

Pure Nash Equilibria

Algorithmic Game Theory

Winter 2019/20

Congestion Games

Convergence Time in Congestion Games

Complexity of Pure Nash equilibria

Stable Matching, Ordinal Potentials, Weakly Acyclic Games

Congestion Games (Rosenthal 1973)

A **congestion game** is a tuple $\Gamma = (\mathcal{N}, \mathcal{R}, (\Sigma_i)_{i \in \mathcal{N}}, (d_r)_{r \in \mathcal{R}})$ with

- ▶ $\mathcal{N} = \{1, \dots, n\}$, set of players
- ▶ $\mathcal{R} = \{1, \dots, m\}$, set of resources
- ▶ $\Sigma_i \subseteq 2^{\mathcal{R}}$, strategy space of player i
- ▶ $d_r : \{1, \dots, n\} \rightarrow \mathbb{Z}$, delay function of resource r


For any state $S = (S_1, \dots, S_n) \in \Sigma_1 \times \dots \times \Sigma_n$,

- ▶ $n_r =$ number of players with $r \in S_i$
- ▶ $d_r(n_r) =$ delay of resource r
- ▶ $\delta_i(S) = \sum_{r \in S_i} d_r(n_r) =$ delay of player i

The *cost* of player i in state S is $c_i(S) = \delta_i(S)$, that is, players aim at minimizing their delays.

Example: Network Congestion Games


- ▶ Given a directed graph $G = (V, E)$. Every edge $e \in E$ has a delay function $d_e : \{1, \dots, n\} \rightarrow \mathbb{Z}$.
- ▶ Player i wants to allocate a path of minimal delay between a source s_i and a target t_i .



- ▶ In this example, $\mathcal{N} = \{1, 2, 3\}$, $\mathcal{R} = E$, $\Sigma_i =$ set of s-t paths.
- ▶ This game is **symmetric**: All players have the same set of strategies. In any state S , if we permute the strategy choices of the players, the resulting player costs also get permuted similarly.



Example: Network Congestion Games

A sequence of (best reply) improvement steps: First step ...




Example: Network Congestion Games

A sequence of (best reply) improvement steps: First step ...





Example: Network Congestion Games

... second step ...




Example: Network Congestion Games

... second step ...





Example: Network Congestion Games

... third step ...





Example: Network Congestion Games

... third step ...



Example: Network Congestion Games

... third step ...



Pure Nash Equilibrium – stop!

Questions

- ▶ Does every congestion game have a pure Nash equilibrium?
- ▶ Is every sequence of improvement steps finite?
- ▶ How many steps are needed to reach a (pure) Nash equilibrium?
- ▶ What is the complexity of computing (pure) Nash equilibria in congestion games?

Finite Improvement Property

Theorem (Rosenthal 1973)

For every congestion game, every sequence of improvement steps is finite.

This result immediately implies

Corollary

Every congestion game has at least one pure Nash equilibrium.

Rosenthal's analysis is based on a potential function argument.


For every state S , let

$$\Phi(S) = \sum_{r \in \mathcal{R}} \sum_{k=1}^{n_r(S)} d_r(k) .$$

This function is called *Rosenthal's potential function*.

Proof of Rosenthal's Theorem


Lemma: Let S be any state. Suppose we go from S to a state S' by an improvement step of player i decreasing his delay by $\Delta > 0$. Then $\Phi(S') = \Phi(S) - \Delta$.



In the picture, the value of the potential is the shaded area. If a player changes from r' to r , his delay changes exactly as the potential value.

Proof of Rosenthal's Theorem

Lemma: Let S be any state. Suppose we go from S to a state S' by an improvement step of player i decreasing his delay by $\Delta > 0$. Then $\Phi(S') = \Phi(S) - \Delta$.



In the picture, the value of the potential is the shaded area. If a player changes from r' to r , his delay changes exactly as the potential value.

Proof of Rosenthal's Theorem

Lemma: Let S be any state. Suppose we go from S to a state S' by an improvement step of player i decreasing his delay by $\Delta > 0$. Then $\Phi(S') = \Phi(S) - \Delta$.

Proof:

- ▶ The potential $\phi(S)$ can be calculated by inserting the agents one after the other in any order, and summing the delays of the players at the point of time at their insertion.
- ▶ W.l.o.g., agent i is the last player that we insert when calculating $\Phi(S)$. For this agent i we add his actual delay in state S to the potential $\Phi(S)$.
- ▶ When going from S to S' , the delay of i decreases by Δ , and, hence, Φ decreases by exactly Δ as well. □ (Lemma)

Proof of Rosenthal's Theorem

The lemma shows that Φ is an **exact potential**, i.e., if a single player decreases its latency by a value of $\Delta > 0$, then Φ decreases by exactly the same amount.

Further observe that

- i) the delay values are integers so that, for every improvement step, $\Delta \geq 1$,
- ii) for every state S , $\Phi(S) \leq \sum_{r \in \mathcal{R}} \sum_{i=1}^n |d_r(i)|$,
- iii) for every state S , $\Phi(S) \geq -\sum_{r \in \mathcal{R}} \sum_{i=1}^n |d_r(i)|$.

