Convergence of Nash Dynamics: Equilibria and Nearly-Optimal Solutions

Vahab Mirrokni
Google Research, New York

Heiko Röglin
Department of Quantitative Economics
Maastricht University

ACM Conference on Electronic Commerce (EC'09)
Stanford, July 2009

Outline

1. Introduction: Games, Equilibria, and Dynamics
2. Convergence to Equilibria
   - Potential Games and PLS
   - Non-Potential Games
3. Convergence to Nearly-Optimal Solutions
   - Potential Games
   - Non-Potential Games
4. Other Dynamics
   - Equilibria
   - Nearly-optimal Solutions
5. Conclusion

Outline

1. Introduction: Games, Equilibria, and Dynamics
2. Convergence to Equilibria
   - Potential Games and PLS
   - Non-Potential Games
3. Convergence to Nearly-Optimal Solutions
   - Potential Games
   - Non-Potential Games
4. Other Dynamics
   - Equilibria
   - Nearly-optimal Solutions
5. Conclusion

Games

Game:

- agents $\mathcal{N} = \{1, \ldots, n\}$
- $\forall i \in \mathcal{N}$: finite strategy space Σ_i
- $\forall i \in \mathcal{N}$: cost function $c_i \colon \Sigma_1 \times \cdots \times \Sigma_n \to \mathbb{R}$
  ($S \in \Sigma_1 \times \cdots \times \Sigma_n$ is called state.)

Games

Game:

- agents $\mathcal{N} = \{1, \ldots, n\}$ — drivers
- $\forall i \in \mathcal{N}$: finite strategy space Σ_i — possible paths from s_i to t_i
- $\forall i \in \mathcal{N}$: cost function $c_i \colon \Sigma_1 \times \cdots \times \Sigma_n \to \mathbb{R}$ — travel time

  ($S \in \Sigma_1 \times \cdots \times \Sigma_n$ is called state.)

Example: Network Congestion Games

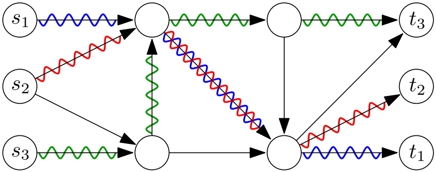

Games

Game:

- agents $\mathcal{N} = \{1, \ldots, n\}$ — drivers
- $\forall i \in \mathcal{N}$: finite strategy space Σ_i — possible paths from s_i to t_i
- $\forall i \in \mathcal{N}$: cost function $c_i \colon \Sigma_1 \times \cdots \times \Sigma_n \to \mathbb{R}$ — travel time

  ($S \in \Sigma_1 \times \cdots \times \Sigma_n$ is called state.)

Example: Network Congestion Games

latency function $\ell_e \colon \mathbb{N} \to \mathbb{R}$ for every edge e

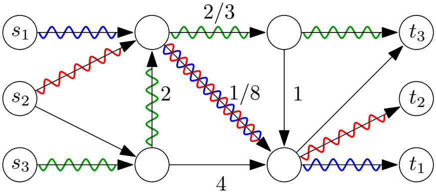

$c_1(S) = 8$

$c_2(S) = 8$

$c_3(S) = 4$

Games

Game:

- agents $\mathcal{N} = \{1, \ldots, n\}$ — drivers
- $\forall i \in \mathcal{N}$: finite strategy space Σ_i — possible paths from s_i to t_i
- $\forall i \in \mathcal{N}$: cost function $c_i \colon \Sigma_1 \times \cdots \times \Sigma_n \to \mathbb{R}$ — travel time

  ($S \in \Sigma_1 \times \cdots \times \Sigma_n$ is called state.)

Example: Network Congestion Games

latency function $\ell_e \colon \mathbb{N} \to \mathbb{R}$ for every edge e

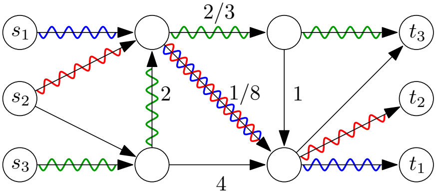

$c_1(S) = 8$

$c_2(S) = 8$

$c_3(S) = 4$

We consider only games with complete information.

Nash Equilibria

$c_1(S) = 4$
$c_2(S) = 1$
$c_3(S) = 5$

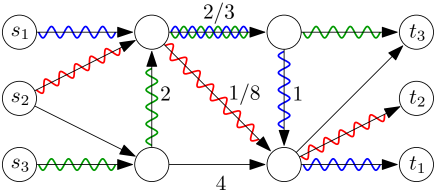

Definition

pure Nash Equilibrium $S \in \Sigma_1 \times \cdots \times \Sigma_n$

$\iff$ no player can unilaterally improve his costs in S

Nash Equilibria

$c_1(S) = 3$
$c_2(S) = 1$
$c_3(S) = 4$

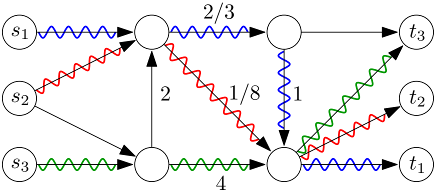

Definition

pure Nash Equilibrium $S \in \Sigma_1 \times \cdots \times \Sigma_n$

$\iff$ no player can unilaterally improve his costs in S

Nash Equilibria

$c_1(S) = 3$
$c_2(S) = 1$
$c_3(S) = 4$

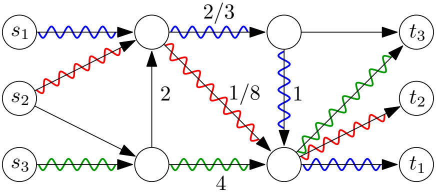

Definition

pure Nash Equilibrium $S \in \Sigma_1 \times \cdots \times \Sigma_n$
$\iff$ no player can unilaterally improve his costs in S

- Nash Equilibrium = stable
  (if players are uncoordinated, rational, selfish)
- We do not consider mixed Nash equilibria in this tutorial.

Properties of Equilibria

A lot of research on **static properties** of equilibria:

How much does society suffer from selfish behavior?

- Let cost be some measure for social cost, e.g., $\mathrm{cost}(S) = \sum_{i \in N} c_i(S)$ or $\mathrm{cost}(S) = \max_{i \in N} c_i(S)$.
- $$\text{price of anarchy} = \max_{S \in \mathrm{NE}} \frac{\mathrm{cost}(S)}{\mathrm{cost}(\mathrm{Opt})}$$

Properties of Equilibria

A lot of research on **static properties** of equilibria:
How much does society suffer from selfish behavior?

- Let cost be some measure for social cost, e.g.,
  $\mathrm{cost}(S) = \sum_{i \in N} c_i(S)$ or $\mathrm{cost}(S) = \max_{i \in N} c_i(S)$.
- $$\text{price of anarchy} = \max_{S \in \mathrm{NE}} \frac{\mathrm{cost}(S)}{\mathrm{cost}(\mathrm{Opt})}$$

Focus of this tutorial: Questions about dynamics

- Do uncoordinated agents reach an equilibrium?
- How long does it take?
- Do they quickly reach a state with small social cost?

Nash Dynamics

- **Nash Dynamics**: Sequence of best responses of players.

$c_1(S) = 8$
$c_2(S) = 8$
$c_3(S) = 2$

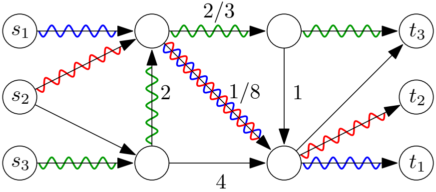

Nash Dynamics

- **Nash Dynamics**: Sequence of best responses of players.

$c_1(S) = 4$
$c_2(S) = 1$
$c_3(S) = 5$

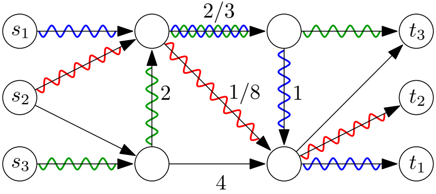

Nash Dynamics

- **Nash Dynamics**: Sequence of best responses of players.

$c_1(S) = 3$
$c_2(S) = 1$
$c_3(S) = 4$

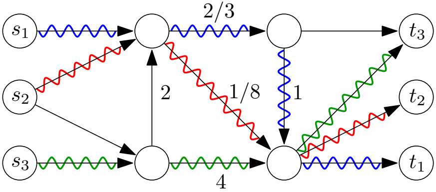

Nash Dynamics

- **Nash Dynamics**: Sequence of best responses of players.

$c_1(S) = 3$
$c_2(S) = 1$
$c_3(S) = 4$

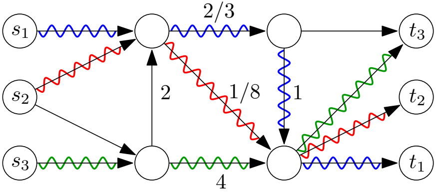

- Nash Dynamics with **Liveness Property**: Each player gets a chance to play his/her best response after at most t steps.

Nash Dynamics

- **Nash Dynamics**: Sequence of best responses of players.

$c_1(S) = 3$
$c_2(S) = 1$
$c_3(S) = 4$

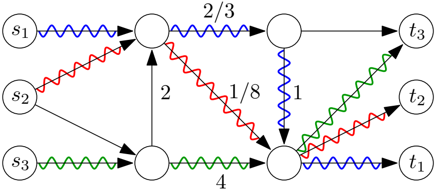

- Nash Dynamics with **Liveness Property**: Each player gets a chance to play his/her best response after at most t steps.
- **Random Nash Dynamics**: Players are chosen uniformly at random.

Nash Dynamics

- **Nash Dynamics**: Sequence of best responses of players.

$c_1(S) = 3$
$c_2(S) = 1$
$c_3(S) = 4$

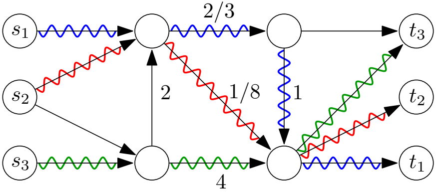

- Nash Dynamics with **Liveness Property**: Each player gets a chance to play his/her best response after at most t steps.
- **Random Nash Dynamics**: Players are chosen uniformly at random.
- **ϵ-Nash Dynamics**: Players change their strategy only if they can improve their own cost by a factor of at least $1 + \epsilon$.

Nash Dynamics

- **Nash Dynamics**: Sequence of best responses of players.

$c_1(S) = 3$
$c_2(S) = 1$
$c_3(S) = 4$

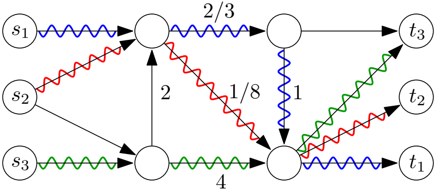

- Nash Dynamics with **Liveness Property**: Each player gets a chance to play his/her best response after at most t steps.
- **Random Nash Dynamics**: Players are chosen uniformly at random.
- **ϵ-Nash Dynamics**: Players change their strategy only if they can improve their own cost by a factor of at least $1 + \epsilon$.

