





The Impact of Spillback on the Price of Anarchy for Flows over Time

Jonas Israel  and Leon Sering ^(✉) 

Technische Universität Berlin, Berlin, Germany
j.israel@tu-berlin.de, sering@math.tu-berlin.de

Abstract. Flows over time enable a mathematical modeling of traffic that changes as time progresses. In order to evaluate these dynamic flows from a game theoretical perspective we consider the *price of anarchy* (PoA). In this paper we study the impact of spillback effects on the PoA, which turn out to be substantial. It is known that, in general, the PoA is unbounded in the spillback setting. We extend this by showing that it is still unbounded even when considering networks with unit edge capacities and that the *Braess ratio* can be arbitrarily large.

In contrast to that, we show that on a fixed network the PoA as a function of the flow amount is bounded by a constant and also upper bound the PoA for the set of networks where the outflow capacities satisfy certain constraints depending on the quickest flow. This upper bound only depends on the worst spillback factor of the Nash flows over time of the given network. It therefore provides a way to quantify the impact of spillback to the quality of the dynamic equilibria.

In addition, we show the surprising fact that the introduction of spillback behavior can actually speed up dynamic equilibria in some networks.

Keywords: Nash flow over time · Dynamic equilibria · Deterministic queuing · Price of anarchy · Spillback · Traffic

Related Version: A full version of this paper including all proofs is available at <https://arxiv.org/abs/2007.04218>.

1 Introduction

Road traffic is an integral part of modern societies, which consists of many users with individual behaviors and goals. For this reason traffic dynamics are very hard to predict and can barely be controlled. However, through recent technologies such as intelligent navigation systems it might be possible to positively

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Grant BR 4744/2-1 and Germany's Excellence Strategy – The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project ID: 390685689).

© Springer Nature Switzerland AG 2020

T. Harks and M. Klimm (Eds.): SAGT 2020, LNCS 12283, pp. 114–129, 2020.

https://doi.org/10.1007/978-3-030-57980-7_8

affect the behavior of traffic, steering it towards shorter travel times leading to less pollution and an overall improved quality of life.

In the following research work we focus on a mathematical traffic flow model called *flows over time with spillback*. Here, the network is depicted as a graph with a source and a sink, and the traffic flow can progress in continuous time from one vertex over an edge to the next vertex. We consider flow as a continuous stream affected by two types of temporal factors. First, flow does not travel instantaneously through the network but needs actual time to traverse an edge, and second, flow on an edge may change over time. Compared to static network flows these temporal components enable us to model traffic realistically through different congestion levels. To model road constraints within the network, we equip each edge with an inflow and an outflow capacity governing with which rate flow can enter and leave the edge, a length characterizing the time it takes a flow-particle to travel from the tail to the head of the edge, and finally, a storage capacity which describes how much flow volume fits on the edge. If the desired outflow exceeds the outflow capacity of an edge the excess flow queues up in front of the bottle-neck at the head of the edge. If at any point in time the queue of an edge is so large that the amount of flow traversing the edge plus the amount of flow in the queue equals the storage capacity, the edge is considered full and new flow can only enter if at least as much flow leaves at the same time. With this mechanic it is possible to model *spillback*, i.e., the phenomenon that traffic congestion at one street can block exits or intersections further upstream. The ability to model spillback within the framework of flows over time is a very recent discovery [22], which has not been studied much yet.

As we experience in our everyday lives traffic is not performing optimal most of the time, but rather consists of agents that behave egoistically. Thus, we are interested in game theoretic aspects of this flow model, particularly in the *price of anarchy* (PoA), the ratio of the worst uncoordinated behavior described via a dynamic equilibrium, and the optimal flow behavior measured by some social cost function. In real-world scenarios that ratio could give us an idea of how much one can possibly improve traffic through optimized traffic control, for example through modern navigation systems or autonomous driving. Even though it has been shown in [22] that the PoA in networks with spillback is unbounded in general we investigate the dependency of the PoA on several parameters, for example, the minimal spillback factor, which measures how much the capacities of an edge are reduced due to spillback. Another interesting phenomenon of selfish road users we study is the well known Braess paradox [2]. It states that the overall travel time of all users might decrease if a frequently used road segment gets closed. In reverse, this means that building new roads between heavily used section of the network might cause more congestion and longer travel times.

Related Work. Flows over time were first introduced by Ford and Fulkerson [8] in the context of an optimization problem to route as much flow as possible in a given time horizon. Gale [9] proved the existence of earliest arrival flows which optimize the amount of flow routed to the sink simultaneously for

all points in time and Wilkinson [26] later presented an algorithm to compute these flows. For an overview on flows over time from an optimization point of view we refer to the survey by Skutella [23]. From a game theoretic point of view, flows over time were first considered by Vickrey [24] in the setting of transportation research. In the last years the theory of Nash equilibria for flow models has been advanced significantly. From the introduction of the price of anarchy by Koutsoupias and Papadimitriou [13, 16] and the congestion games studied by Roughgarden and Tardos [18, 19] (both for static flows), over existence results concerning the dynamic (i.e., time dependent) model by Meunier and Wagner [15], to the constructive approach to dynamic equilibria by Koch and Skutella [12]. Here, the authors present a novel notion of dynamic equilibria, called *Nash flows over time*, which enabled a whole set of proceeding research. This new research includes the study of existence, uniqueness and the long-term behavior of Nash flows over time by Cominetti et al. [4–6], the work by Macko et al. [14] about the Braess paradox for flows over time as well as the extension to multi-terminal settings [21]. Of special interest to the paper at hand are the results by Bhaskar et al. [1] and very recently by Correa et al. [7] about the PoA for flows over time. Since it was already shown that the evacuation-PoA (maximizing the flow amount within some time horizon) is unbounded [12], they focus on the time-PoA (minimizing the completion time for a given flow amount) for which they establish an upper bound of $\frac{e}{e-1}$ under some constraints on the capacities of the network. Sering and Vargas Koch [22] generalized the flows over time model in order to represent spillback and transferred the results about dynamic equilibria to this extension. Very recently, Graf et al. [10] characterized an alternative equilibrium concept for flows over time, where particles do not predict the future evolution of the flow but instead reconsider their route choice on every node. In addition, there is a active research line on packet routing models, where traffic is represented by atomic vehicles that traverses the network in discrete time steps. Recent progress in this area is due to Cao et al. [3], Scarsini et al. [20], Harks et al. [11] and Peis et al. [17].

