RWTH Aachen

Department of Computer Science
Technical Report

Adaptive Routing with Stale Information

Simon Fischer and Berthold Vöcking

ISSN 0935–3232 · Aachener Informatik Berichte · AIB-2005-06

RWTH Aachen · Department of Computer Science · April 2005

Adaptive Routing with Stale Information*

Simon Fischer and Berthold Vöcking

Department of Computer Science, RWTH Aachen University, 52056 Aachen, Germany
{fischer,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.

In our theoretical model, 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.

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

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

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.

One question arising in constrained nonlinear optimization (see, e. g. [3]) is how far and into what direction one can shift a search point improving the objective function but without violating the constraints. Similar questions arise in the study of rerouting policies under stale information where it is crucial not to "overshoot". However, it is unclear how to make use of techniques from convex optimization for our purposes.

2 Preliminaries

When information is always up-to-date, convergence towards Wardrop equilibria can easily be shown using a potential function introduced by Rosenthal in a classical paper [14].

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

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

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

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 $\ell_{\max} \leq (1+\epsilon)c$ we must have

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

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 1. *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

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

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

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

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.

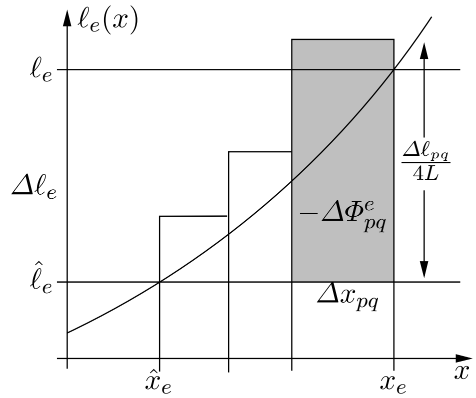

Fig. 1. The plot shows the function $\ell_e(\cdot)$ as a function of the edge load x. The area under ℓ_e in the range $[\hat{x}_e, x_e]$ is its true contribution to the potential difference. The area between $\hat{\ell}_e$ and the graph of $\ell_e(\cdot)$ in the range $[\hat{x}_e, x_e]$ is the error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

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.

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

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

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

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

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

By our choice of e, at least for one pair p, q, $q \ni e$ we have $\Delta\ell_{pq} = \Delta\ell_{\max}$. Comparing $-\Delta\Phi_{pq}^e$ 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 1. *If the conditions from Lemma 1 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. We consider Φ as a function of the edge flows. If the latency functions are strictly increasing, Φ has a unique minimum. By a translation and subtraction of $\min_{\mathbf{x}}\Phi(\mathbf{x})$, we have $\Phi(\mathbf{0}) = 0$. We use a generalized version of Lypunov's second method for differential equations with time delay [2, 7]. We use the notation from Theorem 6 in [7]. Again, by $\hat{t}$ denote the start of the last phase starting before t. Define the Lyapunov functional $\mathcal{V}$ to be half of the accumulated observed potential gain

$$\mathcal{V}(t, \mathbf{x}(\cdot)) = \Phi(\mathbf{x}(0)) + \frac{1}{2}\int_0^t \sum_{e\in E}\ell_e(x_e(\hat{u}))\cdot\dot{x}_e(u)\,du$$

and $w(\mathbf{x}) = \Phi(\mathbf{x})$. We have $\mathcal{V}(t, 0) \equiv 0$ and, $\mathcal{V}$ is continuous in t and Lipschitz with respect to $\mathbf{x}$. Lemma 1 implies that the true potential gain is always at least one half of the observed potential gain, and thus, $\mathcal{V}(t, \mathbf{x}(\cdot)) \geq w(\mathbf{x}(t))$. Also, the "derivative of $\mathcal{V}$ with respect to our differential equation" denoted $\dot{\mathcal{V}}_{(*)}$ in [7] is negative definite by Lemma 1. This proves convergence towards the Wardrop equilibrium. □

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 1 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.

3.1 Advantages of Proportional Sampling

Note that the relative latency proportional sampling rule $\mu(\ell_1, \ell_2) = \max\{0, (\ell_1 - \ell_2)/\ell_1\}$ is not α-smooth for any α. However, we can show that in combination with the proportional sampling rule $\sigma_{pq} = x_q/d_i$, $p, q \in P_i$ it is, if latency functions are bounded linearly from below.

Theorem 2. *Let $\ell_e(x) > \beta' \cdot x$ for all $e \in E$. Then the relative latency proportional migration rule in combination with proportional sampling is $1/(\beta' \min_{i \in \mathcal{I}}\{d_i\})$-smooth.*

Proof. Consider two paths $p, q \in P_i$ for some $i \in \mathcal{I}$. If ℓ_p is small, then ℓ_q must also be small. In particular, $x_q \leq \ell_q/\beta' \leq \ell_p/\beta'$ and we have $\sigma_{pq} \cdot \mu(\ell_p, \ell_q) = x_q/d_i \cdot (\ell_p - \ell_q)/\ell_p \leq (\ell_p - \ell_q)/(\beta' d_i)$. This implies, that the proportional sampling rule in combination with the relative latency proportional migration rule is $1/(\beta' \cdot \min_{i \in \mathcal{I}}\{d_i\})$-smooth. □

3.2 Smoothness is crucial

We can show that for oscillation avoidance α-smoothness is crucial. The following theorem shows that for given α, β, and L there exists a network and policy with these parameters such that the policy produces oscillating behaviour if T exceeds the bound from Theorem 1 by a constant factor. Equivalently, this holds if for given β, L, and T the value of α is by a constant factor below the respective bound.

Definition 2 (Oscillation). *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.*

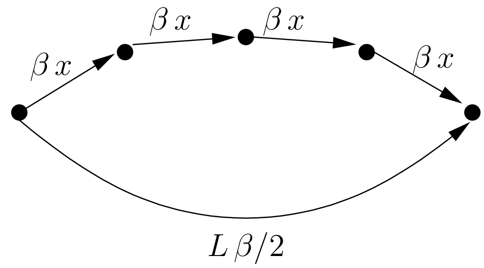

Fig. 2. A dynamics that is not α-smooth oscillates in this network if the update frequency $1/T$ is too small.

