Towards Optimal Bayesian Algorithmic Mechanism Design

Zhiyi Huang (Joint work with Xiaohui Bei)

University of Pennsylvania

September 16, 2010

Outline

1. Background
2. Characterization of Bayesian Incentive Compatibility (BIC)
3. From Algorithm to Mechanism
4. Summary and Future Work

Outline

1. Background
2. Characterization of Bayesian Incentive Compatibility (BIC)
3. From Algorithm to Mechanism
4. Summary and Future Work

A Story About Combinatorial Auction

- Apple, IBM, Google, and Microsoft participate an auction for advertisement slots on CCTV.
- Different companies may have different values for the same slot.
- Different slots may have different values to the companies.
  - e.g. slots in the evening may be more preferable than slots in the morning or midnight.
- Value of combination of slots may not be the sum of the slots' values.
  - e.g. the companies may want slots to be distributed evenly over different time period to cover a broader group of people.
- The values are **private** informations for the companies.

A Story About Combinatorial Auction

- How to design auction (allocate slots, charge prices) so that
  - social welfare (sum of the companies' value of the slots they get) is maximized;
  - the companies will bid their value honestly.

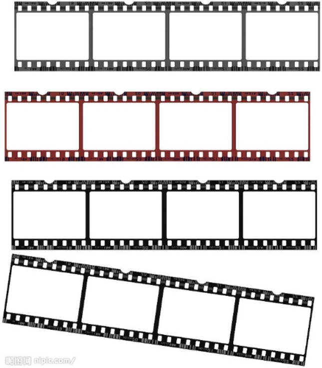

A Story About Combinatorial Auction

- How to design auction (allocate slots, charge prices) so that
  - social welfare (sum of the companies' value of the slots they get) is maximized;
  - the companies will bid their value honestly.

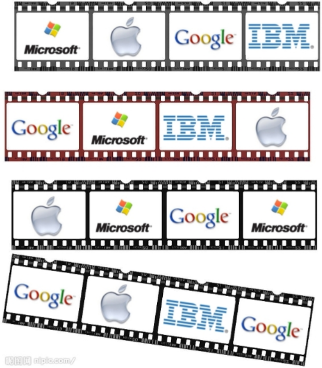

Solution From Economists

VCG Mechanism [Vickrey '61, Clarke '71, Groves '73]

- Always allocate Ad slots using the welfare maximizing scheme.
- There exists pricing scheme that makes the companies tell the truth.

Solution From Economists

VCG Mechanism [Vickrey '61, Clarke '71, Groves '73]

- Always allocate Ad slots using the welfare maximizing scheme.
- There exists pricing scheme that makes the companies tell the truth.
- End of story?

Solution From Economists

VCG Mechanism [Vickrey '61, Clarke '71, Groves '73]

- Always allocate Ad slots using the welfare maximizing scheme.
- There exists pricing scheme that makes the companies tell the truth.
- End of story?

Economists: Hell yeah! It is solved half a century ago!

Solution From Economists

VCG Mechanism [Vickrey '61, Clarke '71, Groves '73]

- Always allocate Ad slots using the welfare maximizing scheme.
- There exists pricing scheme that makes the companies tell the truth.
- End of story?

  Economists: Hell yeah! It is solved half a century ago!

  Computer Scientists: But finding the welfare maximizing allocation and getting proper pricing scheme are both NP-hard. . .

Solution From Economists

VCG Mechanism [Vickrey '61, Clarke '71, Groves '73]

- Always allocate Ad slots using the welfare maximizing scheme.
- There exists pricing scheme that makes the companies tell the truth.
- End of story?

  Economists: Hell yeah! It is solved half a century ago!

  Computer Scientists: But finding the welfare maximizing allocation and getting proper pricing scheme are both NP-hard...

- How about approximation algorithms?

Solution From Economists

VCG Mechanism [Vickrey '61, Clarke '71, Groves '73]

- Always allocate Ad slots using the welfare maximizing scheme.
- There exists pricing scheme that makes the companies tell the truth.
- End of story?

  Economists: Hell yeah! It is solved half a century ago!

  Computer Scientists: But finding the welfare maximizing allocation and getting proper pricing scheme are both NP-hard...

- How about approximation algorithms?
  Combination of VCG and approximation algorithm does **NOT** preserve truthfulness...

Family of Mechanism Design Problems

- **Combinatorial Auction**
- Spectrum Auction
- Digital Auction
- Common Public Project Problem
- Facility Location
- ...

