

On the Evolution of Selfish Routing^{*}

Simon Fischer and Berthold Vöcking

FB Informatik, LS2, University of Dortmund, 44221 Dortmund, Germany
{fischer,voecking}@ls2.cs.uni-dortmund.de

Abstract. We introduce a model to study the temporal behaviour of selfish agents in networks. So far, most of the analysis of selfish routing is concerned with static properties of equilibria which is one of the most fundamental paradigms in classical Game Theory. By adopting a generalised approach of Evolutionary Game Theory we extend the model of selfish routing to study the dynamical behaviour of agents.

For symmetric games corresponding to singlecommodity flow, we show that the game converges to a Nash equilibrium in a restricted strategy space. In particular we prove that the time for the agents to reach an ϵ -approximate equilibrium is polynomial in ϵ and only logarithmic in the ratio between maximal and optimal latency. In addition, we present an almost matching lower bound in the same parameters.

Furthermore, we extend the model to asymmetric games corresponding to multicommodity flow. Here we also prove convergence to restricted Nash equilibria, and we derive upper bounds for the convergence time that are linear instead of logarithmic.

1 Introduction

Presently, the application of Game Theory to networks and congestion games is gaining a growing amount of interest in Theoretical Computer Science. One of the most fundamental paradigms of Game Theory is the notion of Nash equilibria. So far, many results on equilibria have been derived. Most of these are concerned with the ratio between average cost at an equilibrium and the social optimum, mostly referred to as the *coordination ratio* or the *price of anarchy* [2, 7, 8, 13], ways to improve this ratio, e.g. by taxes [1], and algorithmic complexity and efficiency of computing such equilibria [3–5].

Classical Game Theory is based on fully rational behaviour of players and global knowledge of all details of the game under study. For routing games in large networks like the Internet, these assumptions are clearly far away from being realistic. Evolutionary Game Theory makes a different attempt to explain why large populations of agents may or may not “converge” towards equilibrium states. This theory is mainly based on the so-called replicator dynamics, a model of an evolutionary process in which agents revise their strategies from time to time based on local observations.

^{*} Supported in part by the EU within the 6th Framework Programme under contract 001907 (DELIS) and by DFG grant Vo889/1-2.

In this paper, we apply Evolutionary Game Theory to selfish routing. This enables us to study the dynamics of selfish routing rather than only the static structure of equilibria. We prove that the only existing fixed points of the replicator dynamics are Nash equilibria over a restricted strategy space. We prove that these fixed points are “evolutionary stable” which implies that the replicator dynamics converges to one of them. One standard approach in Evolutionary Game Theory to prove stability is based on symmetry properties of payoff matrices. However, we cannot simply cast these results as our model of selfish routing allows for arbitrary latency functions whereas the existing literature on Evolutionary Game Theory assumes affine payoff functions corresponding to linear latency functions with zero offset. In fact our proof of evolutionary stability is based on monotonicity instead of symmetry.

Another aspect that – to our knowledge – has neither been considered in Evolutionary nor in classical Game Theory is the time it takes to reach or come close to equilibria – the speed of convergence. We believe that this is an issue of particular importance as equilibria are only meaningful if they are reached in reasonable time. In fact we can prove that symmetric congestion games – corresponding to singlecommodity flow – converge very quickly to an approximate equilibrium. For asymmetric congestion games our bounds are slightly weaker.

The well established models of selfish routing and Evolutionary Game Theory are described in Section 2. In Section 3 we show how a generalisation of the latter can be applied to the first. In Section 4 we present our results on the speed of convergence. All proofs are collected in Section 5. We finish with some conclusions and open problems.

2 Known Models

2.1 Selfish Routing

Consider a network $G = (V, E)$ and for all $e \in E$, latency functions $l_e : [0, 1] \mapsto \mathbb{R}$ assigned to the edges mapping load to latency, and a set of commodities \mathcal{I} . For commodity $i \in \mathcal{I}$ there is a fixed flow demand of d_i that is to be routed from source s_i to sink t_i . The total flow demand of all commodities is 1. Denote the set of s_i - t_i -paths by $P_i \subseteq \mathcal{P}(E)$, where $\mathcal{P}(E)$ is the power set of E . For simplicity of notation we assume that the P_i are disjoint, which is certainly true if the pairs (s_i, t_i) are pairwise distinct. Then there is a unique commodity $i(p)$ associated with each path $p \in P$. Let furthermore $P = \bigcup_{i \in \mathcal{I}} P_i$.

