Non-Trutfhul Mechanisms and the Smoothness Framework

ETH Zürich Paul Dütting

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, we had to work hard to design truthful polynomial-time mechanisms for multi-parameter settings and the resulting mechanisms turned out to be rather complicated.

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 10.1 (Welfare Maximization Problem)**.** *A* mechanism design problem *is defined by a set $\mathcal{N}$ of n players and a set of feasible outcomes $X \subseteq X_1 \times X_2 \times \cdots \times X_n$. Every player $i \in \mathcal{N}$ has a (private) valuation $v_i : X_i \to \mathbb{R}_{\geq 0}$ from a set of possible valuations V_i. Every player $i \in \mathcal{N}$ has quasi-linear utility $u_i(x_i, p_i) = v_i(x_i) - p_i$. The goal is to choose an outcome 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.*

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 10.2 (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*.*

In what follows we will always consider a fixed mechanism $\mathcal{M} = (f, p)$, so we will simplify notation by writing $u_i(b)$ for $u_i(f_i(b), p_i(b))$.

Next we will define mixed Nash equilibria and coarse correlated equilibria. We will follow our convention and use σ_i to denote a mixed strategy of player i, where a mixed strategy is a probability distribution over bids $b_i \in B_i$. We will use Σ_i to denote all possible probability distributions over B_i, and Σ for all joint probability distributions over B.

Definition 10.3 (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 $b'_i \in B_i$,*

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

Definition 10.4 (Mixed Nash Equilibrium)**.** *A profile of mixed strategies* $\sigma = (\sigma_1, \ldots, \sigma_n) \in \Sigma_1 \times \Sigma_2 \times \cdots \times \Sigma_n$ *is a* mixed Nash equilibrium *(MNE) if for every player* $i \in \mathcal{N}$ *and every deviation* $b'_i \in B_i$

$$\mathbb{E}_{b\sim\sigma}[u_i(b_i, b_{-i})] \geq \mathbb{E}_{b_{-i}\sim\sigma_{-i}}[u_i(b'_i, b_{-i})] \ .$$

Definition 10.5 (Coarse Correlated Equilibrium)**.** *A profile of mixed strategies* $\sigma = (\sigma_1, \ldots, \sigma_n) \in \Sigma$ *is a* coarse correlated equilibrium *(CCE) if for every player* $i \in \mathcal{N}$ *and every deviation* $b'_i \in B_i$*,*

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

We still have the relationship that every pure Nash equilibrium is a mixed Nash equilibrium and every mixed Nash equilibrium is a coarse correlated equilibrium.

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, and write PoA_{PNE}, PoA_{MNE}, and PoA_{CCE} where

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

2 The Smoothness Framework

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

Definition 10.6 (Smooth Mechanism)**.** *Let* $\lambda, \mu \geq 0$*. A mechanism* M *is* (λ, μ)*-smooth if for any valuation profile* $v \in V_i$ *and for any profile of bids* $b \in B$ *there exists a mixed strategy* $\sigma_i^*(v)$ *for each player* $i \in \mathcal{N}$ *such that*

$$\sum_{i\in\mathcal{N}} \mathbb{E}_{b_i^*\sim\sigma_i^*(v)}[u_i(b_i^*, b_{-i})] \geq \lambda \cdot OPT(v) - \mu \sum_{i\in\mathcal{N}} p_i(b) \ .$$

The notation $\sigma_i^*(v)$ means that the mixed strategy may depend on the profile of valuations. Note that the existence of such a randomized bid implies the existence of a deterministic bid that is at least as good.

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

Observation 10.7. *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 10.8 (Syrgkanis and Tardos, 2013)**.** *If a mechanism* $\mathcal{M}$ *is* (λ, μ)*-smooth and players have the possibility to withdraw from the mechanism then*

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

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,

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

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,

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

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

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

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

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

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

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

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

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

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 a $\sqrt{m}$-approximation. 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 10.9 (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* $c_i(S, b_{-i})$ 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 10.10. *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,*

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

Proof. Choose any $\epsilon > 0$. For all i, let b_i' be the single-minded declaration for set X_i^* at value $c_i(X_i^*, b_{-i}) - \epsilon$. 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

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

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 10.9. 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 $v_i(X_i^*)/2 \geq c_i(X_i^*, b_{-i})$, then player i wins the bundle and his utility is $u_i(b_i', b_{-i}) = v_i(X_i^*) - v_i(X_i^*)/2 = v_i(X^*)/2 \geq v_i(X_i^*)/2 - c_i(X_i^*, b_{-i})$. Otherwise, player i does not win and his utility is $u_i(b_i', b_{-i}) = 0 \geq v_i(X_i^*)/2 - c_i(X_i^*, b_{-i})$. So in either case,

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

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

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

where we used Lemma 10.10 in the penultimate step and the last step uses that the mechanism is a first-price mechanism. □

4 Combinatorial Auctions with Item Bidding

As a second application we consider a variant of the combinatorial auction problem in which the players have submodular valuations, but we ask them for additive bids. We then choose the allocation that maximizes the sum of the bids and charge winning players their bids on the items they have won. Notice that alternatively we could view this procedure as running a separate first-price auction for each item. The resulting mechanism is indirect as it does not allow players to express their preferences truthfully.

Definition 10.11. *A valuation function $v_i : 2^M \to \mathbb{R}_{\geq 0}$ is submodular if for every two sets S, T such that $S \subseteq T$ and every item j we have*

$$v_i(S \cup \{j\}) - v_i(S) \geq v_i(T \cup \{j\}) - v_i(T)\ .$$

Submodular valuations are also known as valuations with "decreasing marginals". The reason is that adding an element to a smaller set has a larger effect than adding that same element to a larger set.

