Improved Price of Anarchy via Predictions

Vasilis Gkatzelis
Drexel University

Kostas Kollias
Google Research

Alkmini Sgouritsa
University of Liverpool

Xizhi Tan
Drexel University

Decentralized multiagent systems

- In decentralized multiagent systems, the agents can **strategically** interact with a system aiming to maximize their own utility
- The system designer can use **decentralized mechanisms** that regulate the use of the resources, provided **limited information** regarding the state of the system
- This information limitation, together with agents acting strategically, often leads to very inefficient outcomes, and thus a **high price of anarchy (PoA)** – the ratio between the worst Nash equilibrium and the optimal outcome

Network cost-sharing games

Consider a very general family of games with n agents over an undirected graph $G = (V, E)$

- Each agent i needs to connect a **terminal** $t_i \in V$ to the **source** $s \in V$ by choosing a **path** from t_i to s
- If $e \in E$ is used by ℓ agents, the induced cost is $c_e(\ell)$
- We must decide how to share the cost among the agents to minimize the Price of Anarchy

We focus on two specific settings:

1. **Multicast games**: general graphs + constant cost functions
2. **Scheduling games**: parallel graphs + arbitrary non-decreasing cost functions

Mechanism design problem: Can we approximate the optimal PoA using resource-aware protocols?

1. For multicast games, no resource-aware mechanism can get a PoA better than $O(\log n)$. [Christodoulou and Sgouritsa '16]
2. For scheduling games, no resource-aware mechanism can get a PoA better than $O(\sqrt{n})$. The current best PoA is $O(n)$. [Christodoulou et al. '17]

Our Results

Main Question: Can decentralized protocols, enhanced with predictions, achieve improved price of anarchy bounds? How do these bounds depend on the prediction accuracy?

Problem	Lower bound w/o predictions	Our results: **Consistency** (no error)	Our results: **Robustness** (arbitrary error)	Optimality
Scheduling	$O(n)$	4	$4n$	small constant consistency with best known robustness
Multicast	$\Theta(\log n)$	4	$\log n$	small constant consistency with best possible robustness

Open Questions

- Can we achieve the same improved PoA for more general class of games?
- Can mechanism design with predictions be helpful in other decentralized system design problems?

Scheduling games with predictions

- The system designer has access to a prediction on the **number of players** $\hat{n}$ participating in the game
- The **(unknown) error** of the prediction is $\boldsymbol{\delta = |\hat{n} - n|}$

Cost function modification guided by prediction **+** Ordered protocol

- Our mechanism achieves a PoA of $\mathbf{min\{4(\delta + 1), 4n\}}$.
- The same guarantee holds for **series-parallel graphs** with general non-decreasing cost-functions

Multicast games with predictions

- Prediction: we have predictions over **the set of terminals**
- The **(unknown) error** of the predictions, (δ, D):
  - δ: the total number of "mis-predicted" terminals
  - D: the sum of the distances between the predicted and actual terminals

Ordering of vertices guided by prediction **+** Ordered protocol

- Our mechanism achieves a PoA of $\mathbf{min\{4 + 6D + \log \delta, \log n\}}$
- This result hold for any (δ, D) tuple

Acknowledgments: This work was supported in part by NSF CAREER award CCF2047907. The work of Alkmini Sgouritsa was done while she was a Visiting Researcher at Google Research.