For $p \in P$ denote the amount of flow routed over path p by x_p . We combine the individual values x_p into a vector \mathbf{x} . The set of legal flows is the simplex¹ $\Delta := \{\mathbf{x} | \forall i \in \mathcal{I} : \sum_{p \in P_i} x_p = d_i\}$. Furthermore, for $e \in E$, $x_e := \sum_{p \ni e} x_p$ is the total load of edge e . The latency of edge $e \in E$ is $l_e(x_e)$ and the latency of path p is

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

¹ Strictly speaking, Δ is not a simplex, but a product of $|\mathcal{I}|$ simplices scaled by d_i .

The average latency with respect to commodity $i \in \mathcal{I}$ is

$$\bar{l}_i(\mathbf{x}) = d_i^{-1} \sum_{p \in P_i} x_p l_p(\mathbf{x}).$$

In a more general setting one could abstract from graphs and consider arbitrary sets of resources E and arbitrary non-empty subsets of E as legal strategies. We do not do this for simplicity, though it should be clear that all our considerations below also hold for the general case.

In order to study the behaviour of users in networks, one assumes that there are an infinite number of agents carrying an infinitesimal amount of flow each. In this context the flow vector \mathbf{x} can also be seen as a population of agents where x_p is the non-integral “number” of agents selecting strategy p .

The individual agents strive to choose a path that minimises their personal latency regardless of social benefit. This does not impute maliciousness to the agents, but simply arises from the lack of a possibility to coordinate strategies. One can ask: What population will arise if we assume complete rationality and if all agents have full knowledge about the network and the latency functions?

This is where the notion of an equilibrium comes into play. A flow or population \mathbf{x} , is said to be at Nash equilibrium [11], when no agent has an incentive to change their strategy. The classical existence result by Nash holds if mixed strategies are allowed, i. e., agents may choose their strategy by a random distribution. However, mixed strategies and an infinite population of agents using pure strategies are basically the same. A very useful characterisation of a Nash equilibrium is the Wardrop equilibrium [6].

Definition 1 (Wardrop equilibrium). *A flow \mathbf{x} is at Wardrop equilibrium if and only if for all $i \in \mathcal{I}$ and all $p, p' \in P_i$ with $x_p > 0$ it holds that $l_p(\mathbf{x}) \leq l_{p'}(\mathbf{x})$.*

Several interesting static aspects of equilibria have been studied, among them the famous *price of anarchy*, which is the ratio between average latency at a Nash equilibrium and optimal average latency.

2.2 Evolutionary Game Theory

Classical Game Theory assumes that all agents – equipped with complete knowledge about the game and full rationality – will come to a Wardrop equilibrium. However, these assumptions seem far from realistic when it comes to networks. Evolutionary Game Theory gets rid of these assumptions by modelling the agents’ behaviour in a very natural way that requires the agents simply to observe their own and other agents’ payoff and strategy and change their own strategies based on these observations. Starting with an initial population vector $\mathbf{x}(0)$ – as a function of time – one is interested in its derivative with respect to time $\dot{\mathbf{x}}$.

Originally, Evolutionary Game Theory derives dynamics for large populations of individuals from symmetric two-player games. Any two-player game can be described by a matrix $\mathbf{A} = (a_{ij})$ where a_{ij} is the payoff of strategy i when played

against strategy j . Suppose that two individuals are drawn at random from a large population to play a game described by the payoff matrix \mathbf{A} . Let x_i denote the fraction of players playing strategy i at some given point of time. Then the expected payoff of an agent playing strategy i against the opponent randomly chosen from population \mathbf{x} is the i^{th} component of the matrix product $(\mathbf{A}\mathbf{x})_i$. The average payoff of the entire population is $\mathbf{x} \cdot \mathbf{A}\mathbf{x}$, where \cdot is the inner, or scalar, product.

Consider an initial population in which each agent is assigned a pure strategy. At each point of time, each agent plays against an opponent chosen uniformly at random. The agent observes its own and its opponents payoff and decides to imitate its opponent by adopting its strategy with probability proportional to the payoff difference. One could argue that often it is not possible to observe the opponent's payoff. In that case, consider a random aspiration level for each agent. Whenever an agent falls short of this level, it adopts a randomly observed strategy. Interestingly, both scenarios lead to a dynamics which can be described by the differential equation

$$\dot{x}_i = \lambda(\mathbf{x}) \cdot x_i \cdot ((\mathbf{A}\mathbf{x})_i - \mathbf{x}\mathbf{A}\mathbf{x}) \quad (1)$$

