Adaptive Routing with Stale Information*

Simon Fischer
RWTH Aachen University
Department of Computer Science
52056 Aachen, Germany
fischer@cs.rwth-aachen.de

Berthold Vöcking
RWTH Aachen University
Department of Computer Science
52056 Aachen, Germany
voecking@cs.rwth-aachen.de

ABSTRACT

We investigate adaptive routing policies for large networks in which agents reroute traffic based on old information. It is a well known and practically relevant problem that old information can lead to undesirable oscillation effects resulting in poor performance. We investigate how adaptive routing policies should be designed such that these effects can be avoided.

The network is represented by a general graph with latency functions on the edges. Traffic is managed by a large number of agents each of which is responsible for a negligible amount of traffic. Initially the agents' routing paths are chosen in an arbitrary fashion. From time to time each agent revises her routing strategy by sampling another path and switching with positive probability to this path if it promises smaller latencies. As the information on which the agent bases her decision might be stale, however, this does not necessarily lead to an improvement. The points of time at which agents revise their strategy are generated by a Poisson distribution. Stale information is modelled in form of a bulletin board that is updated periodically and lists the latencies on all edges.

We analyze such a distributed routing process in the so-called fluid limit, that is, we use differential equations describing the fractions of traffic on different paths over time. In our model, we can show the following effects. Simple routing policies that always switch to the better alternative lead to oscillation, regardless at which frequency the bulletin board is updated. Oscillation effects can be avoided, however, when using *smooth adaption policies* that do not always switch to better alternatives but only with a probability depending on the advantage in the latency. In fact, such policies have dynamics that converge to a fixed point corresponding to a Nash equilibrium for the underlying routing game, provided the update periods are not too large.

In addition, we also analyze the speed of convergence towards approximate equilibria of two specific variants of smooth adaptive routing policies, e. g. , for a replication policy adopted from evolutionary game theory.

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

PODC'05, July 17–20, 2005, Las Vegas, Nevada, USA.

Categories and Subject Descriptors

C.2.m [**Computer Systems Organization**]: Computer-Communication Networks—*Miscellaneous*; F.2.m [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity—*Miscellaneous*

General Terms

Algorithms, Performance, Design, Theory

Keywords

Adaptive routing, stale information, (evolutionary) game theory

1. INTRODUCTION

We study a routing game for networks defined by general graphs where infinitesimally small pieces of traffic are managed by selfish agents each of which aims at minimizing her own latency. In contrast to most previous studies for this kind of game, we do not presume that fully rational behavior of the agents necessarily leads to a so-called Wardrop equilibrium. Instead we study the question of how the agents should behave such that the game (quickly) converges to such an equilibrium. Using well-known potential function arguments known from the analysis of congestion games one can show that simple best or better response policies in which agents move their traffic from slower to faster routing paths lead in fact to a Wardrop equilibrium. This kind of convergence result, however, requires that agents always have available the most recent information about the latencies and can react instantaneously to this information. This assumption, however, seems be problematic having in mind real-world traffic or communication networks.

In this paper, we investigate how agents should behave if they have to base their decisions on possibly stale information. In our analysis, we represent stale information using a variant of the bulletin board model introduced by Mitzenmacher [12] in the context of dynamic load balancing: From time to time agents reconsider their routing strategy and possibly map their traffic to a different routing path. The time intervals between the points at which agents rethink their strategy are exponentially distributed. For each individual agent, the reconsideration points are generated at a certain Poisson rate. At these points of time, the agent can decide to map her traffic to a different path based on information published on a bulletin board. This bulletin board it visible to all agents but it does not always contain the most recent information. The board is updated from time to time at least once every T time steps. Using this admittedly simplistic model of stale information, we observe a phenomenon known also from practical studies, namely that naive best and better response strategies do not converge to an equilibrium but result in undesirable oscillation effects [10, 11, 13, 17]. This motivates us to study the question of how agents should behave in order to avoid these effects. In fact, we can show that for every set of non-decreasing latency functions with finite first derivative there is a policy for the agents to reroute their traffic that avoids oscillation and guarantees convergence to the Wardrop equilibrium.

Our adaptive routing policies are inspired by the so-called replicator dynamics known from evolutionary game theory. We study these policies in the fluid limit evaluating differential equations at discrete time steps. We do not only show that the considered dynamics converge to a Wardrop equilibrium but also analyze the time of convergence towards approximate equilibria.

1.1 Description of the routing problem

We are given a network $G = (V, E)$ with latency functions $\ell_e : [0, 1] \mapsto \mathbb{R}_0^+$, $e \in E$, describing the latency of an edge as a function of the traffic mapped to that edge. The domain is restricted to $[0, 1]$ as the total amount of traffic is normalized to 1. Throughout this paper, we assume that latency functions are continuous, non-decreasing and have finite first derivative on the considered domain.

The routing problem is described by a set of commodities $\mathcal{I}$. For commodity $i \in \mathcal{I}$ there is a traffic demand d_i to be routed from node $s_i \in V$ to node $t_i \in V$. W. l. o. g., $\sum_{i\in I} d_i = 1$. Let P_i be the set of paths connecting s_i and t_i and let $P = \cup_{i\in\mathcal{I}} P_i$. For simplicity of notation, assume that the sets P_i are disjoint. A flow vector $(\mathbf{x})_{p\in P}$ is feasible if $x_p \geq 0$ and $\sum_{p\in P_i} = d_i$ for all $i \in \mathcal{I}$. Then the flow on edge e is $x_e = \sum_{p\ni e} x_p$ and the latency of that edge is $\ell_e = \ell_e(x_e)$. The latency of path p is $\ell_p = \ell_p(\mathbf{x}) = \sum_{e\in p} \ell_e(x_e)$. Let $\ell_{\max}$ denote an upper bound on the latency of a path, e. g. $\max_{p\in P} \sum_{e\in p} \ell_e(1)$.

