On the Evolution of Selfish Routing

Simon Fischer
simon.fischer@cs.uni-dortmund.de

University of Dortmund
Dept. Computer Science 2

September 2, 2004

Objective 1:

Stronger Motivation

Objective 1:
Stronger Motivation

Objective 2:
Hurry up!

Classical Game Theory

- We want to analyse games involving large networks, latency functions, many agents, . . .

Classical Game Theory

- We want to analyse games involving large networks, latency functions, many agents, ...
- Our results are mainly concerned with equilibria: price of anarchy, taxes, ...

Classical Game Theory

- We want to analyse games involving large networks, latency functions, many agents, . . .
- Our results are mainly concerned with equilibria: price of anarchy, taxes, . . .
- In order to motivate our results we make assumptions about the agents: rationality, complete and accurate information etc.

Classical Game Theory

- We want to analyse games involving large networks, latency functions, many agents, . . .
- Our results are mainly concerned with equilibria: price of anarchy, taxes, . . .
- In order to motivate our results we make assumptions about the agents: rationality, complete and accurate information etc.
- Computing equilbria seems to be hard

Classical Game Theory

- We want to analyse games involving large networks, latency functions, many agents, . . .
- Our results are mainly concerned with equilibria: price of anarchy, taxes, . . .
- In order to motivate our results we make assumptions about the agents: rationality, complete and accurate information etc.
- Computing equilbria seems to be hard

Problem

How do we justify these assumptions when it comes to games related to something as mind-bogglingly large and inscrutable as the Internet?

Evolutionary Game Theory – Motivation

Goal

Make weaker assumptions about the agents.

Evolutionary Game Theory – Motivation

Goal

Make weaker assumptions about the agents.

- Model the *dynamics* of the population vector $\mathbf{x}(t)$.

Evolutionary Game Theory – Motivation

Goal

Make weaker assumptions about the agents.

- Model the *dynamics* of the population vector $\mathbf{x}(t)$.
- Several models:

Evolutionary Game Theory – Motivation

Goal

Make weaker assumptions about the agents.

- Model the *dynamics* of the population vector $\mathbf{x}(t)$.
- Several models:
  - Agents play against random partner, observe payoff difference, and imitate their opponents' strategy with probability „proportional" to this difference.

Evolutionary Game Theory – Motivation

Goal

Make weaker assumptions about the agents.

- Model the *dynamics* of the population vector $\mathbf{x}(t)$.
- Several models:
  - Agents play against random partner, observe payoff difference, and imitate their opponents' strategy with probability „proportional" to this difference.
  - Agents have a random aspiration level. Whenever they fall short of this level, they adopt a random strategy.

Evolutionary Game Theory – Motivation

Goal

Make weaker assumptions about the agents.

- Model the *dynamics* of the population vector $\mathbf{x}(t)$.
- Several models:
  - Agents play against random partner, observe payoff difference, and imitate their opponents' strategy with probability „proportional" to this difference.
  - Agents have a random aspiration level. Whenever they fall short of this level, they adopt a random strategy.
- Both models lead to the same dynamics!

Replicator dynamics

Definition (Replicator dynamics)

$$\dot{x}_i = \lambda(\mathbf{x}) \cdot x_i \cdot ((\mathbf{A}\mathbf{x})_i - \mathbf{x} \cdot \mathbf{A}\mathbf{x})$$

Replicator dynamics

Definition (Replicator dynamics)

$$\dot{x}_i = \lambda(\mathbf{x}) \cdot x_i \cdot ((\mathbf{A}\mathbf{x})_i - \mathbf{x} \cdot \mathbf{A}\mathbf{x})$$

- Desirable properties:

Replicator dynamics

Definition (Replicator dynamics)

$$\dot{x}_i = \lambda(\mathbf{x}) \cdot x_i \cdot ((\mathbf{Ax})_i - \mathbf{x} \cdot \mathbf{Ax})$$

- Desirable properties:
  - *homogeneity:* $\dot{x}_i = x_i \cdot g_i(\mathbf{x})$.

Replicator dynamics

Definition (Replicator dynamics)

$$\dot{x}_i = \lambda(\mathbf{x}) \cdot x_i \cdot ((\mathbf{Ax})_i - \mathbf{x} \cdot \mathbf{Ax})$$

- Desirable properties:
  - *homogeneity:* $\dot{x}_i = x_i \cdot g_i(\mathbf{x})$.
  - *monotonicity:* $(\mathbf{Ax})_i > (\mathbf{Ax})_j \Leftrightarrow g_i(\mathbf{x}) > g_j(\mathbf{x})$

Replicator dynamics

Definition (Replicator dynamics)

$$\dot{x}_i = \lambda(\mathbf{x}) \cdot x_i \cdot ((\mathbf{Ax})_i - \mathbf{x} \cdot \mathbf{Ax})$$

- Desirable properties:
  - *homogeneity:* $\dot{x}_i = x_i \cdot g_i(\mathbf{x})$.
  - *monotonicity:* $(\mathbf{Ax})_i > (\mathbf{Ax})_j \Leftrightarrow g_i(\mathbf{x}) > g_j(\mathbf{x})$
  - *aggregate monotonicity:* $\mathbf{y} \cdot \mathbf{Ax} > \mathbf{z} \cdot \mathbf{Ax} \Leftrightarrow \mathbf{y} \cdot g(\mathbf{x}) > \mathbf{z} \cdot g(\mathbf{x})$

