Truthful Multi-Parameter Mechanisms and Black-Box Reductions

ETH Zürich Paul Dütting

The design of truthful mechanism in single-parameter settings was hugely simplified by Myerson's Lemma. It told us that allocation rules that can be implemented in a truthful mechanism are precisely those allocation rules that are monotone, and monotonicity turned out to be a rather friendly constraint. Many natural approximation algorithms are already monotone or they can be tweaked and made monotone.

The picture changes when we turn to multi-parameter problems. While there is still a succinct characterization of implementable allocation rules, the resulting requirement "cyclic monotonicity" is very hard to check and ensure. We will therefore take a "bottom up" approach and present a general technique due to Lavi and Swamy for turning polynomial-time algorithms into truthful mechanisms.

1 Combinatorial Auctions with Multi-Minded Bidders

As in the previous lecture our running example will be combinatorial auctions. This time we will be primarily concerned with the case where bidders are interested in a fixed number of bundles.

Definition 9.1 (Combinatorial Auction)**.** *In a* combinatorial auction *a set of m items M shall be allocated to a set of n bidders $\mathcal{N}$. The bidders have private values for bundles of items. The goal is to maximize social welfare.*

- *Feasible allocations:* $A = \{(S_1, \ldots, S_n) \subseteq M^n \mid S_i \cap S_j = \emptyset, i \neq j\}$
- *Valuation functions:* $v_i : 2^M \to \mathbb{R}_{\geq 0}$, $i \in \mathcal{N}$ *(private)*
- *Objective: Maximize social welfare* $\sum_{i=1}^n v_i(S_i)$

We will generally assume free disposal, *i.e.,* $v_i(S) \geq v_i(T)$ *for* $T \subseteq S$, *and that valuations are* normalized, *i.e.,* $v_i(\emptyset) = 0$.

Definition 9.2 (Multi-Minded Bidders)**.** *Bidders are called k-*minded *if, for every bidder $i \in \mathcal{N}$, there exists at most k bundles $S_i^1 \subseteq M, \ldots, S_i^k \subseteq M$ and values $v_i^1, \ldots, v_i^k \in \mathbb{R}_{\geq 0}$ such that*

$$v_i(T) = \begin{cases} \max_{j : S_i^j \subseteq T} v_i^j & \text{if } \exists\ j \text{ s.t. } S_i^j \subseteq T, \\ 0 & \text{otherwise.} \end{cases}$$

We can formulate the problem of maximizing social welfare in a combinatorial auction as an integer linear program.

$$\begin{aligned} \max \sum_{i \in \mathcal{N}, S \subseteq M} & x_{i,S} \cdot v_i(S) \\ \text{s.t.} \sum_{S \subseteq M} & x_{i,S} \leq 1 && \forall\ i \in \mathcal{N} \\ \sum_{i \in \mathcal{N}} \sum_{S \subseteq M : j \in S} & x_{i,S} \leq 1 && \forall j \in M \\ & x_{i,S} \in \{0, 1\} && \forall\ i \in \mathcal{N}, S \subseteq M \end{aligned}$$

We will also be interested in the linear program that results from relaxing the constraint $x_{i,S} \in \{0,1\}$ to $x_{i,S} \geq 0$.

Note that the integral problem is NP-hard as it encodes single-minded CAs as a special case. In fact, as we have seen last time, it is even hard to approximate to within a factor better than $\sqrt{m}$. The relaxation, in contrast, is solvable in time polynomial in the LP size, which is polynomial in the case of k-minded bidders.

2 Truthful Mechanism for Exact Welfare Maximization

One of the general positive results that we obtained for truthful single-parameter mechanism transfers to multi-parameter problems. Namely, exact welfare maximization can always be turned into a truthful mechanism.

