Convergence of Natural Game Dynamics

Eyal Even Dar
Vahab Mirrokni
Google Research, New York

Outline

- Equilibria and Game Dynamics
- Convergence to Equilibria
  - Nash Dynamics
  - Regret-Minimization Dynamics
- Convergence to Nearly-Optimal Solutions
  - Nash Dynamics
  - Regret-Minimization Dynamics
- 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 $c_i : \Sigma_1 \times \cdots \times \Sigma_n \to \mathbb{R}$
  ($S \in \Sigma_1 \times \cdots \times \Sigma_n$ 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

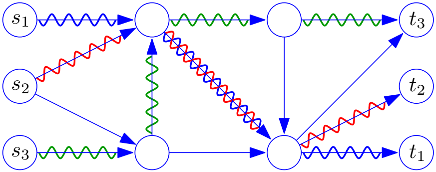

Games

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 $c_i : \Sigma_1 \times \cdots \times \Sigma_n \to \mathbb{R}$ — travel time

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

Example: Network Congestion Games

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

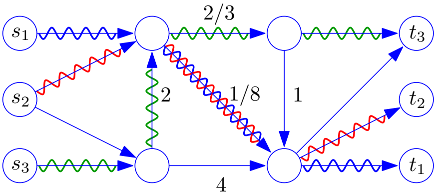

$c_1(S) = 8$

$c_2(S) = 8$

$c_3(S) = 2$

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

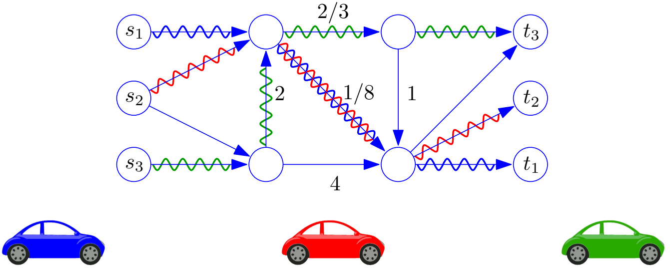

$c_1(S) = 8$ $\quad$ $c_2(S) = 8$ $\quad$ $c_3(S) = 2$

We consider only games with complete information.

Nash Equilibria

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

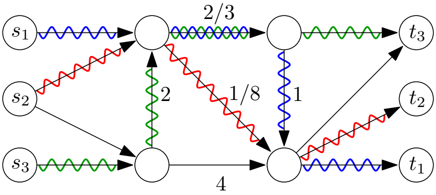

Definition

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

Nash Equilibria

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

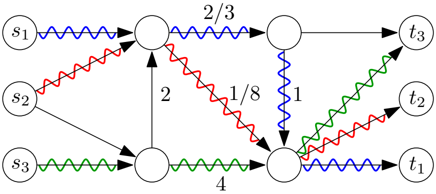

Definition

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

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

Nash Equilibria

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

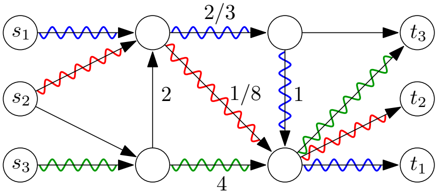

Definition

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

Congestion Games

Congestion Game:

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

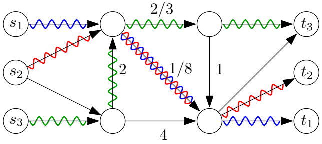

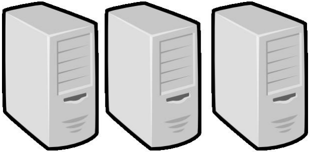

- set of strategies, $\forall i \in \mathcal{N} : \Sigma_i \subseteq 2^{\mathcal{R}}$

Congestion Games

Congestion Game:

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

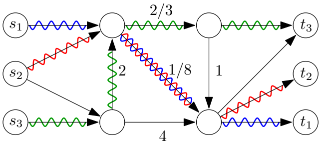

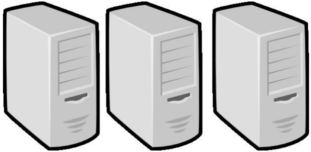

- set of strategies, $\forall i \in \mathcal{N} : \Sigma_i \subseteq 2^{\mathcal{R}}$
  $\Sigma_i = \{P \subseteq \mathcal{R} \mid P \text{ path } s_i \to t_i\}$ (network congestion game)
  $\Sigma_i = \{P \subseteq \mathcal{R} \mid P \text{ path } s \to t\}$ (symmetric congestion game)
  $\Sigma_i =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

Outline

- **Equilibria and Game Dynamics**
- **Convergence to Equilibria**
  - **Nash Dynamics**
  - **Regret-Minimization Dynamics**
- **Convergence to Nearly-Optimal Solutions**
  - **Nash Dynamics**
  - **Regret-Minimization Dynamics**
- **Conclusion**

Nash Dynamics

- Nash Dynamics: Sequence of best responses of players.

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

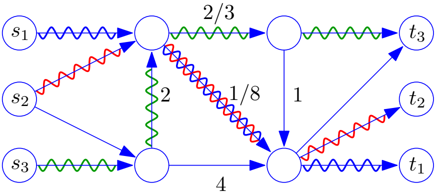

Nash Dynamics

- Nash Dynamics: Sequence of best responses of players.

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

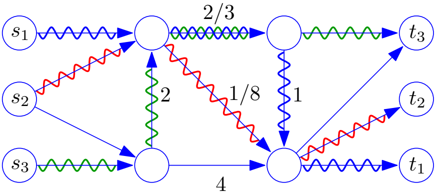

Nash Dynamics

- Nash Dynamics: Sequence of best responses of players.

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

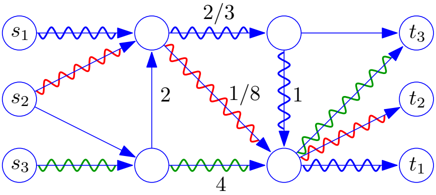

Nash Dynamics

- Nash Dynamics: Sequence of best responses of players.

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

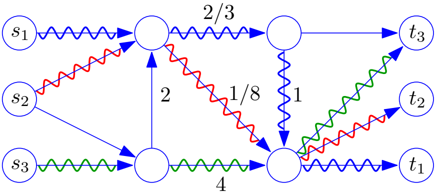

- Nash Dynamics with Liveness Property: Each player gets a chance to play his/her best response after at most t steps.

Nash Dynamics

- Nash Dynamics: Sequence of best responses of players.

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

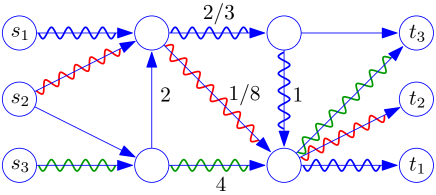

- Nash Dynamics with Liveness Property: Each player gets a chance to play his/her best response after at most t steps.
- Random Nash Dynamics: Players are chosen uniformly at random.