Replicator dynamics

Definition (Replicator dynamics)

$$\dot{x}_i = \lambda(\mathbf{x}) \cdot x_i \cdot ((\mathbf{Ax})_i - \mathbf{x} \cdot \mathbf{Ax})$$

- Desirable properties:
  - *homogeneity:* $\dot{x}_i = x_i \cdot g_i(\mathbf{x})$.
  - *monotonicity:* $(\mathbf{Ax})_i > (\mathbf{Ax})_j \Leftrightarrow g_i(\mathbf{x}) > g_j(\mathbf{x})$
  - *aggregate monotonicity:* $\mathbf{y} \cdot \mathbf{Ax} > \mathbf{z} \cdot \mathbf{Ax} \Leftrightarrow \mathbf{y} \cdot g(\mathbf{x}) > \mathbf{z} \cdot g(\mathbf{x})$
- All dynamics having these properties can be represented this way!

Example: Paper-Scissors-Stone

$$\mathbf{A} = \begin{bmatrix} 1 & 2 & 0 \\ 0 & 1 & 2 \\ 2 & 0 & 1 \end{bmatrix}$$

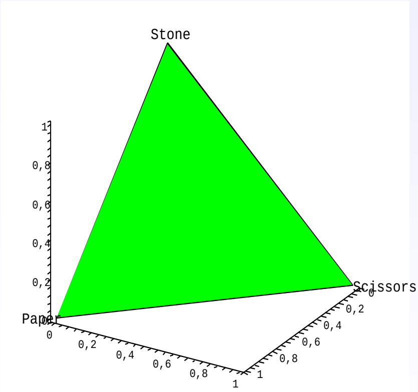

Example: Paper-Scissors-Stone

$$\mathbf{A} = \begin{bmatrix} 1 & 2 & 0 \\ 0 & 1 & 2 \\ 2 & 0 & 1 \end{bmatrix}$$

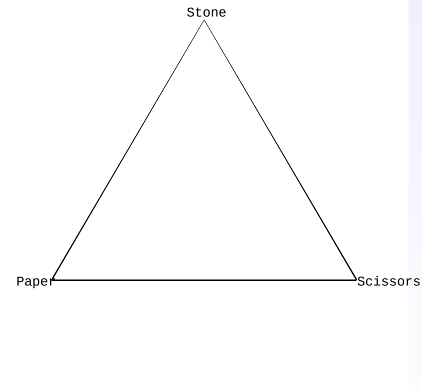

Example: Paper-Scissors-Stone

$$\mathbf{A} = \begin{bmatrix} 1 & 2 & 0 \\ 0 & 1 & 2 \\ 2 & 0 & 1 \end{bmatrix}$$

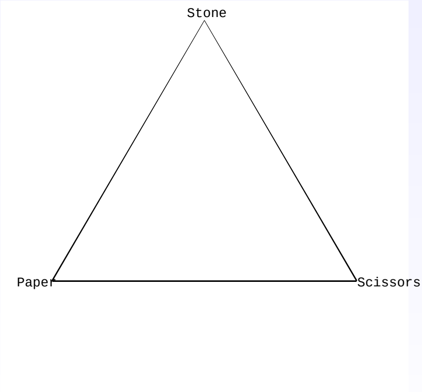

- Where is the Nash equilibrium? What should one expect?

Example: Paper-Scissors-Stone

$$\mathbf{A} = \begin{bmatrix} 1 & 2 & 0 \\ 0 & 1 & 2 \\ 2 & 0 & 1 \end{bmatrix}$$

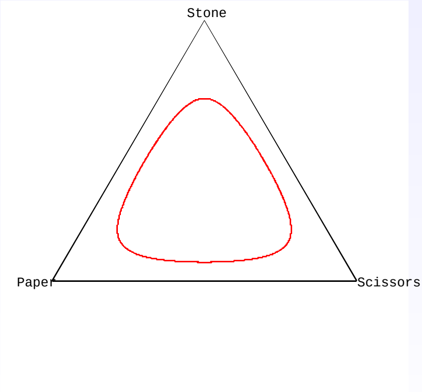

- Where is the Nash equilibrium? What should one expect?
- Stable orbit for the standard case

Example: Paper-Scissors-Stone

$$\mathbf{A} = \begin{bmatrix} 1 & 2+\epsilon & 0 \\ 0 & 1 & 2+\epsilon \\ 2+\epsilon & 0 & 1 \end{bmatrix}$$

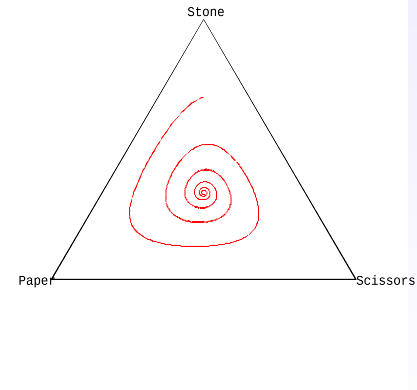

- Where is the Nash equilibrium? What should one expect?
- Stable orbit for the standard case
- Convergent for $\epsilon > 0$

Example: Paper-Scissors-Stone

$$\mathbf{A} = \begin{bmatrix} 1 & 2-\epsilon & 0 \\ 0 & 1 & 2-\epsilon \\ 2-\epsilon & 0 & 1 \end{bmatrix}$$

