CS294 P29: Topics Algorithmic Game Theory November 8, 2011

Lecture 10: Mechanisms, Complexity, and Approximation

Lecturer: Christos Papadimitriou *Scribe: Faraz Tavakoli*

It is possible to characterize a *Mechanism Design (MD)* problem based on weather it's truthful, how fast it is, and closeness of the solution to the optimal. That is we are interested in *Incentive Compatibility (IC), Performance (P), and Approximation (APX)* respectively.

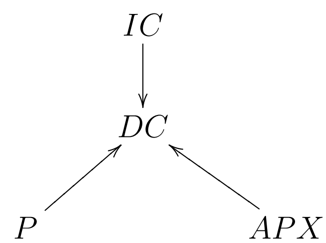

It is possible to make progress if we consider only two of the three; however, with all three we run into difficulties.

1 Combinatorial Auctions

In *Combinatorial Auctions* there are m indivisible objects and n bidders, each with a valuation $v_i : 2^m \to \mathbb{R}$. We are interested in finding an IC mechanism such that,

$$\text{Maximize: } \sum_i v_i(S_i)$$
$$\text{Subject to: } S_i \cap S_j = \emptyset \quad \forall ij$$

VCG can solve the problem; however, this accounts for $n+1$ NP-hard problems. On the other hand, approximation does not result in a truthful mechanism (not IC).

Theorem 1. *Any algorithm that approximates closer than $\frac{1}{2}$, for $n = 2$, must be exponential.*

2 Communication complexity

We want to compute $f(x, y) = \bigvee_{i=1}^{n} x_i \wedge y_i$ where x and y are the input bitstreams. How many bits must be communicated? The $min(|x|, |y|) = n$ bits are needed.

Theorem 2. *The above process requires $\sim n$ bits.*

Now, using the second theorem we can prove the first one.
Proof: [Proof of theorem 1] From Theorem 2, suppose valuations are of the form $v_i(S) = 0$ or 1. Furthermore, $v_i(S) = 0$ unless $|S| = \frac{m}{2}$. So, x is a bit vector of length $\binom{m}{m/2}$,

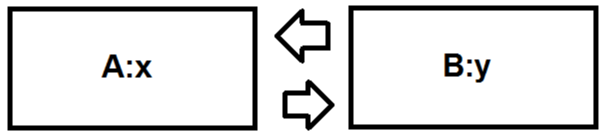

$(v_i = 1)$. Take y_S for some $|S| = \frac{m}{2}$ to be 1 if and only if $v_2([m] - S) = 1$. Then welfare 2 is achievable if and only if x and y intersect. So, $\sim \binom{m}{\frac{m}{2}}$ bits must be communicated which is $\approx \frac{2^m}{\sqrt{m}}$. □

Note that this means that all we can do is to auction $[m]$. Moreover, we have made the assumption that the bidders have exponential information in their possession and are able to communicate it. Instead, suppose they have some *succinct private data* from which they can compute $v_i(S)$ efficiently. Theorem 1 is not applicable to this situation.

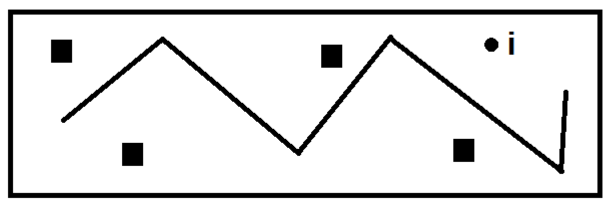

Example 1. *Overlay network routing heuristics, no triangle inequalities.*
Solution: Overlay k (fixed) nodes, route through them optimally.

Every node has a valuation $v_i(S)$ for every overlay set S where $v_i : m^{[k]} \to \mathbb{R}$ valuates how much routing improves by using S. Then, v_i is *submodular*

$$v(S \cup T) + v(S \cap T) \leq v(S) + v(T)$$

or

$$S \subset T \Rightarrow v(T + x) - v(T) \leq v(S + x) - v(S).$$

Combinatorial Public Project Problem (**CPPP**): mechanism for n submodular functions to optimize $\sum_i v_i(S)$.
No IC: greedy algorithm with $1 - \frac{1}{e} = 63\%$
No P: VCG
What if we want to satisfy both criteria, IC and P?

