Computationally efficient approximation mechanisms

Chapter 12 of *Algorithmic Game Theory*, Lavi (2007)

Geelon So, agso@eng.ucsd.edu
Algorithmic game theory reading group — March 1, 2022

Preliminaries

Game setting

- I is a set of n **players**.
- A is a set of **alternatives**.
- Each player $i \in I$ has a **valuation function** $v_i : A \to \mathbb{R}$.
  - $v_i(a)$ is the value of the alternative $a \in A$ to player $i \in I$
- Say alternative a is chosen and player i pays some money m_i. The **utility** is:

$$\text{utility}_i = v_i(a) - m_i.$$

Each player is aiming to maximize their own utility.

Mechanism

Definition (Mechanism)

*A **mechanism** is a:*

- ***social choice function*** $f : V_1 \times \cdots \times V_n \to A$
- *vector of **payment functions*** $p_i : V_1 \times \cdots \times V_n \to \mathbb{R}$

The mechanism works as follows:

1. Each player reports their valuation function v_i.
2. The alternative $a = f(v_1, \ldots, v_n)$ is chosen.
3. Each player pays $p_i(v_1, \ldots, v_n)$, achieving utility:

$$v_i^*(a) - p_i(v_1, \ldots, v_n),$$

   where v_i^* is the ith player's true valuation function.

Incentive-compatible mechanisms

An **incentive-compatible mechanism** is one where for each player, revealing their true valuation function is a dominant strategy. Formally, if:

$$\begin{aligned}\text{utility}_i^{\text{true}} &= v_i\big(f(v_1, \ldots, v_i, \ldots v_n)\big) - p_i\big(v_1, \ldots, v_i, \ldots, v_n\big)\\ \text{utility}_i^{\text{lie}} &= v_i\big(f(v_1, \ldots, v_i', \ldots v_n)\big) - p_i\big(v_1, \ldots, v_i', \ldots, v_n\big)\end{aligned}$$

Then:

$$\text{utility}_i^{\text{true}} \geq \text{utility}_i^{\text{lie}}.$$

Mechanism design

Goal: how can we design a mechanism that **implements** a social choice function f?

- How can a mechanism be designed so that each player is incentivized to reveal their true valuation function, thus collectively achieving the target social choice?

Vickrey-Clarke-Groves (VCG) mechanism

The VCG mechanism $(f, p_1, \dots, p_n)$ is defined so that:

- the social choice function f maximizes the **social welfare**:

$$f(v_1, \dots, v_n) \in \arg\max_{a \in A} \sum_{i \in I} v_i(a).$$

- each player pays the following amount:

$$p_i\big(v_1, \dots, v_n\big) = -\sum_{j \neq i} v_j\big(f(v_1, \dots, v_n)\big).$$

  - Thus, the utility to each player is the social welfare $f(v_1, \dots, v_n)$, which by definition is maximized if each player is truthful.

VCG is incentive compatible

Theorem

The VCG mechanism is incentive compatible.

Incentive-compatible mechanisms

Question. What other social choice functions are incentive compatible?

Examples. Other reasonable social choice functions:

- $f(v_1, \dots, v_n) \in \arg\max_{a \in A} \min_{i \in I} \text{utility}_i(a)$
- $f(v_1, \dots, v_n)$ incorporates fairness into social welfare
- $f(v_1, \dots, v_n)$ efficiently chooses an alternative that approximately maximizes the (possibly computationally intractable) social welfare

Randomized mechanisms

Definition (Randomized mechanism)

*A **randomized mechanism** is a distribution over deterministic mechanisms.*

- A randomized mechanism is incentive-compatible in a **universal sense** if it is supported over incentive-compatible deterministic mechanisms.
- A randomized mechanism is incentive-compatible **in expectation** if truth is a dominant strategy in the game induced by expectation.
  - That is, the utility to player i is:

$$\text{utility}_i = \mathop{\mathbb{E}}_{(f, p_1, \ldots, p_n)} \left[v_i^* \big(f(v_1, \ldots, v_n) \big) - p_i \big(v_1, \ldots, v_n \big) \right].$$

  - Weaker that incentive-compatibility—for example, it does not take into account a player's risk-aversion (assumes players only care about maximizing expected utility).

Mechanism design with computational considerations

Combinatorial auctions (CA)

Game setting:

- We allocate m items Ω to n players I.
- Players value subsets, so:

$$v_i(S) = \text{value of } S \subset \Omega \text{ to player } i.$$

  - Assume valuations are monotonic: $v_i(S) \leq v_i(T)$ if $S \subset T$
  - Assume valuations are normalized: $v_i(\varnothing) = 0$.

- Our goal is to maximize **social welfare**:

$$f(v_1, \dots, v_n) = \arg\max_{(S_1, \cdots, S_n)} \sum_{i \in I} v_i(S_i),$$

where $S_i, S_j \subset \Omega$ are disjoint for $i \neq j$.

CA through VCG

If we can exactly compute:

$$f(v_1, \ldots, v_n) = \underset{(S_1, \cdots, S_n)}{\arg\max} \sum_{i \in I} v_i(S_i),$$

then we can implement VCG to achieve an incentive-compatible mechanism.

- However, f is computationally infeasible.
  - In general, even the valuation v_i is exponential in size.

