

Nash Equilibria and the Price of Anarchy for Flows Over Time *

Ronald Koch Martin Skutella

TU Berlin, Institut für Mathematik, MA 5-2, Str. des 17. Juni 136,
10623 Berlin, Germany, {koch,skutella}@math.tu-berlin.de

Abstract

We study Nash equilibria and the price of anarchy in the context of flows over time. Many results on *static* routing games have been obtained over the last ten years. In flows over time (also called *dynamic* flows), flow travels through a network over time and, as a consequence, flow values on edges change over time. This more realistic setting has not been tackled from the viewpoint of algorithmic game theory yet; on the other hand, there is a rich literature on game theoretic aspects of flows over time in the traffic community.

In this paper, we present the first known results on the price of anarchy for flows over time. We also present algorithms for computing Nash flows over time. Those algorithms have to iteratively solve certain interesting and new static flow problems. Our results are based on a novel characterization of Nash equilibria for flows over time. The underlying flow over time model is a variant of the so-called *deterministic queueing model* that is very popular in road traffic simulation and related fields.

1 Introduction

In a groundbreaking paper, Roughgarden and Tardos [45] (see also Roughgarden’s book [44]) analyzed the price of anarchy for selfish routing games in networks. Such routing games are based on a classical static flow problem with convex latency functions on the arcs of the network. In a Nash equilibrium, flow particles (infinitesimally small flow units) selfishly choose an origin-destination path of minimum latency.

One main drawback of this class of routing games is its restriction to *static* flows. Flow variation over time is, however, an important feature in network flow problems arising in various applications. As examples we mention road or air traffic control, production systems, communication networks (e.g., the Internet), and financial flows. In contrast to static flow models, flow values on edges may change with time in these applications. Moreover, flow does not progress instantaneously but can only travel at a certain pace through the network which is determined by transit times of edges. Both temporal features are captured by *flows over time* (sometimes also called *dynamic* flows) which were introduced by Ford and Fulkerson [17, 18].

Another crucial phenomenon in many of those applications mentioned above is the variation of time taken to traverse an arc with the current (and maybe also past) flow situation on this arc. The latter aspect induces highly complex dependencies and leads to non-trivial mathematical flow models. For a more detailed account and further references we refer to [4, 12, 21, 34, 40, 41]. In particular, all of these *flow over time* models (also called *dynamic flow* models in the literature), have so far resisted a rigorous algorithmic analysis of Nash equilibria and the price of anarchy.

In this paper we identify a suitable flow over time model that is based on the following simplifying assumptions. Every edge of a given network has a fixed free flow transit time and capacity. The capacity of an edge bounds the rate (flow per time unit) at which flow may traverse the edge. The free flow transit time denotes the time that a flow particle needs to travel from the tail to the head of the edge. If, at some point in time, more flow wants to enter an edge than its capacity allows, the flow particles queue up at the tail of the arc and wait in line before they actually enter the edge. When a new flow particle wants to traverse an edge, the time needed to arrive at the head thus consists of the waiting time plus the fixed free flow transit time. In the traffic literature, this flow over time model is known as a “deterministic queueing model”. Similar models are used, for example, in road traffic simulation and related fields.

*This work is supported by DFG Research Center MATHEON “Mathematics for key technologies” in Berlin.

Related Literature. As already mentioned above, flows over time with fixed transit times were introduced by Ford and Fulkerson [17, 18]. For more details and further references on these classical flows over time we refer, for example, to [15, 16, 24, 25, 28, 29].

So far, Nash equilibria for flows over time were mostly studied within the traffic community. Vickrey [50] and Yagar [51] were the first to introduce this topic. Up to the middle of the 1980's, nearly all contributions consider Nash equilibria on given small instances. For example, Vickrey [50], Hendrickson and Kocur [23], and de Palma, Ben-Akiva, Lefevre, and Litinas [14] consider a network consisting of a single edge. Networks consisting of two parallel edges are considered by Mahmassani and Hermann [32].

Starting around the middle of the 1980's, the number of publications in this area increased rapidly and Nash equilibria were modeled mathematically. For a survey see, e.g., [39]. These models can be grouped into four categories: mathematical programming, optimal control, variational inequalities, and simulation-based approaches.

Mathematical programming: The first to come up with an mathematical program for modeling Nash equilibria in the context of flows over time was Janson [26]. His model is an extension of the well known static formulation of Beckmann, McGuire, and Winston [7] and is also used by others; see, e.g., [27]. However Lin and Lo [31] give a counterexample showing that this formulation does not lead to a Nash equilibrium in general. Recently, Han and Heydecker [22] found a mathematical programming formulation modeling Nash equilibria for flows over time.

Optimal control: Such models are very rarely used in the literature. For example Ran, Boyce, and LeBlanc [42] use this kind of model to formulate a Nash equilibrium for flows over time. They consider the problem where each user decides to take a shortest path to its destination based on the current transit times instead of experienced transit times. Moreover every user can switch to another path if its transit time decreases over time. Optimal control formulations are mostly used in order to model system optimal flows over time; see, e.g., [20].

Variational inequalities: Friesz, Bernstein, Smith, Tobin, and Wie [19] were the first to model Nash equilibria in the context of flows over time using variational inequalities. Their approach leads to a very general model. Up to now, variational inequalities are the most common formulation for analyzing Nash equilibria in the context of flows over time, see e.g. [13, 43, 47, 48].

Simulation-based approaches: Yagar [51] for the first times used a computer in order to come up with a flow over time which is approximately a Nash equilibrium. Since then, simulation-based approaches are widely used due to their relative simplicity. On the positive side, they are very flexible in modeling the behavior of individual users/agents in a network. But as a drawback, such models can hardly be analyzed rigorously. Simulation tools are for example DYNASMART [33, 35], DynaMIT [8], CONTRAM [30, 49], and MATSim [5, 6].

Many models mentioned above use a path-based formulation of flows over time. Therefore they are computationally often intractable. Edge-based formulations are, for example, considered in [2, 13, 43]. Realistic assumptions on the underlying flow model with respect to traffic simulation are described by Carey [10, 11].

In this paper a variant of the deterministic queueing model is considered. This model was introduced by Vickrey [50]. His analysis is based on instances consisting of a single bottleneck. A similar approach is used by Hendrickson and Kocur [23]. Smith [46] shows the existence of a departure time equilibrium for this model on instances where each route contains the same single bottleneck, i.e., an edge where a waiting queue is built. He introduces user-dependent desired arrival times and a payoff function for early and late arrival. In this setting, an equilibrium arises out of the choice of the departure times of the network users. Smith gives sufficient conditions for the existence of such an equilibrium.

Akamatsu [1] presents an edge-based variant of the deterministic queueing model on single-source-multi-sink-instances. In [2] he constructs an algorithm for this edge based model. This algorithm is a globally convergent newton method for general nonlinear complementarity problems. Akamatsu and Heydecker [3] study Braess's paradox for single-source-multi-sink-instances. Braess's paradox [9] states (for static flows) that increasing the capacity of one edge can increase the total cost of all users in a Nash flow. It is well known that this paradox is extendable to the dynamic case. Akamatsu and Heydecker [3] give necessary and sufficient conditions for the paradox to occur in the case of saturated networks. In saturated networks there are positive inflow rates on all edges over the entire time period as well as nonzero queues on all edges over the entire time period.

Mounce [36] considers the case where the edge capacities can vary over time. But it is assumed

that the capacity functions are increasing. In the case where there is only one edge per route having a nonzero waiting queue, a Nash equilibrium always exists. In [37] this approach is extended to general but continuous capacity functions.

Our Contribution. In this paper, we characterize and analyze Nash equilibria for a variant of the deterministic queueing model in networks with a single source and sink. We provide algorithms for computing Nash equilibria and obtain results on the price of anarchy for several special cases. Although algorithmic game theory is a flourishing area of research (see, e.g., the recent book [38]), network flows over time have not been studied from this perspective in the algorithms community so far. One purpose of this paper is to stimulate further interesting research in this relevant direction.

A precise description of our flow over time model is given in Section 2.1. We assume that flow particles arrive one after another at the source. Then each flow particle selfishly chooses a route through the network so as to minimize its arrival time at the sink. The resulting model of Nash equilibria along with several equivalent characterizations is discussed in Section 2.2. In Section 3.1 we prove that Nash equilibria always exist. One of our main results is presented in Section 3.2: we show that a Nash equilibrium can be characterized via a sequence of static flows with special properties. We believe that the resulting static flow problem is of interest in its own right. All our results mentioned so far hold for networks with edge capacities that can vary over time. The results discussed so far not only hold for the deterministic queueing model but for a considerably more general class of flow over time models.

In Section 4 we restrict to the case of constant capacities. We present an algorithm for computing a Nash equilibrium that is based on a sequence of static flow computations. For the case of zero free flow transit times or, more generally, for shortest paths networks where all source-sink paths have the same free flow transit time, we prove that a Nash equilibrium can be computed in polynomial time. The algorithm uses a series of sparse cut computations. As our main result we show that the price of anarchy is 1 in this case, that is, a Nash equilibrium is also a system optimum. For networks with arbitrary free flow transit times, however, this nice property no longer holds. The gap between the amounts of flow that have arrived in the sink at a certain time T in a system optimum and a Nash equilibrium, respectively, can be as large as $\Omega(m)$. Another possible objective function is to measure the time needed to send a given amount of flow into the sink. We conjecture that the price of anarchy here is at most $O(1)$. Finally, we show that in a more general multi-commodity flow setting, the price of anarchy can be arbitrarily large.

Due to space limitations, proofs and further details can be found in the appendix.

2 Model

In this section we present a model for Nash equilibria in the context of flows over time with a single source and a single sink. First, in Section 2.1, we introduce an appropriate flow over time model which is closely related to the deterministic queueing model mentioned above. Then, in Section 2.2 we define and characterize Nash equilibria for this flow over time model.

Throughout the paper we often use the term *flow particle* in order to refer to an infinitesimally small flow unit which travels along a single path through the network. The term *flow rate* refers to an amount of flow per time unit.

2.1 Flow Over Time Model

We consider a network (G, u, τ, s, t) consisting of a directed graph $G = (V, E)$, time-dependent Lebesgue integrable capacity functions $u_e : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, $e \in E$, constant *free flow transit times* $\tau_e \in \mathbb{R}_+$, $e \in E$, a source node $s \in V$, and a sink node $t \in V$. The basic concept of our flow over time model are waiting queues which are built at the tail of an edge if, at some point in time, more flow particles want to traverse an edge than the capacity of the edge allows. The free flow transit time of an edge determines the time for traversing an edge after passing the waiting queue. Thus the (*flow-dependent*) *transit time* on an edge is the sum of the current waiting time and the free flow transit time.

Further every flow particle arriving at a node $v \in V$ decides immediately which of the outgoing edges of v it traverses next. Thus one can think of the nodes as rooms with some entries and some exits, that

are large enough for containing all waiting queues of the exits and that are small enough for coming from any entry to any exit in time 0.

In the following we give a precise mathematical description of the model. A flow over time is defined by two families of flow rate functions which have to be Lebesgue integrable: for an edge $e \in E$ we have an inflow rate $f_e^+ : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ meaning that flow arrives at flow rate $f_e^+(\theta)$ at the tail (or at the end of the waiting queue) of e at time $\theta \geq 0$; moreover, the outflow rate $f_e^- : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ describes the rate of flow $f_e^-(\theta)$ arriving at the head of e at time $\theta \geq 0$.