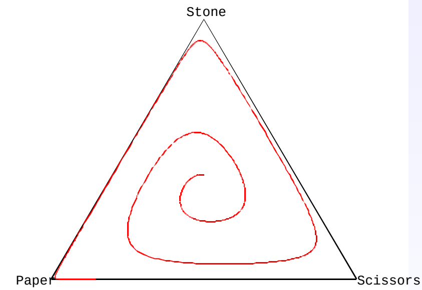

- Where is the Nash equilibrium? What should one expect?
- Stable orbit for the standard case
- Convergent for $\epsilon > 0$
- Divergent for $\epsilon < 0$

Selfish Routing

- Network $G = (V, E)$, latency functions defined for all edges $l_e : [0, 1] \mapsto \mathbb{R}_0^+$.

Selfish Routing

- Network $G = (V, E)$, latency functions defined for all edges $l_e : [0, 1] \mapsto \mathbb{R}_0^+$.
- Commodities $\mathcal{I}$ with demands d_i, s_i-t_i-paths P_i.

Selfish Routing

- Network $G = (V, E)$, latency functions defined for all edges $l_e : [0, 1] \mapsto \mathbb{R}_0^+$.
- Commodities $\mathcal{I}$ with demands d_i, s_i-t_i-paths P_i.
- Flow on path p is x_p; flow on edge $e \in E$ is x_e.

$$l_p(\mathbf{x}) = \sum_{e \in p} l_e(x_e)$$

Selfish Routing

- Network $G = (V, E)$, latency functions defined for all edges $l_e : [0, 1] \mapsto \mathbb{R}_0^+$.
- Commodities $\mathcal{I}$ with demands d_i, s_i-t_i-paths P_i.
- Flow on path p is x_p; flow on edge $e \in E$ is x_e.

$$l_p(\mathbf{x}) = \sum_{e \in p} l_e(x_e)$$

Definition (Wardrop equilibrium)

A flow $\mathbf{x}$ is at a Wardrop equilibrium iff for all $i \in \mathcal{I}$ and all $p, p' \in P_i$ with $x_p > 0$ it holds that $l_p(\mathbf{x}) \leq l_{p'}(\mathbf{x})$.

Selfish Routing and Game Theory

- What has been done so far?

Selfish Routing and Game Theory

- What has been done so far?
  - equilibria

Selfish Routing and Game Theory

- What has been done so far?
  - equilibria
  - *price of anarchy*, *coordination ratio*

Selfish Routing and Game Theory

- What has been done so far?
  - equilibria
  - *price of anarchy*, *coordination ratio*
  - pricing, taxes

Selfish Routing and Game Theory

- What has been done so far?
  - equilibria
  - *price of anarchy*, *coordination ratio*
  - pricing, taxes
  - ...

Selfish Routing and Game Theory

- What has been done so far?
  - equilibria
  - *price of anarchy*, *coordination ratio*
  - pricing, taxes
  - ...
- Here, we are especially interested in dynamical properties.

Synthesis of Models

- Replace average payoff $\mathbf{x} \cdot \mathbf{A}\mathbf{x}$ by $\bar{l}(\mathbf{x})$
- Replace payoff $(\mathbf{A}\mathbf{x})_p$ by latency $-l_p(\mathbf{x})$

Definition (Network dynamics)

$$\dot{x}_p = \lambda(\mathbf{x}) \cdot x_p \cdot (\bar{l}(\mathbf{x}) - l_p(\mathbf{x}))$$

Synthesis of Models

- Replace average payoff $\mathbf{x} \cdot \mathbf{A}\mathbf{x}$ by $\bar{l}(\mathbf{x})$
- Replace payoff $(\mathbf{A}\mathbf{x})_p$ by latency $-l_p(\mathbf{x})$

Definition (Network dynamics)

$$\dot{x}_p = \lambda_{i(p)}(\mathbf{x}) \cdot x_p \cdot (\bar{l}_{i(p)}(\mathbf{x}) - l_p(\mathbf{x}))$$

Synthesis of Models

- Replace average payoff $\mathbf{x} \cdot \mathbf{A}\mathbf{x}$ by $\bar{l}(\mathbf{x})$
- Replace payoff $(\mathbf{A}\mathbf{x})_p$ by latency $-l_p(\mathbf{x})$

Definition (Network dynamics)

$$\dot{x}_p = \lambda_{i(p)}(\mathbf{x}) \cdot x_p \cdot (\bar{l}_{i(p)}(\mathbf{x}) - l_p(\mathbf{x}))$$

- Choose $\lambda_i(\mathbf{x}) = 1/\bar{l}_i(\mathbf{x})$.

Synthesis of Models

- Replace average payoff $\mathbf{x} \cdot \mathbf{A}\mathbf{x}$ by $\bar{l}(\mathbf{x})$
- Replace payoff $(\mathbf{A}\mathbf{x})_p$ by latency $-l_p(\mathbf{x})$

Definition (Network dynamics)

$$\dot{x}_p = x_p \cdot \frac{\bar{l}_{i(p)}(\mathbf{x}) - l_p(\mathbf{x})}{\bar{l}_{i(p)}(\mathbf{x})}$$

- Choose $\lambda_i(\mathbf{x}) = 1/\bar{l}_i(\mathbf{x})$.

