

Nash Equilibria and the Price of Anarchy for Flows Over Time

Ronald Koch · Martin Skutella

Received: date / Accepted: date

Abstract We study Nash equilibria 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 are time-dependent. This more realistic setting has not been tackled from the viewpoint of algorithmic game theory yet; but there is a rich literature on game theoretic aspects of flows over time in the traffic community.

We present a novel characterization of Nash equilibria for flows over time. It turns out that Nash flows over time can be seen as a concatenation of special static flows. The underlying flow over time model is the so-called *deterministic queuing model* that is very popular in road traffic simulation and related fields. Based upon this, we prove the first known results on the price of anarchy for flows over time.

Keywords Flow over time · Nash equilibrium · Routing game · Deterministic queuing model

1 Introduction

In a groundbreaking paper, Roughgarden and Tardos [42] (see also Roughgarden's book [41]) analyze the price of anarchy for selfish routing games in networks. Such routing games are based upon a classical static flow problem with convex latency functions on the edges of the network. In a Nash equilibrium, flow particles (infinitesimal flow units) selfishly choose an origin-destination path of minimum latency.

Supported by DFG Research Center MATHEON "Mathematics for key technologies" in Berlin. An extended abstract of this work has been presented at the 2nd International Symposium on Algorithmic Game Theory, see [25].

TU Berlin, Inst. f. Mathematik, MA 5-2, Str. des 17. Juni 136, 10623 Berlin, Germany
E-mail: {koch,skutella}@math.tu-berlin.de

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; see, e. g., [5,37]. In contrast to static flow models, flow values on edges may change with time in these applications. Moreover, flow does not progress instantaneously but travels at a certain pace through the network which is determined by transit times on the edges. Both temporal features are captured by *flows over time* (sometimes also called *dynamic* flows) which were introduced by Ford and Fulkerson [15,16].

Another crucial phenomenon in many of those applications mentioned above is the variation of time taken to traverse an edge with the current (and maybe also past) flow situation on this edge. 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 [5,11,19,30,37,38]. In particular, all of these *flow over time* models have so far resisted a rigorous algorithmic analysis of Nash equilibria and the price of anarchy.

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 a capacity. The capacity of an edge bounds the rate (flow per time unit) at which flow can leave this 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 traverse an edge than its capacity allows, the flow particles queue up at the end of the edge and wait in line before they actually enter the head node. When a new flow particle wants to traverse an edge, the time needed to arrive at the head thus consists of the fixed free flow transit time plus the waiting time. In the traffic literature, this flow over time model is known as *deterministic queuing model*.

Related Literature. Flows over time with fixed transit times were introduced by Ford and Fulkerson [15,16]. For more details and further references on these classical flows over time we refer, for example, to [14,43].

So far, Nash equilibria for flows over time were mostly studied within the traffic community. Vickrey [48] and Yagar [51] are the first to introduce this topic. Up to the middle of the 1980's, nearly all contributions consider Nash equilibria on given small instances; see, e. g., [21,29,35,48]. Since then, the number of publications in this area has increased rapidly and Nash equilibria were modeled mathematically. Two main models are distinguished: The *route-choice-model* where a player only chooses an *s-t*-path for the controlled flow particle and the *simultaneous departure-time-route-choice-model* where in addition the departure time is also chosen. For a survey see, e. g., [36]. The considered models can be grouped into four categories: mathematical programming (e. g., [20,28]), optimal control (e. g., [18,39]), variational inequalities (e. g., [12,17,40,45,46]), and simulation-based approaches (e. g., [7,6,31,47,51]). Up to now, variational inequalities are the most common formulation for analyzing Nash equilibria in the context of flows over time.

Many models mentioned above use a path-based formulation of flows over time. Therefore they are often computationally intractable. Edge-based formulations are, for example, considered in [2, 12, 40]. Realistic assumptions on the underlying flow model with respect to traffic are described by Carey [9, 10].

In this paper the *deterministic queuing model* is considered. This model was introduced by Vickrey [48] and later by Hendrickson and Kocur [21]. Smith [44] shows the existence of an equilibrium for this model in a special case. Akamatsu [1, 2] presents an edge-based formulation of the deterministic queuing model on restricted single-source-instances. Akamatsu and Heydecker [3] study Braess's paradox for single-source instances. Braess's paradox [8] 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. Mounce [32, 33] considers the case where the edge capacities can vary over time and states some existence results. Again, it should be mentioned that these results are based on strong assumptions.

Recently, Anshelevich and Ukkusuri [4] analyze discrete routing game models for Nash equilibria in the context of flows over time. They consider how a single splittable flow unit present at source s at time 0 would traverse a network assuming every flow particle is controlled by a different player. The underlying flow model allows to send a positive amount of flow over an edge at each integral points in time. Moreover the transit times are assumed to be constant. Hoefer et. al [22] also consider a discrete routing game. They study existence and complexity properties of pure Nash equilibria and best-response strategies.

Our Contribution. In this paper, we characterize and analyze Nash equilibria for flows over time. Although algorithmic game theory is a flourishing area of research (see, e. g., the recent book [34]), network flows over time have not been studied from this perspective in the algorithms community so far. The main purpose of this paper is to make first steps in this relevant direction, present interesting and novel results, and stimulate further interesting research. We consider the deterministic queuing model in networks with a single source and a single sink. A player controls one flow particle and chooses an s - t -path (route-choice-model) but no departure time which is given a priori.

A precise description of a routing game over time and the underlying flow over time model is given in Section 2. The resulting model of Nash equilibria along with several equivalent characterizations is discussed in Section 3. The main technical contribution of this paper is presented in Section 4. Here we show that a Nash equilibrium can be characterized via a sequence of static flows with special properties. The resulting static flow problems are of interest in their own right. The final Section 5 is devoted to results on the price of anarchy. For the important class of shortest paths networks we prove that every Nash equilibrium is a system optimum. Moreover, a Nash flow over time can be computed in polynomial time by a sequence of sparsest cut computations. Surprisingly, for arbitrary networks, the price of anarchy is not bounded by a constant.

2 A model for routing games over time

In this section we present a model for Nash equilibria in the context of flows over time. In Section 2.1 we define a routing game over time and elaborate on the game theoretic aspect of the model. Then, in Section 2.2 we introduce an appropriate flow over time model which is known as the *deterministic queuing model*.

Throughout the paper we often use the term *flow particle* in order to refer to an infinitesimal flow unit which corresponds to one player and travels along a single path through the network. The terms *flow rate* and *supply rate* both refer to an amount of flow per time unit.

2.1 From static routing games to routing games over time

Consider a network consisting of a directed graph $G := (V, E)$ with node set V and edge set E . Further, there is a source $s \in V$ and a sink $t \in V$. Each flow particle is a player and the strategy set of each player is the set \mathcal{P} of all s - t -paths.

In a static routing game, the players' decisions yield a static s - t -flow μ of value d where d is the given supply at the source s . Moreover, there is a continuous cost (or payoff) function ℓ_P for each path $P \in \mathcal{P}$ such that $\ell_P(\mu)$ is the cost that a player choosing path P has to pay. The static flow $\mu = (\mu_P)_{P \in \mathcal{P}}$ is a Nash flow if, for all $P \in \mathcal{P}$ with $\mu_P > 0$, it holds that $\ell_P(\mu) = \min_{P' \in \mathcal{P}} \ell_{P'}(\mu)$.

The situation is considerably more complicated when we turn to routing games over time. Here we assume that supply, i. e., players, occur at the source node s over time at a fixed rate d . We can thus identify each player with the point in time θ at which its corresponding flow particle originates at the source. In particular, and in contrast to static routing games, players are not identical. The routing decisions of players yield a flow over time $\mu = (\mu_P)_{P \in \mathcal{P}}$ where μ_P is a function determining the flow rate $\mu_P(\theta)$ at which flow enters path P at time θ and it holds that $\sum_{P \in \mathcal{P}} \mu_P(\theta) = d$, for all θ . Thus, also $\ell_P(\mu)$ is a function which assigns a cost $\ell_P(\mu)(\theta)$ to every point in time θ . That is, the cost experienced by a flow particle that originates at the source at time θ and chooses path P is equal to $\ell_P(\mu)(\theta)$.

In this paper we restrict to cost functions where $\ell_P(\mu)(\theta)$, $P \in \mathcal{P}$, is the time when a flow originating at s at time θ arrives at the sink t . This time depends upon the particular model of flows over time that we consider which is described in Section 2.2 below.

Like in static routing games, a Nash equilibrium is characterized by a flow over time μ where no player has an incentive to change her chosen path in order to reduce her cost.

Definition 1 (Nash flow over time) Let μ be a flow over time determining the routing decisions of the players in a routing game over time. Then, μ is a Nash equilibrium (Nash flow over time) if, for almost all θ and for all $P \in \mathcal{P}$ with $\mu_P(\theta) > 0$, it holds that $\ell_P(\mu)(\theta) = \min_{P' \in \mathcal{P}} \ell_{P'}(\mu)(\theta)$.