Further we define for an edge e the cumulative in- and outflow at time $\theta \geq 0$ by $F_e^+(\theta) := \int_0^\theta f_e^+(\vartheta) d\vartheta$ and $F_e^-(\theta) := \int_0^\theta f_e^-(\vartheta) d\vartheta$, respectively. Thus the amount of flow which has been assigned to e and which has arrived at the head of e before time θ is $F_e^+(\theta)$ and $F_e^-(\theta)$, respectively. Note that the antiderivatives of nonnegative Lebesgue integrable functions are continuous and monotonically increasing. In particular, F_e^+ and F_e^- are continuous and monotonically increasing for all $e \in E$.

In order to obtain a feasible flow over time (f^+, f^-) (the in- and the outflow rates) must satisfy several conditions. The capacity function of an edge bounds the outflow rate of that edge:

$$f_e^-(\theta) \leq u_e(\theta) \quad \text{for all } e \in E, \theta \in \mathbb{R}_+. \quad (1)$$

More precisely, since we are dealing with Lebesgue integrable functions, we only require that this condition holds almost everywhere for each edge $e \in E$.

Further we have to impose several kinds of flow conservation constraints. Firstly flow can only traverse an edge if it has been assigned to this edge before:

$$F_e^+(\theta) - F_e^-(\theta + \tau_e) \geq 0 \quad \text{for all } e \in E, \theta \in \mathbb{R}_+. \quad (2)$$

Secondly we want flow arriving at an intermediate node $v \in V \setminus \{s, t\}$ to be immediately assigned to an outgoing edge of v . Further, because s is the source of the network, no flow should trickle away at s . Therefore the flow which is assigned to the outgoing edges of s must be greater than the flow arriving at s . Similarly we have to ensure that flow cannot originate at t . Thus the following conditions have to hold (almost everywhere):

$$\sum_{e \in \delta^-(v)} f_e^-(\theta) = \sum_{e \in \delta^+(v)} f_e^+(\theta) \quad \text{for all } v \in V \setminus \{s, t\}, \theta \in \mathbb{R}_+, \quad (3)$$

$$\sum_{e \in \delta^-(s)} f_e^-(\theta) \leq \sum_{e \in \delta^+(s)} f_e^+(\theta) \quad \text{for all } \theta \in \mathbb{R}_+, \quad (4)$$

$$\sum_{e \in \delta^-(t)} f_e^-(\theta) \geq \sum_{e \in \delta^+(t)} f_e^+(\theta) \quad \text{for all } \theta \in \mathbb{R}_+. \quad (5)$$

Note that, unlike the case of static flows, conditions (2)–(4) do not imply (5) since flow might vanish on an edge. The following condition ensures that flow which is assigned to an edge must leave this edge after a finite amount of time:

$$\text{there exists } \Delta < \infty \text{ such that } F_e^+(\theta) \leq F_e^-(\theta + \tau_e + \Delta) \text{ for all } e \in E, \theta \in \mathbb{R}_+. \quad (6)$$

The right hand side of the last inequality is the amount of flow arriving at the head of e before time $\theta + \Delta + \tau_e$. This is equal to the amount of flow which starts to traverse e before time $\theta + \Delta$. Thus, if Δ satisfies this inequality, the total flow which is assigned to e before time θ has started to traverse e before time $\theta + \Delta$. Therefore Δ is an upper bound on the amount of time which the flow particles assigned to e at time θ must stand in line before they start traversing e .

We consider a fixed edge e and a fixed point in time θ . It turns out that the infimum over all values Δ fulfilling the inequality in (6) is in fact a minimum. That is, equality is reached in (6) since F_e^- is a monotonically increasing continuous function. Therefore we define

$$q_e(\theta) := \min\{\Delta \geq 0 \mid F_e^+(\theta) = F_e^-(\theta + \tau_e + \Delta)\} \quad \text{for all } e \in E, \theta \in \mathbb{R}_+ \quad (7)$$

as the waiting time occurring at the tail of e experienced by flow particles assigned to e at time θ . Thus, the (flow-dependent) transit time of e at time θ is equal to $\tau_e + q_e(\theta)$.

Note that the interpretation of $q_e(\theta)$ as the waiting time for flow particles arriving at time θ is based on the assumption that the first-in-first-out (FIFO) condition is fulfilled on edge e . That is, no flow particle overtakes other flow particles in the waiting queue, even if there is free capacity on edge e . Showing formally that our model satisfies the FIFO condition is one purpose of the next lemma.

Lemma 2.1. *For any edge $e \in E$ we have:*

(i) *The function $\mathbb{R}_+ \rightarrow \mathbb{R}_+$, $\theta \mapsto \theta + q_e(\theta)$ is monotonically increasing.*

(ii) *The function q_e is continuous if and only if the following condition is satisfied for almost all $\theta \geq 0$:*

$$f_e^-(\theta + \tau_e + q_e(\theta)) = 0 \quad \Rightarrow \quad f_e^+(\theta) = 0.$$

Statement (ii) shows the following: If flow particles start to traverse e whenever there is a waiting queue at the tail of e or flow particles are assigned to e , then q_e is continuous. If flow particles selfishly travel through the network so as to minimize their arrival time at the sink (as it is the case in a Nash flow; see Section 2.2), this condition is always fulfilled. Of course, this is based on the assumption that the capacity functions u_e is strictly positive (almost everywhere). In the following we assume that all capacity functions u_e are strictly positive and that all edge waiting time functions are continuous.

2.2 Definition of Nash Equilibria for Flows over Time

In this section we define Nash equilibria in the context of flows over time. To simplify matters the approach is based on the deterministic queuing model presented above. But it can be generalized to many other realistic flow models. We use the terms *Nash flow over time* and *dynamic Nash equilibrium* interchangeably.

The main aspect of Nash equilibria in flow models is the selfish routing of flow particles. We assume that flow occurs at the source according to a given supply function $u_0 : \mathbb{R}_+ \rightarrow \mathbb{R}_+$. Here $u_0(\theta)$ is the rate of flow arriving at the source at time $\theta \geq 0$. As soon as a flow particle pops up at the source, it decides by itself how to travel to the sink t . That is, it chooses an s - t -path and travels along this path as early as possible.

We define two classes of flow over times. In the first class, every flow particle travels only along “currently shortest paths”. In the second class, every flow particle tries to overtake as many other flow particles as possible while not be overtaken by others. The latter condition turns out to be a non-overtaking condition. Moreover we show that the two classes of flows over time coincide.

We start by defining *currently shortest s-t-paths* in a given flow over time. To do so, we consider the problem of sending an additional flow particle at time $\theta \geq 0$ from the source to the sink as quickly as possible. Let $\ell_v(\theta)$ be the earliest point in time when this flow particle can arrive at node $v \in V$. Then, for each edge $e = vw \in E$, we have that $\ell_v(\theta) + \tau_e + q_e(\ell_v(\theta)) \geq \ell_w(\theta)$. On the other hand, for each node $w \in V \setminus \{s\}$, there exists an incoming edge $e = vw \in \delta^-(w)$ such that equality holds, i.e. $\ell_v(\theta) + \tau_e + q_e(\ell_v(\theta)) = \ell_w(\theta)$. The latter case suggests that the flow particle uses edge e in order to arrive at node w as early as possible. Moreover, we have $\ell_s(\theta) = \theta$. Therefore we define the *label functions* $\ell_v : \mathbb{R}_+ \rightarrow \mathbb{R}_+ \cup \{\infty\}$ as follows:

$$\ell_s(\theta) := \theta \tag{8}$$

$$\ell_w(\theta) := \min\{\ell_v(\theta) + \tau_e + q_e(\ell_v(\theta)) \mid e = vw \in \delta^-(w)\}. \tag{9}$$

Note that the label functions are computable simultaneously for each time θ by adapting the shortest path algorithm of Bellmann and Ford in the following manner. After setting the label function of each node $v \in V \setminus \{s\}$ to the constant function ∞ and $\ell_s(\theta) := \theta$ for all θ , we have to repeat the following instruction $|V| - 1$ times.

For all edges $e = vw \in E$, set $\ell_w(\theta) := \min\{\ell_w(\theta), \ell_v(\theta) + \tau_e + q_e(\ell_v(\theta))\}$ for all θ .

The correctness of the BELLMANN-FORD algorithm ensures the existence of the label functions ℓ_v , $v \in V$, according to the conditions (8) and (9). Moreover by induction it is not hard to see that the minimum in the instruction is always taken over continuous, monotonically increasing functions. Since the instruction is executed finitely many times we get the following lemma.

Lemma 2.2. *For each node $v \in V$, the function ℓ_v is monotonically increasing and continuous.*

In a Nash equilibrium, flow should always be sent over currently shortest s - t -paths only. We say that an edge $e \in E$ is contained in a shortest path at time $\theta \geq 0$ if and only if $\ell_w(\theta) = \ell_v(\theta) + \tau_e + q_e(\ell_v(\theta))$. Of course, if an edge $e = vw \in E$ does not lie on a shortest s - t -path at a certain time $\theta \geq 0$, then no flow should be assigned to that edge at time $\ell_v(\theta)$ in a Nash flow.

Definition 2.3. We say that *flow is only sent over currently shortest paths* if for every edge $e = vw \in E$ the following condition holds for almost all times $\theta \geq 0$:

$$\ell_w(\theta) < \ell_v(\theta) + \tau_e + q_e(\ell_v(\theta)) \implies f_e^+(\ell_v(\theta)) = 0 .$$

Notice that, in general, the label functions do not have to be *strictly* monotonically increasing. In particular, the label function of t might be constant over a certain time interval $[\theta_1, \theta_2]$. Thus, a flow particle originating at s at time θ_1 might arrive at the sink t at the earliest possible time without necessarily being as early as possible at all intermediate nodes of its path. Definition 2.3 enforces, however, that all subpaths of the s - t -path chosen by a flow particle have to be as short as possible.

Definition 2.3 enforces that flow is sent only over shortest paths by a pointwise condition. The next lemma gives an equivalent definition which is more global and takes a larger time horizon into consideration.

Theorem 2.4. *For a flow over time the following statements are equivalent:*

- (i) *Flow is only sent over currently shortest paths.*
- (ii) *For each edge $e = vw \in E$ and for all $\theta \geq 0$ we have*

$$0 = F_e^-(\ell_v(\theta) + \tau_e + q_e(\ell_v(\theta))) - F_e^-(\ell_w(\theta)) .$$

We next show that the condition of sending flow only over currently shortest paths is equivalent to the condition that every flow particle wants to overtake as much other flow as possible while not being overtaken by other flow. The latter condition is in fact a “non-overtaking condition”, i.e., it is equivalent to the statement that no flow particle can possibly overtake any other flow particle. This is seen as follows. Assume that a flow particle p_2 originating at the source at time θ_2 overtakes an earlier flow particle p_1 originating at the source at time $\theta_1 < \theta_2$. That is, p_2 arrives at the sink before p_1 . Since the function $\theta \mapsto \theta + \tau_e + q_e(\theta)$ is monotonically increasing for each edge e , flow particle p_1 can arrive earlier at the sink and avoid being overtaken by p_2 by choosing the same path as p_2 . Thus flow particle p_1 can improve its situation.