Contribution and Outline. We study the price of anarchy of flows over time with spillback introduced in [22], which is known to be unbounded in general. After introducing the model in Sect. 2, we show in Sect. 3 that the PoA stays unbounded even if we restrict the set of networks to a specific topology but allow arbitrary capacity, or in reverse if we only allow unit capacities but therefore more complex graph structures. Furthermore, we show that the Braess ratio can be arbitrarily large depending only on the minimum edge capacity. Even though it seems that the addition of full edges and spillback only increases completion times this is not a general rule, as we show that there are examples where the completion time of Nash flows over time is larger when disabling spillback. In contrast to the above lower bounds we show in Sect. 4 that if we consider the case of temporal routing games on a fixed network, i.e., only the flow amount that gets routed through the network varies, the PoA is bounded by a constant. In the end we translate the ideas of [1] to the model with spillback and prove

an upper bound of $\frac{ce}{ce-1}$ on the PoA in networks with specific conditions on the capacities in dependency of a maximal flow over time. This upper bound only depends on the worst spillback factor c of the Nash flows over time of the given network, and therefore provides a way to quantify the impact of spillback to the PoA (note that e denotes the Euler constant here). Finally, we give a brief conclusion and outlook for further research in Sect. 5.

2 The Model

In the following we want to recall the essential definitions of the flow over time model with deterministic queuing. We consider the extended version that handles spillback effects, as introduced in [22] and mainly stick to the same notation.

Flow Dynamics. We consider a *network* $\Gamma = (G, s, t, r_0, \tau, \nu^+, \nu^-, \sigma)$ given by a directed graph $G = (V, E)$ with a single *source* s and a single *sink* t , such that every vertex is reachable from s . We have a *network inflow rate* of $r_0 > 0$ determining the constant rate of flow entering the network from time 0 onward. Furthermore, every edge $e \in E$ is equipped with a *transit time* $\tau_e \geq 0$, an *in- and outflow capacity* $\nu_e^+ > 0$ and $\nu_e^- > 0$ as well as a *storage capacity* $\sigma_e > 0$. In order to avoid undefined flow behavior, we require that traversing flow alone can never fill up an edge, i.e., $\sigma_e > \nu_e^+ \cdot \tau_e$ and that the total transit time of every directed cycle is strictly positive. For technical reason we furthermore assume that all properties are rational numbers.

A *flow over time* is given by a family of locally integrable and bounded functions $f = (f_e^+, f_e^-)_{e \in E}$, where $f_e^+, f_e^- : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ denote the in- and outflow rate of edge e at every point in time. The *cumulative in- and outflow* and the *queue size* are given by

$$F_e^+(\theta) := \int_0^\theta f_e^+(\xi) \, d\xi, \quad F_e^-(\theta) := \int_0^\theta f_e^-(\xi) \, d\xi \quad \text{and} \quad z_e(\theta) := F_e^+(\theta - \tau_e) - F_e^-(\theta).$$

We require that flow is preserved at every edge e (*non-deficit constraint*) and at every vertex $v \in V \setminus \{t\}$ (*conservation constraint*), which means, for every point in time θ we have

$$z_e(\theta) \geq 0 \quad \text{and} \quad \sum_{e \in \delta_v^+} f_e^+(\theta) - \sum_{e \in \delta_v^-} f_e^-(\theta) = \begin{cases} 0 & \text{for } v \in V \setminus \{s, t\}, \\ r_0 & \text{for } v = s. \end{cases}$$

Here, δ_v^- is the set of all *incoming* and δ_v^+ the set of all *outgoing* edges of node v . An edge e is *full* at time θ if the total amount of flow on e , called *edge load*, $d_e(\theta) := F_e^+(\theta) - F_e^-(\theta)$ reaches the storage capacity σ_e . The *inflow bound* $b_e^+(\theta)$ denotes that current inflow capacity, which might be smaller than ν_e^+ due to spillback, and the *push rate* $b_e^-(\theta)$ specify the current *desired* outflow rate, which is reached whenever there are no restrictions of following links. Formally, we have

$$b_e^+(\theta) := \begin{cases} \nu_e^+ & \text{if } d_e(\theta) < \sigma_e \\ \min\{f_e^-(\theta), \nu_e^+\} & \text{else,} \end{cases} \quad \text{and} \quad b_e^-(\theta) := \begin{cases} \nu_e^- & \text{if } z_e(\theta) > 0, \\ \min\{f_e^+(\theta - \tau_e), \nu_e^-\} & \text{else.} \end{cases}$$

A flow over time f is *feasible* if for all edges e and all times θ it satisfies $f_e^+(\theta) \leq b_e^+(\theta)$ (*inflow condition*) and if there exists a $c_v \in (0, 1]$ for every $v \in V$ such that for every $e \in \delta^-(v)$ $f_e^-(\theta) = \min\{b_e^-(\theta), \nu_e^- \cdot c_v\}$ (*fair allocation condition*). Furthermore, we require for all time θ that every vertex v with an incoming edge $e_1 \in \delta_v^-$ with $f_{e_1}^-(\theta) < b_{e_1}^-(\theta)$ (called *throttled edge*) there exists an outgoing edge $e_2 \in \delta_v^+$ with $f_{e_2}^+(\theta) = b_{e_2}^+(\theta)$ (*no-slack condition*). Finally, the set of full edges should be cycle free at every point in time (*no-deadlock condition*).