Theorem 3. *For any α, β, and L with $\alpha\,\beta\,L \geq 1$ there exists 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.*

Proof. Consider a Network with two paths p and q. Path p consists of L edges with latency functions $\ell_1(x) = \beta\,x$ and path q consists of a single edge with the constant latency function $\ell_2(x) = L\,\beta/2$. We have $\ell_p(\mathbf{x}) = x_p\,\beta\,L$ and $\ell_q(\mathbf{x}) = L\,\beta/2$. The network is depicted in Figure 2. At the unique Wardrop equilibrium $\mathbf{x}^* = (1/2, 1/2)$, all paths have total latency $L\,\beta/2$. Consider the sampling rule $\sigma_{pq} = \sigma_{qp} = 1$ and $\sigma_{pp} = \sigma_{qq} = 0$ in combination with the migration rule $\mu(u, v) = \min\{\max\{\alpha \cdot (u - v)], 0\}, 1\}$. For simplicity we will also use $\alpha' = \alpha\,\beta\,L$.

We are looking for a starting vector with $x_q(0) = 1/2 + \delta$ such that $x_q(T) = 1/2 - \delta$. As a first guess we choose $\delta = 1/(4\,\alpha')$. Observe, that $\mu(\ell_q, \ell_p) = \alpha\,\beta\,L\,\delta = 1/4$ implying $x_q(t) = (1/2 + \delta)e^{-T/4}$. We start with the weaker condition

$$(1/2 + \delta) \cdot e^{-T/4} \leq 1/2 - \delta, \tag{5}$$

which is equivalent to

$$T/4 \geq \ln\left(\frac{1/2 + \delta}{1/2 - \delta}\right). \tag{6}$$

We have

$$\begin{aligned}
\ln\left(\frac{1/2 + \delta}{1/2 - \delta}\right) &= \ln\left(1 + \frac{2\,\delta}{1/2 - \delta}\right) \\
&= \ln\left(1 + \frac{1/(2\,\alpha')}{1/2 - 1/(4\,\alpha')}\right) \\
&\leq \frac{1/(2\,\alpha')}{1/2 - 1/(4\,\alpha')} \\
&\leq 2/\alpha' \leq T/4
\end{aligned}$$

by our assumption on T. Altogether, this shows that Equation (6) and hence also Equation (5) hold.

If we choose $\delta \leq 1/(4\alpha')$, then $\mu(\ell_q, \ell_p) = \delta\,\alpha'$. At time T we have

$$x_q(T) = (1/2 + \delta) \cdot e^{-\alpha'\,\delta\,T}$$

which is strictly increasing in δ. Therefore, for some $\delta' \leq 1/(4\,\alpha')$, we must have $x_q(0) = 1/2 + \delta'$ and $x_q(T) = 1/2 - \delta'$ with equality. After this phase, a new phase starts with the roles of p and q exchanged. Altogether, we have $x_q(2\,n\,T) = 1/2 + \delta'$ and $x_q((2\,n+1)\,T) = 1/2 - \delta'$ for all $n \in \mathbb{N}$. $\square$

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 3 ((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 (Sketch of 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 $1/|P_i|$ this agents samples a cheapest path with respect to her own commodity i. She migrates to this path with probability at least $((\ell_{\min,i} + \delta) - \ell_{\min,i})/\ell_{\max}$ and 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 $-x_p \le \dot{x}_p \le x_p$ and hence each population share decreases at most exponentially fast. After time T there are still at least $\epsilon \cdot e^{-T}$ δ-unsatisfied agents. We can assume that $T \le 1$. In the worst case, all of these agents belong to the commodity with largest number of paths $|P_i|$. Therefore, at every time within the phase, their virtual potential gain in a time step dt is at most

$$-e^{-1}\epsilon \cdot \frac{1}{m}(\delta/\ell_{\max}) \cdot \delta\, dt$$

and at the end of the phase of length T, the virtual potential gain is at most $-Te^{-1}\epsilon\delta^2/(m\ell_{\max})$. By Lemma 1 the true potential gain is at least half of the virtual latency gain and hence

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

Since Φ is bounded from above by $\ell_{\max}$, we reach $\Phi = 0$ after at most

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

phases which is our claim. $\square$

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 4 (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 (Sketch of 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. Altogether, the virtual potential gain in the phase of length T is at most

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

The rest of the proof is analog to the proof of Theorem 4. $\square$

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, Angelia Nedic, and Asuman E. Ozdaglar. *Convex Analysis and Optimization*. Athena Scientific.
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 *2nd IASTED International Conference 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. of the 12th European Symposium on Algorithms (ESA)*, LNCS 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. A. Khanna and J. Zinky. The revised ARPANET routing metric. In *Proc. ACM SIGCOMM*, pages 45–56, September 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. of the PODC 1997*, 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. of the SODA 2002*, 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 *Proceedings of SIGCOMM*, pages 289–299, Boston, MA, August-September 1999.
18. J. G. Wardrop. Some theoretical aspects of road traffic research. In *Proceedings of the Institute of Civil Engineers, Pt. II*, pages 325–378, 1952.

Aachener Informatik-Berichte

This is a list of recent technical reports. To obtain copies of technical reports please consult http://aib.informatik.rwth-aachen.de/ or send your request to: Informatik-Bibliothek, RWTH Aachen, Ahornstr. 55, 52056 Aachen, Email: `biblio@informatik.rwth-aachen.de`

1987-01 *	Fachgruppe Informatik: Jahresbericht 1986
1987-02 *	David de Frutos Escrig, Klaus Indermark: Equivalence Relations of Non-Deterministic Ianov-Schemes
1987-03 *	Manfred Nagl: A Software Development Environment based on Graph Technology
1987-04 *	Claus Lewerentz, Manfred Nagl, Bernhard Westfechtel: On Integration Mechanisms within a Graph-Based Software Development Environment
1987-05 *	Reinhard Rinn: Über Eingabeanomalien bei verschiedenen Inferenzmodellen
1987-06 *	Werner Damm, Gert Döhmen: Specifying Distributed Computer Architectures in AADL*
1987-07 *	Gregor Engels, Claus Lewerentz, Wilhelm Schäfer: Graph Grammar Engineering: A Software Specification Method
1987-08 *	Manfred Nagl: Set Theoretic Approaches to Graph Grammars
1987-09 *	Claus Lewerentz, Andreas Schürr: Experiences with a Database System for Software Documents
1987-10 *	Herbert Klaeren, Klaus Indermark: A New Implementation Technique for Recursive Function Definitions
1987-11 *	Rita Loogen: Design of a Parallel Programmable Graph Reduction Machine with Distributed Memory
1987-12	J. Börstler, U. Möncke, R. Wilhelm: Table compression for tree automata
1988-01 *	Gabriele Esser, Johannes Rückert, Frank Wagner: Gesellschaftliche Aspekte der Informatik
1988-02 *	Peter Martini, Otto Spaniol: Token-Passing in High-Speed Backbone Networks for Campus-Wide Environments
1988-03 *	Thomas Welzel: Simulation of a Multiple Token Ring Backbone
1988-04 *	Peter Martini: Performance Comparison for HSLAN Media Access Protocols
1988-05 *	Peter Martini: Performance Analysis of Multiple Token Rings
1988-06 *	Andreas Mann, Johannes Rückert, Otto Spaniol: Datenfunknetze
1988-07 *	Andreas Mann, Johannes Rückert: Packet Radio Networks for Data Exchange
1988-08 *	Andreas Mann, Johannes Rückert: Concurrent Slot Assignment Protocol for Packet Radio Networks
1988-09 *	W. Kremer, F. Reichert, J. Rückert, A. Mann: Entwurf einer Netzwerktopologie für ein Mobilfunknetz zur Unterstützung des öffentlichen Straßenverkehrs
1988-10 *	Kai Jakobs: Towards User-Friendly Networking
1988-11 *	Kai Jakobs: The Directory - Evolution of a Standard
1988-12 *	Kai Jakobs: Directory Services in Distributed Systems - A Survey
1988-13 *	Martine Schümmer: RS-511, a Protocol for the Plant Floor

