EECS 495: Combinatorial Optimization Lecture Manolis, Nima
Mechanism Design with Rounding

Motivation

- Make a social choice that (approximately) maximizes the social welfare subject to the economic constraints of truthfulness
- When optimizing the social welfare is NP-hard and the standard solution, the VCG mechanism, cannot be computed efficiently find truthful approximation algorithms

Mechanism Design Basics

- A set of possible outcomes $\mathcal{A}$ and a set of bidders $\mathcal{N}$.
- For each $i \in \mathcal{N}$ a private function $v_i : \mathcal{A} \to \mathbb{R}_0^+$
- A Mechanism elicits bid b_i from each $i \in \mathcal{N}$ and outputs the set $f(b_1, \ldots, b_n) \in \mathcal{A}$ and and charges player $I \in \mathcal{N}$: $p_i(b_1, \ldots, b_n)$
- We relax f and p to be random variables over the possible outcomes and the $\mathbb{R}^n$ payment vectors
- The expected utility of a player $i \in \mathcal{N}$ is given by $E[v_i(f(b_i, \ldots, b_n)) - p_i(b_1, \ldots, b_n)]$
- A mechanism is truthful in expectation if for every $b_i \neq v_i$ and $v_{-i} \in \mathbb{R}^{n-1}$: $E[v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i})] \geq E[v_i(f(b_i, v_{-i})) - p_i(b_i, v_{-i})]$

Problem Formulation

Multi Unit Combinatorial Auctions

- A set $[m]$ of items
- A number B representing how many copies we have of each item
- A set $[n]$ of players
- Player i has valuation function $v_i : 2^{[m]} \to \Re^+$, $v_i(\emptyset) = 0$, $v_i(A) \leq v_i(B)$ for $A \subseteq B$ (monotone)
- Feasible solution: allocation $(S_1, \ldots, S_n)$ such that each item is allocated to at most B players (these sets must be disjoint for $B = 1$)
- i's value for $(S_1, \ldots, S_n)$ is $v_i(S_i)$
- Goal: Maximize Social Welfare: $\sum_i v_i(S_i)$

Lotteries and Value Oracles

- Value oracle takes set and returns value
- Analogous oracle for lotteries
- For $x \in [0,1]^m$, D_x define

$$F_v(x) = E_{S \sim D_x}[v(S)] = \sum_S v(S) \prod_{j \in S} x_j \prod_{j \notin S} (1 - x_j)$$

- That is the value of the lottery can be computed as the expected value of the outcome w.r.t. the distribution of the lottery

A general technique for constructing truthful approximation algorithms

Deterministic Support Mechanism

- Any randomized mechanism can be viewed as a deterministic mechanism with a different outcome space: the set of all lotteries over the original outcomes, the deterministic support mechanism

- Given any truthful in expectation mechanism the corresponding deterministic support mechanism that charges the expected payment vector is truthful

- The reverse is also true. A similar rule Clark-Groves could be used to ensure individual rationality in expectation.

- These transformations preserve the social welfare at any input

The new framework

- The packing property of a problem is that its corresponding natural LP satisfies the following fact: if vector x is a solution to the problem then any $0 \leq x' \leq x$ is also a solution

- The mutli-unit combinatorial auction problem has the packing property:

- The LP formulation we will use is the following

$$\max \sum_{i,S \neq \emptyset} v_i(S) x_{i,S}$$

$$\text{subject to } \sum_{S \neq \emptyset} x_{i,S} \leq 1 \text{ for each player } i$$

$$\sum_i \sum_{S : j \in S} x_{i,S} \leq 1 \text{ for each item } j$$

$$x_{i,S} \geq 0 \text{ for each } i, S$$

Theorem 0.1 *Given any α-approximation algorithm for a problem with the "packing property" that proves also an integrality gap of at most α there is a randomized truthful in expectation α-approximation algorithm*

The mechanism is defined as follows:

1. Compute the fractional optimal solution x^*

2. Scale the solution by α: $x' = x^*/\alpha$

3. Express x' as a convex combination of integer solutions where there are polynomially many non-zero weights

4. Convert the deterministic support allocation into its corresponding randomized version.

Claim: This algorithm returns a truthful in expectation randomized mechanism

Proof:

- Consider the mechanism that did not scale by α and could probably express the best fractional solution as a convex combination of exponentially many integral ones

- This coincides with the VCG mechanism on lotteries therefore its randomization is also truthful

- The randomization obtained in the case where we scale by α has simply scaled the utility of each agent at every input, hence, the truthfulness equation is satisfied

$\square$ **Claim:** We can express the scaled allocation as a convex combination of polynomially many integral solutions

Proof:

- Let $\mathcal{I}$ be the set of integral solutions

- Let x^* be the fractional optimal solution

- Let $E = \{(i, S) \mid x^*_{i,S} > 0\}$

- we define the following primal P and its dual D

Primal

$\min \sum_{l \in \mathcal{I}} \lambda_l$

s.t. $\sum_l \lambda_l x^l_{i,S} = x^*_{i,S}/\alpha$ for all $(i, S) \in E$

$\sum_{l \in \mathcal{I}} \lambda_l \geq 1$

$\lambda_l \geq 0$ for all $l \in \mathcal{I}$

Dual

$\max \frac{1}{\alpha} \sum_{(i,S) \in E} x^*_{i,S} w_{i,S} + z$

s.t. $\sum_{(i,S) \in E} x^l_{i,S} w_{i,S} + z \leq 1$ for all $l \in \mathcal{I}$

$z \geq 0$

$\lambda_l \geq 0$ for all $l \in \mathcal{I}$

- If the optimal solution to the primal is 1 then the variables λ correspond to the weights of the convex combination

- Primal P has exponentially many variables

- It has polynomially many constrains; Since x^* is an extreme point of the original problem, it has at most $m + n$ non zero variables, i.e., $|E| \leq m + n$.

- Equivalently, dual D has exponentially many constraints but polynomially many variables

- We will show how we can solve the Dual restricting the feasible space only to the optimal solutions since then we have a separation oracle

Claim: Let $w^+_{i,S} = \max(w_{i,S}, 0)$. Given any integral x to the original problem one can find integral solution x^l such that

$$\sum_{(i,S) \in E} x^l_{i,S} w_{i,S} = \sum_{(i,S) \in E} \hat{x}_{i,S} w^+_{i,S}$$

Proof:

1. if $w_{i,S} > 0$ set $x^l_{i,S} = x_{i,S}$

2. Otherwise set $x^l_{i,S} = 0$.

$\square$

Claim: For any weight vector w we can compute in polynomial time an integral solution such that

$$\sum_{(i,S) \in E} x^l_{i,S} w_{i,S} \geq \frac{1}{\alpha} \max_{\text{feasible } x} \sum_{(i,S) \in E} x_{i,S} w_{i,S}$$

Proof:

- If some components of w are negative take w^+
- We can monotonize

$$w'_{i,S} = \max_{T \subseteq S:(i,T)\in E} w^+_{i,T}$$

- Apply approximation algorithm with valuations w' and get solution y
- By construction $w'_{i,S} \geq w_{i,S}$ for all $(i,S) \in E$.
- Note that we don't care about as singing weigths for $(i,S) \notin E$ such the respective $w_{i,S}$ are zero.
- We know that

$$\sum_{(i,S)\in E} y_{i,S} w'_{i,S}$$

  is greater than α fraction of the objective
- We need to find another y such that the same holds with weights w^+
- If we put some weight to some (i,S) such that $w'_{i,S} \neq w^+_{i,S}$, i.e., we increased this value due to monotonicity constraints we put all of this weight then we simply put weight 1 to (i,T) where $T \subseteq S$ and $w^+_{i,T} = w'_{i,S}$
- Use the previous claim to conclude that there is an integral solution x^l such that is greater than α fraction of the objective when we use values w

□

Claim: We can find a convex combination of $\frac{x^*}{\alpha}$ with polynomially many non zero weights in polynomial time

Proof:

- First prove that the optimum of the dual is 1: Set $z = 1$ and $w_{i,S} = 0$ for all $(i,S) \in E$.
- Assume we can find something better using solution x^*: Then by the previous claim we can find an integer solution that is α approximate of$x*$ which implies the that corresponding constrain is violated
- Using this argument we add the inequality

$$\frac{1}{\alpha} \cdot \sum_{(i,S)\in E} x^*_{i,S} w_{i,S} + z \geq 1$$

- Run Ellipsoid method on D to identify a dual program with polynomial size set of inequalities that is equivalent to D (the violated inequalities that are returned by the separation oracle that are used to cut the ellipsoid)
- The primal of this program has polynomial number of constraints and variables.
- Therefore non zero variables to the integer solution are at most the number of variables plus the number of the constrains.
- The separation oracle used is the following: If

$$\frac{1}{\alpha} \sum_{(i,S)\in E} x^*_{i,S} w_{i,S} + z > 1$$

  then we identify the violated constrain according to the previous Claim or otherwise we use the half space that is defined by this constraint to cut the ellipsoid

□ □

Applications of this technique

These technique can be used on existing approximation algorithms for the MUCA problem and derive the following results:

- For short valuation (each player is interested in one set and its subsets) $O(m^{\frac{1}{B+!}}$ for any $B \geq 1$ and $(1+\epsilon)$ for $B = \Omega(\log m)$.
- For general valuation functions the same bound but with with ex post Nash equilibrium as a solution concept
- For the special case of multi unit auctions this technique can be proved to be 2 approximate

Convex Rounding Framework

Relaxations

- Π an optimization problem, $\forall(\mathcal{S}, w) \in \Pi$

$$\begin{aligned} &\text{maximize } w(x) \\ &\text{subject to } x \in \mathcal{S} \end{aligned} \tag{1}$$

- Π' relaxation: $\forall(\mathcal{S}, w) \in \Pi$ defines convex and compact relaxed feasible set $\mathcal{R} \in \Re^m$ and an extension $w_{\mathcal{R}} : \mathcal{R} \to \Re$ of the objective
- So we have the following

$$\begin{aligned} &\text{maximize } w_{\mathcal{R}}(x) \\ &\text{subject to } x \in \mathcal{R} \end{aligned}$$

- A rounding scheme $r : \mathcal{R} \to \mathcal{S}$ (possibly randomized)
- If $\forall x \in \mathcal{R}$, $E[w(r(x))] \geq \alpha w_{\mathcal{R}}(x)$, then this is an α-approximation

Convex Rounding and MIDR

- MIDR: Fix some algorithm. Let distribution D_w be outcome for objective function w. Let $\mathcal{D}$ be the set of all possible D_w. The algorithm is MIDR if $\forall w, D \in \mathcal{D}, E_{x \sim D_w}[w(x)] \geq E_{x \sim D}[w(x)]$
- We can convert any MIDR algorithm to a truthful mechanism, with the same approximation guarantee, using VCG payments
- Observation: Instead of solving the relaxation and then rounding, why not optimize over the outcome of the rounding scheme?

$$\begin{aligned} &\text{maximize } E[w(r(x))] \\ &\text{subject to } x \in \mathcal{R} \end{aligned} \tag{2}$$

- We don't know if it is tractable

Lemma 0.2 *Program 2 is MIDR*

- For most typical roundings, 2 is hard to solve
- e.g., if $r(x) = x$ for $x \in \mathcal{S}$, then it is more general than 1
- So we probably should have the unusual property that $r(x) \neq x$ for $x \in \mathcal{S}$
- Rounding scheme r is α-approximate if $w(x) \geq E[w(r(x))] \geq \alpha w(x), \forall x \in \mathcal{S}$

Lemma 0.3 *If r is α-approximate, then 2 is an α-approximation to the original problem 1*

- Call r convex if $E[w(r(x))]$ concave
- For r convex 2 can be solve efficiently

- So the problem is to find α-approximate convex rounding

  Theorem 0.4 *If $\exists$ polynomial, α-approximate, convex r for Π', then $\exists$ truthful-in-expectation, polynomial, α-approximate mechanism for Π*

