

Approximating Wardrop equilibria with finitely many agents

Simon Fischer · Lars Olbrich · Berthold Vöcking

Received: 1 October 2007 / Accepted: 23 January 2008
© Springer-Verlag 2008

Abstract We present efficient algorithms for computing approximate Wardrop equilibria in a distributed and concurrent fashion. Our algorithms are executed by a finite number of agents each of which controls the flow of one commodity striving to balance the induced latency over all utilised paths. The set of allowed paths is represented by a DAG. Our algorithms are based on previous work on policies for infinite populations of agents. These policies achieve a convergence time which is independent of the underlying network and depends mildly on the latency functions. These policies can neither be applied to a finite set of agents nor can they be simulated directly due to the exponential number of paths. Our algorithms circumvent these problems by computing a randomised path decomposition in every communication round. Based on this decomposition, flow is shifted from overloaded to underloaded paths. This way, our algorithm can handle exponentially large path collections in polynomial time. Our algorithms are stateless, and the number of communication rounds depends polynomially on the approximation quality and is independent of the topology and size of the network.

Keywords Adaptive routing · Wardrop model

1 Introduction

We present distributed algorithms for computing approximate equilibria in the Wardrop model. In this model, each of a set of commodities routes a fixed flow demand from a source to a sink via a collection of paths. The latency induced on the edges depends in a non-decreasing fashion on the edge flows. A flow vector is said to be at a Wardrop equilibrium if for all commodities the latencies of all used paths are minimal with respect to this commodity. It is well known that such equilibria can be computed by centralised algorithms in polynomial time (under some mild assumptions on the latency functions) due to a convex programming formulation. In this work, we study *distributed* algorithms to compute Wardrop equilibria.

A common interpretation of the Wardrop model is that flow is controlled by an infinite number of selfish agents each of which carries an infinitesimal amount of flow. Then, a Wardrop equilibrium corresponds to a state in which no infinitesimal agent can improve their situation unilaterally by migrating from their current path to another one. In [13] it was shown that in this setting such a population can approach Wardrop equilibria quickly by following a simple round-based load-adaptive rerouting policy. This policy, called the *replication policy*, is executed by all agents in parallel and proceeds in the following way. Each agent performs two simple selfish steps. First, they sample another agent uniformly at random and, second, if this improves the latency, the agent migrates from their old path to the sampled agent's path with a certain probability that increases with the latency gain. More precisely, it is chosen to be precisely the ratio between the projected latency gain and the agent's old latency.

Supported by DFG grant Vo889/1–3 and by DFG through German excellence cluster UMIC at RWTH Aachen.
Supported by the DFG GK/1298 “AlgoSyn”.
Supported in part by the the EU within the 6th Framework Programme under contract 001907 (DELIS) and by DFG through German excellence cluster UMIC at RWTH Aachen.
A preliminary version of this work has appeared in [14].

S. Fischer (✉) · L. Olbrich · B. Vöcking
RWTH Aachen, Aachen, Germany
e-mail: fischer@cs.rwth-aachen.de

L. Olbrich
e-mail: lars@cs.rwth-aachen.de

B. Vöcking
e-mail: voecking@cs.rwth-aachen.de

In this setting, a natural goal is to reach approximate equilibria in the following bicriterial sense. We say that a flow is at δ - ϵ -equilibrium if at most an ϵ -fraction of the flow utilises paths whose latency exceeds the average latency of their commodity by more than a δ -fraction of the overall average latency. Remarkably, the number of rounds to reach an approximate equilibrium in this sense is independent of the size and the topology of the underlying graph and chiefly depends on the approximation parameters and the elasticity of the latency functions.

In this work, we consider a different setting, in which the flow is controlled by a finite number of agents only, each of which is responsible for the entire flow of one commodity. Each agent has a set of admissible paths among which it may distribute its flow. To be able to represent exponentially large collections of paths we assume that these are represented by an arbitrary DAG connecting the source and the sink of the agent. Each agent aims at balancing its own flow such that the jointly computed allocation will be at Wardrop equilibrium. Let us remark that agents do not aim at minimising the overall latency of their flow, but seek to balance the latencies of their respective commodity. Unfortunately, the replication policy does not yield a feasible distributed algorithm in this setting directly. Simulating an infinite number of agents each of which chooses one out of the given collection of paths would require maintaining one variable for each path and computing a quadratic number of migration rates between pairs of paths. As the number of paths may be exponential in the size of the network this approach is rendered computationally infeasible.

We present two approaches to circumvent this problem. Our first approach exploits the fact that, for a simplified variant of the replication policy, the updates of the edge flows can be expressed in a way that merely uses the edge flow variables themselves (rather than the path flow variables). Thus, the updates can be computed in polynomial time. Unfortunately, the convergence time of this variant is only pseudopolynomial in the latency functions since it depends on the maximum slope of the latency functions.

Since the original replication policy cannot be expressed in this compact way, we consider a second approach to achieve convergence in a polynomial number of communication rounds. Consider a collection of paths for one of the commodities. In a first step, our algorithm samples a polynomial number of paths with probability proportional to their flow (which can be derived implicitly from the edge flows since the paths are represented by a DAG). We thus obtain a randomised path decomposition. We consider paths in this decomposition with above-average latency. From such paths, a fraction of the flow is removed and reallocated proportionally among all admissible paths. If this is done carefully, oscillations can be avoided, and a potential function argument ensures convergence towards Wardrop equilibria. Thus,

we achieve essentially the same convergence rates as in the setting with an infinite number of agents and keep the computation time of one communication round polynomial. Altogether, we can compute approximate Wardrop equilibria in polynomial time.

1.1 Related work

The game theoretic traffic model considered in this paper was introduced by Wardrop [21]. Many aspects of Wardrop equilibria have been studied, the most prominent being the degradation of performance due to the selfish behaviour, called the price of anarchy [20] as well as the inverse, the increase of the maximum latency incurred to an agent due to optimal routing [19]. It has also been shown that the price of anarchy can be decreased by imposing taxes on the edges [9, 15]. Cominetti et al. consider the price of anarchy in a model with finitely many agents aiming at minimising their average latency [10].

For solving the corresponding classical goal of finding a minimum cost multi-commodity flow, several algorithms are known. For an overview see, e.g., [1, 7]. An efficient distributed steepest-descent algorithm for solving multi-commodity flow problems with linear latency functions has been presented recently in [3]. In [2], a stateless algorithm for this problem is presented.

It is also known that the problems of finding an optimal allocation and finding a Wardrop equilibrium are essentially equivalent. Under mild conditions on the latency functions, a flow at Wardrop equilibrium with respect to so-called *marginal-cost* latency functions is optimal with respect to the original latency functions, see e.g. [5, 20].

Several authors (e.g. [4, 8]) consider dynamic routing from an online-learning perspective. Awerbuch and Kleinberg [4] present an algorithm for the online shortest path problem in an end-to-end feedback model. Blum et al. [8] show that approximate Wardrop equilibria defined in a similar way can be attained if the agents follow no-regret algorithms. Their bounds on the convergence time depend polynomially on the regret bounds and network size and depend pseudopolynomially on the maximum slope of the latency functions.

The problem of load-balancing has also been studied in various discrete settings for networks of parallel links. For the case of identical links, both sequential [16] and concurrent distributed algorithms were considered [6]. Even-Dar et al. [11] consider distributed algorithms for load balancing on links with speeds using sampling rules which depend pseudopolynomially on the speed of the links.