Nash Dynamics

- Nash Dynamics: Sequence of best responses of players.

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

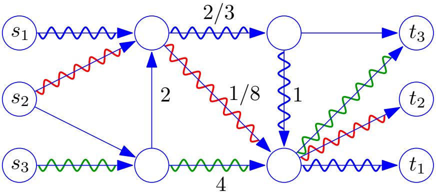

- Nash Dynamics with Liveness Property: Each player gets a chance to play his/her best response after at most t steps.
- Random Nash Dynamics: Players are chosen uniformly at random.
- ϵ-Nash Dynamics: Players change their strategy only if they can improve their own cost by a factor of at least $1 +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

Nash Dynamics

- Nash Dynamics: Sequence of best responses of players.

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

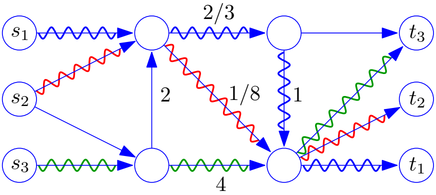

- Nash Dynamics with Liveness Property: Each player gets a chance to play his/her best response after at most t steps.
- Random Nash Dynamics: Players are chosen uniformly at random.
- ϵ-Nash Dynamics: Players change their strategy only if they can improve their own cost by a factor of at least $1 +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

Other dynamics are discussed later.

The State Graph

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

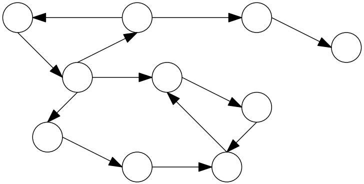

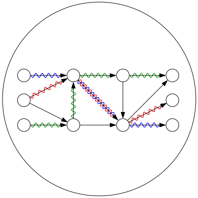

$V =$ **states** $\quad E =$ **better/best responses**

There is an edge from state S to S' with label i if player i improves her cost from S to S'.

The State Graph

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

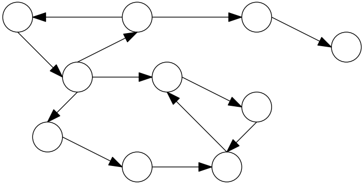

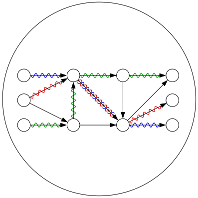

V = **states** E = **better/best responses**

There is an edge from state S to S' with label i if player i improves her cost from S to S'.

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

- pure Nash equilibrium = sink nodes of state graph

The State Graph

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

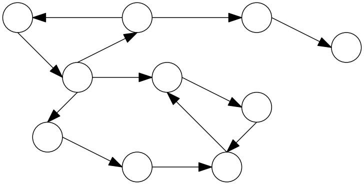

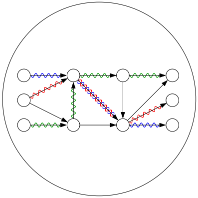

V = **states** E = **better/best responses**

There is an edge from state S to S' with label i if player i improves her cost from S to S'.

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

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

The State Graph

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

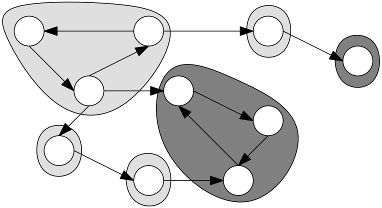

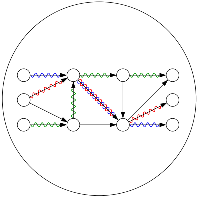

$V =$ **states** $E =$ **better/best responses**

There is an edge from state S to S' with label i if player i improves her cost from S to S'.

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

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

Outline

- **Equilibria and Game Dynamics**
- **Convergence to Equilibria**
  - **Nash Dynamics**
    - **Potential Games**
    - **NonPotential Games**
  - **Regret-Minimization Dynamics**
- **Convergence to Nearly-Optimal Solutions**
  - **Nash Dynamics**
  - **Regret-Minimization Dynamics**
- **Conclusion**

Rosenthal's Potential Function for Congestion Games

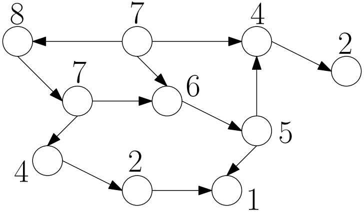

Rosenthal (Int. Journal of Game Theory 1973)

Every congestion game admits an exact potential function.

- $\Phi\colon \Sigma_1 \times \cdots \times \Sigma_n \to \mathbb{N}$ with $0 \le \Phi \le n \cdot m \cdot d_{\max}$
- player decreases his delay by $x \in \mathbb{N} \Rightarrow \Phi$ decreases by x as well

Rosenthal's Potential Function for Congestion Games

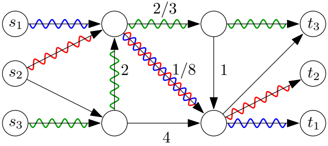

$\phi(S) = 2 + 2 + (1 + 8) = 13$

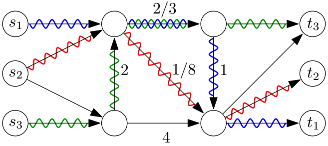

$\phi(S') = 2+(2+3)+1+1 = 9$

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 $0 \leq \Phi \leq n \cdot m \cdot d_{\max}$
- player decreases his delay by $x \in \mathbb{N} \Rightarrow \Phi$ decreases by x as well
- $n_r =$ number of players i with $r \in S_i \in \Sigma_i$

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

Rosenthal's Potential Function for Congestion Games

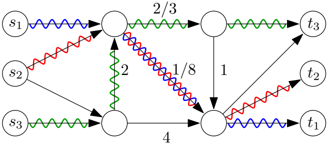

$\phi(S) = 2 + 2 + (1 + 8) = 13$

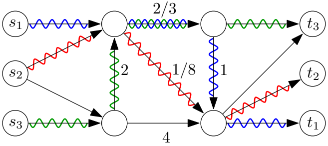

$\phi(S') = 2+(2+3)+1+1 = 9$

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 $0 \leq \Phi \leq n \cdot m \cdot d_{\max}$
- player decreases his delay by $x \in \mathbb{N} \Rightarrow \Phi$ decreases by x as well
- n_r = number of players i with $r \in S_i \in \Sigma_i$

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

$\Rightarrow$ The state graph is acyclic.