for some positive function λ . This equation has interesting and desirable properties. First, the growth rate of strategy i should clearly be proportional to the number of agents already playing this strategy, if we assume homogeneous agents. Then we have $\dot{x}_i = x_i \cdot g_i(\mathbf{x})$. Secondly, the growth rate $g_i(\mathbf{x})$ should increase with payoff, i. e., $(\mathbf{A}\mathbf{x})_i > (\mathbf{A}\mathbf{x})_j$ should imply $g_i(\mathbf{x}) > g_j(\mathbf{x})$ and vice versa. Dynamics having this property are called *monotone*. In order to extend this, we say that \mathbf{g} is *aggregate monotone* if for inhomogeneous (sub)populations the total growth rate of the subpopulations increases with the average payoff of the subpopulation, i. e., for vectors $\mathbf{y}, \mathbf{z} \in \Delta$ it holds that $\mathbf{y} \cdot \mathbf{A}\mathbf{x} < \mathbf{z} \cdot \mathbf{A}\mathbf{x}$ if and only if $\mathbf{y} \cdot \mathbf{g}(\mathbf{x}) < \mathbf{z} \cdot \mathbf{g}(\mathbf{x})$. It is known that all aggregate monotone dynamics can be written in the form of equation (1). In Evolutionary Game Theory the most common choice for λ seems to be the constant function 1. For $\lambda(\mathbf{x}) = 1$, this dynamics is known as the *replicator dynamics*.

Since the differential equation (1) contains cubic terms there is no general method for solving it. It is found that under some reasonable conditions fixed points of this dynamics coincide with Nash equilibria. For a more comprehensive introduction into Evolutionary Game Theory see for example [14].

3 Evolutionary Selfish Flow

We extend the dynamics known from Evolutionary Game Theory to our model of selfish routing. We will see that there is a natural generalisation of equation (1) for our scenario.

3.1 Dynamics for Selfish Routing

First consider symmetric games corresponding to singlecommodity flow. In congestion games latency relates to payoff, but with opposite sign. Unless we have

linear latency functions without offset, we cannot express the payoff by means of a matrix \mathbf{A} . Therefore we replace the term $(\mathbf{A}\mathbf{x})_i$ by the latency $l_i(\mathbf{x})$. The average payoff $\mathbf{x} \cdot \mathbf{A}\mathbf{x}$ is replaced by the average latency $\bar{l}(\mathbf{x})$. Altogether we have

$$\dot{x}_p = \lambda(\mathbf{x}) \cdot x_p \cdot (\bar{l}(\mathbf{x}) - l_p(\mathbf{x})). \quad (2)$$

Evolutionary Game Theory also allows for asymmetric games, i. e., games where each agent belongs to one class determining the set of legal strategies. In principle, asymmetric games correspond to multicommodity flow. The suggested generalisations for asymmetric games, however, are not particularly useful in our context since they assume that agents of the same class do not play against each other. We suggest a simple and natural generalisation towards multicommodity flow. Agents behave exactly as in the singlecommodity case but they compare their own latency to the average latency over the agents in the same class. Although in Evolutionary Game Theory $\lambda = 1$ seems to be the most common choice even for asymmetric games, we suggest to choose the factor λ dependent on the commodity. We will later see why this might be useful. This gives the following variant of the replicator dynamics:

$$\dot{x}_p = \lambda_{i(p)}(\mathbf{x}) \cdot x_p \cdot (\bar{l}_{i(p)}(\mathbf{x}) - l_p(\mathbf{x})). \quad (3)$$

We believe that this, in fact, is a realistic model of communication in networks since agents only need to “communicate” with agents having the same source and destination nodes.

In order for equation (3) to constitute a legal dynamics, we must ensure that for all t the population shares sum up to the total flow demand.

Proposition 1. *Let $\mathbf{x}(t)$ be a solution of equation (3). Then $\mathbf{x}(t) \in \Delta$ for all t .*

This is proved in Section 5. Note that strategies that are not present in the initial population are not generated by the dynamics. Conversely, positive strategies never get completely extinct.

3.2 Stability and Convergence

Maynard Smith and Price [9] introduced the concept of evolutionary stability which is stricter than the concept of a Nash equilibrium. Intuitively, a strategy is evolutionary stable if it is at a Nash equilibrium and earns more against a mutant strategy than the mutant strategy earns against itself. Adopted to selfish routing we can define it in the following way.

Definition 2 (evolutionary stable). *A strategy \mathbf{x} is evolutionary stable if it is at a Nash equilibrium and $\mathbf{x} \cdot \mathbf{l}(\mathbf{y}) < \mathbf{y} \cdot \mathbf{l}(\mathbf{y})$ for all best replies \mathbf{y} to \mathbf{x} .*

A strategy \mathbf{y} is a *best reply* to strategy \mathbf{x} if no other strategy yields a better payoff when played against \mathbf{x} .