The rerouting policy upon which our algorithms are based was introduced in [14] and [13]. It was shown that an infinite number of agents executing this policy can attain a Wardrop equilibrium quickly in a concurrent setting.

2 Model and problem statement

2.1 Wardrop's traffic model

We consider Wardrop's traffic model originally introduced in [21]. We are given a graph $G = (V, E)$ with non-decreasing differentiable latency functions $\ell_e : [0, 1] \rightarrow \mathbb{R}_{\geq 0}$. Furthermore, we are given a set of commodities $[k] = \{1, \dots, k\}$ specified by source-sink-pairs $(s_i, t_i) \in V \times V$, a directed acyclic subgraph G_i of G connecting s_i and t_i and flow demands $r_i \in \mathbb{R}_{\geq 0}$. The total demand is $r = \sum_{i \in [k]} r_i$, and we normalise $r = 1$ for simplicity. Let \mathcal{P}_i denote the admissible paths of commodity i , i.e., all paths connecting s_i and t_i in G_i , and let $\mathcal{P} = \bigcup_{i \in [k]} \mathcal{P}_i$. We may assume that the sets \mathcal{P}_i are disjoint and define i_P to be the unique commodity to which path P belongs.

A non-negative path flow vector $(f_P)_{P \in \mathcal{P}}$ is *feasible* if it satisfies the flow demands $\sum_{P \in \mathcal{P}_i} f_P = r_i$ for all $i \in [k]$. We denote the set of all feasible flow vectors by \mathcal{F} . A path flow vector $(f_P)_{P \in \mathcal{P}}$ induces an edge flow vector $f = (f_{e,i})_{e \in E, i \in [k]}$ with $f_{e,i} = \sum_{P \in \mathcal{P}_i: e \in P} f_P$. The total flow on edge e is $f_e = \sum_{i \in [k]} f_{e,i}$. Furthermore, for $v \in V$ and $i \in [k]$, the total flow of commodity i through node v is $f_{v,i} = \sum_{(u,v) \in E} f_{(u,v),i} = \sum_{(v,w) \in E} f_{(v,w),i}$ for $v \notin \{s_i, t_i\}$ and $f_{s_i,i} = f_{t_i,i} = r_i$. The latency of an edge $e \in E$ is given by $\ell_e(f_e)$ and the latency of a path P is given by the sum of the edge latencies $\ell_P(f) = \sum_{e \in P} \ell_e(f_e)$. Finally, the weighted average latency of commodity $i \in [k]$ is given by $L_i(f) = \sum_{e \in E} \ell_e(f) \cdot (f_{e,i}/r_i)$ and the overall average latency is $L(f) = \sum_{e \in E} \ell_e(f) \cdot f_e/r$. We drop the argument f of $\ell(\cdot)$ and $L(\cdot)$ whenever it is clear from the context.

A flow vector in this model is considered stable when no fraction of the flow can improve its sustained latency by moving unilaterally to another path. This implies that all used paths of a commodity must have the same minimal latency. Unused paths may have larger latency.

Definition 1 (Wardrop equilibrium) A feasible flow vector f is at *Wardrop equilibrium* if for every commodity $i \in [k]$ and paths $P_1, P_2 \in \mathcal{P}_i$ with $f_{P_1} > 0$ it holds that $\ell_{P_1}(f) \leq \ell_{P_2}(f)$.

It is well-known that Wardrop equilibria are exactly those allocations that minimise the following potential function introduced in [5]:

$$\Phi(f) = \sum_{e \in E} \int_0^{f_e} \ell_e(u) du.$$

This potential precisely absorbs progress: If an infinitesimal volume of flow dx is shifted from path ℓ_P to ℓ_Q , thus improving its latency by $(\ell_P - \ell_Q)$, the potential decreases by $(\ell_P - \ell_Q) dx$. We will make use of this fact frequently. The minimum potential is denoted by $\Phi^* = \min_{f \in \mathcal{F}} \Phi(f)$.

Every flow vector f with $\Phi(f) = \Phi^*$ is then at Wardrop equilibrium. We assume that Φ^* is positive. The case that $\Phi^* = 0$ can be treated by adding virtual offsets to the latency functions. For a detailed treatment see [13].

Let us remark, that the problems of computing a Wardrop equilibrium and computing a flow minimising L are equivalent. It is sufficient to replace the latency functions ℓ_e by so-called *marginal-cost* latency functions $h_e(x) = (x \cdot \ell_e(x))' = \ell_e(x) + x \cdot \ell_e'(x)$. If for all $e \in E$, $x \cdot \ell_e(x)$ is convex, then Wardrop equilibria with respect to $(h_e)_{e \in E}$ minimise L [5,20].

The algorithms presented in this paper will compute approximate equilibria in the following bicriterial sense.

Definition 2 (δ - ϵ -equilibrium) Consider a unit flow vector f and let $\mathcal{P}_i^\delta = \{P \in \mathcal{P}_i \mid \ell_P(f) > L_i(f) + \delta L(f)\}$ denote the set of δ -expensive paths of commodity $i \in [k]$. A flow vector is at a δ - ϵ -equilibrium if $\sum_{i \in [k]} \sum_{P \in \mathcal{P}_i^\delta} f_P \leq \epsilon$.

This definition of approximate Wardrop equilibria requires that almost all flow utilises paths with a latency that is close to the average of their own commodity. A similar definition of approximate Nash equilibria is used, e.g., in [8].

2.2 Elasticity of latency functions

Our algorithms take the steepness of the latency functions into account when deciding how much flow to shift from one path to another. In [13] it was shown that the critical parameter in this setting is not the slope but the elasticity.

Definition 3 For any positive differentiable function $\ell : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$, the *elasticity* of ℓ at x is $d(x) = \frac{x \cdot \ell'(x)}{\ell(x)}$.

In other words, the elasticity of a function is bounded from above by d if the (absolute) slope at any point is at most by a factor of d larger than the slope of the line connecting the origin and the point $(x, \ell(x))$. Note that a polynomial with positive coefficients and degree d has elasticity at most d , hence, elasticity can be considered as a generalisation of the degree of such a polynomial. The function $a \cdot \exp(\lambda x)$, $x \in [0, 1]$ has maximum elasticity λ .

2.3 Implicit path decomposition

Wardrop equilibria are defined with respect to path flows. Our algorithms, however, will make use only of the edge flow vectors, which do not determine a vector of path flows uniquely. However, in a DAG, an edge flow vector $(f_e)_{e \in E}$ induces a natural vector of path flows by starting with the flow injected at the source, and splitting the flow at each node v such that the set of paths containing the outgoing edge e receives a flow proportional to f_e . Since the decomposition for one commodity $i \in [k]$ is independent of the flow of other commodities, we can omit the index i for simplicity.

Definition 4 Consider any edge flow vector $(f_e)_{e \in E}$ (for some commodity i). For any path $P = (v_1, \dots, v_l)$ let

$$\tilde{f}_P = f_{v_1} \cdot \prod_{j=1}^{l-1} \frac{f_{(v_j, v_{j+1})}}{f_{v_j}}.$$

It is easily verified by induction on the distance from the source that this is actually a valid flow decomposition of $(f_e)_{e \in E}$, i.e., $f_e = \sum_{P \ni e} \tilde{f}_P$.

2.4 Distributed computation model

Our algorithms operate in the following setting. Agents operate in a synchronous, round-based fashion. We assume that there is a billboard via which the agents are able to share information. On this billboard, each agent can observe the edge flows of its own commodity and the latency values of the paths it uses. Agents know an upper bound d on the elasticity of the latency functions, but they do not know the latency functions themselves. However, it is easily possible to extend our algorithm such that it does not rely on the knowledge of a bound on the elasticity.