Consequently, the number of improvements is upper-bounded by $2 \cdot \sum_{r \in \mathcal{R}} \sum_{i=1}^n |d_r(i)|$ and hence finite. □ (Theorem)

Potential Games

Definition (Potential Game)

A strategic game $\Gamma = (\mathcal{N}, (\Sigma_i)_{i \in \mathcal{N}}, (c_i)_{i \in \mathcal{N}})$ is called **exact potential game** if there exists a function $\Phi: \Sigma \rightarrow \mathbb{R}$ such that for every $i \in \mathcal{N}$, for every $S_{-i} \in \Sigma_{-i}$, and every $S_i, S'_i \in \Sigma_i$:

$$c_i(S_i, S_{-i}) - c_i(S'_i, S_{-i}) = \Phi(S_i, S_{-i}) - \Phi(S'_i, S_{-i}) .$$

Φ is called an **exact potential** function.

Observation

Let $\Gamma = (\mathcal{N}, (\Sigma_i)_{i \in \mathcal{N}}, (c_i)_{i \in \mathcal{N}})$ be an exact potential game. Then Γ has the **finite improvement property** and, hence, there exists a state that is a (pure) Nash equilibrium.

Congestion versus Potential Games

It follows from Rosenthal's potential function that

Corollary

Every congestion game is an exact potential game.

In some sense, the reverse is true as well.

Theorem (Monderer and Shapley, 1996)

Every exact potential game is "isomorphic" to a congestion game.

For every exact potential game, there is another game with the same players, strategies, and costs. The other game uses appropriate resources and delays such that strategies become subsets of resources and costs become sums over resource delays. In this way, one can provide a representation of the potential game as an "isomorphic" congestion game.

Congestion Games

Convergence Time in Congestion Games

Complexity of Pure Nash equilibria

Stable Matching, Ordinal Potentials, Weakly Acyclic Games

Main Question

How many improvement steps are needed to reach a pure Nash equilibrium?

Transition Graph

- ▶ The **transition graph** of a congestion game Γ contains a vertex for every state S and a directed edge (S, S') if S' can be reached from S by an improvement step of a single player.
- ▶ The **best-response transition graph** contains only edges for best response improvement steps.

A sequence of (best response) improvement steps corresponds to a path in the (best response) transition graph.

The sinks of this graph are the Nash equilibria of Γ .

The number of vertices (states) can be as large as 2^{mn} . Thus there might be paths of exponential length.

Singleton Congestion Games – Definition

Definition (Singleton Congestion Game)


A congestion game is called **singleton** if, for every $i \in \mathcal{N}$ and every $R \in \Sigma_i$, it holds that $|R| = 1$.

In a singleton game, every player wants to allocate exactly one single resource from a subset of allowed resources.

Although this constraint on the strategy sets is quite restrictive, there are still up to m^n different states.


Singleton Congestion Games – Example

Consider a “server farm” with three servers a, b, c (resources) and three players 1,2,3. Each player has a single task that needs to be processed by one of the servers. The player chooses the server strategically to minimize the completion time.



Singleton Congestion Games – Example

Consider a “server farm” with three servers a, b, c (resources) and three players 1,2,3. Each player has a single task that needs to be processed by one of the servers. The player chooses the server strategically to minimize the completion time.



Nash equilibrium

Singleton Congestion Games – Convergence

Theorem

In singleton congestion games, every improvement sequence has a length of $O(n^2m)$.

Proof idea:

- ▶ Replace original delays by bounded integer values without changing the preferences of the players.
- ▶ This yields an upper bound on the maximum potential wrt new delays.
- ▶ Due to integer values, the decrease of the potential in an improvement step is at least 1. Hence, the length of every improvement sequence is bounded by the maximum value of the potential.

Proof for Singleton Games

Sort the set of delay values $\{d_r(k) \mid r \in \mathcal{R}, 1 \leq k \leq n\}$ in increasing order. We define alternative, new delay functions:

$$\bar{d}_r(k) := \text{position of } d_r(k) \text{ in sorting.}$$

Example:

The sorted set of delay values from the previous example is

$$15, 16, 17, 20, 30, 50, 70, 90.$$

Hence, the old and new delay functions are

$$\begin{aligned} d_a(1, 2, 3) &= (20, 30, 50) & \bar{d}_a(1, 2, 3) &= (4, 5, 6) \\ d_b(1, 2, 3) &= (30, 70, 90) & \bar{d}_b(1, 2, 3) &= (5, 7, 8) \\ d_c(1, 2, 3) &= (15, 16, 17) & \bar{d}_c(1, 2, 3) &= (1, 2, 3) \end{aligned}$$

The new delay of a player i using resource r in state S is $\bar{\delta}_i(S) = \bar{d}_r(n_r(S))$.

Proof for Singleton Games

Observation:

Let S and S' be two states such that (S, S') is an improvement step for some player i w.r.t. the original delays. Then (S, S') is an improvement step for i w.r.t. the new delays, as well.

Rosenthal's potential function w.r.t. the new delays can be upper bounded as follows:

$$\bar{\Phi}(S) = \sum_{r \in \mathcal{R}} \sum_{k=1}^{n_r(S)} \bar{d}_r(k) \leq \sum_{r \in \mathcal{R}} \sum_{k=1}^{n_r(S)} n m \leq n^2 m .$$

It holds that $\bar{\Phi} \geq 1$. Also, $\bar{\Phi}$ decreases by at least 1 in every step. Therefore, the length of every improvement sequence is at most $n^2 m$. \square (Theorem)

Fast Convergence to Pure Equilibria?

