

Evolutionary Game Theory with Applications to Adaptive Routing*

Simon Fischer[†]

Berthold Vöcking[‡]

Abstract

One of the most important problems in large communication networks like the Internet is the problem of routing traffic through the network. Current Internet technology based on the TCP protocol does not route traffic adaptively to the traffic pattern but uses fixed end-to-end routes and adjusts only the injection rates in order to avoid congestion. A more flexible approach uses load-adaptive rerouting policies that reconsider their routing strategies from time to time depending on the observed latencies. In this manuscript, we survey recent results from [1, 2] about the application of methods from evolutionary game theory to such an adaptive traffic management.

Key words: Evolutionary game theory, Wardrop model, adaptive routing, stale information, analysis and dynamics of complex networks

1 Introduction

Recently, game theoretical analysis of the Internet has attracted a growing amount of interest. One of the models studied in this context is the Wardrop model [3] in which each of an infinite number of selfish users controls an infinitesimal amount of traffic that is to be routed through a network. Due to a lack of central coordination, agents strive to minimise their latency selfishly. Classical game theory predicts that agents will assign their traffic according to a Nash equilibrium, i. e., in a way such that no agent has an incentive to change their routing strategy unilaterally. For this model, a number of interesting results have been found. For example, Beckmann et al. [4] shows existence and essential uniqueness of Nash equilibria. Roughgarden and Tardos [5] prove bounds on the so-called *price of anarchy* which is the worst-case ratio between the social welfare (e. g., the average latency) at a Nash equilibrium and the social welfare at an optimal assignment. Other results show how to impose taxes on the agents such that social optimum and Nash equilibrium coincide [6, 7].

Nash equilibria are interesting from a practical point of view as they represent stable and fair allocations. Classical game theory, however, relies on several assumptions that do not seem to be practical. In particular, players are assumed to have full and accurate information about the game and also about the behaviour of the opponents and must act completely rationally. It is questionable whether these assumptions are satisfied when the game under study should model the Internet. Quite obviously, participants in the Internet have incomplete and inaccurate knowledge. Furthermore, they have bounded rationality.

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

[†]Corresponding author; address: RWTH Aachen, Computer Science 1, D-52056 Aachen, Germany; phone +49 241 80 21108; fax +49 241 80 22216; email: fischer@cs.rwth-aachen.de

[‡]RWTH Aachen, Germany, voecking@cs.rwth-aachen.de

A more reasonable assumption might be that players learn how to perform well in the game by experience. This is where evolutionary game theory (see e. g. [8, 9]) comes into play. Here, all of the above assumptions are dropped and instead it is assumed that the game is played repeatedly against random opponents. Over time, agents have the chance to optimise their personal cost by reacting to simple observations. Describing a process in which agents adapt their behaviour to the overall situation based on the observed payoff, we can formulate a system of differential equations. Our hope is that the solution concepts of differential equations like attractors or asymptotically stable rest points coincide with Nash equilibria. This is not the case for games in general. In order to ensure this, a refinement of Nash equilibria, *evolutionary stability*, is required. We can show that for our class of games, Nash equilibria possess this property, thus enabling us to show that load adaptive rerouting policies converge towards Nash equilibria. This provides additional motivation to the well-studied concept of Nash equilibria and the many analyses based on this concept that have been performed so far. In addition, we give bounds on the speed of convergence. One practically very important aspect in adaptive routing is the effect of stale information. In applications, the latency information that rerouting decisions are based on is typically not up to date. It is well known, that basing routing decisions on stale information can cause oscillation effects and seriously harm network performance. In a simplified model, we show how this can be avoided using a class of adaptive rerouting policies that is not too greedy.

2 The Model

2.1 Selfish Routing

For any of a set of *commodities* $i \in \{1, \dots, k\}$, a fraction of r_i agents wants to route an equivalent amount of flow from source s_i to sink t_i using paths from the set of paths \mathcal{P}_i connecting these two nodes. Let $\mathcal{P} = \cup_{i \in [k]} \mathcal{P}_i$ where $[k] = \{1, \dots, k\}$. A *feasible flow* on this network is a positive real-valued vector $(f_P)_{P \in \mathcal{P}}$ that satisfies the flow demands, i. e., $\sum_{P \in \mathcal{P}_i} f_P = r_i$ for all $i \in [k]$. Given a flow vector f , the flow on an edge $e \in E$ is $f_e = \sum_{P \ni e} f_P$. The latencies in the network are specified by strictly increasing functions $\ell_e : [0, 1] \mapsto \mathbb{R}_0^+$. Given a flow vector f , the latency on edge $e \in E$ is given by $\ell_e(f_e)$ and the latency of path $P \in \mathcal{P}$ is given by $\ell_P(f) = \sum_{e \in P} \ell_e(f_e)$. A flow vector f is said to be at a *Nash equilibrium* if for any commodity $i \in [k]$ and every pair of paths $P, P' \in \mathcal{P}_i$ with $f_P > 0$ it holds that $\ell_P(f) \leq \ell_{P'}(f)$. In other words, at a Nash equilibrium, no agent has an incentive to change their routing strategy.