In every round an agent can update the edge flows of its own commodity on the billboard. These updates become visible to all agents only in the next round. All agents execute the same algorithm in parallel. Therefore, in the descriptions of our algorithms, we may omit the index for the commodity, i.e., f_e refers to the flow $f_{e,i}$ of commodity i on edge e .

Let us remark that the billboard is a purely theoretical construction. It may model systems where this information is collected centrally and polled by or broadcast to the agents at intervals, but may also model scenarios in which agents concurrently perform measurements over finite intervals of time in order to obtain the necessary latency information.

3 A pseudopolynomial algorithm

Our first approach works by simulating the replication policy presented in [13]. We will see that this can be done in polynomial time although this policy operates on an exponential number of paths.

3.1 The replication policy

Let us start by introducing the replication policy formally. We consider an infinite population of agents each of which controls an infinitesimal amount of flow which it assigns to a path. In each round agents may migrate their flow from the current path to another one. Consider an agent in commodity $i \in [k]$ currently using path $P \in \mathcal{P}_i$. Whenever activated, it performs two steps.

1. *Sampling.* Sample another path Q where the probability to sample any path Q' equals $f_{Q'}/r_i$.
2. *Migration.* There are two cases:
 - (a) $\ell_Q \geq \ell_P$. In this case, the agent stays with its old path.
 - (b) $\ell_Q < \ell_P$. The agent migrates to the sampled path Q with probability $\lambda \cdot (\ell_P - \ell_Q)$ for some constant $\lambda > 0$ to be determined later.

Altogether, we can characterise our policy by specifying the rate of agents migrating from one path $P \in \mathcal{P}_i$ to another path $Q \in \mathcal{P}_i$ with $\ell_Q(f) < \ell_P(f)$ within one round. This rate can be obtained by multiplying the probabilities specified in steps (1) and (2) with the volume of agents using path P . For this rate we obtain

$$\rho_{PQ} = \lambda \cdot f_P \cdot \frac{f_Q}{r_i} \cdot (\ell_P - \ell_Q)$$

if $\ell_Q < \ell_P$ and $\rho_{PQ} = 0$ otherwise. Thus, we can compute a sequence of flow vectors $(f_P(t))_{P \in \mathcal{P}}$ generated by this policy by summing over all paths Q :

$$\begin{aligned} f_P(t+1) &= f_P(t) + \sum_{Q \in \mathcal{P}_i} \rho_{QP} - \sum_{Q \in \mathcal{P}_i} \rho_{PQ} \\ &= f_P(t) + \lambda f_P \sum_{Q \in \mathcal{P}_i} \frac{f_Q}{r_i} (\ell_Q - \ell_P) \\ &= f_P(t) + \lambda f_P (L_i - \ell_P). \end{aligned} \quad (1)$$

3.2 Convergence towards equilibria

For the time being assume that agents are migrating in a continuous fashion as described by the above rules. Then, an infinitesimal amount of flow dx migrating from a path P to another path Q improving its latency from ℓ_P to ℓ_Q causes the potential Φ to reduce by $(\ell_P - \ell_Q) dx$. Since we only accept migrations that improve the latency, this implies that the potential always decreases which in turn implies convergence towards a Wardrop equilibrium by Lyapunov's direct method¹ if all paths are used in the initial flow. However, in our concurrent round-based model, flow is not shifted continuously, but in finite chunks. Thus, if these chunks are chosen too large, overshooting and oscillation effects may occur. This issue can be resolved by choosing the migration rate in step 2(b) of the replication policy carefully. In [14] it was

¹ Lyapunov's theorem says that the solution of a dynamic system $\dot{x} = v(x)$ converges towards a fixed point \mathbf{x}_0 if there exists a continuously differentiable function H with the property that \mathbf{x}_0 minimises H and $\dot{H}(\mathbf{x})$ is strictly negative for all $\mathbf{y} \neq \mathbf{x}$. In our case, the potential function Φ satisfies this property.

shown that if we choose $\lambda = \Theta(1/\ell'_{\max})$ small enough with

$$\ell'_{\max} = \max_{P \in \mathcal{P}} \max_{f \in \mathcal{F}} \sum_{e \in P} \ell'_e(f),$$

convergence towards Wardrop equilibria can be guaranteed (provided that initially all paths have non-zero flow and hence positive sampling probability). We may assume that $\ell'_{\max} > 0$ since otherwise all latency functions are constant and our problem can be solved trivially by assigning the entire flow to the path with lowest latency.

Theorem 1 ([12–14]) *If $\lambda = \Theta(1/\ell'_{\max})$ sufficiently small, the replication policy given by Eq. (1) with initial flow $f(0) = f^0$ converges towards a Wardrop equilibrium if $f^0_P > 0$ for all $P \in \mathcal{P}$. Furthermore, the number of rounds in which the flow is not at a δ - ϵ -equilibrium is*

$$\mathcal{O}\left(\frac{1}{\epsilon^2 \delta^2} \cdot \frac{\ell'_{\max}}{\ell_{\min}} \cdot \log\left(\frac{\Phi(f^0)}{\Phi^*}\right)\right).$$

One may observe that the ratio between maximum slope and minimum latency used in this theorem depends on the scale by which we measure flow. This scale, however, is fixed since we have normalised the total flow demand to be $r = 1$.

3.3 Simulating the replication policy

By a naive application of Theorem 1 we can compute a sequence of flow vectors $(f(t))_{t \geq 0}$ according to Eq. (1) to obtain approximate Wardrop equilibria. However, this approach is rendered computationally intractable by the fact that there may be an exponential number of variables f_P .

In the following, we describe an algorithm that computes the iterative change rates of the edge flows according to the implicit flow decomposition \tilde{f} described in the preceding section. To that end, we show that the change rates of the edge flows f_e can be expressed solely in terms of edge flows and edge latencies (i.e., without explicit reference to the f_P variables). It suffices to know the weighted average latencies of all paths containing e defined as

$$L_e = \sum_{P \ni e} \frac{f_P}{f_e} \cdot \ell_P \tag{2}$$

Recall that we have fixed a commodity here, so we may drop the index i .

Lemma 1 *Consider an edge flow vector $(f_e(t))_{e \in E}$ and its path decomposition $\tilde{f}(t)$, and let $\tilde{f}(t+1)$ denote the flow generated by the replication policy in Eq. (1) from $\tilde{f}(t)$. Finally, let $f_e(t+1) = \sum_{P \ni e} \tilde{f}_P(t+1)$. Then,*

$$f_e(t+1) = f_e(t) + \lambda \cdot f_e(t) \cdot (L - L_e).$$

Proof Let $f = f(t)$ and $f' = f(t+1)$. By definition of f_e ,

$$\begin{aligned} f'_e - f_e &= \sum_{P \ni e} (f'_P - f_P) \\ &= \lambda \cdot \sum_{P \ni e} f_P \cdot (L - \ell_P) \\ &= \lambda \cdot f_e \cdot \left(L - \frac{\sum_{P \ni e} f_P \ell_P}{f_e} \right), \end{aligned}$$

where the last term equals L_e . □

In order to obtain the value of L_e , we implicitly compute the path decomposition \tilde{f} , i.e., for every edge e' we compute the flow caused by paths containing e on edge e' . This is done by Algorithm SIMULATED REPLICATION (Algorithm 1) in time $\mathcal{O}(m)$ for every edge $e \in E$. Since there are m edges, each iteration can be performed in time $\mathcal{O}(m^2)$.

