Nash Equilibria and the Price of Anarchy for Flows Over Time

Ronald Koch and Martin Skutella

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

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 change over time. 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 a variant of 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.

1 Introduction

In a groundbreaking paper, Roughgarden and Tardos [35] (see also Roughgarden's book [34]) 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 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; see, e.g., [5, 30]. 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 [15, 16].

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 [5, 11, 19, 24, 30, 31]. 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 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 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". Similar models are used, for example, in road traffic simulation and related fields.

Related Literature. As already mentioned above, 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, 36].

So far, Nash equilibria for flows over time were mostly studied within the traffic community. Vickrey [41] and Yagar [42] 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., [41, 21, 13, 23]. Since then, the number of publications in this area has increased rapidly and Nash equilibria where modeled mathematically. For a survey see, e.g., [29]. The considered models can be grouped into for categories: mathematical programming (e.g., [22, 20]), optimal control (e.g., [32, 18]), variational inequalities (e.g., [17, 12, 33, 38, 39]), and simulation-based approaches (e.g., [42, 25, 7, 40, 6]). 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 computationally often intractable. Edge-based formulations are, for example, considered in [2, 12, 33]. Realistic assumptions on the underlying flow model with respect to traffic simulation are described by Carey [9, 10].

In this paper the deterministic queuing model is considered. This model was introduced by Vickrey [41] and later by Hendrickson and Kocur [21]. Smith [37] 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 [26, 27] 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 a discrete model 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 allowed 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.

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 [28]), 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 sink.

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, the price of anarchy is, in general, unbounded for arbitrary networks. The latter observation marks a clear distinction between static routing games and routing games over time. We conclude with an outlook on further interesting results and challenges.

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. First, in Section 2.1 we define a routing game over time showing 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 mentioned above.

Throughout the paper we often use the term *flow particle* in order to refer to an infinitesimally small 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 node $s \in V$ and a sink node $t \in V$. Each flow particle is a player and the strategy set of players is the set $\mathcal{P}$ of all s-t-paths.

In a static routing game, the players' decisions yield a static s-t-flow $\mu = (\mu_P)_{P \in \mathcal{P}}$ 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 $(\mu_P)_{P \in \mathcal{P}}$ 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 $\ell_P(\mu)(\theta)$; if μ is clear from the context, we write $\ell_P(\theta)$ for short.

In this paper we restrict to payoff functions where $\ell_P(\theta)$, $P \in \mathcal{P}$, is the time when a flow originating at s at time θ arrives at 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. But unlike the static case, this is not as easy to model since we have to define at what time a path is being used.

Definition 1 (Nash flows 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 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 where we assumed that the payoff functions are continuous. A closer look at Definition 1 shows us that the continuity of the payoff functions ℓ_P is also essential here. We skip further technical details du to space restrictions.

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

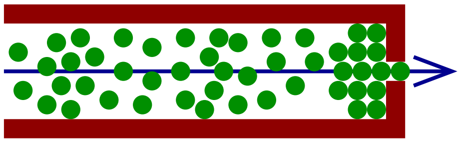

Fig. 1. If more flow particles want to leave an edge than its capacity allows, they form a waiting queue.

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 $\mu_e := \sum_{P \in \mathcal{P}: e \in E(P)} \mu_P$ of that edge which can easily be computed from $(\mu_P)_{P \in \mathcal{P}}$.

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 $(\mu_P)_{P \in \mathcal{P}}$. Notice that the time at which a flow particle that enters path $P \in \mathcal{P}$ at time θ 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. 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 node $s \in V$, and a sink node $t \in V$. We assume without loss of generality that there are no incoming edges at the source node s and no outgoing edges at the sink node t. The capacity u_e of edge e bounds the rate at which flow may leave edge e at its head node. The basic concept of our flow over time model are waiting queues which built 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 exist, which are wide enough for storing all waiting flow particles; 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 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 an 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 outflow rate f_e^- describes the rate of flow $f_e^-(\theta) \geq 0$ leaving the head of e at time θ. Moreover, 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 all $e \in E$.

