Computing Approximate Nash Equilibria of Polymatrix Games

Argyrios Deligkas John Fearnley Rahul Savani Paul Spirakis

Department of Computer Science
University of Liverpool

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

(additive notion of approximation)

Theorem (Rubinstein 2014)

There exists a **constant** ϵ *such that it is* $\mathtt{PPAD}$*-hard to find an* ϵ*-Nash equilibrium of* $\boldsymbol{n}$*-player polymatrix games.*

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)

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: bimatrix games

What is the smallest ϵ such that an ϵ-Nash equilibrium can be computed in polynomial time (payoffs in $[0, 1]$)?

Background: bimatrix games

What is the smallest ϵ such that an ϵ-Nash equilibrium can be computed in polynomial time (payoffs in $[0,1]$)?

HISTORY:

0.5	Daskalakis Mehta Papadimitriou (WINE 06)
0.382	DMP (EC 2007)
0.364	Bosse Byrka Markakis (WINE 07)
0.339	Tsaknakis Spirakis (WINE 07)

Tsaknakis & Spirakis use **gradient descent**

Background: many-player games

- **Two players**: **0.3393** [Tsaknakis and Spirakis]
- ***n* players**: $\mathbf{1 - 1/n}$ [extension of DMP]

Background: many-player games

- **Two players**: **0.3393** [Tsaknakis and Spirakis]
- ***n* players**: $\mathbf{1 - 1/n}$ [extension of DMP]
- DMP idea extends solution for $\mathbf{n - 1}$ players to $\mathbf{n}$ players:
  - **Three players**: **0.6022**
  - **Four players**: **0.7153**
  - Guarantee goes to **1** as $\mathbf{n}$ goes to infinity

Background: many-player games

- **Two players**: **0.3393** [Tsaknakis and Spirakis]
- ***n* players**: $\mathbf{1 - 1/n}$ [extension of DMP]
- DMP idea extends solution for $\mathbf{n - 1}$ players to ***n*** players:
  - **Three players**: **0.6022**
  - **Four players**: **0.7153**
  - Guarantee goes to **1** as ***n*** goes to infinity

Our result is for the class of *n*-player polymatrix games:

$(\mathbf{0.5} + \delta)$ in time polynomial in the input size and $\mathbf{1}/\delta$

Polymatrix games

	a_2	b_2
a_1	1,0	0,1
b_1	1,0	1,1
c_1	1,0	0,1
d_1	0,1	1,0

	a_4	b_4	c_4
a_2	1,0	0,1	0,0
b_2	0,1	1,0	0,1

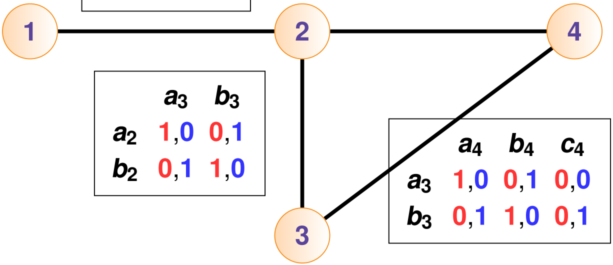

	a_3	b_3
a_2	1,0	0,1
b_2	0,1	1,0

	a_4	b_4	c_4
a_3	1,0	0,1	0,0
b_3	0,1	1,0	0,1

- **Scale** the payoffs s.t. $\mathbf{0} \leq u_i \leq \mathbf{1}$

Polymatrix games

- **Game graph** $G = (V, E)$
- Player $i \in V$ has m_i pure strategies
- $m_i \times m_j$ **bimatrix** game (A_{ij}, A_{ji}) for all $(i, j) \in E$
- A **single strategy** x_i of i is played in all games
- The **payoff** of i is the sum of payoffs in all games

$$u_i(\mathbf{x}) := x_i^T \sum_{j \in N(i)} A_{ij} x_j$$

Polymatrix games known results

- Zero sum polymatrix games are easy [Daskalakis and Papadimitriou]
- Polymatrix games with strictly competitive games on the edges are PPAD-complete to solve [Cai and Daskalakis]
- Pure Nash equilibria in coordination-only polymatrix games is PLS-complete [Cai and Daskalakis]
- Combination of coordination and zero-sum games on the edges, the problem is PPAD-complete [Cai and Daskalakis]

Gradient descent on max regret

- **Regret** for player i: $f_i(x) := u_i^*(x) - u_i(x)$

Gradient descent on max regret

- **Regret** for player i: $f_i(x) := u_i^*(x) - u_i(x)$

Definition (Max regret)

The regret for strategy profile x is:

$$f(x) := \max_{i \in players} f_i(x)$$

Gradient descent on max regret

- **Regret** for player $\boldsymbol{i}$: $f_i(x) := u_i^*(x) - u_i(x)$

Definition (Max regret)