In order to model the non-overtaking condition more formally, we consider again a flow over time and an additional flow particle originating at s at time $\theta \geq 0$. Of course, in order to ensure that no flow particle has the possibility to overtake this particle, it is necessary to take a shortest s - t -path. Therefore we define for each edge $e = vw \in E$ the amount of flow $x_e^+(\theta)$ assigned to e before this particle can reach v and the amount of flow $x_e^-(\theta)$ leaving e before this particle can reach w as follows:

$$x_e^+(\theta) := F_e^+(\ell_v(\theta)), \quad x_e^-(\theta) := F_e^-(\ell_w(\theta)) \quad \text{for all } \theta \geq 0. \quad (10)$$

Thus the amount of flow $b_s(\theta)$ that has originated at s before our flow particle occurs at s and the amount of flow $-b_t(\theta)$ arriving at t before our flow particle can reach t is equal to:

$$b_s(\theta) := \sum_{e \in \delta^+(s)} x_e^+(\theta) - \sum_{e \in \delta^-(s)} x_e^-(\theta), \quad b_t(\theta) := \sum_{e \in \delta^+(t)} x_e^+(\theta) - \sum_{e \in \delta^-(t)} x_e^-(\theta) \quad \text{for all } \theta \geq 0. \quad (11)$$

Note that $b_s(\theta)$ is always nonnegative and $b_t(\theta)$ is always non-positive because of flow over time conditions (4) and (5). Moreover $b_s(\theta) > -b_t(\theta)$ means that the considered flow particle overtakes other flow particles and $b_s(\theta) < -b_t(\theta)$ means that this flow particle is overtaken by other flow particles. This motivates the following definition:

Definition 2.5. We say that *no flow overtakes any other flow* if at all times $\theta \geq 0$ we have:

$$b_s(\theta) = -b_t(\theta).$$

Now we are able to prove the equivalence of the non-overtaking condition and the condition that flow only uses currently shortest paths. In addition, a third equivalent statement is given.

Theorem 2.6. *For a flow over time the following statements are equivalent:*

- (i) *Flow is only sent over currently shortest paths.*
- (ii) *For every edge $e \in E$ and all times $\theta \geq 0$ it holds that $x_e^+(\theta) = x_e^-(\theta)$.*
- (iii) *No flow overtakes any other flow.*

Before presenting the proof of the theorem in the Appendix A.3, we first discuss a intuitive real-world example.

Note that whenever one of the three statements in Theorem 2.6 holds, then x^+ and x^- coincide. Further, we conclude from the proof that $x_e(\theta) := x_e^+(\theta)$ for all $e \in E$ is in fact a static s - t -flow with value $b_s(\theta)$ for all $\theta \geq 0$. In the following, for a flow over time satisfying the non-overtaking condition, we refer to $(x_e(\theta))_{e \in E}$ as *the underlying static flow at time θ* . Also note that, for an arbitrary flow over time, the proof implies that $b_s(\theta) + b_t(\theta) \geq 0$ for all $\theta \geq 0$.

Next we define a Nash equilibrium in the context of flows over time. As already mentioned we assume that the flow particles originate at the source s over time. This is modeled through a Lebesgue integrable function $u_0 : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that the rate at which flow originates at s at a certain time $\theta \geq 0$ is $u_0(\theta)$. Intuitively, all flow initially waits in a queue on an artificial edge pointing into s and the function u_0 is the capacity function of that edge.

So far we did not really take the capacity functions of edges into consideration. But, of course, a Nash flow should use an edge up to its capacity, especially if a nonzero waiting queue occurs at its tail.

Definition 2.7. Let (G, u, s, t, τ, u_0) be a network where flow originates at the source s according to the function u_0 . A flow over time (f^+, f^-) is called *dynamic Nash equilibrium* or *Nash flow over time* if and only if the following conditions are satisfied:

- (i) $u_0 = \sum_{e \in \delta^+(s)} f_e^+ - \sum_{e \in \delta^-(s)} f_e^-$.
- (ii) Flow is only sent over currently shortest paths.
- (iii) For each $e \in E$ and for all $\theta \geq 0$, if $q_e(\theta) > 0$, then $f_e^-(\theta + \tau_e) = u_e(\theta + \tau_e)$.

Note that conditions (i) and (iii) of this definition must only hold almost everywhere because we are working with Lebesgue integrable functions.

As already mentioned at the beginning of this section, the presented approach is extendable to flow over time models satisfying certain weak requirements: For a feasible s - t -flow over time (f^+, f^-) , the in- and outflow rate functions must be nonnegative and must satisfy the flow conservation constraints (2)–(6). In addition, the FIFO-principle must hold for every edge. That is, if $\tau_e(\theta)$ is the current transit time of an edge $e \in E$ experienced by flow particles arriving at the tail of e at time θ , then $\theta \mapsto \theta + \tau_e(\theta)$ must be an increasing function. Moreover, if we assume that the transit time functions are continuous, then all of the results and definitions remain valid. Of course we must delete condition (iii) of Definition 2.7.

3 Analysis of Nash Flows Over Time

The first question one might ask is whether a Nash flow over time as in Definition 2.7 always exists. We give an affirmative answer to this question in Section 3.1. This existence result is generalizable to other flow models satisfying certain weak assumption. Moreover there exists a nice relationship between the underlying static flow and the node labels. This aspect is considered in Section 3.2.

3.1 Existence of Nash Flows Over Time

In this section we prove that a Nash flow over time always exists in the deterministic queuing model introduced in Section 2.1. We split the initial waiting queue in front of s into flow units of size $\epsilon > 0$ and route every flow unit along a currently shortest s - t -path. Thus we get a flow over time which, in general,

violates the non-overtaking condition slightly. Then we let ϵ tend to zero in order to get a sequence of flows over time that converges to a Nash flow over time.

Given a network (G, u, s, t, τ, u_0) and $\epsilon > 0$, the so called GENERAL ITERATIVE ALGORITHM works as follows: In a first step, the initial waiting queue is split into flow units of size ϵ , i.e., the algorithm determines $(\theta_i)_{i \in \mathbb{N}}$ such that $\int_{\theta_i}^{\theta_{i+1}} u_0(\vartheta) d\vartheta = \epsilon$. Then the following instructions are iterated over $i = 1, 2, \dots$

- (i) Compute a shortest s - t -path P at time θ_i according to the flow over time caused by the earlier flow units.
- (ii) Route the flow originating at s within the time interval $[\theta_i, \theta_{i+1}]$ along P such that the resulting flow over time is feasible and satisfies condition (i) and (iii) of Definition 2.7

The next lemma shows that the output of the this algorithm violates the non-overtaking condition only slightly.

Lemma 3.1. *Let (f^+, f^-) be the output of the GENERAL ITERATIVE ALGORITHM on a given network for some $\epsilon > 0$. Then,*

$$|b_s(\theta) + b_t(\theta)| \leq M\epsilon \quad \text{for all } \theta \geq 0$$

where M is the number of simple s - t -paths in G .

Lemma 3.1 shows that the non-overtaking condition is less violated for smaller values of ϵ . Thus, if we let ϵ tend to 0, we can hope to obtain a Nash flow over time. Unfortunately, for each edge e , we get a sequence of in- and outflow functions f_e^+ and f_e^- which are only Lebesgue integrable. Such a sequence has no limit point in general. We even cannot guarantee the existence of an accumulation point which would be enough for proving the existence of a Nash flow over time. Considering the cumulative flow functions (F^+, F^-) instead leads to:

Theorem 3.2. *A Nash flow over time always exists.*

We only remark that this existence result is also valid for flow models satisfying the weak conditions mentioned at the end of Section 2.2, even if we do not assume the continuity of the transit time functions.

3.2 A Special Class of Static Flows

In this section we present a property which will allow us to compute Nash flows over time for the special case of constant capacity functions by only solving a sequence of static flow problems (see Section 4). The property is based on the subnetwork containing all currently shortest s - t -paths.

Definition 3.3 (Current Shortest Paths Network). Let ℓ be the family of label functions and q be the family of waiting time functions of a flow over time on a network (G, u, s, t, τ, u_0) . The current shortest paths network G_θ for a point in time $\theta \geq 0$ is the subnetwork induced by the edges occurring in a currently shortest path, i.e., $V(G_\theta) := V$ and $E(G_\theta) := \{e = vw \in E \mid \ell_v(\theta) + \tau_e + q_e(\ell_v(\theta)) = \ell_w(\theta)\}$.

Note that, assuming the continuity of the label functions for all $\theta \geq 0$, there exists an $\epsilon > 0$ such that $G_{\theta'}$ is a subgraph of G_θ for all $\theta' \in [\theta, \theta + \epsilon)$. Furthermore, we need the following definition of a special static flow:

Definition 3.4 (Thin Flow with Resetting). Let (G, u, s, t, u_0) be a “static” network and $E_1 \subseteq E(G)$ be an edge set. A static flow x' with flow value F is called a *thin flow with resetting on E_1* if there exist node labels ℓ' such that:

$$\ell'_s = \frac{F}{u_0} \tag{12}$$

$$\ell'_w \leq \ell'_v \quad \text{for all } e = vw \in E(G) \setminus E_1, x'_e = 0 \tag{13}$$

$$\ell'_w = \max \left\{ \ell'_v, \frac{x'_e}{u_e} \right\} \quad \text{for all } e = vw \in E(G) \setminus E_1, x'_e > 0 \tag{14}$$

$$\ell'_w = \frac{x'_e}{u_e} \quad \text{for all } e = vw \in E_1 \tag{15}$$

For the special case where $E_1 = \emptyset$, the label ℓ'_v of node v is the congestion of all flow-carrying s - v -path and a lower bound on the congestion of any s - v -path. Here, the congestion of a path is the maximum congestion of its edges. The name “thin flow *with resetting*” refers to the special arcs in E_1 which play the following role. Whenever a path starting at s traverses an edge $e \in E_1$, it “forgets” the congestion of all arcs seen so far and “resets” its congestion to x'_e/u_e . As shown in Section 4, thin flows with resetting are computable in polynomial time for the special case $E = \emptyset$. We conjecture that this also holds for the general case. Without going into details we remark that the general case can be solved iteratively with a fix point approach, where an algorithm for the special case $E_1 = \emptyset$ is used as subroutine.

Next we show that for a Nash flow over time the derivatives of the label functions and of the underlying static flow behave like a thin flow with resetting. The following lemma is only applicable if the derivatives of the label and the underlying static flow functions exist. But both the label functions and the underlying static flow functions are monotonically increasing implying that both families of functions are differentiable almost everywhere. For the next lemma all derivatives are defined as derivatives from the right.

Theorem 3.5. *Consider a Nash flow over time on a network (G, u, s, t, τ, u_0) . Let ℓ be the family of label functions, q be the family of edge time waiting functions, and, for $\theta \geq 0$, let $x(\theta)$ be the underlying static flow. Let $\theta \geq 0$ such that $\frac{dx_e}{d\theta}(\theta)$ and $\frac{d\ell_v}{d\theta}(\theta)$ exist for all $e \in E$ and $v \in V$, respectively. Further assume that the capacity function of each edge is right continuous.*

Moreover, let G_θ be the current shortest path network at time θ . Then we have on G_θ with capacities $u_e(\ell_w(\theta))$, $e = vw \in E(G_\theta)$, and initial capacity $u_0(\theta)$ that $(\frac{dx_e}{d\theta}(\theta))_{e \in E(G_\theta)}$ is a thin flow of value $u_0(\theta)$ with resetting on the waiting edges $E_1 := \{e \in E \mid q_e(\theta) > 0\}$ according to the node labels $(\frac{d\ell_v}{d\theta}(\theta))_{v \in V(G_\theta)}$.

