Convergence in competitive games

Vahab S. Mirrokni

Computer Science and AI Lab. (CSAIL) and Math. Dept., MIT.

This talk is based on joint works with A. Vetta and with A. Sidiropoulos, A. Vetta

Cut game

- Cut game:
  - Players: Nodes of the graph.
  - Player's strategy $\in \{1, -1\}$ (Republican or Democrat)
  - An action profile corresponds to a cut.
  - Payoff: Total Contribution in the cut.
  - Change Party if you gain.

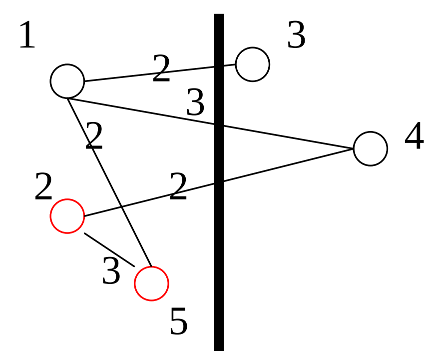

Cut Value: 7
2 and 5 are unhappy.

The Cut Game: Price of Anarchy

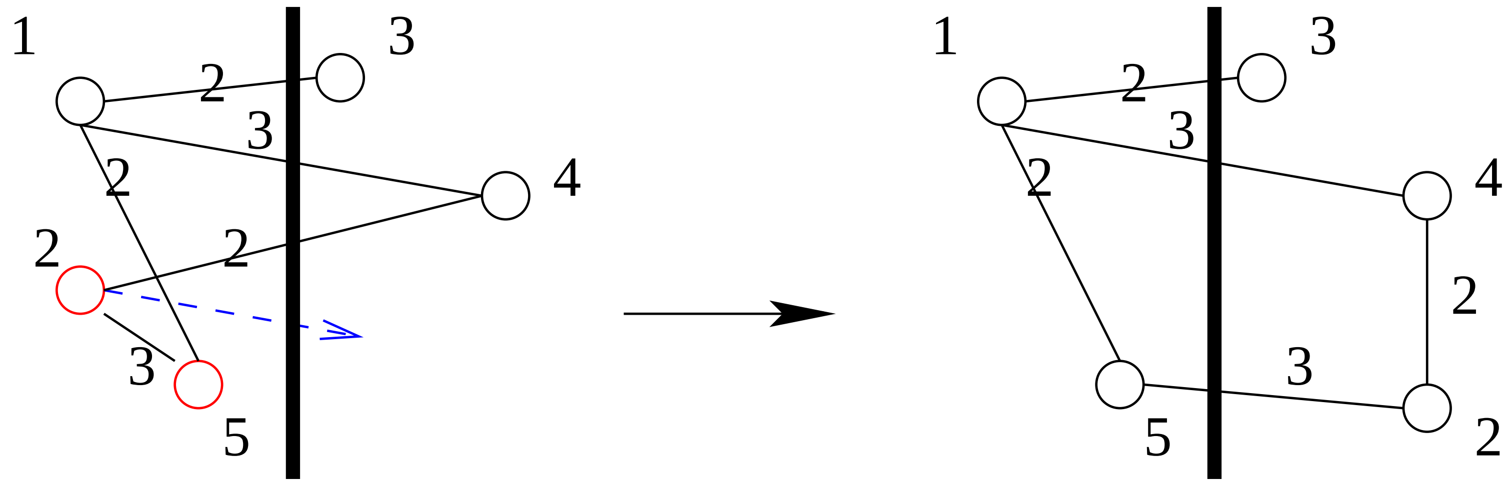

Cut Value: 7
2 and 5 are unhappy.

Cut Value: 8
Pure Nash Equilibrium.

The Cut Game: Price of Anarchy

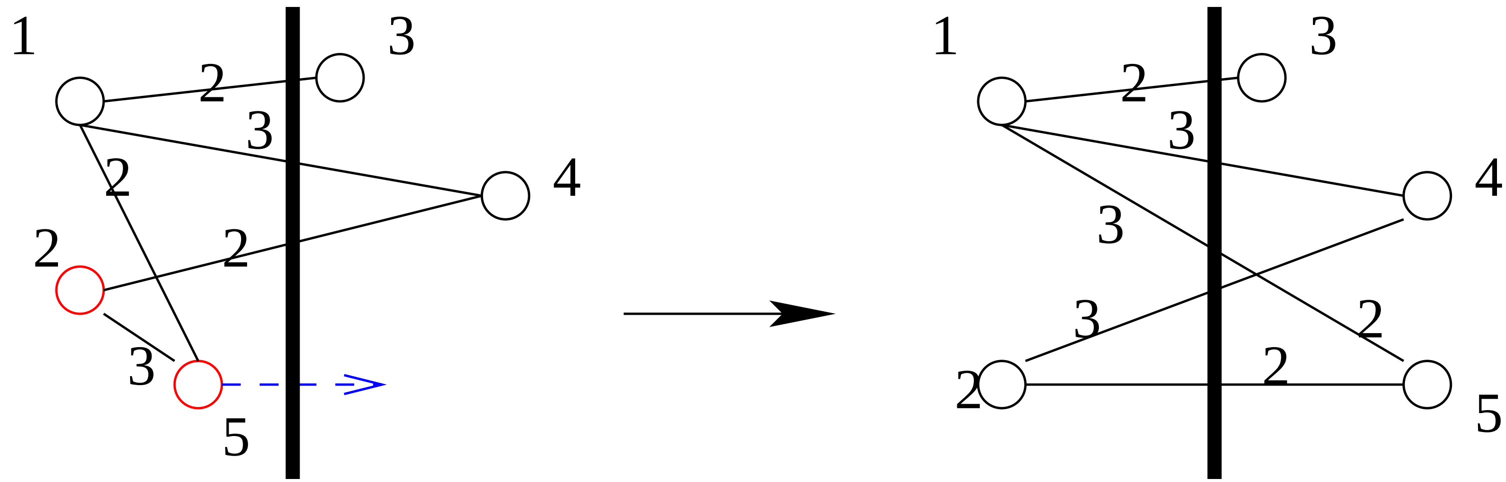

Cut Value: 7
2 and 5 are unhappy.