Proposition 2. *Suppose that latency functions are strictly increasing. If \mathbf{x} is at a Wardrop equilibrium, then $\mathbf{x} \cdot \mathbf{l}(\mathbf{y}) < \mathbf{y} \cdot \mathbf{l}(\mathbf{y})$ for all \mathbf{y} (especially, for all best replies) and hence \mathbf{x} is evolutionary stable.*

Consider a Wardrop equilibrium \mathbf{x} and a population \mathbf{y} such that for some strategy $p \in P$, $y_p = 0$ and $x_p > 0$. The replicator dynamics starting at \mathbf{y} does not converge to \mathbf{x} since it ignores strategy p . Therefore we consider a *restricted strategy space* P' containing only the paths with positive value and *restricted Wardrop equilibria* over this restricted strategy space.

Proposition 3. *Suppose that for all $i, j \in \mathcal{I}$, $\lambda_i = \lambda_j$ and $\lambda_i(\mathbf{x}) \geq \epsilon$ for some $\epsilon > 0$ and any $\mathbf{x} \in \Delta$. Let $\mathbf{y}(t)$ be a solution to the replicator dynamics (3) and let \mathbf{x} be a restricted Wardrop equilibrium with respect to $\mathbf{y}(0)$. Then \mathbf{y} converges towards \mathbf{x} , i. e., $\lim_{t \rightarrow \infty} \|\mathbf{y}(t) - \mathbf{x}\| = 0$.*

Both propositions are proved in Section 5.

4 Speed of Convergence

In order to study the speed of convergence we suggest two modifications to the standard approach of Evolutionary Game Theory. First, one must ensure that the growth rate of the population shares does not depend on the scale by which we measure latency, as is the case if we choose $\lambda(\mathbf{x}) = 1$ or any other constant. Therefore we suggest to choose $\lambda_i(\mathbf{x}) = \bar{l}_i(\mathbf{x})^{-1}$. This choice arises quite naturally if we assume that the probability of agents changing their strategy depends on their relative latency with respect to the current average latency. We call the resulting dynamics the *relative replicator dynamics*.

Secondly, the Euclidian distance $\|\mathbf{y} - \mathbf{x}\|$ is not a suitable measure for approximation. Since “sleeping minorities on cheap paths” may grow arbitrarily slowly, it may take arbitrarily long for the current population to come close to the final Nash equilibrium in Euclidian distance. The idea behind our definition of approximate equilibria is not to wait for these sleeping minorities.

Definition 3 (ϵ -approximate equilibrium). *Let P_ϵ be the set of paths that have latency at least $(1 + \epsilon) \cdot \bar{l}$, i. e., $P_\epsilon = \{p \in P \mid l_p(\mathbf{x}) \geq (1 + \epsilon) \cdot \bar{l}\}$ and let $x_\epsilon := \sum_{p \in P_\epsilon} x_p$ be the number of agents using these paths. A population \mathbf{x} is said to be at an ϵ -approximate equilibrium if and only if $x_\epsilon \leq \epsilon$.*

Note that, by our considerations above, ϵ -approximate equilibria can be left again, when minorities start to grow.

We will give our bounds in terms of maximal and optimal latency. Denote the maximum latency by $l_{\max} := \max_{p \in P} l_p(\mathbf{e}_p)$ where \mathbf{e}_p is the unit vector for path p . Let $l^* := \min_{\mathbf{x} \in \Delta} \bar{l}(\mathbf{x})$ be the average latency at a social optimum.

Theorem 1. *The replicator dynamics for general singlecommodity flow networks and non-decreasing latency functions converges to an ϵ -approximate equilibrium within time $\mathcal{O}(\epsilon^{-3} \cdot \ln(l_{\max}/l^*))$.*

This theorem is robust. The proof does not require that agents behave exactly as described by the replicator dynamics. It is only necessary that a constant fraction of the agents that are by a factor of $(1 + \epsilon)$ above the average move to a better strategy. We can also show that our bound is in the right ballpark.

Theorem 2. For any $r := l_{\max}/l^*$ there exists a network and boundary conditions such that the time to reach an ϵ -approximate equilibrium is bounded from below by $\Omega(\epsilon^{-1} \ln r)$.

For the multicommodity case, we can only derive a linear upper bound. If we define $P_\epsilon := \{p | l_p \geq (1 + \epsilon) \cdot \bar{l}_{i(p)}\}$ then the definition of x_ϵ and Definition 3 translate naturally to the multicommodity case. Let $l_i^* = \min_{\mathbf{x} \in \Delta} \bar{l}_i(\mathbf{x})$ be the minimal average latency for commodity $i \in \mathcal{I}$ and let $l^* = \min_{i \in \mathcal{I}} l_i^*$.