Known Results on Convergence Time

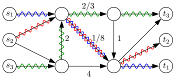

Fabrikant, Papadimitriou, Talwar (STOC 04)

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

Known Results on Convergence Time

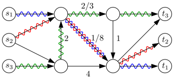

Fabrikant, Papadimitriou, Talwar (STOC 04)

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

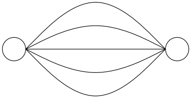

Ieong, McGrew, Nudelman, Shoham, Sun (AAAI 05)

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

Known Results on Convergence Time

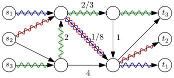

Fabrikant, Papadimitriou, Talwar (STOC 04)

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

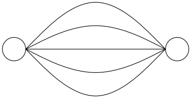

Ieong, McGrew, Nudelman, Shoham, Sun (AAAI 05)

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

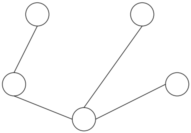

Ackermann, Röglin, Vöcking (FOCS 06)

- 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 delays without affecting the game!

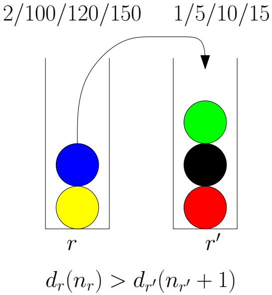

Singleton Games

Singleton Games

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

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

$$d_r(n_r) > d_{r'}(n_{r'} + 1)$$

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

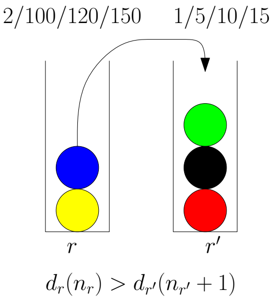

Singleton Games

Singleton Games

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

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

$$d_r(n_r) > d_{r'}(n_{r'} + 1)$$

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

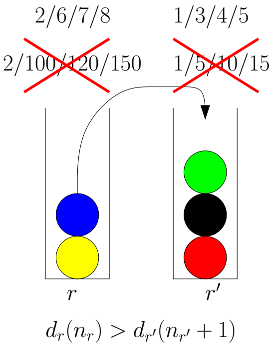

Singleton Games

Singleton Games

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

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

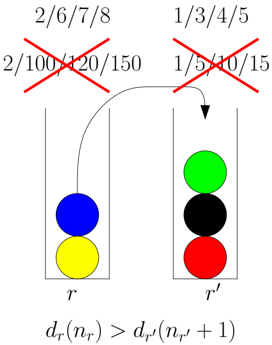

Network Congestion Games

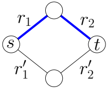

$$d_{r_1}(n_{r_1}) + d_{r_2}(n_{r_2}) > d_{r'_1}(n_{r'_1} + 1) + d_{r'_2}(n_{r'_2} + 1)$$

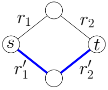

Singleton Games

Singleton Games

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

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

$$d_r(n_r) > d_{r'}(n_{r'} + 1)$$

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

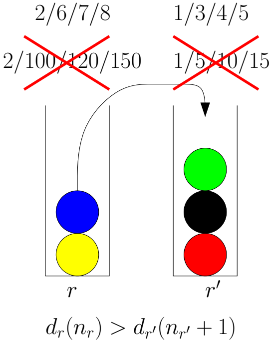

Network Congestion Games

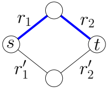

$$d_{r_1}(n_{r_1}) + d_{r_2}(n_{r_2}) > d_{r'_1}(n_{r'_1} + 1) + d_{r'_2}(n_{r'_2} + 1)$$

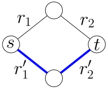

However, delay reduction works also for matroid games.

PLS: Polynomial Local Search Problems

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: Polynomial Local Search Problems

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 $\mathcal{N}(S, I) \subseteq \mathcal{F}(I)$

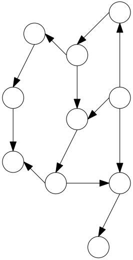

Johnson, Papadimitriou, Yannakakis (FOCS 85)

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

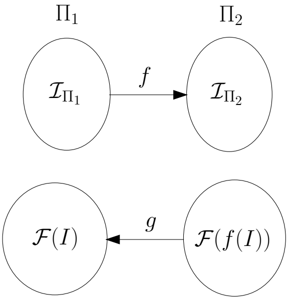

PLS-reductions

PLS-reduction

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

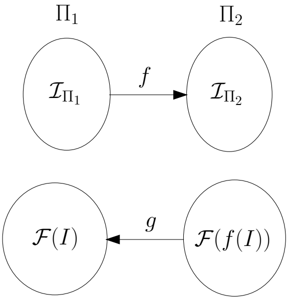

PLS-reductions

PLS-reduction

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

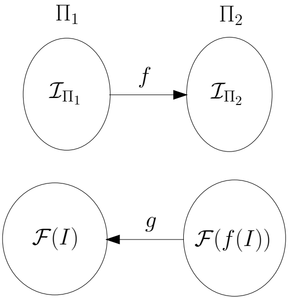

- local opt. of Π_2 easy to find $\Rightarrow$ local opt. of Π_1 easy to find
- local opt. of Π_1 hard to find $\Rightarrow$ local opt. of Π_2 hard to find

PLS-reductions

PLS-reduction

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

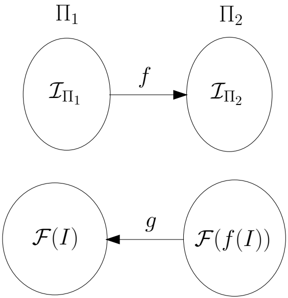

- local opt. of Π_2 easy to find $\Rightarrow$ local opt. of Π_1 easy to find
- local opt. of Π_1 hard to find $\Rightarrow$ local opt. of Π_2 hard to find

- Tight reduction implies exponential running time.

Network Congestion Games and PLS

Fabrikant, Papadimitriou, Talwar (STOC 04),Ackermann, Röglin, Vöcking (FOCS 06)

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

Network Congestion Games and PLS

Fabrikant, Papadimitriou, Talwar (STOC 04),Ackermann, Röglin, Vöcking (FOCS 06)

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

⇒ Computing a pure NE is hard.

Also, the PLS-reduction is **tight**.

⇒ 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 $\forall i \in \mathcal{N}$: delay of player $i \leq$ $(1+\varepsilon) \cdot$ min achievable delay 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 $\forall i \in \mathcal{N}$: delay of player $i \leq$
$(1+\varepsilon) \cdot$ min achievable delay 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 $\mathrm{poly}(n, \alpha, \varepsilon^{-1}, \log(d_{\max}))$ 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 $\forall i \in \mathcal{N}$: delay of player $i \leq$
$(1+\varepsilon) \cdot$ min achievable delay 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 $\mathrm{poly}(n, \alpha, \varepsilon^{-1}, \log(d_{\max}))$ 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 $\forall i \in \mathcal{N}$: delay of player $i \leq$
$(1+\varepsilon) \cdot$ min achievable delay 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.

