation algorithms and convert such allocation algorithms for a given problem into an IC mechanism (i.e. allocation plus revenue) for the given problem while essentially preserving the approximation ratio.

An important early example (due to Lavi and Swamy [2005]) considers linear programming LP for arbitrary packing problems. If the LP has a "verifiable integrality gap" then the LP can be used to derive a *randomized truthful in expectation* mechanism while preserving the social welfare apporoximation.

For quite general problems, Briest et al [2005] and Dughmi and Roughgarden [2010] show how to provide *black box reductions* when given a PTAS algorithm for the underlying allocation.

Bayesian IC mechanisms

Although there are many important special cases of auction design problems and special cases of algorithmic paradigms, there is a basic incompatabiility between the desire for truthfulness and the computational necessity to sometimes settle for approximate algorithms.

Bayesian IC mechanisms

Although there are many important special cases of auction design problems and special cases of algorithmic paradigms, there is a basic incompatabiility between the desire for truthfulness and the computational necessity to sometimes settle for approximate algorithms.

Theoretical computer science is perhaps overly focused on worst case analysis. Economists are much more interested in the Bayesian framework.

Bayesian IC mechanisms

Although there are many important special cases of auction design problems and special cases of algorithmic paradigms, there is a basic incompatabiility between the desire for truthfulness and the computational necessity to sometimes settle for approximate algorithms.

Theoretical computer science is perhaps overly focused on worst case analysis. Economists are much more interested in the Bayesian framework.

The very good news is that we do not have to focus entirely on worst case results.

Results by Hartline and Lucier [2010] (for single parameter problems) and Hartline, Kleinberg and Malekian [2015] (for multi-paramter problems) show that for *any* allocation algorithm $\mathcal{A}$ for any auction problem, a Bayesian incentive compatible (BIC) mechanism for that problem can be obtained by a black box reduction to the underlying allocation algorithm $\mathcal{A}$. This reduction essentially preserves the approximation ratio of $\mathcal{A}$.

By BIC, we mean that the *expected utility* (with respect to the agent distributions) of an agent is maximized by bidding truthfully.