The reverse direction of the lemma also holds. Whenever the derivatives of the underlying static flow functions and the label functions of a flow over time are thin flows with resetting in the current shortest path network for almost all times θ , then the flow over time is in fact a Nash flow over time.

4 Networks with Constant Capacities

In this section we consider dynamic Nash equilibria on networks with constant capacities. Thus a network (G, u, τ, s, t, u_0) is given where the capacity on an edge is a real number $u_e \in \mathbb{R}_+$. Also the rate of flow u_0 originating at the source is constant, i.e., $u_0 \in \mathbb{R}_+$. We show that there exists a Nash flow over time such that all in- and outflow rate functions of edges are piecewise constant. We assume that each flow rate function is right continuous and is encoded as a sequence of supporting points where the flow rate function is discontinuous. The encoding size of the flow over time is the sum of the encoding sizes of the flow rate functions.

We construct an algorithm for computing a Nash flow over time. In this algorithm we iteratively have to find thin flows with resetting and we assume that we can solve these problems, i.e., we use an oracle for solving them. If this oracle is polynomial, the resulting algorithm is polynomial in the input plus output size. Further we consider instances where the network is a shortest path network according to the free flow transit times. In this case we are able to show that a Nash equilibrium sends until every point in time the maximal amount of flow to t .

The idea of the iterative algorithm is based on Theorem 3.5. Initially the flow over time is the constant zero flow. At the beginning of an iteration the flow originating at s up to a given time $\theta \geq 0$ has been routed in a Nash way. That is, it is a Nash flow over time for the instance where the inflow rate at s is zero after time θ . In each iteration we extend the current Nash flow over time for times greater than θ .

Let θ correspond to the beginning of an iteration. We compute a thin flow x' on the current shortest path network G_θ with resetting on the edges having a nonzero waiting queue. Let ℓ' be the corresponding node labels. The flow rates given by x' are used to extend the current flow over time; a detailed description of the algorithm is given in Appendix A.6.

If we increase θ , then the thin flow with resetting instance remains the same if no new edge is added to G_θ and no waiting time decreases to 0. Therefore the flow rate functions are constant in such time intervals. Let α be the largest period such that these requirements are satisfied. Then α must be feasible according to the following conditions. Note that ℓ' are the constant derivatives of the label functions ℓ within the interval $[\theta, \theta + \alpha)$. Further let E_1 be the edge set containing all waiting edges and $\text{dist}_\tau^G(v, w)$

denotes the length of a shortest v - w -path in G according to τ .

$$\ell_w + \alpha \ell'_w - \ell_v - \alpha \ell'_v \leq \text{dist}_\tau^{G \setminus E(G_\theta)}(v, w) \quad \text{for all } v, w \in V(G_\theta), \quad (16)$$

$$\ell_w + \alpha \ell'_w - \ell_v - \alpha \ell'_v \geq \tau_e \quad \text{for all } e = vw \in E_1. \quad (17)$$

Note that deleting edges $e \in E(G_\theta)$ with $x'_e = 0$ does not change the thin flow with resetting on the support of x' . Without going into details, we remark that α can be computed in polynomial time. At the end of each iteration we set the in- and outflow rate of an edge $e = vw$ to $\frac{x'_e}{\ell'_v}$ and $\frac{x'_e}{\ell'_w}$ in the time interval $[\ell_v, \theta + \alpha \ell'_v)$ and $[\ell_w, \theta + \alpha \ell'_w)$, respectively. This is feasible because of the definition of $x^+ = x^-$ in (10). Further if $\alpha < \infty$ we update the current label of node $v \in V$ to $\ell_v + \alpha \ell'_v$ and start the next iteration for $\theta + \alpha$. A detailed description of this NASH FLOW OVER TIME ALGORITHM is given in Appendix A.6.

Theorem 4.1. *Given a polynomial oracle for computing thin flows with resetting, we can compute a Nash flow over time on a network (G, u, τ, s, t, u_0) in time polynomial in the input plus output size.*

In the following we consider a network (G, u, τ, s, t, u_0) with constant capacities and the additional restriction that G is a shortest path network with respect to the free flow transit times τ . In the first iteration of the NASH FLOW OVER TIME ALGORITHM we have that $\alpha = \infty$ because $G_0 = G$ and $E_1 = \emptyset$ implying that (16) and (17) are satisfied for all $\alpha > 0$. Moreover, we only have to find a thin flow *without resetting*, meaning that we can delete condition (15) in Definition 3.4 in this special case. Thin flows without resetting are simply called thin flows and are defined as follows:

Definition 4.2. For a network (G, u, s, t) , a static s - t -flow $x' \in \mathbb{R}^{E(G)}$ is called *thin flow* if for each node v every flow carrying s - v -path has the same congestion ℓ'_v and every s - v -path has congestion at least ℓ'_v . If in addition an initial capacity u_0 is given we initialize $\ell'_s := \frac{F}{u_0}$ where F is the flow value of x .

It is not hard to see that thin flows on a network (G, u, s, t, u_0) are really thin flows with resetting on $E_1 := \emptyset$. In order to study thin flows, we can restrict to instances without an initial capacity, i.e., $\ell'_s = 0$. This is due to the fact that we can model the initial capacity simply by adding a node s_0 and an edge s_0s with capacity u_0 to the network. Then, of course, a thin flow on the new instance corresponds to a thin flow on the original instance and vice versa. Further, the definition of thin flows is directly generalizable to b -flow instances (G, u, b) where only one node s has a positive supply. Next we prove some properties of thin flows also motivating their name. We define an edge label ℓ'_e for each edge $e = vw \in E$ by $\ell'_e := \max\{\ell'_v, \frac{x'_e}{u_e}\}$.

Lemma 4.3. *Let x' be a thin b -flow on a network (G, u, b) where only one node s has a positive supply. Then the maximum label ℓ'_{\max} of any edge is equal to the congestion q^* of a sparsest cut in (G, u, b) , i.e. a node set $X \subset V$ maximizing $\frac{b(X)}{u(\delta^+(X))}$.*

The last lemma shows that a thin s - t -flow of value equal to the maximum flow value is also feasible according to the edge capacities. Moreover, thin flows are unique in some sense and computable in polynomial time:

Lemma 4.4. *Consider a pair of thin flows with the same flow value or the same node balances. Then the node labels are identical. Moreover, a thin flow of given flow value can be computed in polynomial time.*

As already mentioned, the underlying static flow of a Nash flow over time is a thin flow if the given network is a shortest path network with respect to the free flow transit times. As shown above, such a thin flow has maximal value, is unique in some sense and is computable in polynomial time. Thus we get:

Theorem 4.5. *Let (G, u, τ, s, t, u_0) be a network with constant edge capacities such that G is a shortest paths network with respect to free flow transit times τ . Then a Nash flow over time can be computed in polynomial time. Moreover, each Nash flow over time is a universally maximal flow over time, that is, for all $\theta \geq 0$, it sends the maximum possible amount of flow into the sink before time θ .*

Theorem 4.5 shows that a Nash flow over time is a universally maximal flow over time if the underlying network is a shortest path network according to the free flow transit time. As already mentioned a universally maximal flow over time sends the maximum possible amount of flow into the sink until a certain time and this is done for all times simultaneously. We only remark here that a universally maximal flow over time always exists in the case of constant capacities.

In algorithmic game theory it is very common to compare a Nash equilibrium with a system optimum according to some objective. A system optimum can only be obtained by a central authority who routes the flow optimally according to some objective. In the following we consider two different objectives. The first is to route flow such that at every point in time the amount of flow having reached the sink is maximal. This is, for example, the primary goal in an evacuation situation. The second, closely related objective is to route the flow such that a given amount of flow reaches the sink as early as possible. It is not hard to see that a universally maximal flow over time optimizes both objectives.

The ratio between the objective value of a Nash equilibrium and a system optimum is often called *price of anarchy*. Two kinds of prices of anarchy arise out of the two objectives mentioned above: According to the first objective, compute for some given time T the quotient between the amount of flow having reached the sink in a universally maximal flow over time and the amount of flow having reached the sink in a Nash flow over time. Then the *evacuation price of anarchy* is defined to be the largest quotient over all times T . According to the second objective, compute for a given amount of flow F the quotient between the time when this amount of flow has reached the sink in a universally maximal flow over time and the time when this amount of flow has reached the sink in a Nash flow over time. Then the *time price of anarchy* is defined to be the smallest quotient over all amounts F . Thus Theorem 4.5 can be restated as follows.

Corollary 4.6. *Let (G, u, τ, s, t, u_0) be a network with constant edge capacities such that G is a shortest paths network with respect to free flow transit times τ . Then the price of anarchy (w.r.t. both versions) is equal to 1.*

If we drop the assumption that the underlying network is a shortest path network with respect to the free flow transit times, we are able to construct instances with an arbitrarily large evacuation price of anarchy. More precisely, the price of anarchy can get as large as $\Omega(m)$ where m is the number of edges in the graph (see Appendix A.9). But surprisingly it can be shown that the time price of anarchy on the instances considered in Appendix A.9 is lower bounded by $\frac{e-1}{e}$. In fact we believe that the time price of anarchy is constantly bounded:

Conjecture 4.7. *Let (G, u, τ, s, t, u_0) be a network with constant edge capacities. Then the time price of anarchy is $O(1)$.*

Last but not least we mentioned that in a multicommodity setting the evacuation price of anarchy can be arbitrarily large; see Appendix A.10.

References

- [1] T. Akamatsu. A dynamic traffic equilibrium assignment paradox. *Transportation Research Part B: Methodological*, 34(6):515–531, 2000.
- [2] T. Akamatsu. An efficient algorithm for dynamic traffic equilibrium assignment with queues. *Transportation Science*, 35(4):389–404, 2001.
- [3] T. Akamatsu and B. Heydecker. Detecting dynamic traffic assignment capacity paradoxes in saturated networks. *Transportation Science*, 37(2):123 – 138, 2003.
- [4] J. E. Aronson. A survey of dynamic network flows. *Annals of Operations Research*, 20:1–66, 1989.
- [5] M. Balmer, K.W. Axhausen, and K. Nagel. An agent based demand modeling framework for large scale micro-simulations. Technical Report 329, Institute for Transport Planning and Systems (IVT), ETH Zurich, Switzerland, 2005.
- [6] M. Balmer, M. Rieser, K. Meister, D. Charypar, N. Lefebvre, K. Nagel, and K.W. Axhausen. Matsim-t: Architektur und rechenzeiten. submitted to Heureka08, 2008.

