Algorithms for Computing Approximate Nash Equilibria

Vangelis Markakis

Athens University of Economics and Business

Outline

- Introduction to Games
  - The concepts of Nash and ε-Nash equilibrium
- Computing approximate Nash equilibria
  - A subexponential algorithm for any constant $\varepsilon > 0$
  - Polynomial time approximation algorithms
- Conclusions

What is Game Theory?

- Game Theory aims to help us understand situations in which *decision makers* interact
- Goals:
  - Mathematical models for capturing the properties of such interactions
  - Prediction (given a model how should/would a *rational agent* act?)

Rational agent: when given a choice, the agent always chooses the option that yields the highest utility

Models of Games

- Cooperative or noncooperative
- Simultaneous moves or sequential
- Finite or infinite
- Complete information or incomplete information

In this talk:

- Cooperative or noncooperative
- Simultaneous moves or sequential
- Finite or infinite
- Complete information or incomplete information

Noncooperative Games in Normal Form

The Hawk-Dove game

Column Player

Row player

		
	2, 2	0, 4
	4, 0	-1, -1

Example 2: The Bach or Stravinsky game (BoS)



	2, 1	0, 0
	0, 0	1, 2

Example 3: A Routing Game

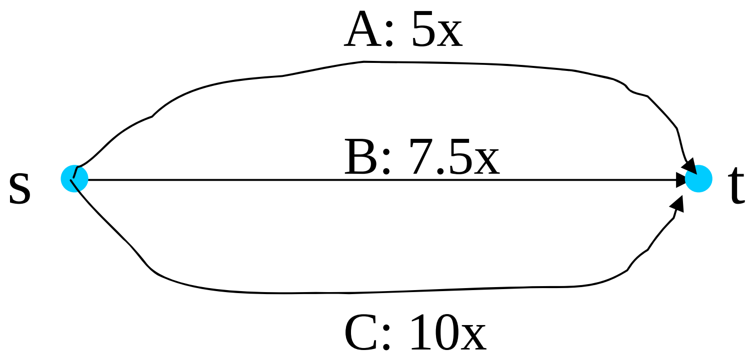

Example 3: A Routing Game

	A	B	C
A	10, 10	5, 7.5	5, 10
B	7.5, 5	15, 15	7.5, 10
C	10, 5	10, 7.5	20, 20

Definitions

- 2-player game (*R*, *C*):
  - *n* available *pure strategies* for each player
  - *n* x *n* payoff matrices *R*, *C*
  - *i*, *j* played $\Rightarrow$ payoffs : R_{ij} , C_{ij}
- Mixed strategy: Probability distribution over [*n*]

$$x = (x_1, \ldots, x_n), \ \sum x_i = 1, \ x_i \geq 0$$

- Expected payoffs : (x, Ry) and (x, Cy)

$$(x, Ry) = \sum_{i,j} x_i y_j R_{ij}$$

Solution Concept

x^*, y^* is a Nash equilibrium if no player has a unilateral incentive to deviate:

$$(x, Ry^*) \le (x^*, Ry^*) \quad \forall\, x$$

$$(x^*, Cy) \le (x^*, Cy^*) \quad \forall\, y$$

[Nash, 1951]: Every finite game has a mixed strategy equilibrium.

(think of it as a steady state)

Proof: Based on Brouwer's fixed point theorem.

Solution Concept

x^*, y^* is a Nash equilibrium if no player has a unilateral incentive to deviate:

$$(x, Ry^*) \le (x^*, Ry^*) \quad \forall\, x$$

$$(x^*, Cy) \le (x^*, Cy^*) \quad \forall\, y$$

[Nash, 1951]: Every finite game has a mixed strategy equilibrium.

(think of it as a steady state)

Proof: Based on Brouwer's fixed point theorem.