Theorem 3. The multicommodity replicator dynamics converges to an ϵ -approximate equilibrium within time $\mathcal{O}(\epsilon^{-3} \cdot l_{\max}/l^*)$.

Our analysis of convergence in terms of ϵ -approximate equilibria uses Rosenthal's potential function [12]. However, this function is not suitable for the proof of convergence in terms of the Euclidian distance since it does not give a general upper bound. Here we use the entropy as a potential function, which in turn is not suitable for ϵ -approximations. Because of our choice of the λ_i in this section, Proposition 3 cannot be translated directly to the multicommodity case.

5 Proofs

First we prove that the relative replicator dynamics is a legal dynamics in the context of selfish routing, i. e., it does not leave the simplex Δ .

Proof (of Proposition 1). We show that for all commodities $i \in \mathcal{I}$ the derivatives \dot{x}_p , $p \in P_i$ sum up to 0.

$$\begin{aligned} \sum_{p \in P_i} \dot{x}_p &= \lambda_i(\mathbf{x}) \left(\sum_{p \in P_i} x_p \bar{l}_i(\mathbf{x}) - \sum_{p \in P_i} x_p l_p(\mathbf{x}) \right) \\ &= \lambda_i(\mathbf{x}) \left(\bar{l}_i(\mathbf{x}) \sum_{p \in P_i} x_p - d_i \cdot d_i^{-1} \sum_{p \in P_i} x_p l_p(\mathbf{x}) \right) \\ &= \lambda_i(\mathbf{x}) (\bar{l}_i(\mathbf{x}) \cdot d_i - d_i \cdot \bar{l}_i(\mathbf{x})) = 0. \end{aligned}$$

Therefore, $\sum_{p \in P_i} x_p = d_i$ is constant. □

Now we prove evolutionary stability.

Proof (of Proposition 2). Let \mathbf{x} be a Nash equilibrium. We want to show that \mathbf{x} is evolutionary stable. In a Wardrop equilibrium all latencies of used paths belonging to the same commodity are equal. The latency of unused paths is equal or even greater than the latency of used paths. Therefore $\mathbf{x} \cdot \mathbf{l}(\mathbf{x}) \leq \mathbf{y} \cdot \mathbf{l}(\mathbf{x})$ for all populations \mathbf{y} . As a consequence,

$$\begin{aligned} \mathbf{y} \cdot \mathbf{l}(\mathbf{y}) &\geq \mathbf{x} \cdot \mathbf{l}(\mathbf{x}) + \mathbf{y} \cdot \mathbf{l}(\mathbf{y}) - \mathbf{y} \cdot \mathbf{l}(\mathbf{x}) \\ &= \mathbf{x} \cdot \mathbf{l}(\mathbf{x}) + \sum_{p \in P} y_p (l_p(\mathbf{y}) - l_p(\mathbf{x})) \\ &= \mathbf{x} \cdot \mathbf{l}(\mathbf{x}) + \sum_{e \in E} y_e (l_e(\mathbf{y}) - l_e(\mathbf{x})). \end{aligned}$$

Consider an edge $e \in E$. There are three cases.

1. $y_e > x_e$. Because of strict monotonicity of l_e , it holds that $l_e(\mathbf{y}) > l_e(\mathbf{x})$. Therefore also $y_e(l_e(\mathbf{y}) - l_e(\mathbf{x})) > x_e(l_e(\mathbf{y}) - l_e(\mathbf{x}))$.
2. $y_e < x_e$. Because of strict monotonicity of l_e , it holds that $l_e(\mathbf{y}) < l_e(\mathbf{x})$. Again, $y_e(l_e(\mathbf{y}) - l_e(\mathbf{x})) > x_e(l_e(\mathbf{y}) - l_e(\mathbf{x}))$.
3. $y_e = x_e$. In that case $y_e(l_e(\mathbf{y}) - l_e(\mathbf{x})) = x_e(l_e(\mathbf{y}) - l_e(\mathbf{x}))$.

There is at least one edge $e \in E$ with $x_e \neq y_e$ and, therefore, $y_e(l_e(\mathbf{y}) - l_e(\mathbf{x})) > x_e(l_e(\mathbf{y}) - l_e(\mathbf{x}))$. Altogether we have

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

which is our claim. Note that this proof immediately covers the multicommodity case. \square

Proof (of Proposition 3). Denote the Nash equilibrium by \mathbf{x} and the current population by \mathbf{y} . We define a potential function $H_{\mathbf{x}}$ by the entropy

