Randomized Mechanism Design: Approximation and Online Algorithms

Part 1: Introduction to Mechanism Design and Multi-unit Auctions

Berthold Vöcking
RWTH Aachen University

August 2012

Multi-unit Auctions

m identical items shall be allocated to n bidders with private valuations such that social welfare is maximized

Definitions:

- feasible allocations: $A = \{(s_1, \ldots, s_n) \in \mathbb{N}^n \mid \sum_i s_i \leq m\}$
- valuation functions: $v_i : \{0, \ldots, m\} \rightarrow \mathbb{R}_{\geq 0}$, $i \in [n]$
- social welfare: $\sum_{i=1}^{n} v_i(s_i)$

Assumptions:

- value queries: What is the valuation of bidder i for k items?
- free disposal: valuations are non-decreasing
- normalization: $v_i(0) = 0$

It is common to assume that the *input length* is $n + \log m$.
We seek for a poly-time "incentive compatible" mechanism.

Lower bound for exact algorithms

Consider the following valuation functions for 2 players:

- Player 1 has valuations $v_1(i) = i$, for $i \in \{0, \ldots, m\}$.
- Player 2 has valuations

$$v_2(i) = \begin{cases} i & \text{for } i \neq k \\ i+1 & \text{for } i = k \end{cases}$$

for some $k \in \{0, \ldots, m\}$.

The unique optimal allocation is $s_1 = m - k$, $s_2 = k$.

Any (randomized) algorithm needs $\Omega(m)$ queries for finding the index k.

Ignoring the aspect of truthfulness ...

A "non-truthful" approximation scheme

- Round down valuations to the nearest power of $(1+\epsilon)$ and consider only the *breakpoints*, i.e., valuations at which the rounded valuations increase.
- The number of breakpoints per bidder is $O(1/\epsilon \cdot \log m)$.
- Use FPTAS for the *multiple-choice knapsack problem* with objects defined by the breakpoints.

Incentive compatibility

Let V be the set of all valuations, and A the set of allocations.

A *mechanism* is a pair (f, p) where

- $f : V^n \to A$ is called *social choice-function*, and
- $p : V^n \to \mathbb{R}^n$ is called a *payment scheme*.

If (f, p) is fixed, then the *utility* of bidder i for valuations $v \in V^n$ is

$$u_i(v) = v_i(f(v)) - p_i(v) \ .$$

Definition

A mechanism (f, p) is *truthful* if for all i, all v_i, $v_i' \in V$ and all $v_{-i} \in V^{n-1}$, we have that $u_i(v_i, v_{-i}) \geq u_i(v_i', v_{-i})$.

In words: A mechanism is called *truthful* if truth-telling is a dominant strategy for every bidder.

Incentive compatibility

Randomized notions of truthfulness:

- *Truthfulness in expectation:* every bidder maximizes his expected utility by bidding truthfully, that is, for all i, all v_i, $v_i' \in V$ and all $v_{-i} \in V^{n-1}$, we have that

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

- *Universal truthfulness:* a universally truthful mechanism is defined by a probability distribution over deterministically truthful mechanisms

Overview

0: Introduction

1: VCG-based mechanisms

– maximal in range (deterministically truthful) –

– maximal in distributional range (truthful in expectation) –

2: A universally truthful approximation scheme

– polynomial query complexity –

– polynomial running time –

Overview

0: Introduction

1: VCG-based mechanisms

– maximal in range (deterministically truthful) –

– maximal in distributional range (truthful in expectation) –

2: A universally truthful approximation scheme

– polynomial query complexity –

– polynomial running time –

VCG-based mechanisms

Vickrey-Clarke-Groves (VCG) mechanism

- Compute an optimal allocation $f(v) = s_1, \ldots, s_n$.
- Set payments by $p_i(v) = \max_{t \in A} \left(\sum_{j \neq i} v_j(t_j) \right) - \sum_{j \neq i} v_j(s_j)$.