Theorem 3. *(Unless $P \approx NP$) No polynomial time mechanism approximates better than* $\frac{1}{\sqrt{m}}$

To prove the theorem we need the following three lemmas,

Lemma 1. *Affine maximizer*

Lemma 2. *Of exponential size*

Lemma 3. *Would 'solve' SAT*

For Lemma 1 (affine maximizer), recall from Roberts' Theorem, in the domain $U = \mathbb{R}^{|A|}$ the only IC mechanisms are of the form

$$\max_{a \in R \subset A} \sum v_i(a)$$

where R is the range. This lemma says the same is true for the U_{CPPP}. So, we assume we have a subset R of $[m]^{[k]}$ which gives a good (better than $\frac{1}{\sqrt{m}}$) approximation, for instance $p = m^{-1/2+\epsilon}$.

Then $R = \Omega(e^{m^\epsilon})$ sample $[m]$ with probability p, and let V be the sample. Consider $v(S) = |S \cap V|$ and $k = \sqrt{m}$.

Claim 1. *Unless* $|R| = \Omega(e^{m^\epsilon})$*, then it approximate* v *worse than* p*.*

What is the optimum $v(S)$? By the Central Limit Theorem (CLT) it is $\sim m^{1/2+\epsilon} = m \cdot p$ What is the optimum in R? $\Pr(|S \cap V| > m^{2\epsilon})$ very small so is $\Pr(\max_{S \in R} S \cap V > m^{\epsilon\epsilon}$ if R is not huge by union bound.

As the consequence of Lemma 2, CPPP is hard the same way as combinatorial auctions are, even if for one bidder.

We need to outline a reduction from SAT to mechanism design for CPPP with one bidder and succinct submodular valuation $v_{F,T}$ where $F = \{A_1, \cdots, A_m\}$ and $T \subset [m]$. Then,

$$v_{F,T}(S)\left(|S| = K = \sqrt{m}\right) = \left|\bigcup_{j \in S} A_j\right| - \frac{\epsilon}{|T|} \max\left(0, |S \cap T| - \frac{|T|}{2}\right)$$

Suppose we have a mechanism for the preceding problem. Then the following consequences hold, if it is IC $\Rightarrow$ it is an affine maximizer $\Rightarrow$ with at least e^{m^ϵ} sets in it.

Let A be any Boolean(0,1) $m \times n$ matrix, then there exist $c = \frac{\log m}{\log n}$ columns such that contain all 2^c bit vectors of length c.

Lemma 4 (Sauer-Shelah)**.** *For all* $z \subset 2^R$*, there is a* $Qq \subset R$*,* $|Q| \geq \frac{\log |Z|}{log |R|}$ *which is* ***shattered*** *by* Z*, that is,* $\forall Q' \subset Q, \exists S \subset Z$ *such that* $Q' = S \cap Q$*.*

Proof: Exercise. □

So, there is a subset $M \subset [m]$ of the items such that all subsets of M are *present* in A (the affine maximizer).

This is *not quite SAT* yet: given a family F of t sets subset of U, can $\frac{t}{2}$ of them be found such that their union covers the universe U? (NP-complete) $t = |M|$.

Now, define the valuation $v = v_{F,M}$ where elements outside M correspond to ϕ

Claim 2. *The mechanism gives optimal valuation* $|U|$ *if and only if there is* $\frac{t}{2}$ *subsets of* F *whose union is* U*.*

This was a reduction form an NP-complete problem to the CPPP mechanism! So the mechanism cannot work in polynomial time. The only problem, the steps were non-constructive. Formally, instead of *unless* $P = NP$ we say *unless* $P/Poly = NP$ where *Poly* is polynomial advise for the non-constructive step.

3 Digital Good Auctions

Suppose, n bidders have (private) valuations $v_1 \geq v_2 \geq v_3 \geq \cdots \geq v_n$ for a digital good such as a movie or an OS. Our intention is to maximize the revenue (maximizing welfare is trivial). What is the ideal? Our goal is to

$$max_i i \cdot v_i$$

Can we approximate it? No! For instance, let $v : 1, 0, 0, 0, \cdots$ in which case no truthful mechanism will get revenue greater than zero. How about benchmark,

$$B = \max_{i \geq 2} i \cdot v_i.$$

Can we approximate this? Yes, but with an expected factor of $15 : 1$. The mechanism through which we can achieve this is known as random sampling (RS) where we divide n bidders into tow sets such that each bidder is equally likely to end up in both sets. Then, get the best price p in the first set and give it to the second set and the other way around. Let $v = (1 + \epsilon, 1, 0, 0, \cdots)$.

