http://pubsonline.informs.org/journal/moor

MATHEMATICS OF OPERATIONS RESEARCH
Articles in Advance, pp. 1–18
ISSN 0364-765X (print), ISSN 1526-5471 (online)

On the Price of Anarchy for Flows over Time

José Correa,[a] Andrés Cristi,[a] Tim Oosterwijk[b]

[a] Department of Industrial Engineering, Universidad de Chile, República 701, 8370439 Santiago, Chile; [b] Vrije Universiteit Amsterdam, De Boelelaan 1105, 1081 HV, Amsterdam, Netherlands

Contact: correa@uchile.cl, https://orcid.org/0000-0002-3012-7622 (JC); andres.cristi@ing.uchile.cl, https://orcid.org/0000-0002-1227-2092 (AC); t.oosterwijk@vu.nl, https://orcid.org/0000-0002-8509-7002 (TO)

Received: July 30, 2019
Revised: October 2, 2020
Accepted: February 25, 2021
Published Online in Articles in Advance: December 22, 2021

MSC2000 Subject Classification: Primary: 91A25, 91A43; secondary: 91A10, 91A13

https://doi.org/10.1287/moor.2021.1173

Abstract. Dynamic network flows, or network flows over time, constitute an important model for real-world situations in which steady states are unusual, such as urban traffic and the internet. These applications immediately raise the issue of analyzing dynamic network flows from a game-theoretic perspective. In this paper, we study dynamic equilibria in the deterministic fluid queuing model in single-source, single-sink networks—arguably the most basic model for flows over time. In the last decade, we have witnessed significant developments in the theoretical understanding of the model. However, several fundamental questions remain open. One of the most prominent ones concerns the price of anarchy, measured as the worst-case ratio between the minimum time required to route a given amount of flow from the source to the sink and the time a dynamic equilibrium takes to perform the same task. Our main result states that, if we could reduce the inflow of the network in a dynamic equilibrium, then the price of anarchy is bounded by a factor, parameterized by the longest path length that converges to $e/(e-1)$, and this is tight. This significantly extends a result by Bhaskar et al. (SODA 2011). Furthermore, our methods allow us to determine that the price of anarchy in parallel-link and parallel-path networks is exactly $4/3$. Finally, we argue that, if a certain, very natural, monotonicity conjecture holds, the price of anarchy in the general case is exactly $e/(e-1)$.

Funding: This work was supported by ANID Chile [Grants PFCHA/Doctorado Nacional/2018-21180347, FONDECYT 1190043, and AFB-170001].

Keywords: flows over time • price of anarchy • dynamic equilibrium

1. Introduction

In the study of traffic in networks, it is often crucial to take the underlying dynamic nature of the problem into account. In some contexts, steady states seem sufficient to deal with the most important situations, and therefore, static models are enough. However, the situation is dramatically different when dealing with networks in which a steady state is rarely observed, such as urban traffic or internet routing. In order to describe the temporal evolution of such systems, one has to consider the propagation of flow across the network by tracking the position of each particle along time.

Probably the most basic model for network flows over time is the so-called *fluid queuing model*. Here, we are given a directed graph $G=(V,E)$, and each edge $e \in E$ is characterized by a nonnegative delay τ_e and a capacity per time unit ν_e. A continuous stream of particles is injected at a source $s \in V$ at constant rate u_0 and travels toward a sink $t \in V$. Flow propagates according to the edge dynamics in which particles arriving to an edge e join a queue with (deterministic) service rate ν_e and, after leaving the queue, move along the edge to reach its head after τ_e time units.

The discrete version of the problem was initially studied from an optimization perspective. Indeed, Ford and Fulkerson [9, 10] consider a fluid queuing model and design an algorithm based on time-expanded networks to compute a flow over time carrying the maximum possible flow from the source s to the sink t within a given time span. Shortly after, Gale [11] shows the existence of a flow pattern that achieves this optimum simultaneously for all time horizons. These results are extended to continuous time by Anderson and Philpott [1] and Fleischer and Tardos [8]. We refer to the survey by Skutella [28] for a detailed exposition of these developments.

When network flows suffer from a lack of coordination among the participating agents, it is natural to consider them from a game-theoretic perspective. In this setting, each infinitesimal inflow particle is interpreted as a player that seeks to complete its journey in the least possible time so that equilibrium occurs when each particle travels along a shortest s,t path. The travel time for a particle entering the network at any given time must take into account the queuing delays induced by other particles on the edges along its path. This requires particles to anticipate the queue lengths by the time when an edge is reached.

This *dynamic equilibrium* model was initially considered, in a very simple network, by Vickrey [29], and shortly after in the transportation science community (Yagar [30]). Since then, it has attracted much attention as a showcase model to understand the surprising behavior of dynamic routing games (Peeta and Ziliaskopoulos [22], Ran and Boyce [23]). In the last decade, there have been significant efforts in understanding the structure and computational properties of dynamic equilibria in the fluid queuing model (Bhaskar et al. [3], Cao et al. [4], Cominetti et al. [5, 6], Graf and Harks [13], Ismaili [15], Kaiser [16], Koch and Skutella [18], Meunier and Wagner [21], Scarsini et al. [24], Sering and Skutella [25], Sering and Vargas Koch [26]). Meunier and Wagner [21] prove, using functional analysis tools, that such dynamic equilibria exist. Unfortunately, this result (and many similar ones) is purely existential and does not shed light on the structure of such equilibria. Later, Koch and Skutella [18] give an elegant characterization of the derivatives (with respect to time) of a dynamic equilibrium and, thus, propose an algorithm to construct a dynamic equilibrium by concatenating static flows. Using this characterization, Cominetti et al. [5] give a constructive proof of existence of equilibria and prove they are essentially unique.[1] Despite these efforts, many fundamental questions remain open, and several apparently obvious properties turn out to be notoriously hard to prove. For instance, it is still unknown whether a dynamic equilibrium can be computed in polynomial time, and furthermore, we do not even know whether the evolution of the equilibrium has finitely many pieces. Indeed, until recently, it was not even known whether the size of the queues remains bounded throughout the evolution of the dynamic equilibrium. Along these lines, Cao et al. [4] establish this property (on a slightly different atomic model that does not influence the result) for series-parallel networks, and Cominetti et al. [6] establish the result for general networks by proving that a steady state is always achieved in finite time (naturally, as long as u_0 is at most the capacity of the minimum cut). Quite surprisingly however, the latter results apply only for constant inflow rate u_0; if the inflow varies over time, say it is u_0 in all intervals of the form $[2i, 2i+1)$ and $u_0/2$ in all intervals of the form $[2i-1, 2i)$ for $i \in \mathbb{N}$, then the boundedness of the queues is still open.