Combinatorial Auctions with Item Bidding

1. Collect additive bids b_i from each player $i \in \mathcal{N}$.
2. For each item $j \in M$:
   (a) Assign the item to the player i with the highest bid $\max i b_i(j)$.
   (b) Make this player pay his bid.

Theorem 10.12 (Christodoulou, Kovacz, Schapira, 2008)**.** *A combinatorial auction with submodular valuations and item bidding is $(1/2, 1)$-smooth.*

We will use the following property of submodular valuations.

Lemma 10.13. *For every submodular valuation v_i and every set T there is an additive function a_i such that*

$$\sum_{j\in S} a_i(j) \begin{cases} = v_i(T) & \text{for } S = T, \text{ and}\\ \leq v_i(S) & \text{otherwise.}\end{cases}$$

Proof of Theorem 10.12. Fix additive bids $b \in B$. Denote the allocation chosen by the mechanism for bids b by X and denote the allocation that maximizes social welfare with respect to v by X^*. For each bidder $i \in \mathcal{N}$ consider the additive bid a_i whose existence is guaranteed by Lemma 10.13 where we require the bid to be exact on X_i^*. That is, for each player $i \in \mathcal{N}$ consider

$$\sum_{j \in T} a_i(j) \begin{cases} = v_i(X_i^*) & \text{for } T = X_i^*, \text{ and} \\ \leq v_i(T) & \text{otherwise.} \end{cases}$$

We can assume that a_i is zero outside of X_i^* for every player $i \in \mathcal{N}$. Suppose player i bids $b_i' = a_i/2$ against b_{-i}. Denote the set that player i wins by Y_i. Since b_i' is zero outside X_i^*, we have $Y_i \subseteq X_i^*$. Then,

$$\begin{aligned} u_i(b_i', b_{-i}) &= v_i(Y_i) - \sum_{j \in Y_i} \frac{a_i(j)}{2} \\ &\geq \sum_{j \in Y_i} \left(a_i(j) - \frac{a_i(j)}{2} \right) \\ &\geq \sum_{j \in Y_i} \frac{a_i(j)}{2} + \sum_{j \in X_i^* \setminus Y_i} \left(\frac{a_i(j)}{2} - \max_{\ell \neq i} b_\ell(j) \right) \\ &= \frac{v_i(X_i^*)}{2} - \sum_{j \in X_i^* \setminus Y_i} \max_{\ell \neq i} b_\ell(j) \\ &\geq \frac{v_i(X_i^*)}{2} - \sum_{j \in X_i^*} \max_{\ell} b_\ell(j) \ , \end{aligned}$$

where the first inequality uses that a_i underbids on $Y_i \subseteq X_i^*$, the second uses that player i lost each of the items j that we add in this step with a bid of $a_i(j)/2$, and the third makes what we subtract larger by increasing the range of the sum and the max.

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

$$\begin{aligned} \sum_{i \in \mathcal{N}} u_i(b_i', b_{-i}) &\geq \frac{OPT(v)}{2} - \sum_{j \in M} \max_{\ell} b_\ell(j) \\ &= \frac{OPT(v)}{2} - \sum_{i \in \mathcal{N}} \sum_{j \in X_i} b_i(j) \\ &= \frac{OPT(v)}{2} - \sum_{i \in \mathcal{N}} p_i(b) \ , \end{aligned}$$

where we used that the mechanism allocates each item to the player with the highest bid and that player has to pay his bid. □

5 Epilogue

The performance of non-truthful mechanisms has been and still is a very active area of research. One can distinguish at least three directions:

The first are better bounds and a more general theory for how "valuation compressions"—such as the one seen in the section on combinatorial auctions with item bidding—affect the Price of Anarchy. Better bounds for additive bids are given in Bhawalkar and Roughgarden (2011) and Feldman et al. (2013).

Then there has been an attempt on identifying classes of algorithms that translate approximation guarantees into Price of Anarchy guarantees. One such class are greedy algorithms, where an argument similar to the one presented above, applies more broadly. See Borodin and Lucier (2010) for details. A second class of algorithms for which this is the case are algorithms that follow the relax-and-round paradigm. See Dütting et al. (2015).

Finally, there have been attempts to obtain a Myerson's lemma-like characterization. So far such a result is only available for single-parameter settings—see Dütting and Kesselheim (2015)—but it is likely that this characterization extends to multi-parameter settings.

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)
- Paul Dütting, Eva Tardos, Thomas Kesselheim. The Price of Anarchy of Relax-and-Round. EC'15. (Results for relax-and-round algorithms)
- George Christodoulou, Annamaria Kovacs, Michael Schapira. Bayesian Combinatorial Auctions. ICALP'08. (The result for combinatorial auctions with item bidding)
- Kshipra Bhawalkar and Tim Roughgarden. Welfare Guarantees for Combinatorial Auctions with Item Bidding. SODA'11. (CAs with item bidding and subadditive players)
- Michal Feldman, Hu Fu, Nick Gravin, Brendan Lucier. Combinatorial Auctions are (Almost) Efficient. STOC'13. (Best known bounds for CAs with item bidding)
- Paul Dütting, Monika Henzinger, Martin Starnberger. Valuation Compressions in VCG-Based Combinatorial Auctions. WINE'13. (PoA bounds for CAs with non-additive bids)