Price of Anarchy of Non-Truthful Mechanisms

Instructor: Thomas Kesselheim

Since we started our discussion of mechanisms we were focused on designing truthful mechanisms. While for single-parameter settings the resulting monotonicity requirement turned out to be rather nice, this was not the case for multi-parameter mechanisms. In fact, very simple algorithms give the best approximation in polynomial time but they cannot be turned into truthful mechanisms.

In this lecture we will explore a different direction, which consists of relaxing the incentive constraint. Instead of asking for truthful mechanisms, we will ask for mechanisms whose equilibria are all close to optimal. To this end we will translate the smoothness concept that we have seen a few lectures ago from games to mechanisms and use it to design simple, near-optimal mechanisms.

1 Basic Definitions

Recall our definition of a mechanism-design problem. We will focus on settings, where the players' preferences are given by valuation functions and the goal is to maximize social welfare.

Definition 12.1. *A* mechanism-design problem *is defined by a set $\mathcal{N}$ of n players and a set of feasible outcomes X. Every player $i \in \mathcal{N}$ has a (private) valuation $v_i \colon X \to \mathbb{R}_{\geq 0}$ from a set of possible valuations V_i.*

Since the valuations are private we will again consider mechanisms, where we slightly generalize our previous definition, so that it can also handle the case where the bids come from a different set than the valuations.

Definition 12.2 (Mechanism)**.** *A* mechanism *consist of $\mathcal{M} = (f, p)$ for the welfare maximization problem defines a set of bids B_i for each player $i \in \mathcal{N}$ and consists of*

- *an* outcome rule $f : B \to X$*, and*
- *a* payment rule $p : B \to \mathbb{R}^n_{\geq 0}$.

We say that the mechanism is direct *if $B_i = V_i$ for all $i \in \mathcal{N}$, otherwise we say it is* indirect.

The utility of bidder i on bid profile $b \in B$ is given as $u_i(b, v_i) = v_i(f(b)) - p_i(b)$. For a fixed choice of v, these utilities define a normal-form maximization game. Today, we will study the equilibria of this game.

A pure Nash equilibrium is a vector of strategies – in this case bids – such that no player wants to unilaterally deviate.

Definition 12.3 (Pure Nash Equilibrium)**.** *A profile of bids $b = (b_1, \ldots, b_n) \in B_1 \times B_2 \times \cdots \times B_n$ is a* pure Nash equilibrium *(PNE) if for every player $i \in \mathcal{N}$ and every deviation $b'_i \in B_i$,*

$$u_i((b_i, b_{-i}), v_i) \geq u_i((b'_i, b_{-i}), v_i) \enspace .$$

Also the concepts of mixed Nash and coarse correlated equilibria still make sense here.

Definition 12.4 (Mixed Nash Equilibrium, Coarse Correlated Equilibrium)**.** *A distribution $\mathcal{D}$ over B is a* coarse correlated equilibrium *(CCE) if for every player $i \in \mathcal{N}$ and every deviation $b'_i \in B_i$,*

$$\mathbb{E}_{b \sim \mathcal{D}}[u_i((b_i, b_{-i}), v_i)] \geq \mathbb{E}_{b \sim \mathcal{D}}[u_i((b'_i, b_{-i}), v_i)] \enspace .$$

If $\mathcal{D}$ is a product distribution then it is called a Mixed Nash Equilibrium.

The goal is to choose an outcome $x \in X$ that maximizes *social welfare* $\sum_{i \in \mathcal{N}} v_i(x_i)$. We use $OPT(v) = \max_{x \in X} \sum_{i \in \mathcal{N}} v_i(x_i)$ to denote the optimal social welfare. For a fixed bid vector b, the mechanism achieves welfare $SW_v(b) = \sum_{i \in \mathcal{N}} v_i(f(b)) = \sum_{i \in \mathcal{N}} u_i(b) + \sum_{i \in \mathcal{N}} p_i(b)$.

