Langrangian Duality for Revenue Approximation [CDW '16]

(Recap.) We use the Lagrangian theory of duality to formulate the general Lagrangian linear program for revenue maximization. The following theory is due to Yang Cai, Nikhil Devanur, and Matt Weinberg [CDW '16]. We will let v_i be a vector which can be indexed for each item j. Similarly for allocation x at bidder i and item j. V_i represents the type space, or the support of the distribution F_i—the possible valuations that v_i can take.

We use $\varnothing$ to denote the type of not participating in the auction. Let $V_i^+ = V_i \cup \{\varnothing\}$. We use $\mathcal{P}$ to denote the polytope of feasible allocation rules.

Decision variables interim allocations $x_{ij}(v_i)$ and payments $p_i(v_i)$.

$$\begin{aligned} \max \quad & \sum_{i=1}^{n} \sum_{v_i \in V_i} f_i(v_i) \cdot p_i(v_i) \\ \text{s.t.} \quad & x_i(v_i) \cdot v_i - p_i(v_i) \geq x_i(v_i') \cdot v_i - p_i(v_i') \qquad \forall i, v_i \in V_i, v_i' \in V_i^+ (\text{dual variable } \lambda_i(v_i, v_i')) \\ & x \in \mathcal{P} \end{aligned}$$

Partial Lagrangian Primal:

$$\max_{x \in \mathcal{P}, p} \min_{\lambda \geq 0} \quad \mathcal{L}(\lambda; x, p)$$

where

$$\begin{aligned} \mathcal{L}(\lambda; x, p) &= \sum_{i=1}^{n} \left(\sum_{v_i \in V_i} f_i(v_i) \cdot p_i(v_i) + \sum_{v_i \in V_i} \sum_{v_i' \in V_i^+} \lambda_i(v_i, v_i') \cdot \big(v_i \cdot \big(x(v_i) - x(v_i') \big) - \big(p_i(v_i) - p_i(v_i') \big) \big) \right) \\ &= \sum_{i=1}^{n} \sum_{v_i \in V_i} p_i(v_i) \left(f_i(v_i) + \sum_{v_i' \in V_i} \lambda_i(v_i', v_i) - \sum_{v_i' \in V_i^+} \lambda_i(v_i, v_i') \right) + \\ & \qquad \sum_{i=1}^{n} \sum_{v_i \in V_i} x_i(v_i) \left(\sum_{v_i' \in V_i^+} v_i \cdot \lambda_i(v_i, v_i') - \sum_{v_i' \in V_i} v_i' \cdot \lambda_i(v_i', v_i) \right) \end{aligned}$$

Partial Lagrangian Dual:

$$\min_{\lambda \geq 0} \max_{x \in \mathcal{P}, p} \quad \mathcal{L}(\lambda; x, p)$$

For the dual to provide a useful (finite) upper bound we need $\max_{x \in \mathcal{P}, p} \mathcal{L}(\lambda; x, p) < \infty$. For this to be true, we must have the coefficient of $p_i(v_i)$ equal to 0, that is:

$$f_i(v_i) + \sum_{v_i' \in V_i} \lambda_i(v_i', v_i) = \sum_{v_i' \in V_i^+} \lambda_i(v_i, v_i').$$

We think of this as a "flow conservation" constraint in the following set-up. A dual solution λ is useful if and only if for each bidder i, λ_i forms a valid flow, i.e., if and only if the following satisfies flow conservation at all nodes except the source and the sink:

- Nodes: A super source s and a super sink $\varnothing$, along with a node v_i for every type $v_i \in V_i$.
- Flow from s to v_i of weight $f_i(v_i)$ for all $v_i \in V_i$.
- Flow from v to v' of weight $\lambda_i(v, v')$ for all $v \in V$ and $v' \in V^+$ (including the sink $\varnothing$).

Then

$$\sum_{i=1}^{n} \sum_{v_i \in V_i} f_i(v_i) p_i(v_i) \leq \sum_{i=1}^{n} \sum_{v_i \in V_i} f_i(v_i) \cdot x_i(v_i) \cdot \Phi_i^{\lambda}(v_i)$$

for

$$\Phi_i^{\lambda}(v_i) = v_i - \frac{1}{f_i(v_i)} \sum_{v_i' \in V_i} \lambda_i(v_i', v_i)(v_i' - v_i).$$

and this holds with equality if and only if x, p, λ are optimal solutions to the primal and dual respectively.

The Canonical Flow

The way that we will set the dual variables, which is in fact optimal in the single-dimensional setting, is as follows: $\lambda_i(v, v+1) = 1 - F_i(v+1) = \Pr_{v_i}[v_i > v+1]$. All other $\lambda_i(v, v') = 0$ except $\lambda_i(0, \varnothing) = 1$. Then

$$\Phi_i^{\lambda}(v_i) = \varphi_i(v_i)$$

is Myerson's virtual value.