Other dynamics (Noisy Nash Dynamics, Fictitious Play, Regret-minimization Dynamics) are discussed later.

Outline

The State Graph

state graph $\mathcal{G} = (V, E)$

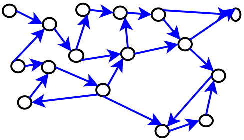

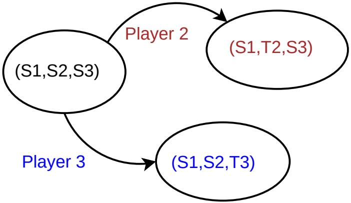

$V =$ **states** $E =$ **better/best responses**

The State Graph

state graph $\mathcal{G} = (V, E)$

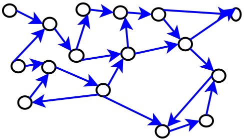

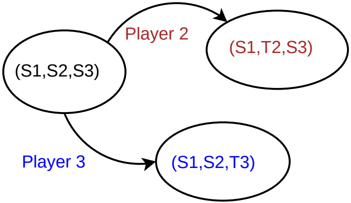

V = **states** $\qquad$ E = **better/best responses**

Properties of dynamics can be phrased in terms of state graph:

- pure Nash equilibrium = sink of state graph

The State Graph

state graph $\mathcal{G} = (V, E)$

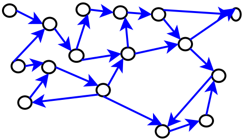

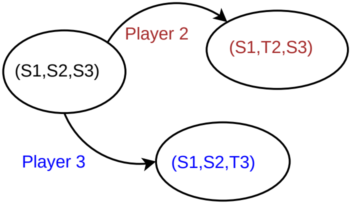

$V =$ **states** $\qquad E =$ **better/best responses**

Properties of dynamics can be phrased in terms of state graph:

- pure Nash equilibrium $=$ sink of state graph
- potential game $=$ acyclic state graph
  $\Rightarrow$ players eventually reach equilibrium.
  Example: Congestion Games

The State Graph

state graph $\mathcal{G} = (V, E)$

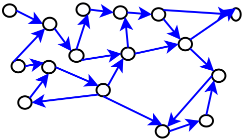

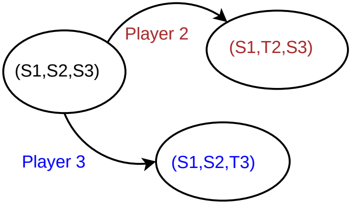

$V =$ **states** $\quad E =$ **better/best responses**

Properties of dynamics can be phrased in terms of state graph:

- pure Nash equilibrium $=$ sink of state graph
- potential game $=$ acyclic state graph
  $\Rightarrow$ players eventually reach equilibrium.
  Example: Congestion Games
- non-potential games $=$ best responses may cycle.

Outline

Congestion Games

Congestion Game:

- set of players $\mathcal{N}$
- set of resources $\mathcal{R}$
  e.g., edges of a graph or set of servers

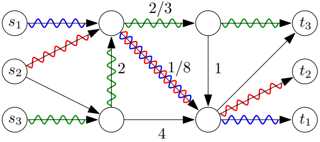

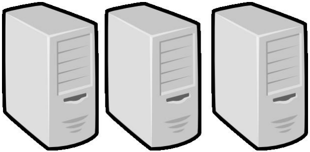

- set of strategies, $\forall i \in \mathcal{N} : \Sigma_i \subseteq 2^{\mathcal{R}}$

Congestion Games

Congestion Game:

- set of players $\mathcal{N}$
- set of resources $\mathcal{R}$

  e.g., edges of a graph or set of servers

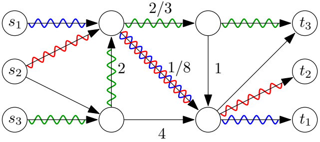

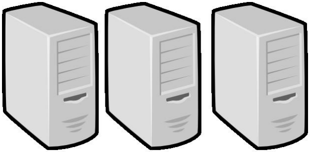

- set of strategies, $\forall i \in \mathcal{N} : \Sigma_i \subseteq 2^{\mathcal{R}}$

  $\Sigma_i = \{P \subseteq \mathcal{R} \mid P \text{ path } s_i \rightarrow t_i\}$ (network congestion game)

  $\Sigma_i = \{P \subseteq \mathcal{R} \mid P \text{ path } s \rightarrow t\}$ (symmetric netw. cong. game)

  $\Sigma_i = \{T \subseteq \mathcal{R} \mid T \text{ spanning tree}\}$

  $\Sigma_i = \{\{r\} \mid r \in \mathcal{R}\}$ (singleton congestion game)

Congestion Games

Congestion Game:

- set of players $\mathcal{N}$
- set of resources $\mathcal{R}$
  e.g., edges of a graph or set of servers

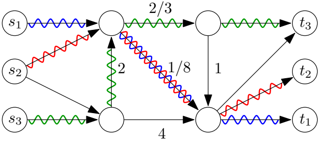

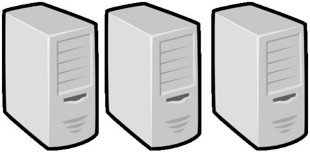

- set of strategies, $\forall i \in \mathcal{N} : \Sigma_i \subseteq 2^{\mathcal{R}}$
  $\Sigma_i = \{P \subseteq \mathcal{R} \mid P \text{ path } s_i \rightarrow t_i\}$ (network congestion game)
  $\Sigma_i = \{P \subseteq \mathcal{R} \mid P \text{ path } s \rightarrow t\}$ (symmetric netw. cong. game)
  $\Sigma_i = \{T \subseteq \mathcal{R} \mid T \text{ spanning tree}\}$
  $\Sigma_i = \{\{r\} \mid r \in \mathcal{R}\}$ (singleton congestion game)
- latency functions $\forall r \in \mathcal{R} \colon \ell_r \colon \mathbb{N} \rightarrow \mathbb{N}$

Rosenthal's Potential Function

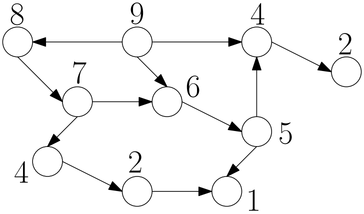

Rosenthal (Int. Journal of Game Theory 1973)

Every congestion game admits an exact potential function.

- $\Phi \colon \Sigma_1 \times \cdots \times \Sigma_n \rightarrow \mathbb{N}$ with $0 \leq \Phi \leq n \cdot m \cdot \ell_{\max}$
- player decreases his latency by $x \in \mathbb{N} \Rightarrow \Phi$ decreases by x as well

Rosenthal's Potential Function

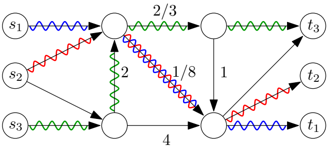

$\phi(S) = 2 + 2 + (1 + 8) = 13$

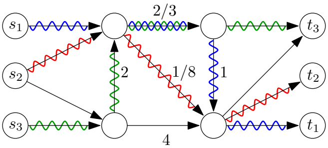

$\phi(S') = 2+(2+3)+1+1 = 9$

Rosenthal (Int. Journal of Game Theory 1973)

Every congestion game admits an exact potential function.

- $\Phi\colon \Sigma_1 \times \cdots \times \Sigma_n \to \mathbb{N}$ with $0 \le \Phi \le n \cdot m \cdot \ell_{\max}$
- player decreases his latency by $x \in \mathbb{N} \Rightarrow \Phi$ decreases by x as well
- $n_r =$ number of players i with $r \in S_i \in \Sigma_i$

$$\phi(S) = \sum_{r \in \mathcal{R}} \sum_{i=1}^{n_r} \ell_r(i)$$

Rosenthal's Potential Function

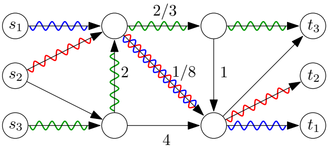

$\phi(S) = 2 + 2 + (1 + 8) = 13$

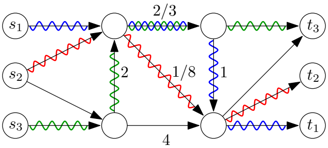

$\phi(S') = 2+(2+3)+1+1 = 9$

Rosenthal (Int. Journal of Game Theory 1973)

Every congestion game admits an exact potential function.

- $\Phi \colon \Sigma_1 \times \cdots \times \Sigma_n \to \mathbb{N}$ with $0 \leq \Phi \leq n \cdot m \cdot \ell_{\max}$
- player decreases his latency by $x \in \mathbb{N} \Rightarrow \Phi$ decreases by x as well
- n_r = number of players i with $r \in S_i \in \Sigma_i$

$$\phi(S) = \sum_{r \in \mathcal{R}} \sum_{i=1}^{n_r} \ell_r(i)$$

$\Rightarrow$ Number of better response at most $n \cdot m \cdot \ell_{\max}$.

Known Results on Convergence Time

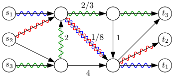

Fabrikant, Papadimitriou, Talwar (STOC 04)

There exist network congestion games with an initial state from which all better response sequences have exponential length.

Known Results on Convergence Time

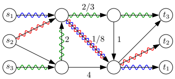

Fabrikant, Papadimitriou, Talwar (STOC 04)

There exist network congestion games with an initial state from which all better response sequences have exponential length.

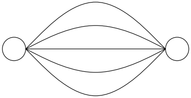

Ieong, McGrew, Nudelman, Shoham, Sun (AAAI 05)

In singleton games all best response sequences have length at most $n^2 \cdot m$.

Known Results on Convergence Time

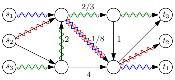

Fabrikant, Papadimitriou, Talwar (STOC 04)

There exist network congestion games with an initial state from which all better response sequences have exponential length.

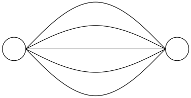

Ieong, McGrew, Nudelman, Shoham, Sun (AAAI 05)

In singleton games all best response sequences have length at most $n^2 \cdot m$.

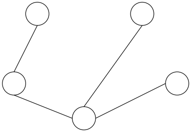

Ackermann, R., Vöcking (FOCS 06)

- In spanning tree congestion games all best response sequences have length at most $n^2 \cdot m \cdot$ number of vertices.
- In matroid congestion games all best response sequences have length at most $n^2 \cdot m \cdot rank$.

Singleton Games

Singleton Games

- Idea: Reduce latencies without affecting the game!

