Lipschitz Continuity and Approximate Equilibria

Argyrios Deligkas John Fearnley Paul Spirakis

Department of Computer Science
University of Liverpool

Background

Strategic-Form game

- M players
- S_i of pure strategies for each player $i \in [M]$
- U_i payoff function for each player $i \in [M]$

Background

Strategic-Form game

- M players
- S_i of pure strategies for each player $i \in [M]$
- U_i payoff function for each player $i \in [M]$

- Player i picks a probability distribution x_i over the set of pure strategies S_i
- Strategy profile $X = (x_1, \ldots, x_M)$
- $U_i(x_i, X_{-i})$ payoff for the player i under the profile X

Background

Strategic-Form game

- M players
- S_i of pure strategies for each player $i \in [M]$
- U_i payoff function for each player $i \in [M]$

- Player i picks a probability distribution x_i over the set of pure strategies S_i
- Strategy profile $X = (x_1, \ldots, x_M)$
- $U_i(x_i, X_{-i})$ payoff for the player i under the profile X
  The payoff function $U_i(x_i, X_{-i})$ is linear in x_i

Background

Definition (ϵ-Nash equilibrium)

A strategy profile is an ϵ-Nash equilibrium if:

no player can gain more than ϵ by a unilateral deviation

(additive notion of approximation)

Background

Definition (ϵ-Nash equilibrium)

A strategy profile is an ϵ-Nash equilibrium if:

no player can gain more than ϵ by a unilateral deviation

Background

Definition (ϵ-Nash equilibrium)

A strategy profile is an ϵ-Nash equilibrium if:

no player can gain more than ϵ by a unilateral deviation

Theorem (Daskalakis, Goldberg, Papadimitriou 2006)

If there is an FPTAS for computing an ϵ-Nash for 4 player games, then $\mathtt{PPAD} = \mathtt{P}$.

Background

Definition (ϵ-Nash equilibrium)

A strategy profile is an ϵ-Nash equilibrium if:

no player can gain more than ϵ by a unilateral deviation

Theorem (Chen, Deng, Teng 2006)

If there is an FPTAS for computing an ϵ-Nash of a bimatrix game, then $\mathtt{PPAD} = \mathtt{P}$.

Background

Definition (ϵ-Nash equilibrium)

A strategy profile is an ϵ-Nash equilibrium if:

no player can gain more than ϵ by a unilateral deviation

Theorem (Rubinstein 2016)

If there is a PTAS for computing an ϵ-Nash of a bimatrix game, then **`EndOfTheLine`** *can be solved faster than exponential time.*

Background

Definition (ϵ-Nash equilibrium)

A strategy profile is an ϵ-Nash equilibrium if:

no player can gain more than ϵ by a unilateral deviation

Theorem (Lipton, Markakis, Mehta (LMM) 2003)

For every constant $\epsilon > \mathbf{0}$, an ϵ-Nash can be computed in **quasi-polynomial time**.

Background

Definition (ϵ-Nash equilibrium)

A strategy profile is an ϵ-Nash equilibrium if:

no player can gain more than ϵ by a unilateral deviation

Theorem (Lipton, Markakis, Mehta (LMM) 2003)

For every constant $\epsilon > 0$, an ϵ-Nash can be computed in **quasi-polynomial time**.

All these results apply on **strategic-form** games.

Lipschitz Games

λ_p-Lipschitz game

- M players
- S_i is the convex hull of n vectors in $\mathbb{R}^d$
- $T_i(x_i, X_{-i})$ utility function for each player $i \in [M]$,

Lipschitz Games

λ_p-Lipschitz game

- M players
- S_i is the convex hull of n vectors in $\mathbb{R}^d$ **(continuous action space)**
- $T_i(x_i, X_{-i})$ utility function for each player $i \in [M]$, where function $T_i(x_i, X_{-i})$ is **λ_p-Lipschitz continuous** w.r.t. x_i when X_{-i} is fixed.

Lipschitz Games

λ_p-Lipschitz game

- M players
- S_i is the convex hull of n vectors in $\mathbb{R}^d$
- $T_i(x_i, X_{-i})$ utility function for each player $i \in [M]$, where function $T_i(x_i, X_{-i})$ is **λ_p-Lipschitz continuous** w.r.t. x_i when X_{-i} is fixed.