- [7] M. Beckmann, C.B. McGuire, and C.B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, New Haven, 1956.
- [8] M. J. Ben-Akiva, M. Bierlaire, H. N. Koutsopoulos, and R. Mishalani. Dynamit: A simulation-based system for traffic prediction and guidance generation. In *Proceedings of the 3rd Triennial Symposium on Transportation Systems*, 1998.
- [9] D. Braess. “über ein paradoxon aus der verkehrsplanung”. *Unternehmensforschung* 12, 12:258–268, 1968.
- [10] M. Carey. Link travel times i: Properties derived from traffic-flow models. *Networks and Spatial Economics*, 4(3):257–268, 2004.
- [11] M. Carey. Link travel times ii: Properties derived from traffic-flow models. *Networks and Spatial Economics*, 4(4):379–402, 2004.
- [12] H.-K. Chen. *Dynamic Travel Choice Models: A Variational Inequality Approach*. Springer, Berlin, 1999.
- [13] H. K. Chen and C. F. Hsueh. A model and an algorithm for the dynamic user-optimal route choice problem. *Transportation Research Part B: Methodological*, 32(3):219–234, 1998.
- [14] A. de Palma, M. Ben-Akiva, C. Lefvre, and N. Litinas. Stochastic equilibrium model of peak period traffic congestion. *Transportation Science*, 17(4):430–453, 1983.
- [15] L. Fleischer and M. Skutella. Quickest flows over time. *SIAM Journal on Computing*, 36:1600–1630, 2007.
- [16] L. K. Fleischer and É. Tardos. Efficient continuous-time dynamic network flow algorithms. *Operations Research Letters*, 23:71–80, 1998.
- [17] L. R. Ford and D. R. Fulkerson. Constructing maximal dynamic flows from static flows. *Operations Research*, 6:419–433, 1958.
- [18] L. R. Ford and D. R. Fulkerson. *Flows in Networks*. Princeton University Press, 1962.
- [19] T. L. Friesz, D. Bernstein, T. E. Smith, R. L. Tobin, and B. W. Wie. A variational inequality formulation of the dynamic network user equilibrium problem. *Operations Research*, 41(1):179–191, 1993.
- [20] T. L. Friesz, J. Luque, R. L. Tobin, and B. W. Wie. Dynamic network traffic assignment considered as a continuous time optimal control problem. *Operations Research*, 37(6):893–901, 1989.
- [21] H. W. Hamacher and S. A. Tjandra. Mathematical modelling of evacuation problems: A state of the art. In M. Schreckenberg and S. D. Sharma, editors, *Pedestrian and Evacuation Dynamics*, pages 227–266. Springer, Berlin, 2002.
- [22] S. Han and B.G. Heydecker. Consistent objectives and solution of dynamic user equilibrium models. *Transportation Research Part B: Methodological*, 40(1):16–34, 2006.
- [23] C. Hendrickson and G. Kocur. Schedule delay and departure time decisions in a deterministic model. *Transportation Science*, 15(1):62–77, 1981.
- [24] B. Hoppe. *Efficient dynamic network flow algorithms*. PhD thesis, Cornell University, 1995.
- [25] B. Hoppe and É. Tardos. The quickest transshipment problem. *Mathematics of Operations Research*, 25:36–62, 2000.
- [26] B. N. Janson. Dynamic traffic assignment for urban road networks. *Transportation Research Part B: Methodological*, 25(2–3):143–161, 1991.
- [27] R. Jayakrishnan, W. K. Tsai, and A. Chen. A dynamic traffic assignment model with traffic-flow relationships. *Transportation Research Part C: Emerging Technologies*, 3(1):51–72, 1995.

- [28] B. Klinz and G. J. Woeginger. Minimum-cost dynamic flows: The series-parallel case. *Networks*, 43:153–162, 2004.
- [29] B. Kotnyek. An annotated overview of dynamic network flows. Rapport de recherche 4936, INRIA Sophia Antipolis, 2003.
- [30] D. R. Leonard, P. Gower, and N. B. Taylor. Contram: structure of the model. Technical report, Transport and Road Research Laboratory, 1989. Report RR 178.
- [31] W. H. Lin and H. K. Lo. Are the objective and solutions of dynamic user-equilibrium models always consistent? *Transportation Research Part A: Policy and Practice*, 34(2):137–144, 2000.
- [32] H. Mahmassani and R. Herman. Dynamic user equilibrium departure time and route choice on idealized traffic arterials. *Transportation Science*, 18(4):362–384, 1984.
- [33] H. S. Mahmassani and R. Jayakrishnan. Dynamic traffic assignment and simulation for advanced network informatics (dynamart). In *Urban Traffic Networks: Dynamic Flow Modeling and Control (Transportation Analysis)*. Springer-Verlag, 1995.
- [34] H. S. Mahmassani and S. Peeta. System optimal dynamic assignment for electronic route guidance in a congested traffic network. In N. H. Gartner and G. Improta, editors, *Urban Traffic Networks. Dynamic Flow Modelling and Control*, pages 3–37. Springer, Berlin, 1995.
- [35] H. S. Mahmassani, H. A. Sbayti, and X. Zhou. *DYNASMART-P Version 1.0 Users Guide*. Maryland Transportation Initiative, College Park, Maryland, 2004.
- [36] R. Mounce. Convergence in a continuous dynamic queueing model for traffic networks. *Transportation Research Part B: Methodological*, 40(9):779–791, 2006.
- [37] R. Mounce. Convergence to equilibrium in dynamic traffic networks when route cost is decay monotone. *Transportation Science*, 41(3):409–414, 2007.
- [38] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [39] S. Peeta and A. K. Ziliaskopoulos. Foundations of dynamic traffic assignment: The past, the present and the future. *Networks and Spatial Economics*, 1(3–4):233–265, 2001.
- [40] W. B. Powell, P. Jaillet, and A. Odoni. Stochastic and dynamic networks and routing. In M. O. Ball, T. L. Magnanti, C. L. Monma, and G. L. Nemhauser, editors, *Network Routing*, volume 8 of *Handbooks in Operations Research and Management Science*, chapter 3, pages 141–295. North-Holland, Amsterdam, The Netherlands, 1995.
- [41] B. Ran and D. E. Boyce. *Modelling Dynamic Transportation Networks*. Springer, Berlin, 1996.
- [42] B. Ran, D. E. Boyce, and L. J. Leblanc. A new class of instantaneous dynamic user-optimal traffic assignment models. *Operations Research*, 41(1):192–202, 1993.
- [43] B. Ran, R. W. Hall, and D. E. Boyce. A link-based variational inequality model for dynamic departure time/route choice. *Transportation Research Part B: Methodological*, 30(1):31–46, 1996.
- [44] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.
- [45] T. Roughgarden and É. Tardos. How bad is selfish routing? In *Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science*, pages 93–102, Redondo Beach, CA, 2000.
- [46] M. J. Smith. The existence of a time-dependent equilibrium distribution of arrivals at a single bottleneck. *Transportation Science*, 18(4):385–394, 1984.
- [47] M. J. Smith. A new dynamic traffic model and the existence and calculation of dynamic user equilibria on congested capacity-constrained road networks. *Transportation Research Part B: Methodological*, 27(1):49–63, 1993.

- [48] W. Y. Szeto and Hong K. Lo. A cell-based simultaneous route and departure time choice model with elastic demand. *Transportation Research Part B: Methodological*, 38(7):593–612, 2004.
- [49] N. B. Taylor. The contrans dynamic traffic assignment model. *Networks and Spatial Economics*, 3(3):297–322, 2003.
- [50] W. S. Vickrey. Congestion theory and transport investment. *The American Economic Review*, 59(2):251–260, 1969.
- [51] S. Yagar. Dynamic traffic assignment by individual path minimization and queuing. *Transportation Research*, 5(3):179–196, 1971.

A Appendix: Omitted Proofs

A.1 Proof of Lemma 2.1

In addition we need the following property thereafter:

(a) Let $\theta > \theta' \geq 0$. If no flow is assigned to e between θ' and θ then q_e decreases, i.e.:

$$\int_{\theta'}^{\theta} f^+(\vartheta) d\vartheta = 0 \quad \implies \quad q_e(\theta) \leq q_e(\theta').$$

Moreover $q_e(\theta) > 0$ implies $\theta + q_e(\theta) = \theta' + q_e(\theta')$.

Proof. In order to prove (i) let $\theta > \theta' \geq 0$. Assuming $\theta + q_e(\theta) < \theta' + q_e(\theta')$ gives us the following. Note that the cumulative in- and outflow functions are monotonically increasing.

$$\begin{aligned} 0 &\leq F_e^+(\theta) - F_e^+(\theta') \\ &\stackrel{(7)}{=} F_e^-(\theta + \tau_e + q_e(\theta)) - F_e^-(\theta' + \tau_e + q_e(\theta')) \\ &\leq 0 \end{aligned}$$

Therefore equality must hold especially for the last inequality which together with (7):

$$F_e^+(\theta') = F_e^-(\theta + \tau_e + q_e(\theta)) = F_e^-(\theta' + \tau_e + (\theta - \theta' + q_e(\theta)))$$

But from our assumption we get $q_e(\theta') > \theta - \theta' + q_e(\theta) \geq 0$ and we have a contradiction according to (7). Therefore (i) must hold.

Claim A.0.1.

The function q_e is left continuous for every edge $e \in E$.

Proof of Claim A.0.1. We have to show $\lim_{\theta' \nearrow \theta} q_e(\theta') = q_e(\theta)$ for all $\theta > 0$. From (i) we get that $\lim_{\theta' \nearrow \theta} \theta' + q_e(\theta')$ exists and therefore also $\lim_{\theta' \nearrow \theta} q_e(\theta')$. The following we get from the continuity of cumulative flow functions.

$$\begin{aligned} F_e^-(\theta + \tau_e + \lim_{\theta' \nearrow \theta} q_e(\theta')) &= F_e^-(\lim_{\theta' \nearrow \theta} (\theta' + \tau_e + q_e(\theta'))) \\ &= \lim_{\theta' \nearrow \theta} F_e^-(\theta' + \tau_e + q_e(\theta')) \stackrel{(7)}{=} \lim_{\theta' \nearrow \theta} F_e^+(\theta') = F_e^+(\theta) \end{aligned}$$

Because of (7) this leads to $\lim_{\theta' \nearrow \theta} q_e(\theta') \geq q_e(\theta)$. From (i) we get $\lim_{\theta' \nearrow \theta} \theta' + q_e(\theta') \leq \theta + q_e(\theta)$ implying $\lim_{\theta' \nearrow \theta} q_e(\theta') \leq q_e(\theta)$ and therefore proving this claim. \square

Claim A.0.2. Let $\bar{q}_e(\theta) := \lim_{\theta' \searrow \theta} q_e(\theta')$. Then it holds $\int_{\theta + q_e(\theta)}^{\theta + \bar{q}_e(\theta)} f_e^-(\vartheta + \tau_e) d\vartheta = 0$.

Proof of Claim A.0.2. First we remark that the existence of $\bar{q}_e(\theta)$ is ensured by (i). The continuity of the cumulative flow functions implies:

$$\begin{aligned} F_e^-(\lim_{\theta' \searrow \theta} (\theta' + \tau_e + q_e(\theta'))) &= \lim_{\theta' \searrow \theta} F_e^-(\theta' + \tau_e + q_e(\theta')) \\ &\stackrel{(7)}{=} \lim_{\theta' \searrow \theta} F_e^+(\theta') = F_e^+(\theta) \stackrel{(7)}{=} F_e^-(\theta + \tau_e + q_e(\theta)). \end{aligned}$$

Subtracting the left most term this directly implies this claim because of $\lim_{\theta' \searrow \theta} \theta' + q_e(\theta') = \theta + \bar{q}_e(\theta)$. \square

For showing (a) we know from (7):

$$F_e^+(\theta) = F_e^+(\theta') = F_e^-(\theta' + \tau_e + q_e(\theta')) = F_e^-(\theta + \tau_e + (\theta' - \theta + q_e(\theta'))).$$

Thus (7) and the monotonicity of F_e^- gives us $q_e(\theta) \leq \max\{0, \theta' - \theta + q_e(\theta')\} \leq q_e(\theta')$. If moreover $q_e(\theta) > 0$ we get $q_e(\theta) = \theta' - \theta + q_e(\theta')$ proving (a).