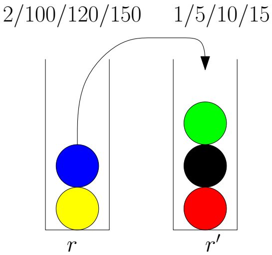

$\ell_r(n_r) > \ell_{r'}(n_{r'} + 1)$

Singleton Games

Singleton Games

- Idea: Reduce latencies without affecting the game!
- equivalent latencies $\overline{\ell}_r(x) \leq n \cdot m$

$$\forall r, r' \in \mathcal{R}, n_r, n_{r'} :$$

$$\begin{aligned} & \ell_r(n_r) > \ell_{r'}(n_{r'} + 1) \\ \iff \quad & \overline{\ell}_r(n_r) > \overline{\ell}_{r'}(n_{r'} + 1) \end{aligned}$$

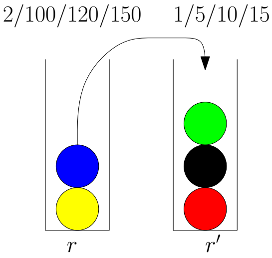

$\ell_r(n_r) > \ell_{r'}(n_{r'} + 1)$

Singleton Games

Singleton Games

- Idea: Reduce latencies without affecting the game!
- equivalent latencies $\overline{\ell}_r(x) \leq n \cdot m$

$$\forall r, r' \in \mathcal{R}, n_r, n_{r'} :$$

$$\begin{aligned} & \ell_r(n_r) > \ell_{r'}(n_{r'}+1) \\ \iff \quad & \overline{\ell}_r(n_r) > \overline{\ell}_{r'}(n_{r'}+1) \end{aligned}$$

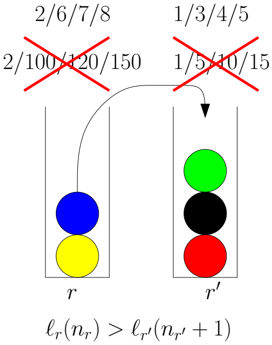

Singleton Games

Singleton Games

- Idea: Reduce latencies without affecting the game!
- equivalent latencies $\overline{\ell}_r(x) \leq n \cdot m$

$$\forall r, r' \in \mathcal{R}, n_r, n_{r'} :$$
$$\ell_r(n_r) > \ell_{r'}(n_{r'} + 1)$$
$$\iff \quad \overline{\ell}_r(n_r) > \overline{\ell}_{r'}(n_{r'} + 1)$$

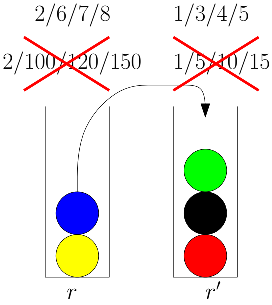

$$\ell_r(n_r) > \ell_{r'}(n_{r'} + 1)$$

Network Congestion Games

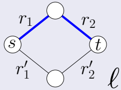

$$\ell_{r_1}(n_{r_1}) + \ell_{r_2}(n_{r_2}) > \ell_{r'_1}(n_{r'_1} + 1) + \ell_{r'_2}(n_{r'_2} + 1)$$

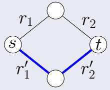

Singleton Games

Singleton Games

- Idea: Reduce latencies without affecting the game!
- equivalent latencies $\overline{\ell}_r(x) \leq n \cdot m$

$$\forall r, r' \in \mathcal{R}, n_r, n_{r'} :$$

$$\ell_r(n_r) > \ell_{r'}(n_{r'} + 1)$$

$$\iff \quad \overline{\ell}_r(n_r) > \overline{\ell}_{r'}(n_{r'} + 1)$$

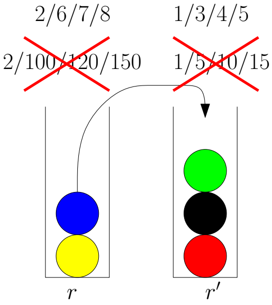

$$\ell_r(n_r) > \ell_{r'}(n_{r'} + 1)$$

Network Congestion Games

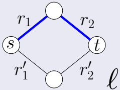

$$\ell_{r_1}(n_{r_1}) + \ell_{r_2}(n_{r_2}) > \ell_{r'_1}(n_{r'_1} + 1) + \ell_{r'_2}(n_{r'_2} + 1)$$

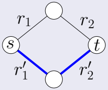

However, latency reduction works also for matroid games.

Matroid Congestion Games

Ackermann, R., Vöcking (FOCS 2006)

Let $(\mathcal{R}, \mathcal{I})$ be any non-matroid anti-chain. Then, for every n, there exists an n-player congestion game with the following properties.

- each Σ_i is **isomorphic to** $\mathcal{I}$,
- there is a **best response sequence of length** $2^{\Omega(n)}$.

Matroid Congestion Games

Ackermann, R., Vöcking (FOCS 2006)

Let $(\mathcal{R}, \mathcal{I})$ be any non-matroid anti-chain. Then, for every n, there exists an n-player congestion game with the following properties.

- each Σ_i is **isomorphic to $\mathcal{I}$**,
- there is a **best response sequence of length $2^{\Omega(n)}$**.

$\Rightarrow$ Matroid property is the maximal property on the individual players' strategy spaces that guarantees polynomial convergence.

PLS

Local Search Problem Π

- set of instances $\mathcal{I}_\Pi$
- for $I \in \mathcal{I}_\Pi$: set of feasible solutions $\mathcal{F}(I)$
- for $I \in \mathcal{I}_\Pi$: objective function $c\colon \mathcal{F}(I) \to \mathbb{Z}$
- for $I \in \mathcal{I}_\Pi$ and $S \in \mathcal{F}(I)$: neighborhood $\mathcal{N}(S, I) \subseteq \mathcal{F}(I)$

PLS

Local Search Problem Π

- set of instances $\mathcal{I}_\Pi$
- for $I \in \mathcal{I}_\Pi$: set of feasible solutions $\mathcal{F}(I)$
- for $I \in \mathcal{I}_\Pi$: objective function $c\colon \mathcal{F}(I) \to \mathbb{Z}$
- for $I \in \mathcal{I}_\Pi$ and $S \in \mathcal{F}(I)$: neighborhood $\mathcal{N}(S, I) \subseteq \mathcal{F}(I)$

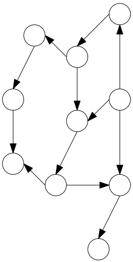

Johnson, Papadimitriou, Yannakakis (FOCS 85)

Π is in **PLS** if polynomial time algorithms exist for

- finding **initial feasible solution** $S \in \mathcal{F}(I)$,
- computing the **objective value** $c(S)$,
- finding a **better solution in the neighborhood** $\mathcal{N}(S, I)$ if S is not locally optimal.

PLS-reductions

PLS-reduction

- Polynomial-time computable function $f\colon \mathcal{I}_{\Pi_1} \to \mathcal{I}_{\Pi_2}$.
- Polynomial-time computable function ($S_2 \in \mathcal{F}(f(I))$) $g\colon S_2 \mapsto S_1 \in \mathcal{F}(I)$

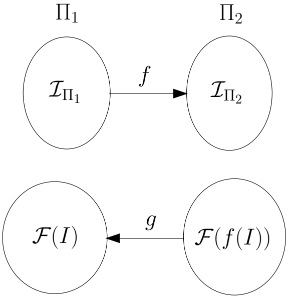

PLS-reductions

PLS-reduction

- Polynomial-time computable function $f\colon \mathcal{I}_{\Pi_1} \to \mathcal{I}_{\Pi_2}$.
- Polynomial-time computable function ($S_2 \in \mathcal{F}(f(I))$) $g\colon S_2 \mapsto S_1 \in \mathcal{F}(I)$
- **S_2 locally optimal $\Rightarrow$ $g(S_2)$ locally optimal.**

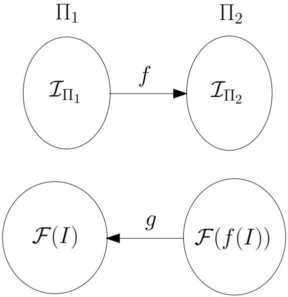

PLS-reductions

PLS-reduction

- Polynomial-time computable function $f\colon \mathcal{I}_{\Pi_1} \to \mathcal{I}_{\Pi_2}$.
- Polynomial-time computable function ($S_2 \in \mathcal{F}(f(I))$) $g\colon S_2 \mapsto S_1 \in \mathcal{F}(I)$
- S_2 **locally optimal** $\Rightarrow$ $g(S_2)$ **locally optimal.**

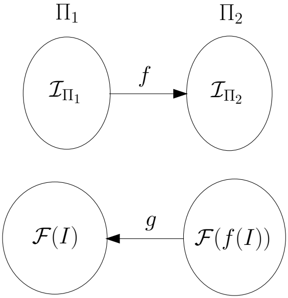

- local opt. of Π_2 easy to find $\Rightarrow$ local opt. of Π_1 easy to find
- local opt. of Π_1 hard to find $\Rightarrow$ local opt. of Π_2 hard to find

PLS-reductions

PLS-reduction

- Polynomial-time computable function $f\colon \mathcal{I}_{\Pi_1} \to \mathcal{I}_{\Pi_2}$.
- Polynomial-time computable function $(S_2 \in \mathcal{F}(f(I)))$
  $g\colon S_2 \mapsto S_1 \in \mathcal{F}(I)$
- **S_2 locally optimal $\Rightarrow$ $g(S_2)$ locally optimal.**

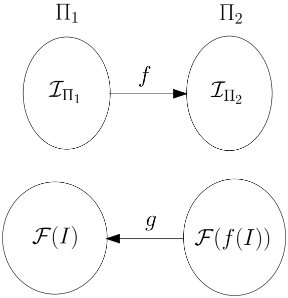

- local opt. of Π_2 easy to find $\Rightarrow$ local opt. of Π_1 easy to find
- local opt. of Π_1 hard to find $\Rightarrow$ local opt. of Π_2 hard to find

A PLS-reduction is called tight if it does not shorten distances in the state graph.
$\Rightarrow$ Exponential lower bounds are preserved.

Party Affiliation Games

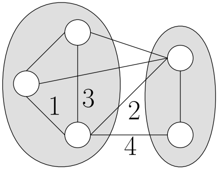

Party Affiliation Game:

- Input: $G(V, E)$ and $w\colon E \to \mathbb{N}$
- agents = nodes, $\Sigma_i = \{\text{left}, \text{right}\}$
- $w(\{u, v\}) =$ antipathy of u and v

Party Affiliation Games

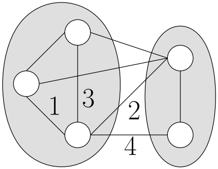

Party Affiliation Game:

- Input: $G(V, E)$ and $w \colon E \to \mathbb{N}$
- agents = nodes, $\Sigma_i = \{\text{left}, \text{right}\}$
- $w(\{u, v\}) =$ antipathy of u and v

Schäffer, Yannakakis (SIAM J. Comput. 1991)

Finding a locally optimal cut is PLS-complete.

Party Affiliation Games

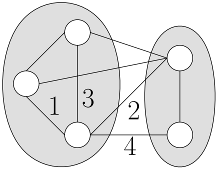

Party Affiliation Game:

- Input: $G(V, E)$ and $w\colon E \to \mathbb{N}$
- agents = nodes, $\Sigma_i = \{\text{left}, \text{right}\}$
- $w(\{u, v\})$ = antipathy of u and v

Schäffer, Yannakakis (SIAM J. Comput. 1991)

Finding a locally optimal cut is PLS-complete.

Very involved reduction from Circuit/Flip.