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

From information theory it is known that this function always exceeds the square of the Euclidean distance $\|\mathbf{x} - \mathbf{y}\|^2$. We can also write this as $H_{\mathbf{x}}(\mathbf{y}) = \sum_{p \in P} (x_p \ln(x_p) - x_p \ln(y_p))$. Using the chain rule we calculate the derivative with respect to time:

$$\dot{H}_{\mathbf{x}}(\mathbf{y}) = - \sum_{p \in P} x_p \dot{y}_p \frac{1}{y_p}.$$

Now we substitute the replicator dynamics for \dot{y}_p , cancelling out the y_p .

$$\begin{aligned} \dot{H}_{\mathbf{x}}(\mathbf{y}) &= \lambda(\mathbf{y}) \sum_{i \in \mathcal{I}} \sum_{p \in P_i} x_p (l_p(\mathbf{y}) - \bar{l}_i(\mathbf{y})) \\ &= \lambda(\mathbf{y}) \sum_{i \in \mathcal{I}} \left(\sum_{p \in P_i} x_p l_p(\mathbf{y}) - \bar{l}_i(\mathbf{y}) \cdot d_i \right) \\ &= \lambda(\mathbf{y}) \sum_{i \in \mathcal{I}} \left(\sum_{p \in P_i} x_p l_p(\mathbf{y}) - \left(d_i^{-1} \sum_{p \in P_i} y_p l_p(\mathbf{y}) \right) \cdot d_i \right) \\ &= \lambda(\mathbf{y}) \sum_{i \in \mathcal{I}} \sum_{p \in P_i} (x_p - y_p) \cdot l_p(\mathbf{y}) \\ &= \lambda(\mathbf{y}) \cdot (\mathbf{x} - \mathbf{y}) \cdot \mathbf{l}(\mathbf{y}). \end{aligned}$$

Since \mathbf{x} is a Wardrop equilibrium, Proposition 2 implies that $(\mathbf{x} - \mathbf{y}) \cdot \mathbf{l}(\mathbf{y}) < 0$. Furthermore, by our assumption $\lambda(\mathbf{x}) \geq \epsilon > 0$. Altogether this implies that $H_{\mathbf{x}}(\mathbf{y})$, and therefore also $\|\mathbf{x} - \mathbf{y}\|^2$ decreases towards the lower bound of $H_{\mathbf{x}}$, which is 0. \square

We will now proof the bound on the time of convergence in the symmetric case.

Proof (of Theorem 1). For our proof we will use a generalisation of Rosenthal's potential function [12]. In the discrete case, this potential function inserts the agents sequentially into the network and sums up the latencies they experience at the point of time they are inserted. In the continuous case this sum can be generalised to an integral. As a technical trick, we furthermore add the average latency at a social optimum l^* :

$$\Phi(\mathbf{x}) := \left(\sum_{e \in E} \int_0^{x_e} l_e(x) dx \right) + l^*. \quad (4)$$

Now we calculate the derivative with respect to time of this potential Φ . Let L_e be an antiderivative of l_e .

$$\dot{\Phi} = \sum_{e \in E} \dot{L}_e(x_e) = \sum_{e \in E} \dot{x}_e \cdot l_e(x_e) = \sum_{e \in E} \sum_{p \ni e} \dot{x}_p \cdot l_e(x_e).$$

Now we substitute the replicator dynamics (2) into this equation and obtain

$$\begin{aligned} \dot{\Phi} &= \sum_{e \in E} \sum_{p \ni e} (\lambda(\mathbf{x}) \cdot x_p \cdot (\bar{l}(\mathbf{x}) - l_p(\mathbf{x}))) \cdot l_e(x_e) \\ &= \lambda(\mathbf{x}) \sum_{p \in P} \sum_{e \in p} x_p \cdot (\bar{l}(\mathbf{x}) - l_p(\mathbf{x})) \cdot l_e(x_e) \\ &= \lambda(\mathbf{x}) \sum_{p \in P} x_p \cdot (\bar{l}(\mathbf{x}) - l_p(\mathbf{x})) \cdot l_p(\mathbf{x}) \\ &= \lambda(\mathbf{x}) \left(\bar{l}(\mathbf{x}) \sum_{p \in P} x_p l_p(\mathbf{x}) - \sum_{p \in P} x_p l_p(\mathbf{x})^2 \right) \\ &= \lambda(\mathbf{x}) \left(\bar{l}(\mathbf{x})^2 - \sum_{p \in P} x_p l_p(\mathbf{x})^2 \right). \end{aligned} \quad (5)$$

By Jensen's inequality this difference is negative.