Synthesis of Models

- Replace average payoff $\mathbf{x} \cdot \mathbf{A}\mathbf{x}$ by $\bar{l}(\mathbf{x})$
- Replace payoff $(\mathbf{A}\mathbf{x})_p$ by latency $-l_p(\mathbf{x})$

Definition (Network dynamics)

$$\dot{x}_p = x_p \cdot \frac{\bar{l}_{i(p)}(\mathbf{x}) - l_p(\mathbf{x})}{\bar{l}_{i(p)}(\mathbf{x})}$$

- Choose $\lambda_i(\mathbf{x}) = 1/\bar{l}_i(\mathbf{x})$.
- Proposition: Solutions of this DE do not leave the simplex.

Stability

- Need to define stability in a new way.

Stability

- Need to define stability in a new way.

Definition (evolutionary stable)

A strategy $\mathbf{x}$ is evolutionary stable, iff it is at a Nash equilibrium and for all best replies $\mathbf{y}$ it holds that $\mathrm{x} \cdot \mathrm{l}(\mathbf{y}) < \mathbf{y} \cdot \mathrm{l}(\mathbf{y})$.

Stability

- Need to define stability in a new way.

Definition (evolutionary stable)

A strategy $\mathbf{x}$ is evolutionary stable, iff it is at a Nash equilibrium and for all best replies $\mathbf{y}$ it holds that $\mathrm{x} \cdot \mathrm{l}(\mathbf{y}) < \mathbf{y} \cdot \mathrm{l}(\mathbf{y})$.

- Best reply: No other strategy yields a better payoff.

Stability

- Need to define stability in a new way.

Definition (evolutionary stable)

A strategy $\mathbf{x}$ is evolutionary stable, iff it is at a Nash equilibrium and for all best replies $\mathbf{y}$ it holds that $\mathrm{x} \cdot \mathfrak{l}(\mathbf{y}) < \mathbf{y} \cdot \mathfrak{l}(\mathbf{y})$.

- Best reply: No other strategy yields a better payoff.

Proposition

For strictly increasing latency functions and single-commodity networks any Nash equilibrium is also evolutionary stable.

Skip proof

Stability – Proof

- In a Wardrop equilibrium $\mathbf{x}$, all latencies of used paths are equal, implying $\mathbf{x} \cdot \mathbf{l}(\mathbf{x}) \leq \mathbf{y} \cdot \mathbf{l}(\mathbf{x})$ and

$$\mathbf{y} \cdot \mathbf{l}(\mathbf{y}) \geq \mathbf{x} \cdot \mathbf{l}(\mathbf{x}) + \sum_{e \in E} y_e(l_e(\mathbf{y}) - l_e(\mathbf{x})).$$

Stability – Proof

- In a Wardrop equilibrium $\mathbf{x}$, all latencies of used paths are equal, implying $\mathbf{x} \cdot \mathbf{l}(\mathbf{x}) \leq \mathbf{y} \cdot \mathbf{l}(\mathbf{x})$ and

$$\mathbf{y} \cdot \mathbf{l}(\mathbf{y}) \geq \mathbf{x} \cdot \mathbf{l}(\mathbf{x}) + \sum_{e \in E} y_e(l_e(\mathbf{y}) - l_e(\mathbf{x})).$$

- For each edge there are three cases:

Stability – Proof

- In a Wardrop equilibrium $\mathbf{x}$, all latencies of used paths are equal, implying $\mathbf{x} \cdot \mathbf{l}(\mathbf{x}) \leq \mathbf{y} \cdot \mathbf{l}(\mathbf{x})$ and

$$\mathbf{y} \cdot \mathbf{l}(\mathbf{y}) \geq \mathbf{x} \cdot \mathbf{l}(\mathbf{x}) + \sum_{e \in E} y_e(l_e(\mathbf{y}) - l_e(\mathbf{x})).$$

- For each edge there are three cases:
  1. $y_e > x_e$. Then $\Rightarrow l_e(\mathbf{y}) > l_e(\mathbf{x})$.

Stability – Proof

- In a Wardrop equilibrium $\mathbf{x}$, all latencies of used paths are equal, implying $\mathbf{x} \cdot \mathbf{l}(\mathbf{x}) \leq \mathbf{y} \cdot \mathbf{l}(\mathbf{x})$ and

$$\mathbf{y} \cdot \mathbf{l}(\mathbf{y}) \geq \mathbf{x} \cdot \mathbf{l}(\mathbf{x}) + \sum_{e \in E} y_e(l_e(\mathbf{y}) - l_e(\mathbf{x})).$$

- For each edge there are three cases:
    1. $y_e > x_e$. Then $\Rightarrow l_e(\mathbf{y}) > l_e(\mathbf{x})$.
    2. $y_e < x_e$. Then $\Rightarrow l_e(\mathbf{y}) < l_e(\mathbf{x})$.

Stability – Proof

- In a Wardrop equilibrium $\mathbf{x}$, all latencies of used paths are equal, implying $\mathbf{x} \cdot \mathbf{l}(\mathbf{x}) \leq \mathbf{y} \cdot \mathbf{l}(\mathbf{x})$ and

$$\mathbf{y} \cdot \mathbf{l}(\mathbf{y}) \geq \mathbf{x} \cdot \mathbf{l}(\mathbf{x}) + \sum_{e \in E} y_e(l_e(\mathbf{y}) - l_e(\mathbf{x})).$$

- For each edge there are three cases:
  1. $y_e > x_e$. Then $\Rightarrow l_e(\mathbf{y}) > l_e(\mathbf{x})$.
  2. $y_e < x_e$. Then $\Rightarrow l_e(\mathbf{y}) < l_e(\mathbf{x})$.
  3. $y_e = x_e$. Then $\Rightarrow l_e(\mathbf{y}) = l_e(\mathbf{x})$.