We define the *Price of Anarchy* for any given equilibrium concept as the worst possible ratio between the optimal social welfare and the (expected) social welfare at equilibrium, that is

$$PoA_{\mathsf{Eq}} = \max_{v \in V} \max_{\mathcal{D} \in \mathsf{Eq}(v)} \frac{OPT(v)}{\mathbb{E}_{b \sim \mathcal{D}}[SW_v(b)]},$$

where $\mathsf{Eq}(v)$ denotes the set of equilibria for the game induced by valuations v.

We now have

$$PoA_{\mathsf{PNE}} \leq PoA_{\mathsf{MNE}} \leq PoA_{\mathsf{CCE}}.$$

2 The Smoothness Framework

We define smooth mechanisms and show how smoothness implies that all equilibria of a mechanism are close to optimal.

Definition 12.5 (Smooth Mechanism, simplified version)**.** *Let $\lambda, \mu \geq 0$. A mechanism M is (λ, μ)-smooth if for any valuation profile $v \in V_i$ for each player $i \in \mathcal{N}$ there exists a bid b_i^* such that for any profile of bids $b \in B$ we have*

$$\sum_{i \in \mathcal{N}} u_i(b_i^*, b_{-i}) \geq \lambda \cdot OPT(v) - \mu \sum_{i \in \mathcal{N}} p_i(b) \enspace .$$

Note that by the order of the quantifiers b_i^* may depend on the profile of valuations but not on the bids

Let's get some intuition for the definition of a (λ, μ)-smooth mechanism by considering a single-item first price auction.

Observation 12.6. *A single-item first-price auction is $(1/2, 1)$-smooth.*

Proof. Consider the player i with the highest value v_i. Suppose this player deviates to $b_i' = v_i/2$. Consider an arbitrary bid profile $b \in B$. Notice that $\sum_{i \in \mathcal{N}} p_i(b) = \max_j b_j$.

Now distinguish two cases: If $\max_{j \neq i} b_j > v_i/2$ then player i does not win the item. In this case, $u_i(b_i', b_{-i}) = 0 > v_i/2 - \max_{j \neq i} b_j \geq v_i/2 - \max_j b_j$. Otherwise, $\max_{j \neq i} b_j \leq v_i/2$ and player i wins the item. His utility in this case is $u_i(b_i', b_{-i}) = v_i - p_i(b_i', b_{-i}) = v_i - v_i/2 = v_i/2 \geq v_i/2 - \max_j b_j$. So in both cases $u_i(b_i', b_{-i}) \geq OPT(v)/2 - \sum_{i \in \mathcal{N}} p_i(b)$.

The claim follows by considering a deviation b_j' for the players $j \neq i$ that gives them a utility of at least zero such as bidding zero. □

Theorem 12.7 (Syrgkanis and Tardos, 2013)**.** *If a mechanism $\mathcal{M}$ is (λ, μ)-smooth and players have the possibility to withdraw from the mechanism then*

$$PoA_{\mathsf{CCE}} \leq \frac{\max\{\mu, 1\}}{\lambda} \enspace .$$

Proof. We will first prove the claim for pure Nash equilibria and then argue that our proof pattern extends to coarse correlated equilibria.

Suppose bid profile b is a pure Nash equilibrium. What does this mean? It means that no player wants to unilaterally deviate from the equilibrium bid to some other bid. That is,

$$u_i(b_i, b_{-i}) \geq u_i(b_i', b_{-i}) \enspace ,$$

for all players $i \in \mathcal{N}$ and bids $b_i' \in B_i$.