In a scenario where routing is not controlled by a central authority but by individually and selfishly acting agents, the concept of Nash equilibria makes sense. Though the concept of Nash equilibria seems to be a natural and appealing stability concept for the above scenario, it neglects the question of how the agents can come to such a state. In fact, this requires very accurate knowledge about the network topology, latency functions, the demands of other commodities, and the behaviour of the other agents as well as full rationality. These assumptions are quite obviously not satisfied if the network under study is the Internet. The following section will present some very simple processes which, in contrast to the assumptions of classical game theory, require only very local knowledge and almost no computational effort at all and will turn out to result in Nash equilibria.

2.2 Rerouting Dynamics

Consider a large population of agents in a routing network, each agent choosing one of the possible routing paths. If the number of agents in this scenario is infinite, then a vector $(x_P)_{P \in \mathcal{P}}$ specifying the fractions of agents using each individual path $P \in \mathcal{P}$ is actually equivalent to a feasible network flow.

Now assume that every agent wants to optimise their personal latency having virtually no knowledge about the flow and demands of other agents. One reasonable behaviour for an agent would be to reconsider their routing strategy from time to time, at Poisson rates, say, and revise their strategy based on the observed performance. An agent might then pick another routing path at random (e. g., with probability proportional to the population share currently using this path) and compare the latency of the own path with the latency of the other. If the other path turns out to be worse than the current routing strategy, nothing happens. However, if the other path offers an improvement with respect to latency, the agent might switch to the new path with a probability proportional to the size of the latency gain. If we take this process to the fluid limit, i. e., letting the number of agents go to infinity and identifying random variables describing the change of the population shares in one step with their expectation values, we obtain an expression for the change rate of the population shares:

$$\dot{f}_P = \lambda_i \cdot f_P \cdot (\bar{\ell}_i - \ell_P) \quad \text{for } i \in [k], P \in \mathcal{P}_i, \quad (1)$$

where $\dot{\cdot}$ indicates the derivative with respect to time, $\bar{\ell}_i$ is the average latency of commodity i , and λ_i is some factor that accounts for proportionality factors needed to ensure that probabilities do not exceed 1 etc. Note that the solution orbit $\{\xi | \exists t \geq 0 : \xi = f(t)\}$ of this system of differential equations is independent of the scale of the vector $(\lambda_i)_{i \in [k]}$ as long as all $\lambda_i > 0$. Scaling all λ_i by the same factor scales the speed at which orbits are traversed by this factor. Equation (1) has several appealing properties and has therefore been studied extensively in the evolutionary game theory literature. It is known as the *replicator dynamics* (for a survey of this and other dynamics, see, e. g., [8]).

A natural generalisation of this dynamics that preserves the aspect of local control is a class of rerouting policies that consists of two steps. Again, agents are activated at Poisson rates. Once activated, an agent performs two steps:

1. Sampling: Pick a path Q at random with probability σ_Q . In the most simple case we have $\sigma_Q = 1/m$ where m is the number of paths of the agent's commodity. For the replicator dynamics we have $\sigma_Q = f_Q$, i. e., the probability to sample path Q is proportional to the fraction of agents using it.
2. Migration: Migrate from the current path P to path Q with probability $\mu(\ell_P, \ell_Q)$. For the class of *better response dynamics* we have $\mu(\ell_P, \ell_Q) = 1$ if $\ell_Q < \ell_P$ and $\mu(\ell_P, \ell_Q) = 0$ otherwise. For the replicator dynamics we have $\mu(\ell_P, \ell_Q) = \max\{(\ell_P - \ell_Q) \cdot \lambda, 0\}$ where we choose λ small enough such that the probability is bounded from above by 1.

Altogether, we can specify the rate r_{PQ} at which agents move from path P to path Q and finally the time derivatives of the population shares. For all commodities $i \in \{1, \dots, k\}$ and all paths $P, Q \in \mathcal{P}_i$ we have

$$r_{PQ} = f_P \cdot \sigma_Q \cdot \mu(\ell_P, \ell_Q) \quad \text{and} \quad \dot{f}_P = \sum_{Q \in \mathcal{P}_i} (r_{QP} - r_{PQ}). \quad (2)$$

Let us remark that by the Picard-Lindelöf-Theorem [10] a unique solution to this system of differential equations exists if the right-hand sides are Lipschitz continuous. We therefore require the latency functions ℓ_e , $e \in E$ as well as σ and μ to be Lipschitz continuous. However, even for linear latency functions, Equation (1) contains cubic terms thus rendering an analytic solution impossible.