For a given $v \in V$ the maximal value c that satisfies the fair allocation condition at a given point in time θ is called *spillback factor* denoted by $c_v(\theta)$. This value denotes the reduction of the outflow capacity due to spillback leading to the *effective outflow capacity* of $\nu_e^- \cdot c_v(\theta)$. If the outflow rate $f_{uv}^-(\theta)$ of an incoming edge is strictly smaller than the push rate $b_{uv}^-(\theta)$, this edge is throttled implying $c_v(\theta) < 1$, which means that there is spillback at v . In this case the no-slack condition ensures that there is a reason for the spillback in form of an outgoing exhausted edge vw : $f_{vw}^+(\theta) = b_{vw}^+(\theta)$. The spillback factor will play an important role throughout this paper. For more details and further intuition on the definitions of a feasible flow over time in this setting we refer to [22].

Nash Flows Over Time. In order to define Nash flows over time we need to define the arrival time of every particle of the flow. To simplify the notation we identify every particle with the point in time θ when it enters the network at the source. For every edge e we define the *waiting time function* $q_e: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ by $q_e(\theta) := \min \left\{ q \geq 0 \mid \int_{\theta+\tau_e}^{\theta+\tau_e+q} f_e^-(\xi) d\xi = z_e(\theta + \tau_e) \right\}$, i.e., $q_e(\theta)$ denotes the time a particle entering e at time θ waits in the queue. For every vertex v the *earliest arrival time function* $\ell_v: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ denotes the earliest point in time the particle θ (which enters the network at time θ) can reach v :

$$\ell_v(\theta) := \begin{cases} \theta & \text{if } v = s, \\ \min_{e=uv \in E} \ell_u(\theta) + \tau_e + q_e(\ell_u(\theta)) & \text{else.} \end{cases}$$

For a given particle θ the *current shortest path network* $G'_\theta = (V, E'_\theta)$ is the network of all edges $e = uv$ that are *active for* θ , i.e., for which $\ell_v(\theta) = \ell_u(\theta) + \tau_e + q_e(\ell_u(\theta))$. It contains all s - v -paths that particle θ can use to be at v at the earliest possible point in time. Furthermore, we denote the *resetting edges* E_θ^* as the set of edges for which particle θ encounters a queue when taking a current shortest path and \bar{E}_θ denotes the set of edges which are full when particle θ reaches its tail. More precisely, $E_\theta^* := \{e = uv \in E \mid q_e(\ell_u(\theta)) > 0\}$ and $\bar{E}_\theta := \{e = uv \in E \mid d_e(\ell_u(\theta)) = \sigma_e\}$.

Definition 1 (Nash flow over time). We call a feasible flow over time f a Nash flow over time, or dynamic equilibrium, if almost every particle uses a current shortest s - t -path, i.e., if $f_e^+(\theta) > 0$ implies $\theta \in \ell_u(\Theta_e)$ for all $e = uv \in E$

and almost all $\theta \in \mathbb{R}_{\geq 0}$, where $\Theta_e := \{\theta \in \mathbb{R}_{\geq 0} \mid e \in E'_\theta\}$ denotes all particles for which e is active.

Equivalently, it has been shown [22, Lemma 4.1] that a feasible flow over time is a dynamic equilibrium if and only if $F_e^+(\ell_u(\theta)) = F_e^-(\ell_v(\theta))$ for all $e = uv \in E$ and all $\theta \in \mathbb{R}_{\geq 0}$. By setting $x_e(\theta) := F_e^+(\ell_u(\theta)) = F_e^-(\ell_v(\theta))$ we observe that $(x_e(\theta))_{e \in E}$ form a static s - t -flow of value $r_0 \cdot \theta$. Since the x_e are absolute continuous, their derivatives

$$x'_e(\theta) = f_e^+(\ell_u(\theta)) \cdot \ell'_u(\theta) = f_e^-(\ell_v(\theta)) \cdot \ell'_v(\theta) \tag{1}$$

exist almost everywhere and can be seen as the strategy of particle θ (as for every θ it is a static s - t -flow of value r_0). For a fixed θ the derivatives $x'_e := x'_e(\theta)$ and $\ell'_v := \ell'_v(\theta)$ together with the spillback factors $c_v := c_v(\ell_v(\theta))$ are called *spillback thin flows* and with $b_e^+ := b_e^+(\ell_u(\theta))$ for all $e = uv$ satisfy the following equations:

$$\begin{aligned} \ell'_s &= 1, \\ \ell'_v &= \min_{e=uv \in E'_\theta} \rho_e(\ell'_u, x'_e, c_v) \quad \text{for } v \in V \setminus \{s\}, \\ \ell'_v &= \rho_e(\ell'_u, x'_e, c_v) \quad \text{for } e = uv \in E'_\theta \text{ with } x'_e > 0, \\ \ell'_v &\geq \max_{e=vw \in E'_\theta} \frac{x'_e}{b_e^+} \quad \text{for } v \in V, \\ \ell'_v &= \max_{e=vw \in E'_\theta} \frac{x'_e}{b_e^+} \quad \text{for } v \in V \text{ with } c_v < 1, \end{aligned}$$

where

$$\rho_e(\ell'_u, x'_e, c_v) := \begin{cases} \frac{x'_e}{c_v \cdot \nu_e} & \text{if } e = uv \in E_\theta^*, \\ \max \left\{ \ell'_u, \frac{x'_e}{c_v \cdot \nu_e} \right\} & \text{if } e = uv \in E'_\theta \setminus E_\theta^*. \end{cases}$$

It turns out that the particles of a Nash flow over time f can be divided into intervals, so called *phases*, for which the derivatives (and thus the inflow and outflow rates) stay constant. We denote the set of phases by \mathcal{I}_f . The transition points between two phases correspond to one or multiple *events*: A new edge (and therefore new s - t -paths) can become active, a queue can deplete, an edge can become full or the outflow rate (and hence the inflow bound) of a full edge might change. Note however, that an event at edge $e = uv$ for a particle θ does not happen at time θ itself but rather at time $\ell_u(\theta)$ when the particle entering the network at time θ reaches vertex u (while taking a shortest s - u path).