Solution Concept

It suffices to consider only deviations to *pure strategies*

Let $x^i = (0, 0, \ldots, 1, 0, \ldots, 0)$ be the ith pure strategy

x^*, y^* is a Nash equilibrium if no player has a unilateral incentive to deviate to a pure strategy:

$$(x^i, Ry^*) \leq (x^*, Ry^*) \quad \forall\, x^i$$

$$(x^*, Cy^j) \leq (x^*, Cy^*) \quad \forall\, y^j$$

Example: The Hawk-Dove Game

Column Player

			
Row player		2, 2	0, 4
		4, 0	-1, -1

Example 2: The Bach or Stravinsky game (BoS)

		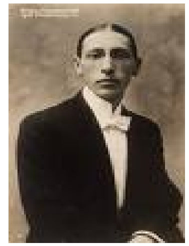
	2, 1	0, 0
	0, 0	1, 2

3 equilibrium points:

1. (B, B)
2. (S, S)
3. ((2/3, 1/3), (1/3, 2/3))

Complexity issues

- m = 2 players, known algorithms: worst case exponential time
  [Kuhn ’61, Lemke, Howson ’64, Mangasarian ’64, Lemke ’65]

- If **NP**-hard $\Rightarrow$ **NP** = co-**NP** [Megiddo, Papadimitriou ’89]
- **NP**-hard if we add more constraints (e.g. maximize sum of payoffs)
  [Gilboa, Zemel ’89, Conitzer, Sandholm ’03]

- Representation problems

  m = 3, there exist games with rational data BUT irrational equilibria
  [Nash ’51]

- **PPAD**-complete even for m = 2
  [Daskalakis, Goldberg, Papadimitriou ’06, Chen, Deng, Teng ’06]
  Poly-time equivalent to:
  - finding approximate fixed points of continuous maps on convex and compact domains

Approximate Nash Equilibria

- Recall definition of Nash eq. :

$$(x, Ry^*) \le (x^*, Ry^*) \quad \forall\, x$$

$$(x^*, Cy) \le (x^*, Cy^*) \quad \forall\, y$$

- ε-Nash equilibria (incentive to deviate $\le \varepsilon$) :

$$(x, Ry^*) \le (x^*, Ry^*) + \varepsilon \quad \forall\, x$$

$$(x^*, Cy) \le (x^*, Cy^*) + \varepsilon \quad \forall\, y$$

Normalization: entries of R, C in [0,1]

Searching for Approximate Equilibria

Definition: A k-uniform strategy is a strategy where all probabilities are integer multiples of $1/k$

e.g. $(3/k, 0, 0, 1/k, 5/k, 0, \ldots, 6/k)$

[Lipton, M., Mehta '03]: For any ε in (0,1), and for every $k \geq 9 \log n/\varepsilon^2$, there exists a pair of k-uniform strategies x, y that form an ε-Nash equilibrium.

A Subexponential Algorithm (Quasi-PTAS)

Definition: A *k*-uniform strategy is a strategy where all probabilities are integer multiples of $1/k$

e.g. $(3/k, 0, 0, 1/k, 5/k, 0, \ldots, 6/k)$

[Lipton, M., Mehta '03]: For any ε in (0,1), and for every $k \geq 9 \log n/\varepsilon^2$, there exists a pair of *k*-uniform strategies *x*, *y* that form an ε-Nash equilibrium.

Corollary : We can compute an ε-Nash equilibrium in time $n^{O(\log n/\epsilon^2)}$

Proof: There are $n^{O(k)}$ pairs of strategies to look at. Verify ε-equilibrium condition.

Proof of Existence

Based on the probabilistic method (sampling)

Let x^*, y^* be a Nash equilibrium.

- Sample k times from the set of pure strategies of the row player, independently, at random, according to x^* $\Rightarrow$ k-uniform strategy x

- Same for column player $\Rightarrow$ k-uniform strategy y

Suffices to show $\Pr[x, y \text{ form an } \varepsilon\text{-Nash eq.}] > 0$

Proof (cont'd)

Enough to consider deviations to pure strategies

$(x^i, Ry) \leq (x, Ry) + \varepsilon \quad \forall i$

(x^i, Ry): sum of k random variables with mean (x^i, Ry^*)

Chernoff-Hoeffding bounds $\Rightarrow (x^i, Ry) \approx (x^i, Ry^*)$ with high probability

$(x^i, Ry) \approx (x^i, Ry^*) \quad \leq \quad (x^*, Ry^*) \quad \approx \quad (x, Ry)$

Finally when $k = \Omega(\log n/\varepsilon^2)$:

Pr[$\exists$ deviation with gain more than ε] = $O(n)e^{-k\epsilon^2/8} < 1$

Multi-player Games

For *m* players, same technique:

support size: $k = O(m^2 \log(m^2 n)/\varepsilon^2)$

running time: $\exp(\log n, m, 1/\varepsilon)$

Previously [Scarf '67]: $\exp(n, m, \log(1/\varepsilon))$
(fixed point approximation)