Another seemingly innocent question regarding the dynamic equilibrium is what we call the *monotonicity conjecture* (cf. conjecture 1). This states that, given an instance of the problem, the time it takes for an amount of flow to reach the sink t is a decreasing function of the inflow rate u_0. In other words, if we consider two identical instances, one with constant inflow rate u_0 and the other with constant inflow rate $u_0 - \varepsilon$, then the time it takes for M flow units to arrive at t in the latter instance is at least that in the former. As we show in this paper, this conjecture is intimately connected to one of the most prominent open problems in the area, namely, the quality of the equilibrium (measured as the time required to send a given amount of flow from s to t) when compared with the optimal solution. Our main result, which can be seen as an improvement upon a result of Bhaskar et al. [3], implies that, if the monotonicity conjecture holds for the dynamic equilibrium, then the price of anarchy (PoA), defined as the worst-case ratio of the quality of an equilibrium to that of an optimal solution, is exactly $\mathrm{e}/(\mathrm{e}-1)$.

1.1. The Price of Anarchy

The usual way of quantifying the inefficiency of selfish behavior is the PoA. It is defined as the worst possible ratio between the quality of an optimal solution and the quality of an equilibrium (Koutsoupias and Papadimitriou [19]). In the context of fluid queuing networks, there are two natural and related goals that induce two natural possible definitions for the PoA. On the one hand, we have the *throughput* objective, under which we are given a time window and are asked to maximize the amount of flow that can reach the sink t within that time. On the other hand, we have the *makespan* objective, under which we are given an amount of flow M that needs to be routed to t in the shortest possible time.

The existence of an *earliest arrival flow*, established by Gale [11], implies that, from an optimization viewpoint, both goals are equivalent. Nevertheless, they induce different notions for the PoA. In the former case, the throughput PoA is, as usual, defined as the supremum over all single s,t graphs, all possible inflows, all possible capacities, all possible transit times, and all possible time windows of the ratio between the amount of flow the optimal solution can send and the amount of flow a dynamic equilibrium sends. In the latter case, the makespan PoA is defined as the supremum over all single s,t graphs, all possible inflows, all possible capacities, all possible transit times, and all possible amounts of flow M of the ratio between the time the optimal solution takes to route M units of flow toward t and the time it takes in a dynamic equilibrium.

The first to study the PoA in this context are Koch [17] and Koch and Skutella [18], who prove that the throughput PoA is unbounded. They also show that, if the delays of all edges are zero, then the dynamic equilibria are optimal, implying that both the throughput and the makespan PoA are one. Interestingly, it has long been conjectured that the makespan PoA is bounded by a small constant (Skutella [27]). The study of this makespan PoA measure is the main focus of this paper, which, from now on, we just call PoA for short.

Beyond the zero delay case, Bhaskar et al. [3] study this question from a mechanism design perspective and find that there is a way of reducing the capacities in the network so that the makespan of an equilibrium under the reduced capacities is within a factor $e/(e-1)$ of the optimal solution with the original capacities. Naturally, as the following example demonstrates, this capacity reduction can improve the behavior of a dynamic equilibrium by blocking particles from taking bad routes.

However, the result still requires a subtle analysis because reducing the capacities too much may also block good routes significantly, increasing the makespan of a dynamic equilibrium. More precisely, Bhaskar et al. [3] consider reducing the capacity of every edge e to be exactly the amount of flow rate the optimal solution propagates through e. Our main result implies that the same bound still holds by doing this only for the inflow, that is, leaving all capacities unchanged but only reducing the inflow.

Example 1. Consider the network in Figure 1, in which an edge e is labeled (ν_e, τ_e), let the total flow M to be sent through the network be two, and let $u_0 = 1$. We claim that the optimal flow sends $\frac{1}{2}$ units of flow along both the path (s, u, t) and the path (s, v, t) until time 2. Therefore, the makespan of the optimal flow is three. On the other hand, the equilibrium first sends one unit of flow along the path (s,u,v,t) from time 0 until time 1. Then, from time 1 to 2, it sends $\frac{1}{2}$ units of flow along both the path (s,u,v,t) and the path (s, v, t). Because a particle originating at s at time 2 encounters a queue time of one on edges (s, u) and (v, t), the makespan of the equilibrium is four, and hence, the PoA of this instance is $\frac{4}{3}$. Note that, if we set $\nu_{uv} = 0$ (as in Bhaskar et al. [3]), the equilibrium in the modified network does exactly the same as the optimal flow, and the new PoA is one.

Finally, Cominetti et al. [6] prove the existence of a steady state and, furthermore, establish that the derivative of this steady-state flow is the solution of the static minimum cost problem in which the cost of edge e is given by τ_e and the capacity by ν_e. This result readily implies that the price of anarchy converges to one as the amount of flow to be routed grows to infinity.

1.2. Our Results

As mentioned, our main contribution is to improve upon the result of Bhaskar et al. [3], who prove a bound of $e/(e-1)$ for the price of anarchy if one can control the capacity of every link in the network. To prove this, they find an expression for the price of anarchy that depends on the ratio between the equilibrium flow and the optimal flow in each edge. In this paper, we show that the same bound can be obtained under the milder assumption that the inflow rate of the equilibrium is equal to the (initial) inflow rate of the optimum flow. This is a theoretical improvement because it potentially makes further progress (e.g., on multicommodity settings) on this problem easier. Moreover, it can be of practical relevance because inflow-limiting mechanisms are easier to implement and currently used in many places, such as metered ramps on highways.

For a network G and a total amount of flow M, denote by T_{OPT} the time the optimal flow takes to route the M units from s to t (cf. Section 2.1 for a formal definition). The simplest algorithm to compute this quantity is that of Ford and Fulkerson [9], which we describe in Section 2.1. The basic idea is to guess $T = T_{OPT}$ and then find a static flow f maximizing $|f|\,T - \sum_{e\in E}\tau_e f_e$, where $|f|$ denotes the size of the flow and is constrained to be at most u_0. We denote the inflow rate of the optimal flow by $u_{OPT} = |f|$. Because, in the dynamic equilibrium, particles are selfish, its inflow rate u_{EQ} always equals u_0.