For singleton games we showed that **every improvement sequence** has a length that is **polynomial in n and m** . This bound holds for every arbitrary sequence, as long as the moving player strictly decreases his cost (not only sequences of best responses).

This result can be generalized to so-called **matroid games**. In these games, every player has a strategy set that corresponds to the bases of a matroid over the set of resources. In these games, it can be shown that **sequences of best responses** have a length that is **polynomial in n and m** .

In general, however, there is for every $n \in \mathbb{N}$ at least one congestion game with

- ▶ $O(n)$ players und $O(n)$ resources,
- ▶ non-negative, monotone delays, und
- ▶ an initial state S

such that **every improvement sequence** from S to a pure Nash equilibrium has a length that is **exponential in n** .

Approximate Equilibria

How does the convergence time change if players are happy with a strategy that is *almost* as good as a best response?

Approximate Equilibria

We use $0 \leq \varepsilon < 1$ as a “tolerance” value for each player.

Definition

A best response s'_i is called **ε -improvement** for player i in state s if

$$c_i(s) - c(s'_i, s_{-i}) > \varepsilon \cdot c_i(s) .$$

The pair $(s, (s'_i, s_{-i}))$ is called **ε -improvement step**.

Definition

A state s is called **(pure) ε -approximate Nash equilibrium** if no player $i \in \mathcal{N}$ has an ε -improvement.

An ε -improvement for player i in state s decreases the cost by more than an ε -fraction of his current cost. In an ε -approximate equilibrium s there is no player i that can decrease his $c_i(s)$ by more than $\varepsilon \cdot c_i(s)$ by unilaterally deviating to a different strategy.

Convergence to Approximate Equilibrium

We examine a sequence of ε -improvement steps in a large class of games with a natural restriction on the delay functions.

Definition

A congestion game is **symmetric** if $\Sigma_i = \Sigma_j$ for all players $i, j \in \mathcal{N}$.

Definition

Suppose $\alpha \geq 1$. Let d_r be a positive, monotonically increasing delay function. d_r is called **α -bounded-jump** if for all $k \geq 2$

$$d_r(k) \leq \alpha \cdot d_r(k-1) .$$

By adding a player to resource r the delay of all players on that resource increases by at most a factor α . This implies, in particular, that $d_r(k) \leq d_r(1) \cdot \alpha^{k-1}$.

Convergence to Approximate Equilibria

Let $\Phi^* = \min_{S \in \Sigma} \Phi(S)$ be the smallest value of Rosenthal's potential function over all states of the game. For positive delays given by natural numbers we obviously have $\Phi^* \geq n$.

Theorem (Chien, Sinclair, 2011)

Let Γ be a symmetric congestion game with α -bounded-jump delays in \mathbb{N} . For every state S there is a sequence of ε -improvement steps of length of at most

$$\frac{n\alpha}{\varepsilon} \cdot \log \left(\frac{\Phi(S)}{\Phi^*} \right) = \frac{n\alpha}{\varepsilon} \cdot \log(m \cdot \max_{r \in \mathcal{R}} d_r(n)) ,$$

which ends in a ε -approximate Nash equilibrium.

Convergence to Approximate Equilibria

Proof:

We construct the sequence via **relative ε -dynamics**:

Consider $I(S)$ as the set of players with ε -improvements in S . Let S'_i be an ε -improvement for $i \in I(S)$. The deviating player $i^* \in I(S)$ is chosen as a player with maximum relative cost improvement:

$$i^* \in \arg \max_{i \in I(S)} \frac{c_i(S) - c_i(S'_i, S_{-i})}{c_i(S)} .$$

Now consider an arbitrary other player $j \neq i^*$ in state S :

- ▶ j could have deviated to S'_{i^*} as well (symmetric game)
- ▶ If j would have deviated to S'_{i^*} instead of i^* , then his cost would have been at most:

$$c_j(S'_{i^*}, S_{-j}) \leq \alpha \cdot c_{i^*}(S'_{i^*}, S_{-i^*})$$

(Why?)

Convergence to Approximate Equilibria

- ▶ Relative ϵ -dynamics chose player i^* , thus:

$$\frac{c_{i^*}(S) - c_{i^*}(S'_{i^*}, S_{-i^*})}{c_{i^*}(S)} \geq \frac{c_j(S) - c_j(S'_{i^*}, S_{-j})}{c_j(S)} \geq \frac{c_j(S) - \alpha c_{i^*}(S'_{i^*}, S_{-i^*})}{c_j(S)}$$

- ▶ Transformation of first and third fraction implies: $c_{i^*}(S) \geq c_j(S)/\alpha$.

The *moving player* i^* in S is – up to a factor α – also the player with *maximum cost*: $c_{i^*}(S) \geq \max_j c_j(S)/\alpha$.