1988-14 * U. Quernheim: Satellite Communication Protocols - A Performance Comparison Considering On-Board Processing
1988-15 * Peter Martini, Otto Spaniol, Thomas Welzel: File Transfer in High Speed Token Ring Networks: Performance Evaluation by Approximate Analysis and Simulation
1988-16 * Fachgruppe Informatik: Jahresbericht 1987
1988-17 * Wolfgang Thomas: Automata on Infinite Objects
1988-18 * Michael Sonnenschein: On Petri Nets and Data Flow Graphs
1988-19 * Heiko Vogler: Functional Distribution of the Contextual Analysis in Block-Structured Programming Languages: A Case Study of Tree Transducers
1988-20 * Thomas Welzel: Einsatz des Simulationswerkzeuges QNAP2 zur Leistungsbewertung von Kommunikationsprotokollen
1988-21 * Th. Janning, C. Lewerentz: Integrated Project Team Management in a Software Development Environment
1988-22 * Joost Engelfriet, Heiko Vogler: Modular Tree Transducers
1988-23 * Wolfgang Thomas: Automata and Quantifier Hierarchies
1988-24 * Uschi Heuter: Generalized Definite Tree Languages
1989-01 * Fachgruppe Informatik: Jahresbericht 1988
1989-02 * G. Esser, J. Rückert, F. Wagner (Hrsg.): Gesellschaftliche Aspekte der Informatik
1989-03 * Heiko Vogler: Bottom-Up Computation of Primitive Recursive Tree Functions
1989-04 * Andy Schürr: Introduction to PROGRESS, an Attribute Graph Grammar Based Specification Language
1989-05 J. Börstler: Reuse and Software Development - Problems, Solutions, and Bibliography (in German)
1989-06 * Kai Jakobs: OSI - An Appropriate Basis for Group Communication?
1989-07 * Kai Jakobs: ISO's Directory Proposal - Evolution, Current Status and Future Problems
1989-08 * Bernhard Westfechtel: Extension of a Graph Storage for Software Documents with Primitives for Undo/Redo and Revision Control
1989-09 * Peter Martini: High Speed Local Area Networks - A Tutorial
1989-10 * P. Davids, Th. Welzel: Performance Analysis of DQDB Based on Simulation
1989-11 * Manfred Nagl (Ed.): Abstracts of Talks presented at the WG '89 15th International Workshop on Graphtheoretic Concepts in Computer Science
1989-12 * Peter Martini: The DQDB Protocol - Is it Playing the Game?
1989-13 * Martine Schümmer: CNC/DNC Communication with MAP
1989-14 * Martine Schümmer: Local Area Networks for Manufactoring Environments with hard Real-Time Requirements
1989-15 * M. Schümmer, Th. Welzel, P. Martini: Integration of Field Bus and MAP Networks - Hierarchical Communication Systems in Production Environments
1989-16 * G. Vossen, K.-U. Witt: SUXESS: Towards a Sound Unification of Extensions of the Relational Data Model

1989-17 * J. Derissen, P. Hruschka, M.v.d. Beeck, Th. Janning, M. Nagl: Integrating Structured Analysis and Information Modelling

1989-18 A. Maassen: Programming with Higher Order Functions

1989-19 * Mario Rodriguez-Artalejo, Heiko Vogler: A Narrowing Machine for Syntax Directed BABEL

1989-20 H. Kuchen, R. Loogen, J.J. Moreno Navarro, M. Rodriguez Artalejo: Graph-based Implementation of a Functional Logic Language

1990-01 * Fachgruppe Informatik: Jahresbericht 1989

1990-02 * Vera Jansen, Andreas Potthoff, Wolfgang Thomas, Udo Wermuth: A Short Guide to the AMORE System (Computing Automata, MOnoids and Regular Expressions)

1990-03 * Jerzy Skurczynski: On Three Hierarchies of Weak SkS Formulas

1990-04 R. Loogen: Stack-based Implementation of Narrowing

1990-05 H. Kuchen, A. Wagener: Comparison of Dynamic Load Balancing Strategies

1990-06 * Kai Jakobs, Frank Reichert: Directory Services for Mobile Communication

1990-07 * Kai Jakobs: What's Beyond the Interface - OSI Networks to Support Cooperative Work

1990-08 * Kai Jakobs: Directory Names and Schema - An Evaluation

1990-09 * Ulrich Quernheim, Dieter Kreuer: Das CCITT - Signalisierungssystem Nr. 7 auf Satellitenstrecken; Simulation der Zeichengabestrecke

1990-11 H. Kuchen, R. Loogen, J.J. Moreno Navarro, M. Rodriguez Artalejo: Lazy Narrowing in a Graph Machine

1990-12 * Kai Jakobs, Josef Kaltwasser, Frank Reichert, Otto Spaniol: Der Computer fährt mit

1990-13 * Rudolf Mathar, Andreas Mann: Analyzing a Distributed Slot Assignment Protocol by Markov Chains

1990-14 A. Maassen: Compilerentwicklung in Miranda - ein Praktikum in funktionaler Programmierung (written in german)

1990-15 * Manfred Nagl, Andreas Schürr: A Specification Environment for Graph Grammars

1990-16 A. Schürr: PROGRESS: A VHL-Language Based on Graph Grammars

1990-17 * Marita Möller: Ein Ebenenmodell wissensbasierter Konsultationen - Unterstützung für Wissensakquisition und Erklärungsfähigkeit

1990-18 * Eric Kowalewski: Entwurf und Interpretation einer Sprache zur Beschreibung von Konsultationsphasen in Expertensystemen

1990-20 Y. Ortega Mallen, D. de Frutos Escrig: A Complete Proof System for Timed Observations