Cut Value: 12
The Optimum.

- Social Function:
  - The cut value.

Price of Anarchy for this instance: $\frac{12}{8} = 1.5$.

Outline

- Performance in lack of Coordination: Price of Anarchy.
- Best-Responses, Convergence, and Random Paths.
- A Potential Game: Cut Game
  - Lower Bounds: Long poor paths
  - Upper Bounds: random paths
- Basic-utility and Valid-utility Games
  - Basic-utility Games: Fast Convergence.
  - Valid-utility Games: Poor Sink Equilibria
- Conclusion: Other Games?

Convergence to Approximate Solutions

We can model selfish behavior of players by a sequence of best responses by players.

Convergence to Approximate Solutions

We can model selfish behavior of players by a sequence of best responses by players.

How fast do players converge to a Nash equilibrium?

Convergence to Approximate Solutions

We can model selfish behavior of players by a sequence of best responses by players.

How fast do players converge to a Nash equilibrium?

How fast do players converge to an approximate solution?

Convergence to Approximate Solutions

We can model selfish behavior of players by a sequence of best responses by players.

How fast do players converge to a Nash equilibrium?

How fast do players converge to an approximate solution?

Our goal: How fast do players converge to an approximate solution?

Fair Paths

In a fair path, we should let each player play at least once after each polynomially many steps.

Fair Paths

In a fair path, we should let each player play at least once after each polynomially many steps.

- One-round path: We let each player play once in a round.
- random path: We pick the next player at random.

Fair Paths

In a fair path, we should let each player play at least once after each polynomially many steps.

- One-round path: We let each player play once in a round.
- random path: We pick the next player at random.

We are interested in the Social Value at the end of a fair path.

A Cut game: The Party Affiliation Game

- Cut game:

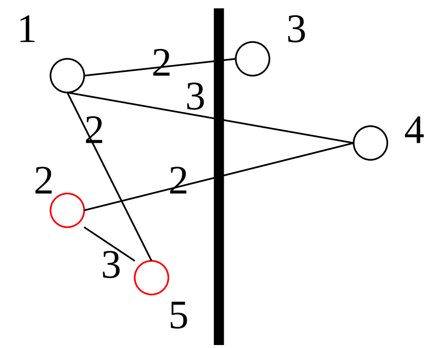

Cut Value: 7
2 and 5 are unhappy.

- Social Function:
  - The Cut Value
  - Total Happiness
- Price of anarchy: at most 2.
- Local search algorithm for Max-Cut!

A Cut game: The Party Affiliation Game

- Cut game:

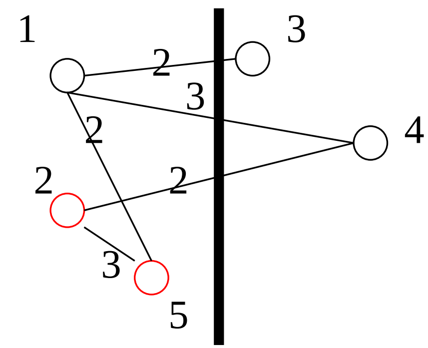

Cut Value: 7
2 and 5 are unhappy.

- Social Function:
  - The Cut Value
- Convergence:
  - Finding local optimum for Max-Cut is PLS-complete (Schaffer, Yannakakis [1991]).

Cut Game: Paths to Nash equilibria

- Unweighted graphs After $O(n^2)$ steps, we converge to a Nash equilibrium.
- Weighted graphs: It is PLS-complete.
  - PLS-Complete problems and tight PLS reduction (Johnson, Papadimitriou, Yannakakis [1988]).
  - Tight PLS reduction from Max-Cut (Schaffer, Yannakakis [1991])
  - There are some states that are exponentially far from any Nash equilibrium.

Question: Are there long poor fair paths?

Cut Game: A Bad Example

- Consider graph G, a line of n vertices. The weight of edges are $1, 1+\frac{1}{n}, 1+\frac{2}{n}, \ldots, 1+\frac{n-1}{n}$. Vertices are labelled $1, \ldots, n$ throughout the line. Consider the round of best responses:

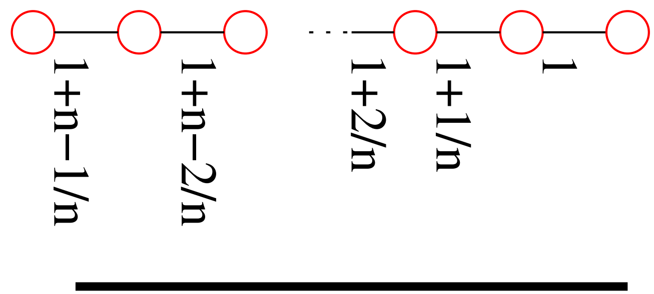

A Bad Example: Illustration

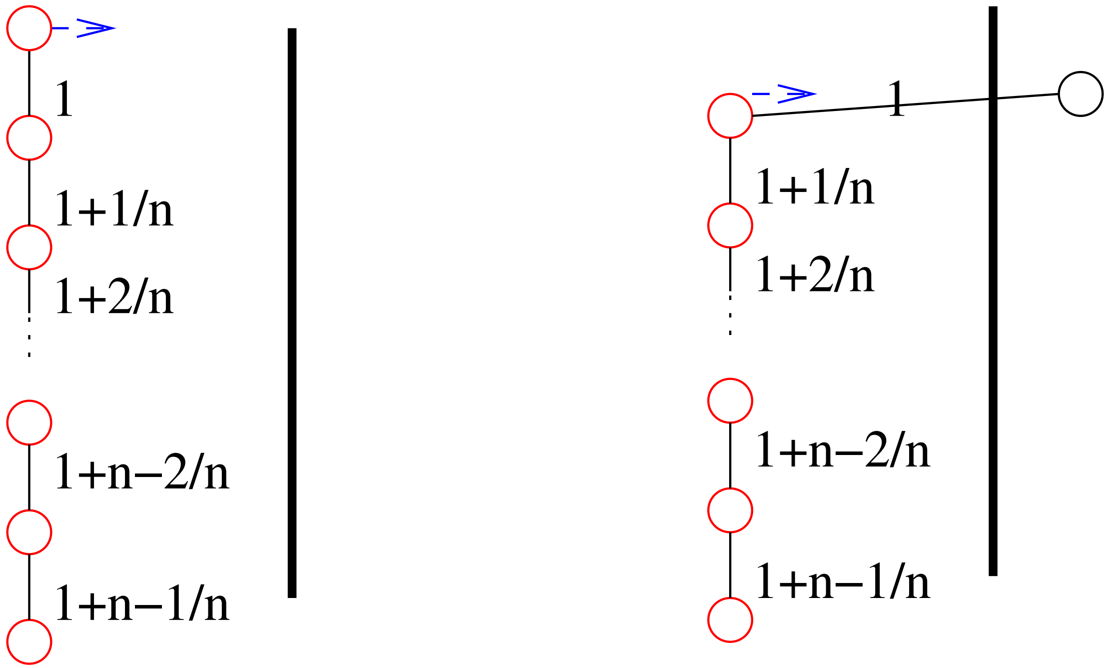

After one move.