First PLS-complete problem: Circuit/Flip

- C: Boolean circuit composed of AND, OR, and NOT gates.
- Input to C: $x_1, \ldots, x_m \in \{0, 1\}$. Output of C: $y_1, \ldots, y_n \in \{0, 1\}$.

Party Affiliation Games

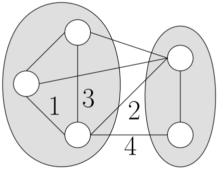

Party Affiliation Game:

- Input: $G(V,E)$ and $w\colon E \to \mathbb{N}$
- agents = nodes, $\Sigma_i = \{\text{left}, \text{right}\}$
- $w(\{u,v\})$ = antipathy of u and v

Schäffer, Yannakakis (SIAM J. Comput. 1991)

Finding a locally optimal cut is PLS-complete.

Very involved reduction from Circuit/Flip.

First PLS-complete problem: Circuit/Flip

- C: Boolean circuit composed of AND, OR, and NOT gates.
- Input to C: $x_1, \ldots, x_m \in \{0,1\}$. Output of C: $y_1, \ldots, y_n \in \{0,1\}$.
- Objective function: $f(x_1, \ldots, x_m) = \sum_{i=1}^{n} 2^{i-1} y_i$.

Party Affiliation Games

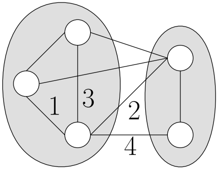

Party Affiliation Game:

- Input: $G(V, E)$ and $w\colon E \to \mathbb{N}$
- agents = nodes, $\Sigma_i = \{\text{left}, \text{right}\}$
- $w(\{u, v\})$ = antipathy of u and v

Schäffer, Yannakakis (SIAM J. Comput. 1991)

Finding a locally optimal cut is PLS-complete.

Very involved reduction from Circuit/Flip.

First PLS-complete problem: Circuit/Flip

- C: Boolean circuit composed of AND, OR, and NOT gates.
- Input to C: $x_1, \ldots, x_m \in \{0, 1\}$. Output of C: $y_1, \ldots, y_n \in \{0, 1\}$.
- Objective function: $f(x_1, \ldots, x_m) = \sum_{i=1}^{n} 2^{i-1} y_i$.
- Neighborhood $=$ Hamming distance 1

Congestion Games and PLS

Finding an equilibrium in a congestion game **belongs to PLS**:

- objective function $=$ Rosenthal's potential function
- $S' \in \mathcal{N}(S)$ if S' is obtained from S by better response of one of the players.

Congestion Games and PLS

Finding an equilibrium in a congestion game **belongs to PLS**:

- objective function = Rosenthal's potential function
- $S' \in \mathcal{N}(S)$ if S' is obtained from S by better response of one of the players.

PLS-completeness follows by reduction from MaxCut:

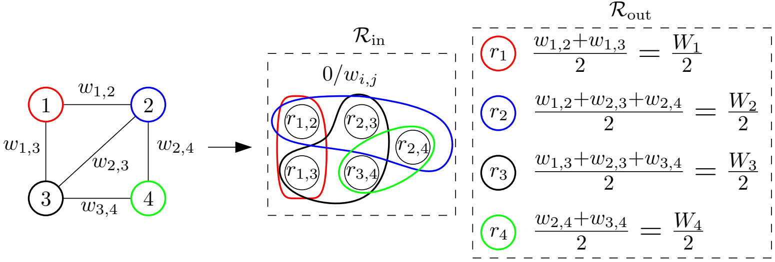

Congestion Games and PLS

Finding an equilibrium in a congestion game **belongs to PLS**:

- objective function $=$ Rosenthal's potential function
- $S' \in \mathcal{N}(S)$ if S' is obtained from S by better response of one of the players.

PLS-completeness follows by reduction from MaxCut:

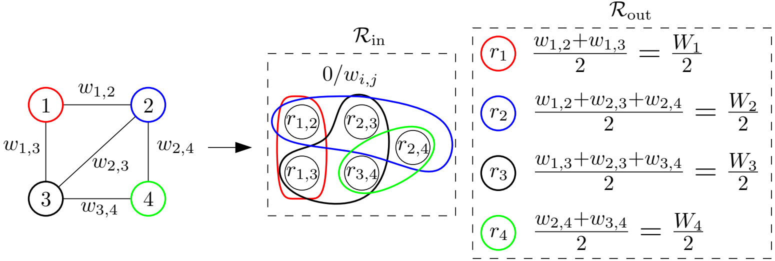

- g: player i on $\mathcal{R}_{\textbf{in}}$ $\Longleftrightarrow$ node i on **left side**

Congestion Games and PLS

Finding an equilibrium in a congestion game **belongs to PLS**:

- objective function $=$ Rosenthal's potential function
- $S' \in \mathcal{N}(S)$ if S' is obtained from S by better response of one of the players.

PLS-completeness follows by reduction from MaxCut:

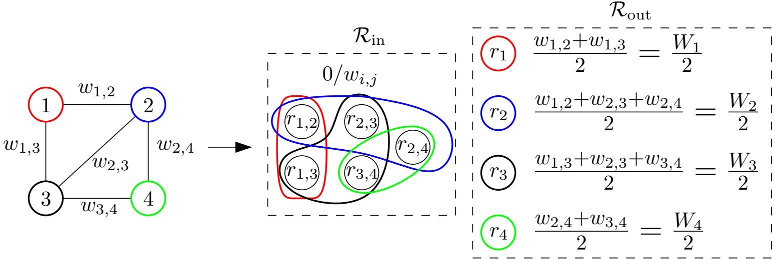

- g: player i on $\mathcal{R}_{\text{in}}$ $\iff$ node i on **left side**
- latency of player i on $\mathcal{R}_{\text{in}}$ $=$ weight of edges from i to the left side

Congestion Games and PLS

Finding an equilibrium in a congestion game **belongs to PLS**:

- objective function $=$ Rosenthal's potential function
- $S' \in \mathcal{N}(S)$ if S' is obtained from S by better response of one of the players.

PLS-completeness follows by reduction from MaxCut:

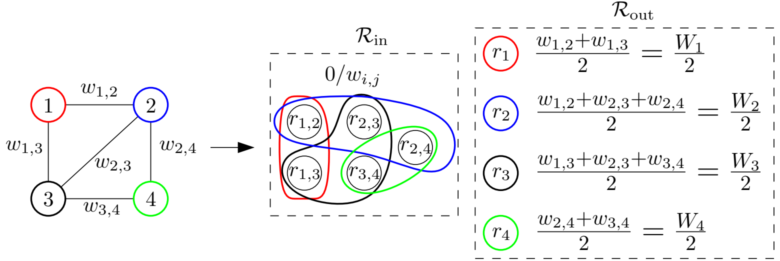

- g: player i on $\mathcal{R}_{\textbf{in}}$ $\iff$ node i on **left side**
- latency of player i on $\mathcal{R}_{\text{in}}$ $=$ weight of edges from i to the left side

node i on **left side**

$\iff$ latency on $\mathcal{R}_{\textbf{in}} \leq W_i/2$

$\iff$ contribution to cut when on **left side** $\geq W_i/2$

Network Congestion Games and PLS

Fabrikant, Papadimitriou, Talwar (STOC 04)

Finding a pure Nash equilibrium in network congestion games is **PLS-complete**.

Reduction from Circuit/Flip that reworks MaxCut reduction.

Network Congestion Games and PLS

Fabrikant, Papadimitriou, Talwar (STOC 04)

Finding a pure Nash equilibrium in network congestion games is **PLS-complete**.

Reduction from Circuit/Flip that reworks MaxCut reduction.

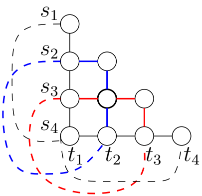

Ackermann, R., Vöcking (FOCS 06)

Network congestion games are PLS-complete for (un)directed networks with **linear latency functions**.

Simple reduction from MaxCut.

Network Congestion Games and PLS

Fabrikant, Papadimitriou, Talwar (STOC 04)

Finding a pure Nash equilibrium in network congestion games is **PLS-complete**.

Reduction from Circuit/Flip that reworks MaxCut reduction.

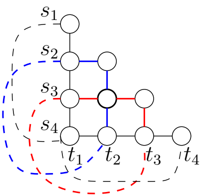

Ackermann, R., Vöcking (FOCS 06)

Network congestion games are PLS-complete for (un)directed networks with **linear latency functions**.

Simple reduction from MaxCut.

All these PLS-reductions are **tight**.

$\Rightarrow$ There exist states exponentially far from all sinks in the state graph.

Approximate Equilibria

What happens if players are **lazy**?

Approximate Equilibria

A state $S = (S_1, \ldots, S_n)$ is called $(1+\varepsilon)$-approximate equilibrium if $\forall i \in \mathcal{N}$: latency of player $i \leq$
$(1+\varepsilon) \cdot$ min achievable latency of player i

Approximate Equilibria

What happens if players are **lazy**?

Approximate Equilibria

A state $S = (S_1, \ldots, S_n)$ is called $(1+\varepsilon)$-approximate equilibrium if $\forall i \in \mathcal{N}$: latency of player $i \leq$ $(1+\varepsilon) \cdot$ min achievable latency of player i

Positive Result:

Chien, Sinclair (SODA 07)

In any symmetric congestion game with α-bounded jump condition, the $(1+\varepsilon)$-Nash dynamics converges after at most $\mathrm{poly}(n, \alpha, \varepsilon^{-1}, \log(\ell_{\max}))$ steps, assuming liveness property.

Approximate Equilibria

What happens if players are **lazy**?

Approximate Equilibria

A state $S = (S_1, \ldots, S_n)$ is called $(1+\varepsilon)$-approximate equilibrium if $\forall i \in \mathcal{N}$: latency of player $i \leq$ $(1+\varepsilon) \cdot$ min achievable latency of player i

Positive Result:

Chien, Sinclair (SODA 07)

In any symmetric congestion game with α-bounded jump condition, the $(1+\varepsilon)$-Nash dynamics converges after at most $\mathrm{poly}(n, \alpha, \varepsilon^{-1}, \log(\ell_{\max}))$ steps, assuming liveness property.

Idea: high-cost player moves $\Rightarrow$ significant potential drop
S not $(1+\varepsilon)$-equilibrium $\Rightarrow \exists$ high-cost player that has an incentive to move. (due to α-bounded jump condition and symmetry)

Approximate Equilibria

What happens if players are **lazy**?

Approximate Equilibria

A state $S = (S_1, \ldots, S_n)$ is called $(1+\varepsilon)$-approximate equilibrium if $\forall i \in \mathcal{N}$: latency of player $i \leq$ $(1+\varepsilon) \cdot$ min achievable latency of player i

Negative Result:

Skopalik, Vöcking (STOC 2008)

It is PLS-hard to compute an $(1+\varepsilon)$-approximate equilibrium for any polynomial-time computable ε.

$\Rightarrow$ Exponentially many steps until $(1+\varepsilon)$-approx. eq. is reached.
Very involved reduction from Circuit/Flip.