Definition (λ_p-Lipschitz)

A function $f : A \to \mathbb{R}$, with $A \subseteq \mathbb{R}^d$ is λ_p-Lipschitz continuous if for every x and y in A, it is true that $|f(x) - f(y)| \leq \lambda \cdot ||x - y||_p$.

Lipschitz Games: examples

Concave Games	Rosen (Econometrica 65)
Risk Games	Fiat, Papadimitriou (SAGT 10)
	Mavronicolas, Monien (TCS 16)
Biased Games	Caragiannis, Kurokawa, Procaccia (AAAI 15)

Lipschitz Games: examples

Concave Games Rosen (Econometrica 65)

Risk Games Fiat, Papadimitriou (SAGT 10)

Mavronicolas, Monien (TCS 16)

Biased Games Caragiannis, Kurokawa, Procaccia (AAAI 15)

- Concave and Biased Games **always** possess an equilibrium
- Equilibrium existence for Risk Games is **NP-complete**

Our results

- Efficient algorithms for computing ϵ-equilibria in λ_p-Lipschitz games for every constant $\epsilon > 0$, or decide that the game does **not** possess one.

Our results

- Efficient algorithms for computing ϵ-equilibria in λ_p-Lipschitz games for every constant $\epsilon > 0$, or decide that the game does **not** possess one.

- Polynomial-time algorithms for computing constant approximate equilibria for three classes of biased games.

The QPTAS of LMM

Existence

In any $n \times n$ bimatrix game, any $\epsilon > 0$ and $k \geq O(\frac{\ln n}{\epsilon^2})$, there exists a k-uniform strategy profile that is an ϵ-NE.

The QPTAS of LMM

Existence

In any $n \times n$ bimatrix game, any $\epsilon > 0$ and $k \geq O(\frac{\ln n}{\epsilon^2})$, there exists a k-uniform strategy profile that is an ϵ-NE.

Computation

Find such a profile in quasi-polynomial time by exhaustive search over all k-uniform strategy profiles (time $n^{O(\frac{\ln n}{\epsilon^2})}$).

k-uniform strategies

$X = \{x_1, x_2, \ldots, x_n\} \subset \mathbb{R}^d$

$conv(X)$: convex hull of X

k-uniform strategies

$X = \{x_1, x_2, \ldots, x_n\} \subset \mathbb{R}^d$

$conv(X)$: convex hull of X

Definition

A vector $y \in conv(X)$ is said to be **k-uniform** with respect to X if there exists a size k multiset S of $[n]$ such that $y = \frac{1}{k} \sum_{i \in S} x_i$.

k-uniform strategies

$X = \{x_1, x_2, \ldots, x_n\} \subset \mathbb{R}^d$

$conv(X)$: convex hull of X

Definition

A vector $y \in conv(X)$ is said to be ***k*-uniform** with respect to X if there exists a size k multiset S of $[n]$ such that $y = \frac{1}{k}\sum_{i \in S} x_i$.

Definition

A strategy profile is said to be k-uniform if the strategy for every player is a k-uniform vector.

Our result

Theorem 1 (Existence)

In any λ_p-Lipschitz game that possess an equilibrium and any $\epsilon > 0$, there is a k-uniform strategy profile, with $k = \frac{16M^2\lambda^2 p}{\epsilon^2}$ that is an ϵ-equilibrium.

Our result

Theorem 1 (Existence)

In any λ_p-Lipschitz game that possess an equilibrium and any $\epsilon > 0$, there is a k-uniform strategy profile, with $k = \frac{16M^2\lambda^2 p}{\epsilon^2}$ that is an ϵ-equilibrium.

- λ_p-Lipschitz continuity is crucial for the existence
- In the proof we utilize a recent result of Barman (STOC 15)

Our result

Theorem 1 (Existence)

In any λ_p-Lipschitz game that possess an equilibrium and any $\epsilon > 0$, there is a k-uniform strategy profile, with $k = \frac{16M^2\lambda^2 p}{\epsilon^2}$ that is an ϵ-equilibrium.

- **Exhaustive search is not trivial in our case.**

Our result

Theorem 1 (Existence)

In any λ_p-Lipschitz game that possess an equilibrium and any $\epsilon > 0$, there is a k-uniform strategy profile, with $k = \frac{16M^2\lambda^2 p}{\epsilon^2}$ that is an ϵ-equilibrium.

- **Exhaustive search is not trivial in our case.**
- Hard to compute the approximation guarantee of a k-uniform strategy profile.