1990-21 * Manfred Nagl: Modelling of Software Architectures: Importance, Notions, Experiences

1990-22 H. Fassbender, H. Vogler: A Call-by-need Implementation of Syntax Directed Functional Programming

1991-01 Guenther Geiler (ed.), Fachgruppe Informatik: Jahresbericht 1990

1991-03 B. Steffen, A. Ingolfsdottir: Characteristic Formulae for Processes with Divergence

1991-04 M. Portz: A new class of cryptosystems based on interconnection networks

1991-05 H. Kuchen, G. Geiler: Distributed Applicative Arrays

1991-06 * Ludwig Staiger: Kolmogorov Complexity and Hausdorff Dimension

1991-07 * Ludwig Staiger: Syntactic Congruences for w-languages

1991-09 * Eila Kuikka: A Proposal for a Syntax-Directed Text Processing System

1991-10 K. Gladitz, H. Fassbender, H. Vogler: Compiler-based Implementation of Syntax-Directed Functional Programming

1991-11 R. Loogen, St. Winkler: Dynamic Detection of Determinism in Functional Logic Languages

1991-12 * K. Indermark, M. Rodriguez Artalejo (Eds.): Granada Workshop on the Integration of Functional and Logic Programming

1991-13 * Rolf Hager, Wolfgang Kremer: The Adaptive Priority Scheduler: A More Fair Priority Service Discipline

1991-14 * Andreas Fasbender, Wolfgang Kremer: A New Approximation Algorithm for Tandem Networks with Priority Nodes

1991-15 J. Börstler, A. Zündorf: Revisiting extensions to Modula-2 to support reusability

1991-16 J. Börstler, Th. Janning: Bridging the gap between Requirements Analysis and Design

1991-17 A. Zündorf, A. Schürr: Nondeterministic Control Structures for Graph Rewriting Systems

1991-18 * Matthias Jarke, John Mylopoulos, Joachim W. Schmidt, Yannis Vassiliou: DAIDA: An Environment for Evolving Information Systems

1991-19 M. Jeusfeld, M. Jarke: From Relational to Object-Oriented Integrity Simplification

1991-20 G. Hogen, A. Kindler, R. Loogen: Automatic Parallelization of Lazy Functional Programs

1991-21 * Prof. Dr. rer. nat. Otto Spaniol: ODP (Open Distributed Processing): Yet another Viewpoint

1991-22 H. Kuchen, F. Lücking, H. Stoltze: The Topology Description Language TDL

1991-23 S. Graf, B. Steffen: Compositional Minimization of Finite State Systems

1991-24 R. Cleaveland, J. Parrow, B. Steffen: The Concurrency Workbench: A Semantics Based Tool for the Verification of Concurrent Systems

1991-25 * Rudolf Mathar, Jürgen Mattfeldt: Optimal Transmission Ranges for Mobile Communication in Linear Multihop Packet Radio Networks

1991-26 M. Jeusfeld, M. Staudt: Query Optimization in Deductive Object Bases

1991-27 J. Knoop, B. Steffen: The Interprocedural Coincidence Theorem

1991-28 J. Knoop, B. Steffen: Unifying Strength Reduction and Semantic Code Motion

1991-30 T. Margaria: First-Order theories for the verification of complex FSMs

1991-31 B. Steffen: Generating Data Flow Analysis Algorithms from Modal Specifications

1992-01 Stefan Eherer (ed.), Fachgruppe Informatik: Jahresbericht 1991

1992-02 * Bernhard Westfechtel: Basismechanismen zur Datenverwaltung in strukturbezogenen Hypertextsystemen

1992-04 S. A. Smolka, B. Steffen: Priority as Extremal Probability

1992-05 * Matthias Jarke, Carlos Maltzahn, Thomas Rose: Sharing Processes: Team Coordination in Design Repositories

1992-06 O. Burkart, B. Steffen: Model Checking for Context-Free Processes
1992-07 * Matthias Jarke, Klaus Pohl: Information Systems Quality and Quality Information Systems
1992-08 * Rudolf Mathar, Jürgen Mattfeldt: Analyzing Routing Strategy NFP in Multihop Packet Radio Networks on a Line
1992-09 * Alfons Kemper, Guido Moerkotte: Grundlagen objektorientierter Datenbanksysteme
1992-10 Matthias Jarke, Manfred Jeusfeld, Andreas Miethsam, Michael Gocek: Towards a logic-based reconstruction of software configuration management
1992-11 Werner Hans: A Complete Indexing Scheme for WAM-based Abstract Machines
1992-12 W. Hans, R. Loogen, St. Winkler: On the Interaction of Lazy Evaluation and Backtracking
1992-13 * Matthias Jarke, Thomas Rose: Specification Management with CAD
1992-14 Th. Noll, H. Vogler: Top-down Parsing with Simultaneous Evaluation on Noncircular Attribute Grammars
1992-15 A. Schuerr, B. Westfechtel: Graphgrammatiken und Graphersetzungssysteme(written in german)
1992-16 * Graduiertenkolleg Informatik und Technik (Hrsg.): Forschungsprojekte des Graduiertenkollegs Informatik und Technik
1992-17 M. Jarke (ed.): ConceptBase V3.1 User Manual
1992-18 * Clarence A. Ellis, Matthias Jarke (Eds.): Distributed Cooperation in Integrated Information Systems - Proceedings of the Third International Workshop on Intelligent and Cooperative Information Systems
1992-19-00 H. Kuchen, R. Loogen (eds.): Proceedings of the 4th Int. Workshop on the Parallel Implementation of Functional Languages
1992-19-01 G. Hogen, R. Loogen: PASTEL - A Parallel Stack-Based Implementation of Eager Functional Programs with Lazy Data Structures (Extended Abstract)
1992-19-02 H. Kuchen, K. Gladitz: Implementing Bags on a Shared Memory MIMD-Machine
1992-19-03 C. Rathsack, S.B. Scholz: LISA - A Lazy Interpreter for a Full-Fledged Lambda-Calculus
1992-19-04 T.A. Bratvold: Determining Useful Parallelism in Higher Order Functions
1992-19-05 S. Kahrs: Polymorphic Type Checking by Interpretation of Code
1992-19-06 M. Chakravarty, M. Köhler: Equational Constraints, Residuation, and the Parallel JUMP-Machine
1992-19-07 J. Seward: Polymorphic Strictness Analysis using Frontiers (Draft Version)
1992-19-08 D. Gärtner, A. Kimms, W. Kluge: pi-Red^+ - A Compiling Graph-Reduction System for a Full Fledged Lambda-Calculus
1992-19-09 D. Howe, G. Burn: Experiments with strict STG code
1992-19-10 J. Glauert: Parallel Implementation of Functional Languages Using Small Processes
1992-19-11 M. Joy, T. Axford: A Parallel Graph Reduction Machine
1992-19-12 A. Bennett, P. Kelly: Simulation of Multicache Parallel Reduction