This implies:

$$\begin{aligned} \Phi(S) - \Phi(S'_{i^*}, S_{-i^*}) &\geq \epsilon \cdot c_{i^*}(S) \geq \frac{\epsilon}{\alpha} \cdot \max_j c_j(S) \\ &\geq \frac{\epsilon}{\alpha} \cdot \frac{1}{n} \cdot \sum_j c_j(S) \geq \frac{\epsilon}{n\alpha} \cdot \Phi(S) \end{aligned}$$

Convergence to Approximate Equilibria

The potential decreases in every improvement step of the relative ε -dynamics by a factor of $\varepsilon/(n\alpha)$. Let S^0, S^1, S^2, \dots be the resulting sequence of states (with $S = S^0$). Then

$$\Phi(S^k) \leq \Phi(S^0) \cdot \left(1 - \frac{\varepsilon}{n\alpha}\right)^k .$$

We now express the value k as a multiple of $n\alpha/\varepsilon$ and observe

$$\Phi(S) \cdot \left(1 - \frac{\varepsilon}{n\alpha}\right)^k = \Phi(S) \cdot \left(1 - \frac{\varepsilon}{n\alpha}\right)^{\ell \cdot n\alpha/\varepsilon} \leq \Phi(S) \cdot \frac{1}{e^\ell}$$

since $(1-x)^{1/x} \leq (e^{-x})^{1/x} = 1/e$ for all $0 < x \leq 1$. When ℓ becomes large, then at some point it holds

$$\Phi(S)/e^\ell \leq \Phi^* .$$

Then no more improvement is possible.

Convergence to Approximate Equilibria

This implies that when

$$\ell \geq \ln \frac{\Phi(S)}{\Phi^*}$$

no more improvement step can occur. Thus, the number k of steps in the relative ε -dynamics from state S to a ε -approximate Nash equilibrium is at most

$$k = \frac{n\alpha}{\varepsilon} \cdot \ell \leq \frac{n\alpha}{\varepsilon} \cdot \ln \frac{\Phi(S)}{\Phi^*}$$



This result does not hold when we have (1) *asymmetric* games with α -bounded-jump delays or (2) symmetric games *without* α -bounded-jump delays.

Congestion Games

Convergence Time in Congestion Games

Complexity of Pure Nash equilibria


Stable Matching, Ordinal Potentials, Weakly Acyclic Games

We investigate the complexity of finding Nash equilibria in different kinds of congestion games.

Our study is restricted to congestion games with non-decreasing delay functions.

Symmetric Network Congestion Games

- ▶ Given a directed graph $G = (V, E)$ with delay functions $d_e : \{1, \dots, n\} \rightarrow \mathbb{Z}$, $e \in E$.
- ▶ Player i wants to allocate a path of minimal delay between a source s and a target t .



- ▶ In more general *asymmetric* network congestion games, different players might have different source-destination pairs.


Symmetric Network Congestion Games

- ▶ It is known that there are instances of symmetric network congestion games in which there are states such that every improvement sequence from this state to a Nash equilibrium has exponential length.
- ▶ Hence, applying improvement steps is not an *efficient* (i.e. polynomial time) algorithm for computing Nash equilibria in these games.
- ▶ However, there is another algorithm which finds Nash equilibria in polynomial time ...

Complexity in Symmetric Network Congestion Games

Efficient algorithm via a reduction to min-cost flow:
(Fabrikant, Papadimitriou, Talwar 2004)

- ▶ Each edge is replaced by n parallel edges of capacity 1 each.
- ▶ The i th copy of edge e has cost $d_e(i)$, $1 \leq i \leq n$.



- ▶ We compute a min-cost-flow, i.e., a network flow of value n from s to t that minimizes the total cost of the edge copies that are used.
- ▶ Optimal solution minimizes Rosenthal's potential function and, hence, is a pure Nash equilibrium.

Relationship to Local Search

Rosenthal's potential function allows us to interpret congestion games as local search problems:

Nash equilibria are local optima w.r.t. the potential function.

How difficult is it to compute local optima?

The complexity class PLS

Definition (PLS (Polynomial Local Search))

PLS contains search problems with an objective function and a specified neighborhood relationship Γ . It is required that there is a poly-time algorithm that, given any solution s ,

- ▶ computes a solution in $\Gamma(s)$ with better objective value, or
- ▶ certifies that s is a local optimum.

Some examples for problems in PLS

- ▶ FLIP (circuit evaluation with Flip-neighborhood)
- ▶ TSP with 2-Opt-neighborhood
- ▶ Pos-NAE- k Sat with Flip-neighborhood
- ▶ Max-Cut with Flip-neighborhood
- ▶ Congestion games w.r.t. improvement steps

Max-Cut – Definition

Input:

A graph $G = (V, E)$ with edge weights $w : E \rightarrow \mathbb{N}$.

- ▶ A **cut** partitions V into two sets *Left* and *Right*.
- ▶ Two cuts are **neighboring** if one can obtain one from the other by moving only one vertex from *Left* to *Right* or vice versa.
- ▶ The **value** of a cut is the weighted number of edges with one endpoint in *Left* and one endpoint in *Right*.

Task:

Find a **local optimum**, i.e., a cut without neighboring cut of higher value.

Fact:

Max-Cut is PLS-complete.

The Complexity Class PLS

Definition (PLS-reduction)

Given two PLS problems Π_1 and Π_2 find a mapping from the instances of Π_1 to the instances of Π_2 such that

- ▶ the mapping can be computed in polynomial time,
- ▶ the local optima of Π_1 are mapped to local optima of Π_2 , and
- ▶ given any local optimum of Π_2 , one can construct a local optimum of Π_1 in polynomial time.

PLS-Hardness for General Congestion Games

We prove a PLS-reduction from Max-Cut to congestion games.