As the last part of this proof we have to show (ii). Firstly we assume that q_e is continuous. Then for every $\epsilon > 0$ and for every time $\theta \geq 0$ there exists an $0 < \epsilon_1 \leq \frac{\epsilon}{2}$ such that $q_e(\theta + \epsilon_1) - \frac{\epsilon}{2} \leq q_e(\theta) \leq q_e(\theta - \epsilon_1) + \frac{\epsilon}{2}$ implying $\theta + \epsilon_1 + q_e(\theta + \epsilon_1) \leq \theta + q_e(\theta) + \epsilon$ and $\theta - \epsilon_1 + q_e(\theta - \epsilon_1) \geq \theta + q_e(\theta) - \epsilon$. Thus if no measurable amount of flow reach the head of e within the time interval $(\theta + \tau_e + q_e(\theta) - \epsilon, \theta + \tau_e + q_e(\theta) + \epsilon)$ we get from the definition of q_e :

$$0 = \int_{\theta + q_e(\theta) - \epsilon}^{\theta + q_e(\theta) + \epsilon} f_e^-(\vartheta + \tau_e) d\vartheta \geq \int_{\theta - \epsilon_1 + q_e(\theta - \epsilon_1)}^{\theta + \epsilon_1 + q_e(\theta + \epsilon_1)} f_e^-(\vartheta + \tau_e) d\vartheta = \int_{\theta - \epsilon_1}^{\theta + \epsilon_1} f_e^+(\vartheta) d\vartheta. \quad (18)$$

But this shows that the implication of (ii) is valid.

For the other direction it is enough to consider only the right continuity because of Claim (A.0.1). Because of Claim (A.0.2) we know for a time $\theta \geq 0$ that q_e is (right) continuous at θ whenever $\theta + q_e(\theta) + \tau_e$ is contained in an left open interval where the outflow rate is strictly positive. Thus it is enough to show that whenever there exists a time $\theta \geq 0$ and an $\epsilon > 0$ such that $\int_{\theta}^{\theta + \epsilon} f_e^+(\vartheta) d\vartheta = 0$ holds then q_e is right continuous at θ . Because of (a) we get in this case that $q_e(\theta) \geq q_e(\theta')$ for all $\theta' \in [\theta, \theta + \epsilon]$ implying $\lim_{\theta' \searrow \theta} q_e(\theta') \leq q_e(\theta)$. On the other hand we know from (i) that $\lim_{\theta' \searrow \theta} \theta' + q_e(\theta') \geq \theta + q_e(\theta)$ such that we can conclude $\lim_{\theta' \searrow \theta} q_e(\theta') \geq q_e(\theta)$. But this proves the right continuity at θ and therefore (ii). \square

A.2 Proof of Theorem 2.4

Proof. Because the condition in (ii) is satisfied obviously if e is contained in a shortest path at time θ it is enough to consider only edges e and times θ such that e does not lie on a shortest path at time θ .

(i) \Rightarrow (ii) Let $\theta \geq 0$ and $e = vw \in E$ be an edge which is not contained in a shortest path at time θ , i.e. $\ell_w(\theta) < \ell_v(\theta) + \tau_e + q_e(\ell_v(\theta))$. Defining $\theta_1 := \max\{0, \sup\{\theta' \geq 0 \mid \ell_v(\theta') + \tau_e + q_e(\ell_v(\theta')) \leq \ell_w(\theta)\}\}$ we have $\ell_w(\theta') \leq \ell_w(\theta) < \ell_v(\theta') + \tau_e + q_e(\ell_v(\theta'))$ for all $\theta' \in (\theta_1, \theta]$. Thus e does not occur in a shortest path within the time interval $(\theta_1, \theta]$. Because of Definition 2.3 we get:

$$0 = F_e^+(\ell_v(\theta)) - F_e^+(\ell_v(\theta_1)) = F_e^-(\ell_v(\theta) + \tau_e + q_e(\ell_v(\theta))) - F_e^-(\ell_v(\theta_1) + \tau_e + q_e(\ell_v(\theta_1)))$$

We know $\ell_v(\theta_1) + \tau_e + q_e(\ell_v(\theta_1)) \leq \ell_w(\theta) < \ell_v(\theta) + \tau_e + q_e(\ell_v(\theta))$. Thus we get (ii) because F_e^- is monotonically increasing.

(ii) \Rightarrow (i) Let $\theta \geq 0$ and $e = vw \in E$ be an edge such that e is not contained in a shortest path at time θ , i.e. $\ell_w(\theta) < \ell_v(\theta) + \tau_e + q_e(\ell_v(\theta))$. Assuming the continuity of the edge waiting time functions and therefore the continuity of the label functions there exists an $\epsilon > 0$ such that $\ell_w(\theta + \epsilon) < \ell_v(\theta + \epsilon) + \tau_e + q_e(\ell_v(\theta + \epsilon))$. Thus we get from the nonnegativity of the flow rate functions:

$$0 \leq \int_{\ell_v(\theta - \epsilon)}^{\ell_v(\theta + \epsilon)} f_e^+(\vartheta) d\vartheta = \int_{\ell_v(\theta - \epsilon) + \tau_e + q_e(\ell_v(\theta - \epsilon))}^{\ell_v(\theta + \epsilon) + \tau_e + q_e(\ell_v(\theta + \epsilon))} f_e^-(\vartheta) d\vartheta \leq \int_{\ell_w(\theta + \epsilon)}^{\ell_v(\theta + \epsilon) + \tau_e + q_e(\ell_v(\theta + \epsilon))} f_e^-(\vartheta) d\vartheta = 0$$

which proves statement (i). \square

A.3 Proof of Theorem 2.6

Before presenting the proof of the theorem, we first discuss the following intuitive real-world example.

Example A.1. Suppose you are at the airport and, since you are already late, you want to get to the gate as quickly as possible. But first you have to check-in. Afterwards, you head for the security check in order to finally get to your gate and board the aircraft. But there is a waiting queue in front of the check-in counter and another queue at the security check. The question is how quickly you should approach the end of the waiting queue at the check-in counter. Of course, as long as the last person in line remains the same, i.e., no one else enters the line, it does not matter at what time you line up—you always leave the check-in counter at the same time. However, Definition 2.3 of routing along shortest path tells you that you should line up as early as possible and Theorem 2.6 also tells you why. It says that if you do not reach the end of the waiting queue as early as possible then other persons may overtake you. And, of course, if there are people behind you who wants to check in at the same counter, they could overtake you if you do not line up immediately.



Figure 1: Construction of the b -flow instance used in the proof of Theorem 2.6. Below the edges the in- and outflow of the dynamic Nash equilibrium (left) and the flow value of the b -flow (right) are shown. Above nodes the corresponding b -values are displayed.

Proof of Theorem 2.6. The main observation we need in order to prove this lemma is the following equation which we get from the definitions of x_e^+ , x_e^- , and q_e in (10) and (7), respectively.

$$x_e^+(\theta) - x_e^-(\theta) = F_e^+(\ell_v(\theta)) - F_e^-(\ell_w(\theta)) = F_e^-(\ell_v(\theta) + \tau_e + q_e(\ell_v(\theta))) - F_e^-(\ell_w(\theta)) . \quad (19)$$

Because of Theorem 2.4 this equation implies directly the equivalence of (i) and (ii).

In order to prove the equivalence of (ii) and (iii) we construct a static b -flow instance (see Figure 1). We split every edge $e = vw \in E$ by inserting a node v_e into two consecutive edges vv_e and $v_e w$. The supply-demand vector of the corresponding b -flow instance is defined as follows. For every node $v \in V \setminus \{s, t\}$ we set $b_v(\theta) := 0$ and for every new node v_e with $e \in E$ we define $b_{v_e}(\theta) := x_e^+(\theta) - x_e^-(\theta)$. Note that we have defined $b_s(\theta)$ and $b_t(\theta)$ in (11). Further we know from (5) and (19) together with the nonnegativity of the outflow rate functions that only node s has a supply, i.e., a positive b -value.

Consider the following static flow. For every edge $e = vw \in E$ we set the flow value on edge vv_e to $x_e^+(\theta)$ and the flow value on edge $v_e w$ to $x_e^-(\theta)$. We claim that this static flow is a feasible b -flow. To prove this we need to check the flow conservation conditions. By construction and (11), flow conservation is fulfilled at s , t , and also at the new nodes v_e , $e \in E$. It remains to verify flow conservation at nodes $v \in V \setminus \{s, t\}$. The following equation follows from linearity of the integral operator and condition (3).

$$\begin{aligned} \sum_{e \in \delta^-(w)} x_e^-(\theta) &= \sum_{e \in \delta^-(w)} \left(\int_0^{\ell_w(\theta)} f_e^-(\vartheta) d\vartheta \right) = \int_0^{\ell_w(\theta)} \left(\sum_{e \in \delta^-(w)} f_e^-(\vartheta) \right) d\vartheta \\ &= \int_0^{\ell_w(\theta)} \left(\sum_{e \in \delta^+(w)} f_e^+(\vartheta) \right) d\vartheta = \sum_{e \in \delta^+(w)} \left(\int_0^{\ell_w(\theta)} f_e^+(\vartheta) d\vartheta \right) = \sum_{e \in \delta^+(w)} x_e^+(\theta). \end{aligned}$$

Thus we have a feasible b -flow on the constructed instance. In particular, the sum over all supplies and demands must be 0, i.e. $\sum_{v \in V} b_v(\theta) + \sum_{e \in E} b_{v_e}(\theta) = 0$. But, as already mentioned, the source s is the only node with a supply, i.e., a positive b -value. Therefore, the supply of s is equal to the demand of t if and only if all other nodes have neither a supply nor a demand meaning that their b -values are 0. This proves the equivalence of (ii) and (iii). \square

A.4 Proof of Theorem 3.1

Proof. Since we know that $b_s(\theta) + b_t(\theta) \geq 0$, it is enough to show $b_s(\theta) + b_t(\theta) \leq M\epsilon$. First we observe that whenever two flow units use the same s - t -path, then the unit which starts later at s cannot overtake flow of the other, i.e., they traverse the network one after another. This is easily seen by induction on the number of edges of the used path.

On the other hand, one flow unit can overtake at most one flow unit on every other path. Assume that there is one flow unit using s - t -path P_1 which overtakes more than one flow unit of an s - t -path P_2 . Therefore the first particle of the last flow unit using P_2 reaches a node $v \in V(P_2)$ strictly later than the first particle of the flow unit on P_1 . Thus v is the node where the flow unit on P_1 overtakes more than one flow unit using P_2 . But because the node labels increase monotonically, the length of the subpath of P_1 to v is shorter than the length of the subpath of P_2 to v at the time when the last flow unit on P_2 starts traversing the network. This is a contradiction because the GENERAL ITERATIVE ALGORITHM sends the flow units always along a currently shortest s - t -path. Thus one flow unit can overtake at most $M - 1$ other flow units which proves the lemma. \square

A.5 Proof of Theorem 3.2

Proof. We prove that for every time $T_0 \in \mathbb{R}_+$ there exists a Nash flow over time for the flow which arrives at s until time T_0 . Thus we assume that u_0 is 0 from time T_0 on. Therefore there exists a time T_e for every edge $e \in E$ such that for all flow over times satisfying conditions (i) and (iii) of the Nash flow over time definition we have that the in- and outflow rate of e is equal to 0 for all times greater than T_e .