VCG is truthful since, for every bidder i,

$$\underbrace{v_i(s_i) - p_i}_{\text{utility of } i} = \underbrace{\sum_{j \in [n]} v_j(s_j)}_{\text{social welfare}} - \underbrace{\max_{t \in A} \left(\sum_{j \neq i} v_j(t_j) \right)}_{\text{independent of } v_i}$$

That is, maximizing social welfare maximizes the bidder's utility (provided that the bidder reports her true valuation).

VCG-based mechanisms

Definition (affine maximizer)

A social choice function (allocation algorithm) f is an **affine maximizer** if there exists a set of allocations $A' \subseteq A$, a constant $\alpha_i \geq 0$, for $i \in \{1, \ldots, n\}$, and a constant $\beta_s \in \mathbb{R}$, for every $s \in A'$, such that

$$f(v) = \mathrm{argmax}_{s \in A'} \left(\sum_{i=1}^{n} \alpha_i v_i(s_i) + \beta_s \right) .$$

VCG-based mechanisms achieve truthfulness by combining an affine maximizer f with generalized VCG payments.

Maximal In Range (MIR)

A mechanism is called *MIR* if it maximizes over a subrange $A' \subset A$.

MIR 1/2-approximation algorithm [Dobzinski and Nisan, 2007]

Split the items into

- n^2 equally-sized bundles of size $b = \lfloor \frac{m}{n^2} \rfloor$ and
- a single extra bundle of size $r = m - n^2 b$.

Optimally allocate these whole bundles among the n bidders.

Observation

An optimal bundle allocation can be found in time polynomial in n using dynamic programming.

MIR 1/2-approximation

Lemma

Let $(a_1, \ldots, a_n)$ be an optimal bundle allocation and $(o_1, \ldots, o_n)$ an optimal unrestricted allocation. Then $\sum_i v_i(a_i) \geq \frac{1}{2} \sum_i v_i(o_i)$.

Proof:

- W.l.o.g., $\sum_i o_i = m$.
- There exists a bidder i with $o_i \geq \frac{m}{n}$.
- If $v_i(o_i) \geq \frac{1}{2} \sum_j v(o_j)$ then assigning all items to bidder i gives a $\frac{1}{2}$-approximation.
- Otherwise, rounding up all bidders $j \neq i$ to full bundles of size b gives a $\frac{1}{2}$-approximation. $\square$

Limitations of VCG-based mechanisms

$A' \subseteq A$ is called a *true sub-range* if there exists $s \in A$ with $\sum_i s_i = m$ and $s \notin A'$.

Theorem

There does not exist a MIR algorithm that optimizes over a true subrange A' and achieves an approximation factor better than $1/2$.

Proof:

- Suppose there are only two bidders.
- Let (s_1, s_2) be an allocation with $s_1 + s_2 = m$ and $(s_1, s_2) \notin A'$.
- Suppose $v_1(k) = 1$, for $k \geq s_1$, and $v_1(k) = 0$, otherwise.
- Suppose $v_2(k) = 1$, for $k \geq s_2$, and $v_2(k) = 0$, otherwise.
- The optimal allocation over A has a value of 2 while the optimal allocation over A' has a value of 1. $\square$

Limitations of VCG-based mechanisms

Corollary

Any deterministic VCG-based mechanism with an approximation factor better than $1/2$ needs an exponential number of queries.

Some more results about deterministic mechanisms

Restricted valuations:

- FPTAS for single-minded valuations using monotonicity [Briest, Krysta, V., 2005]
- PTAS for k-minded valuations based on the MIR approach [Dobzinski, Nisan, 2007]
- There does not exist a MIR-FPTAS for k-minded valuations [Dobzinski, Nisan, 2007]

Multi-dimensional valuations:

- Any "scalable" deterministically truthful mechanism that guarantees a c-approximation, for $c > \frac{1}{2}$, needs to make an exponential number of queries. [Dobzinski, Nisan, 2011]

Overview

0: Introduction

1: VCG-based mechanisms

– maximal in range (deterministically truthful) –

– maximal in distributional range (truthful in expectation) –

2: A universally truthful approximation scheme

– polynomial query complexity –

– polynomial running time –

Maximal In Distributional Range (MIDR)

Let $\mathcal{D}(A)$ denote a set of probability distributions $D : A \to [0, 1]$.