Definition 9.3 (The VCG Mechanism)**.** *The VCG mechanism $M = (f, p)$ is defined as follows:*

- $f(v) \in \arg \max_{\textit{feasible } x} \sum_{i \in \mathcal{N}} x_{i,S} \cdot v_i(S)$ *(f maximizes social welfare)*
- $p_i(v) = \max_{\textit{feasible } x} \sum_{j \neq i} x_{i,S} \cdot v_i(S) - \sum_{j \neq i} v_i(f(v))$ *(bidder i pays his externality)*

Theorem 9.4 (Vickrey, Clarke, Groves 1961-1973)**.** *The VCG mechanism is dominant-strategy incentive compatible or truthful.*

Proof. Consider i, v_{-i}, v_i, v_i. We want to show that if bidder i's true value is v_i then his utility when he reports v_i is at least as high as his utility when he reports v_i'. Define

$$a := f(v_i, v_{-i}) \quad \text{and} \quad a' := f(v_i', v_{-i})$$

Bidder i's utility when he bids v_i resp. v_i' is

$$u_i(v_i, v_{-i}) = v_i(a) - \left(\max_{\text{feasible } x} \sum_{j \neq i} v_j(x) - \sum_{j \neq i} v_j(f(v_i, v_{-i})) \right) \text{ , and}$$

$$u_i(v_i', v_{-i}) = v_i(a') - \left(\max_{\text{feasible } x} \sum_{j \neq i} v_j(x) - \sum_{j \neq i} v_j(f(v_i', v_{-i})) \right) \text{ .}$$

Since f maximizes welfare

$$v_i(a) + \sum_{j \neq i} v_j(f(v_i, v_{-i})) \geq v_i(a') + \sum_{j \neq i} v_j(f(v_i', v_{-i})) \text{ .}$$

Subtracting $\max_{\text{feasible } x} \sum_{j \neq i} v_j(x)$ from both sides shows

$$u_i(v_i, v_{-i}) \geq u_i(v_i', v_{-i}) \text{ ,}$$

as desired. □

A closer inspection of the proof of the preceding theorem reveals that it applies equally well in the integral and fractional domain.

Is this our foothold for designing truthful polynomial-time mechanisms for the integral case? In general, breaking bundles into pieces to assign them fractionally will not be an option. But couldn't we solve the relaxation of the above integer program optimally and then use randomized rounding to obtain an integral solution, which is close to optimal?

3 Intermezzo: Oblivious Rounding is not Enough

Without having a concrete algorithm in mind, let us think about which guarantee we would need from a randomized mechanism to obtain truthfulness. To be able to formally reason about this, we should define what we mean by truthfulness.

Definition 9.5 (Truthfulness in Expectation)**.** *A randomized mechanism* $M = (f, p)$ *is* truthful in expectation *if for every bidder* i *and all* v_i, v_i', v_{-i} *we have*

$$\mathbb{E}[u_i(v_i, v_{-i})] \geq \mathbb{E}[u_i(v_i', v_{-i})] \ ,$$

where the expectation is over the randomness in the mechanism.

As a warm-up we will talk abstractly about randomized rounding schemes and their properties, and whether these ensure truthfulness. A very natural guarantee that is satisfied by many rounding schemes is the following.

Definition 9.6. *For* $\alpha \in [0, 1]$*, a rounding algorithm* R *for the welfare-maximization problem in combinatorial auctions is an* α-approximate oblivious rounding scheme *if for every fractional solution* y*, every bidder* i*, and every bundle* S

$$\Pr[i \textit{ gets } S] \geq \alpha \cdot y_{i,S} \ .$$

A key property of such rounding schemes is that they ensure an α-approximation to the fractional and hence integral optimum.

$$\mathbb{E}\left[\sum_{i=1}^{n} v_i(R(y))\right] = \sum_{i=1}^{n} \sum_{S \subseteq M} \Pr[i \text{ gets } S] \cdot v_i(S) \geq \alpha \cdot \sum_{i=1}^{n} \sum_{S \subseteq M} y_{i,S} \cdot v_i(S) \ .$$

Can we couple an α-approximate oblivious rounding scheme with VCG payments to obtain a truthful mechanism? It turns out that in general we cannot. Intuitively, the problem is the inequality in the above definition.

A possible way out, as we will see shortly, would be to require the following stronger property of a rounding scheme.

Definition 9.7. *For* $\alpha \in [0, 1]$*, and oblivious rounding algorithm* R *for the welfare-maximization problem in combinatorial auctions is an* α-scaling algorithm *if for every fractional solution* y*, every bidder* i*, and every bundle* S

$$\Pr[i \textit{ gets } S] = \alpha \cdot y_{i,S} \ .$$

The problem with this definition, however, is that it is not clear whether natural rounding schemes exist that have this property. In fact, most rounding schemes do *not*. What we will show next is that for certain problem domains there is a general construction for turning algorithms based on oblivious rounding into scaling algorithms.