As long as we are not at an ϵ -approximate equilibrium, there must be a population share of magnitude at least ϵ with latency at least $(1 + \epsilon) \cdot \bar{l}(\mathbf{x})$. For fixed $\bar{l}(\mathbf{x})$ the term $\sum_{p \in P} x_p l_p(\mathbf{x})^2$ is minimal when the less expensive paths all have equal latency l' . This follows from Jensen's inequality as well. We have $\bar{l} = \epsilon \cdot (1 + \epsilon) \cdot \bar{l} + (1 - \epsilon) \cdot l'$ and

$$l' = \bar{l} \cdot \frac{1 - \epsilon - \epsilon^2}{1 - \epsilon}. \quad (6)$$

According to equation (5) we have

$$\dot{\Phi} = \lambda(\mathbf{x}) \cdot (\bar{l}(\mathbf{x})^2 - (\epsilon \cdot ((1 + \epsilon) \cdot \bar{l}(\mathbf{x}))^2 + (1 - \epsilon) \cdot l'^2)). \quad (7)$$

Substituting (6) into (7) and doing some arithmetics we get

$$\begin{aligned}\dot{\Phi} &= -\lambda(\mathbf{x}) \frac{\epsilon^3}{1-\epsilon} \bar{l}(\mathbf{x})^2 \\ &\leq -\lambda(\mathbf{x}) \cdot \epsilon^3 \cdot \bar{l}(\mathbf{x})^2 / 2 = -\epsilon^3 \cdot \bar{l}(\mathbf{x}) / 2.\end{aligned}\tag{8}$$

We can bound \bar{l} from below by $\Phi/2$:

$$\begin{aligned}\bar{l}(\mathbf{x}) &= \sum_{p \in P} x_p l_p(\mathbf{x}) = \sum_{p \in P} \sum_{e \in p} x_p l_e(x_e) \\ &= \sum_{e \in E} \sum_{p \ni e} x_p l_e(x_e) = \sum_{e \in E} x_e l_e(x_e) \\ &\geq \sum_{e \in E} \int_0^{x_e} l_e(x) dx\end{aligned}\tag{9}$$

The inequality holds because of monotonicity of the latency functions. By definition of l^* , also $\bar{l} \geq l^*$. Altogether we have $\bar{l} + \bar{l} \geq l^* + \sum_{e \in E} \int_0^{x_e} l_e(x) dx$, or $\bar{l} \geq \Phi/2$. Substituting this into inequality (8) we get the differential inequality

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

which can be solved by standard methods. It is solved by any function

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

where $\Phi_{init} = \Phi(0)$ is given by the boundary conditions. This does only hold as long as we are not at a ϵ -approximate equilibrium. Hence, we must reach a ϵ -approximate equilibrium at the latest when Φ falls below its minimum Φ^* . We find that the smallest t fulfilling $\Phi(t) \leq \Phi^*$ is

$$t = 4\epsilon^{-3} \ln \frac{\Phi_{init}}{\Phi^*}.$$

Clearly, $\Phi^* \geq l^*$ and $\Phi_{init} \leq 2 \cdot l_{max}$ which establishes our assertion. \square

Where is this proof pessimistic? There are only two estimates. We will give an example where these inequalities almost hold with equality.

1. Inequality (9) holds with equality if we use constant latency functions.
2. The considerations leading to equation (7) are pessimistic in that they assume that there are always very few agents in x_e and that they are always close to \bar{l} .

Starting from these observations we can construct a network in which we can prove our lower bound.

Proof (of Theorem 2, Sketch). Our construction is by induction. Consider a network with m parallel links numbered 0 through $m-1$. We show that the time of convergence for m links is at least $(m-1) \cdot \Omega(\epsilon^{-1})$.

For $i \in \{1 \dots, m\}$ define the (constant) latency functions $l_i(x) = (1 + c\epsilon)^{-i+1}$, c a constant large enough. Initially, some agents are on the most expensive link 1, very few agents are on links $3, \dots, m$, and the majority is on link 2. More precisely, $x_1(0) = 2\epsilon$, $x_2(0) = 1 - 2\epsilon - \gamma$, and $\sum_{i=3}^m x_i(0) = \gamma$, where γ is some small constant. In the induction step we will define the values of the $x_i(0)$ for $i > 2$. Initially, assume $\gamma = 0$.

First consider the case where $m = 2$. Clearly, the average latency \bar{l} is dominated by link 2, i. e., l_1 is by a factor of at least $(1 + \epsilon)$ more expensive than \bar{l} . This implies that we cannot be at an ϵ -approximate equilibrium as long as $x_1(t) > \epsilon$. Since link 1 is by a factor of $\Theta(1 + \epsilon)$ more expensive than \bar{l} we have $\dot{x}_1 = -\Theta(\epsilon^2)$ and it takes time $\Omega(\epsilon^{-1})$ for x_1 to decrease from 2ϵ to ϵ .