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:

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

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:

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

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:

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

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), $F_e^+(\theta)$ 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, $F_e^-(\theta + \tau_e)$ 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 the capacity u_e. Therefore the waiting time spent by a flow particle entering the tail of e at time θ is equal to

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

The interpretation of $q_e(\theta)$ as the waiting time for flow particles arriving at time θ on arc 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.

We state the following proposition which follows directly from (4) and the continuity of F_e^+ and F_e^-.

Proposition 2. *For any edge $e \in E$, the function $\theta \mapsto \theta + q_e(\theta)$ is monotonically increasing and continuous.*

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 on that path.

We consider two 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 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 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,

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

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 in (5). That is, the flow particle can use edge e in order to arrive at node w as early as possible (at time $\ell_w(\theta)$). Moreover, we have $\ell_s(\theta) = \theta$ for all θ. Therefore, we define the *label functions* $\ell_w : \mathbb{R}_+ \to \mathbb{R}_+$ as follows:

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

The label functions can be computed simultaneously for each time θ by adapting the shortest path algorithm of Bellman and Ford. The following proposition follows from (6) and Proposition 2.

Proposition 3. *For each node $v \in V$, the label 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 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 4. *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 $\theta \geq 0$:*

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

We emphasize the following aspect of Definition 4: In general, it is not clear that the label functions are *strictly* monotonically increasing. In particular, the label function of the sink t might possibly be constant over a certain time interval $[\theta_1, \theta_2]$ with $\theta_1 < \theta_2$. 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 4 enforces, however, that all subpaths of the s-t-path chosen by a flow particle have to be as short as possible.

The condition in Definition 4 is equivalent to the condition that every flow particle tries 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". That is, it is equivalent to the statement that no flow particle can possibly overtake any other flow particle.

In order to model the non-overtaking 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:

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

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

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

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

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 6. *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) It is a Nash flow over time.

Before we prove Theorem 6, we state the following lemma, which gives a more global characterization of when flow is being sent only along currently shortest paths (Definition 4 gives only a pointwise characterization).

Lemma 7. *For a given flow over time, the following statements are equivalent:*

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

$$F_e^-\big(\ell_v(\theta) + \tau_e + q_e(\ell_v(\theta))\big) \;=\; F_e^-\big(\ell_w(\theta)\big) \;. \tag{9}$$

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)⇒(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

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

By definition of θ_1, $\ell_w(\theta') \leq \ell_w(\theta) < \ell_v(\theta') + \tau_e + q_e(\ell_v(\theta'))$, 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 4 we get

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

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)⇒(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 2 and 3, 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 yield:

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

This yields statement (i). □

Due to space limitations, the proof of Theorem 6 has been moved to Appendix A.1.

Note that whenever one of the four statements in Theorem 6 holds, then x^+ and x^- coincide. Further, for all $\theta \geq 0$, setting $x_e(\theta) := x_e^+(\theta)$ for all $e \in E$, yields a static s-t-flow $x(\theta)$ of value $b_s(\theta)$. In the following, for a flow over time satisfying the non-overtaking 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.

Definition 8 (Current Shortest Paths Network). *Consider a flow over time on a network (G, u, s, t, τ, d). Let $(\ell_v)_{v\in V}$ be the corresponding family of label functions and $(q_e)_{e\in E}$ the family of waiting time functions. For $\theta \geq 0$, the* current shortest paths network G_θ *is the subnetwork induced by the edges occurring in a currently shortest path.*

By Propositions 2 and 3, there exists an $\epsilon > 0$ such that $G_{\theta'}$ is a subgraph of G_θ for all $\theta' \in [\theta, \theta + \epsilon)$.

