

Cost-Sharing Methods for Scheduling Games under Uncertainty

Vasilis Gkatzelis, Drexel University

Joint work with:

**Giorgos Christodoulou,
Alkmini Sgouritsa,**




**University of Liverpool
Max-Planck Institute**

- We have a set M of m machines and set N of n jobs
- Each job i has weight w_i and needs to be processed by some machine
- Each machine j has a non-decreasing cost function $c_j(\ell)$
 - This cost depends on the load $\ell = \sum_i w_i$ of the agents using it
 - For this talk, just assume that $w_i = 1$ for every agent i
 - The cost function satisfies $c_j(0) = 0$ for every machine j
 - Each cost function can be convex, concave, or more complicated
- A schedule \mathbf{s} assigns each job to a machine
 - Let $S_j(\mathbf{s})$ be the set of jobs assigned to machine j in schedule \mathbf{s}
 - Let $\ell_j(\mathbf{s}) = \sum_{i \in S_j(\mathbf{s})} w_i$ be the total load of the jobs using j in \mathbf{s}
 - Then the cost on each machine j is $c_j(\mathbf{s}) = c_j(\ell_j(\mathbf{s}))$
- Our goal is to output a schedule \mathbf{s} that minimizes $C(\mathbf{s}) = \sum_{j \in M} c_j(\mathbf{s})$


Scheduling Games




Scheduling Games




Scheduling Games



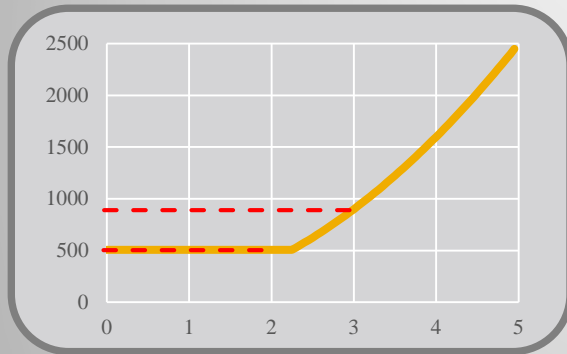
Equal Sharing




Equal Sharing




Scheduling Games




Equal Sharing




Equal Sharing




Scheduling Games



Equal Sharing




Equal Sharing




Oblivious Cost-Sharing

Oblivious Cost-Sharing






Each machine's policy sees only agents using it and is independent of other machines in the system



Omnipotent Cost-Sharing

Omnipotent Cost-Sharing




Each machine's policy sees and can depend on all the agents and all the machines in the system



Resource-Aware Cost-Sharing

Resource-aware Cost-Sharing

Each machine's policy sees only agents using it but can depend on other machines in the system



Scheduling Games Model

- Set M of m machines and set N of n agents
- Agent i needs **one machine** to process a job with weight w_i
- Machine j has a cost function $c_j(\ell)$
 - This cost depends on the load $\ell = \sum_i w_i$ of the agents using it
 - The cost function satisfies $c_j(0) = 0$ for every machine j
- Strategy $s_i \in M$ from each job i leads to profile \mathbf{s}
- Let $S_j(\mathbf{s})$ be the set of jobs using machine j in profile \mathbf{s}
- Let $\ell_j(\mathbf{s}) = \sum_{i \in S_j(\mathbf{s})} w_i$ be the total load of the jobs using j in \mathbf{s}
- Cost-sharing method $\xi_{ij}(\mathbf{s})$ defines cost of i in profile \mathbf{s}
 - **Budget-balanced** if for every \mathbf{s} and every j : $\sum_{i \in S_j(\mathbf{s})} \xi_{ij}(\mathbf{s}) = c_j(\mathbf{s})$
 - **Stable** if a pure Nash equilibrium exists for all sets M and N

- We measure the efficiency of a schedule \mathbf{s} using $C(\mathbf{s}) = \sum_{j \in M} c_j(\mathbf{s})$
- Given a class of games \mathcal{G} , for each game $G \in \mathcal{G}$:
 - Let $F(G)$ be the set of all possible schedules
 - Let $E(G) \subseteq F(G)$ be the set of pure Nash equilibria
- **Price of anarchy (PoA)** of a class \mathcal{G} is:
$$\sup_{G \in \mathcal{G}} \frac{\max_{\mathbf{s} \in E(G)} C(\mathbf{s})}{\min_{\mathbf{s}^* \in F(G)} C(\mathbf{s}^*)}$$
- We may **overcharge** so that $\sum_{i \in S_j(\mathbf{s})} \xi_{ij}(\mathbf{s}) = \hat{C}(\mathbf{s}) > C(\mathbf{s})$
- With overcharging, the PoA of a class \mathcal{G} becomes:
$$\sup_{G \in \mathcal{G}} \frac{\max_{\mathbf{s} \in E(\hat{G})} \hat{C}(\mathbf{s})}{\min_{\mathbf{s}^* \in F(G)} C(\mathbf{s}^*)}$$

- We measure the efficiency of a schedule \mathbf{s} using $C(\mathbf{s}) = \sum_{j \in M} c_j(\mathbf{s})$
- Given a class of games \mathcal{G} , for each game $G \in \mathcal{G}$:
 - Let $F(G)$ be the set of all possible schedules
 - Let $E(G) \subseteq F(G)$ be the set of pure Nash equilibria

Price of Stability (PoS)
if we change max to min

- **Price of anarchy (PoA)** of a class \mathcal{G} is:
$$\sup_{G \in \mathcal{G}} \frac{\max_{\mathbf{s} \in E(G)} C(\mathbf{s})}{\min_{\mathbf{s}^* \in F(G)} C(\mathbf{s}^*)}$$
- We may **overcharge** so that $\sum_{i \in S_j(\mathbf{s})} \xi_{ij}(\mathbf{s}) = \hat{C}(\mathbf{s}) > C(\mathbf{s})$