This definition is an immediate generalization of the definition of static Nash flows under the assumption that the payoff functions are continuous. A closer look at Definition 1 shows that the continuity of the payoff functions ℓ_P is also essential here since it ensures the following. If an actual routing μ does not fit Definition 1, then some measurable set of players must have an incentive to switch to a path P of minimum latency. Hence, the switching of the players must not imply a large increase of the payoff function ℓ_P .

2.2 An appropriate flow over time model

Although Definition 1 is an immediate generalization of static Nash flows, it is still a highly nontrivial problem to come up with an appropriate flow over time model. Here the main issue are the cost functions ℓ_P , $P \in \mathcal{P}$. For static routing games, these cost functions are not given explicitly, but implicitly via edge latency functions. The cost of a path $P \in \mathcal{P}$ is the sum of the latencies of its edges. The latency of an edge e is a function of the load μ_e of that edge which can easily be computed as follows: $\mu_e := \sum_{P \in \mathcal{P}: e \in E(P)} \mu_P$.

The situation is considerably more complicated for flows over time. Here, it is usually a highly nontrivial problem to compute the flow rate function μ_e of edge e from given flow rate functions $(\mu_P)_{P \in \mathcal{P}}$. Consider a flow particle that enters a path $P \in \mathcal{P}$ at a certain time θ . Notice that the time at which this particle arrives at an edge $e \in E(P)$ depends on the latencies experienced on the predecessor edges on path P . This fact induces involved dependencies among the flow rate functions $(\mu_e)_{e \in E}$ of the edges. As a consequence, given a flow over time $(\mu_P)_{P \in \mathcal{P}}$, determining the cost (overall latency) of a flow particle entering path P at time θ is, in general, a highly nontrivial task. For more details on this so-called *dynamic network loading problem* we refer to [49,50]. Nevertheless, for the deterministic queuing model described below, these difficulties can be handled at least for the case of Nash flows over time.

Let (G, u, τ, s, t) be a network consisting of a directed graph $G := (V, E)$, edge capacities $u_e \in \mathbb{R}_+$, $e \in E$, constant *free flow transit times* $\tau_e \in \mathbb{R}_+$, $e \in E$, a source $s \in V$, and a sink $t \in V$. We assume without loss of generality that there are no incoming edges at the source s and no outgoing edges at the sink t and that every node v is reachable from s . The capacity u_e of an edge e bounds the rate at which flow leaves edge e at its head node. The basic concept of the considered flow over time model are waiting queues which build up at the head (exit) of an edge if, at some point in time, more flow particles want to leave an edge than the capacity of the edge allows. The free flow transit time of an edge determines the time for traversing an edge if the waiting queue is empty. Thus, the (*flow-dependent*) *transit time* on an edge is the sum of the free flow transit time and the current waiting time. We think of the edges as corridors with large entries and small exits, which are wide enough for storing all waiting flow particles (point-queue-model); see Fig. 1.

Every flow particle arriving at an intermediate node v immediately enters the next edge on its path without any delay. In the following we give a more

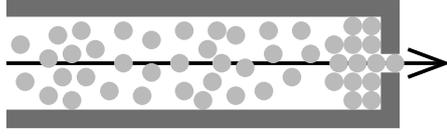


Fig. 1 If more flow particles want to leave an edge than its capacity allows, they form a waiting queue.

precise mathematical description of the flow over time model. A flow over time is defined by two families of flow rate functions. For an edge e we have a Lebesgue integrable inflow rate f_e^+ meaning that the rate at which flow enters the tail of e at time θ is $f_e^+(\theta) \geq 0$; moreover, the Lebesgue integrable outflow rate f_e^- describes the rate of flow $f_e^-(\theta) \geq 0$ leaving at the head of e at time θ . 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 that has entered e before time θ is $F_e^+(\theta)$ and the amount of flow which has traversed e completely before time θ is $F_e^-(\theta)$. Note that F_e^+ and F_e^- are (absolutely) continuous and monotonically increasing, for each $e \in E$.

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

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

We also have to impose several kinds of flow conservation constraints. Firstly, flow can only traverse an edge if it has previously been assigned to this edge:

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

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 :

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

In order to ensure that flow which is assigned to an edge must leave this edge again at some point in time, we proceed as follows: Regarding condition (2), the value $F_e^+(\theta)$ is the amount of flow entering edge e before time θ which is equal to the flow arriving at the end of the waiting queue of e until time $\theta + \tau_e$. Moreover, the value $F_e^-(\theta + \tau_e)$ is the amount of flow arriving at the head node of e until time $\theta + \tau_e$. Thus, $F_e^+(\theta) - F_e^-(\theta + \tau_e)$ is the amount of flow in the waiting queue at time $\theta + \tau_e$. We impose the natural condition that, whenever the waiting queue on edge e is nonempty, the flow rate leaving e at its head equals its capacity u_e . Therefore the waiting time spent by a flow particle entering the tail of e at time θ is equal to

$$q_e(\theta) := \frac{F_e^+(\theta) - F_e^-(\theta + \tau_e)}{u_e} \quad \text{for all } e \in E, \theta \in \mathbb{R}_+. \quad (4)$$

The interpretation of $q_e(\theta)$ as the waiting time for flow particles arriving at time θ on edge e is based on the assumption that the first-in-first-out (FIFO) property holds on edge e . That is, no flow particle overtakes any other flow particle within the waiting queue. Since the free flow transit times are constant, the FIFO property holds for the entire edge.

Proposition 1 *For any edge $e \in E$, the following statements are valid:*

- (i) *The function $\theta \mapsto \theta + q_e(\theta)$ is monotonically increasing and continuous.*
- (ii) *$F_e^+(\theta) = F_e^-(\theta + \tau_e + q_e(\theta))$ for all $\theta \in \mathbb{R}_+$.*
- (iii) *Consider two points in time $\theta_2 > \theta_1 \geq 0$ such that $\int_{\theta_1}^{\theta_2} f^+(\vartheta) d\vartheta = 0$ and $q_e(\theta_2) > 0$. Then we have $\theta_1 + q_e(\theta_1) = \theta_2 + q_e(\theta_2)$.*

Proof Let $\theta_2 > \theta_1 \geq 0$. For proving (i) note that $F_e^-(\theta_2 + \tau_e) - F_e^-(\theta_1 + \tau_e) \leq u_e(\theta_2 - \theta_1)$ because of (1). Hence $q_e(\theta_1) - q_e(\theta_2) \leq \theta_2 - \theta_1$ is implied by $F_e^+(\theta_1) - F_e^+(\theta_2) \leq 0$ and (4). The continuity follows from (4) and the continuity of F_e^+ and F_e^- .

Since we assume that, whenever the waiting queue on edge e is nonempty, the flow rate leaving e at its head equals the capacity u_e , statement (ii) is directly implied by (4).

It remains to prove statement (iii). Since we have $F_e^+(\theta_1) = F_e^+(\theta_2)$, we obtain $F_e^+(\theta') = F_e^+(\theta_2)$ for all $\theta' \in [\theta_1, \theta_2)$ because F_e^+ is monotonically increasing. Since F_e^- is also monotonically increasing and $q_e(\theta_2) > 0$, this implies $F_e^+(\theta') - F_e^-(\theta' + \tau_e) \geq F_e^+(\theta_2) - F_e^-(\theta_2 + \tau_e) > 0$ for all $\theta' \in [\theta_1, \theta_2)$. Hence, there is a nonzero waiting queue on e for all times in $[\theta_1, \theta_2)$, implying that f_e^- is equal to u_e for all times in $[\theta_1, \theta_2)$. Thus, we conclude that $q_e(\theta_1) - q_e(\theta_2) = \theta_2 - \theta_1$. \square

As already mentioned, the flow-dependent transit time $\tau_e(\theta)$ experienced by flow particles entering e at time θ is $\tau_e(\theta) := \tau_e + q_e(\theta)$. Note, that the mapping $f_e^+ \mapsto \tau_e$, which maps an inflow rate f_e^+ to the transit time function τ_e , is continuous. This is due to the fact that additional flow of value $\epsilon > 0$ can cause at most an additional delay of $\frac{\epsilon}{u_e}$. Therefore, for the case of the deterministic queuing model, a Nash equilibrium of a routing game over time is well-defined in terms of Definition 1.

3 Characterizing Nash flows over time

The main aspect of Nash equilibria in flow models is the selfish routing of flow particles which are identified with players. As mentioned in Section 2.1, we assume that flow occurs at the source s according to a fixed supply rate $d \in \mathbb{R}_+$. 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 immediately enters the first edge of that path.

We consider two – apparently related – classes of flows over time. In the first class, every flow particle travels along “currently shortest paths” only.

In the second class, every flow particle tries to overtake as many other flow particles as possible while not being overtaken by others. It turns out that the latter condition leads to a flow where no particle overtakes any other particle. Moreover, we show that the two classes of flows over time coincide and are, in fact, Nash flows over time.