[Lipton, M. '04]: $\exp(n, m)$ but $\text{poly}(\log(1/\varepsilon))$
(using algorithms for polynomial equations)

Outline

- Introduction to Games
  - The concepts of Nash and ε-Nash equilibrium
- Computing approximate Nash equilibria
  - A subexponential algorithm for any constant $\varepsilon > 0$
  - Polynomial time approximation algorithms
- Conclusions

Polynomial Time Approximation Algorithms

For $\varepsilon = 1/2$:

- Pick arbitrary row i
- Let j = best response to i
- Find k = best response to j, play i or k with prob. 1/2

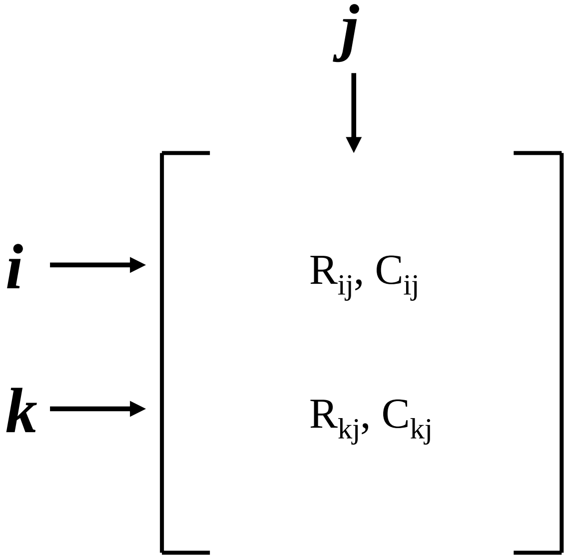

Feder, Nazerzadeh, Saberi '07: For $\varepsilon < 1/2$, we need support at least $\Omega(\log n)$

Polynomial Time Approximation Algorithms

Daskalakis, Mehta, Papadimitriou (EC '07):
in **P** for $\varepsilon = 1\text{-}1/\varphi = (3\text{-}\sqrt{5})/2 \approx 0.382$ (φ = golden ratio)

- Based on sampling + Linear Programming

- Need to solve polynomial number of linear programs

Bosse, Byrka, M. (WINE '07): a different LP-based method

1. Algorithm 1: $1\text{-}1/\varphi$
2. Algorithm 2: 0.364

Running time: need to solve one linear program

Approach

0-sum games: games of the form (*R*, -*R*)

Fact: 0-sum games can be solved in polynomial time (equivalent to linear programming)

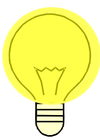

- Start with an equilibrium of the 0-sum game (*R*-*C*, *C*-*R*)

- If incentives to deviate are “high”, players take turns and adjust their strategies via best response moves

Similar idea used in [Kontogiannis, Spirakis ’07] for a different notion of approximation

Algorithm 1

Parameters: α, $\delta_2 \in [0,1]$

1. Find an equilibrium x^*, y^* of the 0-sum game $(R - C, C - R)$
2. Let g_1, g_2 be the incentives to deviate for row and column player respectively. Suppose $g_1 \geq g_2$
3. If $g_1 \leq \alpha$, output x^*, y^*
4. Else: let b_1 = best response to y^*, b_2 = best response to b_1
5. Output:

$$x = b_1$$

$$y = (1 - \delta_2)\, y^* + \delta_2\, b_2$$

Theorem: Algorithm 1 with $\alpha = 1\text{-}1/\varphi$ and $\delta_2 = (1\text{-}\ g_1) / (2\text{-}\ g_1)$ achieves a $(1\text{-}1/\varphi)$-approximation

Analysis of Algorithm 1

Why start with an equilibrium of $(R - C, C - R)$?

Intuition: If row player profits from a deviation from x^* then column player also gains at least as much