A mechanism that chooses a probability distribution from $\mathcal{D}(A)$ such that the expected social welfare is maximized is called Maximal In Distributionan Range (MIDR).

Simplified MIDR mechanism inspired by [Dobzinski,Dughmi, 2009]

For an integer $t \geq 1$, let $q(t)$ denote the number of trailing 0's in the binary representation, e.g., $q(101000) = 3$.

Obviously, $q(t) \leq \lfloor \log m \rfloor$, for $1 \leq t \leq m$. Let $q(0) = \lfloor \log m \rfloor + 1$.

Probabilistic allocations

For $(s_1, \ldots, s_n) \in A$, let $[s_1, \ldots, s_n]_{\mathcal{D}}$ denote the following distribution: Bidder i gets allocated s_i items with probability

$$(1 - \epsilon)^{q(0) - q(s_i)} \quad ,$$

for some given $\epsilon \in [0, 1]$; and 0 items, otherwise.

Simplified MIDR mechanism

The "simplified MIDR mechanism" ...

... outputs a probabilistic allocation $[s_1, \ldots, s_n]_{\mathcal{D}}$ that maximizes expected social welfare among all $(s_1, \ldots, s_n) \in A$, and uses VCG prices over this range.

Perturbed valuations

The expected value of $[s_1, \ldots, s_n]_{\mathcal{D}}$ for bidder i is thus

$$v'(s_i) = v_i(s_i) \cdot (1 - \epsilon)^{q(0) - q(s_i)} \quad .$$

$* * *$ Maximizing wrt to v' yields a $(1-\epsilon)^{q(0)}$-approximation wrt to v $* * *$

Simplified MIDR mechanism

Lemma

The optimal allocation wrt v' can be found with a number of queries bounded polynomially in $\log m$ and $1/\epsilon$ per bidder.

Proof:

- Consider bidder i. Let $V_i = (v_i(0), v_i(1), \ldots, v_i(m))$.
- Partition V_i into subsequences V_i^k, for $0 \leq k \leq q(0)$, such that V_i^k contains the valuations $v_i(t)$ with $q(t) = k$.
- The *k-breakpoints* of bidder i are defined to be those entries from V_i^k at which the value increases by a factor of at least $(1-\epsilon)^{-1}$ in comparison to the preceding k-breakpoint.
- $\#\text{breakpoints} = \text{poly}(n, \log m, 1/\epsilon)$
- Breakpoints can be found efficiently using binary search.

Simplified MIDR mechanism

Lemma

Let $o_1, \ldots, o_n$ denote an optimal allocation wrt v'.
For every $i \in [n]$, o_i is a $q(o_i)$-breakpoint of bidder i.

Proof (sketch):

Simplified MIDR mechanism

Lemma

Let $o_1, \ldots, o_n$ denote an optimal allocation wrt v'.
For every $i \in [n]$, o_i is a $q(o_i)$-breakpoint of bidder i.

Proof (sketch):

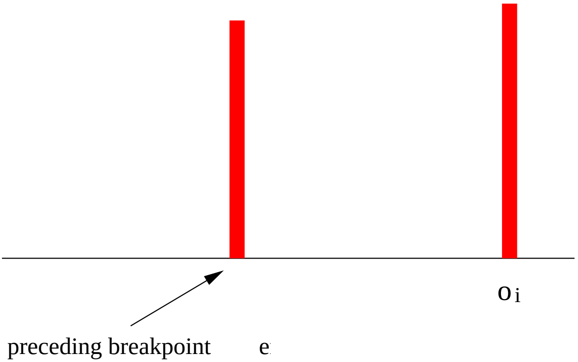

Simplified MIDR mechanism

Lemma

Let $o_1, \ldots, o_n$ denote an optimal allocation wrt v'.
For every $i \in [n]$, o_i is a $q(o_i)$-breakpoint of bidder i.