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 s to the sink t as quickly as possible. Let $\ell_v(\theta)$ be the earliest point in time at which this flow particle can arrive at node $v \in V$. Then,

$$\ell_v(\theta) + \tau_e + q_e(\ell_v(\theta)) \geq \ell_w(\theta) \quad \text{for each } e = vw \in E. \quad (5)$$

On the other hand, for each node $w \in V \setminus \{s\}$, there exists at least one incoming edge $e = vw \in \delta^-(w)$ such that equality holds in (5). That is, the flow particle can use edge e in order to arrive at node w as early as possible, i. e., at time $\ell_w(\theta)$. Moreover, we have $\ell_s(\theta) = \theta$ for all $\theta \geq 0$. Therefore, we define the *label functions* $\ell_w : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ as follows:

$$\ell_w(\theta) := \begin{cases} \theta & \text{for } w = s, \\ \min_{e=vw} \ell_v(\theta) + \tau_e + q_e(\ell_v(\theta)) & \text{for } w \in V \setminus \{s\}. \end{cases} \quad (6)$$

The label functions can be computed simultaneously for all times θ by adapting the shortest path algorithm of Bellman and Ford¹. The next proposition follows from (6) and Proposition 1.

Proposition 2 *For each node $v \in V$, the label function ℓ_v is monotonically increasing and continuous.*

Before proceeding with the main discussion of this section, we first consider an intuitive real-world example. This example will also illustrate the subsequent definitions and results.

Example 1 Suppose you are at the airport and, since you are already late, you want to get to your departure 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. 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, if there are people behind you who want to check in at the same counter, they could overtake you if you do not line up immediately.

¹ The update procedure of Bellman-Ford for a certain label $\ell_w(\theta)$ is applied for all times θ simultaneously and, hence, is seen as an operation on functions. If we use Dijkstra's algorithm instead, we have to maintain the set of already finalized nodes separately for each time θ . Thus, we also have to apply the update procedure of Dijkstra separately for each θ .

In a Nash equilibrium, flow should always be sent over currently shortest s - t -paths only. We say that 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 We say that flow is only sent along currently shortest paths if, for each edge $e = vw \in E$, the following condition holds for almost all times $\theta \geq 0$:

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

We emphasize the following aspect of a routing satisfying Definition 2: In a static shortest path all subpaths are also shortest paths if the transit times are positive. But this is no longer true if we consider the dynamic case as illustrated in Example 1. Here, as long as the last person in line remains the same you always leave the check-in counter at the same time. So, in principle, you could decide to make a detour – maybe for buying a small present for your family – and you can still leave the check-in counter as early as possible. However, in this case you will use at least one edge which does not lie on a currently shortest path. Since Definition 2 forbids entering that edge, you have to line up at the check-in counter as early as possible

As we see below the condition in Definition 2 is equivalent to the condition that every particle tries to overtake as much other flow as possible while not being overtaken. The latter condition is in fact a *universal FIFO condition*. That is, it is equivalent to the statement that no flow particle can possibly overtake any other flow particle.

In order to model the universal FIFO condition more formally, we consider again 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, for each edge $e = vw \in E$, we define 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:

$$x_e^+(\theta) := F_e^+(\ell_v(\theta)) , \quad x_e^-(\theta) := F_e^-(\ell_w(\theta)) \quad \text{for all } \theta \geq 0. \quad (7)$$

Thus, the amount of flow $b_s(\theta) := d \cdot \theta$ 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 satisfy

$$b_s(\theta) = \sum_{e \in \delta^+(s)} x_e^+(\theta) \quad \text{and} \quad b_t(\theta) = - \sum_{e \in \delta^-(t)} x_e^-(\theta) . \quad (8)$$

By definition, $b_s(\theta)$ is always nonnegative and $b_t(\theta)$ is always non-positive. If $b_s(\theta) > -b_t(\theta)$, then the considered flow particle overtakes other flow particles. And if $b_s(\theta) < -b_t(\theta)$, then the flow particle is overtaken by other flow particles. This motivates the following definition.

Definition 3 We say that *no flow overtakes any other flow* if, for each point in time $\theta \geq 0$, it holds that $b_s(\theta) = -b_t(\theta)$.

Intuitively, this definition must be satisfied by a Nash flow over time: 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 . Because the function $\theta \mapsto \theta + \tau_e + q_e(\theta)$ is monotonically increasing for each edge e (see Proposition 1), flow particle p_1 can avoid being overtaken by p_2 and improve its cost (arrival time at the sink) by choosing the same path as p_2 .

Now we are able to prove the equivalence of the universal FIFO condition and the condition that flow only uses currently shortest paths. Further, both conditions characterize Nash flows over time. In addition, a further equivalent characterization is given. With respect to Example 1, Theorem 1 also tells you why you should line up immediately. It says that if you do not reach the end of the waiting queue as early as possible, then other persons may overtake you.

Theorem 1 *For a given flow over time, the following statements are equivalent:*

- (i) *Flow is only sent along currently shortest paths.*
- (ii) *For each edge $e \in E$ and at all times $\theta \geq 0$, it holds that $x_e^+(\theta) = x_e^-(\theta)$.*
- (iii) *No flow overtakes any other flow.*
- (iv) *The given flow over time is a Nash flow over time.*

In the proof of Theorem 1, the following lemma plays an important role. It gives a more global characterization of when flow is being sent only along currently shortest paths (Definition 2 gives only a pointwise characterization).

Lemma 1 *For a given flow over time, the following statements are equivalent:*

- (i) *Flow is only being sent along currently shortest paths.*
- (ii) *For each edge $e = vw \in E$ and for all $\theta \geq 0$, it holds that*

$$F_e^-(\ell_v(\theta) + \tau_e + q_e(\ell_v(\theta))) = F_e^-(\ell_w(\theta)) . \quad (9)$$

Proof Equation (9) is obviously fulfilled if edge e is contained in a shortest path at time θ . In the following, it is thus 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., $\ell_w(\theta) < \ell_v(\theta) + \tau_e + q_e(\ell_v(\theta))$. Let

$$\theta_1 := \max\{0, \sup\{\theta' \geq 0 \mid \ell_w(\theta) \geq \ell_v(\theta') + \tau_e + q_e(\ell_v(\theta'))\}\} .$$

be the latest point in time at which a flow particle is able to arrive at w via e before time $\ell_w(\theta)$ (and 0 if no such point in time exists). By definition of θ_1 ,

$$\ell_w(\theta') \leq \ell_w(\theta) < \ell_v(\theta') + \tau_e + q_e(\ell_v(\theta')) \quad \text{for all } \theta' \in (\theta_1, \theta].$$

Thus, edge e does not occur in a shortest path within the time interval $(\theta_1, \theta]$. Because of Definition 2 and Proposition 1(ii), we get

$$\begin{aligned} 0 &= F_e^+(\ell_v(\theta)) - F_e^+(\ell_v(\theta_1)) \\ &= F_e^-(\ell_v(\theta) + \tau_e + q_e(\ell_v(\theta))) - F_e^-(\ell_v(\theta_1) + \tau_e + q_e(\ell_v(\theta_1))) . \end{aligned} \quad (10)$$

Equation (10) implies (ii) because

$$\ell_v(\theta_1) + \tau_e + q_e(\ell_v(\theta_1)) \leq \ell_w(\theta) < \ell_v(\theta) + \tau_e + q_e(\ell_v(\theta))$$

and because F_e^- is monotonically increasing.

(ii) \Rightarrow (i): Let $\theta \geq 0$ and $e = vw \in E$ an edge that is not contained in a shortest path at time θ , i. e., $\ell_w(\theta) < \ell_v(\theta) + \tau_e + q_e(\ell_v(\theta))$. By Propositions 1 and 2, there exists an $\epsilon > 0$ such that $\ell_w(\theta + \epsilon) < \ell_v(\theta - \epsilon) + \tau_e + q_e(\ell_v(\theta - \epsilon))$. Thus, the nonnegativity of the flow rate functions yields:

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

This yields statement (i). \square

Within the proof of Theorem 1 we construct a static b -flow. For convenience of the reader, we give the definition of b -flows here (for more details we refer to [27]).

Definition 4 Let H be a graph with node set $V(H)$ and edge set $E(H)$ and $b'_v \in \mathbb{R}$ be a real value for each node $v \in V(H)$. The *supply-demand vector* $(b'_v)_{v \in V(H)}$ has to satisfy $\sum_{v \in V(H)} b'_v = 0$. A static flow $x' := (x'_e)_{e \in E(H)} \in \mathbb{R}_+^{E(H)}$ is called *b' -flow* if it satisfies the *flow conservation constraint*

$$\sum_{e \in \delta^+(v)} x'_e - \sum_{e \in \delta^-(v)} x'_e = b'_v \quad \text{for each } v \in V(H).$$

If in addition, for each edge $e \in E(H)$, an edge capacity $u_e \in \mathbb{R}$ is given, a *feasible b -flow* has to satisfy $x'_e \leq u_e$ for each $e \in E(H)$.

Proof (of Theorem 1) The main observation we need is the following equation which we get from Proposition 1(ii) and the definitions of x_e^+ , x_e^- in (7):

$$\begin{aligned} x_e^+(\theta) - x_e^-(\theta) &= F_e^+(\ell_v(\theta)) - F_e^-(\ell_w(\theta)) \\ &= F_e^-(\ell_v(\theta) + \tau_e + q_e(\ell_v(\theta))) - F_e^-(\ell_w(\theta)) . \end{aligned} \quad (11)$$

Because of Lemma 1, this equation implies the equivalence of (i) and (ii).