In the game theoretic model [1, 9, 16, 18], the flow vector $\mathbf{x}$ is interpreted as a population vector, i. e. , we consider an infinite number of agents carrying an infinitesimal load each, and x_p is the fraction of agents sending their load over path p. Agents aim at minimizing their personal latency selfishly without considering the impact on the global situation. One usually assumes that the agents will come to a flow allocation in which no agent can improve her latency by deviating unilaterally. A useful characterization of such an equilibrium goes back to Wardrop [9, 18].

DEFINITION 1 (WARDROP EQUILIBRIUM). *A feasible flow vector* $\mathbf{x}$ *is at a* Wardrop equilibrium *iff for all commodities* $i \in \mathcal{I}$ *and all paths* $p, p' \in P_i$ *with* $x_p > 0$ *it holds that* $\ell_p(\mathbf{x}) \leq \ell_{p'}(\mathbf{x})$.

We are interested in devising simple adaptive routing policies describing how the agents should behave so that, starting from any feasible allocation of routing paths, the system converges to a Wardrop equilibrium. We think that approaching a Wardrop equilibrium is a reasonable objective as such an equilibrium corresponds to a fair allocation of routing paths in the sense that no single agent wants to deviate from this allocation. Recent studies on the price of anarchy for different kinds of latency functions have also shown that a Wardrop equilibrium gives reasonable approximation ratios with respect to the global objective of minimizing average latencies [15, 16]. Let us further remark that the Wardrop equilibrium with respect to virtual latency functions $h_e(x) = x\frac{d\ell_e(x)}{dx} + \ell_e(x)$ minimizes average latency with respect to the original latencies. This shows that our adaptive routing policies applied with virtual latency functions can also be used in order to converge to a routing plan minimizing average latency, provided that $x\frac{d\ell_e(x)}{dx}$ is non decreasing, i. e. , the latency functions ℓ_e are convex.

1.2 Adaptive routing policies

Our main focus in this paper is on a very broad class of adaptive routing policies for the agents that can be described in the following way. Each agent revises her *routing strategy*, i. e. , her routing path, at discrete points of time called the agent's *revision points*. These revision points are generated at a fixed Poisson rate for each agent independently. At each revision point an agent performs the following two steps. Suppose the agents currently uses path $p \in P_i$.

1. *Sampling:* The agent samples a path $q \in P_i$ with probability σ_{pq}.
2. *Migration:* The agent switches to the sampled path q with probability $\mu(\ell_p, \ell_q)$.

Let us first give some examples for possible sampling policy. *Uniform sampling* means that we assume $\sigma_{pq} = 1/|P_i|$ for all $i \in \mathcal{I}$, $p, q \in P_i$. Instead of choosing strategies uniformly at random, one can also imagine that agents choose other agents from the same commodity uniformly at random in order to compare their latencies with the latencies experienced by other agents. This kind of sampling is usually assumed in evolutionary game theory. We call this approach *proportional sampling* as the probability that a strategy is chosen is proportional to the share of the population using this strategy, i. e. , $\sigma_{pq} = x_q$ for all $p, q \in P_i$. In case of proportional sampling, the sampling probability σ_{pq} is a function of the population vector $\mathbf{x}$. There might be other useful generalizations of this kind. In general, let us assume that σ_{pq} is a continuous function of $\mathbf{x}$. Observe that this also covers uniform sampling, in which case σ_{pq} is constant with respect to $\mathbf{x}$.

Similarly, there is a variety of strategies for the migration policy μ. For example, the *better response* policy is defined by

$$\mu(\ell_p, \ell_q) = \begin{cases} 1 & \text{if } \ell_p > \ell_q \\ 0 & \text{otherwise.} \end{cases}$$

All migration probabilities $\mu(\ell_p, \ell_q)$ that we consider in this paper are non-decreasing in the latency difference $\ell_p - \ell_q$. We say that a migration policy is α*-smooth* if the ratio $\mu(\ell_p, \ell_q)$ is upper bounded by $\alpha \cdot (\ell_p - \ell_q)$, and a migration policy is *smooth* if there exists an α such that the policy is α-smooth. (If μ can be written as a function of $\ell_p - \ell_q$ than smoothness corresponds to Lipschitz continuity at 0.) The better response policy is an example for a migration policy that is not smooth. An example for smooth migration policy is the *linear migration policy* $\mu(\ell_p, \ell_q) = \max\left\{\frac{\ell_p - \ell_q}{\ell_{\max}}, 0\right\}$ with

$\alpha = 1/\ell_{\max}$. The linear migration policy in combination with proportional sampling corresponds to the *replicator dynamics* considered evolutionary game theory and analyzed in [8].

The sample and migration probabilities that are applied by every agent at Poisson rate 1 translate into a migration rate r_{pq} at which agents change from path p to path q.

$$r_{pq} = x_p \cdot \sigma_{pq} \cdot \mu(\ell_p, \ell_q).$$

Assuming that the amount of traffic carried by each agent is infinitesimally small, we can now describe the behavior of the population in terms of the so-called *fluid-limit*. Summing up the migration rates for each a p, we obtain a differential equation describing the growth rate (or derivative with respect to time) of the population share for p.

$$\dot{x}_p = \sum_{q \in P} (r_{qp} - r_{pq}) \tag{1}$$

For all policies discussed above the right hand side of this equation is Lipschitz continuous which, for any initial configuration $\mathbf{x}(0)$, guarantees the existence of a unique solution $\mathbf{x}(t)$, $t \geq 0$, to this differential equation. The function $\mathbf{x}(t)$ is the *trajectory* of the population vector over time. Recall that we assumed that the sampling probabilities σ_{pq} are continuous functions of $\mathbf{x} = \mathbf{x}(t)$. In order to ensure that the dynamics can converge to a global equilibrium, we additionally need to assume that $\sigma_{pq} > 0$, for every $t \geq 0$.