A Bad Example: Illustration

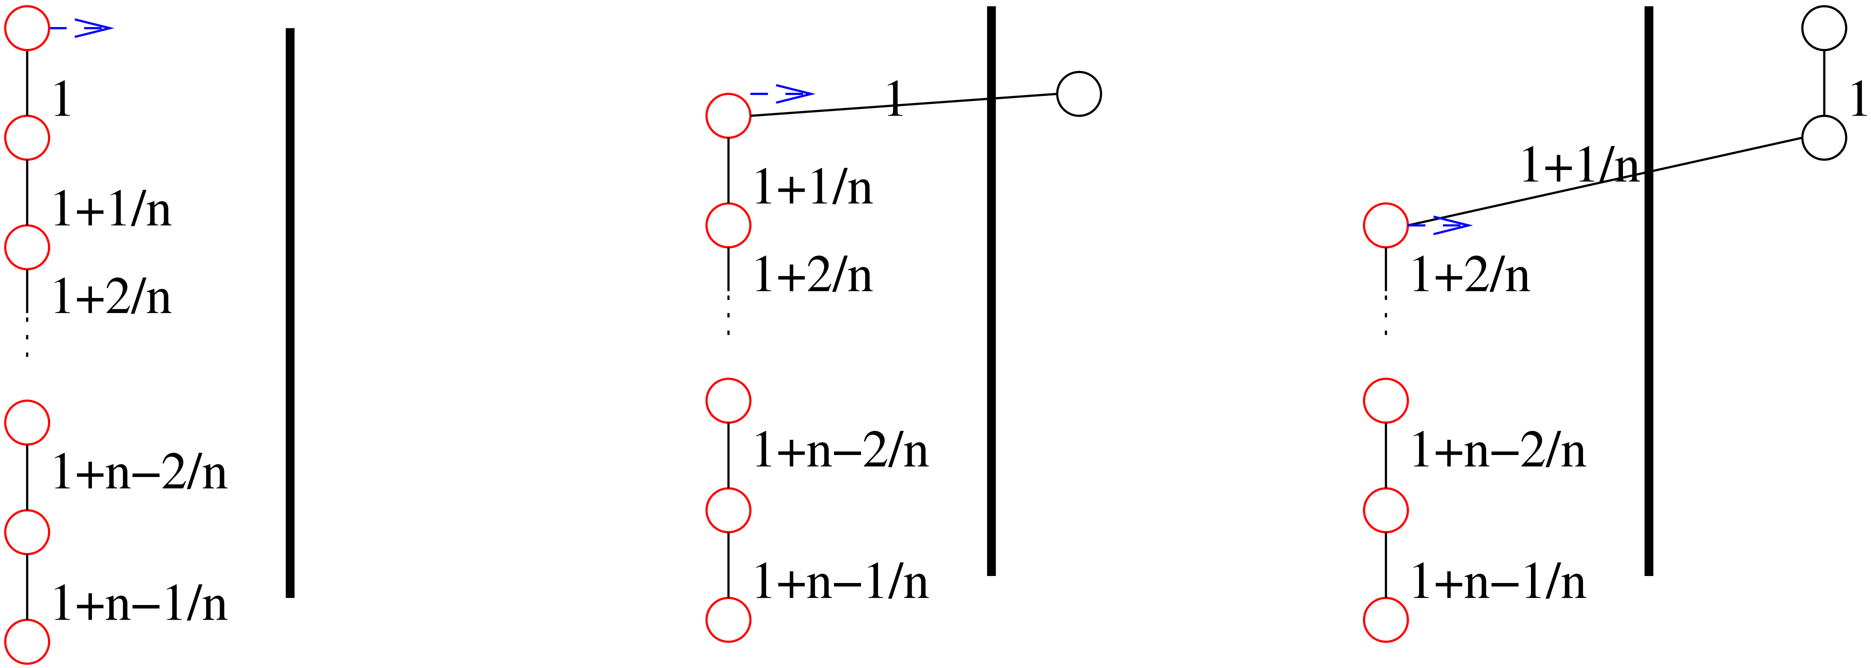

After two moves.

A Bad Example: Illustration

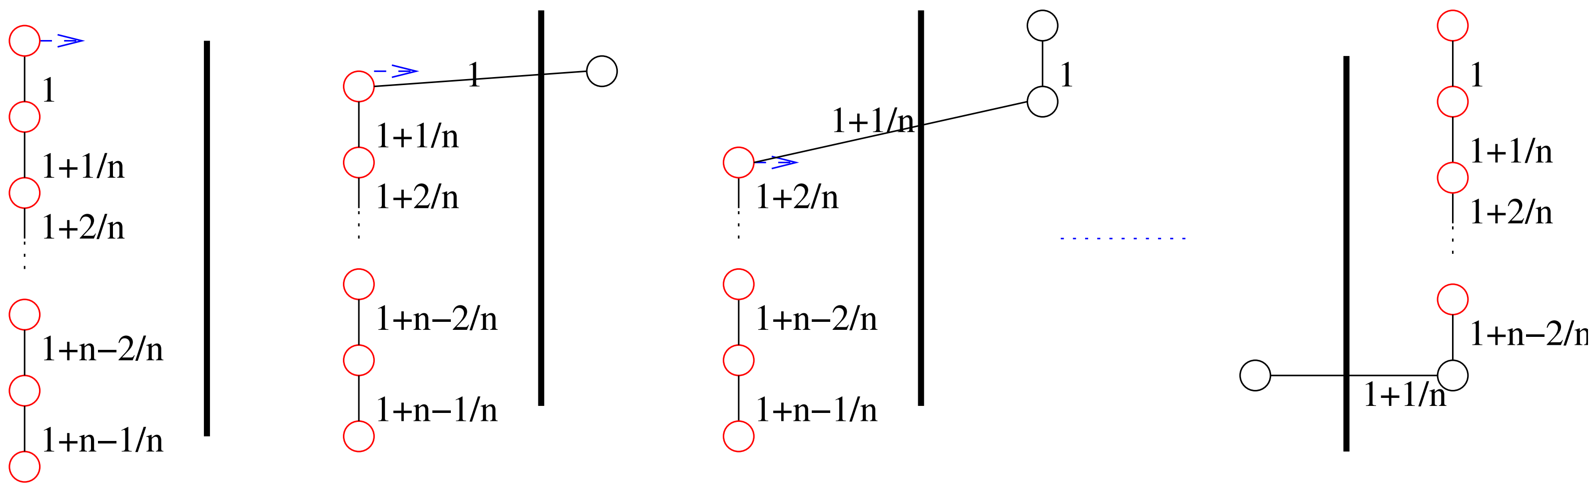

After n moves (one round)