The regret for strategy profile $\boldsymbol{x}$ is:

$$f(x) := \max_{i \in players} f_i(x)$$

$f(x) = \epsilon \Leftrightarrow$ $\boldsymbol{x}$ is an ϵ-Nash equilibrium

Natural definition of gradient

Definition

Given profiles $\mathbf{x}, \mathbf{x}'$ and $\lambda \in [0, 1]$, we define:

$$Df(\mathbf{x}, \mathbf{x}', \lambda) := f((1 - \lambda) \cdot \mathbf{x} + \lambda \cdot \mathbf{x}') - f(\mathbf{x}).$$

Then, we define the gradient of f at $\mathbf{x}$ in the direction $\mathbf{x}' - \mathbf{x}$ as:

$$Df(\mathbf{x}, \mathbf{x}') = \lim_{\lambda \to 0} \frac{1}{\lambda} Df(\mathbf{x}, \mathbf{x}', \lambda).$$

Useful definition of gradient

For each $\mathbf{x}, \mathbf{x}'$, and for each player $i \in [n]$, we define:

$$Df_i(\mathbf{x}, \mathbf{x}') := \max_{k \in \mathrm{Br}_i(\mathbf{x})} \left\{ \left(v_i(\mathbf{x}') \right)_k \right\} - u_i(x_i, \mathbf{x}') + u_i(x_i - x_i', \mathbf{x})$$

$\mathcal{K}(\mathbf{x})$ set of players that have maximum regret under strategy profile $\mathbf{x}$

Useful definition of gradient

For each $\mathbf{x}, \mathbf{x}'$, and for each player $i \in [n]$, we define:

$$Df_i(\mathbf{x}, \mathbf{x}') := \max_{k \in \mathrm{Br}_i(\mathbf{x})} \left\{ \left(v_i(\mathbf{x}') \right)_k \right\} - u_i(x_i, \mathbf{x}') + u_i(x_i - x_i', \mathbf{x})$$

$\mathcal{K}(\mathbf{x})$ set of players that have maximum regret under strategy profile $\mathbf{x}$

Lemma

The gradient of $\mathbf{f}$ *at point* $\mathbf{x}$ *along direction* $\mathbf{x}' - \mathbf{x}$ *is:*

$$Df(\mathbf{x}, \mathbf{x}') = \max_{i \in \mathcal{K}(\mathbf{x})} Df_i(\mathbf{x}, \mathbf{x}') - f(\mathbf{x})$$

Useful definition of gradient

Lemma

The gradient of $\mathbf{f}$ *at point* $\mathbf{x}$ *along direction* $\mathbf{x}' - \mathbf{x}$ *is:*

$$Df(\mathbf{x}, \mathbf{x}') = \max_{i \in \mathcal{K}(\mathbf{x})} Df_i(\mathbf{x}, \mathbf{x}') - f(\mathbf{x})$$

- We can compute the steepest descent direction with an LP

Steepest descent LP

Definition (Steepest descent linear program)

Given a strategy profile $\mathbf{x}$, find $\mathbf{x}' \in \boldsymbol{\Delta}$, $l_1, l_2, \ldots, l_{|\mathcal{K}(\mathbf{x})|}$, and w such that:

$$
\begin{aligned}
\text{minimize} \quad & w \\
\text{subject to} \quad & \left(v_i(\mathbf{x}')\right)_k \leq l_i && \forall k \in \mathrm{Br}_i^{\delta}(\mathbf{x}), \quad \forall i \in \mathcal{K}(\mathbf{x}) \\
& l_i - u_i(x_i, \mathbf{x}') - u_i(x'_i, \mathbf{x}) + u_i(\mathbf{x}) \leq w && \forall i \in \mathcal{K}(\mathbf{x}) \\
& \mathbf{x}' \in \boldsymbol{\Delta}.
\end{aligned}
$$

The algorithm

1. Choose an arbitrary strategy profile $\mathbf{x}$
2. Solve **steepest descent LP** with input $\mathbf{x}$ to obtain $\mathbf{x}'$
3. Set $\mathbf{x} := \mathbf{x} + \lambda(\mathbf{x}' - \mathbf{x})$, where $\lambda = \frac{\delta}{\delta + 2}$
4. If $f(\mathbf{x}) \leq 0.5 + \delta$ then stop, otherwise go to step 2

The algorithm (differences with TS)

1. Choose an arbitrary strategy profile $\mathbf{x}$
2. Solve **steepest descent LP** with input $\mathbf{x}$ to obtain $\mathbf{x}'$
3. Set $\mathbf{x} := \mathbf{x} + \lambda(\mathbf{x}' - \mathbf{x})$, where $\lambda = \frac{\delta}{\delta+2}$
4. If $f(\mathbf{x}) \leq 0.5 + \delta$ then stop, otherwise go to step 2

- TS equalize the regrets before each iteration of steepest descent LP
- 0.3393 guarantee is achieved by considering some extra points too
- Not obvious how to do these in our case

Our result

Theorem

A $(0.5 + \delta)$-Nash equilibrium of a polymatrix game can be found in time polynomial in the size of the game and in $1/\delta$.