Proof (sketch):

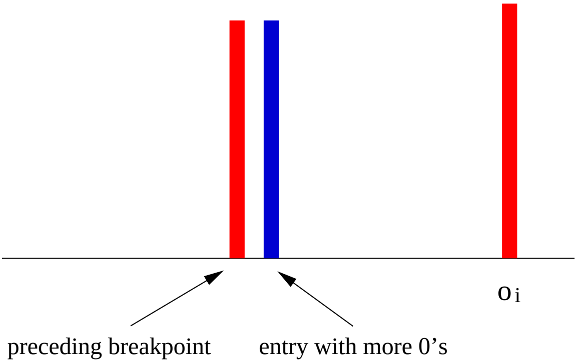

Simplified MIDR mechanism

Lemma

Let $o_1, \ldots, o_n$ denote an optimal allocation wrt v'.
For every $i \in [n]$, o_i is a $q(o_i)$-breakpoint of bidder i.

Proof (sketch):

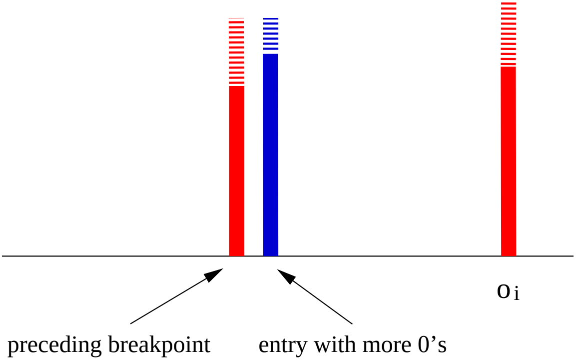

Simplified MIDR mechanism

Lemma

Let $o_1, \ldots, o_n$ denote an optimal allocation wrt v'.
For every $i \in [n]$, o_i is a $q(o_i)$-breakpoint of bidder i.

Proof (sketch):

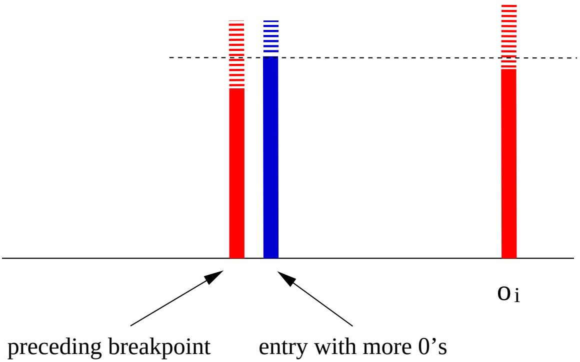

$\square$

The power of randomized mechanism design

[Dobzinski, Dughmi, 2009]

Approximation scheme

There is a truthful-in-expectation FPTAS for multi-unit auctions.

Separation result

A certain (technical) variant of multi-unit auctions

- admits a truthful-in-expectation FPTAS, but
- does not admit a universally truthful algorithm achieving an approximation factor better than 2 with a sub-exponential number of queries.

Overview

0: Introduction

1: VCG-based mechanisms

– maximal in range (deterministically truthful) –

– maximal in distributional range (truthful in expectation) –

2: A universally truthful approximation scheme

– polynomial query complexity –

– polynomial running time –

A universally truthful approximation scheme

Theorem (V., SODA 2012)

There exists a universally truthful polynomial-time approximation scheme for multi-unit auctions.

The approximation scheme corresponds is randomized PTAS. The expected social welfare is lower bounded by $(1 - \epsilon)$ of the optimal social welfare.

We first present a simplified approximation scheme with polynomially bounded query complexity.

Idea: apply small additive perturbations to the valuations

Δ-perturbed maximizer

Let $\Delta > 0$. For $1 \leq i \leq n$, $0 \leq j \leq m$, set

$$v'_i(j) = v_i(j) + q(j)\Delta$$

with $q(j)$ denoting the number trailing 0's in the binary representation of j (as defined before).

Choose an allocation $s \in A$ maximizing $v'(s) = \sum_{i=1}^{n} v'_i(s_i)$.

Claim:

If Δ is set equal to $\epsilon v_{\max}/(n \log m)$ then

- the additive error due to perturbation is $O(\epsilon OPT)$, and
- the allocation maximizing v' can be computed with poly$(\log m, n, {}^1\!/\!_\epsilon)$ queries.