In order to prove the equivalence of (ii) and (iii), we construct a static b -flow instance. We replace each edge $e = vw \in E$ by a new node v_e and two edges vv_e and $v_e w$; see Fig. 2. The supply-demand vector of the corresponding b -flow

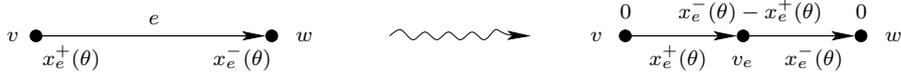


Fig. 2 Construction of the b -flow instance used in the proof of Theorem 1. Below the edges the in- and outflow of the dynamic Nash equilibrium (left) and the flow value of the b -flow (right) are shown. Above nodes the corresponding b -values are displayed.

instance is defined as follows. For each node $v \in V \setminus \{s, t\}$ we set $b_v(\theta) := 0$ and for each new node v_e , $e \in E$, we define $b_{v_e}(\theta) := x_e^-(\theta) - x_e^+(\theta)$. An intuitive explanation of $b_{v_e}(\theta)$ is as follows. Recall that, at the time when a flow particle originating at the source at time θ gets to node v along a shortest path, the amount of flow having previously entered edge $e = vw$ is $x_e^+(\theta)$. Similarly, when the same flow particle travels along a shortest path to w , the amount of flow that has previously arrived at w via edge $e = vw$ is $x_e^-(\theta)$. Hence, a flow particle which arrives at w via e can increase the amount of overtaken flow by $-b_{v_e}(\theta)$ if it goes directly along a currently shortest path to w . However, the condition in (ii) states that this flow particle cannot improve its situation, i. e., $b_{v_e}(\theta) = 0$.

Note that we have defined $b_s(\theta)$ and $b_t(\theta)$ in (8). It follows from (11) and 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 each edge $e = vw \in E$, 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 constraints. By construction and (8), flow conservation is fulfilled at nodes 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).

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

Thus we have a feasible b -flow on the constructed instance. In particular, the sum over all supplies and demands is equal to zero. That is,

$$\sum_{v \in V} b_v(\theta) + \sum_{e \in E} b_{v_e}(\theta) = 0 .$$

Note that this shows $b_s(\theta) \geq -b_t(\theta)$ for all $\theta \geq 0$ which we use later in this proof. Because the source s is the only node with a positive b -value, the supply

of s is equal to the demand of t if and only if the b -values of all other nodes are 0.

This proves the equivalence of (ii) and (iii). It remains to prove that (iv) is equivalent to the other statements.

(i) \Rightarrow (iv): The cost $\ell_P(\theta)$ of a shortest s - t -path P at time θ (see Definition 1) is equal to the label $\ell_t(\theta)$ of t at time θ . This is due to the fact, that the deterministic queuing model satisfies the FIFO condition on each edge. Therefore $\ell_t(\theta)$ cannot be influenced by flow particles originating at the source after time θ . Thus, a flow over time which sends flow only along currently shortest paths is a Nash flow over time.

(iv) \Rightarrow (iii): Assume that (iii) is violated, i. e., there exist a point in time θ such that $b_s(\theta) \neq -b_t(\theta)$. As mentioned above, this implies $b_s(\theta) > -b_t(\theta)$. Hence, there exists flow of value $b_s(\theta) + b_t(\theta)$ which originates at s until time θ but arrives at t strictly later than $\ell_t(\theta)$. Since ℓ_t is monotonically increasing, this flow is not routed along paths with minimum latency, thus yielding a contradiction to (iv). \square

Note that, whenever one of the four statements in Theorem 1 holds, then x^+ and x^- coincide. In this case, for all $\theta \geq 0$, setting $x_e(\theta) := x_e^+(\theta)$ for each $e \in E$, yields a static s - t -flow $x(\theta)$ of value $b_s(\theta)$. In the following, for a flow over time satisfying the universal FIFO condition, we refer to $(x_e(\theta))_{e \in E}$ as the *underlying static flow at time θ* . This flow will be studied in more detail in the next section.

4 A special class of static flows

In this section we study the underlying static flows of a Nash flow over time. It turns out that these static flows have a special structure that can be used to characterize, compute, and analyze Nash flows over time. Further, the network on which these flows are considered is a special subnetwork of the original network.

Definition 5 (Current Shortest Paths Network) Consider a flow over time on a network (G, u, s, t, τ, d) . For $\theta \geq 0$, the *current shortest paths network* G_θ is the subnetwork induced by the edges occurring in a currently shortest path, i. e., edges $e = vw$ with $\ell_v(\theta) + \tau_e + q_e(\ell_v(\theta)) = \ell_w(\theta)$.

Note, that every node v is contained in any current shortest path network since we assume that every node is reachable from s . But in general there are edges which are not contained in a current shortest path network.

Definition 6 (Thin Flow with Resetting) Let (G, u, s, t, d) be a *static* network and $E_1 \subseteq E(G)$ a subset of edges. A *static* flow x' with flow value F

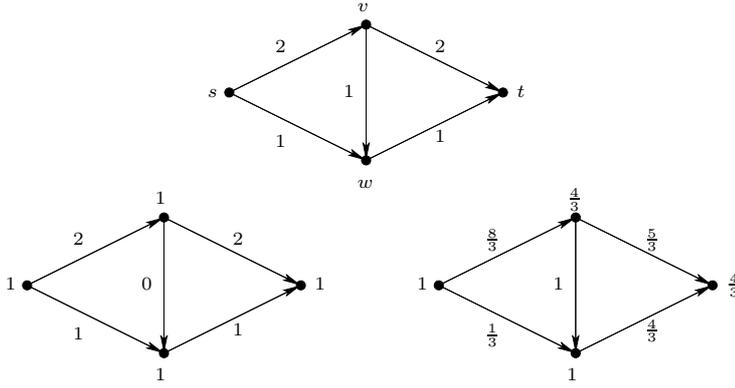


Fig. 3 Thin flow example: On top the network with corresponding edge capacities is shown. In the lower part two thin flows of value 3 with $d = 3$ are shown. On the left instance there is no resetting edge and on the right instance there is resetting on $E_1 = \{sv, vw\}$. The flow is shown on the edges and the corresponding node labels are shown at the nodes.

is a *thin flow with resetting* on E_1 if there exist node labels ℓ' such that:

$$\ell'_s = F/d \quad (12)$$

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

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

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

Further, for an edge e the *edge congestion*² of e is defined by x'_e/u_e .

Let the congestion of a path P be the maximum congestion of its edges, i.e., $\max_{e \in E(P)} x'_e/u_e$. Then, if $E_1 = \emptyset$, the label ℓ'_v of node v is the congestion of all flow-carrying s - v -paths and a lower bound on the congestion of any s - v -path. The name “thin flow with resetting” refers to the special edges 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 edges seen so far and “resets” its congestion to x'_e/u_e . As we will see in Section 5, for the special case $E_1 = \emptyset$, a thin flow with resetting can be computed in polynomial time. Before we proceed with the main discussion, we give an example of a thin flow.

Example 2 Consider the network shown in the upper part of Fig. 3. We assume that $d := 3$ and want to find a thin flow of value 3. Note that the maximal value of a static s - t -flow in this network is also 3. If there is no resetting edge, then we obtain the thin flow shown in the left part of Fig. 3, which is also a maximum s - t -flow. In contrast, on the right side a thin flow with resetting on $E_1 := \{sv, vw\}$ is shown. Verifying the thin flow conditions (12)–(15), we observe that:

² In game theory, “congestion” is mostly associated with congestion games. However, in this paper, “congestion” is used for the percental load of an edge because this is common in network flow theory.

- the label of s satisfies (12),
- the edges sw and wt satisfy (14), where the maximum is attained by the congestion of the particular edge.
- the edge vt satisfies (14), where the maximum is attained by the label of v .
- the resetting edges sv and vw satisfy (15).

Also note that assuming there is no resetting edge, the right flow is not a thin flow *without* resetting, since vw does not satisfy (14). In fact, both thin flows are unique on their particular instance.

Next we show that, for a Nash flow over time, the derivatives of the label functions and of the underlying static flow define a thin flow with resetting. The following theorem is only applicable if the derivatives of the label and the underlying static flow functions exist. However, both the label functions and the underlying static flow functions are monotonically increasing implying that both families of functions are differentiable almost everywhere.

Theorem 2 *Consider a Nash flow over time on a network (G, u, s, t, τ, d) with corresponding label functions $(\ell_v)_{v \in V}$, edge waiting time functions $(q_e)_{e \in E}$ and underlying static flow $(x_e)_{e \in E}$. Let $\theta \geq 0$ such that $\frac{dx_e}{d\theta}(\theta)$ and $\frac{d\ell_v}{d\theta}(\theta)$ exist for each $e \in E$ and $v \in V$. Then, on the current shortest paths network G_θ , the derivatives $(\frac{dx_e}{d\theta}(\theta))_{e \in E(G_\theta)}$ form a thin flow of value d with resetting on the waiting edges $E_1 := \{e \in E \mid q_e(\ell_v(\theta)) > 0\}$. A corresponding set of node labels fulfilling (12) to (15) is given by the derivatives $(\frac{d\ell_v}{d\theta}(\theta))_{v \in V(G_\theta)}$.*

In order to prove Theorem 2, we need the following lemma.