Theorem 4. *RS gets expected revenue* $\geq \frac{B}{15}$

Create two partitions, *Good* and *Bad*. Let,

$$x_i : L \text{ is in the Good partition}$$
$$S_i = \sum_{j \leq i} x_j \text{ the number of bidders with bids} \geq v_i \text{ on the other side.}$$

Then in the Good partition we have,

$$S_i \leq \frac{3}{4} \forall i.$$

Lemma 5. $\Pr($*Good Partition*$) \geq \frac{9}{10}$

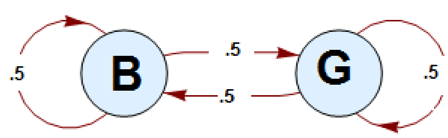

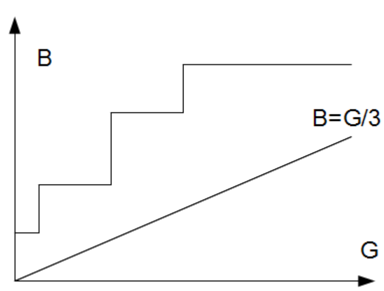

Let,

- L^* be the price on the good side, then $v_{L^*} S_{i^*} \geq v_j S_j$.
- L' be the overall winning price ≥ 2, thus $L' v_{L'} \geq j b_j$.

Consequently, Revenue $\geq (L^* - S_{L^*})v_{L^*}$.

$$\geq_{(G)} \frac{1}{4}L^* v_{L^*} \geq_{(G)} \frac{1}{4}\frac{4}{3}L^* v_{L^*}$$
$$= \frac{1}{3}L^* v_{L^*} \geq \frac{1}{3}L' v_{L'}$$
$$\geq_{(H)} \frac{1}{6}L'^* v_{L'} = \frac{B}{6}$$

So,

$$\text{revenue} \geq (\text{revenue}|G, H)(1 - \Pr(\hat{G}) - \Pr(\hat{H}))$$
$$\geq \frac{B}{6}(1 - \frac{1}{10} - \frac{1}{2}) = \frac{B}{6}\frac{2}{5}$$
$$= \frac{B}{15}$$

Theorem 5. *Nothing better than $\frac{B}{2.42}$ is possible.*

4 Sharing the cost of Multicast Transmissions

Suppose we have a shared delivery tree trough which a content owner sends a single packet to multiple receivers. Given S receive some content, we define the social welfare

$$SW(S) = \sum_{i \in S} v_i - \sum_{l \in T(S)} c(l)$$

where $T(S)$ are all the trees that has user S as a node. We prefer a mechanism that is Incentive Compatible, Welfare maximizer, has a low message complexity, balances the budgets. However, we have the following theorem,

Theorem 6 (Green-Laffont 1977)**.** *No mechanism achieves IC + WM + BB*

We can employ VCG as follows, if rejected pay 0, otherwise pay $v_i - (OSW - OSW_{-i})$

Theorem 7. *VCG can be implemented with 2 messages per link.*

The following algorithm outlines the code for node α,

- want for all children β
$$W^\alpha \leftarrow u^\alpha + \sum_\beta M^\beta - C^\alpha$$
where W is the welfare of subtree, M is the messages, and C is the cost of up link.
- if $W^\alpha < 0$ due , send W^α to the parent.
- Else, root sends W^α (+content) to all live children

We want message M from parent,let $A \leftarrow min(W^{\alpha}, M)$ and pay $max(0, v^i - A)$. Send A (and content) to all live children. This mechanism achieves 3/4; however, it is not BB.

On the other hand, Shapley mechanism which is both BB and IC achieves 2/4. The algorithm for this mechanism is,

- $S \rightarrow \{1, 2, \cdots, n\}$ and while S has changed repeat:
- Send to each node $i \in S$ the proposed payment,

$$p_i = \sum_l \frac{c(l)}{|\text{subtrees of } S \text{ below } l|}$$

- $S \rightarrow S - \{\text{nodes refusing to pay}\}$

While the mechanism is BB and IC, it is not WM and has n messages per link.

Theorem 8. *Any BB mechanism requires n messages per link*