Outline

- **Equilibria and Game Dynamics**
- **Convergence to Equilibria**
  - **Nash Dynamics**
    - **Potential Games**
    - **NonPotential Games**
  - **Regret-Minimization Dynamics**
- **Convergence to Nearly-Optimal Solutions**
  - **Nash Dynamics**
  - **Regret-Minimization Dynamics**
- **Conclusion**

Non-potential Games

Sink equilibrium: strongly connected comp. of state graph w/o outgoing edges [Goemans, M., Vetta]

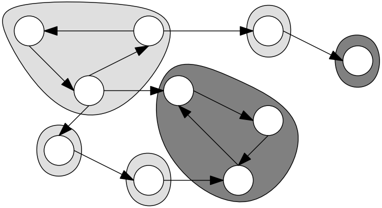

⇒ random Nash dynamics eventually reaches sink equilibrium

Non-potential Games

Sink equilibrium: strongly connected comp. of state graph w/o outgoing edges [Goemans, M., Vetta]

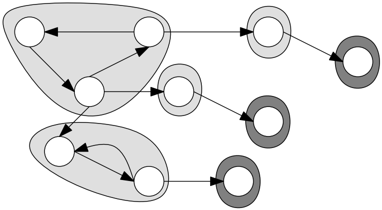

⇒ random Nash dynamics eventually reaches sink equilibrium
Interesting class: Games with only singleton sink equilibria
Example: player-specific singleton congestion games.

Milchtaich, Games and Economics Behaviour, 1996

In player-specific singleton congestion games the best-response dynamics can cycle. From every state there is a sequence of best-responses to a pure 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 $\mathcal{Y}$

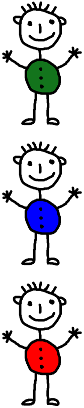

The Stable Marriage Problem

Set of women $\mathcal{X}$

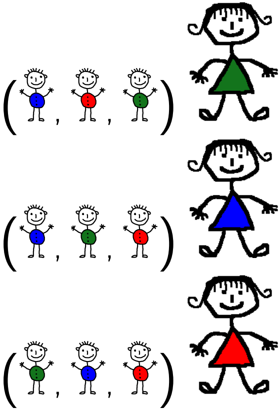

Set of men $\mathcal{Y}$

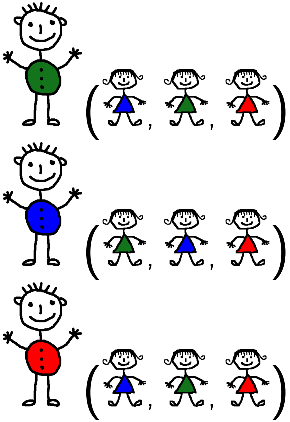

Every person has a preference list.

The Stable Marriage Problem

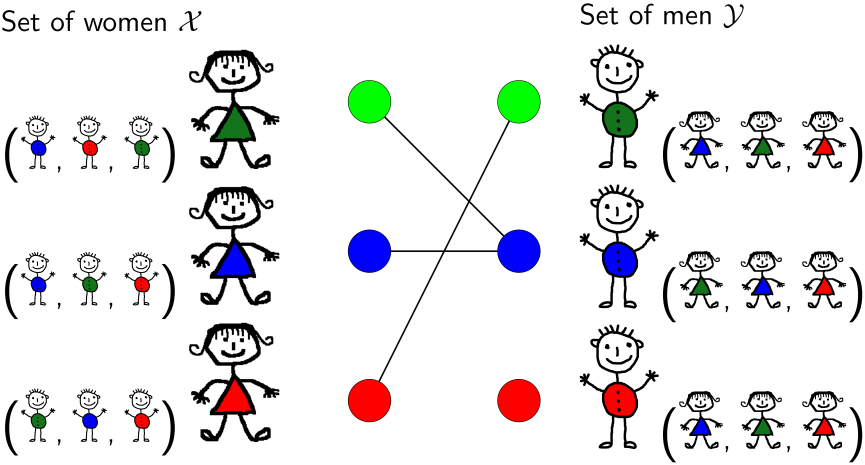

Every person has a preference list.

The Stable Marriage Problem

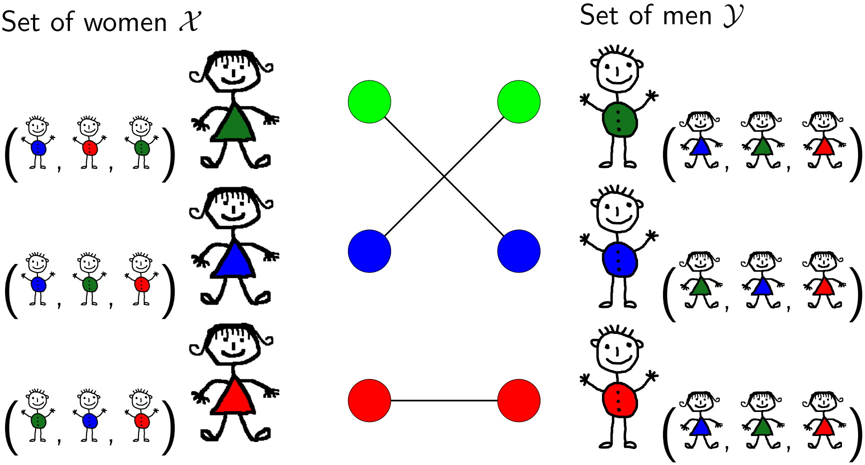

Every person has a preference list.

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.

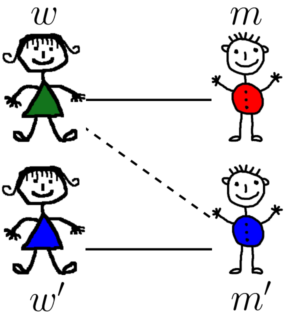

(w, m') is blocking pair

$\Longleftrightarrow$

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.

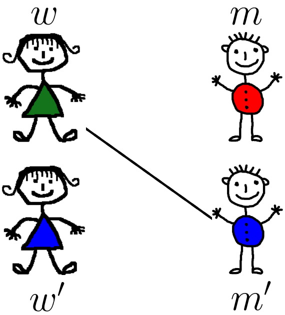

(w, m') is blocking pair

$\Longleftrightarrow$

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.

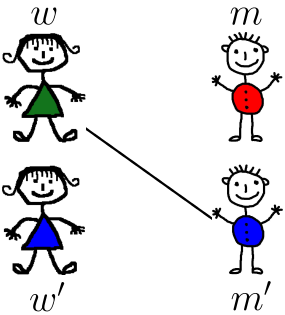

(w, m') is blocking pair

$\Longleftrightarrow$

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.