Stability – Proof

- In a Wardrop equilibrium $\mathbf{x}$, all latencies of used paths are equal, implying $\mathbf{x} \cdot \mathbf{l}(\mathbf{x}) \leq \mathbf{y} \cdot \mathbf{l}(\mathbf{x})$ and

$$\mathbf{y} \cdot \mathbf{l}(\mathbf{y}) \geq \mathbf{x} \cdot \mathbf{l}(\mathbf{x}) + \sum_{e \in E} y_e(l_e(\mathbf{y}) - l_e(\mathbf{x})).$$

- For each edge there are three cases:
  1. $y_e > x_e$. Then $\Rightarrow l_e(\mathbf{y}) > l_e(\mathbf{x})$.
  2. $y_e < x_e$. Then $\Rightarrow l_e(\mathbf{y}) < l_e(\mathbf{x})$.
  3. $y_e = x_e$. Then $\Rightarrow l_e(\mathbf{y}) = l_e(\mathbf{x})$.
- In all cases, $y_e(l_e(\mathbf{y}) - l_e(\mathbf{x})) \geq x_e(l_e(\mathbf{y}) - l_e(\mathbf{x}))$, at least once with strict inequality.

Stability – Proof

- In a Wardrop equilibrium $\mathbf{x}$, all latencies of used paths are equal, implying $\mathbf{x} \cdot \mathbf{l}(\mathbf{x}) \leq \mathbf{y} \cdot \mathbf{l}(\mathbf{x})$ and

$$\mathbf{y} \cdot \mathbf{l}(\mathbf{y}) \geq \mathbf{x} \cdot \mathbf{l}(\mathbf{x}) + \sum_{e \in E} y_e(l_e(\mathbf{y}) - l_e(\mathbf{x})).$$

- For each edge there are three cases:
  1. $y_e > x_e$. Then $\Rightarrow l_e(\mathbf{y}) > l_e(\mathbf{x})$.
  2. $y_e < x_e$. Then $\Rightarrow l_e(\mathbf{y}) < l_e(\mathbf{x})$.
  3. $y_e = x_e$. Then $\Rightarrow l_e(\mathbf{y}) = l_e(\mathbf{x})$.
- In all cases, $y_e(l_e(\mathbf{y}) - l_e(\mathbf{x})) \geq x_e(l_e(\mathbf{y}) - l_e(\mathbf{x}))$, at least once with strict inequality.
- Altogether,

$$\mathbf{y} \cdot \mathbf{l}(\mathbf{y}) > \mathbf{x} \cdot \mathbf{l}(\mathbf{x}) + \sum_{e \in E} x_e(l_e(\mathbf{y}) - l_e(\mathbf{x})) = \mathbf{x} \cdot \mathbf{l}(\mathbf{y}). \quad \square$$

Convergence

- Consider the single-commodity case.

Convergence

- Consider the single-commodity case.
- The replicator dynamics does not find any initially unused strategies.

Convergence

- Consider the single-commodity case.
- The replicator dynamics does not find any initially unused strategies.
- Therefore, we restrict the strategy space to the set of paths P', which is the support of the initial population.

Convergence

- Consider the single-commodity case.
- The replicator dynamics does not find any initially unused strategies.
- Therefore, we restrict the strategy space to the set of paths P', which is the support of the initial population.

Proposition

Solutions of the replicator dynamics with initial population $\mathbf{x}_0$ converge against a restricted Wardrop equilibrium $\mathbf{x}^*$ w. r. t. $\mathbf{x}_0$, i. e. $\lim_{t\to\infty} \|\mathbf{x}(t) - \mathbf{x}^*\| = 0$.

Skip proof

Convergence – Proof

- Use conditional entropy as a potential function:

$$H_{\mathbf{x}^*}(\mathbf{x}) := \sum_{p \in P} x_p^* \ln \frac{x_p^*}{x_p}.$$

Convergence – Proof

- Use conditional entropy as a potential function:

$$H_{\mathbf{x}^*}(\mathbf{x}) := \sum_{p \in P} x_p^* \ln \frac{x_p^*}{x_p}.$$

- It holds that $\|\mathbf{x}^* - \mathbf{x}\|^2 \leq H_{\mathbf{x}^*}(\mathbf{x})$.

Convergence – Proof

- Use conditional entropy as a potential function:

$$H_{\mathbf{x}^*}(\mathbf{x}) := \sum_{p \in P} x_p^* \ln \frac{x_p^*}{x_p}.$$

- It holds that $\|\mathbf{x}^* - \mathbf{x}\|^2 \leq H_{\mathbf{x}^*}(\mathbf{x})$.
- Calculate the derivative $\dot{H}_{\mathbf{x}^*}(\mathbf{x}) = -\sum_{p \in P} x_p^* \frac{\dot{x}_p}{x_p}$.