Algorithm 1 SIMULATED REPLICATION() (executed by all commodities in parallel; $(f_e)_{e \in E}$ denotes the edge flows vector of commodity i)

- 1: **for** all edges $e \in E$ **do**
 - 2: sort all edges (v, w) in the subgraph reachable from e topologically
 - 3: compute total flow of all paths containing e and (v, w) :
 $\tilde{f}_e^{(v,w)} = \sum_{(u,v) \in E} \tilde{f}_e^{(u,v)} \cdot \frac{f_{(v,w)}}{f_v}$
 - 4: reverse all edges and repeat steps 2 and 3 for edges between e and s
 - 5: compute $L_e = \sum_{e'} \frac{\tilde{f}_e^{e'}}{f_e} \ell_{e'}$
 - 6: $f'_e \leftarrow f_e + \lambda \cdot f_e \cdot (L - L_e)$ with $\lambda = 1/\ell'_{\max}$
 - 7: **end for**
 - 8: replace $(f_e)_{e \in E}$ on the billboard with $(f'_e)_{e \in E}$
-

Corollary 1 *The sequence of flow vectors computed by Algorithm SIMULATED REPLICATION converges towards the set of Wardrop equilibria. Furthermore, the number of rounds in which the flow is not at a δ - ϵ -equilibrium with respect to \tilde{f} , is bounded by*

$$\mathcal{O}\left(\frac{1}{\epsilon^2 \delta^2} \cdot \frac{\ell'_{\max}}{\ell_{\min}} \cdot \log\left(\frac{\Phi(f^0)}{\Phi^*}\right)\right),$$

where f^0 is the initial flow vector. Each iteration takes time $\mathcal{O}(m^2)$.

Proof First note that for any edge e' , the value of $\tilde{f}_e^{e'}$ computed in line 3 of the algorithm equals the volume of all paths containing e and e' with respect to our implicit decomposition \tilde{f} , i.e., $\tilde{f}_e^{e'} = \sum_{P \ni \{e, e'\}} \tilde{f}_P$. Thus, the value L_e computed in line 5 equals the definition of L_e in Eq. (2). Hence, Lemma 1 implies that the edge flow vector computed by our algorithm after one round equals the edge flow vector obtained by applying the replication policy given by Eq. (1) to the path decomposition $(\tilde{f})_{P \in \mathcal{P}}$. Combining this with the upper bounds on the convergence time given in [13, 12], the claim follows. □

4 The polynomial time algorithm

The migration probability specified for step 2(b) of the replication policy can get very small since the latency difference $\ell_P - \ell_Q$ may become small in relation to ℓ'_{\max} if $\lambda = 1/\ell'_{\max}$ is chosen constant. This causes the algorithm to obtain only a pseudopolynomial convergence time depending on the maximum slope of the latency functions. In this section we present an approach that gets rid of this dependence.

To this end, we choose the amount of flow removed from a path proportional to its relative deviation $(\ell_P - L_{i_P})/\ell_P$ from the average and the reciprocal of the elasticity d to obtain a polynomial number of communication rounds. Whereas in the preceding section the amount of flow removed or added to a path within one round could be expressed in a nice closed form as $\lambda \cdot f_P \cdot (L - \ell_P)$ [Eq. (1)], this is now no longer possible.

To compute flow updates in polynomial time we use a randomised flow decomposition. First, we sample a path at random according to the implicit path decomposition \tilde{f} , i.e., the probability to sample path P is \tilde{f}_P/r_{i_P} . Since the length of a path is bounded by n , this is possible in time $n \log n$ by representing adjacent nodes and their flows in a binary tree. Now, the path is assigned a certain flow volume f_P . For the time being, assume that we assign the entire bottleneck flow to P . Then, if P has latency above L_{i_P} , we remove a portion of

$$x = \Theta \left(f_P \cdot \frac{\ell_P - L_{i_P}}{d \ell_P} \right)$$

of its flow and distribute it proportionally among all admissible paths, i.e., after removing a flow of x from path P , the flow on every edge $e \in E$ is increased by $(f_{e,i}/r_i) \cdot x$. Thus, the computed flow remains feasible.

Why does this process decrease the potential quickly? As long as we are not at a δ - ϵ -equilibrium, the probability of sampling a δ -expensive path is at least ϵ . In this case, the latency gain and thus the potential gain *per flow unit* will be large and proportional to \tilde{f}_P . If we sample only a single path, we may in fact assign the entire bottleneck flow to it. We can lower bound the probability that this bottleneck flow is not too small (Lemma 2). To increase the potential gain further, we repeat this process T times. Doing this, we can no longer assign the entire bottleneck flow to a path since it may happen that an edge is sampled several times. Consider an edge e . If this edge is sampled k times, we may consume at most f_e/k of its flow in every round. If $f_e = \Theta(1)$, we will have $k = \Theta(T)$, so in this case we have limit the amount of flow consumed in one round to $f_e/k = \Theta(1/T)$. For edges with less flow, however, we may consume more than f_e/k per round, since these edges are sampled less often. It turns out that we can increase the potential gain by a factor of $\Omega(m)$ if we choose $T = \Theta(m \log m)$ and set the flow assigned to

a sampled path to an $\Theta(1/\log m)$ fraction of the bottleneck flow. More precisely, let

$$\Delta_e = \min \left\{ \frac{1}{7m \log m}, \frac{f_e}{7 \log m} \right\}.$$

We start with an empty decomposition. In a round in which path P is sampled we increase f_P by Δ_{e^*} where e^* is a bottleneck edge in P . We say that an edge is *alive* if the overall flow assigned to paths containing e is at most $f_e - \Delta_e$ (i.e. there is still a flow of Δ_e remaining, so it can safely be sampled one more time without having our decomposition exceeding the flow of e). Our algorithm terminates as soon as there are any edges that are not alive. The final algorithm RANDOMISED BALANCING (d) is described in Algorithm 2.

Under the assumption that the latency functions are constant, we can thus show that the potential decreases in every round by a factor that only depends on ϵ and δ , and the elasticity d (Lemma 5). We furthermore show that due to our careful migration rate the potential gain with respect to the true latency functions is still at least half of the potential gain with respect to constant latencies (Lemma 4). Finally, we show that the expected potential gain implies a bound on the time to reach a minimum potential (Lemma 7). Altogether, this yields the following upper bound for our algorithm.

Theorem 2 *Assume that for the initial flow vector f_0 it holds that $f_e > 0$ for all $e \in E$. Then, the sequence of flow vectors computed by Algorithm RANDOMISED BALANCING converges towards the set of Wardrop equilibria. Furthermore, the expected number of rounds in which the flow is not at a δ - ϵ -equilibrium with respect to \tilde{f} , is bounded by*

$$\mathcal{O} \left(\frac{d}{\epsilon^3 \delta^2} \log \left(\frac{\Phi(f_0)}{\Phi^*} \right) \right),$$

if d is an upper bound on the elasticity of the latency functions. The computation time of each round is bounded by $\mathcal{O}(n \log n \cdot m \log m)$.

We present the proof after establishing the necessary lemmas.

Note that our algorithm can be easily modified for the case that the elasticity of the latency functions is not known to the algorithm in advance. In that case, we can find an upper bound d on the maximum elasticity by using an exponential search technique. We continue doubling the value of d until for the first time it holds that for all edges relevant to the commodity, the elasticity of the latency functions is bounded by d within the interval defined by the old and new flow values.