Mainly for reasons of comparison, we also consider the *best response policy* which is not based on sampling but shortest paths computations. Here, each agent reconsidering her strategy chooses a *best reply*, i. e. , a probability distribution of shortest path, belonging to her commodity. Denote the set of all possible best replies by $\beta(\mathbf{x})$. Since the shortest path need not be unique, this leads to a differential inclusion instead of a differential equation.

$$\dot{\mathbf{x}} \in \{\mathbf{y} - \mathbf{x} | \mathbf{y} \in \beta(\mathbf{x})\}.$$

We show that this dynamics does not converge towards an equilibrium even for simple networks with two links in the model of stale information introduced in the following section.

1.3 Models of stale information

For the above derivation of the fluid limit, we assumed that the information on which the agents base their decisions is always up to date. In order to incorporate the effects due to stale information, we enhance the model as follows. We use a bulletin board model similar to the one described by Mitzenmacher in [12]. All information relevant to the dynamics is posted on a bulletin board at the beginning of a phase of fixed length T. We strongly emphasize that this is a purely theoretical model that should reflect the effect of stale information in an as simple as possible way. Similar effects might occur in networks where latency information is broadcasted from time to time or uploaded to a server from which it can be polled by clients.

The edge latencies ℓ_e, $e \in E$, are posted on the bulletin board. We assume that the selection rule is based on the latency information on the bulletin board instead of the true latencies. This way, our differential equation for $\dot{\mathbf{x}}(t)$ now depends not only on $\mathbf{x}(t)$ but also on $\mathbf{x}(\hat{t})$ where $\hat{t}$ marks the beginning of the respective phase. This may lead to discontinuities of $\dot{\mathbf{x}}$ at those points of time when the bulletin board is updated. Given that the right hand side of the differential equation (1) is Lipschitz continuous there exists a unique solution in each update period and, hence, on $\mathbb{R}_0^+$ as well.

In contrast to the migration probabilities, the sampling probabilities do not depend on the latencies, but possibly on $\mathbf{x}$. If this the case then the information on $\mathbf{x}$ may be also posted on the bulletin board and thus stale. Our analysis covers both current and stale information about $\mathbf{x}$.

1.4 Summary of our results

As a first result we show that any routing policy following the rules for sampling and migration based on up-to-date information as described in Section 1.2 defines a dynamics converging to a Wardrop equilibrium. This result holds for the best response policy as well. In fact, given up-to-date information, the best and better response policies seem to be the most natural and effective policies. Assuming the bulletin board model with any update period T, however, neither the best nor the better response policy converge to an equilibrium but both policies show oscillation effects even in a simple network consisting only of two parallel links.

This raises the question whether there exist adaptive routing policies whose dynamics in the bulletin board model converge to a Wardrop equilibrium. We can answer this question positively by proving a sufficient condition ensuring convergence provided that the update periods for the bulletin board are sufficiently small. The condition is that the migration policy needs to be smooth. Suppose the adaption policy is α-smooth, the first derivative of the latency functions is upper-bounded by β, and the length of the routing paths is at most L. We show that it is sufficient to update the bulletin board at a frequency $\frac{1}{T} = \mathcal{O}(\alpha\,\beta\,L)$ in order to ensure convergence to the Wardrop equilibrium. Observe that the product $\alpha\,\beta$ is independent of the scaling of the latency functions.

Furthermore, we study the performance of some specific variants of smooth routing policies. We say that the routing allocation is *approximately balanced* at the beginning of an update period of the bulletin board if the latency of almost every agent is close to the average or, alternatively, best latency within the agent's commodity in some well defined sense. We say that an allocation is *unbalanced* if it is not approximately balanced. Let us remark that the allocation might become approximately balanced at some point of time and later lose this property again. Our analysis, however, upper-bounds the total number of unbalanced periods. We show that the linear migration policy in combination with uniform sampling admits only $\mathcal{O}(\max_{i\in I}\{|P_i|\}\,\alpha\,\beta\,L\,\ell_{\max}^2)$ unbalanced periods, when choosing $T = \Theta(1/(L\,\alpha\,\beta))$. If the number of alternative routing paths is large then this bound might not be satisfying. We show that a different sampling rule can guarantee a much better bound. The replication policy, i. e. , the linear migration policy with proportional sampling, admits only $\mathcal{O}(\alpha\,\beta\,L\,\ell_{\max}^2)$ unbalanced periods.

1.5 Related work

Mitzenmacher [12] studies a load balancing scenario where jobs are assigned to machines based on old information. The main technical differences between his and our work are the following. In Mitzenmacher's model both the number of agents (jobs) and the number of resources (machines) go to infinity, but the ratio between these numbers is finite. In contrast, in our model the set of resources (the considered network) is finite, but the number of agents is infinite. Furthermore, in Mitzenmacher's model, agents enter the system at Poisson rates, allocate their load to a server, and are removed from the system as soon as their task is processed whereas in our model, tasks are permanent and agents only reassign the load. Mitzenmacher comes to the conclusion that placing decisions based on old information can degrade the performance even in comparison to fully random decisions. This phenomenon is even more extreme when the used information is global. Also in our model using old information can lead to a degradation of the performance. We can show, however, that the system can come to a perfect equilibrium solely based on old information.

Dahlin [6] shows that using stale load information in combination with its age and thus employing a greater computational effort can help to improve network performance. In our scenario, the information about the age of the information is not available. Cao and Nyberg [5] present an approach which routes the jobs to the machine with expected shortest queue, but their result is also negative. They show that this does not always minimize the average waiting time.