Convergence – Proof

- Use conditional entropy as a potential function:

$$H_{\mathbf{x}^*}(\mathbf{x}) := \sum_{p \in P} x_p^* \ln \frac{x_p^*}{x_p}.$$

- It holds that $\|\mathbf{x}^* - \mathbf{x}\|^2 \leq H_{\mathbf{x}^*}(\mathbf{x})$.
- Calculate the derivative $\dot{H}_{\mathbf{x}^*}(\mathbf{x}) = -\sum_{p \in P} x_p^* \frac{\dot{x}_p}{x_p}$.
- Substituting the replicator dynamics yields $\dot{H}_{\mathbf{x}^*}(\mathbf{x}) = \lambda(\mathbf{x}) \cdot (\mathbf{x}^* - \mathbf{x}) \cdot \mathbf{l}(\mathbf{x})$.

Convergence – Proof

- Use conditional entropy as a potential function:

$$H_{\mathbf{x}^*}(\mathbf{x}) := \sum_{p \in P} x_p^* \ln \frac{x_p^*}{x_p}.$$

- It holds that $\|\mathbf{x}^* - \mathbf{x}\|^2 \leq H_{\mathbf{x}^*}(\mathbf{x})$.
- Calculate the derivative $\dot{H}_{\mathbf{x}^*}(\mathbf{x}) = -\sum_{p \in P} x_p^* \frac{\dot{x}_p}{x_p}$.
- Substituting the replicator dynamics yields $\dot{H}_{\mathbf{x}^*}(\mathbf{x}) = \lambda(\mathbf{x}) \cdot (\mathbf{x}^* - \mathbf{x}) \cdot \mathbf{l}(\mathbf{x})$.
- This difference is negative because of evolutionary stability. □

Approximating Equilibria

- Need to define approximations.

Approximating Equilibria

- Need to define approximations.
- Euclidian distance not suitable here.

Approximating Equilibria

- Need to define approximations.
- Euclidian distance not suitable here.
- We do not want to wait for minorities, since they can increase time of convergence arbitrarily.

Approximating Equilibria

- Need to define approximations.
- Euclidian distance not suitable here.
- We do not want to wait for minorities, since they can increase time of convergence arbitrarily.

Definition (ϵ-equilibrium)

Let P_ϵ be the set of all paths with latency $(1+\epsilon)\cdot\bar{l}$ or above. Let x_ϵ be the number of agents on these paths. A population is at an ϵ-equilibrium, iff $x_\epsilon \leq \epsilon$.

Upper Bound – Potential Function

Theorem

For strictly increasing latency functions and single commodity networks the replicator dynamics reaches an ϵ-equilibrium in time $\mathcal{O}(\epsilon^{-3} \cdot \ln(l_{\max}/l^*))$.

Skip proof

Upper Bound – Potential Function

Theorem

For strictly increasing latency functions and single commodity networks the replicator dynamics reaches an ϵ-equilibrium in time $\mathcal{O}(\epsilon^{-3} \cdot \ln(l_{\max}/l^*))$.

Skip proof

- Use generalisation of the Rosenthal's potential function:

$$\Phi(\mathbf{x}) := \left(\sum_{e \in E} \int_0^{x_e} l_e(x)\, dx \right) + l^*.$$

Upper Bound – Potential Function

Theorem

For strictly increasing latency functions and single commodity networks the replicator dynamics reaches an ϵ-equilibrium in time $\mathcal{O}(\epsilon^{-3} \cdot \ln(l_{\max}/l^*))$.

Skip proof

- Use generalisation of the Rosenthal's potential function:

$$\Phi(\mathbf{x}) := \left(\sum_{e \in E} \int_0^{x_e} l_e(x)\, dx \right) + l^*.$$

- Calculate derivative and plug in replicator dynamics:

$$\dot{\Phi} = \lambda(\mathbf{x}) \left(\bar{l}(\mathbf{x})^2 - \sum_{p \in P} x_p l_p(\mathbf{x})^2 \right) < 0.$$

Upper Bound – Some Estimates

- At least ϵ agents have latency of $(1+\epsilon)\bar{l}$ or above.

Upper Bound – Some Estimates

- At least ϵ agents have latency of $(1+\epsilon)\bar{l}$ or above.
- Decrease of potential is minimal when other agents are distributed uniformly among the remaining strategies.

Upper Bound – Some Estimates

- At least ϵ agents have latency of $(1+\epsilon)\bar{l}$ or above.
- Decrease of potential is minimal when other agents are distributed uniformly among the remaining strategies.
- This yields $\dot{\Phi} \leq -\epsilon^3 \cdot \bar{l}(\mathbf{x})/2$.

Upper Bound – Some Estimates

- At least ϵ agents have latency of $(1+\epsilon)\bar{l}$ or above.
- Decrease of potential is minimal when other agents are distributed uniformly among the remaining strategies.
- This yields $\dot{\Phi} \leq -\epsilon^3 \cdot \bar{l}(\mathbf{x})/2$.
- Since $\bar{l} \geq \sum_e \int_0^{x_e} l_e(x)\,dx$ and $\bar{l} \geq l^*$, it holds that $2 \cdot \bar{l}(\mathbf{x}) \geq \Phi$ and therefore

$$\dot{\Phi} \leq -\epsilon^3 \Phi/4.$$

Upper Bound – Solution of the DE

- Solving the differential inequality

$$\dot{\Phi} \leq -\epsilon^3 \Phi / 4$$

we get

$$\Phi(t) \leq \Phi(0) e^{-\epsilon^3/4 \cdot t}.$$

Upper Bound – Solution of the DE

- Solving the differential inequality

$$\dot{\Phi} \leq -\epsilon^3 \Phi / 4$$

we get

$$\Phi(t) \leq \Phi(0) e^{-\epsilon^3/4 \cdot t}.$$

- Since $\Phi \geq l^*$ and $\Phi(0) \leq 2 \cdot l_{\max}$, Φ reaches its minimum in time at most

$$t = \mathcal{O}\left(\epsilon^{-3} \cdot \ln \frac{l_{\max}}{l^*}\right).$$

□

Lower Bound

- In what respect are the assumptions of the analysis of the upper bound pessimistic?