Similarly, let T_{EQ} be the time it takes for the equilibrium to route the M units of flow. Denote by m_G the maximum number of nodes in a simple s, t path in G. Our main result (cf. Theorem 2) establishes that $T_{EQ}\cdot(1-(u_{EQ}/u_{OPT})/e) \le T_{EQ}\cdot\left(1-(u_{EQ}/u_{OPT})(1-1/m_G)^{m_G}\right) \le T_{OPT}$. In particular, if $u_{EQ} = u_{OPT}$, then $T_{EQ} \le [e/(e-1)]\cdot T_{OPT}$. We also establish that this bound is the best possible. Note that $u_{EQ} = u_0 \ge u_{OPT}$; therefore, to establish that the price of anarchy equals $e/(e-1)$, the missing case is when $u_{EQ} > u_{OPT}$. We note that the bound holds in general if the following intuitive conjecture holds.

Figure 1. An illustration of the network of Example 1.

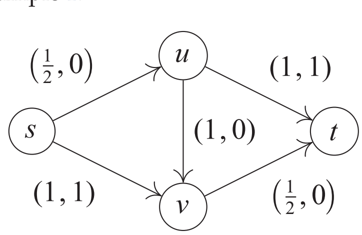

Conjecture 1: Consider a network G and two fixed inflow rates $u_1 < u_2$ with their corresponding dynamic equilibria in G and their corresponding makespans T^1_{EQ} and T^2_{EQ} for routing M units of flow. Then, $T^1_{EQ} \geq T^2_{EQ}$.

For the special case of series composition of parallel-path networks, in which all s,t paths are edge-disjoint, it turns out that the monotonicity conjecture holds (cf. Theorem 3). As a consequence of this result, we obtain that, in parallel-path networks, the PoA is equal to 4/3. We do this by showing that they behave essentially as if each path comprises only one direct link from s to t, in which case we can apply the general bound with $m_G = 2$.

The proof of the main result proceeds in three basic steps. First, we establish that the difference between the makespans $T_{EQ} - T_{OPT}$ is upper bounded by the overall sum of the queues at equilibrium divided by its inflow. This follows from the linear program that computes the optimal solution, combined with the equilibrium conditions stating that particles are routed through (currently) shortest paths. Second, we establish a formula for computing this sum of the queues at equilibrium in terms of the derivatives of the dynamic equilibrium (thin flows). Finally, the formula can be used to upper bound the sum of the queues at equilibrium by an expression in terms of m_G, u_{EQ} and u_{OPT}.

1.3. Further Related Literature

We wrap up this section by mentioning some further related work and variants of the model.

Hoefer et al. [14] study a similar atomic model with multiple sources and sinks and different policies (edge dynamics) and establish different existential and computational results for pure Nash equilibria. Ismaili [15] considers a similar atomic model with the first in, first out (FIFO) policy and establishes that even deciding the existence of a pure Nash equilibrium is hard.

Although most work about dynamic equilibrium in the fluid queuing model, including ours, applies to single-source, single-sink networks, there are some recent efforts to carry over the results to more general multicommodity networks. In particular, Garrido [12] is able to extend some of the results for dynamic equilibria to the case of multiple sinks, and Sering and Skutella [25] do it for the much more involved multisource, multisink case. However, we are still lacking a good understanding of the general multicommodity case.

As mentioned earlier, the issue of bounded queues is studied by Cao et al. [4], who prove that, in the atomic model and series-parallel networks, queues do remain bounded throughout the evolution of the dynamic equilibria. For the precise model of this paper, Cominetti et al. [6] establish this result in general networks by proving the existence of a steady state that is achieved in finite time. On a different line, Macko et al. [20] study new types of Braess' paradox appearing in the dynamic equilibrium.

Some very recent work considers other aspects of the problem. In particular, Sering and Vargas Koch [26] consider spillback effects, which is the study of how an a priori bound on the amount of flow that can be waiting on a queue affects the equilibrium behavior. Graf and Harks [13] consider a related model in which flow particles are myopic in that they make *local* routing decisions based on the current status of the network without anticipating the whole future evolution. Finally, Scarsini et al. [24] consider a discrete variant of the problem and look at the simpler parallel-link networks but add the complication that the inflow varies over time in a periodic fashion.

To close these comments, we note that a remarkable open problem concerns the polynomial time computation of the dynamic equilibria. By the work of Koch and Skutella [18] and Cominetti et al. [5], this boils down to computing in polynomial time a *normalized thin flow*, a special type of static flow with some complementary constraints (see Section 2.2). This problem can be solved in polynomial time in some special cases (Koch and Skutella [18]) by parametric flow techniques, and in general, it can be written as a nonlinear complementarity problem (Cominetti et al. [5], Koch [17]). Very recently, Kaiser [16] noted that the problem is actually a linear complementarity problem and that it can be solved efficiently in series-parallel networks.

1.4. Outline of the Paper

In Section 2, we describe the model and the behavior of the dynamic equilibrium and the optimal flow. Then, in Section 3, we prove our main result and establish its tightness. Section 4 shows how the monotonicity conjecture implies the general price of anarchy result and outlines some difficulties in trying to establish it. Section 5 contains a proof for the monotonicity conjecture for series compositions of parallel-path networks. Finally, in Section 6, we present computational experiments, alternative conjectures, and the implications of our results on the total delay price of anarchy.

2. The Model

Let $G = (V, E)$ be a directed graph, in which each edge $e \in E$ has a positive capacity ν_e and a nonnegative delay τ_e. Let $s, t \in V$ be two vertices that we refer to as the source and the sink, respectively. A total amount of flow M has to travel from s to t; flow departs from s at a network inflow rate denoted by u_0.[2] The flow propagates through the network as described by the following edge dynamics.