The replication policy considered in this manuscript has also been studied in [8]. It was shown that this policy quickly converges towards a Wardrop equilibrium provided the agents always have up-to-date information about the latencies. The effect of stale information has not been considered in that work. Comparing the time of convergence towards approximately balanced states, we see that the dynamics considered in [8] must be slowed down by a factor of $L\,\alpha\,\beta\,\ell_{\max}$ in order to guarantee convergence.

Bertsekas and Tsitsiklis [3] consider a distributed algorithm for optimizing a global cost function defined on a network flow with asynchronous agents. As has been described at the end of Section 1.1, the problems of optimizing a global cost function and finding a Wardrop equilibrium can be cast into one other. The main difference between their and our work is that they consider a specific distributed algorithm which is comparable to a careful best response rerouting policy in our scenario whereas we consider a broad class of very simple rerouting policies. Their algorithm requires much more computational effort than our simple sampling policies. Also, in the model of Bertsekas and Tsitsiklis, traffic is rerouted at discrete time steps. In contrast to this, our model is continuous. Furthermore, we also give bounds on the update frequency depending on network parameters and consider the time of convergence.

2. PRELIMINARIES

When information is always up-to-date, convergence towards Wardrop equilibria can easily be shown using a potential function used by Beckmann [1] in order to prove existence of flows at equilibrium.

$$\Phi = \sum_{e\in E}\int_0^{x_e} \ell_e(x)\,dx.$$

This potential function precisely absorbs progress. Whenever an infinitesimal agent moves from path p to path q thus improving her latency by $(\ell_p - \ell_q)$, the value of the potential changes by precisely $(\ell_q - \ell_p)\,dx$. Hence, a population $\mathbf{x}$ minimizes this potential function iff no agent can improve her situation unilaterally, i. e., $\mathbf{x}$ is a Wardrop equilibrium.

PROPOSITION 1. *Assume that the functions ℓ_e are strictly increasing for all $e \in E$. Then the dynamics described by equation (1) converges towards a Wardrop equilibrium.*

PROOF. The proof is an application of Lyapunov's second method (see, e. g. [4]). Consider the potential function introduced by Rosenthal in a classical paper [14].

$$\Phi = \sum_{e\in E}\int_0^{x_e} \ell_e(x)\,dx.$$

The Wardrop equilibrium can be described by a convex program [1] with Φ as the objective function. If the ℓ_e are strictly increasing, then this program has a unique solution in the x_e variables and therefore, Φ is positive definite up to a translation and subtraction of $\Phi_{\min}$. For the derivative with respect to time we have

$$\dot\Phi \quad = \quad \sum_{e\in E}\dot x_e \ell_e(x_e) = \sum_{p\in P}\dot x_p \ell_p(\mathbf{x}) = \sum_{p,q\in P} r_{pq}(\ell_q - \ell_p).$$

We see that r_{pq} is positive if $\ell_q - \ell_p$ is negative and zero otherwise. Hence, every term is negative and $\dot\Phi < 0$ as long as we are not at a Wardrop equilibrium. Therefore, $\dot\Phi$ is also negative-definite (up to the same translation).

By continuity of the ℓ_e, Φ is continuous and has continuous first-order partial derivatives. Furthermore $\dot\Phi$ is continuous by continuity of σ and μ. Thus, Φ fulfills all requirements for a Lyapunov function [4]. Every vector $\dot{\mathbf{x}}$ points strictly into every contour set of Φ. Therefore, all solutions converge towards the Wardrop equilibrium. □

A similar result holds for the best response dynamics, when information is up to date. However, the following example shows that best response and better response dynamics perform poorly in the bulletin board model with update period T. We consider a simple network with two links 1 and 2 with latency functions $\ell_1(x) = c$, $c \in (0, 1)$ and $\ell_2(x) = x^d$. The population $(1 - x^*, x^*)$ with $x^* := c^{1/d}$ is a Wardrop equilibrium. We are mainly interested in the population share on link 2 which we denote by x. Then there are $1 - x$ agents on link 1. The solution to the differential inclusion is unique and is

$$x(t) = \begin{cases} x(0)\cdot e^{-t} & \text{if } x(0) > x^* \\ 1-(1-x(0))\cdot e^{-t} & \text{if } x(0) < x^*. \end{cases}$$

Given that $x^* \in (\alpha, \beta)$, with $\alpha = \frac{1}{e^T+1}$ and $\beta = \frac{e^T}{e^T+1}$, and starting with $x(0) = \alpha$ we oscillate between α and β, i. e., $x(2Tn) = \alpha$, $x(2T(n + 1)) = \beta$ for $n \in \mathbb{N}$. This yields a maximum latency of $\ell_{\max} = (x(\beta))^d = e^{T\cdot d}/(e^T + 1)^d$, which approaches 1 as d approaches ∞. If we choose $c = 1/(e^T + 1)^d + \delta$ then $x^* = c^{1/d} = \alpha + \delta'$ where we can make δ' arbitrarily small by choosing δ arbitrarily small. In order to reach a state where $\ell_{\max} \le (1 + \epsilon)c$ we must have

$$T \le \frac{\ln(1+\epsilon)}{d} = \mathcal{O}(\epsilon/d)$$

This shows that the best response dynamics cannot guarantee convergence however small we choose the update period T. Note that for this simple network, all better response dynamics and best response dynamics behave roughly in the same way.

3. CONVERGENCE UNDER STALE INFORMATION

In this section we will establish the main result of this paper. We will give an upper bound for the update period T such that our policies still converge. We cannot expect all policies of the form of equation (1) to converge in the bulletin board model if we allow arbitrarily steep latency functions. Also we cannot expect convergence if the migration rule μ is very sensitive to small latency differences.

We show that if the steepness of latency functions and of the migration function μ is bounded and if the update period T is chosen sufficiently small (depending on these parameters), we can guarantee convergence in the bulletin board model. The following lemma shows that the actual potential gain in one round is at least one half of the virtual potential gain "observed" by the agents based on stale information and is non-positive.