Summary of Convergence Results

	Nash Dynamics	ε-Nash Dynamics
Matroid	poly	poly
Symmetric Network	exp	poly
Asymmetric Network	exp, PLS-complete	exp
Symmetric General	exp, PLS-complete	poly
Asymmetric General	exp, PLS-complete	exp, PLS-complete
Cut Games	exp, PLS-complete	?

Outline

Non-potential Games

State Graph

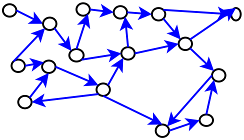

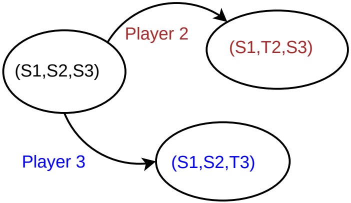

Sink equilibrium: [Goemans, M., Vetta (FOCS 2005)]
strongly connected comp. of state graph without outgoing edges

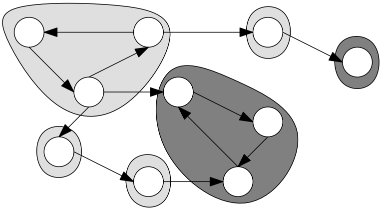

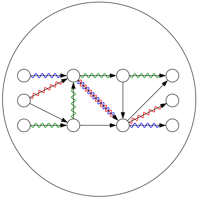

$\Rightarrow$ random Nash dynamics eventually reaches sink equilibrium

An Example

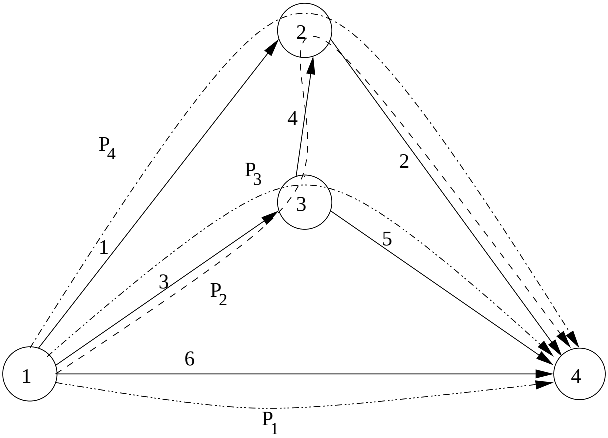

- Two agents: ($r_1 = 1$, $r_2 = 2$).
- $l_1(x) = x + 33$, $l_2(x) = 13x$, $l_3(x) = 3x^2$, $l_4(x) = 6x^2$, $l_5(x) = x^2 + 44$, and $l_6(x) = 47x$.

An Example

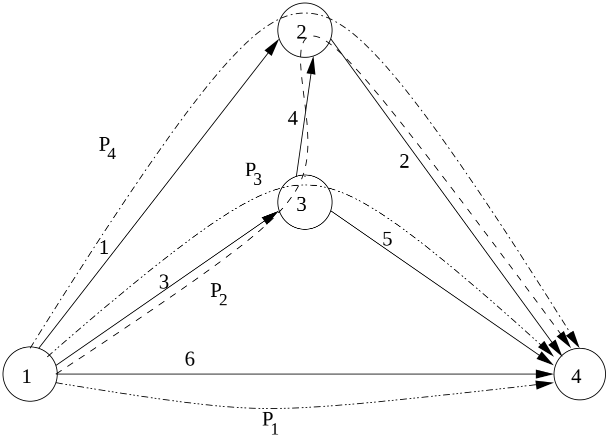

- Two agents: $(r_1 = 1, r_2 = 2)$.
- Only Sink equilibrium: $\{(P_1, P_2), (P_3, P_2), (P_3, P_4), (P_1, P_4)\}$.
- **No Pure Nash equilibrium.**

Non-potential Games

Sink equilibrium: [Goemans, M., Vetta (FOCS 2005)]

strongly connected comp. of state graph without outgoing edges

Non-potential Games

Sink equilibrium: [Goemans, M., Vetta (FOCS 2005)]
strongly connected comp. of state graph without outgoing edges
Complexity Questions about Nash Dynamics and sink equilibria:

1. Given a state in a game, is it in a sink equilibrium?
2. Given a game, determine if it has a pure Nash equilibrium?
3. Given a game, determine if it has any non-singleton sink equilibrium?

Non-potential Games

Sink equilibrium: [Goemans, M., Vetta (FOCS 2005)]
strongly connected comp. of state graph without outgoing edges
Complexity Questions about Nash Dynamics and sink equilibria:

1. Given a state in a game, is it in a sink equilibrium?
2. Given a game, determine if it has a pure Nash equilibrium?
3. Given a game, determine if it has any non-singleton sink equilibrium?

Theorem (M., Skopalik, EC 2009)

For many classes of games with succinet representation, it is PSPACE-hard to answer questions 1 and 3, and it is NP-hard to answer question 2.

- *Player-specific and weighted congestion games.*
- *Anonymous Games and Graphical Games.*
- *Many-to-one two-sided market games.*

Games with Singleton Sink Equilibria

Interesting subclass: Games with only singleton sink equilibria

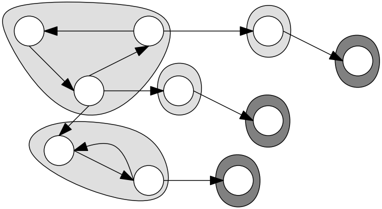

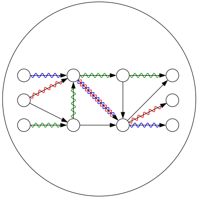

Milchtaich (Games and Economics Behaviour, 1996)

Player-specific singleton congestion games:

- Pure Nash equilibria exist, but best resp. dyn. can cycle.
- From every state there is a sequence of best responses to an equilibrium.

How to find a stable marriage?

Let's get to the really important problems. . .

The Stable Marriage Problem

Set of women $\mathcal{X}$

Set of men $\mathcal{Y}$

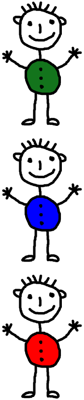

The Stable Marriage Problem

Set of women $\mathcal{X}$

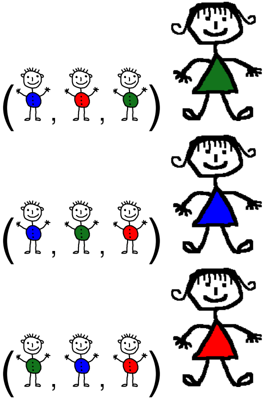

Set of men $\mathcal{Y}$

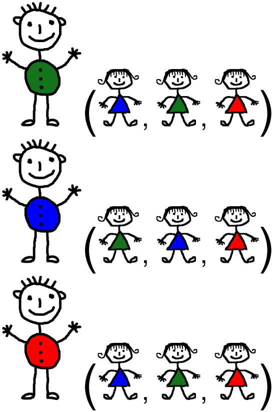

Every person has a preference list.

The Stable Marriage Problem

Set of women $\mathcal{X}$

Set of men $\mathcal{Y}$

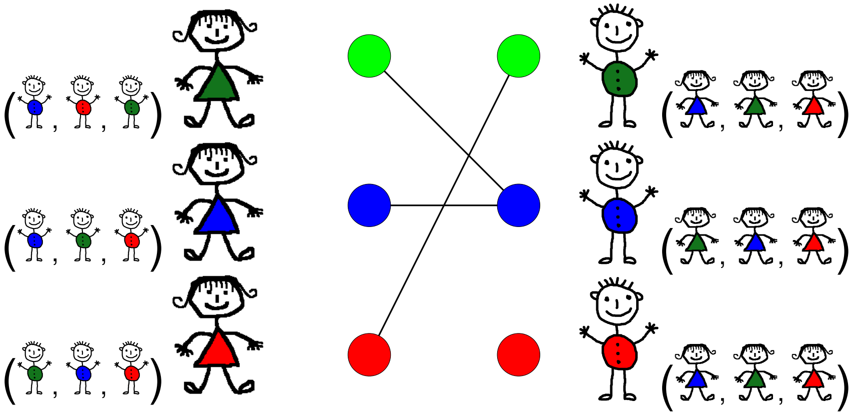

Every person has a preference list.

The Stable Marriage Problem

Set of women $\mathcal{X}$

Set of men $\mathcal{Y}$

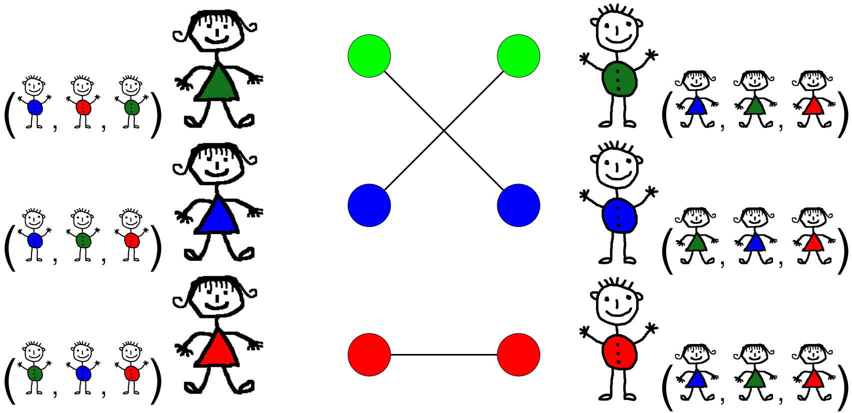

Every person has a preference list.

Formal Definition

Stable Matching

A matching is stable if there does not exist a blocking pair.

Formal Definition

Stable Matching

A matching is stable if there does not exist a blocking pair.

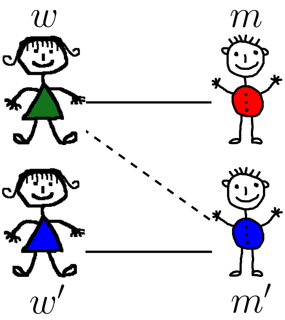

(w, m') is blocking pair

$$\Longleftrightarrow$$

1) w prefers m' to m

2) m' prefers w to w'

Formal Definition

Stable Matching

A matching is stable if there does not exist a blocking pair.

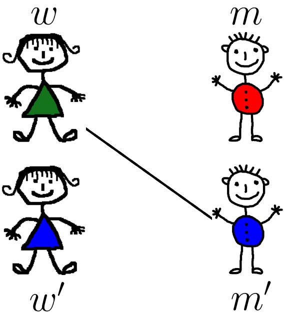

(w, m') is blocking pair

$\Longleftrightarrow$

1) w prefers m' to m

2) m' prefers w to w'

Formal Definition

Stable Matching

A matching is stable if there does not exist a blocking pair.

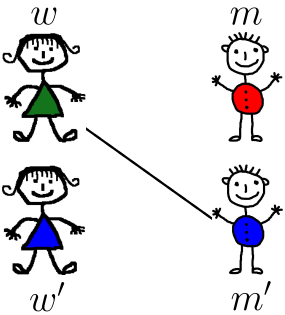

(w, m') is blocking pair

$\iff$

1) w prefers m' to m
2) m' prefers w to w'

Theorem [Gale, Shapley 1962]

A stable matching can be computed efficiently.

Applications and Previous Work

- Many Applications: Interns/Hospitals, College Admission, Labor market.