Let $f_e^+ : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be the function associated with an edge $e \in E$ that maps a nonnegative time θ to the inflow rate into e at time θ. In case the inflow rate $f_e^+(\theta)$ exceeds the edge capacity ν_e, a queue grows at the tail of the edge at rate $f_e^+(\theta) - \nu_e$. The queue mass at time θ is denoted by $z_e(\theta)$, and if $f_e^+(\theta) < \nu_e$, the queue depletes at a rate equal to $f_e^+(\theta) - \nu_e$ until the inflow rate changes again or until $z_e = 0$. Therefore, a particle that enters edge e at time θ waits in the queue $z_e(\theta)/\nu_e$ units of time and, subsequently, travels across the edge, taking time τ_e. Hence, this particle has link exit time

$$T_e(\theta) = \theta + \frac{z_e(\theta)}{\nu_e} + \tau_e\,.$$

This determines outflow rate functions $f_e^- : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ as follows.

$$f_e^-(\theta + \tau_e) = \begin{cases} \nu_e & \text{if } z_e(\theta) > 0, \\ \min\{f_e^+(\theta), \nu_e\} & \text{if } z_e(\theta) = 0. \end{cases}$$

Moreover, the evolution of the queues can be characterized by the following equation.

$$\frac{\mathrm{d}z_e(\theta)}{\mathrm{d}\theta} = \begin{cases} f_e^+(\theta) - \nu_e & \text{if } z_e(\theta) > 0, \\ \max\{f_e^+(\theta) - \nu_e, 0\} & \text{if } z_e(\theta) = 0. \end{cases} \tag{1}$$

A *flow over time* is a collection of edge inflow rates $(f_e^+)_{e\in E}$ that satisfy the following flow conservation constraints for all vertices $V\backslash\{t\}$ and for almost all $\theta \geq 0$. For a vertex $v \in V$, define $\delta^+(v) = \{w \in V : (v,w) \in E\}$ and $\delta^-(v) = \{u \in V : (u,v) \in E\}$.

$$\sum_{e\in\delta^+(v)} f_e^+(\theta) - \sum_{e\in\delta^-(v)} f_e^-(\theta) = \begin{cases} u_0 & \text{if } v = s, \\ 0 & \text{if } v \neq s,t. \end{cases} \tag{2}$$

Finally, for a time θ, we define $F_e^+(\theta) = \int_0^\theta f_e^+(\xi)\mathrm{d}\xi$ and $F_e^-(\theta) = \int_0^\theta f_e^-(\xi)\mathrm{d}\xi$.

2.1. Optimal Flows over Time

In a directed graph $G = (V,E)$ with edge capacities ν_e and source and sink $s,t \in V$, a *static flow* is a function $\boldsymbol{f} : E \to \mathbb{R}_{\geq 0}$ of flow values f_e that satisfies $f_e \leq \nu_e$ for all $e \in E$ and the following flow conservation constraints:

$$\sum_{e\in\delta^+(v)} f_e - \sum_{e\in\delta^-(v)} f_e = 0 \quad \text{for all } v \neq s,t.$$

The size of such a flow $\boldsymbol{f}$ is denoted $|\boldsymbol{f}| = \sum_{e\in\delta^+(s)} f_e$. Because we have an inflow of u_0 in our model, we restrict the size of the flow to be at most this quantity, that is, $|\boldsymbol{f}| \leq u_0$.[3] If G is acyclic and $\mathcal{P}$ denotes the set of all s,t paths, a static flow $\boldsymbol{f}$ can be decomposed into path flows $(f_p)_{p\in\mathcal{P}}$ such that $f_e = \sum_{p\in\mathcal{P}:e\in p} f_p$ (Ahuja et al. [2]).

In the *maximum flow over time* problem (Ford and Fulkerson [9, 10]) with throughput objective, a time horizon T is given, and the objective is to maximize the amount of flow that arrives at t by time T. An optimal solution can be obtained by computing a static flow $\hat{\boldsymbol{f}}$ that solves the following linear program (Ford and Fulkerson [9]).

$$\begin{aligned} \max \quad & T\,|\boldsymbol{f}| - \sum_{e\in E} \tau_e f_e \\ \text{s.t.} \quad & 0 \leq f_e \leq \nu_e, \\ & |\boldsymbol{f}| \leq u_0. \end{aligned} \tag{3}$$

This solution can be decomposed into a path decomposition $\mathcal{P}$ such that flow enters every path $p \in \mathcal{P}$ at rate $\hat{f}_p$ until time $T - \tau_p$, where $\tau_p = \sum_{e\in p} \tau_e$ is the total travel time of the path without queues. Such a flow pattern is called a *temporally repeated flow*, and $\hat{\boldsymbol{f}}$ is called its *underlying static flow*. We define the flow rate or inflow of this temporally repeated flow as $|\hat{\boldsymbol{f}}|$.

For the makespan objective, we are given an amount of flow M, and a *quickest flow* is a flow over time that minimizes the time at which all flow arrives to t. A quickest flow can also be obtained by a temporally repeated flow (see, e.g., Bhaskar et al. [3]), whose underlying static flow is the static flow $\hat{f}$ that solves the following minimization problem in which T is now a variable.

$$\begin{aligned} \min \quad & T \\ \text{s.t.} \quad & T\,|f| - \sum_{e\in E} \tau_e f_e \geq M, \\ & 0 \leq f_e \leq \nu_e, \\ & |f| \leq u_0. \end{aligned} \tag{4}$$

We define T_{OPT} as the optimal value of this problem. Note that, even though this is a nonlinear problem, its solution can be found with a binary search. First guess a time T and solve (3). Decrease T if the objective function value exceeds M; otherwise, increase it. The minimum value of T such that the maximum flow over time with time horizon T routes M units of flow is, thus, the optimal solution that we denote by T_{OPT}.

Hence, there is a quickest flow that is a temporally repeated flow. We also refer to a quickest flow as an optimal flow over time $\hat{f}$. Finally, throughout the paper, we refer to the inflow or flow rate of this quickest flow over time $\hat{f}$ as its size $|\hat{f}|$, and we denote it by u_{OPT}. Therefore, u_{OPT} is equal to the inflow rate of this quickest flow for the limit of θ to zero. Note, in particular, that $u_{OPT} \leq u_0$.

An *earliest arrival flow* is a flow over time that maximizes the amount of flow that arrives at t by time θ for all $\theta \leq T$. An interesting fact is that such a flow always exists (Gale [11]), which justifies the binary search procedure. We refer the interested reader to the survey by Skutella [28] for more details regarding earliest arrival flows.