3 Equilibria and Convergence

The first natural question is whether the above dynamics actually converge towards Nash equilibria in the long run. More formally, we want to show that Nash equilibria are global attractors of our system of differential equations. Since the replicator dynamics is not *innovative*, i. e., $f_P(t) = 0$ implies that $f_P(t') = 0$ for all $t' \geq t$, the replicator dynamics can never discover such unused paths. We therefore assume for the rest of the paper that there are no unused links in the initial population.

For the purpose of showing convergence, *evolutionary stability* has been introduced as an equilibrium concept which is stricter than the concept of Nash equilibria. For the single-commodity case, evolutionary stability can be characterised as follows [9].

Definition 1 (evolutionary stable). *A flow vector $f \in \Delta$ is called evolutionary stable iff (1) it is a Nash equilibrium and (2) for all best replies \tilde{f} to f , $\tilde{f} \neq f$ it holds that $\tilde{f} \cdot \ell(\tilde{f}) > f \cdot \ell(\tilde{f})$.*

In our scenario, a *best reply* to a flow vector f corresponds to a flow vector \tilde{f} that uses only minimum latency paths with respect to the latency vector ℓ induced by f . Since Nash equilibria are not in general unique, but only unique with respect to the edge-flows $(f_e)_{e \in E}$, we say that a flow vector f is *essentially evolutionary stable* if condition (2) above holds for all best replies \tilde{f} that differ from f for at least one edge $e \in E$ (instead of one path $P \in \mathcal{P}$). Then we can show the following lemma.

Lemma 1 ([1]). *For single-commodity networks, Nash equilibria are essentially evolutionary stable.*

Given this property, we can prove convergence of the replicator dynamics towards Nash equilibria for the single-commodity case using standard techniques of evolutionary game theory [9]. For the multi-commodity case and for general dynamics of the form of Equation (2), the proof is an application of Lyapunov's second method [10] in conjunction with a Potential function introduced by Beckmann et al. [4].

Theorem 2 ([1, 2]). *In terms of edge flows $(f_e)_{e \in E}$, the replicator dynamics (1) (provided that $f_P(0) > 0$ for all $P \in \mathcal{P}$) and, more generally, all dynamics of the form of Equation (2) (provided that σ_Q is always positive and σ and μ are continuous) converge towards a Nash equilibrium.*

From the computer scientists' perspective we are interested in the time until our dynamics reach equilibria. Clearly, in the continuous fluid limit model, equilibria cannot be reached exactly but merely approximated. Considering the single-commodity case, we define approximate equilibria as follows. Let \mathcal{P}_ϵ be the set of paths that have latency at least $(1 + \epsilon) \cdot \bar{\ell}$, i. e., $\mathcal{P}_\epsilon = \{P \in \mathcal{P} \mid \ell_P(f) \geq (1 + \epsilon) \cdot \bar{\ell}\}$ and let $f_\epsilon := \sum_{P \in \mathcal{P}_\epsilon} f_P$ be the fraction of agents using these paths. A population f is said to be at an ϵ -approximate equilibrium if and only if $f_\epsilon \leq \epsilon$.

Note that by definition of the replicator dynamics, we cannot, in general, expect to reach a state where $f_\epsilon = 0$ since a population share using a path with constant high latency

will never completely vanish. Similarly, it can take arbitrarily long, until many agents are within a factor of $(1 + \epsilon)$ of the minimum latency path since initially there may be an arbitrarily small fraction of agents on this link. Our definition takes care of these two effects.

In order for the time to reach an approximate equilibrium to be meaningful, the dynamics must not depend on the scale by which we measure latency. We use the parameter λ to normalise our dynamics and choose $\lambda = 1/\bar{\ell}$. For single-commodity networks the replicator dynamics with $\lambda = 1/\bar{\ell}$ converges to an ϵ -approximate equilibrium within time $\mathcal{O}(\epsilon^{-3} \cdot \ln(\ell_{\max}/\ell^*))$ where ℓ^* is the optimal average latency and ℓ_{\max} is the maximum latency of a path over all possible flows. Somewhat weaker bounds on the speed of convergence in multi-commodity games can be found in [1] as well.

4 Stale Information

One of the most important problems in load adaptive routing is the fact that information about latency, delay, or bandwidth may be out of date by the time it is gathered. It is well-known that this can cause oscillation effects and seriously harm performance. Mitzenmacher [11] introduced the bulletin board model to study these effects. In this model, all information relevant to the rerouting process is stored on a centralised bulletin board that is accessible to all agents. However, information on the bulletin board is not always up to date but merely updated every T units of time. The bulletin board can be a model for a scenario where latency information is broadcasted to the agents at intervals or where this information is stored on a server from which it can be polled by the agents. Let us remark, that this is a purely theoretical model which, however, exhibits the effects we want to study.