Proof sketch:

- We do not get stuck at a bad point: Every δ-stationary point $\mathbf{x}^*$ of f is a $(0.5 + \delta)$-NE, i.e., $f(\mathbf{x}^*) \leq 0.5 + \delta$
- Each descent step makes enough progress in reducing f, so that after polynomially many iterations $f(\mathbf{x}) \leq 0.5 + \delta$

Approximate stationary points

Definition (δ-stationary point)

For $\delta > 0$ we say that $\mathbf{x}^*$ is a δ-stationary point if for all $\mathbf{x}'$:

$$Df^{\delta}(\mathbf{x}^*, \mathbf{x}') \geq -\delta$$

Proof that x^* is a $(0.5 + \delta)$-NE

Proof that $\mathbf{x}^*$ is a $(0.5 + \delta)$-NE

$\mathbf{x}^*$ is a δ-stationary point if for all $\mathbf{x}'$:

$$Df^{\delta}(\mathbf{x}^*, \mathbf{x}') = \max_{i \in \mathcal{K}(\mathbf{x}^*)} Df_i^{\delta}(\mathbf{x}^*, \mathbf{x}') - f(\mathbf{x}^*) \geq -\delta$$

Proof that $\mathbf{x}^*$ is a $(0.5 + \delta)$-NE

$\mathbf{x}^*$ is a δ-stationary point if for all $\mathbf{x}'$:

$$Df^{\delta}(\mathbf{x}^*, \mathbf{x}') = \max_{i \in \mathcal{K}(\mathbf{x}^*)} Df_i^{\delta}(\mathbf{x}^*, \mathbf{x}') - f(\mathbf{x}^*) \geq -\delta$$

So we have **for all** $\mathbf{x}'$: $f(\mathbf{x}^*) \leq \max_{i \in \mathcal{K}(\mathbf{x}^*)} Df_i^{\delta}(\mathbf{x}^*, \mathbf{x}') + \delta$

Proof that $\mathbf{x}^*$ is a $(0.5 + \delta)$-NE

$\mathbf{x}^*$ is a δ-stationary point if for all $\mathbf{x}'$:

$$Df^{\delta}(\mathbf{x}^*, \mathbf{x}') = \max_{i \in \mathcal{K}(\mathbf{x}^*)} Df_i^{\delta}(\mathbf{x}^*, \mathbf{x}') - f(\mathbf{x}^*) \geq -\delta$$

So we have **for all** $\mathbf{x}'$: $f(\mathbf{x}^*) \leq \max_{i \in \mathcal{K}(\mathbf{x}^*)} Df_i^{\delta}(\mathbf{x}^*, \mathbf{x}') + \delta$

Lemma

For every $\mathbf{x}^*$, $\exists \mathbf{x}'$ *such that:* $\max_{i \in \mathcal{K}(\mathbf{x}^*)} Df_i^{\delta}(\mathbf{x}^*, \mathbf{x}') \leq 0.5$

Proof that $\mathbf{x}^*$ is a $(0.5 + \delta)$-NE

Lemma

For every $\mathbf{x}^*$, $\exists \mathbf{x}'$ *such that:* $\max_{i \in \mathcal{K}(\mathbf{x}^*)} Df_i^{\delta}(\mathbf{x}^*, \mathbf{x}') \leq 0.5$

Let $\bar{\mathbf{x}}$ be the best response profile against $\mathbf{x}^*$. Set $\mathbf{x}' = \frac{\bar{\mathbf{x}}+\mathbf{x}^*}{2}$.