4 The Lavi-Swamy Reduction

In fact, the construction that we will present will work for for a broader class of algorithms, namely those that have an oblivious-rounding type guarantee "in aggregate". That is, in sum over all bidders they guarantee that the integral solution is a c-approximation to the optimal fractional solution.

Definition 9.8. *Algorithm A* verifies a c-integrality gap *for the linear program for maximizing social welfare in a combinatorial auction if it receives as input real numbers $v_i(S)$ and outputs a feasible integral point $\tilde{x}$ such that*

$$\sum_{i,S} \tilde{x}_{i,S} \cdot v_i(S) \geq \frac{1}{c} \cdot \max_{feasible\ x} \sum_{i,S} x_{i,S} \cdot v_i(S) \ .$$

Intuitively, what the construction does is to decompose a fractional solution x, which has been scaled down by c, into a convex combination of integral solutions. That it is a convex combination will allow us to draw one of these integral solutions with probability proportional to the coefficient.

Lemma 9.9 (The Decomposition Lemma)**.** *Suppose algorithm A verifies a c-integrality gap for the linear program for maximizing social welfare in a combinatorial auction in polynomial time. Then for any feasible point x one can decompose x/c into a convex combination of integral feasible points in polynomial time.*

Before we prove this lemma in the next section, let us first see how we will use it. To this end we need to be able to refer to the convex combination whose existence and computability is established in the lemma. We will henceforth use $\{x^\ell\}_{\ell\in\mathcal{I}}$ to refer to all integral solutions. The proof will find $\{\lambda_\ell\}_{\ell\in\mathcal{I}}$ such that (1) for all $\ell \in \mathcal{I}$, $\lambda_\ell \geq 0$, (2) $\sum_{\ell\in\mathcal{I}} \lambda_\ell = 1$, and (3) $\sum_{\ell\in\mathcal{I}} \lambda_\ell \cdot x^\ell = x/c$.

Definition 9.10 (The Decomposition-Based Mechanism)**.**

1. *Compute an optimal fractional solution x^* and VCG prices $p_i^F(v)$.*
2. *Obtain a decomposition $x^*/c = \sum_{\ell\in\mathcal{I}} \lambda_\ell \cdot x^\ell$.*
3. *With probability λ_ℓ: Choose allocation x^ℓ and set prices $p_i^R(v) = [v_i(x^\ell)/v_i(x^*)] \cdot p_i^F(v)$.*

The idea behind the prices is to ensure that in expectation over the randomness in the mechanism each bidder pays the fractional VCG payments scaled down by c. We will verify this in the proof of the following theorem.

Theorem 9.11. *The decomposition-based mechanism is truthful in expectation and obtains a c-approximation to the social welfare.*

Proof. To see that the mechanism obtains the claimed approximation guarantee observe that the expected social welfare of the mechanism is exactly $1/c \cdot \sum_{i,S} x^*_{i,S} v_i(S)$, where $\sum_{i,S} x^*_{i,S} v_i(S)$ can only be higher than the best integral solution.

For truthfulness in expectation we will first argue that the expected payment of any bidder i is equal to the respective fractional payment divided by c:

$$\begin{aligned}\mathbb{E}_{\lambda_\ell}[p_i^R(v)] &= \sum_{\ell\in\mathcal{I}} \lambda_\ell \cdot [v_i(x^\ell)/v_i(x^*)] \cdot p_i^F(v) \\ &= [p_i^F(v)/v_i(x^*)] \cdot \sum_{\ell\in\mathcal{I}} \lambda_\ell \cdot v_i(x^\ell) \\ &= [p_i^F(v)/v_i(x^*)] \cdot v_i(x^*)/c = p_i^F(v)/c \ . \end{aligned} \tag{1}$$