2.2. Equilibrium Flows

In our definitions, we follow the refined notion of dynamic equilibria from Cominetti et al. [5]. An equilibrium flow is a flow over time such that no flow particle can choose another route and arrive earlier at t given the fixed flow pattern of all other flow particles. More formally, consider a particle departing from s at time θ. We denote by $\ell_v(\theta)$ the earliest time at which this particle can arrive at node v.[4] Hence, $\ell_s(\theta) = \theta$, and for all $v \neq s$, we have

$$\ell_v(\theta) = \min_{u:e=(u,v)\in E} T_e(\ell_u(\theta)).$$

For any time θ, these labels induce a *dynamic shortest path network* G_θ with edge set

$$E'_\theta = \{e = (u,v) \in E : \ell_v(\theta) = T_e(\ell_u(\theta))\}.$$

The edges in E'_θ are called the *active* edges at time θ. We also define the set of edges that have a queue at time θ as $E^*_\theta = \{e = (u,v) \in E : z_e(\ell_u(\theta)) > 0\}$.

A feasible flow over time is called a dynamic equilibrium if and only if, for all $e = (v,w) \in E$ and almost all $\theta \in \mathbb{R}_{\geq 0}$, we have $f^+_e(\ell_v(\theta)) > 0 \Rightarrow e \in E'_\theta$. In other words, in a dynamic equilibrium flow is sent along shortest paths.

Cominetti et al. [5] prove that, for a dynamic equilibrium, we can equivalently write the sets of active edges and edges with queue as the simpler expressions

$$\begin{aligned} E'_\theta &= \{e = (u,v) \in E : \ell_v(\theta) \geq \ell_u(\theta) + \tau_e\}, \text{ and} \\ E^*_\theta &= \{e = (u,v) \in E : \ell_v(\theta) > \ell_u(\theta) + \tau_e\}. \end{aligned}$$

From this, it is immediate that $E^*_\theta \subseteq E'_\theta$.

It turns out that an equivalent characterization of a dynamic equilibrium is given by the condition that, for each $e = (v,w) \in E$ and all θ, we have $F^+_e(\ell_v(\theta)) = F^-_e(\ell_w(\theta))$ (Cominetti et al. [5]). It is convenient to define the *cumulative flow* induced by an equilibrium $\boldsymbol{f}$ on an edge $e = (v,w) \in E$ at time $\theta \in \mathbb{R}_{\geq 0}$ as

$$x_e(\theta) = F^+_e(\ell_v(\theta)) = F^-_e(\ell_w(\theta)).$$

Integrating the flow conservation constraints in Equation (2) over the interval $[0, \ell_v(\theta)]$ yields that the cumulative flow $\boldsymbol{x}(\theta)$ is a static s,t flow of value $u_0\theta$ for every $\theta \in \mathbb{R}_{\geq 0}$. Now define

$$x'_e(\theta) = \frac{\mathrm{d}x_e(\theta)}{\mathrm{d}\theta} = f^+_e(\ell_v(\theta))\ell'_v(\theta),$$

where

$$\ell'_v(\theta) = \frac{\mathrm{d}\ell_v(\theta)}{\mathrm{d}\theta} = \begin{cases} 1 & \text{if } v = s, \\ \min\limits_{(u,v)\in E} T'_{uv}(\ell_u(\theta))\ell'_u(\theta) & \text{if } v \neq s. \end{cases}$$

Observe that, because $\boldsymbol{x}(\theta)$ is a static flow of value $u_0\theta$, for almost all $\theta \in \mathbb{R}_{\geq 0}$, $\boldsymbol{x}'(\theta) = (x'_e(\theta))_{e\in E}$ is a static s,t flow of value u_0, for which $x'_e(\theta) = 0$ for all $e \in E'_\theta$. $\boldsymbol{x}'(\theta)$ is called a *normalized thin flow with resetting*, and the following theorem states some important properties.

Theorem 1 (Characterization of Dynamic Equilibrium (Cominetti et al. [5], Koch and Skutella [18])). *Consider a dynamic equilibrium* $\boldsymbol{f}$ *and a time* θ *such that* $x'_e(\theta)$ *and* $\ell'_v(\theta)$ *exist for all* $e \in E$ *and* $v \in V$. *Then, the static flow* $\boldsymbol{x}'(\theta)$ *satisfies*

$$\ell'_w(\theta) \leq \ell'_v(\theta) \qquad \forall e = (v,w) \in E'_\theta \backslash E^*_\theta : x'_{vw}(\theta) = 0, \tag{5}$$

$$\ell'_w(\theta) = \max\left\{\ell'_v, \frac{x'_e(\theta)}{\nu_e}\right\} \qquad \forall e = (v,w) \in E'_\theta \backslash E^*_\theta : x'_{vw}(\theta) > 0, \tag{6}$$

$$\ell'_w(\theta) = \frac{x'_e(\theta)}{\nu_e} \qquad \forall e = (v,w) \in E^*_\theta. \tag{7}$$

Moreover, it turns out that, for a given pair (E'_θ, E^*_θ),[5] if $E^*_\theta \subseteq E'_\theta \subseteq E$, E'_θ is acyclic and for all $v \in V$, E'_θ contains an s,v path, then there always exists a pair $(\boldsymbol{\ell}', \boldsymbol{x}')$ that satisfies Equations (5)–(7) such that $\boldsymbol{x}' = (x'_e(\theta))_{e\in E}$ is a static s, t flow of value u_0 with support in E'_θ. Furthermore, the ℓ'_v labels are unique (Cominetti et al. [5]).

Therefore, the derivatives $\boldsymbol{\ell}' = (\ell'_v(\theta))_{v\in V}$ only change if the shortest path network changes or if the set of edges with positive queue changes. This can be used to prove that the shortest path labels are unique throughout the evolution of the dynamic equilibrium (Cominetti et al. [5]).[6] The dynamic equilibrium, thus, consists of a sequence of phases, in which the edge inflow rates and the dynamic shortest path network are constant during each phase. These phases last a positive amount of time, and one can show that phase transitions only happen when new paths enter the dynamic shortest path network or when queues deplete. The rate at which the lengths of the paths and the queues change within one phase are completely determined by the $\boldsymbol{\ell}'$ labels, and therefore, the length of each phase can be computed, integrating the derivatives, with the α-extension algorithm of Koch and Skutella [18].