Games, Optimal Flows and the Price of Anarchy. For a *temporal routing game* we consider a finite volume of flow $M \in (0, \infty)$ entering the network. For a Nash flow over time f the last particle enters the network at time $\frac{M}{r_0}$ and leaves the network at time $\ell_t(\frac{M}{r_0})$. As the network satisfy the first-in-first-out-principle (FIFO), ℓ_t is non-decreasing, which means that $T_f := \ell_t(\frac{M}{r_0})$ denotes the *completion time* when the entire flow of volume M has reached t . Most of the time we identify a network Γ with its corresponding temporal routing game (i.e., Γ and M). In contrast to dynamic equilibria, optimal *quickest flows* can

be computed by determining a time horizon T with $\ell_t(\frac{M}{r_0}) = T$ by applying a binary search framework to the maximum flow over time problem. Hereby, a maximum flow over time for time horizon T can be constructed via a feasible static flow y maximizing $T \cdot |y| - \sum_{e \in E} \tau_e \cdot y_e$. This *underlying static flow* y is then temporally repeated, which means a rate of y_p is sent into every s - t path $p \in \mathcal{P}_{st}$ over time $[0, T - \tau_p]$. The arrival time of the last particle at t (i.e., the optimal completion time) is denoted by $T_{\text{opt}}(M)$. For more details on optimal flows over time we refer to Skutella's survey [23].

In this paper we consider the *time price of anarchy* (which we simply refer to as "price of anarchy"). For a given temporal routing game Γ it measures the worst ratio between the arrival time at t for the last particle in a Nash flow over time and the arrival time in an optimal flow: $\text{PoA}(\Gamma) := \frac{T_{EQ}(\Gamma)}{T_{\text{opt}}(\Gamma)}$.¹ As it is unknown whether the arrival time functions ℓ_t are unique over all Nash flows over time, we need to consider the worst dynamic equilibrium, i.e., $T_{EQ}(\Gamma) := \sup_{f \in \mathcal{F}(\Gamma)} T_f$, where $\mathcal{F}(\Gamma)$ denotes the set of Nash flows over time in Γ .

Further Notation. We enumerate the event points by the order of their occurrence seen by particles at the source, i.e., $\theta_i < \theta_{i+1}$ and say phase i is given by (θ_{i-1}, θ_i) (using $\theta_0 = 0$).² In addition, we consider the point in time $\frac{M}{r_0}$ when the last particle enters the network as the last event r , i.e., $\theta_r := \frac{M}{r_0}$. Since the edge sets E'_θ , E_θ^* , \bar{E}_θ , the inflow bound, and hence, the spillback thin flow stay constant within each phase i we use the following notation for $\theta \in (\theta_{i-1}, \theta_i)$

$$G'_i := G'_\theta, E'_i := E'_\theta, E_i^* := E_\theta^*, \bar{E}_i := \bar{E}_\theta, x'_i := x'(\theta), \ell'_{i,v} := \ell'_v(\theta), c_{i,v} := c_v(\ell_v(\theta)).$$

The inflow at the sink is also constant in a phase. We denote this by the *capacity* $\kappa_i := f_t^+(\ell_t(\theta))$ for some $\theta \in (\theta_{i-1}, \theta_i)$ where we use $f_t^+(\theta) := \sum_{vt \in \delta^-(t)} f_{vt}^-(\theta)$. Finally, the derivatives of the waiting times $(q_e(\ell_u(\theta)))'$ stay constant within a phase as they are either $\ell'_v(\theta) - \ell'_u(\theta)$ if $e = uv$ is active or 0 otherwise. For $\theta \in (\theta_{i-1}, \theta_i)$ we write $q'_{i,e} := (q_e(\ell_u(\theta)))'$ and $q'_{i,p} := \sum_{e \in p} q'_{i,e}$ for an s - t path p .

3 Lower Bounds on the Price of Anarchy

We first show in 3 that the PoA can be unbounded even on very simple graphs (an observation first made in [22]) and that the same is true for graphs with unit capacities. Afterwards, in 3 and 3, we use similar constructions to investigate the Braess paradox for flows over time with spillback and to show that

¹ All results from Sect. 4 can also be translated to the *total delay price of anarchy* measuring the arrival times of all particles combined, similarly as is done in [1].

² We imagine i as a natural number. But since it is an open question whether the event point converges to a finite limit, it is possible to expand the index set to the ordinal numbers up to ω^ω . In this case the i -th phase should be defined as (θ_i, θ_{i+1}) as it is not possible to determine a predecessor of an ordinal number. For the sake of simplicity however, we stick to the definition where $(0, \theta_1)$ is the first phase.

there exist networks on which Nash flows over time with spillback are faster than their respective counterparts without spillback.

PoA Depending on Graph Structure or Capacities. Consider the network Γ given in Fig. 1 and a Nash flow over time f of it. Since in the first phase the shortest path is (e_1, e_2) , edge e_2 fills up quickly. Once this happens flow already queues up at the end of edge e_1 , and thus, e_3 is never used by f . An optimal flow can use e_3 and therefore routes flow to the sink much faster, resulting in an unbounded PoA. This construction can easily be generalized to all graphs that have the graph given in Fig. 1 as a minor.

Theorem 1. (cf. [22, introductory example]) *Let G be any graph that has the graph given in Fig. 1 as a minor, then there exists a temporal routing game Γ on G with $PoA(\Gamma) \in \Omega(\frac{1}{\nu_{min}^-})$ where $\nu_{min}^- := \min_{e \in E} \{\nu_e^- : \nu_e^- > 0\}$.*




Fig. 1. This is a network on which Nash flows over time with spillback have unbounded price of anarchy (see Theorem 1). A similar example was first given in [22, Fig. 2].