Further let $(\epsilon_i)_{i \in \mathbb{N}}$ be a zero sequence, for all $i \in \mathbb{N}$ let f_i be the flow over time computed by the GENERAL ITERATIVE ALGORITHM on (G, u, τ, s, t, u_0) with $\epsilon := \epsilon_i$ and let (F_i^+, F_i^-) be the corresponding pair of families of cumulative flow functions. Since each cumulative flow function is the antiderivative of a nonnegative Lebesgue integrable function they are absolutely continuous and of bounded variation. Moreover since the flow rate function of an edge is bounded by a Lebesgue integrable capacity function the corresponding sequence of cumulative in- and outflow functions, respectively, has uniformly bounded total variation. Further each cumulative flow function is 0 at time 0.

Now Helly's selection theorem¹ ensures the existence of an accumulation point (F^+, F^-) of these (F_i^+, F_i^-) 's, i.e. there exists a subsequence which converges pointwise again (F^+, F^-) . But this also implies that every cumulative flow function of (F^+, F^-) is also absolute continuous. Further Kōmura's theorem² states that every absolute continuous function has a density. This means that there exists corresponding flow rate functions (f^+, f^-) to (F^+, F^-) .

Now it is not hard to see that f is a flow over time satisfying conditions (i) and (iii) of the Nash flow over time definition. Moreover because of Lemma 3.1 and Theorem 2.6 condition (ii) is also satisfied. Thus (f^+, f^-) is a Nash flow over time. \square

In order to prove Theorem 3.5 we need the following lemma.

Lemma A.2. *Let f be a flow over time which sends flow only over currently shortest paths on a network (G, u, τ, s, t, u_0) . Further let $e = vw \in E$ be an edge and $\theta \geq 0$ be a time such that there exists a nonzero waiting queue in front of e at time $\ell_v(\theta)$, i.e. $q_e(\ell_v(\theta)) > 0$. Then e is contained in a shortest path at time θ .*

Proof. We have to show $\ell_v(\theta) + \tau_e + q_e(\ell_v(\theta)) = \ell_w(\theta)$. For it let θ_1 be the smallest time such that no measurable amount of flow is assigned to e within the time interval $[\ell_v(\theta_1), \ell_v(\theta)]$. Then for each $\epsilon > 0$ there exists an $\theta_\epsilon \in [\theta_1 - \epsilon, \theta_1)$ such that flow is assigned to e at time $\ell_v(\theta_\epsilon)$. This means that e is contained in a shortest path at time θ_ϵ . Let ϵ tend to zero. Since the label and edge waiting time functions are continuous we get $\ell_v(\theta_1) + \tau_e + q_e(\ell_v(\theta_1)) = \ell_w(\theta_1)$. But this implies $\ell_v(\theta) + \tau_e + q_e(\ell_v(\theta)) = \ell_w(\theta_1)$ because of Lemma 2.1.(a) (Statement found in the Appendix A.1). Further we know that the label functions are increasing which completes the proof because of (9). \square

Proof of Theorem 3.5. From Theorem 2.6 and Definition 2.7 we know that $x(\theta + \epsilon) - x(\theta)$ is a static flow of value $\int_\theta^{\theta+\epsilon} u_0(\vartheta) d\vartheta$ for every $\epsilon > 0$. Since the flow conservation equation for static flows is differentiable $\frac{dx_e}{d\theta}(\theta)$ is also a static flow with flow value $u_0(\theta)$. This follows from the right-continuity of the initial capacity functions are.

Thus we have to show that $(\frac{dx_e}{d\theta}(\theta))_{e \in E(G_\theta)}$ and $(\frac{d\ell_v}{d\theta}(\theta))_{v \in V(G_\theta)}$ satisfy the thin flow with resetting conditions (12) - (15) according to the edge set $E_1 := \{e \in E \mid q_e(\theta) > 0\}$. Because the label and the edge waiting time functions are right-continuous there exists an $\epsilon > 0$ such that for all $\theta' \in [\theta, \theta + \epsilon)$ we have that $G_{\theta'}$ is a subgraph of G_θ and $q_e(\theta') > 0$ for all edges $e \in E_1$.

The condition (12) for the label of s is implied by the equation (8) defining the label ℓ_s from which we get:

$$\frac{d\ell_s}{d\theta}(\theta) = \frac{d\theta}{d\theta}(\theta) = 1 = \frac{u_0(\theta)}{u_0(\theta)}.$$

¹ The statement of Helly's theorem is the following: Let $U \subseteq \mathbb{R}$ be an open set and let $(g_n : U \rightarrow \mathbb{R}_+)_{n \in \mathbb{N}}$ be a sequence of functions. Assume that the total variation of each g_i on every compact set $W \subseteq U$ is bounded by the same constant (only depending on W). Further assume that there exist a point $\mu \in U$ such that the set $\{g_i(\mu) \mid i \in \mathbb{N}\}$ is bounded. Then there exists a subsequence which converges pointwise.

² The statement of Kōmura's theorem is the following: Every absolute continuous function defined on a closed interval of \mathbb{R} has a density.

For proving the other conditions we consider three cases and show that the conditions (13)-(15) are satisfied in every case. For it let $e = vw \in E(G_\theta)$ be an edge which is contained in a currently shortest s - t -path at a time $\theta \geq 0$.

Case 1 The edge e fits this case if there exists an $\epsilon > 0$ such that for all $\theta' \in (\theta, \theta + \epsilon]$ we have $q_e(\ell_v(\theta')) > 0$. That means a waiting queue is built or occurs which does not decrease to zero over a small time interval. Note that if $e \in E_1$ then e belongs to this case. Condition (iii) of Definition 2.7 gives us:

$$x_e(\theta + \epsilon) - x_e(\theta) = \int_{\ell_w(\theta)}^{\ell_w(\theta + \epsilon)} f_e^-(\vartheta) d\vartheta = \int_{\ell_w(\theta)}^{\ell_w(\theta + \epsilon)} u_e(\vartheta) d\vartheta.$$

Further the mean value theorem of the integral calculus ensures the existence of a number $u_\epsilon \in [\min\{u_e(\theta') \mid \theta' \in (\ell_w(\theta), \ell_w(\theta + \epsilon))\}, \max\{u_e(\theta') \mid \theta' \in (\ell_w(\theta), \ell_w(\theta + \epsilon))\}]$ such that $x_e(\theta + \epsilon) - x_e(\theta) = (\ell_w(\theta + \epsilon) - \ell_w(\theta)) \cdot u_\epsilon$. Now we divide both sides of the last equation by ϵ and let ϵ tend to 0. With the right continuity of the capacity functions we obtain:

$$\frac{d\ell_w}{d\theta}(\theta) = \frac{dx_e}{d\theta}(\theta) \cdot \frac{1}{u_\epsilon(\ell_w(\theta))}$$

Therefore condition (15) is satisfied in this case. Further condition (13) is also satisfied because the label functions are monotonically increasing. In order to show that condition (14) is also valid in this case we have to show that $\frac{d\ell_v}{d\theta}(\theta) \leq \frac{d\ell_w}{d\theta}(\theta)$ if there is no waiting queue in front of e at time θ meaning $\ell_v(\theta) + \tau_e = \ell_w(\theta)$. Because we know that e is contained in a shortest path for all times in $(\theta, \theta + \epsilon]$ we can conclude that $\ell_v(\theta + \epsilon) - \ell_v(\theta) = \ell_w(\theta + \epsilon) - \ell_w(\theta) - q_e(\ell_v(\theta + \epsilon)) \leq \ell_w(\theta + \epsilon) - \ell_w(\theta)$. But this gives us the desired result if we divide both sides by ϵ and let ϵ tend to 0.

Case 2 Here we consider the case that there exists an ϵ such that for all $\theta' \in (\theta, \theta + \epsilon]$ we have $\ell_v(\theta') + \tau_e + q_e(\ell_v(\theta')) > \ell_w(\theta')$, i.e. edge e is not contained in a shortest path for all times in $(\theta, \theta + \epsilon]$. Note that this case is disjoint to Case 1 because of Lemma A.2. Further we know that $q_e(\ell_v(\theta')) = 0$ for all $\theta' \in [\theta, \theta + \epsilon]$. Therefore we have $\ell_v(\theta + \epsilon) - \ell_v(\theta) > \ell_w(\theta + \epsilon) - \ell_w(\theta)$. Further we know that no flow is assigned to e during the time interval $(\ell_v(\theta), \ell_v(\theta + \epsilon)]$, i.e. $x_e(\theta + \epsilon) - x_e(\theta) = 0$. Thus dividing both sides of the last inequality and of the last equation by ϵ we get

$$\frac{d\ell_w}{d\theta}(\theta) \leq \frac{d\ell_v}{d\theta}(\theta) \quad \text{and} \quad \frac{dx_e}{d\theta}(\theta) = 0 \quad (20)$$

if we let ϵ tend to 0. Thus condition (13) is satisfied and the two other conditions do not fall in this case.

Case 3 In this case we firstly consider the complement of Case 2. This means for every $\epsilon > 0$ there exists an $\theta_\epsilon \in (\theta, \theta + \epsilon]$ such that $\ell_v(\theta_\epsilon) + \tau_e + q_e(\ell_v(\theta_\epsilon)) = \ell_w(\theta_\epsilon)$. Because we can use the fact that we need not consider situations which fall in Case 1 we can assume further that there exists an $\theta' \in (\theta, \theta_\epsilon]$ such that $q_e(\ell_v(\theta')) = 0$. Let $\theta'_\epsilon \in (\theta, \theta_\epsilon]$ be the supremum over these θ' . Because we assume that the edge waiting time functions are continuous θ'_ϵ is in fact a maximum implying $q_e(\ell_v(\theta'_\epsilon)) = 0$. Further we know that between the times θ'_ϵ and θ_ϵ there is always a nonzero waiting queue. Lemma A.2 and the continuity of the label functions show that e occurs also in a shortest path at time θ'_ϵ , i.e. $\ell_v(\theta'_\epsilon) + \tau_e = \ell_w(\theta'_\epsilon)$. But this leads to $\ell_v(\theta'_\epsilon) - \ell_v(\theta) = \ell_w(\theta'_\epsilon) - \ell_w(\theta)$. If we divide both sides of the last equation by $\theta'_\epsilon - \theta$ and let ϵ tend to 0 we get the following. Note that we assume that the derivatives exist.

$$\frac{d\ell_w}{d\theta}(\theta) = \frac{d\ell_v}{d\theta}(\theta)$$

Therefore condition (13) is satisfied. Because condition (15) does not belong to this case we only have to show that condition (14) is valid. For it we have to show that $\frac{dx_e}{d\theta}(\theta) \cdot \frac{1}{u_\epsilon(\ell_w(\theta))} < \frac{d\ell_w}{d\theta}(\theta)$. To prove this we apply like in Case 1 the mean value theorem of the integral calculus to the right most side of the following inequality which we get from condition (1) in the flow over time model.

$$x_e(\theta + \epsilon) - x_e(\theta) = \int_{\ell_w(\theta)}^{\ell_w(\theta + \epsilon)} f_e^-(\vartheta) d\vartheta \leq \int_{\ell_w(\theta)}^{\ell_w(\theta + \epsilon)} u_e(\vartheta) d\vartheta.$$

If we divide for the last time both sides by ϵ and let ϵ tend to 0 we get the desired result and complete the proof. \square

A.6 Nash Flow Over Time Algorithm

NASH FLOW OVER TIME ALGORITHM

Input: A network (G, u, τ, s, t, u_0) .