First of all, we observe that Max-Cut can be represented as a game:

Party Affiliation Game (Max-Cut)

Players correspond to vertices in a weighted graph $G = (V, E)$.

- ▶ Every player has 2 strategies: *left* or *right*.
- ▶ A state of the game yields a **cut**, i.e., a partition of V into left and right vertices.
- ▶ Edge weights represent antisympathy among players.
- ▶ Players choose a strategy to maximize the sum of weights of incident edges crossing the cut.
- ▶ Pure Nash equilibria correspond to local optima of Max-Cut.

PLS-Hardness for General Congestion Games

Minimization Variant of the Party Affiliation Game

The strategies of a vertex are

- ▶ **left:** choose the left hand side of the cut
- ▶ **right:** choose the right hand side of the cut

The costs for these strategies are

- ▶ **left:** sum of the weights of the incident edges to the left
- ▶ **right:** sum of the weights of the incident edges to the right

Both games have the same transition graph:

For each player, minimizing the weights of incident edges on “her side”, is equivalent to maximizing the sum of edges leading to the “other side”.

PLS-Hardness for General Congestion Games

Now the minimization variant can be described in terms of a congestion game.

Party Affiliation Congestion Game:

- ▶ Represent each edge e by two resources e_{left}, e_{right} with delay functions $d(1) = 0$ and $d(2) = w_e$.
- ▶ For each player the strategy S_{left} contains resources e_{left} for all incident edges; strategy S_{right} contains resources e_{right} for all incident edges.

Players in this congestion game have exactly the same cost as players in the minimization variant of the party affiliation game.

Hence, the pure Nash equilibria of this congestion game coincide with local optima of the Max-Cut instance. Hence, we obtain a PLS-reduction from Max-Cut to congestion games. □ (PLS completeness)

Congestion Games

Convergence Time in Congestion Games

Complexity of Pure Nash equilibria

Stable Matching, Ordinal Potentials, Weakly Acyclic Games

Potential Games

An exact potential implies the finite improvement property and, thus, the existence of a pure Nash equilibrium. The definition of potential game can be made much more general without losing the finite improvement property.

Definition (Ordinal Potential Game)

A strategic game $\Gamma = (\mathcal{N}, (\Sigma_i)_{i \in \mathcal{N}}, (c_i)_{i \in \mathcal{N}})$ is called an **ordinal potential game** if there exists a function $\Phi: \Sigma \rightarrow \mathbb{R}$ such that for every $i \in \mathcal{N}$, for every $S_{-i} \in \Sigma_{-i}$, and every $S_i, S'_i \in \Sigma_i$:


$$c_i(S_i, S_{-i}) > c_i(S'_i, S_{-i}) \quad \Rightarrow \quad \Phi(S_i, S_{-i}) > \Phi(S'_i, S_{-i}) .$$

Φ is called an **ordinal potential** function.

Observation:

Let Γ be an ordinal potential game. Then Γ has the **finite improvement property**, and, hence, a pure Nash equilibrium.

Stable Matching



Every person has a **preference list** (left/right is most/least preferred). No polygamy – at most one match per person.


Stable Matching

- ▶ Set \mathcal{X} of m **men**, set \mathcal{Y} of n **women**
- ▶ Each $x \in \mathcal{X}$ has a **preference order** \succ_x over all matches $y \in \mathcal{Y}$.
- ▶ Each $y \in \mathcal{Y}$ has a **preference order** \succ_y over all matches $x \in \mathcal{X}$.
- ▶ For each person being unmatched is the least preferred state, i.e., each person wants to be **matched rather than unmatched**.
- ▶ For matching M let $M(x) \in \mathcal{Y}$ be the match of man $x \in \mathcal{X}$ in M , similarly let $M(y)$ be the match of woman $y \in \mathcal{Y}$.
- ▶ Let $M(x) = *$ if x is unmatched in M . Similar for $M(y) = *$.

Stable Matching

When is a matching stable? What is a hazard to stability?

- ▶ In a matching M , a pair $\{x, y\}$ is **blocking pair** if and only if x and y prefer each other to $y' = M(x)$ and $x' = M(y)$, respectively.
- ▶ M is a **stable matching** if and only if it admits no blocking pair.



Stable matching is a central concept in many applications.

Residents/Hospitals



College Admission



Job Market



etc.

Stable Matching as Strategic Game

We interpret the model as a **strategic game**:

- ▶ Players are the men \mathcal{X} . The strategy set of a man $x \in \mathcal{X}$ is the set of women $\Sigma_x = \mathcal{Y}$. A man picks a woman as strategy and “proposes to her”.
- ▶ We express preference order by costs: Every match $\{x, y\}$ yields cost values $c_x(y) > 0$ and $c_y(x) > 0$ for the involved agents, which satisfy

$$\begin{aligned}
 c_x(y) > c_x(y') &\Leftrightarrow y' \succ_x y, \\
 c_y(x) > c_y(x') &\Leftrightarrow x' \succ_y x, \\
 c_y(*) = c_x(*) = \infty &\quad \text{für alle } x \in \mathcal{X}, y \in \mathcal{Y}.
 \end{aligned}$$

Stable Matching as Strategic Game


- ▶ In a state S every man $x \in \mathcal{X}$ chooses a woman $y_x \in \mathcal{Y}$. In S every woman $y \in \mathcal{Y}$ receives a (possibly empty) set $A_y(S)$ of proposals.
- ▶ A match emerges only when $x_y^* = \arg \min\{c_y(x) \mid x \in A_y(S)\}$, i.e., y matches to the man from $A_y(S)$ that she likes best.
- ▶ In S man x obtains cost $c_x(S) = c_x(M_S(x))$, where $M_S(x)$ is his match in state S (Note: $M_S(x) = *$ is possible).