A Bad Example: Illustration

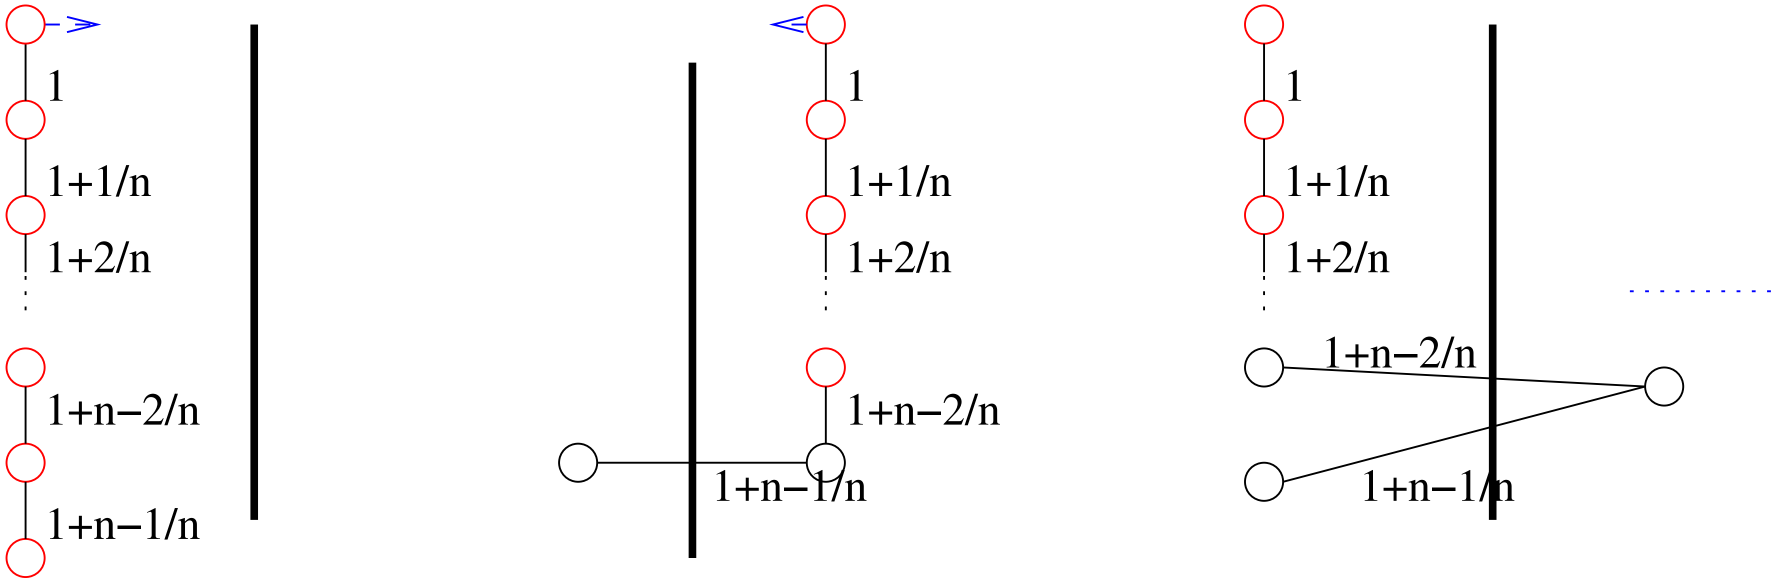

After two rounds.

- Theorem: In the above example, the cut value after k rounds is $O(\frac{k}{n})$ of the optimum.

Random One-round paths

- Theorem:(M., Sidiropoulos[2004]) The expected value of the cut after a random one-round path is at most $\frac{1}{8}$ of the optimum.

Random One-round paths

- Theorem:(M., Sidiropoulos[2004]) The expected value of the cut after a random one-round path is at most $\frac{1}{8}$ of the optimum.
- Proof Sketch: The sum of payoffs of nodes after their moves is $\frac{1}{2}$-approximation. In a random ordering, with a constant probability a node occurs after $\frac{3}{4}$ of its neighbors. The expected contribution of a node in the cut is a constant-factor of its total weight.

Exponentially Long Poor Paths

- Theorem: (M., Sidiropoulos[2004]) There exists a weighted graph $G = (V(G), E(G))$, with $|V(G)| = \Theta(n)$, and exponentially long fair path such that the value of the cut at the end of $\mathcal{P}$, is at most $O(1/n)$ of the optimum cut.

Exponentially Long Poor Paths

- Theorem: (M., Sidiropoulos[2004]) There exists a weighted graph $G = (V(G), E(G))$, with $|V(G)| = \Theta(n)$, and exponentially long fair path such that the value of the cut at the end of $\mathcal{P}$, is at most $O(1/n)$ of the optimum cut.
- Proof Sketch:
  Use the example for the exponentially long paths to the Nash equilibrium in the cut game. Find a player, v, that moves exponentially many times. Add a line of n vertices to this graph and connect all the vertices to player v.