LEMMA 2. *Assume that the following three conditions hold.*

1. *the slope of the edge latency functions $\ell_e(\cdot)$ is bounded by β for all $e \in E$,*
2. *the migration rule μ is α-smooth, i. e., $\mu(\ell_1, \ell_2) \leq \alpha \cdot (\ell_1 - \ell_2)$ for all $\ell_1 > \ell_2$, and*
3. *the length of all paths $p \in P$ is bounded by L.*

Then define $T = 1/(4\,L\,\beta\,\alpha)$ and consider a phase beginning at time t with an update of the bulletin board and ending at time $t+\tau$, $\tau \leq T$. For any solution of a routing policy of the form of equation (1) it holds that

$$\Delta\Phi = \Phi(t+\tau) - \Phi(t) \leq \frac{1}{2}\sum_{e\in E}(x_e(t+\tau) - x_e(t)) \cdot \hat{\ell}_e \leq 0$$

while the latter inequality is strict unless $\mathbf{x}(t)$ is a Wardrop equilibrium.

PROOF. In the following, we omit the index t and write $x_e = x_e(t+\tau)$ and $\hat{x}_e = x_e(t)$. Similarly, we add a ˆ to all symbols referring to the stale information that was valid at the beginning of the phase. The potential change within one phase is

$$\Delta\Phi = \Phi - \hat{\Phi} = \sum_{e\in E}\int_{\hat{x}_e}^{\hat{x}_e+\Delta x_e} \ell_e(x)dx.$$

In order to see the claim we rearrange $\Delta\Phi$ in a suitable manner. Consider the agents moving from path p to path q within one phase. Denote this fraction of agents by Δx_{pq}. These agents "see" a virtual potential gain of

$$\Delta\Phi_{pq} = -\Delta x_{pq} \cdot \Delta\ell_{pq}$$

where $\Delta\ell_{pq} = \hat{\ell}_p - \hat{\ell}_q$. By definition, $\Delta x_{pq} \geq 0$ with strict inequality if $\Delta\ell_{pq} > 0$ and thus $\Delta\Phi_{pq} \leq 0$ for all $p, q \in P$. As long as we are not at a Wardrop equilibrium, a pair p, q with $\Delta\ell_{pq} > 0$ exists and the inequality is strict. Unfortunately, this information is not up to date and the true latency difference may be smaller. The "error" that all the agents make on edge e altogether is

$$\Delta\Phi_e = \int_{\hat{x}_e}^{\hat{x}_e+\Delta x_e} (\ell_e(x) - \hat{\ell}_e)\,dx.$$

By monotonicity of ℓ_e, $e \in E$, this error term is positive for all $e \in E$. In the following we show that our intuitive partitioning of the potential into observed potential gains $\Delta\Phi_{pq}$ and error terms $\Delta\Phi_e$ correctly sums up to the true potential gain $\Delta\Phi$.

$$\begin{aligned}
\Delta\Phi &= \sum_{e\in E}\int_{\hat{x}_e}^{\hat{x}_e+\Delta x_e} \ell_e(x)\,dx \\
&= \sum_{e\in E}\left(\int_{\hat{x}_e}^{\hat{x}_e+\Delta x_e} (\ell_e(x) - \hat{\ell}_e)\,dx + \Delta x_e \cdot \hat{\ell}_e\right) \\
&= \sum_{e\in E}\left(\Delta\Phi_e + \sum_{p\in P, q\ni e}(\Delta x_{pq} - \Delta x_{qp}) \cdot \hat{\ell}_e\right).
\end{aligned}$$

Now observe that the term $\Delta x_{pq}\hat{\ell}_e$ occurs positively for all $e \in q$ and negatively for all $e \in p$. Therefore, we obtain

$$\begin{aligned}
\Delta\Phi &= \sum_{e\in E}\Delta\Phi_e + \sum_{p,q\in P}\Delta x_{pq}\cdot\left(\sum_{e\in q}\hat{\ell}_e - \sum_{e'\in p}\hat{\ell}_{e'}\right) \\
&= \sum_{e\in E}\Delta\Phi_e + \sum_{p,q\in P}\Delta x_{pq}\cdot(\hat{\ell}_q - \hat{\ell}_p) \\
&= \sum_{e\in E}\Delta\Phi_e + \sum_{p,q\in P}\Delta\Phi_{pq} \\
&= \sum_{e\in E}\Delta\Phi_e + \frac{1}{2}\sum_{p,q\in P}\Delta\Phi_{pq} + \frac{1}{2}\sum_{p,q\in P}\Delta\Phi_{pq}. \qquad (2)
\end{aligned}$$

We have seen that the last term is non-positive and negative as long as the phase does not start at a Wardrop equilibrium. Hence, in order to prove the claim, it is sufficient to show that the first two terms in equation (2) are negative. In order to see that this is true, we distribute at most half of the potential gain of the $\Delta\Phi_{pq}$ over all edges $e \in p \cup q$. Since these may be at most $2L$ edges, we can safely credit each edge $e \in p \cup q$ with a potential gain of $\Delta\Phi_{pq}/(4L)$.

$$\Delta\Phi \leq \sum_e\left(\Delta\Phi_e + \sum_{p,q\in P}\Delta\Phi_{pq}^e\right) + \frac{1}{2}\sum_{p,q\in P}\Delta\Phi_{pq},$$

where $\Delta\Phi_{pq}^e = \Delta\Phi_{pq}/4L$ if $e \in p$ or $e \in q$. By our considerations above, the $\Delta\Phi_{pq}^e$ terms exhaust at most one half of each $\Delta\Phi_{pq}$. The situation is depicted in Figure 1.