Problem Setting (Combinatorial Auction)

- n heterogeneous items.
- m different agents.
- Each agent i has a private valuation function $v_i : 2^{[n]} \mapsto \mathbb{R}^+$.
- **Bayesian Version:** The prior distribution F_i of v_i is publicly known (we assume poly-size and discrete support).
- An auction mechanism is an (randomized) algorithm $\mathcal{M}$
    - Input: the agents' bidding values $\widetilde{v}_i$
    - Output: an allocation of items $(S_1, S_2, \dots, S_m)$; a collection of prices $(p_1, p_2, \dots, p_m)$.
- Each agent i aims to maximize its utility $v_i(S_i) - p_i$.

Incentive Compatibility (a.k.a. Truthfulness)

- Ex-Post Incentive Compatible **(IC)**:

$$\forall i \; \forall v_{-i} \;\; v_i \in \arg\max_{v} \mathbb{E}_{S_i, p_i \sim \mathcal{M}(v_{-i}, v)}[v_i(S_i) - p_i]$$

- Interim Bayesian Incentive Compatible **(BIC)**:

$$\forall i \;\; v_i \in \arg\max_{v} \mathbb{E}_{v_{-i}}[\mathbb{E}_{S_i, p_i \sim \mathcal{M}(v_{-i}, v)}[v_i(S_i) - p_i]]$$

- Simple observation: **IC** $\Rightarrow$ **BIC**

What We Know Algorithmically

- In general, the best poly-time algorithm achieves $O(\sqrt{n})$ approximation to the optimal social welfare.
- **NO** efficient algorithm can do better than this, even in the restricted single-minded case. [Sandholm '02]

- If we restrict ourself to sub-additive or sub-modular agents, there exists efficient algorithm that provides **constant** approximation. [Feige '05, DS '06, FV '06]

What We Know About IC and Poly-Time Mechanism

- For the general case, there exists $O(\sqrt{n})$ approximation, **IC**, and poly-time mechanism, matching the known computational result! [LS '05, DNS '06]

- But. . . for the restricted sub-additive or sub-modular case, the best known **IC** and poly-time mechanism only gives $O(\log n)$ approximation ratio. [Dobzinski '07]
- Whether one can get to $o(\log n)$ remains open.

Our Contribution

- It would be great if: given any poly-time algorithm $\mathcal{A}$, we can convert it into a **IC** and poly-time mechanism $\mathcal{M}^{\mathcal{A}}$ whose performance is as good as $\mathcal{A}$.
  - Impossible for Common Public Project Problem. [PSS '08]

Theorem

Given any poly-time algorithm $\mathcal{A}$ that computes the allocation of items, we can convert it into a **BIC** *and poly-time mechanism $\mathcal{M}^{\mathcal{A}}$ whose performance is as good as $\mathcal{A}$.*

- Prior to our work, Hartline and Lucier develop similar result for the more restricted single-parameter mechanism design problems. [HL '10]
- Hartline, Kleinberg, and Malekian independently discover similar results.

Outline

1. Background
2. Characterization of Bayesian Incentive Compatibility (BIC)
3. From Algorithm to Mechanism
4. Summary and Future Work

BIC

- Let $\{v_i^1, v_i^2, \ldots, v_i^\ell\}$ be the support of v_i's prior distribution. Then BIC means for any $1 \le s, t \le \ell$:

$$\mathbb{E}_{v_{-i}, S_i, p_i \sim \mathcal{M}(v_{-i}, v_i^s)}[v_i^s(S_i) - p_i] \ge \mathbb{E}_{v_{-i}, S_i, p_i \sim \mathcal{M}(v_{-i}, v_i^t)}[v_i^s(S_i) - p_i]$$

- Let us simplify the notation...

$$\widetilde{p}_s = \mathbb{E}_{v_{-i}, p_i \sim \mathcal{M}(v_{-i}, v_i^s)}[p_i] \, , \; w_{st} = \mathbb{E}_{v_{-i}, S_i \sim \mathcal{M}(v_{-i}, v_i^t)}[v_i^s(S_i)]$$

- **BIC** is equivalent to

$$\forall s, t \;\; w_{ss} - \widetilde{p}_s \ge w_{st} - \widetilde{p}_t$$

Induced Matching Problem

- **BIC** $\Leftrightarrow \forall s, t \;\; w_{ss} - \widetilde{p}_s \geq w_{st} - \widetilde{p}_t$

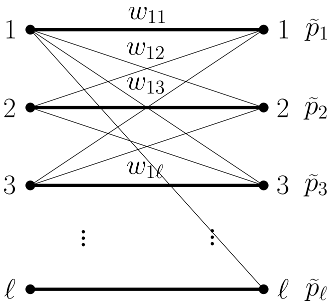

- This is the *envy-free* condition for the above **induced matching problem** of agent i:
  - There exists prices $\widetilde{p}$ such that $\forall\, s \;\; s \in \arg\max_t w_{st} - \widetilde{p}_t$.