To avoid the above unboundedness one could ask for the PoA for temporal routing games on graphs with restricted edge capacities. By restrictions of the model we have to set the inflow and storage capacities of all edges $e \in \delta^+(s)$ to $\nu_e^+ > r$ and $\sigma_e = \infty$, respectively. We say a network has *unit edge capacities* if for all edges $e \notin \delta^+(s)$ it holds that $\nu_e^+ = \nu_e^- = \sigma_e = 1$ and further for all edges $e \in \delta^+(s)$ also $\nu_e^- = 1$. Unfortunately, even when restricting to networks with unit edge capacities the PoA is unbounded.

Theorem 2. *capacities There exists a family of networks with unit edge capacities and $\tau_e \in \{0, 1\}$ for all edges for which the PoA is linear in the number of edges.*

This can be seen by considering the network given in Fig. 1 and exchanging e_1 and e_3 with bunches of unit-capacity parallel edges and setting $\nu_{e_2}^- = \sigma_{e_2} = 1$. If we use enough parallel edges we can generate a similar flow behavior as we encountered when lowering the capacity of edge e_2 in the proof of Theorem 1.

Nevertheless, we show another way of constraining edge capacities to achieve an interesting upper bound on the PoA in Sect. 4.2.

Braess Ratio. In his work on selfish routing with static flows [2] Braess showed that there are networks where adding an edge can paradoxically increase congestion leading to a worse equilibrium. In line with the paper of Macko et al. [14]

we define the *Braess ratio* for flows over time with spillback as follows. Let Γ be a temporal routing game on a graph G and let $\Gamma(H)$ be the same instance restricted to some subgraph $H \subseteq G$. Then the Braess ratio of Γ is

$$\text{BR}(\Gamma) = \max_{H \subseteq G} \frac{T_{\text{EQ}}(\Gamma)}{T_{\text{EQ}}(\Gamma(H))}.$$

We say graph G admits a *Braess paradox* if there is a temporal routing game Γ on G with $\text{BR}(\Gamma) > 1$. In [14] it is shown that the Braess ratio for flows over time without spillback (for a slightly different cost function instead of the last completion time) is arbitrarily large depending linearly on the number of edges of the underlying graph. The authors furthermore show that a graph G or its transpose (the graph where every edge uv is replaced by the edge vu and s and t are swapped) admit a Braess paradox if and only if G contains at least one of the following graphs as a topological minor.



When considering flows over time with spillback and the graph in Fig. 1 it is easy to see that this graph admits a Braess paradox with arbitrarily large Braess ratio even though it does not have one of the graphs above as a topological minor (and neither does its transpose). To see this choose H to be the subgraph where from the graph in Fig. 1 we delete edge e_2 .

Corollary 1. *For any $a \in \mathbb{R}$ there exists a temporal routing game Γ on the graph given in Fig. 1 such that the Braess ratio satisfies $\text{BR}(\Gamma) > a$.*

Spillback Can Improve Completion Time. The following proposition shows that there are temporal routing games where Nash flows with spillback perform better than Nash flows without spillback. This might at first be surprising, as spillback seems to only be obstructive to routing flow fast. But it is indeed possible to construct networks where spillback leads to shorter completion times. In the network depicted in Fig. 2 there are two parallel edges, namely e_3 and $e_{3'}$, for which it holds that the completion time of a Nash flow is worse if the edges are present compared to the same network without those edges. We exploit this in our construction: In the spillback model one of these ‘bad’ edges becomes full nearly instantaneously yielding the other ‘bad’ edge to never get active. Thus, the spillback Nash flow routes flow only over one of those ‘bad’ edges. Since in the Koch-Skutella model without spillback both of these parallel edges get active at some point, the Nash flow over time here uses both of them resulting in a worse completion time.

Proposition 1. *In the network Γ given in Fig. 2 the completion time of any Nash flow over time with spillback is less than the completion time of the Nash flow over time without spillback on the same network using $\nu_e := \min\{\nu_e^+, \nu_e^-\}$.*




Fig. 2. This example shows that Nash flows over time with spillback can be faster than Nash flows over time without spillback, see Proposition 1.

4 Upper Bounds on the Price of Anarchy

In the following we prove two upper bounds on the price of anarchy. First, we show that for a single, fixed network the PoA is bounded by a constant in the long run. After that we show that if for a given network we are allowed to decrease the outflow capacities by a certain amount then the PoA only depends on the worst spillback factor of the Nash flows over time.

4.1 Price of Anarchy for a Fixed Network

Until now we have studied the PoA depending on the structure of the underlying graph or its capacities. For both questions we constructed games satisfying strong constraints that still have unbounded PoA. Now we are interested in the PoA of a network where every parameter is fixed except for the target amount M , i.e., we ask the question of how the PoA behaves in the long run on a single network.

Lemma 1. *For a temporal routing game Γ on a fixed network the completion time of the optimal flow depending on the target amount M is bounded by $T_{opt}(M) \in \Theta(M)$.*

This result is mainly due to the fact that the optimal flow does not build up any queues. Therefore its completion time depends mainly on M and the minimum edge-capacity, which we consider to be fixed.

For the classification of the asymptotic long term behavior of $T_{EQ}(M)$ we use the following auxiliary lemma that gives us a lower bound on the spillback factors of a Nash flow. The lemma follows by an application of [22, Lemma 3].

Lemma 2. *For a temporal routing game Γ there exists an $\varepsilon > 0$ such that for any Nash flow over time $f \in \mathcal{F}(\Gamma)$ the spillback factors satisfy $\min\{c_v(\theta) : v \in V, \theta \in \mathbb{R}_{\geq 0}\} > \varepsilon$.*

To get an asymptotic bound on $T_{\text{EQ}}(M) = \sup_{f \in \mathcal{F}(\Gamma)} T_f(M)$ we first argue that seen as a function in M , T_f is a piece-wise linear and non-decreasing function. We can then use Lemma 2 to bound its derivative and with that obtain the desired result.