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

Proof We have to show that $\ell_v(\theta) + \tau_e + q_e(\ell_v(\theta)) = \ell_w(\theta)$. Let θ_1 be the earliest 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 a $\theta_\epsilon \in [\theta_1 - \epsilon, \theta_1)$ 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 Proposition 1(iii). Further, we know that the label functions are increasing which completes the proof because of the definition of the label functions in (6). \square

Proof (of Theorem 2) We have to show that $(\frac{dx_e}{d\theta}(\theta))_{e \in E(G_\theta)}$ and the labels $(\frac{d\ell_v}{d\theta}(\theta))_{v \in V(G_\theta)}$ satisfy the thin flow with resetting conditions (12) to (15) with respect to the edge set $E_1 := \{e \in E \mid q_e(\theta) > 0\}$.

Condition (12) for the label of s is implied by equation (6) defining the label ℓ_s . In order to prove the other conditions, we distinguish three cases and show that conditions (13) to (15) are satisfied in every case. Consider an

edge $e = vw \in E(G_\theta)$ which is contained in a currently shortest s - t -path at time $\theta \geq 0$.

Case 1: Edge e fits this case if there exists an $\epsilon > 0$ such that, for all $\theta' \in (\theta, \theta + \epsilon]$, we have $q_e(\ell_v(\theta')) > 0$. This means a waiting queue is built or occurs which does not decrease to zero over a small time interval. In particular, if $e \in E_1$, then e belongs to this case. Because e is used up to its capacity in this case, we get:

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

Dividing both sides of the last equation by ϵ and letting ϵ tend to zero, we obtain

$$\frac{d\ell_w}{d\theta}(\theta) = \frac{dx_e}{d\theta}(\theta) \cdot \frac{1}{u_e} .$$

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 on e , i. e., $\ell_v(\theta) + \tau_e = \ell_w(\theta)$. Because we know that e is contained in a shortest path for all times in $(\theta, \theta + \epsilon]$, we can conclude that

$$\ell_v(\theta + \epsilon) - \ell_v(\theta) = \ell_w(\theta + \epsilon) - \ell_w(\theta) - q_e(\ell_v(\theta + \epsilon)) \leq \ell_w(\theta + \epsilon) - \ell_w(\theta) .$$

This yields the desired result if we divide both sides by ϵ and let ϵ tend to zero.

Case 2: Here we consider the case that there exists an ϵ such that, for all $\theta' \in (\theta, \theta + \epsilon]$, we have $\ell_v(\theta') + \tau_e + q_e(\ell_v(\theta')) > \ell_w(\theta')$. That is, edge e is not contained in a shortest path for all times in $(\theta, \theta + \epsilon]$. Note that this case is disjoint from Case 1 because of Lemma 2. Further, we know that $q_e(\ell_v(\theta')) = 0$ for all $\theta' \in [\theta, \theta + \epsilon]$. Therefore, it holds that $\ell_v(\theta + \epsilon) - \ell_v(\theta) \geq \ell_w(\theta + \epsilon) - \ell_w(\theta)$. Moreover, 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 and letting ϵ tend to zero, yields

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

Thus, condition (13) is satisfied and the two other conditions are not relevant in this case.

Case 3: We first 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 a $\theta' \in (\theta, \theta_\epsilon]$ such that $q_e(\ell_v(\theta')) = 0$. Let $\theta'_\epsilon \in (\theta, \theta_\epsilon]$ be the supremum over these θ' . Because 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 2 and the continuity of the label functions show that e occurs also in a shortest path at time θ'_ϵ , that is, $\ell_v(\theta'_\epsilon) + \tau_e = \ell_w(\theta'_\epsilon)$. 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 zero, we get

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

Therefore, condition (13) is satisfied. Because condition (15) does not apply in this case, we only show that condition (14) is valid. For this, we show that $\frac{dx_e}{d\theta}(\theta) \cdot \frac{1}{u_e} \leq \frac{d\ell_w}{d\theta}(\theta)$. From condition (1) in the flow over time model we get

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

If we divide both sides by ϵ and let ϵ tend to zero, we get the desired result. This completes the proof. \square

The reverse direction of Theorem 2 also holds. If, for all times θ , 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 paths network, then the flow over time is in fact a Nash flow over time. The following Theorem 3 is not the direct conversion of Theorem 2 but rather a corollary of the reverse direction. It shows that a Nash flow over time can be seen as the concatenation of thin flows with resetting, which are *static flows*.

The scenario is the following. We assume that we already have a flow over time $f := (f^+, f^-)$ showing the selfish routing behavior of flow particles originating at s until a certain time θ . That is, f is a Nash flow for the “supply” function

$$d(\vartheta) = \begin{cases} d & \text{for } \vartheta < \theta \\ 0 & \text{for } \vartheta \geq \theta \end{cases} .$$

Note that this does not really fit the model with a constant supply rate. But since the deterministic queuing model satisfies universal FIFO, all of the previous results carry over to this case. We call such a flow f a *restricted Nash flow over time* on $[0, \theta)$.

In order to extend f , we compute a thin flow x' on the current shortest path network (G_θ, s, t, u) of value d with resetting on the waiting edges given by $E_1 := \{e \in E \mid q_e(\ell(v)) > 0\}$. Let ℓ' be the corresponding node labels of the thin flow x' . For an $\alpha > 0$, we extend the node label functions ℓ of f :

$$\ell_v(\vartheta) := \ell_v(\theta) + (\vartheta - \theta) \cdot \ell'_v \quad \text{for all } v \in V \text{ and } \vartheta \in [\theta, \theta + \alpha) . \quad (16)$$

Then, we also extend the flow rate functions:

$$f_e^+(\vartheta) := \frac{x'_e}{\ell'_v} \quad \text{for all } e = vw \in E \text{ and } \vartheta \in [\ell_v(\theta), \ell_v(\theta + \alpha)) , \quad (17)$$

$$f_e^-(\vartheta) := \frac{x'_e}{\ell'_w} \quad \text{for all } e = vw \in E \text{ and } \vartheta \in [\ell_w(\theta), \ell_w(\theta + \alpha)) . \quad (18)$$

The result is called an α -extension of f .

We want to show that an α -extension is a feasible Nash flow over time. For this we have to choose α appropriately such that:

$$\ell_w(\theta) - \ell_v(\theta) + \alpha(\ell'_w - \ell'_v) \geq \tau_e \quad \text{for all } e \in E_1, \quad (19)$$

$$\ell_w(\theta) - \ell_v(\theta) + \alpha(\ell'_w - \ell'_v) \leq \tau_e \quad \text{for all } e \in E \setminus E(G_\theta). \quad (20)$$

It is essential to show that α can be chosen strictly positive. In order to prove this, first observe that $\ell_w(\theta) - \ell_v(\theta) - \tau_e = q_e(\ell_v(\theta)) > 0$ for all edges $e \in E_1$. Hence, there exists an $\alpha_1 > 0$ such that (19) is satisfied for all $\alpha \leq \alpha_1$. The second condition (20) refers to edges $e = vw$ which are not contained in the shortest path network at time θ , i. e., $\ell_w(\theta) < \ell_v(\theta) + \tau_e$. Thus, there exists an $\alpha_2 > 0$ such that the second condition is satisfied for all $\alpha \leq \alpha_2$. This shows the existence of an $\alpha > 0$ satisfying both conditions simultaneously.

In order to extend a restricted flow as far as possible, α must be chosen as large as possible. That is, $\alpha > 0$ can be interpreted as the largest number such that no waiting queue decreases to 0 and no new edge is added to the current shortest path network. In particular, all edges satisfying (19) with equality are removed from the set of resetting edges E_1 , and all edges satisfying (20) with equality are added to the current shortest path network at time $\theta + \alpha$.

If we insert the extension of the node labels in the two conditions on α , we obtain the following for all $\vartheta \in [\theta, \theta + \alpha)$ (using (14) in addition):

$$\ell_w(\vartheta) \geq \ell_v(\vartheta) + \tau_e \quad \text{for all } e = vw \in E(G_\theta) \text{ with } x'_e > 0,$$

$$\ell_w(\vartheta) \leq \ell_v(\vartheta) + \tau_e \quad \text{for all } e = vw \in E \setminus E(G_\theta).$$

This is used in the proof of the following theorem.

Theorem 3 *Let f be a restricted Nash flow over time on $[0, \theta)$ and let $\alpha > 0$ be a positive real number satisfying (19) and (20). Then, the α -extension of f is a restricted Nash flow over time on $[0, \theta + \alpha)$.*

Proof We show first that the α -extension of f is a feasible flow over time. Afterwards we show that the α -extension of f is also a Nash flow over time.