Now consider a network with $m > 2$ edges. By induction hypothesis we know that it takes time $t_m = (m - 2) \cdot \Omega(\epsilon^{-1})$ for the agents to shift their load to link $m - 1$. We carefully select $x_m(0)$ such that it fulfils the following conditions:

- For $t < t_m$ the population share on link m does not lower \bar{l} significantly such that our assumptions on the growth rates still hold.
- For $t = t_m$, $x_m(t)$ exceeds a threshold such that \bar{l} decreases below $l_{m-1}(1 + \epsilon)$.

Although we do not calculate it, there surely exists a boundary condition for $x_m(0)$ having these properties. Because of the second condition, the system exactly fails to enter an ϵ -approximate equilibrium at time t_m . Because of this, we must wait another phase of length $\Omega(\epsilon^{-1})$ for the agents on link $m - 1$ to switch to link m . Note that there may be agents on link $m - 2$ and above. These move faster towards cheaper links, but this does not affect our lower bound on the agents on link $m - 1$.

We have $l_{\max} = 1$ and $l^* = (1 + c\epsilon)^{-m+1}$. For arbitrary ratios $r = l_{\max}/l^*$ we choose $m = \ln(r)$ yielding a lower bound of $\Omega(\epsilon^{-1} \cdot m) = \Omega(\epsilon^{-1} \cdot \ln r)$ on the time of convergence. \square

Proof (of Theorem 3). In the multicommodity case, equation (5) takes the form

$$\dot{\Phi} = \sum_{i \in \mathcal{I}} \lambda_i(\mathbf{x}) \left(\bar{l}_i^2 - \sum_{p \in P_i} x_p l_p(\mathbf{x})^2 \right).$$

Let $i^* = \arg \min_{i \in \mathcal{I}} \bar{l}_i$. In the worst case, all agents in x_ϵ belong to commodity i^* and equation (8) takes the form $\dot{\Phi} \leq -\epsilon^{-3} \bar{l}_{i^*} / 2$. Then Theorem 3 follows directly by the trivial estimate $\bar{l}_{i^*} > l^*$. \square

6 Conclusions and Open Problems

We introduced the replicator dynamics as a model for the dynamic behaviour of selfish agents in networks. For the symmetric case we have given an essentially tight bound for the time of convergence of this dynamics that is polynomial in the degree of approximation and logarithmic in network parameters. For the multicommodity case, we derived an upper bound which is linear in the network

parameters. This model can also be used in the design of distributed load balancing algorithms, since one can reasonably assume that an algorithm based on this model would be accepted by network users.

Several interesting problems remain open:

- The replicator dynamics is based on random experiments performed by agents playing against each other. By this “fluid limit” we can ensure that the outcomes of the experiments meet their expectation values. What happens if we go back to the discrete process?
- How can one improve the upper bound in the multicommodity scenario?
- There is a delay between the moment the agents observe load and latency in the network and the moment they actually change their strategy. What effects does the use of old information have? Similar questions are, e. g., studied in [10].
- A question of theoretical interest is: What is the convergence time for the final Nash equilibrium?

References

1. Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. Pricing network edges for heterogeneous selfish users. In *STOC 2003*, pages 521–530, 2003.
2. Arthur Czumaj and Vöcking Berhold. Tight bounds for worst-case equilibria. In *Proc. 13th SODA*, pages 413–420, San Francisco, 2002.
3. Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. The complexity of pure Nash equilibria, 2004. to appear.
4. D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis. The structure and complexity of Nash equilibria for a selfish routing game. In *Proc. of the ICALP*, pages 123–134, Malaga, Spain, 2002.
5. Martin Gairing, Thomas Luecking, Marios Mavronicolas, and Burkhard Monien. Computing Nash equilibria for scheduling on restricted parallel links. In *Proc. of the STOC 2004*, 2004.
6. Alain B. Haurie and Patrice Marcotte. On the relationship between Nash-Cournot and Wardrop equilibria. *Networks*, 15:295–308, 1985.
7. E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *STACS '99*, 1999.
8. Marios Mavronicolas and Paul G. Spirakis. The price of selfish routing. In *Proc. of STOC 2001*, pages 510–519, 2001.
9. J. Maynard Smith and G. R. Price. The logic of animal conflict. *Nature*, 246:15–18, 1973.
10. Michael Mitzenmacher. How useful is old information? In *Proc. of the PODC 1997*, pages 83–91, 1997.
11. John F. Nash. Equilibrium points in n -person games. In *Proc. of National Academy of Sciences*, volume 36, pages 48–49, 1950.
12. Robert W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
13. Tim Roughgarden and Eva Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.
14. Jörgen W. Weibull. *Evolutionary Game Theory*. MIT press, 1995.