Observation:

A state S in the game is a pure Nash equilibrium

\Leftrightarrow The matching M_S is a stable matching.

Representation with Costs



In this instance the preference orders can be represented by correlated cost values.

Correlated Preferences

An intuitive case of matching is when both players receive the same cost from a match. Then each match has a single positive edge cost, and this cost is assigned to both players if they match along this edge. This is referred to as **correlated** or **weighted matching**.

In a **correlated matching game**, we have $c_x(y) = c_y(x) = c(x, y)$ for all $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, and $c_x(*) = c_y(*) = \infty$.

Preferences are now correlated among agents – the smaller the edge cost, the better the match **for both partners**.

Correlated Matching: Ordinal Potential Game

Theorem

Every correlated matching game is an ordinal potential game. If all edge costs are pairwise distinct, the pure Nash equilibrium is unique.


Proof:

For a state S , define the following function

$$\Phi(S) = (c_{x_1}(S), \dots, c_{x_n}(S)),$$

where the men are sorted in non-decreasing order of cost, i.e., for $i \leq j$ it holds $c_{x_i}(S) \leq c_{x_j}(S)$. This is a **lexicographic potential**.


Exmample



Consider an improvement sequence from state S^0 . The sorted vector of player costs decreases lexicographically in every step.

State	Sorted Costs
S^0	18, 20, 23


Exmample



Consider an improvement sequence from state S^0 . The sorted vector of player costs decreases lexicographically in every step.

State	Sorted Costs
S^0	18, 20, 23
S^1	14, 20, ∞


Exmample



Consider an improvement sequence from state S^0 . The sorted vector of player costs decreases lexicographically in every step.

State	Sorted Costs
S^0	18, 20, 23
S^1	14, 20, ∞
S^2	12, 18, ∞


Exmple



Consider an improvement sequence from state S^0 . The sorted vector of player costs decreases lexicographically in every step.

State	Sorted Costs
S^0	18, 20, 23
S^1	14, 20, ∞
S^2	12, 18, ∞
S^3	12, 18, 22


Exmample



Consider an improvement sequence from state S^0 . The sorted vector of player costs decreases lexicographically in every step.

State	Sorted Costs
S^0	18, 20, 23
S^1	14, 20, ∞
S^2	12, 18, ∞
S^3	12, 18, 22
S^4	12, 16, ∞

Exmample



Consider an improvement sequence from state S^0 . The sorted vector of player costs decreases lexicographically in every step.

State	Sorted Costs
S^0	18, 20, 23
S^1	14, 20, ∞
S^2	12, 18, ∞
S^3	12, 18, 22
S^4	12, 16, ∞
S^5	12, 16, 28
	pure NE

Correlated Matching: Ordinal Potential Game

Consider man x that deviates from strategy y to y' and improves strictly. Let S and S' be the resulting states. Since x strictly improves, he is matched to $y' \neq y$ in S' . Depending on x and y' being matched in S , we obtain four cases (see below).

In every case the smallest cost value added is $c(x, y')$, and it is strictly smaller than the smallest cost value removed. Hence, the sorted vector of costs decreases lexicographically.

1. $\{x, y\} \in M_S$, y' unmatched in S :
Then $c_x(S) = c(x, y) > c(x, y') = c_x(S')$. In the sorted vector of costs $c(x, y)$ is replaced by $c(x, y')$. If y has another proposal in S beyond x , then she has a proposal in S' . In this case, for the best man $x'' \in A_y(S')$ in the sorted vector we replace ∞ by $c(x'', y)$.

Correlated Matching: Ordinal Potential Game

2. $\{x, y\} \notin M_S$, y' unmatched in S :
Then $c_x(S) = \infty > c(x, y') = c_x(S')$. In the sorted vector of costs ∞ is replaced by $c(x, y')$.

3. $\{x, y\} \in M_S$, $\exists x' \in \mathcal{X}$ with $\{x', y'\} \in M_S$:
Then $c_x(S) = c(x, y) > c(x, y') = c_x(S')$. Since $\{x, y'\} \in M_{S'}$ it holds $c(x', y') > c(x, y')$. In the sorted vector of costs $c(x, y)$ is replaced by $c(x, y')$ and $c(x', y')$ is replaced by ∞ . If y has another proposal in S beyond x , then she has a proposal in S' . In this case, for the best man $x'' \in A_y(S')$ in the sorted vector we replace ∞ by $c(x'', y)$.

4. $\{x, y\} \notin M_S$, $\exists x' \in \mathcal{X}$ with $\{x', y'\} \in M_S$:
Then $c_x(S) = \infty > c(x, y') = c_x(S')$. Since $\{x, y'\} \in M_{S'}$ it holds $c(x', y') > c(x, y')$. In the sorted vector of costs ∞ is replaced by $c(x, y')$ and $c(x', y')$ is replaced by ∞ .