In order to show that the α -extension is a feasible flow over time, we have to show the validity of (1), (2), and (3). Since f is a restricted Nash flow over time on $[0, \theta)$, it is enough to check these conditions from time $\ell_v(\theta)$ on. Further, conditions (1) and (2) are obviously satisfied if $x'_e = 0$. Let $e = vw$ be an edge with $x'_e > 0$ and $\vartheta \in [\theta, \theta + \alpha)$ be a point in time. From Definition 6 we know that $\ell'_w \geq x'_e/u_e$ if $x'_e > 0$, and hence

$$f_e^-(\ell_v(\vartheta)) = \frac{x'_e}{\ell'_w} \leq u_e.$$

This proves (1). In order to prove (2), we observe that

$$F_e^+(\ell_v(\vartheta)) = F_e^+(\ell_v(\theta)) + (\ell_v(\vartheta) - \ell_v(\theta)) \cdot \frac{x'_e}{\ell'_v}$$

and

$$F_e^-(\ell_w(\vartheta)) = F_e^-(\ell_w(\theta)) + (\ell_w(\vartheta) - \ell_w(\theta)) \cdot \frac{x'_e}{\ell'_w}.$$

Since f is a restricted Nash flow on $[0, \theta]$, we know that $F_e^+(\ell_v(\theta)) = F_e^-(\ell_w(\theta))$. In addition, we obtain from the definition of the α -extension that

$$\begin{aligned} (\ell_v(\vartheta) - \ell_v(\theta)) \cdot \frac{x'_e}{\ell'_v} &= (\vartheta - \theta) \cdot \ell'_v \cdot \frac{x'_e}{\ell'_v} \\ &= (\vartheta - \theta) \cdot \ell'_w \cdot \frac{x'_e}{\ell'_w} = (\ell_w(\vartheta) - \ell_w(\theta)) \cdot \frac{x'_e}{\ell'_w}. \end{aligned}$$

This yields $F_e^+(\ell_v(\vartheta)) = F_e^-(\ell_w(\vartheta))$ and proves (2) since $\ell_w(\vartheta) \geq \ell_v(\vartheta) + \tau_e$ if $x'_e > 0$ (which is implied by (19)). Finally, (3) is satisfied because x' satisfies the static flow conservation constraints which implies

$$\sum_{e \in \delta^-(v)} f_e^-(\ell(\vartheta)) = \sum_{e \in \delta^-(v)} \frac{x'_e}{\ell'_v} = \sum_{e \in \delta^+(v)} \frac{x'_e}{\ell'_v} = \sum_{e \in \delta^+(v)} f_e^+(\ell(\vartheta)).$$

Hence, the α -extension of f is a feasible flow over time.

It remains to show that the α -extension is also a Nash flow over time. For this, we show that the extended label functions coincide with label functions defined by (6) and we show that condition (ii) of Theorem 1 is satisfied. If $x'_e > 0$, Condition (ii) of Theorem 1 is already proved since, as shown above, we know that

$$x_e^+(\vartheta) := F_e^+(\ell_v(\vartheta)) = F_e^-(\ell_w(\vartheta)) =: x_e^-(\vartheta).$$

For all other edges, condition (ii) of Theorem 1 is still valid because f is a restricted Nash flow.

Since $\ell'_s = 1$ because of (12), the first condition of (6) regarding the label of s is obviously satisfied. If $x'_e > 0$, the equations $F_e^+(\ell_v(\vartheta)) = F_e^-(\ell_w(\vartheta))$ imply $\ell_v(\vartheta) + \tau_e + q_e(\ell_v(\vartheta)) = \ell_w(\vartheta)$. Next, we consider edges $e = vw$ contained in the shortest path network at time θ with $x'_e = 0$. It is not hard to see that for these edges we have $q_e(\ell_v(\theta)) = 0$. But then condition (13) implies that $\ell_v(\vartheta) + \tau_e \geq \ell_w(\vartheta)$ for $\vartheta \geq \theta$. Hence, it remains to show $\ell_v(\vartheta) + \tau_e \geq \ell_w(\vartheta)$ for all $\vartheta \geq \theta$ for edges which are not in the shortest path network at time θ . But this follows directly from the definition of α .

We have thus shown that the α -extension of f is a restricted Nash flow over time on $[0, \theta + \alpha)$. \square

Theorem 3 shows that a restricted Nash flow over time is extendable by a particular thin flow with resetting. Hence, a dynamic Nash flow can be seen as the concatenation of such static flows. In what follows, we exemplarily show how the α -extension can be used in order to construct a Nash flow over time.

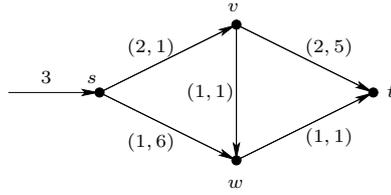


Fig. 4 The network used in Example 3 for which we construct a Nash flow over time. On the incoming edge of the source s the supply rate is shown. For all other edges e , the pair (u_e, τ_e) denotes the capacity u_e and the free flow transit time τ_e .

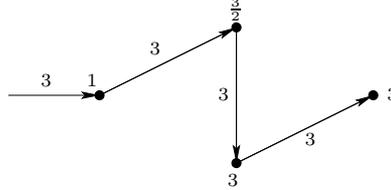


Fig. 5 The thin flow on the zigzag path.

Example 3 Consider the network (G, u, τ, s, t, d) shown in Fig. 4. In order to construct a Nash flow over time, we first start at time 0 with the zero flow and compute an α -extension for some suitable $\alpha > 0$. Recalling Theorem 3, the α -extension results in a restricted Nash flow over time f . Next, we compute again an α -extension for f implying that f remains a restricted Nash flow, but now on a longer time interval. In the following we do not evaluate the in- and outflow rate function explicitly since we want to focus on the interesting part of this Nash flow computation. Note, however, that they can be easily computed using (17) and (18).

Before computing an α -extension for the zero flow, we have to identify the current shortest path network G_0 and need to evaluate the node label functions at time 0. It is quite obvious that G_0 is given by the zigzag path $svwt$ and that $(\ell_s(0), \ell_v(0), \ell_w(0), \ell_t(0)) = (0, 1, 2, 3)$. Since no waiting queue exists at this initial state, we have to find a thin flow x' of value $d = 3$ on (G_0, u, s, t) without resetting, i. e., $E_1 = \emptyset$. Fig. 5 shows x' together with the corresponding node labels ℓ' . Next, we have to choose an α satisfying (19) and (20). Since there is no waiting edge, we have to verify condition (20) for the edges sw and vt , which are not contained in G_0 . Hence, α has to satisfy:

$$\begin{aligned} 2 + 2\alpha &= \ell_w(0) - \ell_s(0) + \alpha(\ell'_w - \ell'_s) \leq \tau_{sw} = 6, \\ 2 + \frac{3}{2}\alpha &= \ell_t(0) - \ell_v(0) + \alpha(\ell'_t - \ell'_v) \leq \tau_{vt} = 5. \end{aligned}$$

Of course, we choose α as big as possible. Therefore we set $\alpha := 2$. Thus the α -extension results in a restricted Nash flow on the time interval $[0, 2)$.

In order to extend this flow further, first note that both inequalities are satisfied with equality and hence, at time 2 both edges, sw and vt , enter the current shortest path network, i. e., $G_2 = G$. Using (16), the label functions at

time 2 are $(\ell_s(2), \ell_v(2), \ell_w(2), \ell_t(2)) = (2, 4, 8, 9)$. This shows that on sv and vw the experienced travel time is greater than the free flow transit time. Thus, there must be a waiting queue on sv and vw at this state. Recalling Example 2, the thin flow x' on (G_2, u, s, t) with resetting on $E_1 = \{sv, vw\}$ is shown in the right part of Fig. 3, and the corresponding node labels are $(\ell'_s, \ell'_v, \ell'_w, \ell'_t) = (1, \frac{4}{3}, 1, \frac{4}{3})$. Like in the first stage, we have to choose an α satisfying (19) and (20). But since $G_2 = G$, we now have to verify condition (19) for the edges sv and vw . Hence, α needs to satisfy

$$\begin{aligned} 2 + \frac{1}{3}\alpha &= \ell_v(2) - \ell_s(2) + \alpha(\ell'_v - \ell'_s) \geq \tau_{sv} = 1, \\ 4 - \frac{1}{3}\alpha &= \ell_w(2) - \ell_v(2) + \alpha(\ell'_w - \ell'_v) \geq \tau_{vw} = 1. \end{aligned}$$

Note that the first inequality is valid for all nonnegative α . Hence, we set α such that the second inequality is satisfied with equality, i. e., $\alpha := 9$. Thus, the α -extension is a restricted Nash flow over time up to time 11. Using (16), the label functions at time 11 are $(\ell_s(11), \ell_v(11), \ell_w(11), \ell_t(11)) = (11, 16, 17, 21)$. This shows that the flow particle which start traversing the network at time 11 has to wait on sv and wt . On the other hand, the waiting queue on vw disappears completely within the time interval $[2, 11)$.

