ETH zürich   .IfA AUTOMATIC CONTROL LABORATORY   LAS Learning & Adaptive Systems

Bounding Inefficiency of Equilibria in Continuous Actions Games using Submodularity and Curvature

Pier Giuseppe Sessa, Maryam Kamgarpour, Andreas Krause

Motivation and Problem Set-Up

- Games with continuous actions arise in several domains. Their (in)efficiency, however, is less understood than in games with finitely many actions.
- We consider N-player continuous games $\mathcal{G}$ described by:
  - strategy sets: $\mathcal{S}_i \subseteq \mathbb{R}^d_+$
  - payoffs: $\pi_i : \mathcal{S} = \prod_{i=1}^N \mathcal{S}_i \to \mathbb{R}$
  - social function: $\gamma : \mathbb{R}^{Nd}_+ \to \mathbb{R}$
- A **coarse correlated equilibrium** (CCE) is a probability distribution σ over the outcomes $\mathcal{S}$ that satisfies
$$\mathbb{E}_{\mathbf{s}\sim\sigma}[\pi_i(\mathbf{s})] \geq \mathbb{E}_{\mathbf{s}\sim\sigma}[\pi_i(\mathbf{s}'_i, \mathbf{s}_{-i})], \qquad \forall i, \forall \mathbf{s}'_i \in \mathcal{S}_i$$
- **No-regret learning** dynamics converge to CCEs of the repeated game.

- Efficiency of the game is measured with the **Price of Anarchy** of any CCE:
$$PoA_{CCE} = \frac{\max_{\mathbf{s}\in\mathcal{S}} \gamma(\mathbf{s})}{\min_{\sigma\in\Delta} \mathbb{E}_{\mathbf{s}\sim\sigma}[\gamma(\mathbf{s})]}$$
- PoA_{CCE} has two important implications:
  - In multi-agent systems, bounds the inefficiency of no-regret dynamics followed by selfish agents.
  - In distributed optimization, certifies approximation guarantees of distributed no-regret algorithms.

Main Results

Def. A function $f : \mathcal{X} \subseteq \mathbb{R}^n \to \mathbb{R}$ is **DR-submodular** [1] if, $\forall \mathbf{x} \leq \mathbf{y} \in \mathcal{X}, \forall i \in [n]$, $\forall k \in \mathbb{R}_+$ s.t. $(\mathbf{x} + k\mathbf{e}_i)$ and $(\mathbf{y} + k\mathbf{e}_i) \in \mathcal{X}$,
$$f(\mathbf{x} + k\mathbf{e}_i) - f(\mathbf{x}) \geq f(\mathbf{y} + k\mathbf{e}_i) - f(\mathbf{y})$$

Def. $\mathcal{G}$ is a valid utility game with continuous strategies if:

i. γ is monotone **DR-submodular**

ii. $\pi_i(\mathbf{s}_i, \mathbf{s}_{-i}) \geq \gamma(\mathbf{s}) - \gamma(\mathbf{0}, \mathbf{s}_{-i})$ for each i and $\mathbf{s}$.

iii. $\gamma(\mathbf{s}) \geq \sum_{I=1}^N \pi_i(\mathbf{s})$ for each $\mathbf{s}$.

Extends [2,3] to continuous domains.

Def. *Curvature* of a monotone DR-submodular $f : \mathcal{X} \subseteq \mathbb{R}^n_+ \to \mathbb{R}$, w.r.t. $\mathbf{0} \in \mathcal{Z} \subseteq \mathcal{X}$:
$$\alpha(\mathcal{Z}) = 1 - \inf_{\substack{\mathbf{x}\in\mathcal{Z}, i\in[n]: \\ \mathbf{x}+k\mathbf{e}_i\in\mathcal{Z}}} \lim_{k\to 0^+} \frac{f(\mathbf{x} + k\mathbf{e}_i) - f(\mathbf{x})}{f(k\mathbf{e}_i) - f(\mathbf{0})} \quad \in [0,1]$$

Generalizes *total curvature* of set functions. $\alpha(\mathcal{Z}) = 0$ iff f is affine.