Applications and Previous Work

- Many Applications: Interns/Hospitals, College Admission, Labor market.

Main Question

What happens without central authority?

- Do players reach a stable matching?
- How long does it take?

Applications and Previous Work

- Many Applications: Interns/Hospitals, College Admission, Labor market.

Main Question

What happens without central authority?

- Do players reach a stable matching?
- How long does it take?

Consecutive resolving of blocking pairs:

- Knuth observed a cycle.
- Roth and Vande Vate showed that there is no non-trivial sink equilibrium (Econometrica 1990).

Best Response Dynamics

Matching not stable $\Rightarrow$ Choose woman, let her play best response.

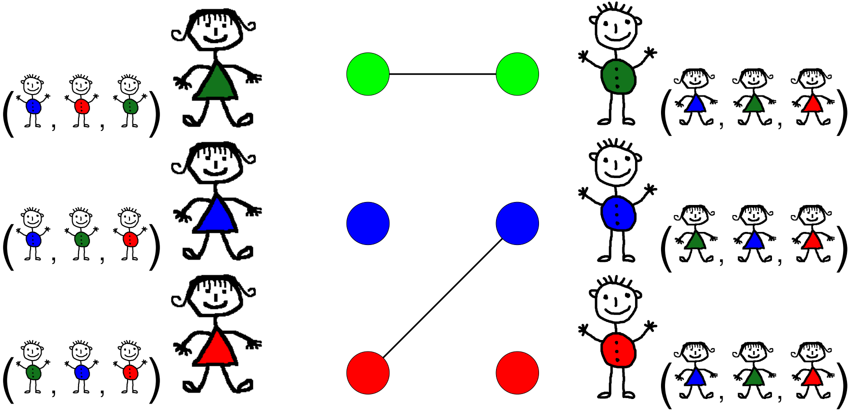

Best Response Dynamics

Matching not stable $\Rightarrow$ Choose woman, let her play best response.

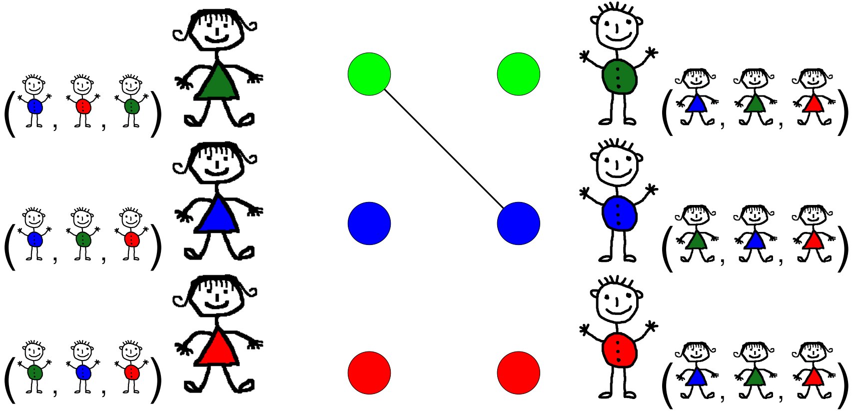

Best Response Dynamics

Matching not stable $\Rightarrow$ Choose woman, let her play best response.

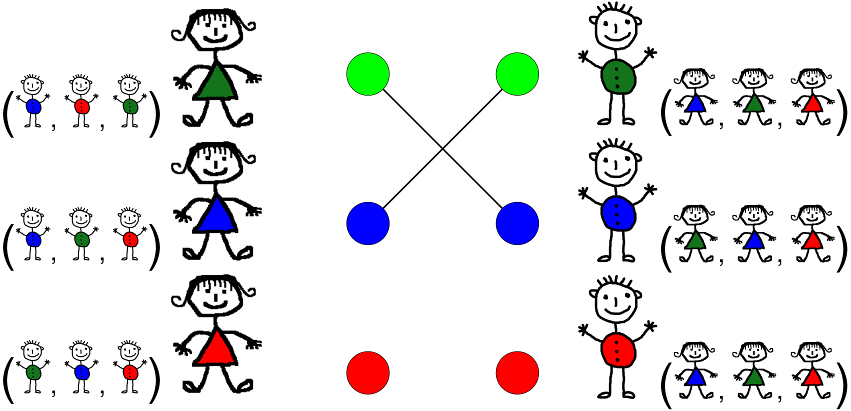

Best Response Dynamics

Matching not stable $\Rightarrow$ Choose woman, let her play best response.

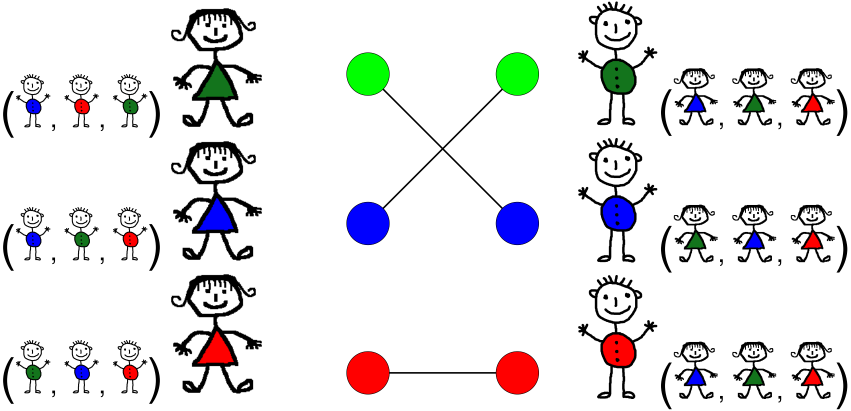

Best Response Dynamics

Ackermann, Goldberg, M., R., Vöcking (EC 2008)

The best response dynamics can cycle.

Was shown for better response dynamics by Knuth.

Best Response Dynamics

Ackermann, Goldberg, M., R., Vöcking (EC 2008)

The best response dynamics can cycle.

Was shown for better response dynamics by Knuth.

Ackermann, Goldberg, M., R., Vöcking (EC 2008)

From every matching there exists a sequence of $2n^2$ best responses to a stable matching.

Was shown for better response dynamics by Roth and Vande Vate.
$\Rightarrow$ Random best response dynamics reaches a stable matching with probability 1.

Best Response Dynamics

Ackermann, Goldberg, M., R., Vöcking (EC 2008)

The best response dynamics can cycle.

Was shown for better response dynamics by Knuth.

Ackermann, Goldberg, M., R., Vöcking (EC 2008)

From every matching there exists a sequence of $2n^2$ best responses to a stable matching.

Was shown for better response dynamics by Roth and Vande Vate.
$\Rightarrow$ Random best response dynamics reaches a stable matching with probability 1.

Ackermann, Goldberg, M., R., Vöcking (EC 2008)

There exist instances such that the expected number of best responses is $\Omega(c^n)$ for some constant $c > 1$.

Similar exponential bound holds for better response dynamics.

Best Response Dynamics – Upper Bound

Theorem

From every matching there exists a sequence of $2n^2$ best responses to a stable matching.

Claim 1

If only married women play best responses, after at most n^2 steps every married woman is happy.

Claim 2

If every married woman is happy, every sequence of best responses terminates after at most n^2 steps.

Best Response Dynamics

Claim 1

If only married women play best responses, after at most n^2 steps every married woman is happy.

Proof.

Use the following potential function:

$$\Phi = \sum_{\text{married woman } w} \text{rank of } w\text{'s current partner}$$

$0 \le \Phi \le n^2$ and Φ decreases with every best response. $\square$

Best Response Dynamics

Claim 1

If only married women play best responses, after at most n^2 steps every married woman is happy.

Proof.

Use the following potential function:

$$\Phi = \sum_{\text{married woman } w} \text{rank of } w\text{'s current partner}$$

$0 \leq \Phi \leq n^2$ and Φ decreases with every best response. $\square$

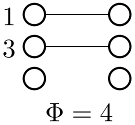

Best Response Dynamics

Claim 1

If only married women play best responses, after at most n^2 steps every married woman is happy.

Proof.

Use the following potential function:

$$\Phi = \sum_{\text{married woman } w} \text{rank of } w\text{'s current partner}$$

$0 \leq \Phi \leq n^2$ and Φ decreases with every best response. $\square$

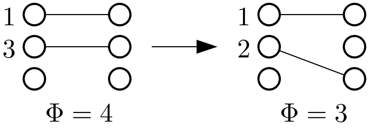

Best Response Dynamics

Claim 1

If only married women play best responses, after at most n^2 steps every married woman is happy.

Proof.

Use the following potential function:

$$\Phi = \sum_{\text{married woman } w} \text{rank of } w\text{'s current partner}$$

$0 \leq \Phi \leq n^2$ and Φ decreases with every best response. $\square$

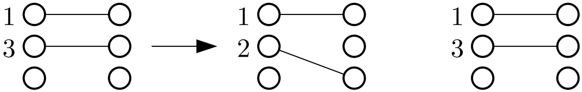

$\Phi = 4$ $\qquad$ $\Phi = 3$ $\qquad$ $\Phi = 4$

Best Response Dynamics

Claim 1

If only married women play best responses, after at most n^2 steps every married woman is happy.

Proof.

Use the following potential function:

$$\Phi = \sum_{\text{married woman } w} \text{rank of } w\text{'s current partner}$$

$0 \leq \Phi \leq n^2$ and Φ decreases with every best response. $\square$

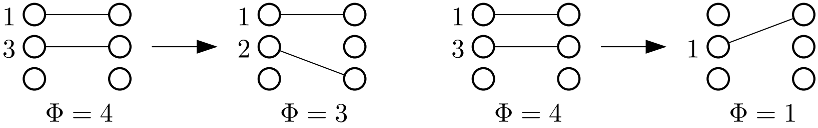

Best Response Dynamics – Upper Bound

Claim 2

If every married woman is happy, every sequence of best responses terminates after at most n^2 steps.

Proof.

Invariant: No married woman can improve.

Best Response Dynamics – Upper Bound

Claim 2

If every married woman is happy, every sequence of best responses terminates after at most n^2 steps.

Proof.

Invariant: No married woman can improve.

$\Rightarrow$ Men are never dumped.

Best Response Dynamics – Upper Bound

Claim 2

If every married woman is happy, every sequence of best responses terminates after at most n^2 steps.

Proof.

Invariant: No married woman can improve.

$\Rightarrow$ Men are never dumped.

Use the following potential function:

$$\Psi = \sum_{\text{married man } m} n + 1 - \text{rank of } m\text{'s current partner}$$

$0 \leq \Psi \leq n^2$ and Ψ increases with every best response. $\square$

Best Response Dynamics – Upper Bound

Claim 2

If every married woman is happy, every sequence of best responses terminates after at most n^2 steps.

Proof.

Invariant: No married woman can improve.

$\Rightarrow$ Men are never dumped.

Use the following potential function:

$$\Psi = \sum_{\text{married man } m} n + 1 - \text{rank of } m\text{'s current partner}$$

$0 \leq \Psi \leq n^2$ and Ψ increases with every best response. $\square$

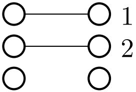

$\Psi = 5$

Best Response Dynamics – Upper Bound

Claim 2

If every married woman is happy, every sequence of best responses terminates after at most n^2 steps.