1992-19-13 K. Langendoen, D.J. Agterkamp: Cache Behaviour of Lazy Functional Programs (Working Paper)

1992-19-14 K. Hammond, S. Peyton Jones: Profiling scheduling strategies on the GRIP parallel reducer

1992-19-15 S. Mintchev: Using Strictness Information in the STG-machine

1992-19-16 D. Rushall: An Attribute Grammar Evaluator in Haskell

1992-19-17 J. Wild, H. Glaser, P. Hartel: Statistics on storage management in a lazy functional language implementation

1992-19-18 W.S. Martins: Parallel Implementations of Functional Languages

1992-19-19 D. Lester: Distributed Garbage Collection of Cyclic Structures (Draft version)

1992-19-20 J.C. Glas, R.F.H. Hofman, W.G. Vree: Parallelization of Branch-and-Bound Algorithms in a Functional Programming Environment

1992-19-21 S. Hwang, D. Rushall: The nu-STG machine: a parallelized Spineless Tagless Graph Reduction Machine in a distributed memory architecture (Draft version)

1992-19-22 G. Burn, D. Le Metayer: Cps-Translation and the Correctness of Optimising Compilers

1992-19-23 S.L. Peyton Jones, P. Wadler: Imperative functional programming (Brief summary)

1992-19-24 W. Damm, F. Liu, Th. Peikenkamp: Evaluation and Parallelization of Functions in Functional + Logic Languages (abstract)

1992-19-25 M. Kesseler: Communication Issues Regarding Parallel Functional Graph Rewriting

1992-19-26 Th. Peikenkamp: Charakterizing and representing neededness in functional loginc languages (abstract)

1992-19-27 H. Doerr: Monitoring with Graph-Grammars as formal operational Models

1992-19-28 J. van Groningen: Some implementation aspects of Concurrent Clean on distributed memory architectures

1992-19-29 G. Ostheimer: Load Bounding for Implicit Parallelism (abstract)

1992-20 H. Kuchen, F.J. Lopez Fraguas, J.J. Moreno Navarro, M. Rodriguez Artalejo: Implementing Disequality in a Lazy Functional Logic Language

1992-21 H. Kuchen, F.J. Lopez Fraguas: Result Directed Computing in a Functional Logic Language

1992-22 H. Kuchen, J.J. Moreno Navarro, M.V. Hermenegildo: Independent AND-Parallel Narrowing

1992-23 T. Margaria, B. Steffen: Distinguishing Formulas for Free

1992-24 K. Pohl: The Three Dimensions of Requirements Engineering

1992-25 * R. Stainov: A Dynamic Configuration Facility for Multimedia Communications

1992-26 * Michael von der Beeck: Integration of Structured Analysis and Timed Statecharts for Real-Time and Concurrency Specification

1992-27 W. Hans, St. Winkler: Aliasing and Groundness Analysis of Logic Programs through Abstract Interpretation and its Safety

1992-28 * Gerhard Steinke, Matthias Jarke: Support for Security Modeling in Information Systems Design

1992-29 B. Schinzel: Warum Frauenforschung in Naturwissenschaft und Technik

1992-30 A. Kemper, G. Moerkotte, K. Peithner: Object-Orientation Axiomatised by Dynamic Logic

1992-32 * Bernd Heinrichs, Kai Jakobs: Timer Handling in High-Performance Transport Systems

1992-33 * B. Heinrichs, K. Jakobs, K. Lenßen, W. Reinhardt, A. Spinner: Euro-Bridge: Communication Services for Multimedia Applications

1992-34 C. Gerlhof, A. Kemper, Ch. Kilger, G. Moerkotte: Partition-Based Clustering in Object Bases: From Theory to Practice

1992-35 J. Börstler: Feature-Oriented Classification and Reuse in IPSEN

1992-36 M. Jarke, J. Bubenko, C. Rolland, A. Sutcliffe, Y. Vassiliou: Theories Underlying Requirements Engineering: An Overview of NATURE at Genesis

1992-37 * K. Pohl, M. Jarke: Quality Information Systems: Repository Support for Evolving Process Models

1992-38 A. Zuendorf: Implementation of the imperative / rule based language PROGRES

1992-39 P. Koch: Intelligentes Backtracking bei der Auswertung funktional-logischer Programme

1992-40 * Rudolf Mathar, Jürgen Mattfeldt: Channel Assignment in Cellular Radio Networks

1992-41 * Gerhard Friedrich, Wolfgang Neidl: Constructive Utility in Model-Based Diagnosis Repair Systems

1992-42 * P. S. Chen, R. Hennicker, M. Jarke: On the Retrieval of Reusable Software Components

1992-43 W. Hans, St.Winkler: Abstract Interpretation of Functional Logic Languages

1992-44 N. Kiesel, A. Schuerr, B. Westfechtel: Design and Evaluation of GRAS, a Graph-Oriented Database System for Engineering Applications

1993-01 * Fachgruppe Informatik: Jahresbericht 1992

1993-02 * Patrick Shicheng Chen: On Inference Rules of Logic-Based Information Retrieval Systems

1993-03 G. Hogen, R. Loogen: A New Stack Technique for the Management of Runtime Structures in Distributed Environments

1993-05 A. Zündorf: A Heuristic for the Subgraph Isomorphism Problem in Executing PROGRES

1993-06 A. Kemper, D. Kossmann: Adaptable Pointer Swizzling Strategies in Object Bases: Design, Realization, and Quantitative Analysis

1993-07 * Graduiertenkolleg Informatik und Technik (Hrsg.): Graduiertenkolleg Informatik und Technik

1993-08 * Matthias Berger: k-Coloring Vertices using a Neural Network with Convergence to Valid Solutions

1993-09 M. Buchheit, M. Jeusfeld, W. Nutt, M. Staudt: Subsumption between Queries to Object-Oriented Databases

1993-10 O. Burkart, B. Steffen: Pushdown Processes: Parallel Composition and Model Checking

1993-11 * R. Große-Wienker, O. Hermanns, D. Menzenbach, A. Pollacks, S. Repetzki, J. Schwartz, K. Sonnenschein, B. Westfechtel: Das SUKITS-Projekt: A-posteriori-Integration heterogener CIM-Anwendungssysteme