Poor Long Path: Illustration

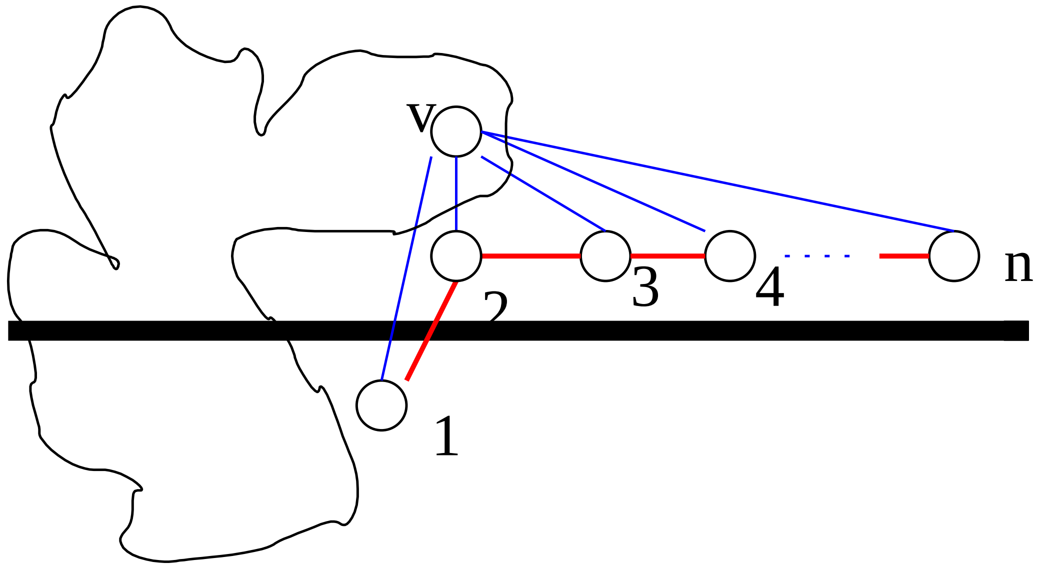

Poor Long Path: Illustration

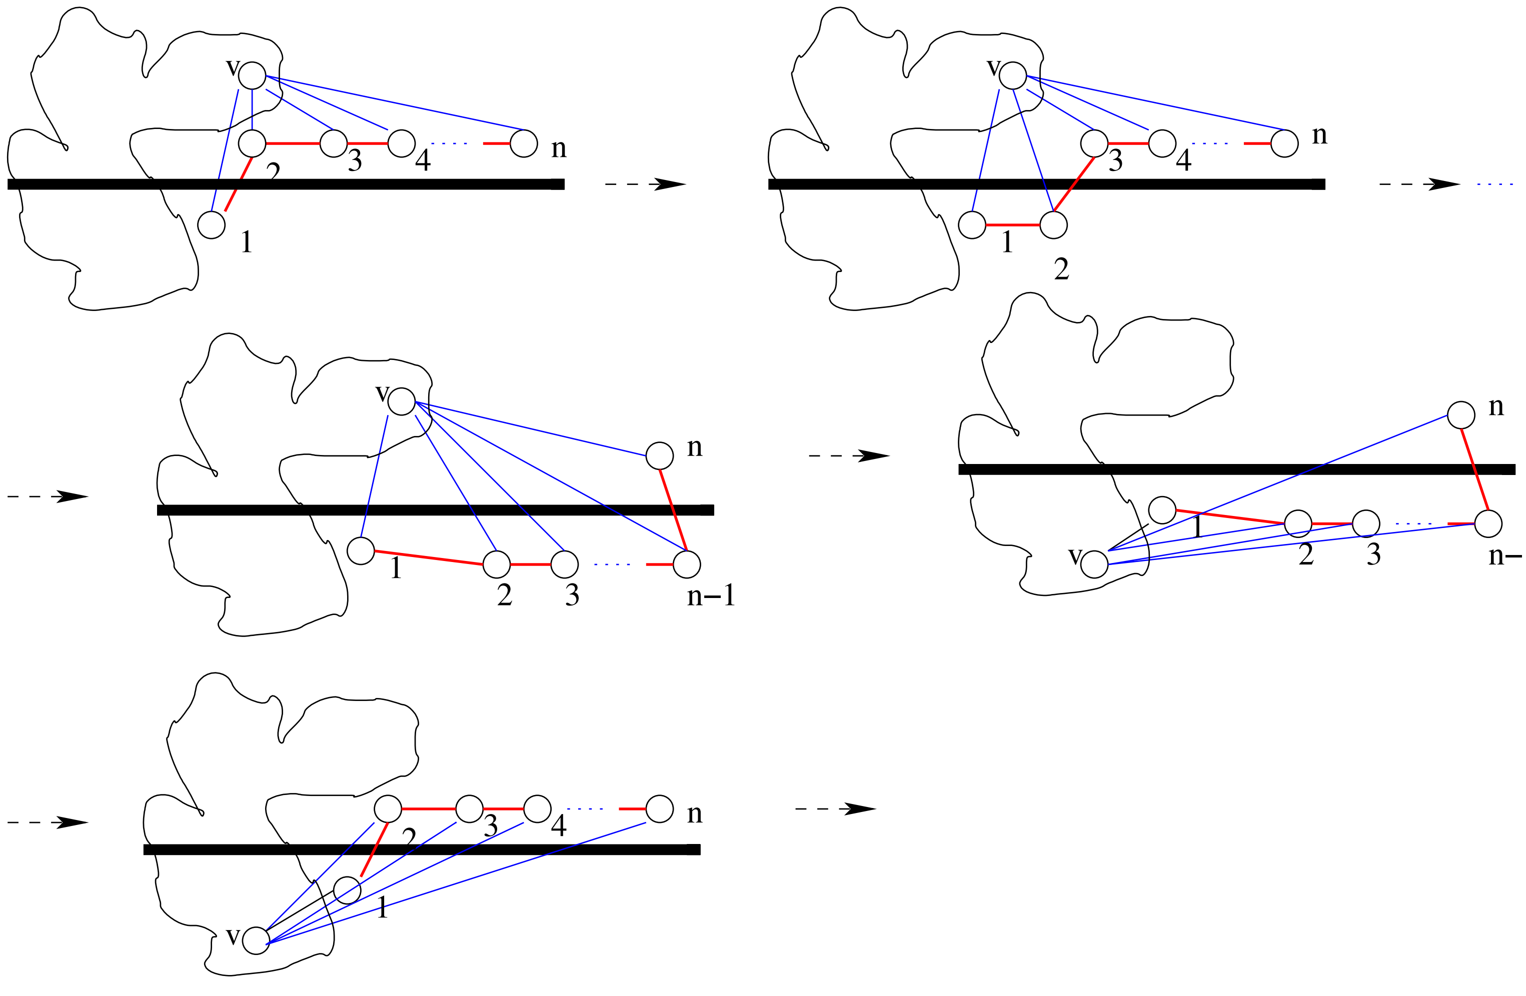

Mildly Greedy Players

A Player is 2-greedy, if she does not move if she cannot double her payoff.

Mildly Greedy Players

- A Player is 2-greedy, if she does not move if she cannot double her payoff.
- Theorem:(M., Sidiropoulos[2004]) One round of selfish behavior of 2-greedy players converges to a constant-factor cut.
- Proof Idea: If a player moves it improves the value of the cut by a constant factor of its contribution in the cut.

Mildly Greedy Players

- A Player is 2-greedy, if she does not move if she cannot double her payoff.
- Theorem:(M., Sidiropoulos[2004]) One round of selfish behavior of 2-greedy players converges to a constant-factor cut.
- Proof Idea: If a player moves it improves the value of the cut by a constant factor of its contribution in the cut.
- Message: Mildly Greedy Players converge faster.

Mildly Greedy Players

- A Player is 2-greedy, if she does not move if she cannot double her payoff.
- Theorem:(M., Sidiropoulos[2004]) One round of selfish behavior of 2-greedy players converges to a constant-factor cut.
- Proof Idea: If a player moves it improves the value of the cut by a constant factor of its contribution in the cut.
- Message: Mildly Greedy Players converge faster.