4.1 Randomised decomposition

Our algorithm generates a randomised flow decomposition using a sampling process based on \tilde{f} . In this section, we lower bound the probability that the bottleneck flows of the

Algorithm 2 RANDOMISED BALANCING (d) (executed by all commodities in parallel; $(f_e)_{e \in E}$ denotes the edge flows vector of commodity i)

```

1:  $F \leftarrow 0$ 
2: for  $T = m \log m$  times do
3:   sample a path  $P$  where  $\mathbb{P}[P] = \frac{\tilde{f}_P}{r_i}$ 
4:   let  $e^*$  denote the bottleneck edge of  $P$ ; let  $f_P = \Delta_{e^*}$ 
5:   if  $\ell_P > L_i$  then
6:     reduce the flow on all edges  $e \in P$  by  $\Delta f_P = f_P \cdot \frac{\ell_P - L_i}{4d \ell_P}$ 
7:      $F \leftarrow F + \Delta f_P$ 
8:     if for any  $e \in P$ ,  $e$  is not alive then
9:       abort loop and continue in line 2
10:    end if
11:  end if
12: end for
13: increase the flow on all edges  $e \in E$  proportionally by  $\frac{f_e}{r_i} \cdot F$ 

```

sampled paths are not too small. Furthermore, we show that the total flow volume removed from every edge is at most f_e with high probability.

For a unit flow f , let $\mathbb{P}_{P \sim f}[X(P)]$ denote the probability that event $X(P)$ occurs if the probability to draw a path P equals f_P .

Lemma 2 Consider a flow vector f of volume 1 and a set of paths \mathcal{P}_ϵ with $\sum_{P \in \mathcal{P}_\epsilon} \tilde{f}_P = \epsilon$. Then, $\mathbb{P}_{P \sim \tilde{f}}[P] \in \mathcal{P}_\epsilon \wedge \min_{e \in P} f_e \geq \frac{\epsilon}{2m} \geq \frac{\epsilon}{2}$.

Proof We consider a scaled flow vector which supports only paths in \mathcal{P}_ϵ .

$$f'_P = \begin{cases} \frac{\tilde{f}_P}{\epsilon} & P \in \mathcal{P}_\epsilon \\ 0 & P \notin \mathcal{P}_\epsilon \end{cases}$$

Observe that the total volume of f' is 1 again, hence $\mathbb{P}_{P \sim f'}[P = Q] = \mathbb{P}_{P \sim \tilde{f}}[P = Q \mid P \in \mathcal{P}_\epsilon]$. Now,

$$\begin{aligned} & \mathbb{P}_{P \sim \tilde{f}} \left[P \in \mathcal{P}_\epsilon \wedge \min_{e \in P} f_e \geq \frac{\epsilon}{2m} \right] \\ &= \mathbb{P}_{P \sim \tilde{f}} [P \in \mathcal{P}_\epsilon] \cdot \mathbb{P}_{P \sim \tilde{f}} \left[\min_{e \in P} f'_e \geq \frac{1}{2m} \mid P \in \mathcal{P}_\epsilon \right] \\ &= \epsilon \cdot \mathbb{P}_{P \sim f'} \left[\min_{e \in P} f'_e \geq \frac{1}{2m} \right], \end{aligned}$$

where the first equality uses the definition of f' and the second one uses the above observation. It remains to show that

$$\mathbb{P}_{P \sim f'} \left[\min_{e \in P} f'_e \geq \frac{1}{2m} \right] \geq 1/2. \tag{3}$$

To see this, let $d(x, y)$ denote the number of edges of a shortest path connecting x and y . We can show that $\mathbb{P}[e = (v, w) \in P] = f_e$ by induction on $d(s, v)$. This holds for $d(s, v) = 0$ by definition of \tilde{f} . Now, assume that the statement holds for all edges (u, v) with $d(s, u) = k$ and consider

an edge $e = (v, w)$ with $d(s, v) = k + 1$.

$$\begin{aligned} \mathbb{P}[e \in P] &= \mathbb{P}[v \in P] \cdot \mathbb{P}[e \in P \mid v \in P] \\ &= \sum_{(u,v)} \mathbb{P}[(u, v) \in P] \cdot \frac{f_e}{f_v} \\ &= \sum_{(u,v)} f_{(u,v)} \cdot \frac{f_e}{f_v} = f_e. \end{aligned}$$

With $E' = \{e \in E \mid f_e \leq 1/(2m)\}$,

$$\mathbb{P}[P \ni e : e \in E'] \leq \sum_{e \in E'} \mathbb{P}[e \in P] \leq \sum_{e \in E'} f_e \leq \frac{|E'|}{2m} \leq \frac{1}{2}.$$

Thus, Eq. (3) holds which completes the proof. \square

We now consider a sequence of $T = m \log m$ rounds. Observe that Δ_e is an upper bound on the flow removed from a path containing e by our algorithm, since for the bottleneck edge e^* , $\Delta_{e^*} = \min_{e \in P} \{\Delta_e\}$.

The flow on e may decrease to below zero only if it is contained in the sampled path at least f_e/Δ_e times. In the following we show that this is unlikely.

Lemma 3 With probability $1 - o(1)$, after a sequence of $T = m \log m$ iterations, all edges are still alive.

Proof In the proof of Lemma 2 we have seen that the probability to hit edge e in one round equals f_e . Let the random variable X denote the number of hits in T rounds. We have $\mathbb{E}[X] = T f_e$. An edge is alive if $X \leq f_e/\Delta_e - 1$. There are two cases:

1. $f_e < \frac{1}{m}$ implying $\Delta_e = f_e/(7 \log m)$. Then,

$$\begin{aligned} \mathbb{P} \left[X > \frac{f_e}{\Delta_e} - 1 \right] &= \mathbb{P} \left[X > \mathbb{E}[X] \cdot \left(\frac{7}{f_e m} - \frac{1}{T f_e} \right) \right] \\ &\leq \mathbb{P} \left[X > \mathbb{E}[X] \cdot \left(\frac{6}{f_e m} \right) \right] \\ &\leq 2^{-\mathbb{E}[X] \cdot \frac{6}{f_e m}} = m^{-6}. \end{aligned}$$

The first inequality is the definition of T and Δ_e and uses our assumption that $f_e \cdot m < 1$, and the second inequality is Chernoff's inequality (which asserts that $\mathbb{P}[X \geq r \cdot \mathbb{E}[X]] \leq 2^{-r \cdot \mathbb{E}[X]}$ for $r \geq 6$ for a random variable X that is the sum of 0-1 random variables, see [17]).

2. $f_e \geq \frac{1}{m}$ implying $\Delta_e = 1/(7T)$. Then, with the same arguments,

$$\begin{aligned} \mathbb{P} \left[X > \frac{f_e}{\Delta_e} - 1 \right] &= \mathbb{P} \left[X > \mathbb{E}[X] \cdot \left(\frac{1}{T \Delta_e} - \frac{1}{T f_e} \right) \right] \\ &\leq \mathbb{P} \left[X > \mathbb{E}[X] \cdot \left(7 - \frac{1}{\log(m)} \right) \right] \\ &\leq 2^{-6 \cdot \mathbb{E}[X]} \leq 2^{-6 \cdot \log m} = m^{-6}. \end{aligned}$$

This time, the first inequality uses our assumption $f_e \geq 1/m$.

In both cases, the probability that edge e is not alive at the end of a sequence of T iterations is bounded by m^{-6} . Using a union bound, the probability that at least one edge does not survive is at most m^{-5} and consequently the probability that all edges survive the sequence is at least $1 - m^{-5}$. \square

4.2 Lower bounding the potential gain

We use a potential function argument to prove convergence. In order to show that our algorithm avoids oscillations, we consider the potential gain achieved within one round. We show that this potential gain is at least half of the potential gain that would occur if latencies values were fixed at the beginning of a round. A second lemma shows that, in expectation, the potential decreases by a factor in every round, as long as we are not yet at an approximate equilibrium.