1993-12 * Rudolf Mathar, Jürgen Mattfeldt: On the Distribution of Cumulated Interference Power in Rayleigh Fading Channels
1993-13 O. Maler, L. Staiger: On Syntactic Congruences for omega-languages
1993-14 M. Jarke, St. Eherer, R. Gallersdoerfer, M. Jeusfeld, M. Staudt: ConceptBase - A Deductive Object Base Manager
1993-15 M. Staudt, H.W. Nissen, M.A. Jeusfeld: Query by Class, Rule and Concept
1993-16 * M. Jarke, K. Pohl, St. Jacobs et al.: Requirements Engineering: An Integrated View of Representation Process and Domain
1993-17 * M. Jarke, K. Pohl: Establishing Vision in Context: Towards a Model of Requirements Processes
1993-18 W. Hans, H. Kuchen, St. Winkler: Full Indexing for Lazy Narrowing
1993-19 W. Hans, J.J. Ruz, F. Saenz, St. Winkler: A VHDL Specification of a Shared Memory Parallel Machine for Babel
1993-20 * K. Finke, M. Jarke, P. Szczurko, R. Soltysiak: Quality Management for Expert Systems in Process Control
1993-21 M. Jarke, M.A. Jeusfeld, P. Szczurko: Three Aspects of Intelligent Cooperation in the Quality Cycle
1994-01 Margit Generet, Sven Martin (eds.), Fachgruppe Informatik: Jahresbericht 1993
1994-02 M. Lefering: Development of Incremental Integration Tools Using Formal Specifications
1994-03 * P. Constantopoulos, M. Jarke, J. Mylopoulos, Y. Vassiliou: The Software Information Base: A Server for Reuse
1994-04 * Rolf Hager, Rudolf Mathar, Jürgen Mattfeldt: Intelligent Cruise Control and Reliable Communication of Mobile Stations
1994-05 * Rolf Hager, Peter Hermesmann, Michael Portz: Feasibility of Authentication Procedures within Advanced Transport Telematics
1994-06 * Claudia Popien, Bernd Meyer, Axel Kuepper: A Formal Approach to Service Import in ODP Trader Federations
1994-07 P. Peters, P. Szczurko: Integrating Models of Quality Management Methods by an Object-Oriented Repository
1994-08 * Manfred Nagl, Bernhard Westfechtel: A Universal Component for the Administration in Distributed and Integrated Development Environments
1994-09 * Patrick Horster, Holger Petersen: Signatur- und Authentifikationsverfahren auf der Basis des diskreten Logarithmusproblems
1994-11 A. Schürr: PROGRES, A Visual Language and Environment for PROgramming with Graph REwrite Systems
1994-12 A. Schürr: Specification of Graph Translators with Triple Graph Grammars
1994-13 A. Schürr: Logic Based Programmed Structure Rewriting Systems
1994-14 L. Staiger: Codes, Simplifying Words, and Open Set Condition
1994-15 * Bernhard Westfechtel: A Graph-Based System for Managing Configurations of Engineering Design Documents
1994-16 P. Klein: Designing Software with Modula-3
1994-17 I. Litovsky, L. Staiger: Finite acceptance of infinite words

1994-18 G. Hogen, R. Loogen: Parallel Functional Implementations: Graphbased vs. Stackbased Reduction

1994-19 M. Jeusfeld, U. Johnen: An Executable Meta Model for Re-Engineering of Database Schemas

1994-20 * R. Gallersdörfer, M. Jarke, K. Klabunde: Intelligent Networks as a Data Intensive Application (INDIA)

1994-21 M. Mohnen: Proving the Correctness of the Static Link Technique Using Evolving Algebras

1994-22 H. Fernau, L. Staiger: Valuations and Unambiguity of Languages, with Applications to Fractal Geometry

1994-24 * M. Jarke, K. Pohl, R. Dömges, St. Jacobs, H. W. Nissen: Requirements Information Management: The NATURE Approach

1994-25 * M. Jarke, K. Pohl, C. Rolland, J.-R. Schmitt: Experience-Based Method Evaluation and Improvement: A Process Modeling Approach

1994-26 * St. Jacobs, St. Kethers: Improving Communication and Decision Making within Quality Function Deployment

1994-27 * M. Jarke, H. W. Nissen, K. Pohl: Tool Integration in Evolving Information Systems Environments

1994-28 O. Burkart, D. Caucal, B. Steffen: An Elementary Bisimulation Decision Procedure for Arbitrary Context-Free Processes

1995-01 * Fachgruppe Informatik: Jahresbericht 1994

1995-02 Andy Schürr, Andreas J. Winter, Albert Zündorf: Graph Grammar Engineering with PROGRES

1995-03 Ludwig Staiger: A Tight Upper Bound on Kolmogorov Complexity by Hausdorff Dimension and Uniformly Optimal Prediction

1995-04 Birgitta König-Ries, Sven Helmer, Guido Moerkotte: An experimental study on the complexity of left-deep join ordering problems for cyclic queries

1995-05 Sophie Cluet, Guido Moerkotte: Efficient Evaluation of Aggregates on Bulk Types

1995-06 Sophie Cluet, Guido Moerkotte: Nested Queries in Object Bases

1995-07 Sophie Cluet, Guido Moerkotte: Query Optimization Techniques Exploiting Class Hierarchies

1995-08 Markus Mohnen: Efficient Compile-Time Garbage Collection for Arbitrary Data Structures

1995-09 Markus Mohnen: Functional Specification of Imperative Programs: An Alternative Point of View of Functional Languages

1995-10 Rainer Gallersdörfer, Matthias Nicola: Improving Performance in Replicated Databases through Relaxed Coherency

1995-11 * M.Staudt, K.von Thadden: Subsumption Checking in Knowledge Bases

1995-12 * G.V.Zemanek, H.W.Nissen, H.Hubert, M.Jarke: Requirements Analysis from Multiple Perspectives: Experiences with Conceptual Modeling Technology

1995-13 * M.Staudt, M.Jarke: Incremental Maintenance of Externally Materialized Views

1995-14 * P.Peters, P.Szczurko, M.Jeusfeld: Oriented Information Management: Conceptual Models at Work