Definition 9 (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 is a* thin flow with resetting on E_1 *if there exist node labels ℓ' such that:*

$$\ell'_s = F/d \tag{11}$$

$$\ell'_w \leq \ell'_v \qquad \text{for all } e = vw \in E(G) \setminus E_1 \text{ with } x'_e = 0 \tag{12}$$

$$\ell'_w = \max\{\ell'_v, x'_e/u_e\} \qquad \text{for all } e = vw \in E(G) \setminus E_1 \text{ with } x'_e > 0 \tag{13}$$

$$\ell'_w = x'_e/u_e \qquad \text{for all } e = vw \in E_1 \tag{14}$$

Notice that, if $E_1 = \emptyset$, 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. It is not difficult to see that, for the special case $E_1 = \emptyset$, a thin flow with resetting can be computed in polynomial time; see Appendix A.3 for details.

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. But both the label functions and the underlying static flow functions are monotonically increasing implying that both families of functions are differentiable almost everywhere. In the following, all derivatives are defined as derivatives from the right.

Theorem 10. *Consider a Nash flow over time on a network (G, u, s, t, τ, d) with corresponding label functions $(\ell_v)_{v\in V}$ and waiting time functions $(q_e)_{e\in E}$. 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 $v \in V$. Then, on the current shortest paths network $G_{\theta'}$, 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 $E_1 = \{e \in E \mid q_e(\theta') > 0\}$. *A corresponding set of node labels fulfilling (11) to (14) is given by the derivatives $(\frac{d\ell_v}{d\theta}(\theta))_{v\in V(G_\theta)}$.*

The reverse direction of Theorem 10 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 paths network for almost all times θ, then the flow over time is in fact a Nash flow over time. We skip further details due to space restrictions.

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 free flow travel times of all edges are zero. We study the price of anarchy which, in general, is the worst case ratio of the cost of a Nash equilibrium to 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:[1] For given θ, maximize the amount of flow arriving at the sink until time θ. In particular, according to this definition, earliest arrival flows that maximize the amount of flow at the sink simultaneously for each point in time are the system optima.

Theorem 11. *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.*

Proof (sketch). For a Nash flow over time, it is not difficult to see that the behavior of the underlying static flow and the corresponding label functions can be described via essentially the same thin flow with resetting for all points in time. The key observation is that the current shortest paths network remains unchanged throughout the whole algorithm and is thus equal to the given network for all times θ. Since at time zero there are no waiting edges, the underlying thin flow is a thin flow *without resetting*, i.e., $E_1 = \emptyset$. It is shown in Appendix A.3 that thin flows without resetting of maximum flow value are in some sense unique and can be computed in polynomial time. As an immediate consequence of this, a Nash flow over time is an earliest arrival flow. □

In contrast to static routing games, there exist instances of the routing game over time where the price of anarchy is unbounded; see Appendix A.4.

Proposition 12. *There exists a family of instances for which the price of anarchy is $\Omega(m)$ where m is the number of edges.*

We conclude with a short outlook to further results which, due to space constraints, are beyond the scope of this paper.

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

– For general constant free flow transit times, we can show, that the underlying “sequence” of thin flows with resetting is piecewise constant. Moreover, the solutions to the thin flow instances occurring in this setting are unique with respect to node labels. Thus, a Nash flow over time can be seen as a concatenation of static flows.
– There is an algorithm for computing thin flows with resetting which we conjecture is polynomial. The algorithm uses a fixed point approach and iteratively calls an algorithm for the special case $E_1 = \emptyset$ as a subroutine.
– Throughout this paper we have always assumed that capacities of edges are constant. The models and results presented in this paper (with the exception of Theorem 11) can be generalized to the case of time-varying edge capacities. More generally, most results and approaches not only hold for the deterministic queuing model but for a considerably more general class of flow over time models.

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 fo dynamic traffic equilibrium assignmen 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. Elliot Anshelevich and Satish Ukkusuri. Equilibria in dynamic selfish routing. Presented on the International Symposium on Dynamic Traffic Assignment 2008, 2008.
5. J. E. Aronson. A survey of dynamic network flows. *Annals of Operations Research*, 20:1–66, 1989.
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. 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.
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, Berlin, 1999.
12. 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.
13. 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.