Our result

Theorem 1 (Existence)

In any λ_p-Lipschitz game that possess an equilibrium and any $\epsilon > 0$, there is a k-uniform strategy profile, with $k = \frac{16M^2\lambda^2 p}{\epsilon^2}$ that is an ϵ-equilibrium.

- **Exhaustive search is not trivial in our case.**
- Hard to compute the approximation guarantee of a k-uniform strategy profile.
- Hard to compute best responses ($\max_{x_i} T_i(x_i, X_{-i})$)

Our result

Theorem 1 (Existence)

In any λ_p-Lipschitz game that possess an equilibrium and any $\epsilon > 0$, there is a k-uniform strategy profile, with $k = \frac{16M^2\lambda^2 p}{\epsilon^2}$ that is an ϵ-equilibrium.

- **Exhaustive search is not trivial in our case.**
- Hard to compute the approximation guarantee of a k-uniform strategy profile.
- Hard to compute best responses ($\max_{x_i} T_i(x_i, X_{-i})$)
- We compute **approximate** best responses using k-uniform strategies.

Our result

Theorem 2 (Computation)

For any λ_p-Lipschitz game L in time $O\left(Mn^{Mk+l}\right)$, we can either compute a ϵ-equilibrium, or decide that L does not posses an exact equilibrium, where $k = O\left(\frac{\lambda^2 Mp}{\epsilon^2}\right)$ and $l = O\left(\frac{\lambda^2 p}{\epsilon^2}\right)$.

Bimatrix games

I \ II	c	d
a	1, 0.3	0, 0.8
b	0.1, 1	1, 0.5

- Two players: row, column
- Each player has n pure strategies
- R, C are $n \times n$ matrices
- Row player plays x, column player plays y
- Payoff functions
  - Row: $x^T R y$
  - Column: $x^T C y$

Penalty Games

- Extension of bimatrix games
- Payoff functions
  - Row: $T_r(x, y) = x^T R y - f_r(x)$
  - Column: $T_c(x, y) = x^T C y - f_c(y)$

Penalty Games

- Extension of bimatrix games
- Payoff functions
  - Row: $T_r(x, y) = x^T R y - f_r(x)$
  - Column: $T_c(x, y) = x^T C y - f_c(y)$
- $f_r(x)$ and $f_c(y)$ are continuous, non-linear functions

Penalty Games

- Extension of bimatrix games
- Payoff functions
  - Row: $T_r(x, y) = x^T Ry - f_r(x)$
  - Column: $T_c(x, y) = x^T Cy - f_c(y)$
- $f_r(x)$ and $f_c(y)$ are continuous, non-linear functions

$\mathcal{P}_{\lambda_p}$: penalty games where $f_r(x)$ and $f_c(y)$ are λ_p-Lipschitz continuous

Our result

Theorem 3 (Existence)

For any penalty game in the class $\mathcal{P}_{\lambda_p}$ that possesses an equilibrium, any $\epsilon > 0$, and any $k \in \frac{\Omega(\lambda^2 \log n)}{\epsilon^2}$, there exists a k-uniform strategy profile that is an ϵ-equilibrium.

Our result

Theorem 3 (Existence)

For any penalty game in the class $\mathcal{P}_{\lambda_p}$ that possesses an equilibrium, any $\epsilon > 0$, and any $k \in \frac{\Omega(\lambda^2 \log n)}{\epsilon^2}$, there exists a k-uniform strategy profile that is an ϵ-equilibrium.

Theorem 4 (Computation)

In any penalty game $\mathcal{P}_{\lambda_p}$ and any $\epsilon > 0$, in quasi polynomial time we can either compute a ϵ-equilibrium, or decide that $\mathcal{P}_{\lambda_p}$ does not posses an exact equilibrium.

Biased Games

- Subclass of penalty games
- Base strategies: $\boldsymbol{s}$ and $\boldsymbol{t}$
- Payoff functions
  - Row: $T_r(x, y) = x^T R y - d_r \cdot ||x - s||_p$
  - Column: $T_c(x, y) = x^T C y - d_c \cdot ||y - t||_p$

Biased Games

- Subclass of penalty games
- Base strategies: $\boldsymbol{s}$ and $\boldsymbol{t}$
- Payoff functions
  - Row: $T_r(x, y) = x^T R y - d_r \cdot ||x - s||_p$
  - Column: $T_c(x, y) = x^T C y - d_c \cdot ||y - t||_p$

$$T_r(x, y) = x^T R y - \frac{1}{3} \cdot \left\| x - [\tfrac{1}{2}, \tfrac{1}{2}]^T \right\|_2^2$$