- Many further results since the 60s: roommates, ties, incomplete preferences, many-to-many matchings, etc.

Applications and Previous Work

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

- Many further results since the 60s: roommates, ties, incomplete preferences, many-to-many matchings, etc.
- Mechanism Design Questions: Can players benefit from lying? [Roth 1982]

Applications and Previous Work

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

- Many further results since the 60s: roommates, ties, incomplete preferences, many-to-many matchings, etc.
- Mechanism Design Questions: Can players benefit from lying? [Roth 1982]

Main Question

What happens without central authority?

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

Best Response Dynamics

Matching not stable ⇒ Choose woman, let her play best response.

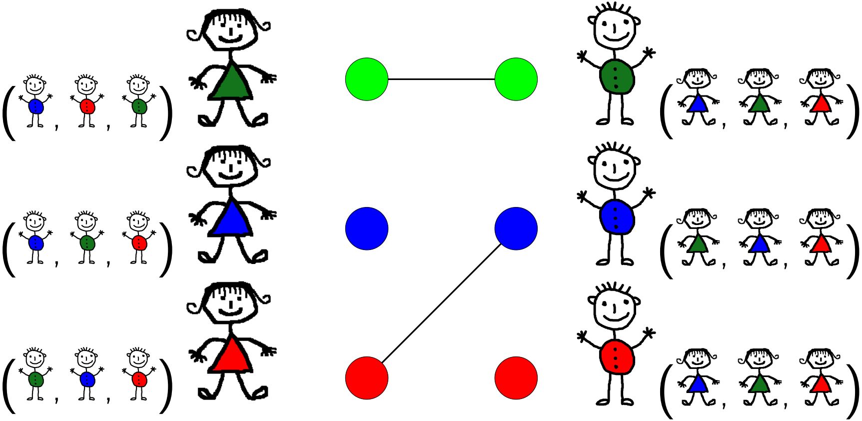

Best Response Dynamics

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

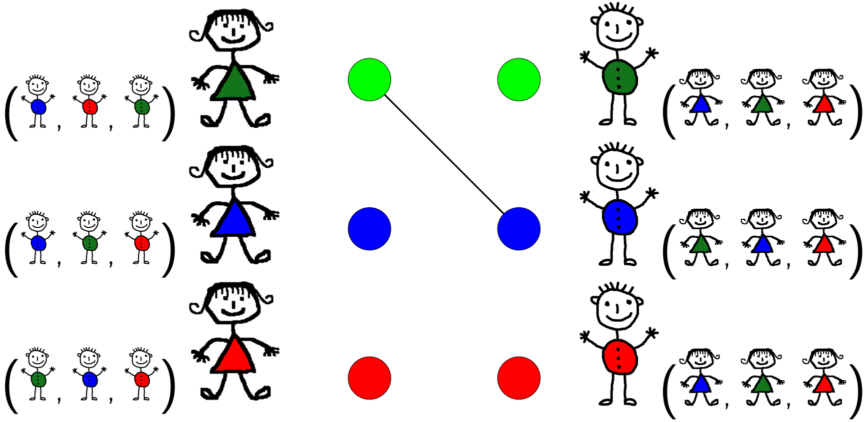

Best Response Dynamics

Matching not stable ⇒ Choose woman, let her play best response.

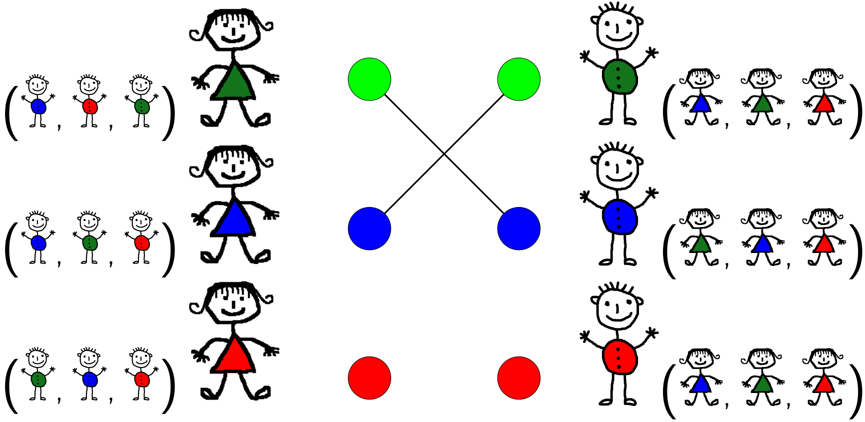

Best Response Dynamics

Matching not stable ⇒ Choose woman, let her play best response.

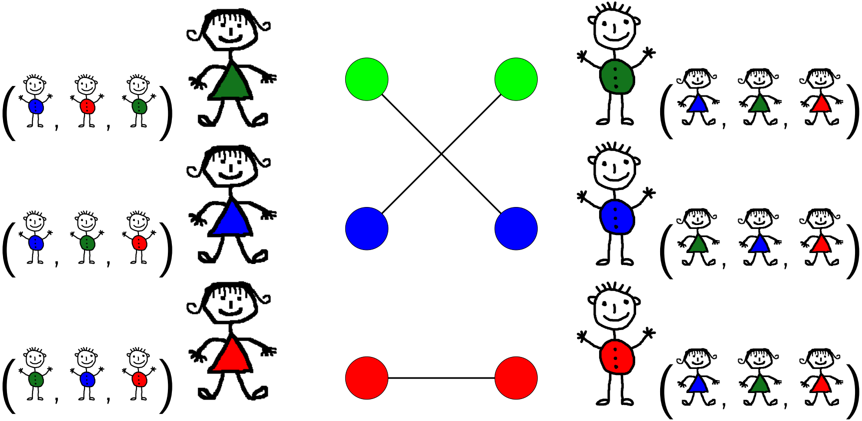

Best Response Dynamics

Good news:

Theorem

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

$\Rightarrow$ Random best-response dynamics reaches a stable matching with probability 1.