Next fix any bidder i and valuations v_{-i} of the bidders other than i. Suppose that when i declares v_i the fractional optimum is x^* and when i declares v_i' it is z^*. Since fractional VCG is truthful we have

$$v_i(x^*) - p_i^F(v_i, v_{-i}) \geq v_i(z^*) - p_i^F(v_i', v_{-i}). \tag{2}$$

Now consider the decompositions $x^*/c = \sum_\ell \lambda_\ell^{x^*} \cdot x^\ell$ and $z^*/c = \sum_\ell \lambda_\ell^{z^*} \cdot x^\ell$. Dividing (2) by c and substituting the scaled down fractional VCG payments $p_i^F(v_i, v_{-i})/c$ and $p_i^F(v_i', v_{-i})/c$ using (1) yields

$$\left[\sum_{\ell \in \mathcal{I}} \lambda_\ell \cdot v_i(x^\ell)\right] - \mathbb{E}_{\lambda_\ell}[p_i^R(v_i, v_{-i})] \geq \left[\sum_{\ell \in \mathcal{I}} \lambda_\ell \cdot v_i(z^\ell)\right] - \mathbb{E}_{\lambda_\ell}[p_i^R(v_i', v_{-i})] \ .$$

The left-hand side is the expected utility for declaring v_i and the right-hand side the expected utility for declaring v_i'. This concludes the proof. □

5 Existence and Computation of the Decomposition

We still have to prove that from any given algorithm A that verifies a c-integrality gap we can compute a decomposition of any fractional solution x that we scale down by c. We will achieve this by writing down the problem of computing such a decomposition as a linear program and by arguing that we can solve it in polynomial time.

The key ingredient in our construction are the following linear program P and its dual linear program D. Consider any feasible fractional solution x we want to obtain a decomposition $x/c = \sum_{\ell \in \mathcal{I}} \lambda_\ell \cdot x^\ell$. Let $E = \{(i, S) \mid x_{i,S} > 0\}$. Recall that E has polynomial size, whenever we can compute x in polynomial time.

$$\begin{array}{llll}
\min & \sum_{\ell \in \mathcal{I}} \lambda_\ell & & (P) \\
s.t. & \sum_{\ell} \lambda_\ell \cdot x_{i,S}^\ell = \frac{x_{i,S}}{c} & \forall\ (i,S) \in E & \\
& \sum_{\ell} \lambda_\ell \geq 1 & & \\
& \lambda_\ell \geq 0 & \forall\ \ell \in \mathcal{I} &
\end{array}$$

$$\begin{array}{llll}
\max & \frac{1}{c} \cdot \sum_{(i,S) \in E} x_{i,S} \cdot v_i(S) + z & & (D) \\
s.t. & \sum_{(i,S) \in E} x_{i,S}^\ell \cdot v_i(S) + z \leq 1 & \forall\ \ell \in \mathcal{I} & \\
& z \geq 0 & & \\
& v_{i,S} \text{ unconstrained} & \forall\ (i,S) \in E &
\end{array}$$

The idea behind behind the primal program P is that it yields the desired decomposition if its optimal value is 1 because then $\sum_{\ell in \mathcal{I}} \lambda_\ell = 1$ and $\sum_\ell \lambda_\ell \cdot x_{i,S}^\ell = x_{i,S}/c$ for all $(i, S) \in E$ by the first constraint. The primal program P, however, has exponentially many variables, so we have to argue that it can be solved in polynomial time.

The dual linear program D will help us to show that (a) the optimal value of the *primal* program is 1 and (b) that a solution to the *primal* program can be computed efficiently. For this we will use what it is known as *strong duality*, namely that the optimal value of the primal and dual program coincide.

We will use this directly for (a). By showing that the optimal value of the dual program is 1 we can show that the optimal value of the primal program is 1. We will use it indirectly for (b). For this we will start from the observation that the dual program has polynomially many variables, but exponentially constraints. We will then argue that we can use the ellipsoid method to solve it in polynomial time. This argument involves the construction of a so-called *separation oracle* from the verifier A. This separation oracle tells us for each point whether it is feasible or not, and if it s not it gives us the violated constraint. Our argument will be completed by observing that by restricting ourselves to these constraints yields an equivalent dual program with polynomially many constraints. The dual of this dual program yields an alternative primal program with polynomially many constraints. The solution to this primal program is what we will ultimately use.