$$
\begin{aligned}
& \max_{k \in \mathrm{Br}_i^{\delta}(\mathbf{x}^*)} \left\{ \left(v_i\left(\frac{\bar{\mathbf{x}}+\mathbf{x}^*}{2}\right)\right)_k \right\} - u_i\left(x_i^*, \frac{\bar{\mathbf{x}}+\mathbf{x}^*}{2}\right) - u_i\left(\frac{\bar{x}_i + x_i^*}{2}, \mathbf{x}^*\right) + u_i(x_i^*, \mathbf{x}^*) \\
&= \frac{1}{2} \cdot \max_{k \in \mathrm{Br}_i^{\delta}(\mathbf{x}^*)} \left\{ \left(v_i(\bar{\mathbf{x}}+\mathbf{x}^*)\right)_k \right\} - \frac{1}{2} \cdot u_i(x_i^*, \bar{\mathbf{x}}) - \frac{1}{2} \cdot u_i(\bar{x}_i, \mathbf{x}^*) \\
&\leq \frac{1}{2} \cdot \left(\max_{k \in \mathrm{Br}_i^{\delta}(\mathbf{x}^*)} \left\{ \left(v_i(\bar{\mathbf{x}})\right)_k \right\} + \max_{k \in \mathrm{Br}_i^{\delta}(\mathbf{x}^*)} \left\{ \left(v_i(\mathbf{x}^*)\right)_k \right\} - u_i(x_i^*, \bar{\mathbf{x}}) - u_i(\bar{x}_i, \mathbf{x}^*) \right) \\
&= \frac{1}{2} \cdot \left(\max_{k \in \mathrm{Br}_i^{\delta}(\mathbf{x}^*)} \left\{ \left(v_i(\bar{\mathbf{x}})\right)_k \right\} - u_i(x_i^*, \bar{\mathbf{x}}) \right) \quad \text{because } \bar{x}_i \text{ is a b.r. to } x^* \\
&\leq \frac{1}{2} \cdot \max_{k \in \mathrm{Br}_i^{\delta}(\mathbf{x}^*)} \left\{ \left(v_i(\bar{\mathbf{x}})\right)_k \right\} \\
&\leq \frac{1}{2}.
\end{aligned}
$$

Proof: enough progress

Lemma

Fix $\mathbf{0} < \delta \leq \mathbf{0.5}$. *Either* $\mathbf{x}$ *is a* δ*-stationary point or:*

$$f(\mathbf{x}_{new}) \leq \left(1 - \left(\frac{\delta}{\delta + 2}\right)^2\right) f(\mathbf{x})$$

Proof: enough progress

Lemma

Fix $0 < \delta \leq 0.5$. *Either* $\mathbf{x}$ *is a* δ*-stationary point or:*

$$f(\mathbf{x}_{new}) \leq \left(1 - \left(\frac{\delta}{\delta + 2}\right)^2\right) f(\mathbf{x})$$

Lemma

After $O(\frac{1}{\delta^2})$ *iterations we have* $f(\mathbf{x}) \leq 0.5 + \delta$

2-player Bayesian games

- Row player
  - k_1 pure strategies
  - m possible types
- Column player
  - k_2 pure strategies
  - n possible types
- A type is chosen for each player according to a known distribution
- When Row player has type i and Column has type j, then the bimatrix game (R_{ij}, C_{ij}) is played.
- Each player choose a strategy according to his type

2-player Bayesian games

- Each player choose a strategy according to his type
  - Row player plays x_i when has type $i \in [m]$:

$$\mathbf{x} = (x_1, \ldots, x_m)$$

  - Column player plays y_j when has type $j \in [n]$:

$$\mathbf{y} = (y_1, \ldots, y_n)$$

Definition (Approximate Bayes Nash Equilibrium (ϵ-BNE))

The profile $(\mathbf{x}, \mathbf{y})$ is an ϵ-BNE iff the following conditions hold:

$$u_R(x_i, \mathbf{y}) \geq u_R(x_i', \mathbf{y}) - \epsilon \quad \text{for all } x_i' \quad \text{for all } i \in [m],$$
$$u_C(\mathbf{x}, y_j) \geq u_C(\mathbf{x}, y_j') - \epsilon \quad \text{for all } y_j' \quad \text{for all } j \in [n].$$

Application: 2-player Bayesian games

Can be written as a **complete bipartite polymatrix game**

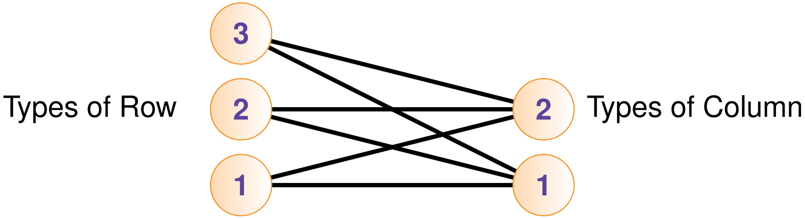

Theorem

Every ϵ-NE of the bipartite polymatrix game constructed above is and ϵ-BNE for the original Bayesian game.

Open questions

Better **upper bounds**:

- Constant number of players or strategies
- ϵ-well-supported approximate equilibria

Lower bounds:

It is $\mathtt{PPAD}$-hard to find an ϵ-Nash equilibrium of a polymatrix game for a **constant but very small ϵ** [Rubinstein]

Improve the value of ϵ in such a lower bound