Best Response Dynamics – Good News

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 – Good News

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 – Good News

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. □

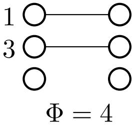

Best Response Dynamics – Good News

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. □

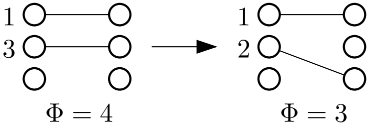

Best Response Dynamics – Good News

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. □

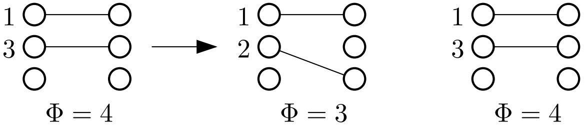

Best Response Dynamics – Good News

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. □

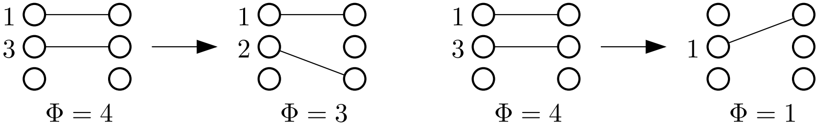

Best Response Dynamics – Good News

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 – Good News

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 – Good News

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 – Good News

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. $\square$

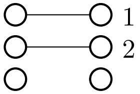

$\Psi = 5$

Best Response Dynamics – Good News

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. □

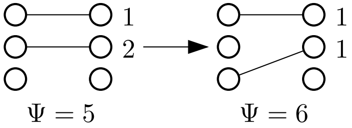

Best Response Dynamics – Good News

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. $\square$

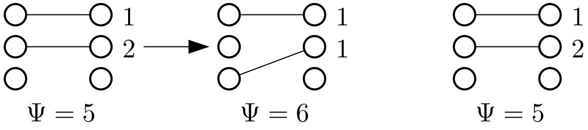

Best Response Dynamics – Good News

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. □

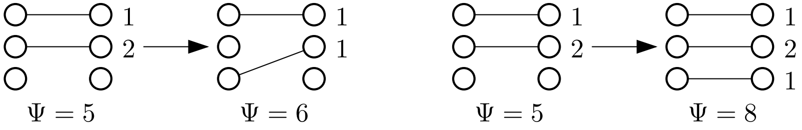

Lower Bound for Random Best Responses

Bad news:

Theorem

The best-response dynamics can cycle.

Lower Bound for Random Best Responses

Bad news:

Theorem

The best-response dynamics can cycle.

Theorem

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

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.

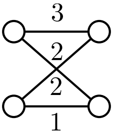

Further Results – Correlated Instances

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.

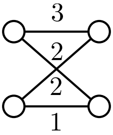

Theorem

Random best/better responses converge in polynomial time whp.

Outline

- **Equilibria and Game Dynamics**
- **Convergence to Equilibria**
  - **Nash Dynamics**
  - **Regret-Minimization Dynamics**
- **Convergence to Nearly-Optimal Solutions**
  - **Nash Dynamics**
  - **Regret-Minimization Dynamics**
- **Conclusion**

Natural Distributed/Synchronous Dynamics

- Fictitious Play
- Replicator dynamics
- No regret

Natural Distributed/Synchronous Dynamics

- **Fictitious Play**
  - **play a best response to the empirical distribution of the opponents.**
  - **Nash equilibrium is an "absorbing state"**
- Replicator dynamics
- No regret

Natural Distributed/Synchronous Dynamics

- Fictitious Play
  - play a 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 [Racke et al. STOC06]**
- No regret

Natural Distributed/Synchronous Dynamics

- Fictitious Play
  - play a 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 [Racke et al. STOC 06]
- **No regret**
  - **Today**

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 [Racke et al. STOC 06]
- No regret

- **Known to converge in specific games to Nash equilibrium**
- **There exist games on which uncoupled dynamics do not converge [Hart and Mas-Collel] a simple example for no regret [Zinkevich 03]**

No regret in Congestion Games

- 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

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

  No External Regret

No regret in Congestion Games

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

  No External 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

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

  No External 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 Internal Regret

No regret in Congestion Games

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

  No External 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 Internal Regret

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

No Regret - Motivation

These properties can influence a rational user to adapt these algorithms (note that in stochastic setting these algorithms will converge to the optimal strategy)

No Regret convergence outline

- **No internal regret convergence to Correlated equilibrium**

No Regret convergence outline

- **No internal regret convergence to Correlated equilibrium**
- **No external regret and zero sum games**

No Regret convergence outline

- **No internal regret convergence to Correlated equilibrium**
- **No external regret and zero sum games**
- No external regret - non convergence examples

No Regret convergence outline

- **No internal regret convergence to Correlated equilibrium**
- **No external regret and zero sum games**
- No external regret - non convergence examples
- **No external regret and Routing games**

No Regret convergence outline

- **No internal regret convergence to Correlated equilibrium**
- **No external regret and zero sum games**
- No external regret - non convergence examples
- **No external regret and Routing games**
- **No external regret and socially concave games**

Equilibria Types

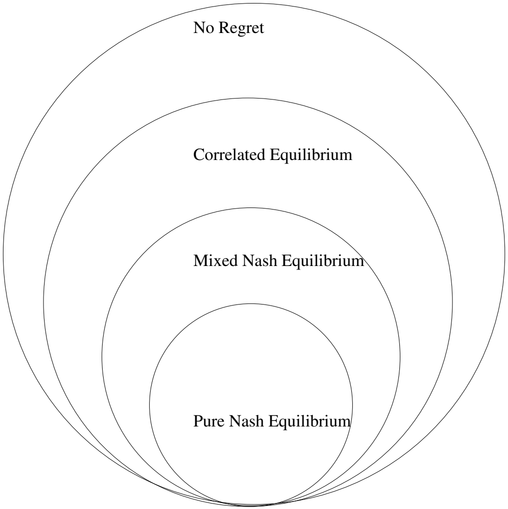

Correlated Equilibria [Aumann 1974]

- **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**

No internal regret convergence theorem

[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 internal regret simple algorithm

Regret Matching [Hart and Mas-Collel]

No internal regret simple algorithm

Regret Matching [Hart and Mas-Collel]

- **Inertia**

No internal regret simple algorithm

Regret Matching [Hart and Mas-Collel]

- **Inertia**
- **Switching probability**

No internal regret simple algorithm

Regret Matching [Hart and Mas-Collel]

- **Inertia**
- **Switching probability**
  - $R(i,k)$ - **regret of not playing** k **instead of** i
  - **Switching to action** j **from action** i **is proportional to** $R(i,j)$

No internal regret simple algorithm

Regret Matching [Hart and Mas-Collel]

- **Inertia**
- **Switching probability**
  - $R(i,k)$ - **regret of not playing** k **instead of** i
  - **Switching to action** j **from action** i **is proportional to** $R(i,j)$

There exists many others algorithms, see Foster and Vohra, Blum and Mansour, Lugosi and Stoltz ,...

No internal regret simple algorithm

Regret Matching [Hart and Mas-Collel]

- **Inertia**
- **Switching probability**
  - $R(i, k)$ - **regret of not playing** k **instead of** i
  - **Switching to action** j **from action** i **is proportional to** $R(i, j)$

There exists many others algorithms, see Foster and Vohra, Blum and Mansour, Lugosi and Stoltz ,...

Computational side:

- **All implementation requires space which is number of actions^2**
- **No efficient implementation for continuous case**
- **Influence convergence rate as well.**

No external regret and online learning

No external regret and online learning

How to measure online algorithms?