Now in particular players do not want to deviate to the (derandomized) bid b_i' whose existence is guaranteed by smoothness. Considering, for each player $i \in \mathcal{N}$ the deviation to b_i' and summing over all players,

$$\sum_{i \in \mathcal{N}} u_i(b_i, b_{-i}) \geq \sum_{i \in \mathcal{N}} u_i(b_i', b_{-i}) \geq \lambda \cdot OPT(v) - \mu \cdot \sum_{i \in \mathcal{N}} p_i(b) \enspace .$$

Since players have quasi-linear utilities $u_i(b) = v_i(b) - p_i(b)$ or $v_i(b) \geq u_i(b) + p_i(b)$. Using this we obtain

$$\sum_{i \in \mathcal{N}} v_i(b) \geq \lambda \cdot OPT(v) + (1 - \mu) \cdot \sum_{i \in \mathcal{N}} p_i(b) \enspace .$$

Notice that the left-hand side is precisely the social welfare at equilibrium. So if $\mu \leq 1$ we can bound $(1 - \mu) \cdot \sum_{i \in \mathcal{N}} p_i(b) \geq 0$ and obtain

$$\sum_{i \in \mathcal{N}} v_i(b) \geq \lambda \cdot OPT(v) \enspace ,$$

which shows a Price of Anarchy of $1/\lambda = \max\{1, \mu\}/\lambda$.

On the other hand, if $\mu > 1$, we can use that players have the right to withdraw from the mechanism and obtain a utility of zero to argue that $u_i(b) = v_i(b) - p_i(b) \geq 0$ and so $p_i(b) \leq v_i(b)$. Since $(1 - \mu) < 0$ we obtain

$$\sum_{i \in \mathcal{N}} v_i(b) \geq \lambda \cdot OPT(v) + (1 - \mu) \cdot \sum_{i \in \mathcal{N}} v_i(b) \enspace .$$

Subtracting $(1 - \mu) \cdot \sum_{i \in \mathcal{N}} v_i(b)$ and dividing by $\mu > 1$ we obtain

$$\sum_{i \in \mathcal{N}} v_i(b) \geq \lambda/\mu \cdot OPT(v) \enspace ,$$

which again shows a Price of Anarchy bound of $\mu/\lambda = \max\{1, \mu\}/\lambda$.

Now does this argument extend to more general equilibrium concepts such as coarse correlated equilibria? The only point where we used the equilibrium condition is when we argued that players do not want to deviate from the equilibrium bid b_i to some other bid b_i'. In fact, the specific deviations that we considered only depended on the valuation profile v and did not depend on the bids b. Hence the exact same argument applies to coarse correlated equilibria and shows a Price of Anarchy of $\max\{1, \mu\}/\lambda$. □

In the remainder we will use the smoothness framework to show that certain simple mechanisms achieve near-optimal performance. Both results were discovered before the smoothness result was discovered, but the basic arguments were already present in the original publications. So we cite those.

3 Multi-Minded Combinatorial Auctions

As a first application of the smoothness framework, we will use it to show that the greedy mechanism for k-minded combinatorial auctions combined with a first-price payment rule has Price of Anarchy $O(\sqrt{m})$.

First-Price Greedy Mechanism for Multi-Minded CAs

1. Collect bids b.
2. Sort the player-bundle pairs (i, S) by non-increasing score $\frac{b_i(S)}{\sqrt{|S|}}$.
3. Go through the sorted list and assign S to player i unless

(a) player i has already been allocated a bundle or

(b) one or more of the items in S has already been allocated.

4. Charge each player i his bid $b_i(S)$ on the bundle S he is allocated.

The allocation algorithm that this mechanism is based upon is an $O(\sqrt{m})$-approximation, even if a bidder can report multiple pairs $(S, b_i(S))$. We will now show that all equilibria of this mechanism achieve social welfare within a factor of $O(\sqrt{m})$ of the optimal social welfare.

Theorem 12.8 (Borodin and Lucier, 2010)**.** *The first-price greedy mechanism for multi-minded CAs is* $(1/2, \sqrt{m})$*-smooth.*

Define for each player i and bundle S the *critical bid* $\tau_i(S, b_{-i})$ as the smallest bid with which player i wins bundle S against bids b_{-i}.