Lower Bound

- In what respect are the assumptions of the analysis of the upper bound pessimistic?
  - The estimate $\bar{l} \geq \sum_{e \in E} \int_0^{x_e} l_e(x)\,dx$ holds with equality when latency functions are constant.

Lower Bound

- In what respect are the assumptions of the analysis of the upper bound pessimistic?
  - The estimate $\bar{l} \geq \sum_{e \in E} \int_0^{x_e} l_e(x)\,dx$ holds with equality when latency functions are constant.
  - We assume that always exactly ϵ agents are by a factor of exactly $1 + \epsilon$ worse than the average.

Lower Bound

- In what respect are the assumptions of the analysis of the upper bound pessimistic?
  - The estimate $\bar{l} \geq \sum_{e \in E} \int_0^{x_e} l_e(x)\, dx$ holds with equality when latency functions are constant.
  - We assume that always exactly ϵ agents are by a factor of exactly $1 + \epsilon$ worse than the average.
- No more estimates.

Lower Bound

- In what respect are the assumptions of the analysis of the upper bound pessimistic?
  - The estimate $\bar{l} \geq \sum_{e \in E} \int_0^{x_e} l_e(x)\,dx$ holds with equality when latency functions are constant.
  - We assume that always exactly ϵ agents are by a factor of exactly $1+\epsilon$ worse than the average.
- No more estimates.
- Can we find a scenario where all inequalities hold with equality?

Lower Bound – Construction

- n parallel links $\{0, \ldots, n-1\}$.
- Link i has constant latency $(1 + c\epsilon)^{-i}$.
- Start with $x_0 = 2\epsilon$, $x_1 = 1 - 2\epsilon - \gamma$, $\sum_{i=2}^{n-1} x_i = r$ for some very small rest γ.
- Link 1 dominates the average, $x_\epsilon = x_0 > \epsilon$.

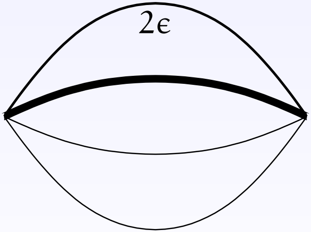

Lower Bound – Construction

- n parallel links $\{0, \ldots, n-1\}$.
- Link i has constant latency $(1 + c\epsilon)^{-i}$.
- Start with $x_0 = 2\epsilon$, $x_1 = 1 - 2\epsilon - \gamma$, $\sum_{i=2}^{n-1} x_i = r$ for some very small rest γ.
- Link 1 dominates the average, $x_\epsilon = x_0 > \epsilon$.

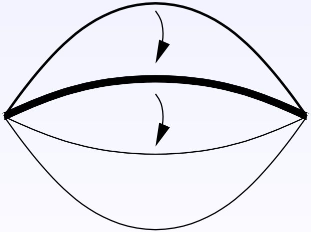

Lower Bound – Construction

- n parallel links $\{0, \ldots, n-1\}$.
- Link i has constant latency $(1 + c\epsilon)^{-i}$.
- Start with $x_0 = 2\epsilon$, $x_1 = 1 - 2\epsilon - \gamma$, $\sum_{i=2}^{n-1} x_i = r$ for some very small rest γ.
- Link 1 dominates the average, $x_\epsilon = x_0 > \epsilon$.

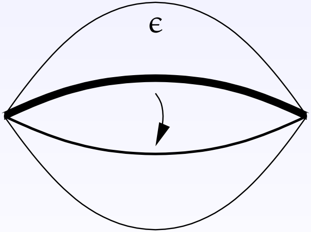

Lower Bound – Construction

- n parallel links $\{0, \ldots, n-1\}$.
- Link i has constant latency $(1 + c\epsilon)^{-i}$.
- Start with $x_0 = 2\epsilon$, $x_1 = 1 - 2\epsilon - \gamma$, $\sum_{i=2}^{n-1} x_i = r$ for some very small rest γ.
- Link 1 dominates the average, $x_\epsilon = x_0 > \epsilon$.

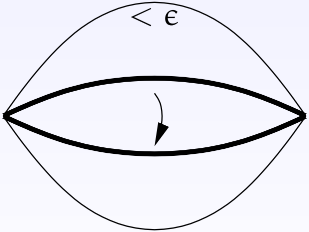

Lower Bound – Construction

- n parallel links $\{0, \ldots, n-1\}$.
- Link i has constant latency $(1 + c\epsilon)^{-i}$.
- Start with $x_0 = 2\epsilon$, $x_1 = 1 - 2\epsilon - \gamma$, $\sum_{i=2}^{n-1} x_i = r$ for some very small rest γ.
- Link 1 dominates the average, $x_\epsilon = x_0 > \epsilon$.