No external regret and online learning

How to measure online algorithms?

- Number of mistakes
- Regret to a best hypothesis in a class

No external Regret - generic algorithm

Follow the Regularized Leader

Let ℓ_τ be the loss function at time τ

$$w_{t+1} = argmin_w [\sum_{\tau=1}^{T} \eta \ell_\tau(w) + \texttt{Regulizer}(w)]$$

No external Regret - generic algorithm

Follow the Regularized Leader

Let ℓ_τ be the loss function at time τ

$$w_{t+1} = argmin_w [\sum_{\tau=1}^{T} \eta \ell_\tau(w) + \texttt{Regulizer}(w)]$$

Includes, gradient descent, weight majority and more.

No External Regret - History

Evolution of Bounds

Author and Year	Rate	Remarks
Hannan 56	$O(\sqrt{NT})$	Adapted by KV
Blackwell 57	$O(\sqrt{NT})$	Sufficient conditions
Littlestone and Warmuth 89	$O(\sqrt{\log(N)T})$	weighted majority
Cesa Bianchi et al. 93	$O(\sqrt{log(N)T})$	Optimal

No External Regret - History

For the bandit setting

Lai and Robinns 85	$O(\log T)$	Normal dist.
Auer et al. 95	$O(\sqrt{T})$	Simplex
Bartlett et al. 08	$O(\sqrt{T})$	More sets less efficcent
Aberenthy et al. 09	$O(\sqrt{T})$	Convex sets efficient

Applications to special cases

Author	Settings
Helmbbold and Schapire	Prunning Decision trees
Takimoto and Warmuth	shortest path
Kalai and Vempala	Hannan's algorithm for many settings
E. et al.	MDPs
Zinkevich	Convex functions
Aggarwarl at al	strongly convex function
Lugosi et al.	Bin Packing
E. et al	Load balancing

convergence in two players zero sum games

[Freund and Schapire Game and Economic Behavior 98]

- M - the first player loss matrix.
- Minmax/Maxmin strategies: p^*, q^*
- v value of the game, $M(p^*, q^*) = v$.

convergence in two players zero sum games

[Freund and Schapire Game and Economic Behavior 98]

- M - the first player loss matrix.
- Minmax/Maxmin strategies: p^*, q^*
- v value of the game, $M(p^*, q^*) = v$.

Let p_t, q_t be the strategies taken at time t:

$$\sum_{t=1}^{T} M(p_t, q_t) \quad \leq$$

convergence in two players zero sum games

[Freund and Schapire Game and Economic Behavior 98]

- M - the first player loss matrix.
- Minmax/Maxmin strategies: p^*, q^*
- v value of the game, $M(p^*, q^*) = v$.

Let p_t, q_t be the strategies taken at time t:

$$\sum_{t=1}^{T} M(p_t, q_t) \quad \leq \quad \min_p \sum_{t=1}^{T} M(p, q_t) + R(T)$$

convergence in two players zero sum games

[Freund and Schapire Game and Economic Behavior 98]

- M - the first player loss matrix.
- Minmax/Maxmin strategies: p^*, q^*
- v value of the game, $M(p^*, q^*) = v$.

Let p_t, q_t be the strategies taken at time t:

$$\sum_{t=1}^{T} M(p_t, q_t) \quad \leq \quad \min_p \sum_{t=1}^{T} M(p, q_t) + R(T) \leq \sum_{t=1}^{T} M(p^*, q_t) + R(T)$$

convergence in two players zero sum games

[Freund and Schapire Game and Economic Behavior 98]

- M - the first player loss matrix.
- Minmax/Maxmin strategies: p^*, q^*
- v value of the game, $M(p^*, q^*) = v$.

Let p_t, q_t be the strategies taken at time t:

$$\begin{aligned}\sum_{t=1}^{T} M(p_t, q_t) &\leq \min_p \sum_{t=1}^{T} M(p, q_t) + R(T) \leq \sum_{t=1}^{T} M(p^*, q_t) + R(T) \\ &\leq T \cdot v + R(T)\end{aligned}$$

No External Regret and Routing Games

- **Atomic games specific update rule[Kleinberg, Piliouras and Tardos STOC 09], Parallel links [Blum, E. and Ligett PODC 06]**
- **Splittable traffic [E., Mansour and Nadav STOC 09]**
- **Infinitesimal users (Wardrop model) [Blum, E. and Ligett PODC 06]**

No External Regret and Routing Games

- Atomic games specific update rule[Kleinberg, Piliouras and Tardos STOC 09], Parallel links [Blum, E. and Ligett PODC 06]
- Splittable traffic [E., Mansour and Nadav STOC 09]
- Infinitesimal users (Wardrop model) [Blum, E. and Ligett PODC 06]

Wardrop Model

A change in the model

- **Infinitesimal users, assume over all traffic is 1**
- **All latency function $d_e(x)$ are non decreasing**
- **Multi commodity flow with K types**

Wardrop Model

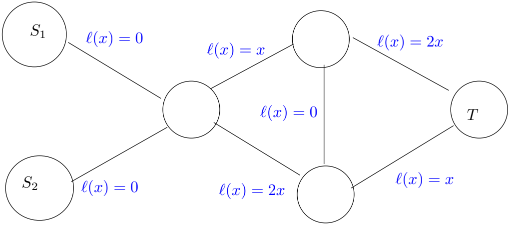

Convergence type

- L_1 **convergence**
- **All users converge to a pure Nash equilibrium**

Convergence type

- L_1 convergence
- All users converge to a pure Nash equilibrium

A flow f is at equilibrium if and only if for every player type i, and paths $P_1, P_2 \in \mathcal{P}_i$ with $f_{P_1} > 0$, $\ell_{P_1}(f) \leq \ell_{P_2}(f)$.