The dilemma

On the one hand:

In order to get a polynomial time approximation scheme, Δ needs to be chosen in a way depending on the valuations.

On the other hand:

In order to obtain truthfulness, Δ must be chosen independently of the valuations.

We introduce a subjective variant of VCG in order to overcome this problem.

Description of the mechanism

Let $L : \mathbb{R}_{>0} \to \mathbb{R}_{>0} \cup \{\perp\}$ denote a suitable function with $L(x) \leq x$, unless $L(x) = \perp$, called drop-out consensus function.

For every bidder i, compute s_i as follows:

- Let $v_{\max}^{(-i)}$ denote the maximum valuation of the other bidders.
- Compute a lower bound $L_i = L(v_{\max}^{(-i)})$.
- If $L_i = \perp$ then the algorithm sets $s_i = 0$. ("player i drops out")
- Otherwise, compute an allocation $s^{(i)} \in \{0, \ldots, m\}^n$ by calling the Δ_i-perturbed maximizer with $\Delta_i = L_i/N$ (with $N = (\lceil \log m \rceil + 1)n/\epsilon$) and set $s_i = s_i^{(i)}$.

Observe that there are only two different outcomes of $v_{\max}^{(-i)}$.

Drop-out consensus

Ideally, we seek for a consensus function $\ell : \mathbb{R} \to \mathbb{R}$ with the following properties:

- For any $a \in \mathbb{R}$, $\ell(a) \in [a - \frac{1}{\epsilon}, a]$.
- For any $a, b \in \mathbb{R}$ with $|a - b| \leq 1$, $\ell(a) = \ell(b)$.

Exercise: Show that such a consensus function does not exist.

We will present a randomized consensus procedure $\ell : \mathbb{R} \to \mathbb{R} \cup \{\perp\}$ with the following properties:

- For any $v \in \mathbb{R}$, $\ell(a) \in [a - \frac{1}{\epsilon}, a]$, unless $\ell(a) = \perp$.
- For any $a, b \in \mathbb{R}$ with $|a - b| \leq 1$, $\ell(a) = \ell(b)$, unless $\ell(a) = \perp$ or $\ell(b) = \perp$.

In particular, $\mathbf{Pr}[\ell(a) = \perp] \leq \epsilon$, for any $a \in \mathbb{R}$.