Theorem. Let $\tilde{\mathcal{S}} = \{\mathbf{x} \in \mathbb{R}^{Nd} \,|\, \mathbf{0} \leq \mathbf{x} \leq \mathbf{s} + \mathbf{s}', \forall \mathbf{s}, \mathbf{s}' \in \mathcal{S}\}$.
If $\mathcal{G}$ is a valid utility game where γ has curvature $\alpha(\tilde{\mathcal{S}}) \leq \alpha$, then
$$PoA_{CCE} \leq 1 + \alpha \ .$$

See extension to a class of non-submodular functions in [4].

Designing Games for Distributed Optimization

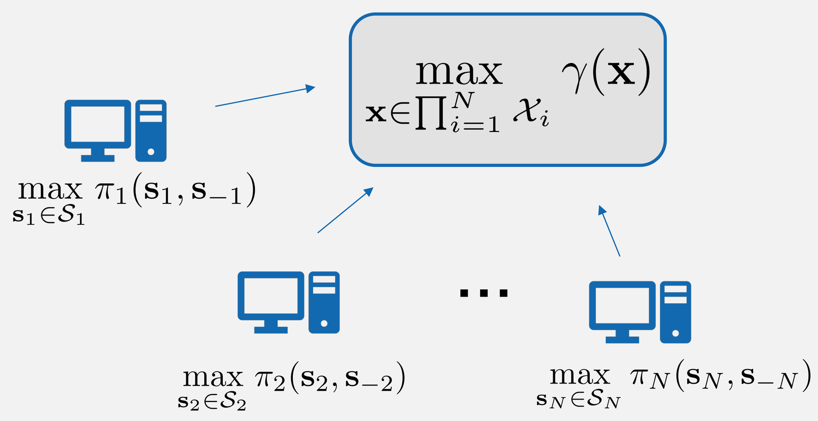

- γ is monotone DR-submodular
- Disjoint constraints set

Idea: set-up a repeated game $\hat{\mathcal{G}}$ with
- strategy sets: $\mathcal{S}_i = \mathcal{X}_i$
- payoffs $\pi_i(\mathbf{s}) = \gamma(\mathbf{s}) - \gamma(\mathbf{0}, \mathbf{s}_{-i})$, for $i = 1, \dots, N$.

Was done in [5] for binary strategy sets.

Fact. $\hat{\mathcal{G}}$ is a valid utility game with continuous strategies.

Assume a **no-regret learning algorithm** exists for each player.
Let D-NoRegret be the simultaneous implementation of such algorithms.

Corollary. D-NoRegret converges to a distribution σ over $\mathcal{X}$ such that
$$\mathbb{E}_{\mathbf{x}\sim\sigma}[\gamma(\mathbf{x})] \geq 1/(1+\alpha)\, \gamma(\mathbf{x}^\star) \ .$$

Can improve the available $(1 - e^{-1})$ approximation by [1].

Examples and Experiments

Continuous Budget Allocation game

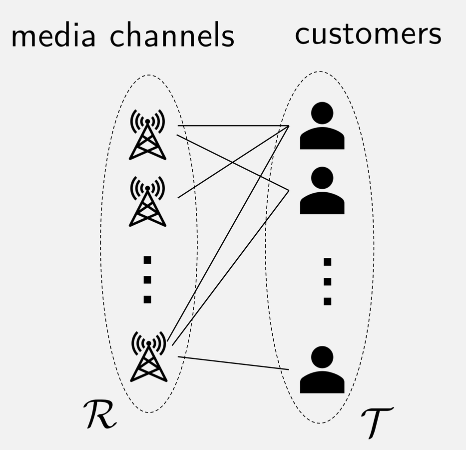

- N advertisers invest in R media channels to attract the maximum number of customers.
- $[\mathbf{s}_i]_r :=$ amount invested by advertiser i in channel r.
- $p_i(r,t) :=$ probability that advertiser i attracts customer t via channel r.
- Market analyst aims to maximize the average number of total attracted customers:
$$\gamma(\mathbf{s}) = \sum_{t\in\mathcal{T}} \left(1 - \prod_{i=1}^N \prod_{r\in\Gamma(t)} (1 - p_i(r,t))^{[\mathbf{s}_i]_r}\right)$$