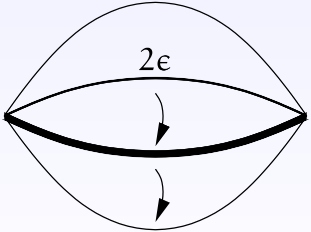

Lower Bound – Analysis

- For all (except the first) links we need time $\Omega(1/\epsilon)$.

Lower Bound – Analysis

- For all (except the first) links we need time $\Omega(1/\epsilon)$.
- We calculate the number of links n, which can be inserted into the network for some fixed $r = l_{\max}/l^* = (1 + c\epsilon)^{n-1}$.

Lower Bound – Analysis

- For all (except the first) links we need time $\Omega(1/\epsilon)$.
- We calculate the number of links n, which can be inserted into the network for some fixed $r = l_{\max}/l^* = (1 + c\epsilon)^{n-1}$.
- Solving this yields $n = \frac{\ln r}{\ln(1+c\epsilon)} + 1 \geq c^{-1}\epsilon^{-1} \ln r$.

Lower Bound – Analysis

- For all (except the first) links we need time $\Omega(1/\epsilon)$.
- We calculate the number of links n, which can be inserted into the network for some fixed $r = l_{\max}/l^* = (1 + c\epsilon)^{n-1}$.
- Solving this yields $n = \frac{\ln r}{\ln(1+c\epsilon)} + 1 \geq c^{-1}\epsilon^{-1}\ln r$.

Theorem

For any fixed value of $r = \frac{l_{\max}}{l^*}$ there exists a network, in which the time of reaching am ϵ-equilibrium is at least $\Omega(\epsilon^{-2} \cdot \ln r)$.

Some remarks

- This analysis is fairly robust – agents are not required to behave exactly as described by the dynamics.

Some remarks

- This analysis is fairly robust – agents are not required to behave exactly as described by the dynamics.
  - We only require an ϵ-fraction of the agents to improve their strategy when they are above $(1+\epsilon)\bar{l}$

Some remarks

- This analysis is fairly robust – agents are not required to behave exactly as described by the dynamics.
  - We only require an ϵ-fraction of the agents to improve their strategy when they are above $(1+\epsilon)\bar{l}$
  - One might chose λ in a different way, e. g. the prob. to change $\frac{l_i - l_j}{l_i}$

Some remarks

- This analysis is fairly robust – agents are not required to behave exactly as described by the dynamics.
  - We only require an ϵ-fraction of the agents to improve their strategy when they are above $(1+\epsilon)\bar{l}$
  - One might chose λ in a different way, e. g. the prob. to change $\frac{l_i - l_j}{l_i}$
- In the multicommodity case

Some remarks

- This analysis is fairly robust – agents are not required to behave exactly as described by the dynamics.
  - We only require an ϵ-fraction of the agents to improve their strategy when they are above $(1+\epsilon)\bar{l}$
  - One might chose λ in a different way, e. g. the prob. to change $\frac{l_i - l_j}{l_i}$
- In the multicommodity case
  - (non-unique) Nash equilibria are *essentially evolutionary stable* and the replicator dynamics converges against one of them

Some remarks

- This analysis is fairly robust – agents are not required to behave exactly as described by the dynamics.
  - We only require an ϵ-fraction of the agents to improve their strategy when they are above $(1+\epsilon)\bar{l}$
  - One might chose λ in a different way, e. g. the prob. to change $\frac{l_i - l_j}{l_i}$
- In the multicommodity case
  - (non-unique) Nash equilibria are *essentially evolutionary stable* and the replicator dynamics converges against one of them
  - we can show an upper bound of $\mathcal{O}(\epsilon^{-3} \ln l_{\max}/l^*)$ using weaker estimates.

Conclusion

- For strictly increasing latency functions, Wardrop equilibria are evolutionary stable.
- Solutions of the replicator dynamics converge towards restricted Wardrop equilibria.
- The time until almost all are almost happy is bounded by
  - $\mathcal{O}(\epsilon^{-3} \cdot \ln \frac{l_{max}}{l^*})$ for single commodity networks
  - $\mathcal{O}(\epsilon^{-2} \cdot \ln \frac{l_{max}}{l^*})$ for multicommodity parallel links.
  - $\mathcal{O}(\epsilon^{-3} \cdot \frac{l_{max}}{l^*})$ for general multicommodity networks, and
  - $\Omega(\epsilon^{-2} \cdot \ln \frac{l_{max}}{l^*})$ for all of these.

Open Questions

- How can we improve the bound for the multi-commodity case?

Open Questions

- How can we improve the bound for the multi-commodity case?
- What happens in a discrete process?

Open Questions

- How can we improve the bound for the multi-commodity case?
- What happens in a discrete process?
- How can we use this to implement routing algorithms?

Simon Fischer and Berthold Vöcking.
On the evolution of selfisch routing.
In *Proc. of ESA 2004*, 2004.

J. Maynard Smith and G. R. Price.
The logic of animal conflict.
Nature, 246:15–18, 1973.

Robert W. Rosenthal.
A class of games possessing pure-strategy Nash equilibria.
International Journal of Game Theory, 2:65–67, 1973.

Jörgen W. Weibull.
Evolutionary Game Theory.
MIT press, 1995.

Thank you for your attention!