To be more precise, fix a time θ, and let $(\boldsymbol{\ell}', \boldsymbol{x}')$ be a solution to Conditions (5)–(7). Then, for the pair (E'_θ, E^*_θ), there exists an $\alpha > 0$ such that, if one integrates the $\boldsymbol{\ell}'$ labels, all inactive edges remain inactive and positive queues remain positive. In other words, for all $\Delta \in [0, \alpha]$,

$$\ell_w(\theta) + \Delta\ell'_w - \ell_v(\theta) - \Delta\ell'_v \leq \tau_e, \text{ for all } e = (v,w) \notin E'_\theta,$$
$$\ell_w(\theta) + \Delta\ell'_w - \ell_v(\theta) - \Delta\ell'_v \geq \tau_e, \text{ for all } e = (v,w) \in E^*_\theta.$$

Note that if Equation (5) holds with strict inequality for an edge $e = (v,w)$, integrating $\boldsymbol{\ell}'$ makes e inactive immediately; that is, if $\ell_w(\theta) - \ell_v(\theta) = \tau_e$ and $\ell'_w - \ell'_v < 0$, then $\ell_w(\theta) + \Delta\ell'_w - \ell_v(\theta) - \Delta\ell'_v < \tau_e$ for any $\Delta > 0$. If this happens, $(\boldsymbol{\ell}', \boldsymbol{x}')$ is still a solution at time $\theta + \Delta$ because there are no conditions on inactive edges. Also, if $\ell'_v < x'_e/\nu_e$ for an edge e satisfying Equation (6), a queue starts to grow immediately after θ. This does not pose a problem either because, in this case $\ell'_w = x'_e/\nu_e$, so e also satisfies Equation (7). As a result, the derivatives in $[0, \alpha]$ are constant and equal to $\boldsymbol{\ell}'$, so the equilibrium can be extended to $[\theta, \theta + \alpha]$ by integration.

Taking the maximum possible value of α, the current phase lasts until time $\theta + \alpha$, and the same procedure can be iterated. Therefore, assuming that the dynamic equilibrium does not exhibit Zeno-type behavior—that is, that the sequence defined by the α-extension algorithm does not have accumulation points—we can enumerate all the phases as $0, 1, 2, \ldots$, where each phase i lasts from time θ_i to θ_{i+1}. Within the interval (θ_i, θ_{i+1}), the configuration (E'_θ, E^*_θ), the $\boldsymbol{\ell}'$ labels, and the static flow $\boldsymbol{x}'$ remain constant. Our main results hold, however, without this assumption of the absence of Zeno-type behavior in the dynamic equilibrium.

3. The Price of Anarchy

In this section, we present our main result. First, in Section 3.1, we prove the upper bound on the price of anarchy in terms of the ratio between the inflow rate of the equilibrium u_{EQ} and the inflow rate of the optimal flow u_{OPT} and obtain as a corollary that, if $u_{EQ} = u_{OPT}$, then the price of anarchy is at most $\mathrm{e}/(\mathrm{e}-1)$. Later, in Section 3.2, we present a family of instances that match this upper bound.

3.1. Upper Bound

For a single-source, single-sink network G with inflow u_0 and a total amount of flow M, denote by T_{OPT} the time the quickest flow takes to route the M units from s to t. Denote the inflow rate of the quickest flow over time by u_{OPT}. For the dynamic equilibrium with inflow rate $u_{EQ} = u_0$, denote by $\hat{\theta}$ the first time at which M flow units have departed from the source s, that is, $\hat{\theta} = M/u_{EQ}$. Thus, because dynamic equilibria satisfy FIFO (Koch and

Skutella [18]), the time at which M units of flow have arrived at the sink t is $\ell_t(\hat{\theta}) = T_{EQ}$. Our result about the price of anarchy is the following.

Theorem 2. *Let m_G be the maximum number of nodes in a simple s,t path in G. It holds that*

$$T_{EQ} \cdot \left(1 - \frac{u_{EQ}}{u_{OPT}} \cdot \frac{1}{\mathrm{e}}\right) \le T_{EQ} \cdot \left(1 - \frac{u_{EQ}}{u_{OPT}}\left(1 - \frac{1}{m_G}\right)^{m_G}\right) \le T_{OPT}.$$

Corollary 1. *If $u_{EQ} = u_{OPT}$, then*

$$T_{EQ} \le T_{OPT} \cdot \left(1 - \left(1 - \frac{1}{m_G}\right)^{m_G}\right)^{-1} \le T_{OPT} \cdot \frac{\mathrm{e}}{\mathrm{e}-1}.$$

The corollary as well as the first inequality in the theorem are direct. To prove the second, we prove two main lemmata that together form the heart of the proof of Theorem 2. The first lemma relates the completion time of the optimal and equilibrium flow. It assumes the inflow rate of the optimum and equilibrium flow are equal.

Lemma 1. *The completion time of the optimal flow T_{OPT} and of the equilibrium T_{EQ} are related as follows:*

$$T_{EQ} - T_{OPT} \le \frac{1}{u_{OPT}} \sum_{e=(v,w)\in E} z_e(\ell_v(\hat{\theta})). \tag{8}$$

Proof. Consider a path decomposition $\mathcal{P}$ of the optimal flow. From Linear Program (3), it follows that $M = u_{OPT}T_{OPT} - \sum_{p\in\mathcal{P}} \hat{f}_p \tau_p$, where $\tau_p = \sum_{e\in p} \tau_e$. Moreover, from the equilibrium flow, we know $M = u_{EQ}\hat{\theta}$. Therefore,

$$u_{OPT}T_{OPT} - u_{EQ}\hat{\theta} = \sum_{p\in\mathcal{P}} \hat{f}_p \tau_p. \tag{9}$$

We rewrite the right-hand side as follows. Note that, for an edge $e = (v,w) \in E$, $\ell_w(\theta) \le \ell_v(\theta) + z_e(\ell_v(\theta))/\nu_e + \tau_e$, and hence, $\tau_e \ge \ell_w(\theta) - \ell_v(\theta) - z_e(\ell_v(\theta))/\nu_e$. Considering a path p, summing over all edges $e \in p$ gives $\tau_p \ge \ell_t(\theta) - \ell_s(\theta) - \sum_{e\in p} z_e(\ell_v(\theta))/\nu_e$. Applying this inequality for $\theta = \hat{\theta}$ to Equation (9) and using that $\ell_t(\hat{\theta}) = T_{EQ}$, yields

$$u_{OPT}T_{OPT} - u_{EQ}\hat{\theta} \ge \sum_{p\in\mathcal{P}} \hat{f}_p \left(T_{EQ} - \hat{\theta} - \sum_{e=(v,w)\in p} \frac{z_e(\ell_v(\hat{\theta}))}{\nu_e}\right). \tag{10}$$