The Characterization Lemma

- From the economics literature, there exists envy-free prices if and only if it is a max-matching. [GS '99]

Lemma (Characterization Lemma)

A mechanism is **BIC** *if and only if for each induced matching problem, the identity matching is a max-matching and the prices are the corresponding envy-free prices.*

The Problem with Approximation Algorithms

Consider the following example:

- 2 agents Alice and Bob, 2 items a and b. $v_B(S) = \epsilon, \forall S \neq \phi$,

$$\begin{cases} v_A(\{a,b\}) = 1, v_A(a) = 1, v_A(b) = 1/2 & \text{w.p. } 1/2 \\ v_A(\{a,b\}) = 1, v_A(a) = 1/2, v_A(b) = 1 & \text{w.p. } 1/2 \end{cases}$$

- Optimal social welfare is $1 + \epsilon$.
- A 2-approximation algorithm may gives a to Alice when $v_A(a) = 1/2$ and gives b to her otherwise.
- The induced matching problem for Alice looks like

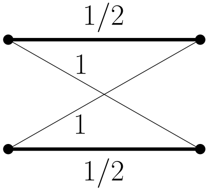

Outline

1 Background

2 Characterization of Bayesian Incentive Compatibility (BIC)

3 From Algorithm to Mechanism

4 Summary and Future Work

High-Level Idea

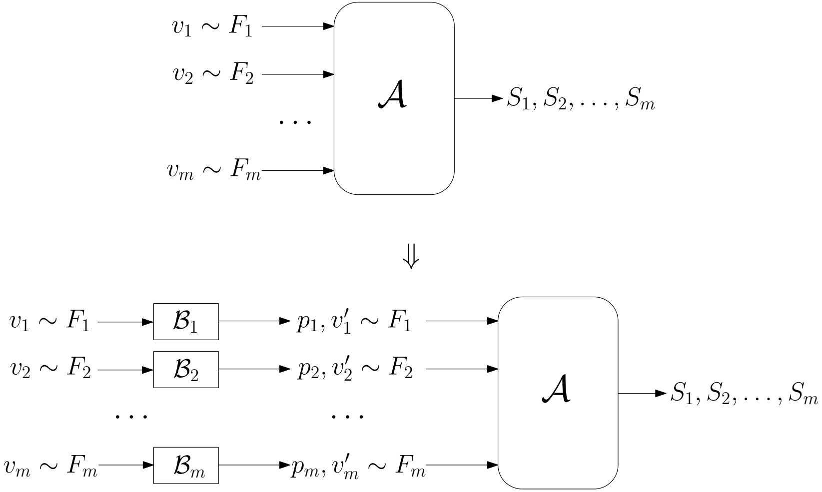

Open the Blackbox

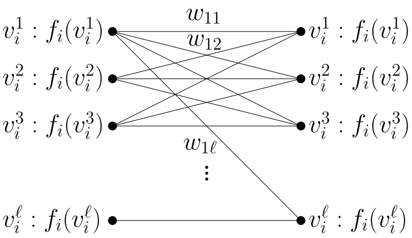

- Solve primal and dual LPs for the induced matching problems (with weights)

$$
\begin{array}{rcll}
\max & \sum_{s=1}^{\ell}\sum_{t=1}^{\ell} x_{st} w_{st} & & \text{s.t.} \\
\sum_{t=1}^{\ell} x_{st} & \leq & f_i(v_i^s) & \forall s \\
\sum_{s=1}^{\ell} x_{st} & \leq & f_i(v_i^t) & \forall t \\
x_{st} & \geq & 0 & \forall s, t
\end{array}
\qquad
\begin{array}{rcll}
\min & \sum_{s=1}^{\ell} f_i(v_i^s) u_s + \sum_{t=1}^{\ell} f_i(v_i^t) p_t & & \text{s.t.} \\
u_s + p_t & \geq & w_{st} & \forall s, t \\
u_s & \geq & 0 & \forall s \\
p_t & \geq & 0 & \forall t
\end{array}
$$