Output: A Nash flow over time (f^+, f^-) .

- (1) Initialize the flow rate functions $f_e^+, f_e^- : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ for all edges $e \in E$:

$$f_e^{+/-}(\theta) := 0 \quad \text{for all } \theta.$$

- (2) Initialize the node labels $\ell \in \mathbb{R}_+^V$ with $\ell_v := \text{dist}_\tau^G(s, v)$ and set $q := 0 \in \mathbb{R}_+^E$, $\theta := 0 \in \mathbb{R}_+$.

- (3) Define the shortest path network G_θ as the subgraph of G induced by the edge set $E(G_\theta) := \{e = vw \in E \mid \ell_w - \ell_v \geq \tau_e\}$ and $E_1 := \{e \in E \mid q_e \geq 0\}$.

- (4) Compute a thin flow x' of value u_0 with resetting on E_1 on (G_θ, u, s, t, u_0) and let ℓ' be the corresponding node labels.

- (5) Find the largest $\alpha > 0$ such that:

$$\begin{aligned} \ell_w + \alpha \ell'_w - \ell_v - \alpha \ell'_v &\leq \text{dist}_\tau^{G \setminus E(G_\theta)}(v, w) && \text{for all } v, w \in V(G_\theta) \\ \ell_w + \alpha \ell'_w - \ell_v - \alpha \ell'_v &\geq \tau_e && \text{for all } e = vw \in E_1 \end{aligned}$$

- (6) Update the flow over time parameters in this order:

$$\begin{aligned} f_e^+(\ell_v) &:= \frac{x'_e}{\ell'_v} && \text{for all } e = vw \in E(G_\theta) \\ f_e^-(\ell_w) &:= \frac{x'_e}{\ell'_w} && \text{for all } e = vw \in E(G_\theta) \\ \ell_v &:= \ell_v + \alpha \ell'_v && \text{for all } v \in V(G_\theta) \\ q_e &:= \max\{0, \ell_w - \ell_v - \tau_e\} && \text{for all } e = vw \in E(G_\theta) \\ \theta &:= \theta + \alpha \end{aligned}$$

- (7) Update the label functions for nodes in $V \setminus V(G_\theta)$.

- (8) If $\alpha < \infty$ go to (3).

A.7 Proof of Theorem 4.3

Proof. The relation $\ell'_{\max} \geq q^*$ is obvious because at least one edge in a sparsest cut must have congestion at least q^* in any b -flow. Thus we have to show that $\ell_{\max} \leq q^*$. Consider the cut $s \in X \subsetneq V$ defined by $X := \{v \in V \mid \ell'_v < \ell'_{\max}\}$. Since the labels of the nodes in X are strictly smaller than the labels of the nodes not in X , there is no flow on any edge in $\delta^-(X)$. Further the congestion of any edge in $\delta^+(X)$ is at least ℓ_{\max} . This leads to $\ell_{\max} \leq \frac{x(\delta^+(X))}{u(\delta^+(X))} = \frac{b(X)}{u(\delta^+(X))} \leq q^*$ because q^* is the congestion of a sparsest cut. \square

A.8 Proof of Theorem 4.4

Proof. Consider two thin flows $x', \tilde{x} \in \mathbb{R}^{E(G)}$. We prove by induction on the number nodes that the corresponding edge labels $\ell', \tilde{\ell}$ are identical. Then this must also hold for the corresponding node labels. If there is only one node s nothing has to be proved. Thus let us assume that there are several nodes. Lemma 4.3 shows that the maximum edge label ℓ'_{\max} is unique and equal to the congestion of a sparsest cut. Thus the edges with this maximal label are the same since x', \tilde{x} are static flows minimizing the maximal edge congestion. Thus we know the flow on edges contained in $\delta^+(X)$ where X is a sparsest cut because the label of such edges must be defined by their congestion, i.e. $x'_e = \tilde{x}_e = \ell'_{\max} u_e$.

Now we delete the node set $V \setminus X$. Then, x' and $\tilde{x} \in \mathbb{R}^{E(G)}$ are thin b -flows on the induced subgraph $G[X]$ according to the new node balances $b'(v) := b(v) - x'(\delta_{G'}^+(v) \cap \delta^+(X))$, $v \in X$. Since the graph $G[X]$ has less nodes than G , we can apply the induction hypothesis and conclude this part of the proof.

To see that we can compute a thin flow of given flow value in polynomial time, note that the induction is constructive and results in an algorithm where in any iteration we have to find a sparsest cut for a b -flow instance. This can be done in polynomial time. Moreover the number of iterations is bounded by the number of nodes and thus this algorithm is polynomial. \square

A.9 A Bad Example




Figure 2: An example with an arbitrary high price of anarchy.

The following example shows that the price of anarchy can be arbitrarily large, i.e., the price of anarchy of the presented example is $\Omega(m)$ where m equals the number of edges. The graph of the underlying instance is shown in Figure 2.

For the edge capacities we define $u_k := u_{e_k}$ and $u^k := u_{e^k}$ for all $k = 1, \dots, m$ and assume $u^k = \sum_{i=1}^k u_i$ for all $k = 1, \dots, m$. The s - t path P_k is defined by the edge set $\{e^m, e^{m-1}, \dots, e^k, e_k\}$ for all $k = 1, \dots, m$. Further let $\tau_k := \sum_{e \in E(P_k)} \tau_e$ be the free flow transit time of the path P_k for each $k = 1, \dots, m$.

In a dynamic Nash flow the free flow transit times should ensure the following behavior. At time zero the first flow particles use only path P_1 (implying that $\tau_1 < \tau_k$ for all $k = 2, \dots, m$). Since $u^m > u^{m-1} > \dots > u^1$ in front of every edge of P_1 a waiting queue is built, which increase linearly. Therefore the times for traversing the paths P_k , $k = 1, \dots, m$ increase monotonically with the time when flow is originated at s and for $m \geq k > l \geq 1$ the slope of the transit time of P_k is smaller than the slope of the transit time of P_l , especially the transit time of P_1 has the greatest slope. Thus at a certain time α the time for traversing P_1 becomes equal to the time for traversing one other path $P \in \{P_k \mid k \in \{2, \dots, m\}\}$ for flow particles originated at s at time α . This means that the next flow particles have to choose P_1 and P in a dynamic Nash flow. Now the important aspect of this instance is, that at time α not only the transit time of P becomes equal to the transit time of P_1 but also the transit times of all other paths P_2, \dots, P_m . (Note that therefore the free flow transit times τ_k of the paths P_k must be monotonically increasing.) This means that from time α flow particles use the whole network in order to reach t . Summarizing the free flow transit times ensure that the first flow particles use only P_1 and suddenly from time α on the whole network is used in a dynamic Nash equilibrium.

Next we model this expected behavior precisely. For $k = 1, \dots, m$ let $\ell_t^k(\alpha)$ be the arrival time using P_k for flow originated at s at time α by assuming that up to this time all flow units use only P_1 in order

to reach t . Thus (because of the NASH FLOW OVER TIME ALGORITHM for constant transit times — first iteration):

$$\ell_t^k(\alpha) = \tau_k + \frac{u^m}{u^k} \cdot \alpha \quad \iff \quad u^m \alpha = (\ell_t^k(\alpha) - \tau_k) u^k$$

Because for flow originated at s at time α all s - t -paths must have the same duration we get:

$$\begin{aligned} \tau_1 + \frac{u^m}{u^1} \cdot \alpha &= \tau_k + \frac{u^m}{u^k} \cdot \alpha =: \ell_t(\alpha) && \text{for all } k = 2, \dots, m \\ \iff \tau_1 &= \tau_k + u^m \alpha \left(\frac{1}{u^k} - \frac{1}{u^1} \right) && \text{for all } k = 2, \dots, m \end{aligned}$$

Note that for edge capacities satisfying the conditions of this example there exists edge free flow transit times ensuring the last equalities. Simply set $\tau_{e^k} := 0$ for all $k = 1, \dots, m$. Then we have $\tau_{e_1} = \tau_1 = 0$ and $\tau_{e_k} = \tau_k$. Thus the remaining edge transit times are computable with the last equalities.

Let $F_{\text{NE}}(\theta)$ be the amount of flow arriving at t until time θ in a Nash flow over time and $F_{\text{SO}}(\theta)$ be the amount of flow arriving at t until time θ in a system optimum. It is not hard to see, that in a system optimum the inflow rate on each path P_k is equal to u_k from time zero on. Thus the path P_k contributes an inflow rate of u_k to s from time τ_k on. We call this flow an *earliest arrival flow*. Because a Nash flow over time has to satisfy the non-overtaking condition we know $F_{\text{NE}}(\ell_t(\alpha)) = u^m \alpha$. The corresponding value for the earliest arrival flow is $F_{\text{SO}}(\ell_t(\alpha)) = \sum_{k=1}^m (\ell_t(\alpha) - \tau_k) u_k$. Thus an lower bound of the price of anarchy is given by (In fact this is the price of anarchy):

$$\frac{F_{\text{SO}}(\ell_t(\alpha))}{F_{\text{NE}}(\ell_t(\alpha))} = \frac{\sum_{k=1}^m (\ell_t(\alpha) - \tau_k) u_k}{u^m \alpha} = \sum_{k=1}^m \frac{(\ell_t(\alpha) - \tau_k) u_k}{u^m \alpha} = \sum_{k=1}^m \frac{u_k}{u^k}$$

This shows that the price of anarchy can increase linearly in the number of edges (set $u^k := 2^k$ for example). Further consider the price of anarchy restricted to instances with unit capacities this shows that the price of anarchy increase logarithmically in the number of edges (Set $u_k = 1$ and replace e^k by k parallel edges. Then the right sum is equal to the Harmonic series and the number of edges is quadratically in m).

A.10 A very bad Multicommodity Example




Figure 3: A multicommodity example were no measurable amount of flow reaches its sink in a Nash flow.

Consider the network shown in Figure 3. Assume that the free flow transit time of each edge is equal to 1 and that the capacity is also equal to 1 over the entire time period. Further there are three kinds of flow particles which want to travel from source s_i to sink t_i for $i = 1, 2, 3$. Moreover assume that all flow particles which wants to traverse the network are present at their corresponding source at time 0. Thus in front of every edge there is a waiting queue of infinite size at time 0 because all flow units line up at the tail of the edge they want to traverse first. Since the transit times are 1 every flow particle has to wait at the end of the waiting queue of the second edge it has to traverse. But the waiting time is equal to ∞ because of the infinite length of the queue and the bounded capacities. Thus no flow particle arrive at their corresponding source.