It is clear that two conditions can cause the policy to oscillate. First, if small changes in the flow on an edge can cause a large change in latency, then agents must migrate to this edge very carefully in order not to overshoot the balanced state. Hence, we consider only latency functions of bounded slope, i.e., we consider some β such that $\ell'_e(x) \leq \beta$ for all $e \in E$, $x \in [0, 1]$. Second, agents must not move too fast if the observed latency difference between the two considered paths is small as this could otherwise cause the same effect. We consider a number $\alpha > 0$ such that $\mu(\ell_1, \ell_2) \leq \alpha(\ell_1 - \ell_2)$ for $\ell_1 \geq \ell_2$ and $\mu(\ell_1, \ell_2) = 0$ if $\ell_1 < \ell_2$. We call rerouting policies satisfying this property α -smooth. Finally, let L be the length of the longest path in the network. Given these properties we can show that our policies converge towards a Nash equilibrium provided that the bulletin board is updated frequently enough.

Theorem 3 ([2]). *If the update frequency $1/T \geq 4L\alpha\beta$ and σ_P assigns non-zero probabilities to all paths $P \in \mathcal{P}$, then the solution of dynamics (2) converges towards a Nash equilibrium in the bulletin board model.*

We can also show that the parameters that go into our upper bound are actually necessary in the following sense. We say that a function $\mathbf{x}(\cdot)$ *oscillates* if for some $\tau > 0$ and some t_0 it holds that $\mathbf{x}(t_0) = \mathbf{x}(t_0 + n\tau)$ for all $n \in \mathbb{N}$. A differential equation $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$ *oscillates* if there exists a boundary condition $\mathbf{x}(0) = \mathbf{x}_0$ such that the solution orbit oscillates.

Theorem 4. *For any α , β , and L with $\alpha\beta L \geq 1$ there exist a network in which the longest path has length L , latency functions whose slope is bounded by β , and an α -smooth migration rule such that for any $T > 8/(\alpha\beta L)$ the differential equation related to the dynamics oscillates.*

5 Conclusion

We have argued that the assumptions of full information and unbounded rationality made by classical game theory in order to motivate the central concept of Nash equilibria are not realistic with respect to routing in large networks. Instead, we have described a simple class of rerouting policies that, in the fluid limit, converge towards Nash equilibria. These policies do not rely upon these assumptions but merely on very simple observations and little computation. Thus, we have strengthened the motivation of Nash equilibria.

We have also given upper bounds on the time of convergence towards approximate equilibria. Furthermore, we have studied the effects of stale information which often imposes significant performance degradation on practical applications. We have shown bounds on the necessary update time for the information depending on smoothness parameters of the network such that these effects can be avoided.

References

- [1] Simon Fischer and Berthold Vöcking. On the evolution of selfish routing. In *Proc. 12th Ann. European Symp. on Algorithms (ESA)*, number 3221 in Lecture Notes in Comput. Sci., pages 323–334, Bergen, Norway, September 2004. Springer-Verlag.
- [2] Simon Fischer and Berthold Vöcking. Adaptive routing with stale information. In *Proc. 24th Ann. ACM SIGACT-SIGOPS Symp. on Principles of Distributed Computing (PODC)*. ACM, July 2005.
- [3] John Glen Wardrop. Some theoretical aspects of road traffic research. In *Proc. of the Institute of Civil Engineers, Pt. II*, pages 325–378, 1952.
- [4] M. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics and Transportation*. Yale University Press, 1956.
- [5] Tim Roughgarden and Éva Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.
- [6] Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. How much can taxes help selfish routing? In *Proc. 4th ACM Conference on Electronic Commerce*, pages 98–107, 2003.
- [7] Lisa Fleischer. Linear tolls suffice: New bounds and algorithms for tolls in single source networks. In *Proc. 31st Int. EATCS Coll. on Automata, Languages and Programming (ICALP)*, pages 544–554, 2004.
- [8] Josef Hofbauer and Karl Sigmund. Evolutionary game dynamics. *Bulletin of the American Mathematical Society*, 40(4):479–519, 2003.
- [9] Jörgen W. Weibull. *Evolutionary Game Theory*. MIT press, 1995.
- [10] William Boyce. *Elementary Differential Equations and Boundary Value Problems*. John Wiley & Sons, 1965.
- [11] Michael Mitzenmacher. How useful is old information? In *Proc. 16th Ann. ACM SIGACT-SIGOPS Symp. on Principles of Distributed Computing (PODC)*, pages 83–91, 1997.