14. L. K. Fleischer and É. Tardos. Efficient continuous-time dynamic network flow algorithms. *Operations Research Letters*, 23:71–80, 1998.
15. L. R. Ford and D. R. Fulkerson. Constructing maximal dynamic flows from static flows. *Operations Research*, 6:419–433, 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, Berlin, 2002.
20. 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.
21. C. Hendrickson and G. Kocur. Schedule delay and departure time decisions in a deterministic model. *Transportation Science*, 15(1):62–77, 1981.
22. 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.
23. H. Mahmassani and R. Herman. Dynamic user equilibrium departure time and route choice on idealized traffic arterials. *Transportation Science*, 18(4):362–384, 1984.
24. 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.
25. H. S. Mahmassani, H. A. Sbayti, and X. Zhou. *DYNASMART-P Version 1.0 Users Guide*. Maryland Transportation Initiative, College Park, Maryland, 2004.
26. R. Mounce. Convergence in a continuous dynamic queueing model for traffic networks. *Transportation Research Part B: Methodological*, 40(9):779–791, 2006.
27. R. Mounce. Convergence to equilibrium in dynamic traffic networks when route cost is decay monotone. *Transportation Science*, 41(3):409–414, 2007.
28. N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
29. 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.
30. 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.
31. B. Ran and D. E. Boyce. *Modelling Dynamic Transportation Networks*. Springer, Berlin, 1996.
32. 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.

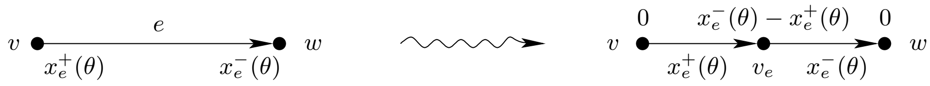

Fig. 2. Construction of the b-flow instance used in the proof of Theorem 6. Below the edges the in- and outflow of the dynamic Nash equilibrium (left) and the flow value of the b-flow (right) are shown. Above nodes the corresponding b-values are displayed.

33. 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.
34. T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.
35. 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.
36. M. Skutella. An introduction to network flows over time. In W. Cook, L. Lovász, and J. Vygen, editors, *Research Trends in Combinatorial Optimization*, pages 451–482. Springer, Berlin, 2009.
37. M. J. Smith. The existence of a time-dependent equilibrium distribution of arrivals at a single bottleneck. *Transportation Science*, 18(4):385–394, 1984.
38. 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.
39. 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.
40. N. B. Taylor. The contram dynamic traffic assignment model. *Networks and Spatial Economics*, 3(3):297–322, 2003.
41. W. S. Vickrey. Congestion theory and transport investment. *The American Economic Review*, 59(2):251–260, 1969.
42. S. Yagar. Dynamic traffic assignment by individual path minimization and queuing. *Transportation Research*, 5(3):179–196, 1971.

A Appendix

A.1 Proof of Theorem 6

Proof (of Theorem 6). The main observation we need is the following equation which we get from the definitions of x_e^+, x_e^-, and q_e in (7) and (4), respectively.

$$\begin{aligned} x_e^+(\theta) - x_e^-(\theta) &= F_e^+\big(\ell_v(\theta)\big) - F_e^-\big(\ell_w(\theta)\big) \\ &= F_e^-\big(\ell_v(\theta) + \tau_e + q_e(\ell_v(\theta))\big) - F_e^-\big(\ell_w(\theta)\big) \ . \end{aligned} \tag{15}$$

Because of Lemma 7, 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_ew; see Fig. 2. 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, $e \in E$, 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 (8). It follows from (15) 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).

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

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

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

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 proof 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. Thus, a flow over time which sends flow only over currently shortest paths is a Nash flow over time.