Open the Blackbox (cont'd)

$$
\begin{array}{rcll}
\max & \sum_{s=1}^{\ell}\sum_{t=1}^{\ell} x_{st} w_{st} & & \text{s.t.} \\
\sum_{t=1}^{\ell} x_{st} & \leq & f_i(v_i^s) & \forall s \\
\sum_{s=1}^{\ell} x_{st} & \leq & f_i(v_i^t) & \forall t \\
x_{st} & \geq & 0 & \forall s,t
\end{array}
\qquad
\begin{array}{rcll}
\min & \sum_{s=1}^{\ell} f_i(v_i^s) u_s + \sum_{t=1}^{\ell} f_i(v_i^t) p_t & & \text{s.t.} \\
u_s + p_t & \geq & w_{st} & \forall s,t \\
u_s & \geq & 0 & \forall s \\
p_t & \geq & 0 & \forall t
\end{array}
$$

- Suppose x^*, u^*, and p^* are optimal solutions.
- By duality theorem

$$x^*_{st} > 0 \Leftrightarrow u^*_s + p^*_t = w^*_{st} \Leftrightarrow w^*_{st} - p^*_t = u^*_s \geq w^*_{sr} - p^*_r \quad \forall r \quad .$$

- So p^* are envy-free prices.

Open the Blackbox (cont'd)

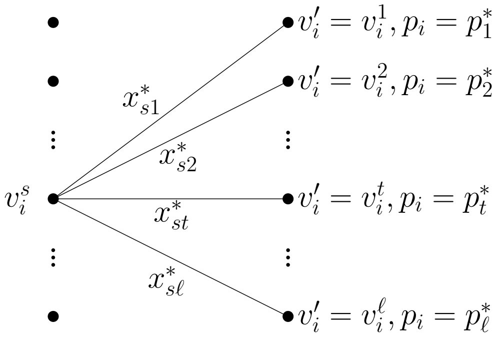

- Suppose agent i claims $v_i = v_i^s$.
- Then w.p. $x^*_{st}/f_i(v_i^s)$, let $v_i' = v_i^t$ and $p_i = p_t^*$.

Our Result

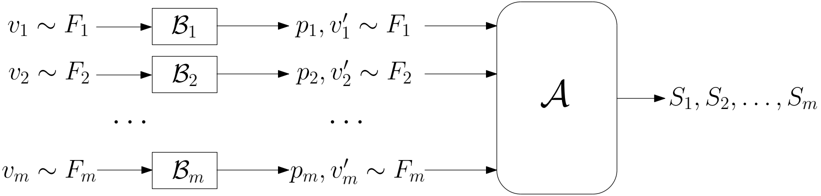

Theorem

Given any poly-time algorithm $\mathcal{A}$ that computes the allocation of items, we can convert it into a **BIC** *and poly-time mechanism $\mathcal{M}^{\mathcal{A}}$ whose performance is as good as $\mathcal{A}$.*

I Cheated You

- When I say let's simplify the notation

$$\widetilde{p}_s = \mathbb{E}_{v_{-i}, p_i \sim \mathcal{M}(v_{-i}, v_i^s)}[p_i] \, , \; w_{st} = \mathbb{E}_{v_{-i}, S_i \sim \mathcal{M}(v_{-i}, v_i^t)}[v_i^s(S_i)]$$

  I hide the fact that computing $\widetilde{p}_s$ and w_{st} exactly is difficult...
- Fortunately, we can sample and estimate these values good enough.

Theorem

*Given any poly-time algorithm $\mathcal{A}$ that computes the allocation of items, we can convert it into a ϵ-**BIC** and poly-time mechanism $\mathcal{M}^{\mathcal{A}}$ whose performance is as good as $\mathcal{A}$.*

Outline

1. Background
2. Characterization of Bayesian Incentive Compatibility (BIC)
3. From Algorithm to Mechanism
4. Summary and Future Work

Summary

- We give a non-trivial characterization of **BIC** mechanism (the Characterization Lemma).
- For any mechanism design problem for social welfare maximization, we propose a reduction which converts any poly-time approximation algorithm into a poly-time and **BIC** mechanism whose performance is as good as the algorithm.
- Our work implies that there exists **constant**-approximation, **BIC**, and poly-time mechanisms for combinatorial auction with sub-additive and sub-modular agents.

Future Works

- Can we get weaker assumption on the prior distributions?
- Is there a reduction converting approximation algorithm to **IC** mechanism for combinatorial auction?
- Can we get better characterization for **IC** mechanisms?

Q & A

Thank You!