A Cut game: Total Happiness

- Cut game:
  - The happiness of player v is equal to his total contribution in the cut minus the weight of its adjacent edges not in the cut.
- Social Function:
  - Total Happiness: Sum of happiness of players

A Cut game: Total Happiness

- Cut game:
  - The happiness of player v is equal to his total contribution in the cut minus the weight of its adjacent edges not in the cut.
- Social Function:
  - Total Happiness: Sum of happiness of players
- In the context of correlation clustering: Maximizing agreement minus disagreement (Bansal, Blum, Chawla[2002]).
- $\log n$-approximation algorithm is known. (Charikar, Wirth[2004]).

A Cut game: Total Happiness

- Cut game:
  - The happiness of player v is equal to his total contribution in the cut minus the weight of its adjacent edges not in the cut.
- Social Function:
  - Total Happiness: Sum of happiness of players
- Price of anarchy: unbounded in the worst case.
- A bad example: a cycle of size four.

A Cut game: Total Happiness

- Cut game:
  - The happiness of player v is equal to his total contribution in the cut minus the weight of its adjacent edges not in the cut.
- Social Function:
  - Total Happiness: Sum of happiness of players
- The expected happiness of a random cut is zero.
- Our result: For unweighted graphs of large girth, if we start from a random cut, then after a random one-round path, the expected happiness is a sublogarthmic-approximation.

Cut Game: Total Happiness

For some $\delta > 0$, we call an edge of G, δ-good, if at least one of its end-points, has degree at most δ.

Cut Game: Total Happiness

For some $\delta > 0$, we call an edge of G, δ-good, if at least one of its end-points, has degree at most δ.

For a pair $u, v \in V(G)$, let $\mathcal{E}_{u,v}$ denote the event that there exists a path $p = x_1, x_2, \ldots, x_{|p|}$, with $u = x_1$, and $v = x_{|p|}$, and for any i, with $1 \leq i < |p|$, $x_i \prec x_{i+1}$.

Cut Game: Total Happiness

For some $\delta > 0$, we call an edge of G, δ-good, if at least one of its end-points, has degree at most δ.

For a pair $u, v \in V(G)$, let $\mathcal{E}_{u,v}$ denote the event that there exists a path $p = x_1, x_2, \ldots, x_{|p|}$, with $u = x_1$, and $v = x_{|p|}$, and for any i, with $1 \leq i < |p|$, $x_i \prec x_{i+1}$.

- Lemma: Let $\{u, v\}, \{v, w\} \in E(G)$, such that $u \prec w \prec v$. There exists a constant C, such that if the girth of G is at least $C \frac{\log n}{\log \log n}$, then $\mathbf{Pr}[\mathcal{E}_{u,w}] < n^{-3}$.

Cut Game: Total Happiness

For some $\delta > 0$, we call an edge of G, δ-good, if at least one of its end-points, has degree at most δ.

For a pair $u, v \in V(G)$, let $\mathcal{E}_{u,v}$ denote the event that there exists a path $p = x_1, x_2, \ldots, x_{|p|}$, with $u = x_1$, and $v = x_{|p|}$, and for any i, with $1 \leq i < |p|$, $x_i \prec x_{i+1}$.

- Lemma: Let $\{u, v\}, \{v, w\} \in E(G)$, such that $u \prec w \prec v$. There exists a constant C, such that if the girth of G is at least $C \frac{\log n}{\log \log n}$, then $\mathbf{Pr}[\mathcal{E}_{u,w}] < n^{-3}$.
- Lemma: For any $e \in E(G)$, we have $\mathbf{Pr}[e \text{ is cut }] \geq 1/2 - o(1)$.

Cut Game: Total Happiness

- Lemma: Let $e = \{u, v\} \in E(G)$, with $u \prec v$, and $\deg(v) \leq \delta$. Then, $\mathbf{Pr}[e \text{ is cut }] \geq 1/2 + \Omega(1/\sqrt{\delta})$.

Cut Game: Total Happiness

- Lemma: Let $e = \{u, v\} \in E(G)$, with $u \prec v$, and $\deg(v) \leq \delta$. Then, $\mathbf{Pr}[e$ is cut $] \geq 1/2 + \Omega(1/\sqrt{\delta})$.
- Theorem: (M., Sidiropoulos[2004]) There exists a constant C', such that for any $C > C'$, and for any unweighted simple graph of girth at least $C\frac{\log n}{\log\log n}$, if we start from a random cut, the expected value of the happiness at the end of a random one-round path, is within a $\frac{1}{(\log n)^{O(1/C)}}$ factor from the maximum happiness.

Outline

- Performance in lack of Coordination: Price of Anarchy.
- State Graphs, Convergence, and Fair Paths.
- Cut Games: Party Affiliation Games
  - Lower Bounds: Long poor paths
  - Upper Bounds: random paths
  - Total Happiness: Cut minus Other Edges
- Basic-utility and Valid-utility Games.
  - Basic-utility Games: Fast Convergence.
  - Valid-utility Games: Poor Sink Equilibria!
- Conclusion: Other Games?

Valid-Utility Games

- Ground Set of Markets: $V = \{v_1, v_2, \ldots, v_n\}$.
- Player i can provide a subset of V. $\mathcal{S}_i$ is a family of subsets of V feasible for player i.
- $S_i \subset V$ is the strategy of player i. $S_i \in \mathcal{S}_i$.