By the definition of $\hat{f}$, $\sum_p \hat{f}_p = u_{OPT}$. Taking this out of the sum for the first two terms, we get

$$u_{OPT}T_{OPT} - u_{EQ}\hat{\theta} \ge u_{OPT}T_{EQ} - u_{OPT}\hat{\theta} - \sum_{p\in\mathcal{P}} \hat{f}_p \sum_{e\in p} \frac{z_e(\ell_v(\hat{\theta}))}{\nu_e},$$

and hence,

$$\begin{aligned} u_{OPT}(T_{EQ} - T_{OPT}) + \hat{\theta}(u_{EQ} - u_{OPT}) &\le \sum_{p\in\mathcal{P}} \hat{f}_p \sum_{e\in p} \frac{z_e(\ell_v(\hat{\theta}))}{\nu_e} \\ &= \sum_{e\in E} \hat{f}_e \frac{z_e(\ell_v(\hat{\theta}))}{\nu_e} \\ &\le \sum_{e\in E} z_e(\ell_v(\hat{\theta})). \end{aligned}$$

The equality follows by summing over all edges instead of all paths, and the last inequality is implied by $\hat{f}_e \le \nu_e$. The result follows from the fact that $u_{OPT} \le u_{EQ}$. □

To complete the proof of Theorem 2, it remains to bound the sum in the right-hand side of Equation (8). The following lemma does exactly this.

Lemma 2. *In the dynamic equilibrium, for all $\theta \ge 0$,*

$$\sum_{e=(v,w)\in E} z_e(\ell_v(\theta)) \le u_{EQ} \cdot \left(1 - \frac{1}{m_G}\right)^{m_G} \cdot (\ell_t(\theta) - \ell_t(0)).$$

We prove this lemma using two technical claims. For those claims, we need the following two auxiliary propositions. The first proposition states that, if flow is sent along an edge, the derivatives of the distance labels of both its vertices are positive. The second proposition is a bit technical.

Proposition 1. *In the dynamic equilibrium, for all $\theta \geq 0$ and all $e = (v, w) \in E'_\theta$ such that $x'_e(\theta) > 0$, we have that both $\ell'_v(\theta) > 0$ and $\ell'_w(\theta) > 0$.*

Proof. Consider some $\theta \geq 0$ and an edge $e = (v, w) \in E'_\theta$ with $x'_e(\theta) > 0$. If $e \in E^*_\theta$, then $\ell'_w(\theta) = x'_e/\nu_e > 0$ by Thin Flow Condition (7). If $e \in E'_\theta \backslash E^*_\theta$, then $\ell'_w(\theta) = \max\{\ell'_v(\theta), x'_e/\nu_e\} \geq x'_e/\nu_e > 0$ by Thin Flow Condition (6). The claim is proved for vertex w. Now, if $v = s$, then the result follows immediately because $\ell'_s(\theta) = 1$. On the other hand, if $v \neq s$, because $\boldsymbol{x}'(\theta)$ satisfies the flow conservation constraints, there must be an edge $e' = (u, v) \in E'_\theta$ with $x'_{e'}(\theta) > 0$. Following the same reasoning as before, we conclude that $\ell'_v(\theta) > 0$. □

Proposition 2. *For all $k \in \mathbb{N}\backslash\{0\}$ and all reals $y_0, y_1, \ldots, y_k > 0$ such that $y_0 = 1$, it holds that*

$$\sum_{i=1}^{k}\left(1 - \frac{y_{i-1}}{y_i}\right) \leq y_k \cdot \left(1 - \frac{1}{k+1}\right)^{k+1}.$$

Proof. Define $y_{k+1} := 1/(1 - 1/(k+1))^{k+1}$, and let us prove the equivalent statement $\sum_{i=1}^{k+1} y_{i-1}/y_i \geq k$. Indeed, by the arithmetic-geometric inequality,

$$\sum_{i=1}^{k+1}\frac{y_{i-1}}{y_i} \geq (k+1)\left(\prod_{i=1}^{k+1}\frac{y_{i-1}}{y_i}\right)^{1/(k+1)} = (k+1)\left(\frac{1}{y_{k+1}}\right)^{1/(k+1)} = (k+1)\frac{k}{k+1} = k. \quad \square$$

With these two propositions at hand, we can prove the two claims that imply Lemma 2.

Claim 1 states that, in the dynamic equilibrium, for all $\theta \geq 0$,

$$\sum_{e=(v,w)\in E} z_e(\ell_v(\theta)) \leq \int_0^\theta \sum_{e=(v,w)\in E'_\xi} x'_e(\xi)\left(1 - \frac{\ell'_v(\xi)}{\ell'_w(\xi)}\right)\mathrm{d}\xi.$$

In the following proof of claim 1, by an edge e we mean an edge $e = (v, w)$ unless indicated otherwise. We begin by writing the queue length in terms of its derivative by using Equation (1).

$$\sum_{e=(v,w)\in E} z_e(\ell_v(\theta)) = \sum_{e\in E}\int_0^\theta \frac{\mathrm{d}z_e(\ell_v(\xi))}{\mathrm{d}\xi}\,\mathbb{1}_{z_e(\ell_v(\xi))>0}\mathrm{d}\xi = \int_0^\theta \sum_{e\in E^*_\xi}\frac{\mathrm{d}z_e(\ell_v(\xi))}{\mathrm{d}\xi}\mathrm{d}\xi. \tag{11}$$

Denoting the flow underlying the dynamic equilibrium by f, for $e \in E^*_\xi$, we have $z'_e(\xi) = f^+_e(\xi) - \nu_e$ and $\ell'_w(\xi) = x'_e(\xi)/\nu_e$. Then, we can write

$$\begin{aligned}
\frac{\mathrm{d}z_e(\ell_v(\xi))}{\mathrm{d}\xi} &= z'_e(\ell_v(\xi))\ell'_v(\xi)\\
&= f^+_e(\ell_v(\xi))\ell'_v(\xi) - \nu_e\ell'_v(\xi)\\
&= x'_e(\xi) - \nu_e\ell'_v(\xi)\\
&\leq (x'_e(\xi) - \nu_e\ell'_v(\xi))\,\mathbb{1}_{x'_e(\xi)>0}\\
&= x'_e(\xi)\left(1 - \frac{\nu_e\ell'_v(\xi)}{x'_e(\xi)}\right)\mathbb{1}_{x'_e(\xi)>0}\\
&= x'_e(\xi)\left(1 - \frac{\ell'_v(\xi)}{\ell'_w(\xi)}\right)\mathbb{1}_{x'_e(\xi)>0}.
\end{aligned}$$