In the last iteration, we have to consider the current shortest path network G_{11} at time 11, which is again equal to G . Next we have to compute a thin flow x' on (G_{11}, u, s, t) with resetting on sv and wt . We can easily see, that x' is equal to the maximum s - t -flow, i. e., $x'_e = u_e$ for all edges $e \neq vw$, and the corresponding node labels are all equal to 1. Without going into details, we note that we can set α to $+\infty$ without violating (19) and (20). Hence, this α -extension results in a Nash flow over time on the network shown in Figure 4 (defined for all nonnegative points in time). Since all node labels of x' are equal to 1, the experienced travel time of each edge remains constant from time 11 on. In particular, the lengths of the waiting queues on sv and wt do not vary over time.

5 Nash flows over time and the price of anarchy

The characterization of Nash flows over time via thin flows with resetting enables us to completely analyze shortest paths networks where every s - t -path has the same total free flow transit time. An important subclass of shortest paths networks are networks where the free flow travel times of each edge is zero. We study the price of anarchy which, in general, is the worst case ratio of the cost of a Nash equilibrium and the cost of a system optimum. In the context of routing games over time, we define the price of anarchy of an instance as the worst case ratio over all points in time θ regarding the following objective:³ For given θ , maximize the amount of flow arriving at the sink until time θ . In particular, according to this definition, earliest arrival

³ This objective is well motivated if we think of, e. g., modeling an evacuation situation.

flows that maximize the amount of flow at the sink simultaneously for all points in time, are system optima. In addition, an earliest arrival flow optimizes also the following objective functions simultaneously for all times θ and for each demand value F , respectively: the average inflow rate in t until time θ , the total time needed to send a demand of value F to t , the average arrival time for a demand of value F (see, e. g., [23]). Hence, the following theorem shows that, on shortest paths networks, the price of anarchy is 1 for each of these objective functions.

Theorem 4 *For shortest paths networks, each Nash flow over time is an earliest arrival flow and thus a system optimum. Moreover, a Nash flow over time can be computed in polynomial time.*

For proving this, we analyze thin flows without resetting. Therefore we first discuss the results related to this part before we give a proof of Theorem 4. The following is an equivalent definition of thin flows without resetting, i. e., thin flows with resetting on $E_1 = \emptyset$.

Definition 7 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, a supply d is given, we initialize $\ell'_s := \frac{|x'|}{d}$ where $|x'|$ is the flow value of x' .

In order to study thin flows, we can restrict to instances with infinite supply rate d , i. e., $\ell'_s = 0$. This is due to the fact that we can model a finite supply simply by adding a dummy source s_0 and an edge s_0s with capacity d 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, i. e., $b_s > 0$ and $b_v \leq 0$ for all $v \in V \setminus \{s\}$. Next we prove some properties of thin flows. 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 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 $\frac{b(X)}{u(\delta^+(X))}$, where $u(\delta^+(X)) := \sum_{e \in \delta^+(X)} u_e$.*

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 the set of nodes $X := \{v \in V \mid \ell'_v < \ell'_{\max}\}$. 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

The last lemma shows that a thin s - t -flow of value equal to the maximum s - t -flow value is also feasible with respect to edge capacities. The next lemma shows that thin flows are unique with respect to node labels. Moreover, thin flows without resetting can be computed in polynomial time by a sequence of sparsest cut computations.

Lemma 4 *Consider a pair of thin flows with the same s - t -flow value or the same node balances b . Then, the node labels ℓ'_v are identical. Moreover, a thin flow of given flow value can be computed in polynomial time.*

Proof Consider two thin flows $x', \tilde{x} \in \mathbb{R}_+^{E(G)}$. We prove by induction on the number of 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. Let us thus assume that there are several nodes. Lemma 3 shows that the maximum edge label ℓ'_{\max} is unique for x' and \tilde{x} and equal to the congestion of a sparsest cut. Let $\delta^+(X)$ be the sparsest cut where $X \subset V$ is inclusionwise minimal. Then, the labels of the edges on and behind $\delta^+(X)$ coincide for x' and \tilde{x} , and are equal to ℓ'_{\max} . Further, we know the flow values on edges contained in $\delta^+(X)$, because the labels of such edges must be defined by their congestion, i. e., $x'_e = \tilde{x}_e = \ell'_{\max} \cdot u_e$.

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

$$b'(v) := b(v) - x'(\delta_G^+(v) \cap \delta^+(X)) \quad \text{for all } v \in X.$$

Since the graph $G[X]$ has less nodes than G , we can apply the induction hypothesis and conclude this part of the proof.

We finally argue that we can compute a thin flow of given flow value in polynomial time. Note that the induction above is constructive and describes an algorithm where, in each iteration, we have to compute 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. Note that an algorithm for computing a sparsest cut usually also computes a flow minimizing the maximum edge congestion. This flow is used to obtain the flow values for edges behind the sparsest cut. \square

Next we give the proof of Theorem 4.

Proof (Proof of Theorem 4) First note that the current shortest path network at time 0 is the entire network. Since at time 0 no waiting occurs, Theorem 3 is valid for $\alpha = \infty$ and shows that a thin flow x' without resetting on the shortest path network results in a Nash flow over time f (defined for all times). Lemma 3 implies that f sends, at each point in time, the maximum possible flow rate into t . This can be seen as follows. Let ℓ' be the corresponding node labels of x' and let F^* be the maximum static flow value on (G, u, s, t) . Then, Lemma 3 shows that the label of t is $\ell'_t = \frac{d}{F^*}$. Hence, the definition of the flow rate functions in an α -extension shows that the inflow rate in t is $\frac{|x'_t|}{\ell'_t} = F^*$.

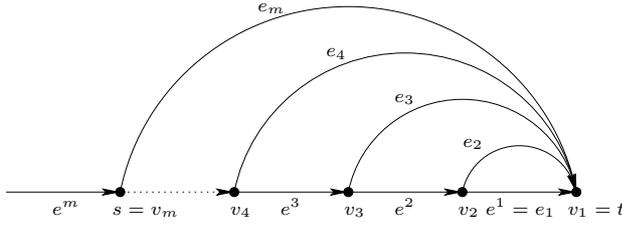


Fig. 6 A family of instances with unbounded price of anarchy

This implies that f is a system optimum since this inflow rate is achieved as early as possible. Lemma 4 shows that this Nash flow over time can be computed in polynomial time. \square

In contrast to static routing games, there exist instances of the routing game over time where the price of anarchy is not constant.

Proposition 3 *There exists a family of instances for which the price of anarchy is $\Omega(m)$, where m is the number of edges. If all edge capacities are equal to 1, there exists a family of instances for which the price of anarchy is still $\Omega(\log(m))$.*

Proof The instances are defined as follows. The underlying graph is shown in Fig. 6. Regarding the capacities $u_{e_k} := u_k$ of the bow edges, for $k = 1, \dots, m$, let $u_k > 0$ be an arbitrary positive real number. The capacities $u_{e^k} := u^k$ of the lower straight edges are given by $u^k := \sum_{i=1}^k u_i$, for $k = 1, \dots, m$. Note that the supply is represented by the edge e^m . Hence, the supply rate is equal to u^m .

The transit times of the lower straight edges are all equal to 0, i. e., $\tau_{e^k} := 0$ for all $k = 1, \dots, m$. Further, for $k = 2, \dots, m$, the transit time $\tau_{e_k} := \tau_k$ of the bow edge e_k is set to

$$\tau_k := \alpha \cdot u^m \left(\frac{1}{u^1} - \frac{1}{u^k} \right)$$

for some $\alpha > 0$. Note that, for $k = 2, \dots, m$, the definition of the capacities implies $u^1 < u^{k-1}$ and hence $\tau_k > 0$. Therefore, the current shortest path network at time 0 is the lower straight path P_1 .

In order to compute a Nash flow over time, let x' be the thin flow on the lower straight path, i.e. $x'_{e^k} = u^m$ for all $k = 1, \dots, m$. Then, corresponding node labels are given by $\ell'_{v_k} := \frac{u^m}{u^k}$. We observe that

$$\alpha'(\ell'_{v_1} - \ell'_{v_k}) \leq \tau_k \quad \text{for all } k = 2, \dots, m \text{ and } 0 \leq \alpha' \leq \alpha$$

where equality holds for $\alpha' = \alpha$. Since $\ell_{v_k}(0) = 0$ for all $k = 2, \dots, m$, condition (20) is satisfied with equality on all bow edges for α . Hence, in terms of Theorem 3, the α -extension of the zero flow over time is a restricted Nash flow over time on $[0, \alpha)$. In addition, at time α , the entire network is the current

shortest path network and a thin flow is given by the edge capacities. Hence, from time α on, the Nash flow over time uses the entire network. Further considering the label functions of the node v_k , we obtain

$$\alpha \frac{u^m}{u^k} = \ell_{v_k}(\alpha) = \ell_t(\alpha) - \tau_k .$$

The left hand side follows from the definition of an α -extension and the right hand side from the definition of the label functions. We rewrite this equality as

$$\alpha \cdot u^m = (\ell_t(\alpha) - \tau_k) \cdot u^k . \quad (21)$$

In order to get a lower bound on the price of anarchy, we consider the point in time $\ell_t(\alpha)$. In the Nash flow described above, the total amount of flow at the sink is $u^m \cdot \alpha$. Since a system optimum, which is an earliest arrival flow, uses the entire network right from the beginning, the total amount of flow at the sink in an earliest arrival flow is $\sum_{k=1}^m (\ell_t(\alpha) - \tau_k) \cdot u_k$. Thus, with (21) we obtain the following lower bound on the price of anarchy:

$$\frac{\sum_{k=1}^m (\ell_t(\alpha) - \tau_k) u_k}{u^m \alpha} = \sum_{k=1}^m \frac{(\ell_t(\alpha) - \tau_k) u_k}{u^m \alpha} = \sum_{k=1}^m \frac{u_k}{u^k} .$$

In fact, this is the exact price of anarchy for this example.