Uniqueness

Now suppose all values $c(x, y)$ are pairwise distinct. Consider pair $\{x_0, y_0\}$ with smallest edge cost. In every state S woman y_0 is the unique best response for player x_0 . The match yields smallest cost, and y_0 will always accept the proposal. Hence, x_0 must play y_0 in every pure Nash equilibrium.

Among the remaining possible pairs again consider the remaining one $\{x_1, y_1\}$ with smallest cost. Conditioned on x_0 and y_0 being matched, y_1 is the unique best response for player x_1 . The match yields smallest cost (conditioned on $\{x_0, y_0\}$ being matched), and y_1 will accept the proposal. Hence, given that x_0 plays y_0 , we know that x_1 must play y_1 in every pure Nash equilibrium.

Applied inductively the argument shows uniqueness of the pure Nash equilibrium. □

Improvement Sequences

The last proof defines an algorithm to build a short improvement sequence: In every step let the player deviate that obtains the smallest cost by deviating. Then x_0 first deviates to y_0 , then x_1 to y_1 etc. (unless they are already matched, resp.)

The greedy algorithm applies this strategy to the empty matching and tries to insert edges in the order of non-decreasing cost. In this way, the greedy algorithm computes a pure Nash equilibrium in polynomial time.

Corollary

1. *For every state S there is an improvement sequence that reaches a pure Nash equilibrium in at most n steps.*
2. *A pure Nash equilibrium can be computed with a greedy algorithm in time $O(n^2 \log n)$.*

The corollary continues to hold even when the values $c(x, y)$ are not pairwise distinct. (Why?)

When preferences are not correlated...

Consider the general class of matching games, where player costs are not correlated via a single cost value per edge. For simplicity we now assume that for every player the **cost is the position of the partner in the preference order**.

Let $\succ_x = (y_1, \dots, y_n)$ be the preference order of man $x \in \mathcal{X}$. Then $c_x(y_k) = k$ for $k \in \{1, \dots, n\}$ and $c_x(*) = n + 1$. The costs $c_y(x)$ are given similarly. In state S we obtain the matching M_S as described before and the resulting costs $c_x(S) = c_x(M_S(x))$.


When preferences of the agents are not correlated, then there exist matching games without the finite improvement property. Even when all players must play **best responses**, there can be **cyclic improvement sequences**. Hence, in particular, there are matching games without an ordinal potential function.

Potential Game

Proposition

There are matching games without the finite improvement property.

Consider the following matching game with a cyclic sequence of (best-response) improvement steps.




Potential Game

Proposition

There are matching games without the finite improvement property.

Consider the following matching game with a cyclic sequence of (best-response) improvement steps.




Potential Game

Proposition

There are matching games without the finite improvement property.

Consider the following matching game with a cyclic sequence of (best-response) improvement steps.




Potential Game

Proposition

There are matching games without the finite improvement property.

Consider the following matching game with a cyclic sequence of (best-response) improvement steps.




Potential Game

Proposition

There are matching games without the finite improvement property.

Consider the following matching game with a cyclic sequence of (best-response) improvement steps.




Potential Game

Proposition

There are matching games without the finite improvement property.

Consider the following matching game with a cyclic sequence of (best-response) improvement steps.




Potential Game

Proposition

There are matching games without the finite improvement property.

Consider the following matching game with a cyclic sequence of (best-response) improvement steps.



Existence and Computation

While there are cyclic improvement sequences, we always have a pure Nash equilibrium (i.e., a stable matching)

Algorithm 1: Deferred Acceptance (DA) Algorithm with Man-Proposal

Initialize $\succ'_x = \succ_x$ for all $x \in \mathcal{X}$

while there is an unmatched man $x \in \mathcal{X}$ with $\succ'_x \neq \emptyset$ **do**

 Every man $x \in \mathcal{X}$ proposes to topmost woman in \succ'_x

 Every woman $y \in \mathcal{Y}$ keeps most preferred man from $A_y(S)$

y rejects all other men from $A_y(S)$

 If his current proposal is rejected, man x removes topmost entry from

\succ'_x

Theorem (Gale, Shapley 1962)

A stable matching always exists and can be computed in time $O(nm)$.

Convergence

Proof:

The DA algorithm can be implemented to run in time $O(nm)$. It computes a matching M , as each man proposes to at most one woman at a time and each woman keeps at most one proposal.

It is straightforward to verify that over the run of the algorithm

- ▶ for a man, the preference of proposed women is strictly decreasing, and
- ▶ for a woman, the preference of matched partners is strictly increasing.

Assume for contradiction M has a blocking pair $\{x, y\}$ with $y \succ_x M(x)$ and $x \succ_y M(y)$. x must have proposed to y and got rejected, so y must keep a proposal of some better man $x' \succ_y x$. Hence, her match in M can only be better than x' . Thus, $M(y) \succeq_y x' \succ_y x$, a contradiction. \square

A reformulation of this idea implies an even stronger property in matching games.

Convergence

Theorem

For every matching game and every initial state S_0 , there is a sequence of $2nm$ best-response improvement steps to a pure Nash equilibrium.

Proof:

The sequence has two phases.

In **Phase 1**, only **matched men** are allowed to play best responses. Let X be the set of matched men in M_S . The following function keeps decreasing over phase 1:

$$\Phi(S) = \sum_{x \in X} c_x(S) + \sum_{x \in \mathcal{X} \setminus X} c_x(S_x). \quad (\text{rank of } x\text{'s partner in } \succ_x)$$

$\Phi(S)$ sums for each matched man the rank of his partner, and for each unmatched man the rank of the “virtual” partner, i.e., the rank if he would be matched to his choice S_x .