- With overcharging, the PoA of a class \mathcal{G} becomes:
$$\sup_{G \in \mathcal{G}} \frac{\max_{\mathbf{s} \in E(\hat{G})} \hat{C}(\mathbf{s})}{\min_{\mathbf{s}^* \in F(G)} C(\mathbf{s}^*)}$$

Many papers on **cost-sharing** and **coordination mechanisms**

- **Chen, Roughgarden, and Valiant 2010**
 - Network design games (constant cost functions)
 - Agents can choose multiple machines
 - Characterization of stable cost-sharing protocols
- **von Falkenhausen and Harks 2013**
 - Studied general cost functions
 - Also considered extension to matroids
- Both of these papers are restricted to **budget-balanced** protocols and the **omnipotent** and **oblivious** models


Convex Cost Functions

- Assume all the cost functions are convex
- How inefficient can the outcome be if we use equal sharing?
- E.g., **2 machines** $c_1(\ell) = 50 \cdot 2^{\ell-1}$ and $c_2(\ell) = 350 \ell$ and **5 agents**



Convex Cost Functions


- Given global ordering π over the universe of agents
- **Incremental cost-sharing protocol** [Moulin '99]
 - Order the agents using a machine based on π
 - Charge each agent for a cost equal to its marginal contribution
 - $\xi_{ij}(\mathbf{s}) = c_j(\ell_j^{<i}(\mathbf{s}) + \mathbf{w}_i) - c_j(\ell_j^{<i}(\mathbf{s}))$
- This protocol is **stable**, **budget-balanced**, and **oblivious**, and it achieves a **PoA of 1** for unweighted agents and convex functions



Theorem: Every **stable, budget-balanced, resource-aware** mechanism has a **PoA of $\Omega(m)$** for general cost functions

- What if we allow the use of overcharging?


Theorem: Every **stable, (non-budget-balanced,) resource-aware** mechanism has a **PoA of $\Omega(\sqrt{m})$** for general cost functions



Concave Cost Functions

Theorem: Any **stable**, **oblivious**, and **budget-balanced** cost-sharing policy has $\text{PoA} \geq n$ for strictly concave valuations

- **Observation:** optimal solution assigns all jobs to one machine, but which machine this is depends on the total load of the jobs
- The lower bound above is for unweighted, but our mechanism works for general weights as well



Concave Cost Functions

- $\phi(\ell) = \min_{j \in M} c_j(\ell)$: smallest cost over all machines at load ℓ
- $X_{\min}(\ell) = \arg \min_{j \in M} c_j(\ell)$: set of machines with cost $\phi(\ell)$ at load ℓ
- $h_j(s) = \arg \min_{i' \in S_j(s)} \{\pi(i')\}$: highest priority agent on machine j

$$\xi_{ij}(s) = \begin{cases} c_j(\ell_j(s)) & \text{if } j \notin X_{\min}(\ell_j(s)) \text{ and } i = h_j \\ 0 & \text{if } j \notin X_{\min}(\ell_j(s)) \text{ and } i \neq h_j \\ \phi(w_i) & \text{if } j \in X_{\min}(\ell_j(s)) \text{ and } i = h_j \\ w_i \frac{c_j(\ell_j(s)) - \phi(w_{h_j})}{\ell_j(s) - w_{h_j}} & \text{if } j \in X_{\min}(\ell_j(s)) \text{ and } i \neq h_j \end{cases}$$

Theorem: This cost-sharing mechanism is **stable**, **budget-balanced**, **resource-aware**, and it achieves **PoA of 1** for concave cost functions

- No **budget-balanced** mechanism can guarantee a PoS better than $O(\log m)$ **even if it is omnipotent** [vFH 13]
- If $j \in M_{convex}$ and $\ell \geq n_j^{max}$, instead of $c_j(\ell)$, we use cost functions $\hat{c}(\ell) = \max \left\{ \max_{j' \in M_{concave}} c_{j'}(1), c_j(\ell) \right\}$
- **VC mechanism:** Using the over-charged cost functions above
 - For convex machines use incremental cost-sharing protocol
 - For concave machines use our concave cost-sharing protocol

Theorem: This mechanism is **stable**, **resource-aware** and achieves a **PoA of 2** for instances with convex and concave functions

Two-Machine Instances

- Let $\alpha(s)$ be highest priority agent using the first machine
- Let $\beta(s)$ be lowest priority agent using second machine
- **Increasing-Decreasing mechanism:** For any profile s ,
 - Charge agent $\alpha(s)$ for the whole cost of the first machine
 - Charge agent $\beta(s)$ for the whole cost of the second machine.

Theorem: This mechanism is **stable**, **budget-balanced**, **resource-aware**, and it achieves a **PoA of 2** for arbitrary cost functions

Note that this mechanism is stable, but not a Shapley value variant

Theorem: Any **stable**, **(non-budget-balanced)**, **resource-aware** mechanism has **PoA > 1.36** , even for instances with just two machines with convex and concave cost functions

- **Resource-aware** cost-sharing
 - Well motivated middle-ground between omnipotent and oblivious
 - Non-trivial use of extra information may enable improvements
- Power of **over-charging**
 - Leads to improvements despite the additional costs
 - For omnipotent protocols, over-charging can yield PoA of 1
 - Budget-balanced omnipotent protocols have PoA $\Omega(\log n)$ [vFH 13]
- **Open problems**
 - Weighted upper bounds in our setting
 - Applying the resource-aware model in other settings

Thank you!