This shows that the price of anarchy increases linearly in the number of edges (set $u^k := 2^k$ for example). If we restrict to instances with unit edge capacities, the price of anarchy still increases logarithmically in the number of edges — set $u_k := 1$ and replace e^k by k parallel edges. Then the sum on the right hand side is equal to the harmonic series and the number of edges is quadratic in m . \square

6 Conclusion and Outlook

We have studied a routing game over time which is based on the deterministic queuing model. In particular, we deduced equivalent characterizations of Nash flows over time. Theorem 1 shows that the routing behavior of flow particles results in a static flow, which we call underlying static flow. Based upon this, we showed in Section 4 that a Nash flow over time can be seen as a sequence of static flows fulfilling special properties. As a result, we were able to prove several bounds on the price of anarchy. In particular, we gave a complete analysis of Nash flows on shortest paths networks where each s - t -path has the same length with respect to free flow transit times.

It is an interesting question whether this approach is applicable to other flow models as well. Flow models satisfying the FIFO condition are of particular importance. It turns out that, even for this general class of flow models, underlying static flows of a given Nash flow exist and can be defined (we refer to the upcoming PhD thesis [24] for details). We believe that studying these

special flows is a promising direction for analyzing Nash flows in the corresponding flow model. Natural candidates for future research are flow models with linear load-dependent transit times (see [26]) as well as the flow model of Daganzo (see [13]). The model of Daganzo is a generalization of the deterministic queuing model in which the size of the waiting queues are bounded by given constants.

Finally, we remark that the presented results remain true with tiny modifications if edge capacities are allowed to vary over time. This model is harder to analyze since the effective transit time of an edge is in general discontinuous if the capacity vanishes over a time interval.

Acknowledgments. The authors are much indebted to Jose Correa for inspiring discussions on the topic of this paper. The authors would also like to thank the referees of [25] and the referees of this paper for numerous valuable suggestions and comments that helped to improve the presentation of the paper.

References

1. T. Akamatsu. A dynamic traffic equilibrium assignment paradox. *Transportation Research B*, 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. E. Anshelevich and S. Ukkusuri. Equilibria in dynamic selfish routing. In M. Mavronicolas, editor, *Proceedings of the 2nd International Symposium on Algorithmic Game Theory*, volume 5814 of *Lecture Notes in Computer Science*, pages 171–182. Springer, 2009.
5. J. E. Aronson. A survey of dynamic network flows. *Annals of Operations Research*, 20(1–4):1–66, 1989.
6. M. Balmer, M. Rieser, K. Meister, D. Charypar, N. Lefebvre, and K. Nagel. MATSim-T: Architecture and simulation times. In A. Bazzan and F. Klügl, editors, *Multi-Agent Systems for Traffic and Transportation Engineering*, chapter III. Information Science Reference, 2009.
7. 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.
8. D. Braess. Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung* 12, 12:258–268, 1968. In German.
9. M. Carey. Link travel times I: Properties derived from traffic-flow models. *Networks and Spatial Economics*, 4(3):257–268, 2004.
10. M. Carey. Link travel times II: Properties derived from traffic-flow models. *Networks and Spatial Economics*, 4(4):379–402, 2004.
11. H.-K. Chen. *Dynamic Travel Choice Models: A Variational Inequality Approach*. Springer, 1999.
12. H.-K. Chen and C.-F. Hsueh. A model and an algorithm for the dynamic user-optimal route choice problem. *Transportation Research B*, 32(3):219–234, 1998.
13. C. F. Daganzo. Queue spillovers in transportation networks with a route choice. *Transportation Science*, 32(1):3–11, 1998.
14. L. K. Fleischer and É. Tardos. Efficient continuous-time dynamic network flow algorithms. *Operations Research Letters*, 23(3-5):71–80, 1998.
15. L. R. Ford and D. R. Fulkerson. Constructing maximal dynamic flows from static flows. *Operations Research*, 6(3):419–433, 1958.

-
16. L. R. Ford and D. R. Fulkerson. *Flows in Networks*. Princeton University Press, 1962.
 17. 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.
 18. 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.
 19. 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, 2002.
 20. S. Han and B. G. Heydecker. Consistent objectives and solution of dynamic user equilibrium models. *Transportation Research B*, 40(1):16–34, 2006.
 21. C. Hendrickson and G. Kocur. Schedule delay and departure time decisions in a deterministic model. *Transportation Science*, 15(1):62–77, 1981.
 22. M. Hoefer, V. Mirrokni, H. Röglin, and S.-H. Teng. Competitive routing over time. In S. Leonardi, editor, *Proceedings of the 5th International Workshop on Internet and Network Economics*, volume 5929 of *Lecture Notes in Computer Science*, pages 18–29. Springer, 2009.
 23. J. J. Jarvis and H. D. Ratliff. Some equivalent objectives for dynamic network flow problems. *Management Science*, 28:106–108, 1982.
 24. R. Koch. PhD thesis, TU Berlin. In preparation.
 25. R. Koch and M. Skutella. Nash equilibria and the price of anarchy for flows over time. In M. Mavronicolas, editor, *Proceedings of the 2nd International Symposium on Algorithmic Game Theory*, volume 5814 of *Lecture Notes in Computer Science*, pages 323–334. Springer, 2009.
 26. E. Köhler and M. Skutella. Flows over time with load-dependent transit times. *SIAM Journal on Optimization*, 15:1185–1202, 2005.
 27. B. Korte and J. Vygen. *Combinatorial Optimization: Theory and Algorithms*. Springer, 4th edition, 2008.
 28. W. H. Lin and H. K. Lo. Are the objective and solutions of dynamic user-equilibrium models always consistent? *Transportation Research A*, 34(2):137–144, 2000.
 29. H. S. Mahmassani and R. Herman. Dynamic user equilibrium departure time and route choice on idealized traffic arterials. *Transportation Science*, 18(4):362–384, 1984.
 30. 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.
 31. H. S. Mahmassani, H. A. Sbayti, and X. Zhou. *DYNASMART-P Version 1.0 Users Guide*. Maryland Transportation Initiative, College Park, Maryland, 2004.
 32. R. Mounce. Convergence in a continuous dynamic queueing model for traffic networks. *Transportation Research B*, 40(9):779–791, 2006.
 33. R. Mounce. Convergence to equilibrium in dynamic traffic networks when route cost is decay monotone. *Transportation Science*, 41(3):409–414, 2007.
 34. N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
 35. A. de Palma, M. Ben-Akiva, C. Lefèvre, and N. Litinas. Stochastic equilibrium model of peak period traffic congestion. *Transportation Science*, 17(4):430–453, 1983.
 36. 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.
 37. 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, 1995.
 38. B. Ran and D. E. Boyce. *Modelling Dynamic Transportation Networks*. Springer, 1996.
 39. 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.
 40. B. Ran, R. W. Hall, and D. E. Boyce. A link-based variational inequality model for dynamic departure time/route choice. *Transportation Research B*, 30(1):31–46, 1996.

41. T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.
42. 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, 2000.
43. M. Skutella. An introduction to network flows over time. In W. Cook, L. Lovász, and J. Vygen, editors, *Research Trends in Combinatorial Optimization*, chapter 21, pages 451–482. Springer, 2009.
44. M. J. Smith. The existence of a time-dependent equilibrium distribution of arrivals at a single bottleneck. *Transportation Science*, 18(4):385–394, 1984.
45. 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 B*, 27(1):49–63, 1993.
46. W. Y. Szeto and Hong K. Lo. A cell-based simultaneous route and departure time choice model with elastic demand. *Transportation Research B*, 38(7):593–612, 2004.
47. N. B. Taylor. The CONTRAM dynamic traffic assignment model. *Networks and Spatial Economics*, 3(3):297–322, 2003.
48. W. S. Vickrey. Congestion theory and transport investment. *The American Economic Review*, 59(2):251–260, 1969.
49. J. H. Wu, Y. Chen, and M. Florian. The continuous dynamic network loading problem: a mathematical formulation and solution method. *Transportation Research Part B: Methodological*, 32(3):173–187, 1998.
50. Y. W. Xu, J. H. Wu, M. Florian, P. Marcotte, and D. L. Zhu. Advances in the continuous dynamic network loading problem. *Transportation Science*, 33(4):341–353, 1999.
51. S. Yagar. Dynamic traffic assignment by individual path minimization and queuing. *Transportation Research*, 5(3):179–196, 1971.