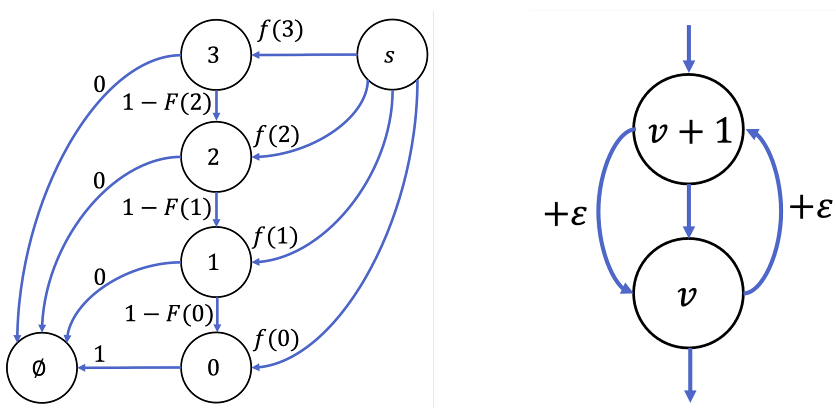

Figure 1: Left: The single-dimensional canonical dual resulting in Myersonian virtual values. Right: The process for ironing.

Ironing

For a non-monotone interval $[L, H]$ in which $\Phi_i^{\lambda}(L) > \cdots > \Phi_i^{\lambda}(H)$, we augment the following dual variables until $\Phi_i^{\lambda}(L) = \cdots = \Phi_i^{\lambda}(H)$ by increasing $\lambda_i(v+1, v)$ and $\lambda_i(v, v+1)$ by ε for $v \in [L, H-1]$.

The Unit-Demand Setting

Let $P_{ij}(v_{-i})$ denote the price that bidder i could pay to receive exactly item j in the VCG mechanism against other bidders with values v_{-i}. We then let $R_j^{v_{-i}}$ contain all types v_i such that $j \in \text{argmax}_k\{(v_{ik} - P_{ik}(v_{-i}))^+\}$. That is, $R_j^{v_{-i}}$ is the set of valuations under which bidder i prefers item j at the VCG price, breaking ties lexicographically (by smallest item index)—if $v_i \in R_j^{v_{-i}}$ then item j is bidder i's *favorite item* under valuation profile $\mathbf{v}$. $R_0^{v_{-i}}$ is the set of valuations such that bidder i prefers no item—all prices lead to negative utility.

Then our "canonical flow" is as follows for bidder i.

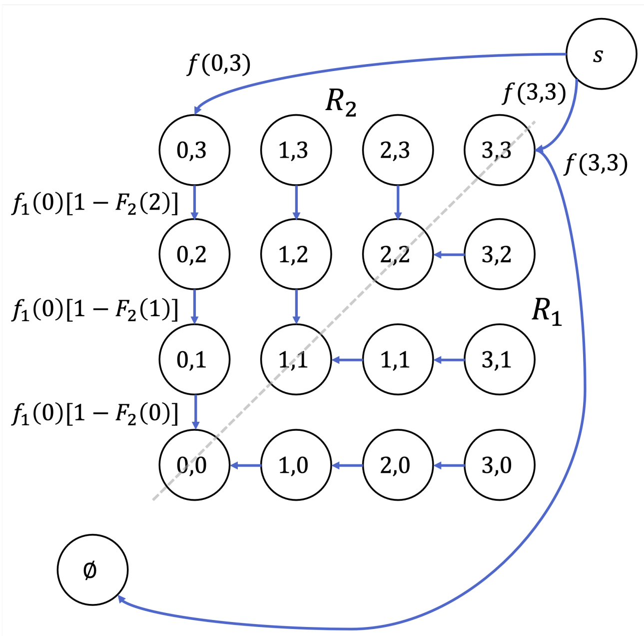

Figure 2: Left: The single-dimensional canonical dual resulting in Myersonian virtual values. Right: The process for ironing.

(We drop the subscripts i in what follows as we focus just on a specific bidder i when v_{-i} is fixed. To achieve an actual multi-bidder dual, we average over v_{-i}.)

Then for $v \in R_j$, $\lambda\left((v_j + 1, v_{-j}), (v_j + 1, v_{-j})\right) = f_{-j}(v_{-j})[1 - F_j(v_j)]$, and all other λ are 0 except $\lambda(v, \varnothing)$ used to ensure flow conservation.

Claim 1. Under the above dual variables, we get that:

- For any type $v_i \in R_j^{v_{-i}}$, its corresponding virtual value $\Phi_{ik}^{v_{-i}}(v_i)$ for item k is exactly its value v_{ik} for all non-favorite $k \neq j$.
- For any type $v_i \in R_j^{v_{-i}}$, its corresponding virtual value $\Phi_{ij}^{v_{-i}}(v_i)$ for favorite item j is exactly $\varphi_{ij}(v_{ij}) = v_{ij} - \frac{1-F_{ij}(v_{ij})}{f_{ij}(v_{ij})}$.