Lemma 4 *Let d denote an upper bound on the elasticity of the latency functions. For a flow vector f consider a flow vector f' generated by Algorithm RANDOMISED BALANCING (d) (Algorithm 2) with positive probability. For any $P \in \mathcal{P}$ let Δf_P denote the amount of flow removed from path P . Then, $\Phi(f) - \Phi(f') \geq \frac{1}{2} \cdot \sum_{P \in \mathcal{P}} (\ell_P(f) - L_{i_P}) \cdot \Delta f_P$.*

To prove the lemma, we need the following fact about functions with bounded elasticity:

Fact 1 ([13]) *If the elasticity of a function ℓ is bounded from above by d , then for $0 \leq \delta \leq 1/(2d)$, it holds that $\ell((1 + \delta) \cdot x) \leq (1 + 2d\delta) \cdot \ell(x)$.*

Proof of Lemma 4 Throughout this proof, whenever we write ℓ and L without an argument, we refer to $\ell(f)$ and $L(f)$. Let $V_P = (\ell_P - L_{i_P}) \Delta f_P$ denote the virtual potential gain of any path P with $\ell_P \geq L_i$. It can easily be checked (see [14]) that

$$\Phi(f) - \Phi(f') = \sum_{P \in \mathcal{P}} V_P - \sum_{e \in E} \int_{f_e}^{f'_e} (\ell_e(u) - \ell_e) du.$$

The true potential gain $\Phi(f) - \Phi(f')$ would be achieved if the latency values did not change due to the changing flow. The terms in the second sum can be understood as *error terms* that reduce the potential gain to account for this idealistic assumption. We show that the error terms are at most half of the virtual potential gain. To that end, we attribute

$$V_P^e = V_P \cdot \frac{f_{e,i_P} \cdot \ell_e}{2 \cdot r_{i_P} \cdot L} = (\ell_P - L_{i_P}) \cdot \Delta f_P \cdot \frac{f_{e,i_P} \cdot \ell_e}{2 \cdot r_{i_P} \cdot L}$$

of the virtual potential gain made by path P to edge e . Note that summing over all edges e , this consumes precisely half of

the virtual potential gain of path P , i.e., $\sum_{e \in E} V_P^e = V_P/2$, thus by a reordering of the terms,

$$\Phi(f) - \Phi(f') \leq \sum_{P \in \mathcal{P}} \frac{V_P}{2} + \sum_{e \in E} \left(\sum_{P \in \mathcal{P}} V_P^e - \int_{f_e}^{f'_e} (\ell_e(u) - \ell_e) du \right).$$

Hence, to prove the theorem it suffices to show that for any edge $e \in E$,

$$\int_{f_e}^{f'_e} (\ell_e(u) - \ell_e) du \leq \sum_{P \in \mathcal{P}} V_P^e. \tag{4}$$

Fix an edge $e \in E$ and assume $f'_e > f_e$. Edges with $f'_e < f_e$ can be treated symmetrically. We partition the integral over the interval $[f_e, f'_e]$ into segments of width Δf_P^e where $\Delta f_P^e = \Delta f_P \cdot f_{e,i_P} / r_{i_P}$ is the amount of flow moved from P to paths containing e . We consider the sequence of paths P_j sampled by our algorithm in ascending order of $(\ell_P - L_{i_P}) / \ell_P$. In this sequence, a path may occur more than once. Let $f_e^l = f_e + \sum_{j=1}^l \Delta f_{P_j}^e$. To conclude the proof, we prove Eq. (4) by showing that for $l \geq 0$,

$$\int_{f_e}^{f_e^l} (\ell_e(u) - \ell_e) du \leq \sum_{i=1}^l V_{P_i}^e.$$

by induction on l . Then, the increase of flow of edge e caused by the first l paths is

$$\begin{aligned} f_e^l - f_e &\leq \sum_{j=1}^l \Delta f_{P_j} \cdot \frac{f_{e,i_{P_j}}}{r_{i_{P_j}}} \\ &\leq \sum_{j=1}^l f_{P_j} \cdot \frac{\ell_{P_j} - L_{i_{P_j}}}{4d \cdot \ell_{P_j}} \cdot \frac{f_{e,i_{P_j}}}{r_{i_{P_j}}} \\ &\leq \frac{1}{4d} \cdot \frac{\ell_{P_l} - L_{i_{P_l}}}{\ell_{P_l}} \cdot \sum_{i=1}^k f_{e,i} \\ &\leq \frac{1}{4d} \cdot \frac{\ell_{P_l} - L_{i_{P_l}}}{\ell_{P_l}} \cdot f_e. \end{aligned}$$

The third inequality holds since we can separate the sum into sums over paths from only one commodity. Then, the f_P in sum cancel out the r_i 's. Now, due to the bounded elasticity of ℓ_e , we can apply Fact 1 to bound the total increase of latency caused by this increase of flow by

$$\Delta \ell_e^l = \ell_e(f_e^l) - \ell_e \leq 2d \cdot \ell_e \cdot \frac{f_e^l - f_e}{f_e} \leq \ell_e \cdot \frac{\ell_{P_l} - L_{i_{P_l}}}{2 \ell_{P_l}}.$$

Using the definition of $\Delta f_{P_l}^e$,

$$\Delta f_{P_l}^e \cdot \Delta \ell_e^l \leq \frac{\ell_e f_{e,i_{P_l}}}{2 \ell_{P_l} r_{i_{P_l}}} \cdot (\ell_{P_l} - L_{i_{P_l}}) \cdot \Delta f_{P_l} \leq V_{P_l}^e,$$

Using the induction hypothesis and the preceding inequality,

$$\int_{f_e}^{f_e^l} (\ell_e(u) - \ell_e) du \leq \int_{f_e}^{f_e^{l-1}} (\ell_e(u) - \ell_e) du + \Delta f_{P_i}^e \cdot \Delta \ell_e^l$$

$$\leq \sum_{i=1}^l V_{P_i}^e.$$

our claim. □

Lemma 5 *Assume that f is a flow that is not at δ - ϵ -equilibrium and let the random variable f' denote a flow generated by our algorithm. Then, $\mathbb{E}[\Phi(f')] \leq \Phi(f) \cdot (1 - \Omega(\frac{\epsilon^3 \delta^2}{d}))$.*

Proof For the time being, assume that the latency functions are constant. By Markov’s inequality, the total volume of flow in commodities with $L_i > 2 \cdot L/\epsilon$ is at most $\epsilon/2$. We consider only commodities with $L_i \leq 2 \cdot L/\epsilon$. In total, at least a flow volume of ϵ utilises δ -expensive paths and there is still at least a volume of $\epsilon/2$ left in the commodities we consider. Consider such a commodity $i \in [k]$ and denote the flow volume using δ -expensive paths in this commodity by ϵ_i .