Theorem 3. *For a temporal routing game Γ on a fixed network the completion time of any Nash flow over time $f \in \mathcal{F}(\Gamma)$ is bounded by $T_f(M) \in \Theta(M)$.*

We can now use Lemma 1 and Theorem 3 and the fact that $\text{PoA}(M) = \frac{T_{\text{EQ}}(M)}{T_{\text{opt}}(M)}$ to bound the PoA for a fixed network. In order to do so we consider the PoA as a function of the target flow amount M .

Theorem 4. *For a temporal game Γ on a fixed network, i.e. when treating everything except the amount of flow M as a constant, the price of anarchy is bounded by a constant, $\text{PoA}(M) \in \Theta(1)$.*

4.2 Bound on the Price of Anarchy for Saturated Graphs

In this section we focus on networks with an additional constraint on the edge capacities. Given a game Γ we know that the quickest flow of Γ is also a temporally repeated flow, i.e., it has an underlying static flow y . We say that y *saturates every edge* of the given graph if for each edge the outflow capacity is exhausted by y , i.e., for each $e \in E$ we have $\nu_e^- = y_e$ and additionally it holds that $|y| = \sum_{sv \in \delta^+(s)} y_{sv} = r_0$. We call the underlying graph of such a game a *saturated graph*. Even though restricting attention to saturated graphs may seem harsh, note, that every network can be made saturated by lowering the edge capacities. This can be imagined to be done by a system operator in a Stackelberg strategy-like scenario [25] and is applicable in many real-world examples. For one, streets can be narrowed down by a city administration in practice.

For temporal routing games on saturated graphs we will show that the PoA can be bounded by a value that is only dependent on the worst spillback factor of all Nash flows over time. In order to do that we adapt the idea of the proofs given by Bhaskar et al. [1] for the Koch-Skutella model to the spillback model. Note, however, that the proofs given in [1] implicitly assume only finitely many phases, which has not been proven for any of the two models. Our generalization also holds for the case of an infinite number of phases in both models.³

In principle the proof works as follows. For a given game Γ the relation of the completion time of any Nash flow over time of Γ to the optimal completion time can be determined by examining the capacity of the current shortest path network and the derivatives of the waiting times for a single phase of the Nash flow. One can then bound the derivatives of the waiting times and use the fact that the PoA is the maximum over the relation of the optimal completion time to the completion times of all Nash flows. This achieves the desired bound.

³ Note, that in [7] an even more general result is shown for the Koch-Skutella model.

Bound on the Derivatives of the Waiting Times. We start by proving a relation between the derivative of the label-function at the sink and the inflow into the sink. Our proof of this result uses a different idea than the one given in [1] and is considerably shorter.

Lemma 3. (cf. [1, Lemma 15]) *Let Γ be a temporal routing game and let $f \in \mathcal{F}(\Gamma)$ with corresponding labels ℓ . Then for any $\theta \leq \frac{M}{r_0}$ we have*

$$\ell'_t(\theta) = \frac{r_0}{f_t^+(\ell_t(\theta))}.$$

Proof. Let $\theta \leq \frac{M}{r_0}$ be arbitrary. Using that $x'(\theta)$ is a static s - t flow of value r_0 and $x'_{vt}(\theta) = f_{vt}^-(\ell_t(\theta)) \cdot \ell'_t(\theta)$ from Eq. (1) we obtain

$$r_0 = \sum_{vt \in \delta^-(t)} x'_{vt}(\theta) = \sum_{vt \in \delta^-(t)} f_{vt}^-(\ell_t(\theta)) \cdot \ell'_t(\theta) = \ell'_t(\theta) \cdot f_t^+(\ell_t(\theta)).$$

Since $f_t^+(\ell_t(\theta)) > 0$ for all θ , rearranging terms give the desired result. □

We now proceed with a path-wise bound on the derivatives of the waiting times $q'_{i,p}$ for a single phase of the Nash flow over time i using the capacities κ_i .

Lemma 4. (cf. [1, Lemma 18]) *Let Γ be a temporal routing game where the static flow underlying the quickest flow saturates every edge and let $f \in \mathcal{F}(\Gamma)$. For any s - t path p , the travel time is bounded by*

$$\tau_p \geq \ell_t(\theta_r) - \sum_{i \in \mathcal{I}_f} (1 + q'_{i,p}) \cdot \frac{\kappa_i}{r_0} \cdot (\ell_t(\theta_i) - \ell_t(\theta_{i-1})).$$

In the proof we first establish a dependence of the length of a phase as it is experienced at the source and at the sink, respectively. Then we express τ_p in terms of the label functions ℓ and the waiting times q and their derivatives. The result then follows from applying Lemma 3.

Relation of the Completion Times of Nash Flow and Quickest Flow.

The following lemma enables us to give a first relation of the completion times of the optimal quickest flow and a Nash flow over time.

Lemma 5. (cf. [1, Lemma 19]) *Let Γ be a temporal routing game where the static flow y underlying the quickest flow saturates every edge and let $f \in \mathcal{F}(\Gamma)$. Then, the completion time T_{opt} of the optimal flow and the completion time T_f of the Nash flow f are related as*

$$r_0 \cdot T_{opt} = \sum_{p \in \mathcal{P}_{s,t}} y_p \tau_p + \sum_{i \in \mathcal{I}_f} \kappa_i \cdot (\ell_t(\theta_i) - \ell_t(\theta_{i-1})),$$

where $\mathcal{P}_{s,t}$ is the set of all simple s - t paths in G and $\ell_t(\theta_r) = T_f$.

The proof idea is to compare the arrival rates of both flows at the sink t where we use a flow decomposition along paths for the optimal flow and a decomposition by phases for the Nash flow over time.