Case 1: $g_1 \le \alpha \Rightarrow \alpha$-approximation

Case 2: $g_1 > \alpha$

Incentive to deviate:

for row player $\le \delta_2$

for column player $\le (1 - \delta_2)(1 - (b_1, Cy^*))$

$\le (1 - \delta_2)(1 - g_1) = (1 - g_1) / (2 - g_1)$

$\Rightarrow$ $\max\{\alpha, (1 - \alpha)/(2 - \alpha)\}$-approximation

Analysis of Algorithm 1

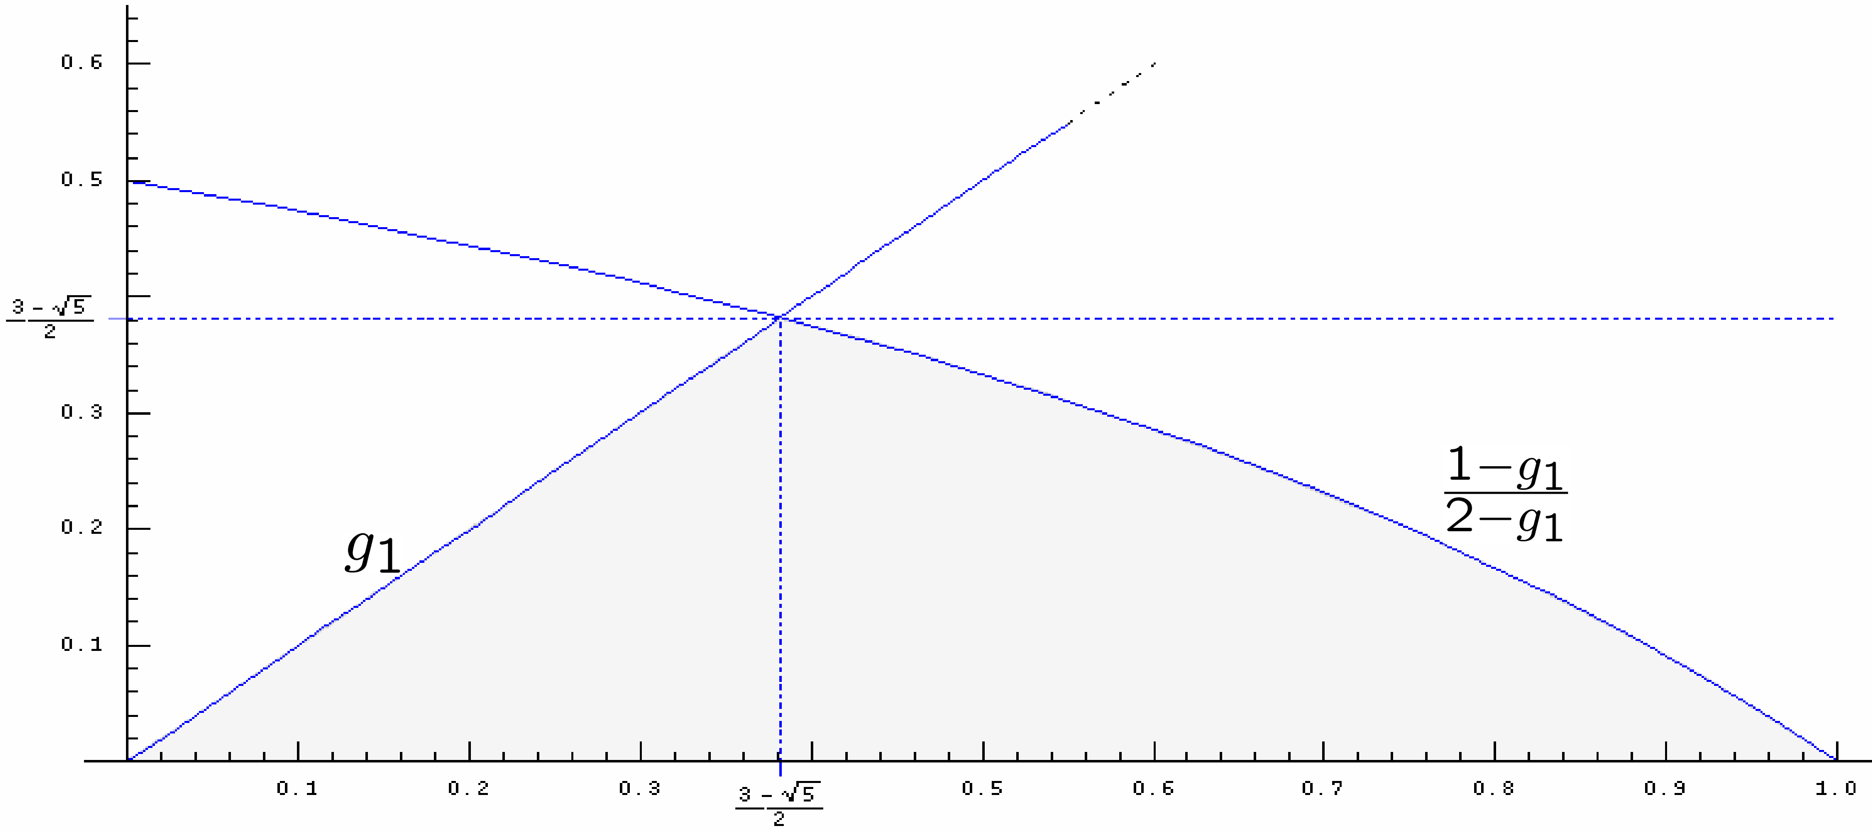

Towards a better algorithm

1. Find an equilibrium x^*, y^* of the 0-sum game $(R - C, C - R)$
2. Let g_1, g_2 be the incentives to deviate for row and column player respectively. Suppose $g_1 \geq g_2$
3. If $g_1 \leq \alpha$, output x^*, y^*
4. Else: let b_1 = best response to y^*, b_2 = best response to b_1
5. Output:

$$x = b_1$$

$$y = (1 - \delta_2)\, y^* + \delta_2\, b_2$$

Algorithm 2

1. Find an equilibrium x^*, y^* of the 0-sum game $(R - C, C - R)$
2. Let g_1, g_2 be the incentives to deviate for row and column player respectively. Suppose $g_1 \geq g_2$