Let $\Delta\ell_{\max} = \max_{p,q}\Delta\ell_{pq}$. We focus our attention on a pair p, q where $\Delta\ell_{pq} = \Delta\ell_{\max}$. Now, we consider an arbitrary link $e \in q$ with observed latency $\hat{\ell}_e$. If $\Delta x_e < 0$ then the link is even cheaper than expected which even increases the potential gain. Therefore, let $\Delta x_e > 0$. Then there must be at least a fraction of Δx_e agents migrating to a path containing e. (There may be even more, since other agents may also move away from link e.) In order to give an upper bound on Δx_e it suffices to count the agents migrating to a path containing edge e and ignoring the ones leaving edge e.

$$\begin{aligned}
\Delta x_e &= \int_t^{t+\tau}\dot{x}_e(u)\,du \\
&\leq \int_t^{t+\tau}\sum_{p\in P}\sum_{q\ni e}x_p(u)\cdot\sigma_{pq}(\mathbf{x}(u))\cdot\mu(\ell_p, \ell_q)\,du \\
&\leq \int_t^{t+\tau}\sum_{p\in P}\sum_{q\ni e}\max_{v\in[t,t+\tau)}\{x_p(v)\cdot\sigma_{pq}(\mathbf{x}(v))\}\cdot \\
&\qquad\qquad \mu(\ell_p, \ell_q)\,du \\
&\leq T\cdot\alpha\cdot\Delta\ell_{\max}.
\end{aligned}$$

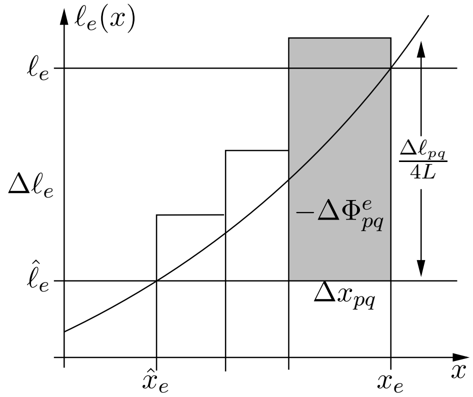

Figure 1: The plot shows the function $\ell_e(\cdot)$ as a function of the edge load x. The area under ℓ_e in the range $[\hat{x}_e, x_e]$ is its true contribution to the potential difference. The area between $\hat{\ell}_e$ and the graph of $\ell_e(\cdot)$ in the range $[\hat{x}_e, x_e]$ is the error term $\Delta\Phi_e$. The gray area is the (negative) observed potential gain $-\Delta\Phi^e_{pq}$ for the agents moving from p to q.

In the last step we used that $\ell_p - \ell_q \leq \Delta\ell_{\max}$, our assumption on the steepness of μ, $\sigma_{pq} \leq 1$, $\sum_p x_p = 1$, and $\tau \leq T$. Altogether, this yields the lower bound

$$\Delta\ell_{\max} \geq \Delta x_e/(T \cdot \alpha). \quad (3)$$

Furthermore, since the slope of ℓ_e is bounded by β, this implies that for the true latency difference we have

$$\Delta\ell_e = \ell_e(x_e) - \ell_e(\hat{x}_e) \leq \Delta x_e \cdot \beta \quad \text{and} \quad \Delta x_e \geq \Delta\ell_e/\beta. \quad (4)$$

By our choice of e, at least for one pair p, q, $q \ni e$ we have $\Delta\ell_{pq} = \Delta\ell_{\max}$. Comparing $-\Delta\Phi^e_{pq}$ with the respective segment of $\Delta\Phi_e$ of width Δx_{pq} shows that the first is greater than the second:

$$\begin{aligned} -\Delta\Phi^e_{pq} &= \frac{\Delta\ell_{\max} \cdot \Delta x_{pq}}{4L} \\ &\geq \frac{\Delta x_e \cdot \Delta x_{pq}}{4L \cdot T \cdot \alpha} \qquad \text{(eq. (3))} \\ &\geq \frac{\Delta\ell_e \cdot \Delta x_{pq}}{4L \cdot T \cdot \alpha \cdot \beta} \qquad \text{(eq. (4))} \\ &= \Delta\ell_e \cdot \Delta x_{pq}. \end{aligned}$$

The last equality holds by our choice of T. We see that we can cover the rightmost segment of $\Delta\Phi_e$ by an appropriate $\Delta\Phi^e_{pq}$.

A similar argument holds for all edges $e' \in p$ with $\Delta x_e < 0$. This is sufficient to show that the true potential gain of the agents moving from p to q is still negative. We now remove these agents from the population and repeat the above analysis to obtain the same result subsequently for all such pairs $p, q \in P$. Altogether we have shown that half of the $\Delta\Phi_{pq}$ suffice to annihilate the error terms $\Delta\Phi_e$ and hence

$$\Delta\Phi \leq \frac{1}{2} \sum_{p,q \in P} \Delta\Phi_{pq} = \frac{1}{2} \sum_{e \in E} (x_e - \hat{x}_e)\hat{\ell}_e$$

which is our claim. The calculation involved here is the reverse of the calculation leading to equation (2). □

This lemma can be used to show global convergence towards a Wardrop equilibrium.

Theorem 3. *If the conditions from Lemma 2 are met, and if furthermore σ_{pq} assigns non-zero probabilities to all paths and ℓ_e is a strictly increasing function for all $e \in E$, then the solution of dynamics (1) converges towards a Wardrop equilibrium in the bulletin board model.*

Proof. As in the proof of Proposition 1, this is an application of Lyapunov's second method [4]. However, now $\dot{\Phi}$ depends not only on $\mathbf{x}$, but also on $\hat{\mathbf{x}}$. Lyapunov's second method can be generalized to differential equations with time delays [2, 7] using the accumulated virtual potential gain as a Lyapuonv functional. We give a simplified proof here. It is an implication of Lemma 2 that if a phase starts at $\Phi(t) = \phi$, then $\Phi(t') \leq \phi$ for all $t' > t$. Note however, that this does not directly imply convergence.