By combining the previous two lemmas we can now derive a lower bound on the inverse of the PoA that we will afterwards use to achieve an upper bound on the actual PoA. But in order to prove that we first need the following.

Lemma 6. *Let $\lambda_i := \frac{\kappa_i}{r_0} \cdot \sum_{p \in \mathcal{P}_{s,t}} y_p q'_{i,p}$ for each phase $i \in \mathcal{I}_f$. Then,*

$$\sum_{i \in \mathcal{I}_f} \lambda_i \cdot (\ell_t(\theta_i) - \ell_t(\theta_{i-1})) \leq (\ell_t(\theta_r) - \ell_t(\theta_0)) \cdot \sup_{i \in \mathcal{I}_f} \lambda_i.$$

In the proof we first establish that the set $\{\lambda_i : i \in \mathcal{I}_f\}$ is bounded and then use this and the telescoping principle to bound the left hand side.

The next lemma establishes the aforementioned bound on the inverse of the PoA. It is in this proof that the number of α -extension phases comes into play. If we assume that the supremum in the statement of Lemma 6 is attained by some phase $i \in \mathcal{I}_f$, which is in particular true if there are only finitely many phases, then we can prove Lemma 7 without the ε error and the proofs go through similar to [1]. But since it is still an open problem whether the number of those phases is always finite (in the Koch-Skutella model as well as the spillback model), we prove it here for the case of infinitely many α -extension phases.

Lemma 7. *Let Γ be a temporal routing game where the static flow y underlying the quickest flow saturates every edge and let $f \in \mathcal{F}(\Gamma)$. Then for every $\varepsilon > 0$ there exists a phase i of f such that*

$$\frac{T_{opt}}{T_f} + \varepsilon \geq 1 - \frac{\kappa_i}{r_0^2} \sum_{e \in E} \nu_e^- q'_{i,e}.$$

The proof idea is to sum $y_p \tau_p$ over all paths $p \in \mathcal{P}_{s,t}$ and using Lemma 4 to bound this from below. Afterwards, we use Lemmas 5 and 6 to obtain a lower bound on $\frac{T_{opt}}{T_f}$ in terms of a supremum of the capacities and derivatives of the queuing delay over all phases. Since we do not know whether this supremum is attained we have to inject the ε error and after rearranging terms we obtain the desired result.

Upper Bound for Saturated Graphs. We can now turn the lower bound in Lemma 7 into an upper bound on the price of anarchy by proving a bound on the sum of the right-hand side of the expression given in Lemma 7. Here for the first time the spillback factors of the Nash flow over time play an important role.

Lemma 8. *Let Γ be a temporal routing game and $f \in \mathcal{F}(\Gamma)$. In any phase i of f where $\frac{r_0}{\kappa_i} \geq 1$ we have*

$$\sum_{e \in E} \nu_e^- q'_{i,e} \leq \frac{r_0}{c_i^f} \ln \left(\frac{r_0}{\kappa_i} \right),$$

where $c_i^f := \min\{c_v(\theta) : v \in V, \theta \in (\theta_{i-1}, \theta_i)\}$ is the minimal c_v of f in phase i .

The proof utilizes [7, Claim 12] and follows the line of argumentation in [1] but incorporates the added complexity of the spillback model. We obtain that $c_v v_e^- q_e' = x_e' \cdot (1 - \frac{\ell_e'}{\ell_v})$ for every edge $e = uv$ and then sum this expression over all edges in the graph. Rearranging and plugging in the above expression then yields the desired result.

We can now obtain the desired upper bound on the price of anarchy.

Theorem 5. *Let Γ be a temporal routing game where the static flow y underlying the quickest flow saturates every edge of the graph. If the minimal spillback factor satisfies $c := \min_{f \in \mathcal{F}(\Gamma)} \min\{c_v(\theta) : v \in V, \theta \in \mathbb{R}_{\geq 0}\} > \frac{1}{e}$, then the price of anarchy is bounded by $\frac{T_{\text{EQ}}}{T_{\text{opt}}} \leq \frac{ce}{ce-1}$.*

Proof. For any $f \in \mathcal{F}(\Gamma)$ with completion time T_f we know that $f_t^+(\theta) = \sum_{vt \in \delta^-(t)} f_{vt}^-(\theta) \leq r_0$ for all $\theta \in \mathbb{R}_{\geq 0}$ since we only consider saturated graphs. Thus, we have $\frac{r_0}{\kappa_i} \geq 1$ in all phases of f . From Lemmas 7 and 8 we obtain that for every $\varepsilon > 0$ there exists a phase i of f such that

$$\frac{T_{\text{opt}}}{T_f} + \varepsilon \geq 1 - \frac{\kappa_i}{r_0^2} \sum_{e \in E} v_e^- q_{i,e}' \geq 1 - \frac{\kappa_i r_0}{r_0^2 c_i^f} \cdot \ln \left(\frac{r_0}{\kappa_i} \right) = 1 - \frac{a_i}{c} \cdot \ln \left(\frac{1}{a_i} \right),$$

where $c := \min_{f \in \mathcal{F}(\Gamma)} \min\{c_v(\theta) : v \in V, \theta \in \mathbb{R}_{\geq 0}\} \leq c_i^f$ and $a_i := \frac{\kappa_i}{r_0}$.

Simple calculus shows that the term $\frac{a_i}{c} \cdot \ln \left(\frac{1}{a_i} \right)$ is maximized for $a_i = \frac{1}{e}$. Using the above inequality, derived from some phase i , for any $\varepsilon > 0$ we obtain

$$\frac{T_{\text{opt}}}{T_f} + \varepsilon \geq 1 - \frac{1}{ce} = \frac{ce-1}{ce}.$$

Since by assumption we have $c > \frac{1}{e}$ we can take the inverse of the inequality to obtain $\frac{T_f}{T_{\text{opt}}} \leq \frac{ce}{ce-1}$. We finish by noting that $T_{\text{EQ}} = \sup_{f \in \mathcal{F}(\Gamma)} T_f$. \square

5 Conclusions

Our work shows that the PoA is highly dependent on spillback effects. Although, even in restricted network classes the completion times of dynamic equilibria can be arbitrarily bad compared to a quickest flow, the PoA can still be bounded in terms of the spillback factors under some constraints on the edge capacities. Transferred to real-world traffic this means the interplay between selfish traffic users is critical in particular in high congested areas.