1995-15 * Matthias Jarke, Sudha Ram (Hrsg.): WITS 95 Proceedings of the 5th Annual Workshop on Information Technologies and Systems
1995-16 * W.Hans, St.Winkler, F.Saenz: Distributed Execution in Functional Logic Programming
1996-01 * Jahresbericht 1995
1996-02 Michael Hanus, Christian Prehofer: Higher-Order Narrowing with Definitional Trees
1996-03 * W.Scheufele, G.Moerkotte: Optimal Ordering of Selections and Joins in Acyclic Queries with Expensive Predicates
1996-04 Klaus Pohl: PRO-ART: Enabling Requirements Pre-Traceability
1996-05 Klaus Pohl: Requirements Engineering: An Overview
1996-06 * M.Jarke, W.Marquardt: Design and Evaluation of Computer–Aided Process Modelling Tools
1996-07 Olaf Chitil: The Sigma-Semantics: A Comprehensive Semantics for Functional Programs
1996-08 * S.Sripada: On Entropy and the Limitations of the Second Law of Thermodynamics
1996-09 Michael Hanus (Ed.): Proceedings of the Poster Session of ALP96 - Fifth International Conference on Algebraic and Logic Programming
1996-09-0 Michael Hanus (Ed.): Proceedings of the Poster Session of ALP 96 - Fifth International Conference on Algebraic and Logic Programming: Introduction and table of contents
1996-09-1 Ilies Alouini: An Implementation of Conditional Concurrent Rewriting on Distributed Memory Machines
1996-09-2 Olivier Danvy, Karoline Malmkjær: On the Idempotence of the CPS Transformation
1996-09-3 Víctor M. Gulías, José L. Freire: Concurrent Programming in Haskell
1996-09-4 Sébastien Limet, Pierre Réty: On Decidability of Unifiability Modulo Rewrite Systems
1996-09-5 Alexandre Tessier: Declarative Debugging in Constraint Logic Programming
1996-10 Reidar Conradi, Bernhard Westfechtel: Version Models for Software Configuration Management
1996-11 * C.Weise, D.Lenzkes: A Fast Decision Algorithm for Timed Refinement
1996-12 * R.Dömges, K.Pohl, M.Jarke, B.Lohmann, W.Marquardt: PRO-ART/CE* — An Environment for Managing the Evolution of Chemical Process Simulation Models
1996-13 * K.Pohl, R.Klamma, K.Weidenhaupt, R.Dömges, P.Haumer, M.Jarke: A Framework for Process-Integrated Tools
1996-14 * R.Gallersdörfer, K.Klabunde, A.Stolz, M.Eßmajor: INDIA — Intelligent Networks as a Data Intensive Application, Final Project Report, June 1996
1996-15 * H.Schimpe, M.Staudt: VAREX: An Environment for Validating and Refining Rule Bases
1996-16 * M.Jarke, M.Gebhardt, S.Jacobs, H.Nissen: Conflict Analysis Across Heterogeneous Viewpoints: Formalization and Visualization
1996-17 Manfred A. Jeusfeld, Tung X. Bui: Decision Support Components on the Internet

1996-18 Manfred A. Jeusfeld, Mike Papazoglou: Information Brokering: Design, Search and Transformation

1996-19 * P.Peters, M.Jarke: Simulating the impact of information flows in networked organizations

1996-20 Matthias Jarke, Peter Peters, Manfred A. Jeusfeld: Model-driven planning and design of cooperative information systems

1996-21 * G.de Michelis, E.Dubois, M.Jarke, F.Matthes, J.Mylopoulos, K.Pohl, J.Schmidt, C.Woo, E.Yu: Cooperative information systems: a manifesto

1996-22 * S.Jacobs, M.Gebhardt, S.Kethers, W.Rzasa: Filling HTML forms simultaneously: CoWeb architecture and functionality

1996-23 * M.Gebhardt, S.Jacobs: Conflict Management in Design

1997-01 Michael Hanus, Frank Zartmann (eds.): Jahresbericht 1996

1997-02 Johannes Faassen: Using full parallel Boltzmann Machines for Optimization

1997-03 Andreas Winter, Andy Schürr: Modules and Updatable Graph Views for PROgrammed Graph REwriting Systems

1997-04 Markus Mohnen, Stefan Tobies: Implementing Context Patterns in the Glasgow Haskell Compiler

1997-05 * S.Gruner: Schemakorrespondenzaxiome unterstützen die paargrammatische Spezifikation inkrementeller Integrationswerkzeuge

1997-06 Matthias Nicola, Matthias Jarke: Design and Evaluation of Wireless Health Care Information Systems in Developing Countries

1997-07 Petra Hofstedt: Taskparallele Skelette für irregulär strukturierte Probleme in deklarativen Sprachen

1997-08 Dorothea Blostein, Andy Schürr: Computing with Graphs and Graph Rewriting

1997-09 Carl-Arndt Krapp, Bernhard Westfechtel: Feedback Handling in Dynamic Task Nets

1997-10 Matthias Nicola, Matthias Jarke: Integrating Replication and Communication in Performance Models of Distributed Databases

1997-11 * R. Klamma, P. Peters, M. Jarke: Workflow Support for Failure Management in Federated Organizations

1997-13 Markus Mohnen: Optimising the Memory Management of Higher-Order Functional Programs

1997-14 Roland Baumann: Client/Server Distribution in a Structure-Oriented Database Management System

1997-15 George Botorog: High-Level Parallel Programming and the Efficient Implementation of Numerical Algorithms

1998-01 * Fachgruppe Informatik: Jahresbericht 1997

1998-02 Stefan Gruner, Manfred Nagel, Andy Schürr: Fine-grained and Structure-Oriented Document Integration Tools are Needed for Development Processes

1998-03 Stefan Gruner: Einige Anmerkungen zur graphgrammatischen Spezifikation von Integrationswerkzeugen nach Westfechtel, Janning, Lefering und Schürr

1998-04 * O. Kubitz: Mobile Robots in Dynamic Environments

1998-05 Martin Leucker, Stephan Tobies: Truth - A Verification Platform for Distributed Systems