Let us for a moment switch to the model of current information. Let ϕ^* denote the potential of the Wardrop equilibrium. Fix an arbitrary $\phi > \phi^*$. We have to show that the potential function finally takes the value ϕ at the beginning of a phase. Consider the set C_ϕ of points with potential at most ϕ. The derivative of Φ is continuous in $\mathbf{x}$, non-positive and takes the value 0 only at a Wardrop equilibrium. The set of Wardrop equilbria and its neighborhood is contained in C_ϕ. Consequently, the derivative of Φ at any point not in C_ϕ is negative and can be bounded from above by some $\epsilon_\phi < 0$.

For the bulletin board model, Lemma 2 implies that the decrease in the potential is at least $k_0\,\epsilon_\phi\,T$ for a suitable constant $k_0 > 0$. This implies the theorem. □

Remark 1. *The maximal slope of the latency function might be an unusual measure for its steepness. In the literature a more common measure is the degree of the polynomial growth. Assume that for all $e \in E$ the growth of the latency functions is polynomially bounded, i. e., $\ell_e(x \cdot (1 + \epsilon)) \leq \ell_e(x) \cdot (1 + \epsilon)^d$ for all $\epsilon > 0$ and some d and that ℓ_e is convex. The steepest function fulfilling these requirements and $\ell(1) \leq \ell_{\max}$ is $\ell(x) = \ell_{\max} \cdot x^d$, which has maximum slope $d \cdot \ell_{\max}$ at $x = 1$. Therefore, restricting the growth polynomially to degree d implies restricting its slope to $d \cdot \ell_{\max}$.*

Remark 2. *Lemma 2 also enables us to transfer bounds on the speed of convergence from the model of current information to the model of stale information. Consider a phase starting at time t and ending at time $t + \tau$. For the round-based model, $\Delta x_e = \tau \cdot \dot{x}_e(t)$. For the bulletin board model (both with current and stale information on the population shares), x_p can grow or shrink at most exponentially, and thus by at most a constant factor in constant time. For $\tau \leq 1$ this implies that also $\Delta\Phi \leq c \cdot \tau \cdot \dot{\Phi}(t)$ for some constant c. Hence, bounds on the speed of convergence that hold in the model with current information and are based on estimates for $\dot{\Phi}$ also hold in the bulletin board model up to a constant factor.*

4. PERFORMANCE ISSUES

4.1 Uniform sampling

We now consider the uniform sampling rule and the linear migration rule, i. e., $\sigma_{pq} = 1/|P|$ for all $p, q \in P$ and $\mu(\ell_1, \ell_2) = (\ell_1 - \ell_2)/\ell_{\max}$. We will give an upper bound for the time to reach approximate equilibria.

DEFINITION 2 (((δ,ϵ)–APPROXIMATE EQUILIBRIUM). *An agent is* δ-unsatisfied *when it uses a path* $p \in P_i$ *with* $\ell_p > \ell_{\min,i} + \delta$ *where* $\ell_{\min,i} := \min_{q \in P_i}\{\ell_q\}$. *A population* $\mathbf{x}$ *is said to be at a* (δ,ϵ)-approximate equilibrium *if at most* ϵ *agents are* δ*-unsatisfied.*

A stronger definition that would require *all* paths to be by an additive term of at most δ more expensive than the cheapest path would be useless, since a strategy with constant, but very high latency will never get completely extinct. In the following we show fast convergence towards a (δ,ϵ)-approximate equilibrium. Let $m = \max_{i\in I} |P_i|$.

THEOREM 4. *Assume that the bulletin board is updated at intervals of length* $T = 1/(4L \cdot \alpha \cdot \beta)$. *For the uniform sampling policy with linear migration policy the number of update periods not starting in a* (δ,ϵ)*-approximate equilibrium is bounded from above by*

$$\mathcal{O}\left(\frac{m \cdot L \cdot \alpha \cdot \beta}{\epsilon} \cdot \left(\frac{\ell_{\max}}{\delta}\right)^2\right).$$

PROOF. We first consider the potential gain of the agents at the beginning of a phase starting at time t. All agents migrating from one path to another add a negative contribution to the potential change. Consider a δ-unsatisfied agent on commodity i. At least with probability $1/|P_i| \geq 1/m$ this agent samples a cheapest path with respect to her own commodity i. She migrates to this path with probability at least $((\ell_{\min,i} + \delta) - \ell_{\min,i})/\ell_{\max}$ since her latency gain is at least δ. As long as we are not at a (δ,ϵ)-approximate equilibrium, there are at least ϵ such agents at the beginning of the phase. Throughout the phase, the fraction of δ-unsatisfied agents may decrease. However, for all paths it holds that $-x_p \leq \dot{x}_p \leq x_p$ and hence each population share decreases at most exponentially fast. After time T the fraction of δ-unsatisfied agents is still at least $\epsilon \cdot \exp(-T)$. We can assume that $T \leq 1$. Altogether, the rate of agents switching their route to a path with minimum latency and gaining a latency difference of at least δ is at least

$$\frac{\exp(-1)\,\epsilon\,\delta}{m\,\ell_{\max}}.$$

We consider the virtual potential gain of the phase as used in Lemma 2:

$$\begin{aligned}\Delta\tilde{\Phi} &= \sum_{e\in E}(x_e(t+T) - x_e(t)) \cdot \hat{\ell}_e \\ &= \sum_{p,q\in P} \Delta x_{pq} \cdot (\hat{\ell}_q - \hat{\ell}_p),\end{aligned}$$

where Δx_{pq} is again the fraction of agents migrating from path p to path q within one phase. By P_i^δ denote the set of expensive paths of commodity i, i. e. , paths with latency δ above the minimum of their commodity. By P_i^* denote the set of paths with minimal latency in P_i. Since Δx_{pq} is positive iff $(\hat{\ell}_q - \hat{\ell}_p)$ is negative and zero otherwise, we have

$$\Delta\tilde{\Phi} \leq \sum_{i\in I} \sum_{p\in P_i^\delta} \sum_{q \in P_i^*} x_{pq} \cdot \delta$$

and by our considerations above

$$\Delta\tilde{\Phi} \leq T\,\frac{\exp(-1)\,\epsilon\,\delta^2}{m\,\ell_{\max}}.$$