Lemma 9.12. *If $\frac{1}{c} \cdot \sum_{(i,S) \in E} x_{i,S} \cdot v_i(S) + z < 1$, one can use verifier A to find a violated constraint of D in polynomial time.*

Proof. Suppose $\frac{1}{c} \cdot \sum_{(i,S)\in E} x_{i,S} \cdot v_i(S) + z > 1$. Let verifier A receive v as input and suppose that the integral allocation that A outputs is x^ℓ. Then,

$$\sum_{(i,S)\in E} x^\ell_{i,S} \cdot v_i(S) \geq \frac{1}{c} \cdot \sum_{(i,S)\in E} x_{i,S} > 1 - z.$$

The first inequality holds because A verifies a c-integrality gap and the second inequality is a consequence of our assumption. Thus the first constraint in the dual linear program is violated for the integral solution x^ℓ found by A. □

Lemma 9.13. *The optimum of the dual linear program D and hence the primal program P is* 1*, and the decomposition* $x/c = \sum_{\ell\in\mathcal{I}} \lambda_\ell \cdot x^\ell$ *can be computed in polynomial time.*

Proof. We first show that the optimum of the dual linear program is 1. From Lemma 9.12 we know that whenever $\frac{1}{c} \cdot \sum_{(i,S)\in E} x_{i,S} \cdot v_i(S) + z > 1$ we can find a violated constraint, so clearly the optimum can be at most one. On the other hand a feasible solution to the dual program is given by $z = 1$ and $v_i(S) = 0$ for all $(i, S) \in E$. So the optimum is at least 1. Together this shows that the optimum is exactly 1.

The preceding discussion shows that we can add the inequality $\frac{1}{c} \cdot \sum_{(i,S)\in E} x_{i,S} \cdot v_i(S) + z \geq 1$ to the dual program D without altering anything. We can then run the ellipsoid method on the the modified dual program to identify an equivalent dual program D' with polynomially many constraints. These inequalities will be the violated inequalities returned by the separation oracle that we construct below. Taking the dual of this alternate dual program D' yields a alternate primal program P' with a polynomial number of variables and constraints, which we can solve to obtain the λ_ℓ's that sum to one.

Recall that a separation oracle has to, given a point v, z, decide whether the point is in the polytope of feasible solutions and if not return a violated inequality. We can therefore construct a separation oracle for our problem as follows: If $\frac{1}{c} \cdot \sum_{(i,S)\in E} x_{i,S} \cdot v_i(S) + z > 1$ we can use verifier A as described in Lemma 9.12 to find a violated inequality. Otherwise, we can simply use the half space $\frac{1}{c} \cdot \sum_{(i,S)\in E} x_{i,S} \cdot v_i(S) + z \geq 1$. □

6 Truthful Mechanism for Multi-Minded CA

As a final step let us return to our problem of finding a truthful mechanism for multi-minded CAs. To use Lavi and Swamy's construction we need to construct a verifier A for the k-minded combinatorial auction problem. We will use the following lemma, which reduces the task of finding such a verifier to the task of finding an approximation algorithm for the integral problem that approximates the fractional optimum.

Lemma 9.14 (cf. AGT Book, Claim 12.19)**.** *Given a c-approximation for general CAs* A' *where the approximation guarantee is with respect to the* fractional *optimum, one can obtain an algorithm A that verifies a c-integrality gap for the linear program for maximizing social welfare in a combinatorial auction, with a polynomial time overhead on top of A.*

Now if we go back to last week's greedy algorithms and the proof that it yields a $\sqrt{m}$ approximation, it is not difficult to see that it in fact approximates the fractional optimum. In fact, this is also true for the natural generalization of this algorithm to k-minded bidders.

Theorem 9.15 (cf. AGT Book, Lemma 12.21)**.** *The greedy mechanism achieves at least a* $1/\sqrt{2m}$ *approximation to the optimal fractional welfare in general combinatorial auctions.*

This is very remarkable because one can show that the greedy mechanism does *not* satisfy the requirement for truthfulness alluded to at the beginning of this chapter, "cyclic monotonicity", so no payment rule makes it truthful!

Recommended Literature

- R. Lavi and C. Swamy. Truthful and Near-Optimal Mechanism Design via Linear Programming. Journal of the ACM, December 2011. (The original publication)
- Chapter 12.3 of the AGT Book (Summary of the results presented in the original publication with less technical details)
- Tim Roughgarden's Lecture Notes for Course CS364B Frontiers in Mechanism Design, Lecture 8 and Lecture 9