Proof.

Invariant: No married woman can improve.

$\Rightarrow$ Men are never dumped.

Use the following potential function:

$$\Psi = \sum_{\text{married man } m} n + 1 - \text{rank of } m\text{'s current partner}$$

$0 \leq \Psi \leq n^2$ and Ψ increases with every best response. $\square$

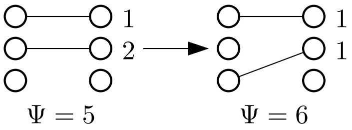

Best Response Dynamics – Upper Bound

Claim 2

If every married woman is happy, every sequence of best responses terminates after at most n^2 steps.

Proof.

Invariant: No married woman can improve.

$\Rightarrow$ Men are never dumped.

Use the following potential function:

$$\Psi = \sum_{\text{married man } m} n + 1 - \text{rank of } m\text{'s current partner}$$

$0 \leq \Psi \leq n^2$ and Ψ increases with every best response. $\square$

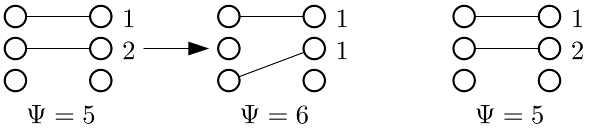

Best Response Dynamics – Upper Bound

Claim 2

If every married woman is happy, every sequence of best responses terminates after at most n^2 steps.

Proof.

Invariant: No married woman can improve.

$\Rightarrow$ Men are never dumped.

Use the following potential function:

$$\Psi = \sum_{\text{married man } m} n + 1 - \text{rank of } m\text{'s current partner}$$

$0 \leq \Psi \leq n^2$ and Ψ increases with every best response. $\square$

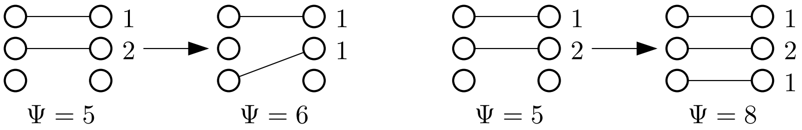

Further Results – Correlated Instances

Good news: Correlation helps!

Monotone Instances

Input: complete, weighted bipartite graph $G = (V, E, w)$.
Every player tries to maximize the weight of her/his relationship.

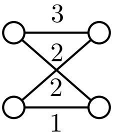

Further Results – Correlated Instances

Good news: Correlation helps!

Monotone Instances

Input: complete, weighted bipartite graph $G = (V, E, w)$.
Every player tries to maximize the weight of her/his relationship.

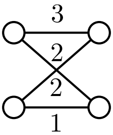

Theorem

Random best/better responses converge in polynomial time whp.

Outline

Price of Anarchy and Convergence

- Price of anarchy $= \dfrac{\text{Social Value of the worst equilibrium}}{\text{Optimal Social Value}}$.

Price of Anarchy and Convergence

- Price of anarchy $= \dfrac{\text{Social Value of the worst equilibrium}}{\text{Optimal Social Value}}$.
- Large Price of Anarchy: Need for Central Regulation.
- Small Price of Anarchy: Does not indicate good performance.

Price of Anarchy and Convergence

- Price of anarchy $= \dfrac{\text{Social Value of the worst equilibrium}}{\text{Optimal Social Value}}$.

- Large Price of Anarchy: Need for Central Regulation.
- Small Price of Anarchy: Does not indicate good performance.

- Players may not converge to those equilibria.
- Convergence to equilibria may take exponential time.

Price of Anarchy and Convergence

- Price of anarchy $= \frac{\text{Social Value of the worst equilibrium}}{\text{Optimal Social Value}}$.

- Large Price of Anarchy: Need for Central Regulation.
- Small Price of Anarchy: Does not indicate good performance.

- Players may not converge to those equilibria.
- Convergence to equilibria may take exponential time.

Question 1: Potential Games: How fast do players converge to approximate solutions? (and not to equilibria).

Question 2 : Non-Potential Games: What is the quality of solutions that players converge to?

Outline

Congestion Games: Conv. to Nearly-Optimal Sol.

- Question 1 (Potential Games): How fast do players converge to approximate solutions? (and not to equilibria).
- Price of anarchy: 2.5 (Koutsoupias, Christoudolou,05 and Awerbuch, Azar, Epstein, 05).

Congestion Games: Conv. to Nearly-Optimal Sol.

- Question 1 (Potential Games): How fast do players converge to approximate solutions? (and not to equilibria).
- Price of anarchy: 2.5 (Koutsoupias, Christoudolou,05 and Awerbuch, Azar, Epstein, 05).
- Congestion games are potential games, but convergence will take exponential time even for approximate Nash Dynamics

Congestion Games: Conv. to Nearly-Optimal Sol.

- Question 1 (Potential Games): How fast do players converge to approximate solutions? (and not to equilibria).
- Price of anarchy: 2.5 (Koutsoupias, Christoudolou,05 and Awerbuch, Azar, Epstein, 05).
- Congestion games are potential games, but convergence will take exponential time even for approximate Nash Dynamics

How about convergence time to constant-factor approximate solutions?

Convergence to Nearly-optimal Solutions

Theorem (Awerbuch, Azar, Epstein, M., Skopalik, EC 2008)

Convergence time of Nash dynamics with liveness property to constant-factor optimal solutions in linear congestion games might be exponential.

Convergence to Nearly-optimal Solutions

Theorem (Awerbuch, Azar, Epstein, M., Skopalik, EC 2008)

Convergence time of Nash dynamics with liveness property to constant-factor optimal solutions in linear congestion games might be exponential.

- This is in contrast to:

Theorem (Goemans, M., Vetta, FOCS 2005)

For Random Nash dynamics, convergence time to constant-factor solutions in linear congestion games is polynomial.

Convergence to Nearly-optimal Solutions

Theorem (Awerbuch, Azar, Epstein, M., Skopalik, EC 2008)

Convergence time of Nash dynamics with liveness property to constant-factor optimal solutions in linear congestion games might be exponential.

- This is in contrast to:

Theorem (Goemans, M., Vetta, FOCS 2005)

For Random Nash dynamics, convergence time to constant-factor solutions in linear congestion games is polynomial.

- Proof Idea: Three lemmas:
  - In any bad state, there exists a player who improves the average by a large margin, thus there is a state.
  - In any bad state, the expected value of the change incurred by players is not too bad.
  - Use induction on the above lemmas.

$\Rightarrow$ The price of anarchy for sink equilibrium is a constant.

Convergence to Nearly-optimal Solutions

Theorem (Awerbuch, Azar,Epstein, M., Skopalik, EC 2008)

For a large class of potential games that are β-nice, and satisfy bounded-jump condition, after polynomial steps of ϵ-Nash dynamics with a liveness property, players converge to a solution with approximation factor of price of anarchy.

Convergence to Nearly-optimal Solutions

Theorem (Awerbuch, Azar,Epstein, M., Skopalik, EC 2008)

For a large class of potential games that are β-nice, and satisfy bounded-jump condition, after polynomial steps of ϵ-Nash dynamics with a liveness property, players converge to a solution with approximation factor of price of anarchy.

- Bounded-jump condition (informal): After a player i plays a best response, the change in the payoff (cost) of other players is bounded by the new payoff (cost) of player i.

Convergence to Nearly-optimal Solutions

Theorem (Awerbuch, Azar,Epstein, M., Skopalik, EC 2008)

For a large class of potential games that are β-nice, and satisfy bounded-jump condition, after polynomial steps of ϵ-Nash dynamics with a liveness property, players converge to a solution with approximation factor of price of anarchy.

- Bounded-jump condition (informal): After a player i plays a best response, the change in the payoff (cost) of other players is bounded by the new payoff (cost) of player i.
- For example:
  - Congestion games with constant-degree polynomial delay functions,
  - Weighted congestion games with linear delay functions,
  - Party affiliation games,
  - Market sharing games.

Summary of Convergence to Nearly-Optimal Solutions

Convergence to Nash equilibria: exponential

Convergence to nearly-optimal solutions:

Game	PoA	Nash	Rand. Nash	ϵ-Nash
Linear Congestion	2.5	expon	poly, 70	poly, $2.5 + \epsilon$
Deg. d Cong.	2.5	expon	poly, $O(2^{2d})$	poly, $O(2^d) + \epsilon$
Wei. Lin. Cong.	2.62	expon	poly, 70	poly, $2.62 + \epsilon$
Cut Games	$\frac{1}{2}$	expon	poly, $\frac{1}{6}$	poly, $\frac{1}{2} - \epsilon$
Market Sharing	$\frac{1}{2}$	poly, $\frac{1}{\log n}$	poly, $\frac{1}{\log n}$	poly, $\frac{1}{2} - \epsilon$

Summary of Convergence to Nearly-Optimal Solutions

Convergence to Nash equilibria: exponential

Convergence to nearly-optimal solutions:

Game	PoA	Nash	Rand. Nash	ϵ-Nash
Linear Congestion	2.5	expon	poly, 70	poly, $2.5 + \epsilon$
Deg. d Cong.	2.5	expon	poly, $O(2^{2^d})$	poly, $O(2^d) + \epsilon$
Wei. Lin. Cong.	2.62	expon	poly, 70	poly, $2.62 + \epsilon$
Cut Games	$\frac{1}{2}$	expon	poly, $\frac{1}{6}$	poly, $\frac{1}{2} - \epsilon$
Market Sharing	$\frac{1}{2}$	poly, $\frac{1}{\log n}$	poly, $\frac{1}{\log n}$	poly, $\frac{1}{2} - \epsilon$

For other games, check the β-nice and bounded jump condition.

Outline

Sink Equilibria and Convergence

- Question 2 (Non-Potential Games): What is the quality of solutions that players converge to?

Sink Equilibria and Convergence

- Question 2 (Non-Potential Games): What is the quality of solutions that players converge to?
- Price of anarchy for mixed NE might be good, but how about convergence to good-quality solutions in non-potential games?

Sink Equilibria and Convergence

- Question 2 (Non-Potential Games): What is the quality of solutions that players converge to?
- Price of anarchy for mixed NE might be good, but how about convergence to good-quality solutions in non-potential games?
- In other words, what is the price of anarchy of sink equilibria?

Price of Anarchy for Sink equilibria

- A sink equilibrium is a set of states.
- Each state has a social value.

Price of Anarchy for Sink equilibria

- A sink equilibrium is a set of states.
- Each state has a social value.

Social Value of a Sink equilibrium?

Social Value of a Sink equilibrium = Average Social value of states on a random best-response walk.

- *Random Best-response* Walk: Choose a player uniformly at random at each step.

Price of Anarchy for Sink equilibria

- A sink equilibrium is a set of states.
- Each state has a social value.

Social Value of a Sink equilibrium?

Social Value of a Sink equilibrium = Average Social value of states on a random best-response walk.

- *Random Best-response* Walk: Choose a player uniformly at random at each step.