Consider any iteration satisfying the precondition that all edges are alive. Let P denote the path sampled by the algorithm. Consider the event that $\ell_P \geq L_i + \delta L$ and the minimum edge flow along P is at least $\epsilon_i/(2m)$. By Lemma 2, the probability of this event is at least $\epsilon_i/(2r_i)$ (we have to scale the flow of this commodity by a factor $1/r_i$ to make it a unit flow). The amount of flow removed from this path by our algorithm is

$$\frac{\epsilon_i}{2m} \cdot \frac{1}{7 \log m} \cdot \frac{\ell_P - L_i}{4d \ell_P} \geq \frac{\epsilon_i \epsilon \delta}{113 d m \log m}$$

where we have used that $\ell_P \geq L_i + \delta L$ and $L_i \leq 2L/\epsilon$. The latency gain of this path is then at least δL and since this event happens with probability $\epsilon_i/(2r_i)$ the expected virtual potential gain of such a path is then at least

$$\frac{\epsilon_i^2 \epsilon \delta^2}{226 d r_i m \log m} L.$$

By Lemma 3 the probability that in this iteration all edges are alive is $1 - o(1)$ and the expected potential gain computed above is independent of this event. Summing up over all $T = m \log m$ iterations and all commodities, the total expected virtual potential gain of one round is at least

$$(1 - o(1)) \cdot \sum_{i \in [k]} \frac{\epsilon_i^2 \epsilon \delta^2}{226 d r_i} L \geq (1 - o(1)) \cdot \frac{\epsilon^3 \delta^2}{226 d} L.$$

For the last inequality we have used the Cauchy–Schwarz inequality which asserts that for two vectors (a_i) and (b_i) , $\sum_i a_i^2 \geq (\sum_i a_i b_i)^2 / \sum_i b_i^2$. Using $a_i = \epsilon_i/\sqrt{r_i}$ and $b_i =$

$\sqrt{r_i}$ yields the result. This implies the claim since L is an upper bound on Φ and Lemma 4 ensures that the true potential gain with respect to the real latency functions is at least half of the potential gain with respect to the constant latency functions. □

4.3 From expected potential gain to expected stopping time

The preceding section has shown that in every round the potential decreases by a factor in expectation. Intuitively, this implies an expected running time that is logarithmic in this factor and the initial values. This intuition is made precise by the following two lemmas. Although it seems likely that similar lemmas have been proven elsewhere before, the authors are not aware of any formulation that can be used here.

Lemma 6 *Let X_0, X_1, \dots denote a sequence of non-negative random variables and assume that for all $i \geq 0$*

$$\mathbb{E}[X_i | X_{i-1} = x_{i-1}] \leq x_{i-1} - 1$$

and let τ denote the first time t such that $X_t = 0$. Then,

$$\mathbb{E}[\tau | X_0 = x_0] \leq x_0.$$

Proof The proof is by induction on x_0 . Let

$$T(r) = \mathbb{E}[\tau | X_0 = r].$$

Clearly, $T(0) = 0$. For $i \in [x_0]$ let $p(i)$ denote the probability that $x_0 - X_1 = i$. By our assumption and definition of $p(i)$,

$$1 \leq x_0 - \mathbb{E}[X_1 | X_0 = x_0] = \sum_{i=0}^{x_0} p(i) \cdot i.$$

By definition of $T(j)$,

$$T(j) = 1 + \sum_{i=0}^j p(i) \cdot T(j - i)$$

$$\leq 1 + p(0) \cdot T(j) + \sum_{i=1}^j p(i) \cdot (j - i)$$

$$= 1 + p(0) \cdot T(j) + j \cdot (1 - p(0)) - \sum_{i=1}^j p(i) \cdot i$$

$$\leq p(0) \cdot T(j) + j \cdot (1 - p(0))$$

where the first inequality uses the induction hypothesis for $1 \leq i < j$.

Hence, $T(j) \leq j$, implying our claim. □

Lemma 7 *Let X_0, X_1, \dots denote a sequence of non-negative random variables and assume that for all $i \geq 0$*

$$\mathbb{E}[X_i | X_{i-1} = x_{i-1}] \leq x_{i-1} \cdot \alpha$$

for some constant $\alpha \in (0, 1)$. Furthermore, fix some constant $x^* \in (0, x_0]$ and let τ be the random variable that describes the smallest t such that $X_t \leq x^*$. Then,

$$\mathbb{E}[\tau \mid X_0 = x_0] \leq \frac{2}{\log(1/\alpha)} \cdot \log\left(\frac{x_0}{x^*}\right).$$

Proof In order to apply Lemma 6 we transform our random variable into a new sequence of random variables

$$Y_i = 2 \cdot \frac{\log(X_i) - \log(x^*)}{\log(1/\alpha)}.$$

Let $x_i = 2^{y_i \cdot \log(1/\alpha) / 2 + \log x^*}$. Then,

$$\begin{aligned} \mathbb{E}[Y_i \mid Y_{i-1} = y_{i-1}] &= \mathbb{E}\left[2 \cdot \frac{\log(X_i) - \log x^*}{\log(1/\alpha)} \mid X_{i-1} = x_{i-1}\right] \\ &= 2 \cdot \frac{\mathbb{E}[\log(X_i) \mid X_{i-1} = x_{i-1}] - \log(x^*)}{\log(1/\alpha)} \\ &\leq 2 \cdot \frac{\log(\mathbb{E}[X_i \mid X_{i-1} = x_{i-1}]) - \log(x^*)}{\log(1/\alpha)} \\ &\leq 2 \cdot \frac{\log(x_{i-1} \cdot \alpha) - \log(x^*)}{\log(1/\alpha)} \\ &= y_{i-1} - 2 \end{aligned}$$

where the first inequality is Jensen’s inequality and the second is our assumption on the sequence X_i . Observe that $X_i = x^*$ if and only if $Y_i = 0$ and $\lfloor Y_i + 1 \rfloor = 0$ implies that $X_i \leq x^*$. Now,

$$\begin{aligned} \mathbb{E}[\lfloor Y_i \rfloor \mid \lfloor Y_{i-1} \rfloor = y_{i-1}] &\leq \mathbb{E}[Y_i \mid \lfloor Y_{i-1} \rfloor = y_{i-1}] \\ &\leq \max_{z: \lfloor y_{i-1} \rfloor = z} \mathbb{E}[Y_i \mid Y_{i-1} = z] \\ &\leq \max_{z: \lfloor y_{i-1} \rfloor = z} z - 2 \\ &\leq y_{i-1} - 1 \end{aligned}$$

implying that the sequence $\lfloor Y_i \rfloor$ satisfies the conditions of Lemma 6. Let $\tilde{\tau}$ denote the smallest i such that $\lfloor Y_i + 1 \rfloor = 0$ and observe that $X_0 = x_0$ implies $\lfloor Y_0 + 1 \rfloor = \lfloor 2 \log(x_0/x^*) / \log(1/\alpha) + 1 \rfloor$ and hence,

$$\begin{aligned} \mathbb{E}[\tau \mid X_0 = x_0] &\leq \mathbb{E}[\tilde{\tau} \mid X_0 = x_0] \\ &\leq \mathbb{E}\left[\tilde{\tau} \mid \lfloor Y_0 + 1 \rfloor = \left\lfloor 2 \cdot \frac{\log(x_0/x^*)}{\log(1/\alpha)} + 1 \right\rfloor\right] \\ &\leq \left\lfloor 2 \cdot \frac{\log(x_0/x^*)}{\log(1/\alpha)} + 1 \right\rfloor, \end{aligned}$$

our desired bound. □

4.4 Convergence time

Finally, we can prove our main result.