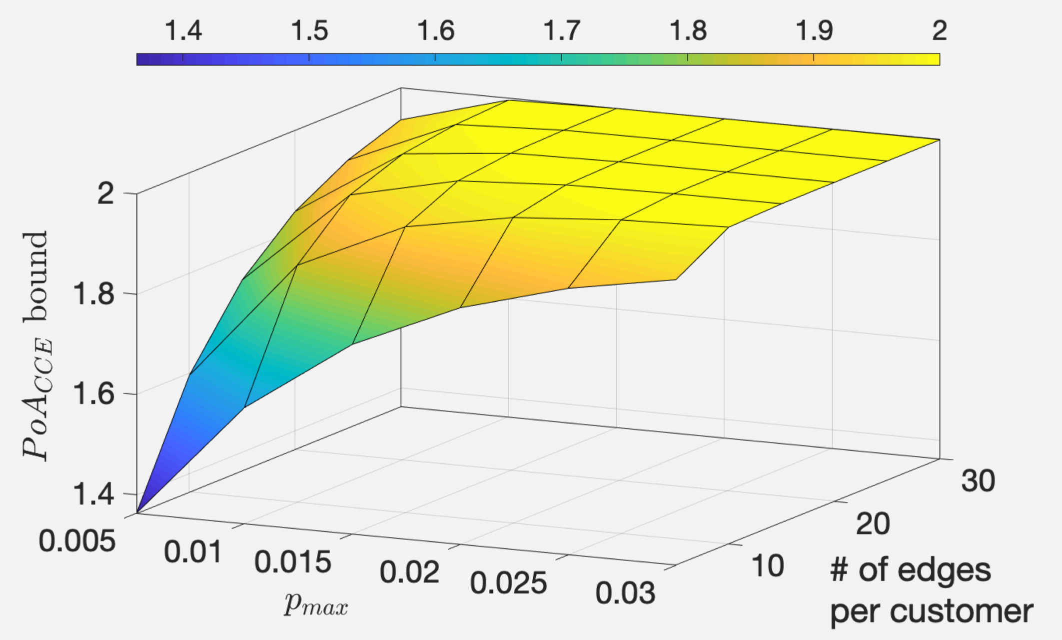

Fig1: For small attraction probabilities and number of edges, the obtained PoA bound strictly improves the bound of 2 by [3] for the discrete setting.

Sensor Coverage Problem with continuous assignments

- Given: N autonomous sensors and d locations.
- $[\mathbf{x}_i]_r :=$ energy of sensor i allocated to location r.
- $1 - (1 - p_i^r)^{[\mathbf{x}_i]_r} :=$ probability that sensor i detects an event in location r.
- $w_r :=$ probability of an event occurring in location r.
- Goal: Assign sensors to locations to maximize the probability of detecting an event:
$$\gamma(\mathbf{x}) = \sum_{r\in[d]} w_r \left(1 - \prod_{i\in[N]} (1 - p_i^r)^{[\mathbf{x}_i]_r}\right)$$
- We can set-up a valid utility game $\hat{\mathcal{G}}$ and implement D-NoRegret (Online Gradient Ascent is no-regret for each player).