Even though we give a substantial analysis of the PoA in the flow over time model with spillback, there are still some open problems remaining. Is the bound we establish in Theorem 5 tight? Are there any bounds in the case of $c \leq \frac{1}{e}$ or is it possible to enforce $c > \frac{1}{e}$ through some Stackelberg-like strategy? Do the results of the recent work of Correa et al. [7] also transfer to the spillback setting? On the more applied side of the research it would also be very interesting to algorithmically identify street segments (edges) which are especially vulnerable

for spillback. In the long run this could help road administrations to decide which roads should be expanded (increasing the storage capacity) or which roads should be narrowed or closed (due to the Braess effect).

References

1. Bhaskar, U., Fleischer, L., Anshelevich, E.: A stackelberg strategy for routing flow over time. *Games Econ. Behav.* **92**, 232–247 (2015)
2. Braess, D.: Über ein paradoxon aus der verkehrsplanung. *Unternehmensforschung* **12**(1), 258–268 (1968). <https://doi.org/10.1007/BF01918335>
3. Cao, Z., Chen, B., Chen, X., Wang, C.: A network game of dynamic traffic. In: *Proceedings of the 2017 ACM Conference on Economic and Computation*, pp. 695–696 (2017)
4. Cominetti, R., Correa, J.R., Larré, O.: Existence and uniqueness of equilibria for flows over time. In: Aceto, L., Henzinger, M., Sgall, J. (eds.) *ICALP 2011. LNCS*, vol. 6756, pp. 552–563. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22012-8_44
5. Cominetti, R., Correa, J., Larré, O.: Dynamic equilibria in fluid queueing networks. *Oper. Res.* **63**(1), 21–34 (2015)
6. Cominetti, R., Correa, J., Olver, N.: Long term behavior of dynamic equilibria in fluid queueing networks. In: Eisenbrand, F., Koenemann, J. (eds.) *IPCO 2017. LNCS*, vol. 10328, pp. 161–172. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-59250-3_14
7. Correa, J., Cristi, A., Oosterwijk, T.: On the price of anarchy for flows over time. In: *Proceedings of the 2019 ACM Conference on Economic and Computation*, pp. 559–577. ACM (2019)
8. Ford Jr., L.R., Fulkerson, D.R.: *Flows in Networks*. Princeton University Press, Princeton (2015)
9. Gale, D.: Transient flows in networks. *Mich. Math. J.* **6**(1), 59–63 (1959)
10. Graf, L., Harks, T., Sering, L.: Dynamic flows with adaptive route choice. *Math Program.* 1–27 (2020). <https://doi.org/10.1007/s10107-020-01504-2>
11. Harks, T., Peis, B., Schmand, D., Tauer, B., Vargas Koch, L.: Competitive packet routing with priority lists. *ACM Trans. Econ. Comput.* **6**(1), 4 (2018)
12. Koch, R., Skutella, M.: Nash equilibria and the price of anarchy for flows over time. *Theory Comput. Syst.* **49**(1), 71–97 (2011). <https://doi.org/10.1007/s00224-010-9299-y>
13. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: Meinel, C., Tison, S. (eds.) *STACS 1999. LNCS*, vol. 1563, pp. 404–413. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-49116-3_38
14. Macko, M., Larson, K., Steskal, Ľ.: Braess’s paradox for flows over time. In: Koutogiannis, S., Koutsoupias, E., Spirakis, P.G. (eds.) *SAGT 2010. LNCS*, vol. 6386, pp. 262–275. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-16170-4_23
15. Meunier, F., Wagner, N.: Equilibrium results for dynamic congestion games. *Transp. Sci.* **44**(4), 524–536 (2010)
16. Papadimitriou, C.H.: Algorithms, games, and the internet. In: Orejas, F., Spirakis, P.G., van Leeuwen, J. (eds.) *ICALP 2001. LNCS*, vol. 2076, pp. 1–3. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-48224-5_1

17. Peis, B., Tauer, B., Timmermans, V., Vargas Koch, L.: Oligopolistic competitive packet routing. In: 18th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (2018)
18. Roughgarden, T.: Selfish Routing and the Price of Anarchy, vol. 174. MIT Press, Cambridge (2005)
19. Roughgarden, T., Tardos, É.: How bad is selfish routing? *J. ACM (JACM)* **49**(2), 236–259 (2002)
20. Scarsini, M., Schröder, M., Tomala, T.: Dynamic atomic congestion games with seasonal flows. *Oper. Res.* **66**(2), 327–339 (2018)
21. Sering, L., Skutella, M.: Multi-source multi-sink Nash flows over time. In: 18th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems, vol. 65, pp. 12:1–12:20 (2018)
22. Sering, L., Vargas Koch, L.: Nash flows over time with spillback. In: Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 935–945. SIAM (2019)
23. Skutella, M.: An introduction to network flows over time. In: Cook, W., Lovász, L., Vygen, J. (eds.) *Research Trends in Combinatorial Optimization*, pp. 451–482. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-540-76796-1_21
24. Vickrey, W.S.: Congestion theory and transport investment. *Am. Econ. Rev.* **59**(2), 251–260 (1969)
25. Von Stackelberg, H.: *Marktform und Gleichgewicht*. J. Springer, Heidelberg (1934)
26. Wilkinson, W.L.: An algorithm for universal maximal dynamic flows in a network. *Oper. Res.* **19**(7), 1602–1612 (1971)