3. If $g_1 \in [0, 1/3]$, output x^*, y^*
4. If $g_1 \in (1/3, \beta]$,
   - let r_1 = best response to y^*, $x = (1 - \delta_1)\, x^* + \delta_1\, r_1$
   - let b_2 = best response to x, $y = (1 - \delta_2)\, y^* + \delta_2\, b_2$
5. If $g_1 \in (\beta, 1]$ output:

$$x = r_1$$

$$y = (1 - \delta_2)\, y^* + \delta_2\, b_2$$

Analysis of Algorithm 2 (Reducing to an optimization question)

- We set δ_2 so as to equalize the incentives of the players to deviate

- Let $h = (x^*, Cb_2) - (x^*, Cy^*)$

Theorem: The approximation guarantee of Algorithm 2 is 0.364 and is given by:

$$\max_{g_1 \in [1/3, 1/2]} \min_{\delta_1 \in [0,1]} \max_{h \in [0, g_1]} F(g_1, \delta_1, h)$$

Analysis of Algorithm 2 (solution)

Optimization yields:

$$\delta_1(g_1) = (1 - g_1)\left(\sqrt{1 + \tfrac{1}{1-2g_1} - \tfrac{1}{g_1}} - 1\right)$$

$$\delta_2(g_1, \delta_1, h) = \tfrac{\delta_1 - g_1 + (1-\delta_1)h}{1+\delta_1 - g_1}$$

Graphically:

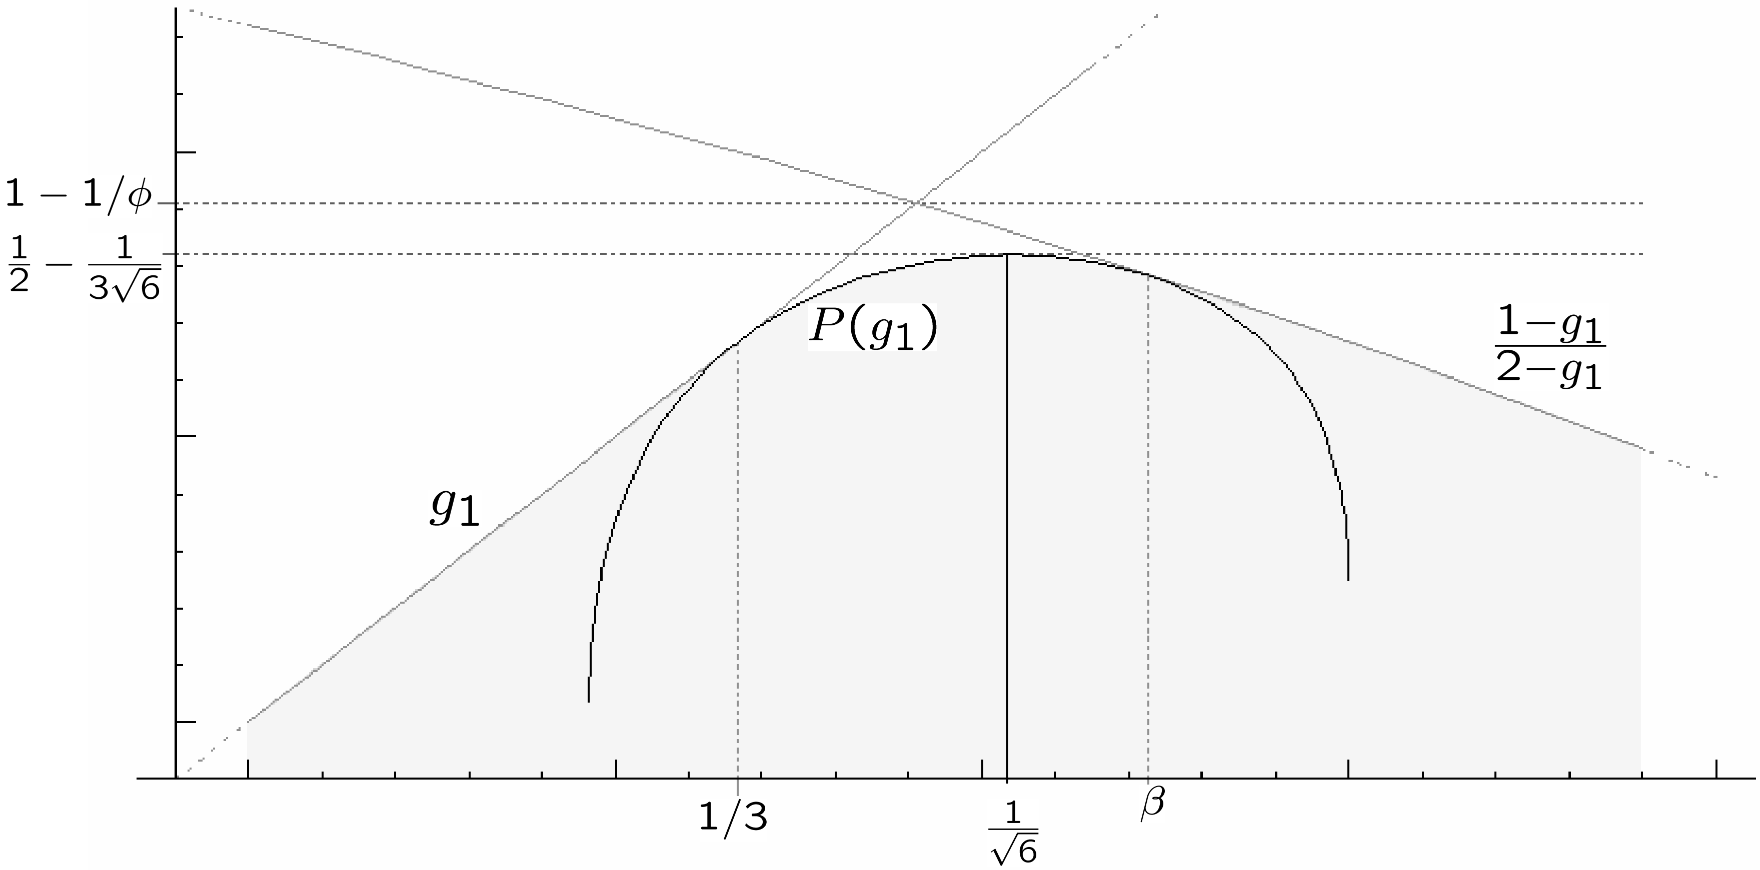

Analysis – tight example

(R, C) =

0, 0	α, α	α, α
α, α	0, 1	1, 1/2
α, α	1, 1/2	0, 1

$$\alpha = 1/\sqrt{6}$$

Remarks and Open Problems

- Spirakis, Tsaknakis (WINE ’07): currently best approximation of 0.339
  - yet another LP-based method

- Polynomial Time Approximation Scheme (PTAS)?

  Yes if:
  - rank(R) = O(1) & rank(C) = O(1) [Lipton, M. Mehta ’03]
  - rank(R+C) = O(1) [Kannan, Theobald ’06]

- **PPAD**-complete for $\varepsilon = 1/n$ [Chen, Deng, Teng ’06]

Other Notions of Approximation

- ε-well-supported equilibria: every strategy in the support is an approximate best response
  - [Kontogiannis, Spirakis ’07]: 0.658-approximation, based also on solving 0-sum games
- Strong approximation: output is geometrically close to an exact Nash equilibrium
  - [Etessami, Yannakakis ’07]: mostly negative results

Thank You!