By Lemma 2, $\Delta\Phi \leq \Delta\tilde{\Phi}/2$ and since Φ is bounded from above by $\ell_{\max}$, we reach the trivial lower bound $\Phi = 0$ after at most

$$\mathcal{O}\left(\frac{m}{T\epsilon}\frac{\ell_{\max}^2}{\delta^2}\right)$$

phases which is our claim. □

4.2 Proportional sampling

We now consider the uniform sampling rule and the linear migration rule, i. e. , $\sigma_{pq} = x_q$ for all $p, q \in P$ and $\mu(\ell_1, \ell_2) = (\ell_1 - \ell_2)/\ell_{\max}$. This dynamics corresponds to the replicator dynamics considered in evolutionary game theory. We will see that using this policy we can get rid of the linear term m (which can actually be exponential in the network size) at the cost of a weaker definition of approximate equilibria.

DEFINITION 3 (WEAK (δ,ϵ)-APPROXIMATE EQUILIBRIUM). *An agent is* δ-unsatisfied *when it uses a path* $p \in P_i$ *with* $\ell_p > (1+\epsilon)\bar{\ell}_i$ *where* $\bar{\ell}_i = \sum_{p\in P_i}(x_p/d_i)\ell_p$ *is the average latency of commodity* i. *A population* $\mathbf{x}$ *is said to be at a* (δ,ϵ)-approximate equilibrium *if at most* ϵ *agents are* δ*-unsatisfied.*

THEOREM 5. *Assume that the bulletin board is updated at intervals of length* $T = 1/(4L \cdot \alpha \cdot \beta)$. *For the proportional sampling policy with linear migration policy the number of update periods not starting in a weak* (δ,ϵ)*-approximate equilibrium is bounded from above by*

$$\mathcal{O}\left(\frac{L \cdot \alpha \cdot \beta}{\epsilon} \cdot \left(\frac{\ell_{\max}}{\delta}\right)^2\right).$$

PROOF. We use similar argument here. Consider a δ-unsatisfied agent in commodity i. The expected latency of the strategy she samples is $\bar{\ell}_i$. The probability to migrate to this path there is at least $((\bar{\ell}_i + \delta) - \bar{\ell}_i)/\ell_{\max}$ and the latency gain is at least δ. Using the same arguments as in the proof of Theorem 4, these probabilities yield a potential gain in the of at least

$$\Delta\Phi \leq -\Omega\left(T\epsilon\frac{\delta^2}{\ell_{\max}}\right)$$

within one phase. The rest of the proof is analog to the proof of Theorem 4. □

5. REFERENCES

[1] M. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics and Transportation*. Yale University Press, 1956.

[2] Richard Ernest Bellman and Kenneth L. Cooke. *Differential-Difference Equiations*, volume 6 of *Mathematics in Science and Engineering*. Academic Press, 1963.

[3] Dimitri P. Bertsekas and John N. Tistiklis. *Parallel and Distributed Computing: Numerical Methods*. Athena Scientific, 1989.

[4] William Boyce. *Elementary Differential Equations and Boundary Value Problems*. John Wiley & Sons, 1965.

[5] Jianhua Cao and Christian Nyberg. An approximate analysis of load balancing using stale state information for servers in parallel. In *Proc. 2nd IASTED Int. Conf. on Communications Internet and Information Technology*, November 2003.

[6] Michael Dahlin. Interpreting stale load information. *IEEE Trans. Parallel Distrib. Syst.*, 11(10):1033–1047, 2000.

[7] Rodney D. Driver. Existence and stability of solutions of a delay-differential system. *Arch. Rational Mech. Anal.*, 10:401–426, 1962.

[8] Simon Fischer and Berthold Vöcking. On the evolution of selfish routing. In *Proc. 12th Ann. European Symp. on Algorithms (ESA)*, Lecture Notes in Comput. Sci. 3221, pages 323–334. Springer-Verlag, September 2004.

[9] Alain B. Haurie and Patrice Marcotte. On the relationship between Nash-Cournot and Wardrop equilibria. *Networks*, 15:295–308, 1985.

[10] Atul Khanna and John A. Zinky. The revised ARPANET routing metric. In *Proc. ACM SIGCOMM*, pages 45–56, 1998.

[11] James F. Kurose and Keith W. Ross. *Computer Networking, a top down approach featuring the Internet, 3rd ed.* Addison-Wesley Longman, 2004.

[12] 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.

[13] Jennifer Rexford. *Handbook of Optimization in Telecommunications*, chapter Route optimization in IP networks. Kluwer Academic Publishers, 2005.

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

[15] Tim Roughgarden. How unfair is optimal routing? In *Proc. 13th Ann. ACM–SIAM Symp. on Discrete Algorithms (SODA)*, pages 203–204, 2002.

[16] Tim Roughgarden and Éva Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.

[17] Stefan Savage, Andy Collins, Eric Hoffman, John Snell, and Thomas E. Anderson. The end-to-end effects of internet path selection. In *Proc. ACM SIGCOMM*, pages 289–299, Boston, MA, 1999.

[18] John Glen Wardrop. Some theoretical aspects of road traffic research. In *Proc. of the Institute of Civil Engineers, Pt. II*, pages 325–378, 1952.