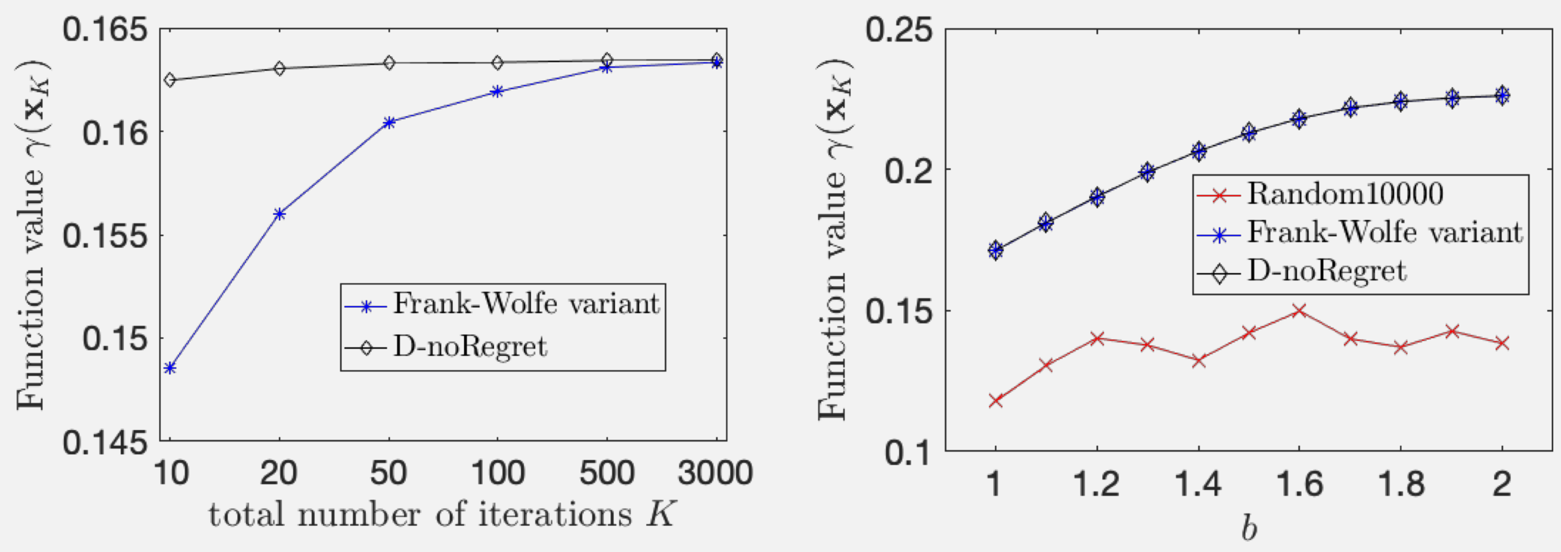

Fig2: D-NoRegret shows faster convergence than Frank-Wolfe variant by [1]. However, for K=3000 iterations the two algorithms perform equally.

Acknowledgements

This work was gratefully supported by Swiss National Science Foundation, under the grant SNSF 200021_172781, and by the European Union's Horizon 2020 ERC grant 815943.

References

[1] A. A. Bian, B. Mirzasoleiman, J. M. Buhmann, and A. Krause. Guaranteed non-convex optimization: Submodular maximization over continuous domains. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, AISTATS 2017.
[2] A. Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *Proceedings of the 43rd Symposium on Foundations of Computer Science*, FOCS '02.
[3] T. Maehara, A. Yabe, and K. Kawarabayashi. Budget allocation problem with multiple advertisers: A game theoretic view. In *Proceedings of the 32nd International Conference on Machine Learning*, ICML'15.
[4] P.G. Sessa, M. Kamgarpour, and A. Krause. Bounding Inefficiency of Equilibria in Continuous Actions Games using Submodularity and Curvature. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics*, AISTATS 2019.
[5] J. R. Marden and A. Wierman. Distributed welfare games. *Operations Research*, 2013.

22nd International Conference on Artificial Intelligence and Statistics (AISTATS). Naha, Okinawa, Japan, 2019.

Emails: {sessap, maryamk, krausea}@ethz.ch