- Price of anarchy for sink equilibrium =

$$\frac{\text{value of the worst sink equilibrium}}{\mathrm{Opt}}.$$

Sink Equilibria and Convergence

Theorem (Goemans, M., Vetta, FOCS 2005)

For weighted congestion games with constant-degree polynomial delay functions, the price of anarchy for sink equilibria is constant.

Related to convergence of random Nash dynamics to constant-factor approximate solutions.

Sink Equilibria and Convergence

Theorem (Goemans, M., Vetta, FOCS 2005)

For weighted congestion games with constant-degree polynomial delay functions, the price of anarchy for sink equilibria is constant.

Related to convergence of random Nash dynamics to constant-factor approximate solutions.

Theorem (Goemans, M., Vetta, FOCS 2005)

For a general class of market sharing games (aka valid-utility games), eventhough the price of anarchy for mixed NE is constant ($1/2$), the price of anarchy for sink equilibria is very poor ($\frac{1}{n}$).

$\Rightarrow$ Players may converge to a bad-quality solution and they may get stuck there.

Sink Equilibria and Convergence

Theorem (Goemans, M., Vetta, FOCS 2005)

For weighted congestion games with constant-degree polynomial delay functions, the price of anarchy for sink equilibria is constant.

Related to convergence of random Nash dynamics to constant-factor approximate solutions.

Theorem (Goemans, M., Vetta, FOCS 2005)

For a general class of market sharing games (aka valid-utility games), eventhough the price of anarchy for mixed NE is constant ($1/2$), the price of anarchy for sink equilibria is very poor ($\frac{1}{n}$).

$\Rightarrow$ Players may converge to a bad-quality solution and they may get stuck there.

What if players follow other dynamics?

Outline

Natural Distributed/Synchronous Dynamics

- Fictitious Play
- Replicator dynamics
- Noisy Nash dynamics
- No-regret dynamics

Natural Distributed/Synchronous Dynamics

- **Fictitious Play**
  - **Best response to the empirical distribution of the opponents.**
  - **Nash equilibrium is an “absorbing state”**
- Replicator dynamics
- Noisy Nash dynamics
- No-regret dynamics

Natural Distributed/Synchronous Dynamics

- Fictitious Play
  - Best response to the empirical distribution of the opponents.
  - Nash equilibrium is an “absorbing state”
- **Replicator dynamics**
  - **Each strategy survives according to its excess payoff**
  - **Most reasonable variants converge in potential games [Sandholm JET 2001]**
  - **Convergence rate [Fischer, Räcke, Vöcking STOC06]**
- Noisy Nash dynamics
- No-regret dynamics

Natural Distributed/Synchronous Dynamics

- Fictitious Play
  - Best response to the empirical distribution of the opponents.
  - Nash equilibrium is an “absorbing state”
- Replicator dynamics
  - Each strategy survives according to its excess payoff
  - Most reasonable variants converge in potential games [Sandholm JET 2001]
  - Convergence rate [Fischer, Räcke, Vöcking STOC06]
- **Noisy Nash dynamics**
  - **At each step, there is a probability of not playing best response.**
  - **Convergence properties in Congestion Games [Asadpour, Saberi 2009].**
- No-regret dynamics

Natural Distributed/Synchronous Dynamics

- Fictitious Play
  - Best response to the empirical distribution of the opponents.
  - Nash equilibrium is an ”absorbing state”
- Replicator dynamics
  - Each strategy survives according to its excess payoff
  - Most reasonable variants converge in potential games [Sandholm JET 01]
  - Convergence rate [Fischer, Räcke, Vöcking STOC 06]
- Noisy Nash dynamics.
  - At each step, there is a probability of not playing best response.
  - Convergence properties in Congestion Games [Asadpour, Saberi 2009].
- **No-regret dynamics.**
  - **Known to converge in specific games to Nash equilibrium.**
  - **There exist games on which uncoupled dynamics do not converge [Hart and Mas-Collel].**

No regret in Congestion Games

No-External-Regret

- Is there a strategy that guarantees that the total routing time will take almost as time as the best fixed path in hindsight?

No regret in Congestion Games

No-External-Regret

- Is there a strategy that guarantees that the total routing time will take almost as time as the best fixed path in hindsight?

No-Internal-Regret

- Is there a strategy that guarantees that the total routing time when it took path P will take almost as time as the best fixed path in hindsight for that time steps?

No regret in Congestion Games

No-External-Regret

- Is there a strategy that guarantees that the total routing time will take almost as time as the best fixed path in hindsight?

No-Internal-Regret

- Is there a strategy that guarantees that the total routing time when it took path P will take almost as time as the best fixed path in hindsight for that time steps?

We say that algorithm is No X-Regret if its regret to best static decision, $R(T)$ is sublinear.

Outline

Correlated Equilibria [Aumann 1974]

- **Distribution over N-tuples.**

Correlated Equilibria [Aumann 1974]

- **Distribution over N-tuples.**
- **Nash Equilibrium with a shared signal**

Correlated Equilibria [Aumann 1974]

- **Distribution over N-tuples.**
- **Nash Equilibrium with a shared signal**
  - **Independent signal - Nash equilibrium**

Correlated Equilibria [Aumann 1974]

- Distribution over N-tuples.
- Nash Equilibrium with a shared signal
  - Independent signal - Nash equilibrium
  - Public signal - Convex combinations of Nash equilibrium

Correlated Equilibria [Aumann 1974]

- Distribution over N-tuples.
- Nash Equilibrium with a shared signal
  - Independent signal - Nash equilibrium
  - Public signal - Convex combinations of Nash equilibrium
  - Private signal - not necessarily convex hull of Nash equilibrium (e.g. chicken game)

Correlated Equilibria [Aumann 1974]

- Distribution over N-tuples.
- Nash Equilibrium with a shared signal
  - Independent signal - Nash equilibrium
  - Public signal - Convex combinations of Nash equilibrium
  - Private signal - not necessarily convex hull of Nash equilibrium (e.g. chicken game)

Properties:

- Contains the convex hull of Nash equilibrium.
- Can be computed efficiently

Equilibria Types

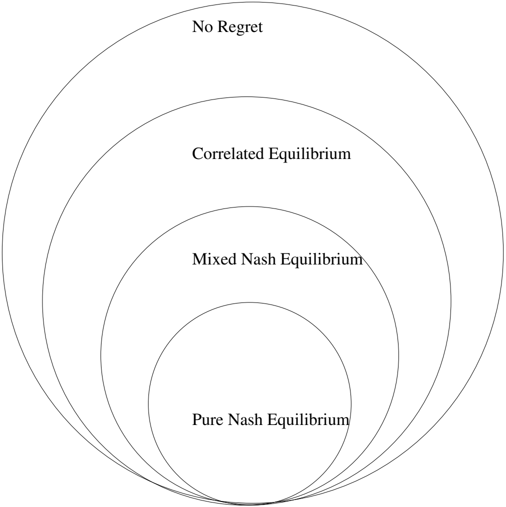

No Regret convergence

- **No-internal-regret convergence to Correlated equilibria**
  - [Hart and Mas-Collel, Foster and Vohra] If every player plays a no internal regret algorithm, then the empirical distributions of play converge almost surely as $t \longrightarrow \infty$ to the set of correlated equilibrium distributions of the game
  - The convergence is of the empirical distributions and not at a specific time.

No Regret convergence

- **No-internal-regret convergence to Correlated equilibria**
  - [Hart and Mas-Collel, Foster and Vohra] If every player plays a no internal regret algorithm, then the empirical distributions of play converge almost surely as $t \longrightarrow \infty$ to the set of correlated equilibrium distributions of the game
  - The convergence is of the empirical distributions and not at a specific time.
- **No-external-regret and zero sum games**
  - [Freund and Schapire Game and Economic Behavior 98]

No Regret convergence

- **No-internal-regret convergence to Correlated equilibria**
  - [Hart and Mas-Collel, Foster and Vohra] If every player plays a no internal regret algorithm, then the empirical distributions of play converge almost surely as $t \to \infty$ to the set of correlated equilibrium distributions of the game
  - The convergence is of the empirical distributions and not at a specific time.
- **No-external-regret and zero sum games**
  - [Freund and Schapire Game and Economic Behavior 98]
- **No-external-regret and Routing games**
  - Atomic games specific update rule[Kleinberg, Piliouras and Tardos STOC 09], Parallel links [Blum, Even-dar and Ligett PODC 06]
  - Splittable traffic [Even-dar, Mansour and Nadav STOC 09]
  - Infinitesimal users (Wardrop model) [Blum, Even-dar and Ligett PODC 06]

Outline

Quality of playing no-regret

- **In congestion games same bounds hold through similar arguments [Roughgarden STOC 09]**
- **Valid utility games and Hotelling games [Blum et al. STOC 08]**

Quality of playing no regret

Recall

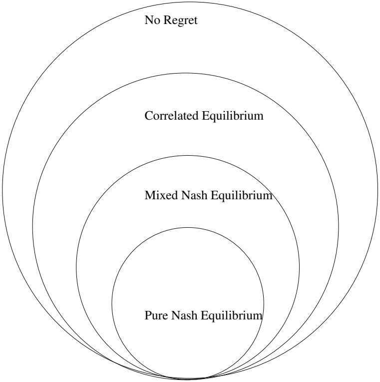

Quality of playing no regret

Recall

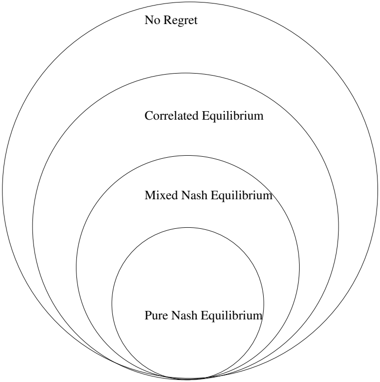

price of **No regret** $\geq$ price of **Correlated** $\geq$ price of **Mixed N.E** $\geq$ price of **Pure N.E**

Load balancing example

Consider n parallel links and n identical users and Makespan metric then:

Load balancing example

Consider n parallel links and n identical users and Makespan metric then:

Correlated Eq. and No regret: PofA $= \sqrt{n}$

Mixed N.E: PofA $= \log n / \log\log n$

Pure N.E and sink : PofA $= \mathbf{1}$

Valid-Utility Games

Consider valid-utility games then:

Valid-Utility Games

Consider valid-utility games then:

Sink Eq.: PofA $\geq$ n

Pure N.E to No Regret : PofA = 2

Outline

Learning Algorithms

- In many realistic games learning algorithms can lead to Nash equilibrium or high quality state.
  - Can be used for computing Nash equilibria.
- What can we say about games where nice behavior is not guaranteed?
- Effect of using machine learning algorithms and game dynamics in (ad) auctions (or everywhere...)

Conclusions and Future Directions

Questions about Dynamics

1. What do players converge to?
   Find potential functions? Characterize sink equilibria?
2. How long does it take?
   PLS-complete?
3. Do they quickly reach a state with small social cost?
   Performance of equilibria? Random or ε-dynamics.
4. Take your favorite game and answer these questions.
   Ad auctions, scheduling games, distributed caching games, ...

Thank You

Special thanks to Eyal Even Dar for sharing his slides with us from another joint tutorial.