- $\mathcal{P}_i$ - **possible paths for type** i
- f_{P_j} - **flow on** P_j

We would like both the average flow and the average cost to converges to Nash equilibrium

Convergence Theorems[Blum, E, Liggett]

Theorem

Let $\epsilon' = \epsilon + 2\sqrt{s\epsilon n}$*. Then for general functions with maximum slope s, for* $T \geq T_\epsilon$*, the time-average flow is* ϵ'*-Nash, that is,*

$$\sum_{e \in E} \ell_e(\hat{f}_e)\hat{f}_e \leq \epsilon + 2\sqrt{s\epsilon n} + \sum_i a_i \min_{P \in \mathcal{P}_i} \sum_{e \in P} \ell_e(\hat{f}_e).$$

Theorem

In general routing games with general delay functions with maximum slope s*, for all but a* $(ms^{1/4}\epsilon^{1/4})$ *fraction of time steps up to time* T_ϵ*,* f^t *is a* $(\epsilon + 2\sqrt{s\epsilon n} + 2m^{3/4}s^{1/4}\epsilon^{1/4})$*-Nash flow.*

Simple theorem and proof for linear latency functions

Theorem

Suppose the delay functions are linear. Then for $T \geq T_\epsilon$, the average flow $\hat{f}$ is ϵ-Nash, i.e.

$$C(\hat{f}) \leq \epsilon + \sum_i a_i \min_{P \in \mathcal{P}_i} \sum_{e \in P} \ell_e(\hat{f}_e).$$

Simple proof for linear delay functions

Linearity:

- $\ell_e(\hat{f}_e) = \frac{1}{T}\sum_{t=1}^{T} \ell_e(f_e^t)$
- $\ell_e(f_e^t)f_e^t$ **is convex**

Simple proof for linear delay functions

Linearity:

- $\ell_e(\hat{f}_e) = \frac{1}{T}\sum_{t=1}^{T} \ell_e(f_e^t)$
- $\ell_e(f_e^t) f_e^t$ **is convex**

Convexity:

$$\ell_e(\hat{f}_e)\hat{f}_e \leq \frac{1}{T}\sum_{t=1}^{T} \ell_e(f_e^t) f_e^t .$$

Simple proof for linear delay functions

Linearity:

- $\ell_e(\hat{f}_e) = \frac{1}{T}\sum_{t=1}^{T} \ell_e(f_e^t)$
- $\ell_e(f_e^t) f_e^t$ **is convex**

Convexity:

$$\ell_e(\hat{f}_e)\hat{f}_e \leq \frac{1}{T}\sum_{t=1}^{T} \ell_e(f_e^t) f_e^t.$$

Combining all:

$$\begin{aligned} C(\hat{f}) &\leq \tfrac{1}{T}\textstyle\sum_{t=1}^{T} C(f^t) \\ &\leq \epsilon + \textstyle\sum_i a_i \min_{P\in\mathcal{P}_i} \tfrac{1}{T}\sum_{t=1}^{T}\sum_{e\in P} \ell_e(f_e^t) \\ &= \epsilon + \textstyle\sum_i a_i \min_{P\in\mathcal{P}_i} \sum_{e\in P} \ell_e(\hat{f}_e). \end{aligned}$$

Socially concave games

A subclass of concave games [Rosen]

- **There exists a combination** $\lambda_1, ..\lambda_n$ **such that** $\sum_{i=1}^{N} \lambda_i U_i(x)$ **is concave**
- $u_i(x_i, x_{-i})$ **is convex in** x_{-i}

Socially concave games

A subclass of concave games [Rosen]

- **There exists a combination** $\lambda_1, ..\lambda_n$ **such that** $\sum_{i=1}^{N} \lambda_i U_i(x)$ **is concave**
- $u_i(x_i, x_{-i})$ **is convex in** x_{-i}

Theorem (E., Mansour and Nadav)

If every player in a socially concave games follows a no regret policy then:

- *The average strategy vector converges to* ϵ*-Nash equilibrium*
- *The average utility converges to the payoff at* ϵ*-Nash equilibrium*

Socially concave games

- Cournot competition **(Best response does not converge)**

Socially concave games

- Cournot competition **(Best response does not converge)**
- Resource allocation. **(Best response does not converge)**

Socially concave games

- Cournot competition **(Best response does not converge)**
- Resource allocation. **(Best response does not converge)**
- Atomic splittable routing

Socially concave games

- Cournot competition **(Best response does not converge)**
- Resource allocation. **(Best response does not converge)**
- Atomic splittable routing
- Congestion control protocols.

Outline

- **Equilibria and Game Dynamics**
- **Convergence to Equilibria**
  - **Nash Dynamics**
  - **Regret-Minimization Dynamics**
- **Convergence to Nearly-Optimal Solutions**
  - **Nash Dynamics**
  - **Regret-Minimization Dynamics**
- **Conclusion**

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

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

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?

Congestion Games: Convergence to Nearly-Optimal Solutions

- 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: Convergence to Nearly-Optimal Solutions

- 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: Convergence to Nearly-Optimal Solutions

- 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($O(2^{2d})$)	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($\frac{1}{2}-\epsilon$)

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

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?

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?

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

For weighted congestion games, the price of anarchy for sink equilibria is constant.

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?

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

For weighted congestion games, the price of anarchy for sink equilibria is constant.

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, the price of anarchy for sink equilibria is very poor.

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

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?

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

For weighted congestion games, the price of anarchy for sink equilibria is constant.

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, the price of anarchy for sink equilibria is very poor.

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

- What if players follow other dynamics?

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

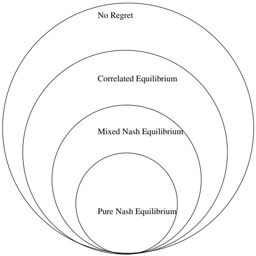

Quality of playing no regret

Recall

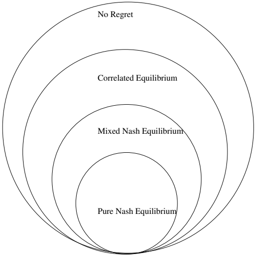

price of **No regret** $\geq$ price of **Correlated** $\geq$ price of **Mixed N.E** $\geq$ price of **Pure N.E**

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:

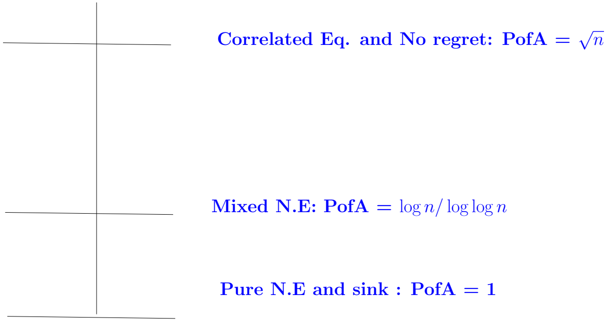

Valid-Utility Games

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

- **Equilibria and Game Dynamics**
- **Congestion Games**
- **Convergence to Equilibria**
  - **Nash Dynamics**
  - **Regret-Minimization Dynamics**
- **Convergence to Nearly-Optimal Solutions**
  - **Nash Dynamics**
  - **Regret-Minimization Dynamics**
- **Conclusion**

Conclusions and Future Directions

- In many realistic games learning algorithms can lead to Nash equilibrium or high quality state (later)
  - Can be used to explain N.E
  - Can be used for computing N.E
- What can we say about games where nice behavior is not guaranteed?
- Different types of regret for computing N.E in large games [Counterfactual, Zinkevich 07]
- Effect of using machine learning algorithms and game dynamics in (ad) Auctions (or everywhere...)

Thank You

Special thanks to Heiko Roeglin for sharing his slides with us from another joint tutorial.