Two models of CA

- **Bidding language model:** v_i must have a succinct representation.
  - It is NP-hard to approximate social welfare within a ratio of $\Omega(m^{1/2-\varepsilon})$ for any $\varepsilon > 0$.
- **Query access model:** the mechanism may query players for values $v_i(a)$.
  - Any algorithm with polynomial communication cannot approximate social welfare within a ratio of $\Omega(m^{1/2-\varepsilon})$ for any $\varepsilon > 0$.

Impossibility of computational-efficiency and incentive-compatibility

Implication. If P $\neq$ NP, then there does not exist a computationally-efficient incentive-compatible mechanism for either models of combinatorial auctions.

Approximation mechanisms through randomness

Goal: a randomized approximation mechanism that is truthful in expectation.

Key idea to design random mechanism:

- Relax problem and solve for optimal fractional allocation.
- Construct distribution over integral allocations such that allocations match the optimal fractional allocation in expectation.
  - Use VCG mechanism to ensure truthfulness in expectation.

Fractional domain

The fractional CA problem is the following linear program:

$$
\begin{aligned}
\max \quad & \sum_{i \in I, S \neq \varnothing} x_{i,S} \cdot v_i(S) \\
\text{subject to} \quad & \sum_{S \neq \varnothing} x_{i,S} \leq 1 \quad \text{for each player } i \in I \\
& \sum_{i \in I, S : j \in S} x_{i,S} \leq 1 \quad \text{for each item } j \in \Omega \\
& x_{i,S} \geq 0 \quad \text{for all } i \in I, S \neq \varnothing
\end{aligned}
$$

Player i receives a $x_{i,S}$ fraction of set S.

- The first constraint ensures each player is allocated at most one unit of subsets.
- The second ensures the amount of each item j allocated out is at most one.

Efficient optimal mechanism for fractional CA

Proposition

For the fractional CA problem, there is a truthful optimal mechanism with efficient computation and communication for the bidding language model and query-access models.

- Solve the linear program (e.g. ellipsoid method) to obtain *optimal fractional allocation*; apply VCG.

Fractional CA to integral CA

Converting a fractional allocation to a distribution over integral allocations

- Suppose we can verify that CA has a *c-integrality gap*. That is, we have an efficient algorithm returns an integral point $\tilde{x}$ such that:

$$c \cdot \sum_{i,S} \tilde{x}_{i,S} \cdot v_i(S) \geq \max_{\text{feasible } x\text{'s}} \sum_{i,S} x_{i,S} \cdot v_i(S).$$

  We can find an integral allocation $\tilde{x}$ no more than c times worse than optimal.
- **Decomposition lemma.** Suppose we can verify a c-integrality gap. If x is a feasible point, then we can efficiently decompose x/c into a convex combination of integral feasible points.

Decomposition-based random mechanism

1. Compute an optimal fractional solution x^* and VCG prices $p_i^*(v)$.
   - Let $v = (v_1, \ldots, v_n)$.
2. Obtain decomposition: $x^*/c = \sum_{\ell \in \mathcal{I}} \lambda_\ell \cdot x^\ell$
   - $\mathcal{I}$ set of feasible integral allocations
3. With probability λ_ℓ:
   (i) choose allocation x^ℓ
   (ii) set prices $p_i(v) = [v_i(x^\ell)/v_i(x^*)]p_i^*(v)$.

Expected welfare and price

Expected welfare:

$$\mathbb{E}[\text{social welfare}] = \frac{1}{c} \sum_{i \in I} v_i(x^*).$$

Expected price per player:

$$\begin{aligned}\mathbb{E}\left[p_i(v)\right] &= \sum_{\ell \in \mathcal{I}} \lambda_\ell \cdot [v_i(x^\ell)/v_i(x^*)] p_i^*(v) \\ &= \left[p_i^*(v)/v_i(x^*)\right] \cdot \sum_{\ell \in \mathcal{I}} \lambda_\ell \cdot v_i(x^\ell) = p_i^*(v)/c.\end{aligned}$$

Thus, truthfulness follows from the optimality of fractional CA.

Impossibilities of dominant strategy implementability

Nothing besides VCGs

Definition (Affine maximizer)

A social choice function f is an ***affine maximizer*** *if there exist weights $k_1, \ldots, k_n$ for each player and constants C_a for each alternative such that:*

$$f(v_1, \ldots, v_n) \in \arg\max_{a \in A} C_a + \sum_{i \in I} k_i v_i(x).$$

Theorem (Nothing but VCGs)

Suppose $|A| \geq 3$ and $v_i \in \mathbb{R}^A$ is unrestricted for all i. Then f is dominant-strategy implementable if and only if it is an affine maximizer.

Implication as impossibility result

1. Often, VCG is computationally intractable (e.g. CAs)
2. We can also seek goals different from welfare maximization.

Alternative solution concept

Definition (Algorithmic implementation)

A mechanism M is an ***algorithmic implementation*** *of a c-approximation in undominated strategies if there exists a set of strategies D such that:*

(i) *M obtains a c-approximation for any combination of strategies from D in polytime*

(ii) *for any strategy not in D, there is a strategy in D that weakly dominates it, and this transition is polynomial-time computable.*

References

Ron Lavi. Computationally efficient approximation mechanisms. *Algorithmic Game Theory*, pages 301–329, 2007.