1998-06 * Matthias Oliver Berger: DECT in the Factory of the Future
1998-07 M. Arnold, M. Erdmann, M. Glinz, P. Haumer, R. Knoll, B. Paech, K. Pohl, J. Ryser, R. Studer, K. Weidenhaupt: Survey on the Scenario Use in Twelve Selected Industrial Projects
1998-08 * H. Aust: Sprachverstehen und Dialogmodellierung in natürlichsprachlichen Informationssystemen
1998-09 * Th. Lehmann: Geometrische Ausrichtung medizinischer Bilder am Beispiel intraoraler Radiographien
1998-10 * M. Nicola, M. Jarke: Performance Modeling of Distributed and Replicated Databases
1998-11 * Ansgar Schleicher, Bernhard Westfechtel, Dirk Jäger: Modeling Dynamic Software Processes in UML
1998-12 * W. Appelt, M. Jarke: Interoperable Tools for Cooperation Support using the World Wide Web
1998-13 Klaus Indermark: Semantik rekursiver Funktionsdefinitionen mit Striktheitsinformation
1999-01 * Jahresbericht 1998
1999-02 * F. Huch: Verifcation of Erlang Programs using Abstract Interpretation and Model Checking — Extended Version
1999-03 * R. Gallersdörfer, M. Jarke, M. Nicola: The ADR Replication Manager
1999-04 María Alpuente, Michael Hanus, Salvador Lucas, Germán Vidal: Specialization of Functional Logic Programs Based on Needed Narrowing
1999-05 * W. Thomas (Ed.): DLT 99 - Developments in Language Theory Fourth International Conference
1999-06 * Kai Jakobs, Klaus-Dieter Kleefeld: Informationssysteme für die angewandte historische Geographie
1999-07 Thomas Wilke: CTL+ is exponentially more succinct than CTL
1999-08 Oliver Matz: Dot-Depth and Monadic Quantifier Alternation over Pictures
2000-01 * Jahresbericht 1999
2000-02 Jens Vöge, Marcin Jurdzinski: A Discrete Strategy Improvement Algorithm for Solving Parity Games
2000-04 Andreas Becks, Stefan Sklorz, Matthias Jarke: Exploring the Semantic Structure of Technical Document Collections: A Cooperative Systems Approach
2000-05 Mareike Schoop: Cooperative Document Management
2000-06 Mareike Schoop, Christoph Quix (eds.): Proceedings of the Fifth International Workshop on the Language-Action Perspective on Communication Modelling
2000-07 * Markus Mohnen, Pieter Koopman (Eds.): Proceedings of the 12th International Workshop of Functional Languages
2000-08 Thomas Arts, Thomas Noll: Verifying Generic Erlang Client-Server Implementations
2001-01 * Jahresbericht 2000
2001-02 Benedikt Bollig, Martin Leucker: Deciding LTL over Mazurkiewicz Traces
2001-03 Thierry Cachat: The power of one-letter rational languages

2001-04 Benedikt Bollig, Martin Leucker, Michael Weber: Local Parallel Model Checking for the Alternation Free mu-Calculus

2001-05 Benedikt Bollig, Martin Leucker, Thomas Noll: Regular MSC Languages

2001-06 Achim Blumensath: Prefix-Recognisable Graphs and Monadic Second-Order Logic

2001-07 Martin Grohe, Stefan Wöhrle: An Existential Locality Theorem

2001-08 Mareike Schoop, James Taylor (eds.): Proceedings of the Sixth International Workshop on the Language-Action Perspective on Communication Modelling

2001-09 Thomas Arts, Jürgen Giesl: A collection of examples for termination of term rewriting using dependency pairs

2001-10 Achim Blumensath: Axiomatising Tree-interpretable Structures

2001-11 Klaus Indermark, Thomas Noll (eds.): Kolloquium Programmiersprachen und Grundlagen der Programmierung

2002-01 * Jahresbericht 2001

2002-02 Jürgen Giesl, Aart Middeldorp: Transformation Techniques for Context-Sensitive Rewrite Systems

2002-03 Benedikt Bollig, Martin Leucker, Thomas Noll: Generalised Regular MSC Languages

2002-04 Jürgen Giesl, Aart Middeldorp: Innermost Termination of Context-Sensitive Rewriting

2002-05 Horst Lichter, Thomas von der Maßen, Thomas Weiler: Modelling Requirements and Architectures for Software Product Lines

2002-06 Henry N. Adorna: 3-Party Message Complexity is Better than 2-Party Ones for Proving Lower Bounds on the Size of Minimal Nondeterministic Finite Automata

2002-07 Jörg Dahmen: Invariant Image Object Recognition using Gaussian Mixture Densities

2002-08 Markus Mohnen: An Open Framework for Data-Flow Analysis in Java

2002-09 Markus Mohnen: Interfaces with Default Implementations in Java

2002-10 Martin Leucker: Logics for Mazurkiewicz traces

2002-11 Jürgen Giesl, Hans Zantema: Liveness in Rewriting

2003-01 * Jahresbericht 2002

2003-02 Jürgen Giesl, René Thiemann: Size-Change Termination for Term Rewriting

2003-03 Jürgen Giesl, Deepak Kapur: Deciding Inductive Validity of Equations

2003-04 Jürgen Giesl, René Thiemann, Peter Schneider-Kamp, Stephan Falke: Improving Dependency Pairs

2003-05 Christof Löding, Philipp Rohde: Solving the Sabotage Game is PSPACE-hard

2003-06 Franz Josef Och: Statistical Machine Translation: From Single-Word Models to Alignment Templates

2003-07 Horst Lichter, Thomas von der Maßen, Alexander Nyßen, Thomas Weiler: Vergleich von Ansätzen zur Feature Modellierung bei der Softwareproduktlinienentwicklung

2003-08 Jürgen Giesl, René Thiemann, Peter Schneider-Kamp, Stephan Falke: Mechanizing Dependency Pairs

2004-01 * Fachgruppe Informatik: Jahresbericht 2003

2004-02	Benedikt Bollig, Martin Leucker: Message-Passing Automata are expressively equivalent to EMSO logic
2004-03	Delia Kesner, Femke van Raamsdonk, Joe Wells (eds.): HOR 2004 – 2nd International Workshop on Higher-Order Rewriting
2004-04	Slim Abdennadher, Christophe Ringeissen (eds.): RULE 04 – Fifth International Workshop on Rule-Based Programming
2004-05	Herbert Kuchen (ed.): WFLP 04 – 13th International Workshop on Functional and (Constraint) Logic Programming
2004-06	Sergio Antoy, Yoshihito Toyama (eds.): WRS 04 – 4th International Workshop on Reduction Strategies in Rewriting and Programming
2004-07	Michael Codish, Aart Middeldorp (eds.): WST 04 – 7th International Workshop on Termination
2004-08	Klaus Indermark, Thomas Noll: Algebraic Correctness Proofs for Compiling Recursive Function Definitions with Strictness Information
2004-09	Joachim Kneis, Daniel Mölle, Stefan Richter, Peter Rossmanith: Parameterized Power Domination Complexity
2004-10	Zinaida Benenson, Felix C. Gärtner, Dogan Kesdogan: Secure Multi-Party Computation with Security Modules
2005-01 *	Fachgruppe Informatik: Jahresbericht 2004
2005-02	Maximillian Dornseif, Felix C. Gärtner, Thorsten Holz, Martin Mink: An Offensive Approach to Teaching Information Security: Aachen Summer School Applied IT Security
2005-03	Jürgen Giesl, René Thiemann, Peter Schneider-Kamp: Proving and Disproving Termination of Higher-Order Functions
2005-04	Daniel Mölle, Stefan Richter, Peter Rossmanith: A Faster Algorithm for the Steiner Tree Problem
2005-05	Fabien Pouget, Thorsten Holz: A Pointillist Approach for Comparing Honeypots

* These reports are only available as a printed version.

Please contact biblio@informatik.rwth-aachen.de to obtain copies.