Convergence

Suppose $x \in X$ deviates from woman $y = S_x$ to a best response y' .

- ▶ x remains matched, improves rank of partner by at least 1.
- ▶ If y is matched to some $x' \in X$, then x' becomes unmatched. Cost $c_{x'}(y')$ moves from the first sum to the second sum. Value of $\Phi(S)$ does not change because of this.
- ▶ If y has other offers (i.e., other men pick y as strategy), then y remains matched after x deviates. Hence, some $x'' \in \mathcal{X} \setminus X$ gets matched to y . Cost term $c_{x''}(y)$ moves from the second sum to first sum. Value of $\Phi(S)$ does not change because of this.

Thus, Φ drops by at least 1 in every iteration. As $1 \leq \Phi(S) \leq nm$, phase 1 terminates after at most nm iterations.

Convergence

In **Phase 2**, only **unmatched men** are allowed to play best responses. Denote by Y the set of matched women in M_S . The following function keeps increasing over phase 2:

$$\Psi(S) = \sum_{y \in Y} (n + 1 - c_y(S)).$$

Suppose an unmatched man x deviates to a best response.

- ▶ x gets matched to $y \in Y$, c_y decreases by at least 1.
- ▶ x gets matched to $y \notin Y$, y enters Y .

Thus, Ψ grows by at least 1 in every iteration. As $1 \leq \Psi(S) \leq nm$, phase 2 terminates after at most nm iterations.

Why is the final state a stable matching? Observe that throughout phase 2 **no matched man can improve**. When unmatched x gets matched to y , this only decreases y 's cost. Assuming that there was no blocking pair with any of the matched men before, there is no blocking pair after x and y are matched – x played a best response and y 's cost is even lower now. Finally, there are no improvements of unmatched men (because phase 2 is finished). □

Weakly Acyclic

We proved that matching games are weakly acyclic.

Definition (Weakly Acyclic Game)

A strategic game $\Gamma = (\mathcal{N}, (\Sigma_i)_{i \in \mathcal{N}}, (c_i)_{i \in \mathcal{N}})$ is **weakly acyclic** if for every state S there is **at least one finite improvement sequence** that leads from S to a pure Nash equilibrium.

Consider **random better-response dynamics**: In every step a player $i \in \mathcal{N}$ is chosen uniformly at random. He chooses a strategy $s'_i \in \Sigma_i$ uniformly at random and deviates if s'_i represents a strict improvement.

Random better-response dynamics emerge from a Markov chain (more concretely: a random walk) over the states of the game. Pure Nash equilibria are the absorbing states.

Random Dynamics – Convergence in the Limit

In weakly acyclic games **random better-response dynamics** converge **with probability 1 in the limit** to a pure Nash equilibrium. A simple consequence from pure equilibria being absorbing states – if the dynamics run long enough, they eventually implement a sequence that leads to (and then remains in) a pure Nash equilibrium.

How **long does it take** until we reach a pure Nash equilibrium? What happens when we also restrict to random **best-response** dynamics? Unfortunately, there are games and starting states such that the convergence time to a pure Nash equilibrium is exponential with high probability.

Theorem (Ackermann, Goldberg, Mirrokni, Röglin, Vöcking 2011)

There is a matching game with n men and n women and an initial state S_0 such that, with probability $1 - 2^{-\Omega(n)}$, random dynamics starting from S_0 need $2^{\Omega(n)}$ steps to reach a stable matching.

This result holds for both random better- and best-response dynamics.

Literature

- ▶ D. Monderer, L. Shapley. Potential Games. Games and Economic Behavior 14:1124–1143, 1996.
- ▶ A. Fabrikant, C. Papadimitriou, K. Talwar. The Complexity of Pure Nash Equilibria. STOC 2004.
- ▶ S. Jeong, R. McGrew, E. Nudelman, Y. Shoham, Q. Sun. Fast and Compact: A Simple Class of Congestion Games. AAAI 2005. (Singleton games)
- ▶ H. Ackermann, H. Röglin, B. Vöcking. On the Impact of Combinatorial Structure on Congestion Games. Journal of the ACM 55(6), 2008. (Matroid Games, PLS)
- ▶ S. Chien, A. Sinclair. Convergence to Approximate Nash Equilibria in Congestion Games. Games and Economic Behavior 71(2):315–327, 2011.

Literature

- ▶ A. Roth, J. Vande Vate. Random Paths to Stability in Two-Sided Matching. *Econometrica* 58(6):1475-1480, 1990. (Weak Acyclicity)
- ▶ D. Abraham, A. Levavi, D. Manlove, G. O'Malley. The Stable Roommates Problem with Globally Ranked Pairs. *Internet Math.* 5(4):493-515, 2008. (Ordinal Potential)
- ▶ F. Mathieu. Self-Stabilization in Preference-based Systems. *Peer-to-Peer Netw. Appl.* 1(2):104-121, 2008. (Ordinal Potential)
- ▶ H. Ackermann, P. Goldberg, V. Mirrokni, H. Röglin, B. Vöcking. Uncoordinated Two-Sided Matching Markets. *SIAM J. Comput.* 40(1):92-106, 2011.