Proof of Theorem 2 Again, convergence follows from Lemma 4 as in [13]. To obtain a bound on the convergence time, let f_0, f_1, \dots denote a sequence of flow vectors generated by Algorithm 2. Lemma 5 implies that

$$\mathbb{E}[\Phi(f_{i+1}) \mid \Phi(f_i) = \phi] \leq \phi \cdot \left(1 - \Omega\left(\frac{\epsilon^3 \delta^2}{d}\right)\right).$$

Thus, the sequence $(\Phi(f_i))_{i \geq 0}$ satisfies the conditions of Lemma 7 and the expected time until $\Phi(f_i)$ reaches its minimum Φ^* implying that f_i is a δ - ϵ -equilibrium is

$$\begin{aligned} &\frac{2}{\log\left(\left(1 - \Omega\left(\frac{\epsilon^3 \delta^2}{d}\right)\right)^{-1}\right)} \log\left(\frac{\Phi(f_0)}{\Phi^*}\right) \\ &= \mathcal{O}\left(\frac{d}{\epsilon^3 \delta^2} \log\left(\frac{\Phi(f_0)}{\Phi^*}\right)\right), \end{aligned}$$

our desired bound.

One path can be sampled in time $\mathcal{O}(n \log n)$, the bottleneck edge can be found in time $\mathcal{O}(n)$, and the flow update can be computed in time $\mathcal{O}(n)$. Altogether, at most $T = m \log m$ iterations have to be computed. Finally, the removed flow can be reinserted in time $\mathcal{O}(m)$. □

5 Conclusions and open problems

5.1 Summary

We have introduced a simple distributed algorithm that computes approximate Wardrop equilibria. The algorithm has a very natural interpretation in terms of agents each of which controls the flow of a single commodity trying to balance its own latency. The algorithm is stateless and does not depend on the desired approximation parameter. In the long run, it converges to an equilibrium. With respect to a bicriterial definition of approximate equilibria, our algorithm behaves like a fully polynomial time approximation scheme (FPTAS). The number of communication rounds depends polynomially on the approximation parameters and is almost independent of the network size. The computation time of each round is polynomial.

5.2 Naive deterministic decomposition fails

Our algorithm works by redistributing flow of overloaded paths. To identify such paths we face the subproblem of finding a flow decomposition that assigns much flow to paths with high latency (induced by the current flow). In our algorithm we have used a randomised path decomposition to achieve this goal. It is a natural question whether this randomisation can be avoided. In a greedy approach we could use a path decomposition that chooses a path with the largest latency, assigns to it a flow equivalent to the bottleneck flow of this

path and removes it from the network. In fact, a simple example shows that this approach does not necessarily maximise the unbalancedness of the decomposition. Recently, however, it has been shown in [18] that the problem of finding an unbalanced decomposition can be reduced to a min-cost-flow problem. At the cost of an increased running time this could be used as a module in our algorithm derandomising it.

5.3 Stronger approximation concepts

In the long run, our algorithm converges towards the set of Wardrop equilibria. A weakness of our notion of approximate equilibria, however, is the fact that the average latency may be arbitrarily far away from the minimum latency. Furthermore, a δ - ϵ -equilibrium, allows some of the commodities to be very out of balance.

There are two alternative, stronger definitions of approximate equilibria. First, one could require all but an ϵ -fraction to deviate from the average of their commodity by at most δL_i rather than δL . Second, one could also consider deviations from the minimum latency rather than from the average latency. It is unclear whether convergence towards approximate equilibria in this sense can be guaranteed in polynomial time.

Finally, it would be desirable to design specialised (not necessarily distributed) algorithms to compute (exact) Wardrop equilibria that improve upon the standard solution via convex programming.

References

- Ahuja, R.K., Magnanti, T.L., Orlin, J.B.: *Network Flows: Theory, Algorithms and Applications*. Prentice-Hall, Englewood Cliffs (1993)
- Awerbuch, B., Khandekar, R.: Greedy distributed optimization of multi-commodity flows. In: *Proceedings of 26th Annual Symposium on Principles of Distributed Computing (PODC)* (2007)
- Awerbuch, B., Khandekar, R., Rao, S.: Distributed algorithms for multicommodity flow problems via approximate steepest descent framework. In: *Proceedings of 18th Annual Symposium on Discrete Algorithms (SODA)* (2007)
- Awerbuch, B., Kleinberg, R.D.: Adaptive routing with end-to-end feedback: distributed learning and geometric approaches. In: *Proceedings of 36th Annual Symposium on Theory of Computation (STOC)*, pp. 45–53 (2004)
- Beckmann, M., Bart McGuire, C., Winsten, C.B.: *Studies in the Economics of Transportation*. Yale University Press, New Haven (1956)
- Berenbrink, P., Friedetzky, T., Goldberg, L.A., Goldberg, P., Hu, Z., Martin, R.: Distributed selfish load balancing. In: *Proceedings of 17th Annual Symposium on Discrete Algorithms (SODA)* (2006)
- Bertsekas, D.P.: *Network Optimization: Continuous and Discrete Models*. Athena Scientific, Belmont (1998)
- Blum, A., Even-Dar, E., Ligett, K.: Routing without regret: on convergence to Nash equilibria of regret-minimizing algorithms in routing games. In: *Proceedings of 25th Annual Symposium on Principles of Distributed Computing (PODC)*, pp. 45–52. ACM, Denver, Colorado (2006)
- Cole, R., Dodis, Y., Roughgarden, T.: How much can taxes help selfish routing? In: *Proceedings of 4th ACM Conference on Electronic Commerce*, pp. 98–107 (2003)
- Cominetti, R., Correa, J.R., Stier Moses, N.E.: Network games with atomic players. In: *Proceedings of 33rd International EATCS Coll. on Automata, Languages and Programming (ICALP)*, pp. 525–536 (2006)
- Even-Dar, E., Mansour, Y.: Fast convergence of selfish rerouting. In: *Proceedings of 16th Annual Symposium on Discrete Algorithms (SODA)*, pp. 772–781 (2005)
- Fischer, S.: *Dynamic selfish routing*. PhD Thesis, Lehrstuhl für Algorithmen und Komplexität, RWTH Aachen (2007)
- Fischer, S., Räcke, H., Vöcking, B.: Fast convergence to Wardrop equilibria by adaptive sampling methods. In: *Proceedings of 38th Annual ACM Symposium on Theory of Computing (STOC)*, pp. 653–662. ACM, Seattle (2006)
- Fischer, S., Vöcking, B.: Adaptive routing with stale information. In: *Proceedings of 24th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pp. 276–283. ACM, Las Vegas (2005)
- Fleischer, L.: Linear tolls suffice: new bounds and algorithms for tolls in single source networks. In: *Proceedings of 31st International EATCS Coll. on Automata, Languages and Programming (ICALP)*, pp. 544–554 (2004)
- Goldberg, P.W.: Bounds for the convergence rate of randomized local search in a multiplayer, load-balancing game. In: *Proceedings of 23rd Annual Symposium on Principles of Distributed Computing (PODC)*, pp. 131–140. ACM (2004)
- Hagerup, T., Rüb, C.: A guided tour of Chernoff bounds. *Inf. Process. Lett.* **33**, 305–308 (1990)
- Ochel, M.: *Effiziente Algorithmen für Gleichgewichtsflüsse*. Master's Thesis, RWTH Aachen (2007)
- Roughgarden, T.: How unfair is optimal routing? In: *Proceedings of 13th Annual Symposium on Discrete Algorithms (SODA)*, pp. 203–204 (2002)
- Roughgarden, T., Tardos, É.: How bad is selfish routing? *J ACM* **49**(2), 236–259 (2002)
- Wardrop, J.G.: Some theoretical aspects of road traffic research. In: *Proceedings of the Institute of Civil Engineers, Pt. II, vol. 1*, pp. 325–378 (1952)