- The following is true:

$$\textsc{Rev}(F) = \sum_i \sum_{v_i \in V_i} f_i(v_i) \cdot p_i(v_i) \leq \sum_i \sum_{v_i \in V_i} \sum_j f_i(v_i) \cdot \hat{x}_{ij}(v_i) \cdot \Phi_{ij}(v_i)$$
$$\leq \sum_i \sum_{v_i \in V_i} \sum_j f_i(v_i) \cdot \hat{x}_{ij}(v_i) \cdot \left(v_{ij} \cdot \Pr_{\mathbf{v}_{-i}} \left[v_i \notin R_j^{(v_{-i})} \right] + \bar{\varphi}_{ij}(v_{ij}) \cdot \Pr_{\mathbf{v}_{-i}} \left[v_i \in R_j^{(v_{-i})} \right] \right)$$

Further, for a single unit-demand bidder, the above quantity is equal to

$$= \sum_{v \in V} \sum_j f(v) \cdot \hat{x}_j(v) \cdot v_j \cdot \mathbb{1}[v \notin R_j] \quad (\textsc{Non-Favorite})$$
$$+ \sum_{v \in V} \sum_j f(v) \cdot \hat{x}_j(v) \cdot \bar{\varphi}_j(v_j) \cdot \mathbb{1}[v \in R_j] \quad (\textsc{Single})$$

We will introduce the "copies setting" to show that for a single unit-demand bidder, we can bound optimal revenue by $2 \cdot \textsc{opt}^{\textsc{Copies}}$.

The Copies Setting. [Chawla Hartline Kleinberg EC '07]

- Description: For any multi-dimensional instance F we can imagine splitting bidder i into m different copies, with bidder i's copy j interested only in receiving item j and nothing else. The copies bidders are now in competition with one another (they no longer coordinate). However, any feasibility constraint still applies—if bidder i is unit demand, then at most one of i's copies can be allocated to.
- This is a *single-dimensional* setting with nm single-dimensional bidders, where copy (i, j)'s value for winning is v_{ij} (just one parameter—which is still drawn from F_{ij}).
- $\textsc{opt}^{\textsc{Copies}}(F)$ = the revenue of Myerson's optimal auction [Mye81] in the copies setting induced by F.

Lemma 1. *For any feasible $\hat{x}(\cdot)$, $\textsc{Single} \leq \textsc{opt}^{\textsc{Copies}}$.*

Proof. Recall that

$$\textsc{Single} = \sum_{v \in V} \sum_j f(v) \cdot \hat{x}_j(v) \cdot \bar{\varphi}_j(v_j) \cdot \mathbb{1}[v \in R_j].$$

Let M be the mechanism that induces $\hat{x}(\cdot)$. Consider the mechanism M' that serves agent j if and only if M would allocate item j in the original setting *and* $v \in R_j$. Then an agent j with type v_j has probability of being served in M':

$$\sum_{v_{-j}} f_{-j}(v_{-j}) \cdot \hat{x}_j(v_j, v_{-j}) \cdot \mathbb{1}[v \in R_j]$$

for all j and v_j. Because the Copies setting is single-dimensional, then $\textsc{Single}$ is the ironed virtual welfare achieved by M' with respect to $\bar{\varphi}(\cdot)$. Then the optimal revenue $\textsc{opt}^{\textsc{Copies}}$ equals the maximum ironed virtual welfare, which can only be larger than $\textsc{Single}$. (Note that this proof makes use of the assumption that item values are independent, as otherwise Myerson's theory doesn't apply.) □

Lemma 2. *When the types are unit-demand, for any feasible* $\hat{x}(\cdot)$*,* $\textsc{Non-Favorite} \leq \textsc{opt}^{\textsc{Copies}}$.

Proof. Recall that

$$\textsc{Non-Favorite} = \sum_{v \in V} \sum_{j} f(v) \cdot \hat{x}_j(v) \cdot v_j \cdot \mathbb{1}[v \notin R_j].$$

Define $S(v)$ to be the second largest number in $\{v_1, \cdots, v_m\}$. When the types are unit-demand, the Copies setting is a single-item auction with m bidders—they bid for the single item that is "winning," and then based on the agent, that item is allocated their preferred item. Therefore, if we run the Vickrey auction in the Copies setting, the revenue is $\sum_{v \in V} f(v) \cdot S(v)$. If $v \notin R_j$, then there exists some $k \neq j$ such that $v_k \geq v_j$, so $v_j \cdot \mathbb{1}[v \notin R_j] \leq S(v)$ for all j. Therefore, $\sum_{v \in V} \sum_j f(v) \cdot \hat{x}_j(v) \cdot v_j \cdot \mathbb{1}[v \notin R_j] \leq \sum_{v \in V} \sum_j f(v) \cdot \hat{x}_j(v) \cdot S(v) \leq \sum_{v \in V} f(v) \cdot S(v)$. The last inequality is because the bidder is unit demand, so $\sum_j \hat{x}_j(v) \leq 1$. □

References

[1] Yang Cai, Nikhil R. Devanur, and S. Matthew Weinberg. A Duality Based Unified Approach to Bayesian Mechanism Design. In *Proceedings of the Forty-eighth Annual ACM Symposium on Theory of Computing*, STOC '16, pages 926–939, New York, NY, USA, 2016. ACM.

[2] Shuchi Chawla, Jason D. Hartline, and Robert Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, EC '07, pages 243–251, 2007.