(iv)$\Rightarrow$(iii): 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 2), 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. □

A.2 Proof of Theorem 10

In order to prove Theorem 10 we need the following lemma.

Lemma 13. *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 $\ell_v(\theta) + \tau_e + q_e(\ell_v(\theta)) = \ell_w(\theta_1)$ since the waiting time is monotonically decreasing if no flow is assigned to e. Further, we know that the label functions are increasing which completes the proof because of (6). ☐

Proof (of Theorem 10). 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 (11) to (14) 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 $G_{\theta'}$ and $q_e(\theta'') > 0$ for all edges $e \in E_1$.

Condition (11) 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 the conditions (12) to (14) are satisfied in every case. 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: 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. 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:

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

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 (14) is satisfied in this case. Further, condition (12) is also satisfied because the label functions are monotonically increasing. In order to show that condition (13) 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 to Case 1 because of Lemma 13. Further, we know that

$q_e(\ell_v(\theta'')) = 0$ for all $\theta'' \in [\theta', \theta'+\epsilon]$. Therefore, it holds that $\ell_v(\theta'+\epsilon) - \ell_v(\theta') > \ell_w(\theta'+\epsilon) - \ell_w(\theta')$. 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

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

Thus, condition (12) 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 13 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 zero, we get

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

Therefore, condition (12) is satisfied. Because condition (14) does not belong to this case, we only show that condition (13) is valid. For this, we show that $\frac{dx_e}{d\theta}(\theta') \cdot \frac{1}{u_e} < \frac{d\ell_w}{d\theta}(\theta')$. 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 for the last time both sides by ϵ and let ϵ tend to zero, we get the desired result. This completes the proof. □

A.3 Thin flows without resetting

The following is an equivalent definition for thin flows without resetting, i.e., thin flows with resetting on $E_1 = \emptyset$.

Definition 14. *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 $\ell'_s := \frac{F}{d}$ where F 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 node s_0 and an edge $s_0 s$ 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. 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 15. *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 $\ell'_{\max}$ 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 $X := \{v \in V \mid \ell'_v < \ell'_{\max}\}$. 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 $\ell_{\max}$. 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. □

The last lemma shows that a thin s-t-flow of value equal to the maximum 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 16. *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.*

Proof. Consider two thin flows $x', \widetilde{x} \in \mathbb{R}_+^{E(G)}$. We prove by induction on the number of nodes that the corresponding edge labels ℓ', $\widetilde{\ell}$ 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 15 shows that the maximum edge label $\ell'_{\max}$ is unique and equal to the congestion of a sparsest cut. Therefore, the sets of edges with this maximal label coincide since x', $\widetilde{x}$ are static flows minimizing the maximal edge congestion. Thus, we know the flow values on edges contained in $\delta^+(X)$, where X defines a sparsest cut, because the labels of such edges must be defined by their congestion, i.e., $x'_e = \widetilde{x}_e = \ell'_{\max} \cdot u_e$.

Now we delete the node set $V \setminus X$. Then, x' and $\widetilde{x} \in \mathbb{R}^{E(G)}$ are thin b-flows on the induced subgraph $G[X]$ according to the new node balances

$$b'(v) \;:=\; b(v) - x'\big(\delta^+_{G'}(v) \cap \delta^+(X)\big) \qquad \text{for all } v \in X.$$

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

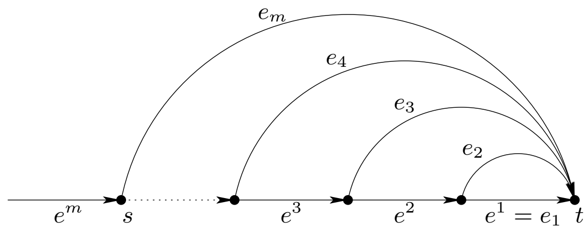

Fig. 3. A family of instances with unbounded price of anarchy

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. □

A.4 A family of instances with unbounded price of anarchy

The following family of instances shows that the price of anarchy can be arbitrarily large, i.e., the price of anarchy is $\Omega(m)$, where m is the number of edges. The graph of the underlying instances is shown in Fig. 3.