Drop-out consensus

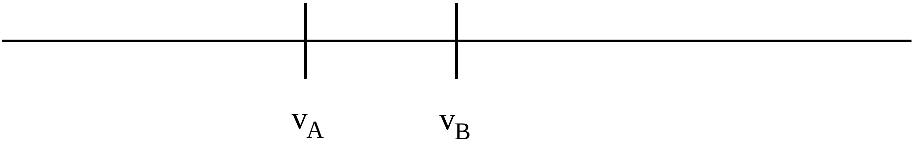

Drop-out consensus

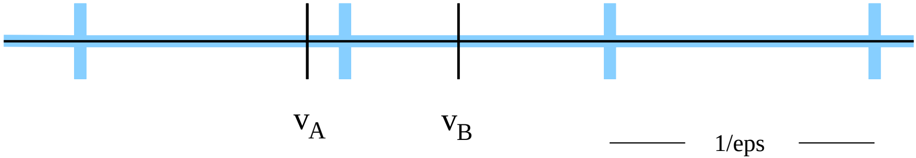

Drop-out consensus

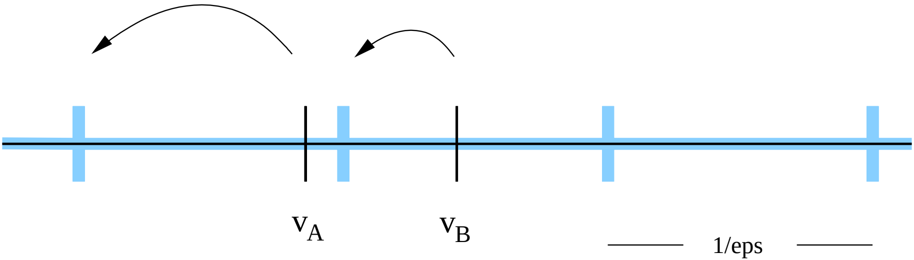

Drop-out consensus

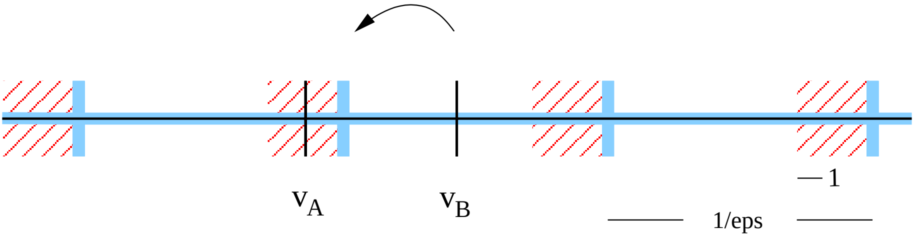

Drop-out consensus

More formally, we define a function $\ell : [0,1] \times \mathbb{R} \to \mathbb{R} \cup [\perp]$ where the first parameter is picked uniformly at random.

Properties of ℓ

1. For every $a > 0$ and τ chosen uniformly at random from $[0,1]$, $\mathbf{Pr}\left[\ell(\tau, a) = \perp\right] = \epsilon$.
2. For every $a \in \mathbb{R}$ and $\tau \in [0,1]$ with $\ell(\tau, a) \neq \perp$, it holds $\ell(\tau, a) \in [a - 1/\epsilon, a]$.
3. For any numbers $a_2 > a_1$, $\tau \in [0,1]$ with $\ell(\tau, a_1) \neq \perp$ and $\ell(\tau, a_2) \neq \perp$, it holds:

   $$\text{If } \ell(\tau, a_1) \neq \ell(\tau, a_2) \text{ then } a_1 \leq \ell(\tau, a_2) - 1.$$

Drop-out consensus

We use the drop-out consensus procedure on an exponential scale. That is $L(\tau, v) = N^{\ell(\tau,v)}$, for $N = (\lceil \log m \rceil + 1)n/\epsilon$.

$L : [0,1] \times \mathbb{R}_{>0} \to \mathbb{R}_{>0} \cup [\perp]$ satisfies

1. For every $a > 0$ and τ chosen uniformly at random from $[0,1]$, $\mathbf{Pr}\left[L(\tau, a) = \perp\right] = \epsilon$.
2. For every $a > 0$ and $\tau \in [0,1]$ with $L(\tau, a) \neq \perp$, it holds $L(\tau, a) \leq a$ and $L(\tau, a) \geq aN^{-1/\epsilon}$.
3. For any numbers $a_2 > a_1 > 0$, $\tau \in [0,1]$ with $L(\tau, a_1) \neq \perp$ and $L(\tau, a_2) \neq \perp$, it holds:

$$\text{If } L(\tau, a_1) \neq L(\tau, a_2) \text{ then } a_1 \leq L(\tau, a_2)/N.$$

Feasibility of the mechanism

Let v_{1st} and v_{2nd} denote the "largest" and the "second largest" valuation, respectively.

Feasibility analysis

1. If $L(v_{1st}) = \perp$ or $L(v_{2nd}) = \perp$ then the solution is feasible as the bidder with the largest bid or all other bidders drop out. Otherwise:
2. If $L(v_{1st}) = L(v_{2nd})$ then all players call the same perturbed maximizer and, hence, the solution is feasible.
3. If $L(v_{1st}) \neq L(v_{2nd})$ then $v_{2nd} < L(v_{1st})/N = \Delta$. This implies
   - $\Delta(q(0) - q(k)) > v_{2nd}$, for $k \in \{1, \ldots, m\}$, so that
   - the mechanism sets $s_i = 0$ for all bidders except the bidder with the maximum bid.

Truthfulness of the mechanism

Composition of quasi-linear maximizers

Let $f^{(1)}, \ldots, f^{(n)}$, $f^{(i)} : V^n \to A$ be a collection of n functions s.t.

$$f^{(i)}(v) = \underset{s \in A}{\operatorname{argmax}}(v_i(s) + g_s^{(i)}(v_{-i}))$$

with $g_s^{(i)} : V^{n-1} \to \mathbb{R}$ being an arbitrary function.