Combinatorial Auctions

Matroid Rank Sum

- Set function $v : 2^{[m]} \to \Re$ is an MRS function if $\exists u_1 \ldots, u_k$ (all matroid rank functions), and weights $w_1, \ldots, w_k \in \Re^+$ such that $v(S) = \sum_{\ell=1}^{k} w_\ell u_\ell(S)$
- Includes weighted coverage functions, matroid weighted rank functions, and all convex combinations of them
- Negative result: no universally truthful, polynomial, VCG-based mechanism achieves constant factor assuming $NP \not\subseteq P|Poly$

Results

Theorem 0.5 $\exists (1-1/e)$*-approximate mechanism for combinatorial auctions with MRS valuations*

- Formulation

$$\begin{array}{lll} \text{maximize} & w(x) = \sum_i v_i(\{j : x_{ij} = 1\}) & \\ \text{subject to} & \sum_i x_{ij} \le 1 & \forall j \\ & x_{ij} \in \{0,1\} & \forall i,j \end{array}$$

- $\mathcal{R}$ is relaxation to $0 \le x_{ij} \le 1$
- For $x \in [0,1]^m$, D_x define the extension of the value function

$$F_v(x) = E_{S \sim D_x}[v(S)] = \sum_S v(S) \prod_{j \in S} x_j \prod_{j \notin S} (1 - x_j)$$

Poisson Rounding

Lemma 0.6 *Consider $f : 2^V \to \Re$ monotone, submodular, and normalizes ($f(\emptyset) = 0$). Consider set $S \subseteq V$ and random set S' by choosing each element of S independently with prob p. Then $E[f(S')] \ge p \cdot f(S)$.*

Proof:

- Fix an ordering on elements of S
- Let S_i be the first i elements of S (similarly for S'_i)
- $f(S) = \sum_{1 \le i \le |S|} f(S_i) - f(S_{i-1})$, where $f(S_0) = 0$

$$\begin{aligned} E[f(S')] &= E[\sum_{1 \le i \le |S'|} f(S'_i) - f(S'_{i-1})] \\ &\ge \sum_{1 \le i \le |S|} p \cdot (f(S_i) - f(S_{i-1})) \\ &= p \cdot f(S) \end{aligned}$$

□

- Now we define the Poisson Rounding
- Given fractional solution x, independently for each item assign j to i with prob $1 - e^{-x_{ij}}$
- Note $1 - e^{-x_{ij}} \le x_{ij}$

Lemma 0.7 *Poisson rounding is* $(1 - 1/e)$*-approximate for submodular valuations*

Proof:

- Rounding applied to integer solution cancels each allocated item with prob $1/e$

– Consider $(S_1, \ldots, S_n)$ integer and corresponding (random) $(S'_1, \ldots, S'_n)$
– S'_i includes any $j \in S_i$ independently with prob $1 - 1/e$
– Then $E[v_i(S'_i)] \geq (1 - 1/e) v_i(S_i)$

□

Lemma 0.8 *Poisson rounding is concave for MRS valuations*

Proof:

- Let $(S_1, \ldots, S_n) = r_{\text{poiss}}(x)$
- Want to prove $E[w(r_{\text{poss}}(x))] = E[\sum v_i(S_i)]$ concave
- Show $E[v_i(S_i)]$ concave
- We prove concavity for a subclass: Coverage functions

Definition 0.1 *A function $f : 2^V \to \Re$ is a coverage function if $\exists$ a set A (different from V) of "activities", a value v_i for each activity $i \in A$, and a set $X_j \subseteq A$ for each j, such that $v(S) = \sum_{i \in c(S)} v_i$, where $c(S) = \bigcup_{j \in S} v_j$.*

Proof:

$$
\begin{aligned}
E[v(S)] &= E[\sum_{i \in c(S)} v_i] \\
&= \sum_{i \in A} v_i Pr[i \in c(S)]
\end{aligned}
$$

For each $i \in A$, define $Y_i = \{j \in S | i \in X_j\}$. We have

$$
\begin{aligned}
Pr[i \in c(S)] &= Pr[S \cap X_i \neq \emptyset] \\
&= 1 - e^{-\sum_{j \in Y_i} x_j}
\end{aligned}
$$

□

□