Langrangian Duality for Revenue Approximation [CDW '16]

Now we will use this 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 $\hat{x}_{ij}(v_i)$ and payments $p_i(v_i)$.

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

Partial Lagrangian Dual:

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

where

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

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

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

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: $\lambda_i(v+1, v) = 1 - F_i(v) = \Pr_{v_i}[v_i > v]$. 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.

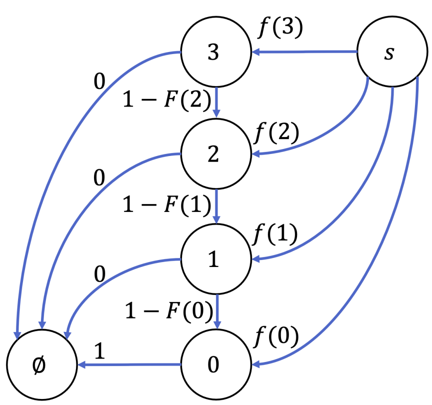

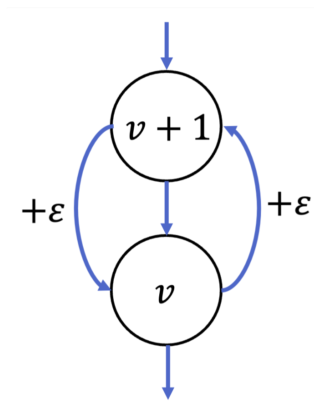

Figure 1: Left: The single-dimensional canonical dual resulting in Myersonian virtual values. Right: The process for ironing.

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.

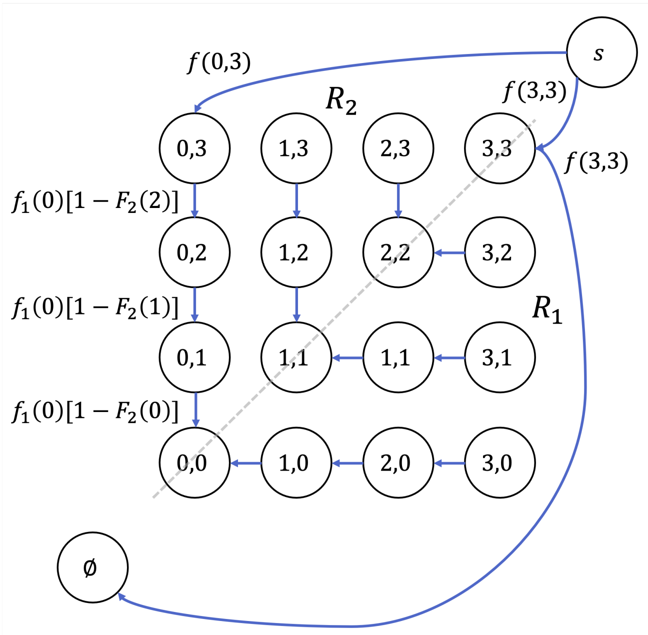

Figure 2: Left: The canonical dual for a multi-parameter setting.

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:

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

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:

- Type of setting?
- $\textsc{opt}^{\textsc{Copies}}(F) =$

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

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

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.