The edge capacities are defined as follows. Let $u_k := u_{e_k}$, $u^k := u_{e^k}$, and assume that $u^k = \sum_{i=1}^{k} u_i$, for $k = 1, \ldots, m$. The s-t-path P_k is defined by the edge sequence $e^m, e^{m-1}, \ldots, e^k, e_k$, for $k = 1, \ldots, m$. Further let $\tau_k := \sum_{e \in E(P_k)} \tau_e$ be the free flow transit time of path P_k.

In a Nash flow over time, 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 $k = 2, \ldots, m$). Since $u^m > u^{m-1} > \cdots > u^1$, a linearly increasing waiting queue builds up on each edge of P_1. Therefore, the time for traversing path P_k, $k = 1, \ldots, m$, increases monotonically with the time when a flow particle originates at s. Moreover, 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. In particular, the transit time of P_1 has the greatest slope. Thus, at a certain point in time α, the time for traversing P_1 becomes equal to the time for traversing some other path $P \in \{P_k \mid k \in \{2, \ldots, m\}\}$ for flow particles originating at s at time α. This means that flow particles start to choose P_1 and P in a Nash flow over time.

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 $P_2, \ldots, P_k$. (Therefore, the free flow transit times τ_k of paths P_k must increase with k.) This means that, from time α on, flow particles use the entire network in order to reach t. Summarizing, the free flow transit times ensure that, in a Nash flow over time, the first flow particles use only P_1 and suddenly from time α on the whole network is used.

We model this expected behavior precisely as follows. For $k = 1, \ldots, m$, let $\ell_t^k(\alpha)$ be the arrival time at t using P_k for flow originated at s at time α under the assumption that, up to this time, all flow units use only P_1 in order to reach t. Thus,

$$\ell_t^k(\alpha) \;=\; \tau_k + \frac{u^m}{u^k}\cdot\alpha \qquad\Longleftrightarrow\qquad u^m\alpha \;=\; (\ell_t^k(\alpha)-\tau_k)u^k \;\;.$$

Because for flow originating at s at time α, all s-t-paths must have the same latency, we get:

$$\begin{aligned} & \tau_1 + \frac{u^m}{u^1}\cdot\alpha \;=\; \tau_k + \frac{u^m}{u^k}\cdot\alpha \;=:\; \ell_t(\alpha) && \text{for } k=2,\ldots,m \\ \Longleftrightarrow\qquad & \tau_1 \;=\; \tau_k + u^m\alpha\left(\frac{1}{u^k}-\frac{1}{u^1}\right) && \text{for } k=2,\ldots,m. \end{aligned}$$

For edge capacities satisfying the conditions of this example, there exist edge free flow transit times ensuring the last equations. Simply set $\tau_{e^k} := 0$, for $k=1,\ldots,m$. 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 $F_{\mathrm{NE}}(\theta)$ and $F_{\mathrm{SO}}(\theta)$ be the amount of flow arriving at t before time θ in a Nash flow over time and in a system optimum, respectively. It is not difficult 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, path P_k contributes an inflow rate of u_k to t from time τ_k on. This flow over time is a so-called *earliest arrival flow*. Because a Nash flow over time has to satisfy the non-overtaking condition, we know that $F_{\mathrm{NE}}(\ell_t(\alpha)) = u^m\alpha$. The corresponding value for the earliest arrival flow is $F_{\mathrm{SO}}(\ell_t(\alpha)) = \sum_{k=1}^m(\ell_t(\alpha)-\tau_k)\cdot u_k$. Thus, a lower bound on the price of anarchy is given by

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

(In fact, this is the exact price of anarchy for the instance under consideration.) This shows that the price of anarchy can increase 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 can still increase 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 quadratical in m).