The function $f : V^n \to \{0, \ldots, m\}^n$ defined by $f(v)_i = f^{(i)}(v)_i$ is called a *composition of quasi-linear maximizers*.

This composition is called *feasible* if $f(V^n) \subseteq A$.

Truthfulness – Subjective VCG

The mechanism calls an affine maximizer for each bidder i.

Let $f^{(i)}$ denote the maximizer of bidder i. This way, the social choice function f of the mechanism is a composition of quasi-linear maximizers $f^{(1)}, \ldots, f^{(n)}$.

For every bidder i, the mechanism uses VCG prices wrt to f_i.

Lemma

The mechanism is truthful.

Proof:

- For every bidder, the mechanism solves an optimization problem that maximizes the bidder's utility (like VCG).
- Hence, it is a dominant strategy to report valuations truthfully.

$\square$

Overview

0: Introduction

1: VCG-based mechanisms

– maximal in range (deterministically truthful) –

– maximal in distributional range (truthful in expectation) –

2: A universally truthful approximation scheme

– polynomial query complexity –

– polynomial running time –

Idea: use two kinds of perturbations

For $1 \leq i \leq n$, $0 \leq j \leq m$, set

$$v_i'(j) = v_i(j) + \beta_i^j \, \Delta$$

with $\beta_i^j = 2q(j) + x_i^j$, where

a) $q(j)$ denotes the number trailing 0's (as before), and
b) x_i^j is a random variable chosen independently, uniformly at random from $[0, 1]$.

Perturbations of type (a) yield that the number of "breakpoints" per bidders is bounded polynomially (as before).

Perturbations of type (b) yield that the number of "Pareto-optimal allocations" is bounded polynomially.