We first show that the greedy mechanisms approximates not only the sum of the bids on the optimal allocation as every approximation algorithm would, but that it also approximates the sum of the critical bids on the optimal allocation.

Lemma 12.9. *Fix bids $b \in B$. Let X be the allocation chosen by the greedy mechanism for bids b and let X^* be the allocation that maximizes soical welfare with respect to v. Then,*

$$\sum_{i \in \mathcal{N}} b_i(X_i) \geq \frac{1}{\sqrt{m}} \sum_{i \in N} \tau_i(X_i^*, b_{-i}).$$

Proof. Choose any $\epsilon > 0$. For all i, let b'_i be the single-minded declaration for set X_i^* at value $\tau_i(X_i^*, b_{-i}) - \epsilon$. Let b_i^* be the point-wise maximum of b_i and b'_i. The allocation chosen by greedy on profile b^* is the same as on b. So

$$\sum_{i \in \mathcal{N}} b_i(X_i) = \sum_{i \in \mathcal{N}} b_i^*(X_i) \geq \frac{1}{\sqrt{m}} \sum_{i \in \mathcal{N}} b_i^*(X_i^*) \geq \frac{1}{\sqrt{m}} \sum_{i \in \mathcal{N}} \tau_i(X_i^*, b_{-i}) - n\epsilon \ ,$$

where the first inequality follows from the fact that greedy is a $\sqrt{m}$-approximation and the second inequality follows from the definition of b^* as the point-wise maximum of b and b'. The claim follows by taking the limit as $\epsilon \to 0$. □

Proof of Theorem 12.8. Consider an arbitrary bid profile b. For each player $i \in \mathcal{N}$ let b'_i be the single-minded declaration for set X_i^* at value $v_i(X_i^*)/2$. Now distinguish two cases: If $v_i(X_i^*)/2 \geq \tau_i(X_i^*, b_{-i})$, then player i wins the bundle and his utility is $u_i(b'_i, b_{-i}) = v_i(X_i^*) - v_i(X_i^*)/2 = v_i(X^*)/2 \geq v_i(X_i^*)/2 - \tau_i(X_i^*, b_{-i})$. Otherwise, player i does not win and his utility is $u_i(b'_i, b_{-i}) = 0 \geq v_i(X_i^*)/2 - \tau_i(X_i^*, b_{-i})$. So in either case,

$$u_i(b'_i, b_{-i}) \geq v_i(X_i^*)/2 - \tau_i(X_i^*, b_{-i}) \ .$$

Summing over all players $i \in \mathcal{N}$ we obtain

$$\begin{aligned}
\sum_{i \in \mathcal{N}} u_i(b'_i, b_{-i}) &\geq \sum_{i \in \mathcal{N}} \left(\frac{v_i(X_i^*)}{2} - \tau_i(X_i^*, b_{-i}) \right) \\
&= \frac{1}{2} \cdot OPT(v) - \sum_{i \in \mathcal{N}} \tau_i(X_i^*, b_{-i}) \\
&\geq \frac{1}{2} \cdot OPT(v) - \sqrt{m} \cdot \sum_{i \in \mathcal{N}} b_i(X_i) \\
&= \frac{1}{2} \cdot OPT(v) - \sqrt{m} \cdot \sum_{i \in \mathcal{N}} p_i(b) \ ,
\end{aligned}$$

where we used Lemma 12.9 in the penultimate step and the last step uses that the mechanism is a first-price mechanism. □

Recommended Literature

- Vasilis Syrgkanis and Eva Tardos. Composable and Efficient Mechanisms. STOC'13. (Smoothness for mechanisms)
- Paul Dütting and Thomas Kesselheim. Algorithms against Anarchy: Understanding Non-Truthful Mechanisms. EC'15. (Characterization of algorithms with small PoA)
- Allan Borodin and Brendan Lucier. Price of Anarchy of Greedy Auctions. SODA'10. (The PoA result for greedy multi-minded CAs, results for general greedy algorithms)