Biased Games

- Subclass of penalty games
- Base strategies: s and t
- Payoff functions
  - Row: $T_r(x, y) = x^T Ry - d_r \cdot ||x - s||_p$
  - Column: $T_c(x, y) = x^T Cy - d_c \cdot ||y - t||_p$

We study three norms: L_1, L_2^2, L_∞

Approximation guarantee

$$L_1 : \quad 2/3$$

$$L_2^2 : \quad 5/7$$

$$L_\infty : \quad 2/3$$

A polynomial time approximation aglorithm

Generalization of DMP algorithm

A polynomial time approximation aglorithm

Generalization of DMP algorithm

The Base Algorithm

1. Compute a best response $\boldsymbol{y}^*$ against $\boldsymbol{s}$.
2. Compute a best response $\boldsymbol{x}$ against $\boldsymbol{y}^*$.
3. Set $\boldsymbol{x}^* = \delta \cdot \boldsymbol{s} + (1 - \delta) \cdot \boldsymbol{x}$, for some $\delta \in [0, 1]$.
4. Return the strategy profile $(\boldsymbol{x}^*, \boldsymbol{y}^*)$.

A polynomial time approximation aglorithm

Generalization of DMP algorithm

The Base Algorithm

1. Compute a best response $\boldsymbol{y}^*$ against $\boldsymbol{s}$.
2. Compute a best response $\boldsymbol{x}$ against $\boldsymbol{y}^*$.
3. Set $\boldsymbol{x}^* = \delta \cdot \boldsymbol{s} + (1 - \delta) \cdot \boldsymbol{x}$, for some $\delta \in [0, 1]$.
4. Return the strategy profile $(\boldsymbol{x}^*, \boldsymbol{y}^*)$.

Not trivial how to compute best responses

A polynomial time approximation aglorithm

Generalization of DMP algorithm

The Base Algorithm

1. Compute a best response $\boldsymbol{y}^*$ against $\boldsymbol{s}$.
2. Compute a best response $\boldsymbol{x}$ against $\boldsymbol{y}^*$.
3. Set $\boldsymbol{x}^* = \delta \cdot \boldsymbol{s} + (1 - \delta) \cdot \boldsymbol{x}$, for some $\delta \in [0, 1]$.
4. Return the strategy profile $(\boldsymbol{x}^*, \boldsymbol{y}^*)$.

Not trivial how to compute best responses

We derive simple combinatorial algorithms
for computing best responses.

Simple best response algorithm

Best Response Algorithm for L_∞ penalty

1. For all $i \in \mathcal{L}$, set $x_i = 0$.
2. If $\mathcal{P} \leq |\mathcal{H}| \cdot s_{\max}$, then set $x_i = s_i + \frac{\mathcal{P}}{|\mathcal{H}|}$ for all $i \in \mathcal{H}$ and $x_j = s_j$ for $j \in \mathcal{M}$.
3. Else if $\mathcal{P} < |\mathcal{H} \cup \mathcal{M}| \cdot s_{\max}$, then
   - Set $x_i = s_i + s_{\max}$ for all $i \in \mathcal{H}$.
   - Set $k = \lfloor \frac{\mathcal{P} - |\mathcal{H}| \cdot p_{\max}}{p_{\max}} \rfloor$.
   - Set $x_i = s_i + s_{\max}$ for all $i \leq |\mathcal{H}| + k$.
   - Set $x_{|\mathcal{H}|+k+1} = s_{|\mathcal{H}|+k+1} + \mathcal{P} - (|\mathcal{H}| + k) \cdot s_{\max}$.
   - Set $x_j = s_j$ for all $|\mathcal{H}| + k + 2 \leq j \leq |\mathcal{H}| + |\mathcal{M}|$.
4. Else set $x_i = s_i + \frac{\mathcal{P}}{|\mathcal{H} \cup \mathcal{M}|}$ for all $i \in \mathcal{H} \cup \mathcal{M}$.

Open questions

- Exact complexity for Lipschitz and biased games?
  PPAD is not suitable. ***FIXP*?**
- Better polynomial-time approximation algorithms
- Tractable cases? (Zero sum biased games)

Open questions

- Exact complexity for Lipschitz and biased games?
  PPAD is not suitable. ***FIXP***?
- Better polynomial-time approximation algorithms
- Tractable cases? (Zero sum biased games)

THANK YOU!