Pareto-optimal allocations

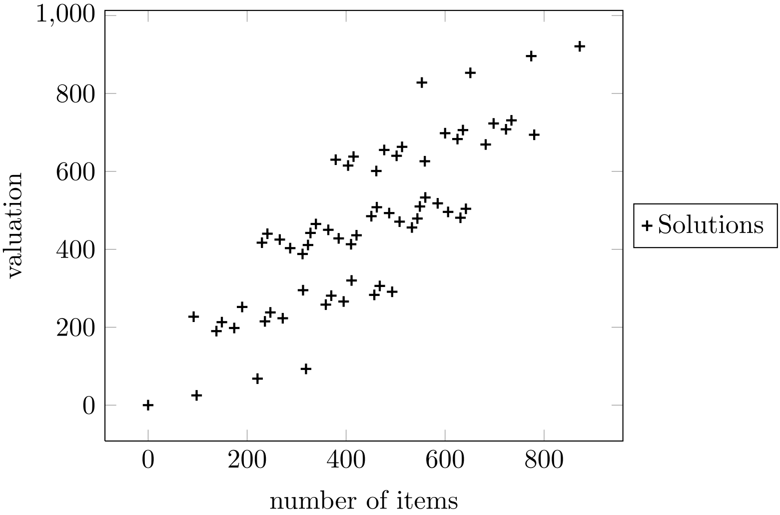

Pareto-optimal allocations

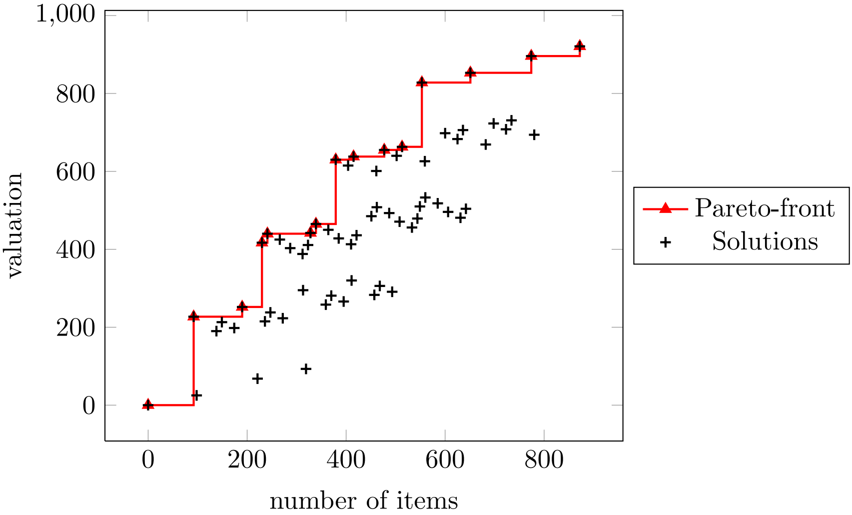

Pareto-optimal allocations

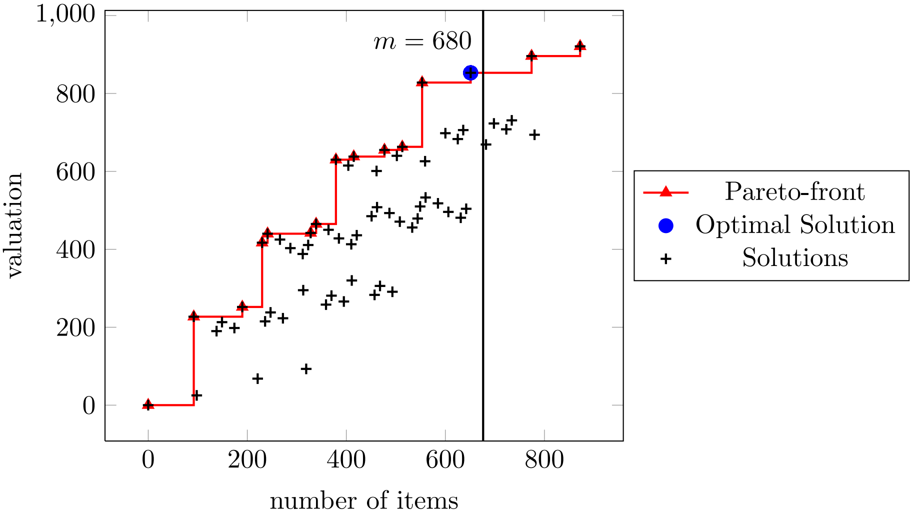

Running time analysis

The Pareto-front can be enumerated in time $O(b\sum_{i=1}^{n-1} k_i)$ where k_i denotes the number of Pareto-optimal solutions restricted to bidders 1 to i.

Smoothed analysis of the knapsack problem [Beier, V., 2003]

Suppose object values are chosen from $[0,1]$ by an adversary and then these values are perturbed by adding numbers that are picked uniformly at random from $[0,\sigma]$. $\mathbf{E}[k_i] = O(b^2 i^2/\sigma)$.

The expected running time is thus $O(b^3 n^3/\sigma)$.
In our context,

$$\begin{aligned} b &= \#\text{ number of breakpoints} \\ \sigma &= \Delta/v_{2nd} \end{aligned}$$

As $\Delta \geq v_{2nd}/N^{1/\epsilon+1}$ and $b = \text{poly}(\log m, n, {}^1\!/\!_\epsilon)$, the expected running time is polynomially bounded. $\square$

Recommended Reading

- Chapter 9 in *"Algorithmic Game Theory," Nisan N., Roughgarden T., Tardos E., Vazirani V. (Eds.)*, 2007.
- Shahar Dobzinski and Shaddin Dughmi. On the power of randomization in algorithmic mechanism design. FOCS 2009.
- Berthold Vöcking. A universally-truthful approximation scheme for multi-unit auctions. SODA 2012.
- Patrick Briest, Piotr Krysta, and Berthold Vöcking. Approximation techniques for utilitarian mechanism design. SIAM J. Comput. 40(6), 2011.
- Shahar Dobzinski and Noam Nisan. Mechanisms for multi-unit auctions. EC 2007.
- Shaddin Dughmi and Tim Roughgarden. Black-box randomized reductions in algorithmic mechanism design. FOCS 2010.
- Shahar Dobzinski and Noam Nisan. Multi-unit auctions: beyond Roberts. EC 2011.