Valid-Utility Games

- Ground Set of Markets: $V = \{v_1, v_2, \ldots, v_n\}$.
- Player i can provide a subset of V. $\mathcal{S}_i$ is a family of subsets of V feasible for player i.
- $S_i \subset V$ is the strategy of player i. $S_i \in \mathcal{S}_i$.
- Social Function:
  - A submodular set function $f : 2^V \rightarrow R$ on union of strategies: $f(\cup_{1 \le i \le n} S_i)$.

Valid-Utility Games

- Ground Set of Markets: $V = \{v_1, v_2, \ldots, v_n\}$.
- Player i can provide a subset of V. $\mathcal{S}_i$ is a family of subsets of V feasible for player i.
- $S_i \subset V$ is the strategy of player i. $S_i \in \mathcal{S}_i$.
- Social Function:
  - A submodular set function $f : 2^V \rightarrow R$ on union of strategies: $f(\cup_{1 \leq i \leq n} S_i)$.
- The payoff of any player is at least the change that he makes in the social function by playing.
- The sum of payoffs is at most the social function.
- Several examples, including the market sharing game and a facility location game

Valid-Utility Games

- Ground Set of Markets: $V = \{v_1, v_2, \ldots, v_n\}$.
- Player i can provide a subset of V. $\mathcal{S}_i$ is a family of subsets of V feasible for player i.
- $S_i \subset V$ is the strategy of player i. $S_i \in \mathcal{S}_i$.
- Social Function:
  - A submodular set function $f : 2^V \rightarrow R$ on union of strategies: $f(\cup_{1 \leq i \leq n} S_i)$.
- The payoff of any player is at least the change that he makes in the social function by playing.
- The sum of payoffs is at most the social function.
- In basic-utility games, the payoff is equal to the change that a player makes.

Example: Market Sharing Game

- Market Sharing Game
  - n markets and m players.
  - Market i has a value q_i and cost C_i.
  - Player j has a budget B_j.
  - Player j's action is to choose a subset of markets of his interest whose total cost is at most B_j.
  - The value of a market is divided equally between players that provide these markets.

Example: Market Sharing Game

- Market Sharing Game
  - n markets and m players.
  - Market i has a value q_i and cost C_i.
  - Player j has a budget B_j.
  - Player j's action is to choose a subset of markets of his interest whose total cost is at most B_j.
  - The value of a market is divided equally between players that provide these markets.

  Social Function: Total query that's satisfied in the market. (submodular.)

Valid-utility Games: Price of Anarchy

- Theorem:(Vetta[2002]) The price of anarchy (of a mixed Nash equilibrium) in valid-utility games is at most 2.
- Theorem:(Vetta[2002]) Basic-utility games are potential games. In particular, best responses will converge to a pure Nash equilibrium.
- Theorem:(Goemans, Li, Mirrokni, Thottan[2004]) Pure Nash equilibria exist for market sharing games and can be found in polynomial time in the uniform case.

Basic-Utility Games : Convergence

- Theorem:(M.,Vetta[2004]) In basic-utility games, after one round of selfish behavior of players, they converge to a $\frac{1}{3}$-optimal solution.

Market Sharing Games : Convergence

- Theorem:(M.,Vetta[2004]) In basic-utility games, after one round of selfish behavior of players, they converge to a $\frac{1}{3}$-optimal solution.
- Theorem: (M., Vetta[2004]) In a market sharing game, after one round of selfish behavior of players, they converge to a $\frac{1}{\log(n)}$-optimal solution and this is almost tight.

Valid-utility Games: Convergence

- Theorem:(M., Vetta[2004]) For any $k > 0$, in valid-utility games, the social value after k rounds might be $\frac{1}{n}$ of the optimal social value.

Sink Equilibria

A sink equilibrium is a minimal set of states such that no best response move of any player goes out of these states.

Sink Equilibria

A sink equilibrium is a minimal set of states such that no best response move of any player goes out of these states.

If we enter a sink equilibrium, we are stuck there. Even random best-response paths cannot help us going out of a sink equilibria.

Price of anarchy for sink equilibria vs. the price of anarchy for Nash equilibria.

Sink Equilibria

- Theorem: (M., Vetta) In valid-utility games, even though the price of anarchy for Nash equilibria is $\frac{1}{2}$, the price of anarchy for sink equilibria is $\frac{1}{n}$.

  The performance of the Nash equilibria (or the price of anarchy for NE) is not a good measure for these games.

- Theorem: (M., Vetta) Finding a sink equilibrium in valid-utility games is PLS-Hard and there are states that are exponentially far from any sink equilibria.

Sink Equilibria

- Theorem: (M., Vetta) In valid-utility games, even though the price of anarchy for Nash equilibria is $\frac{1}{2}$, the price of anarchy for sink equilibria is $\frac{1}{n}$.

  The performance of the Nash equilibria (or the price of anarchy for NE) is not a good measure for these games.

- Theorem: (M., Vetta) Finding a sink equilibrium in valid-utility games is PLS-Hard and there are states that are exponentially far from any sink equilibria.

Conclusion

- Study Speed of convergence to approximates solutions instead of to Nash equilibria.
- Sink equilibria: an alternative measure to study the performance of the systems in lack of coordination.

Open problems

- Are there exponentially long fair paths in Basic-utility games.
- Is finding a 2-approximate Nash equilibrium for the cut game in P? How long does it take that 2-greedy players converge to a (2-approximate) Nash equilibrium? If it is polynomial, then finding a 2-approximate Nash equilibrium is in P.
- Are there exponentially long paths in the market sharing game?
- Study covering and random paths in other games.