Here, the inequality follows because, if $x'_e(\xi) = 0$, the expression is negative. Because $x'_e(\xi) > 0$, we can take it out of the brackets in the next step. Plugging the preceding expression into Equation (11) yields

$$\sum_{e\in E} z_e(\ell_v(\theta)) \leq \int_0^\theta \sum_{e\in E^*_\xi} x'_e(\xi)\left(1 - \frac{\ell'_v(\xi)}{\ell'_w(\xi)}\right)\mathbb{1}_{x'_e(\xi)>0}\mathrm{d}\xi.$$

Note that, for edges $e = (v, w) \in E'_\xi \backslash E^*_\xi$, we have $\ell_w(\xi) = \ell_v(\xi) + z_e(\ell_v(\theta))/\nu_e + \tau_e$, so $\ell'_w(\xi) = \ell'_v(\xi) + z'_e(\ell_v(\theta))$ $\ell'_v(\theta)/\nu_e = \ell'_v(\xi) + 0$ almost everywhere. Using Lemma 1, this gives $1 - \ell'_v(\xi)/\ell'_w(\xi) = 0$ almost everywhere, and hence,

$$\sum_{e\in E} z_e(\ell_v(\theta)) \leq \int_0^\theta \sum_{e\in E'_\xi} x'_e(\xi)\left(1 - \frac{\ell'_v(\xi)}{\ell'_w(\xi)}\right)\mathbb{1}_{x'_e(\xi)>0}\mathrm{d}\xi.$$

The proof of claim 1 is complete. The next claim bounds the integral on the right-hand side.

Claim 2 states that, in the dynamic equilibrium, for almost all $\xi \geq 0$,

$$\sum_{e=(v,w)\in E'_\xi} x'_e(\xi)\left(1-\frac{\ell'_v(\xi)}{\ell'_w(\xi)}\right)\mathbb{1}_{x'_e(\xi)>0} \leq u_{EQ}\cdot\left(1-\frac{1}{m_G}\right)^{m_G}\cdot \ell'_t(\xi).$$

The proof of claim 2 begins by considering a path decomposition $\mathcal{P}$ of the dynamic equilibrium $\boldsymbol{x}'$ at time ξ. Note that paths in $\mathcal{P}$ only traverse edges $e\in E'_\xi$ having $x'_e(\xi)>0$, and note that edges $e\in E'_\xi$ either have $x'_e(\xi)=0$ or are contained in paths in $\mathcal{P}$. Then, we can rewrite the left-hand side as

$$\sum_{e\in E'_\xi} x'_e(\xi)\left(1-\frac{\ell'_v(\xi)}{\ell'_w(\xi)}\right)\mathbb{1}_{x'_e(\xi)>0} \leq \sum_{p\in\mathcal{P}} x'_p(\xi)\sum_{e\in p}\left(1-\frac{\ell'_v(\xi)}{\ell'_w(\xi)}\right).$$

For a path $p\in\mathcal{P}$, denote by $|p|$ the number of links in p. Because of Lemma 1 and the fact that $\ell'_s(\xi)=1$, we can apply Proposition 2 for each path p, taking $k=|p|$ and $y_i=\ell'_w(\xi)$ with w the head of the ith link in the path. Hence, we obtain that

$$\sum_{p\in\mathcal{P}} x'_p(\xi)\sum_{e\in p}\left(1-\frac{\ell'_v(\xi)}{\ell'_w(\xi)}\right) \leq \sum_{p\in\mathcal{P}} x'_p(\xi)\left(1-\frac{1}{|p|+1}\right)^{|p|+1}\cdot \ell'_t(\xi).$$

Now, note that $\left(1-\frac{1}{|p|+1}\right)^{|p|+1} \leq \left(1-\frac{1}{m_G}\right)^{m_G}$ for all $p\in\mathcal{P}$ because the expression is increasing in $|p|$ and $|p|+1\leq m_G$ for all $p\in\mathcal{P}$. Using that $\boldsymbol{x}'(\xi)$ is a flow of size u_{EQ}, we conclude the proof of the claim.

We now show how Lemma 2 follows from these claims.

Proof of Lemma 2. From claims 1 and 2, we see that

$$\sum_{e=(v,w)\in E} z_e(\ell_v(\theta)) \leq u_{EQ}\left(1-\frac{1}{m_G}\right)^{m_G}\int_0^\theta \ell'_t(\xi)\mathrm{d}\xi = u_{EQ}\left(1-\frac{1}{m_G}\right)^{m_G}(\ell_t(\theta)-\ell_t(0)),$$

where the last factor comes from the integration of $\ell'_t(\xi)$. □

Theorem 2 follows in a straightforward manner from the two main lemmata.

Proof of Theorem 2. Applying Lemma 2 for $\theta=\hat{\theta}$ to the bound of Lemma 1 yields

$$T_{EQ}-T_{OPT} \leq \frac{u_{EQ}}{u_{OPT}}\cdot\left(1-\frac{1}{m_G}\right)^{m_G}(\ell_t(\hat{\theta})-\ell_t(0)) \leq \frac{u_{EQ}}{u_{OPT}}\cdot\left(1-\frac{1}{m_G}\right)^{m_G}\ell_t(\hat{\theta}).$$

The result follows by writing $\ell_t(\hat{\theta})=T_{EQ}$ and rearranging this inequality to the desired expression. □

3.2. Tightness

Consider the family of instances described in Koch [17, section 7.4], in which it is proved that the price of anarchy of these instances is at most $\mathrm{e}/(\mathrm{e}-1)$. We prove tightness of our results by showing that, for a given choice of the edge capacities, the price of anarchy of these instances matches the bound of Theorem 2, which tends to $\mathrm{e}/(\mathrm{e}-1)$ in the limit.

For completeness, we describe the family of instances here, and they are illustrated in Figure 2. Fix the parameter $m\in\mathbb{N}$. Denote the capacity of edges e_i and e^i, respectively, by u_i and $u^i=\sum_{k=1}^{i}u_i$. Set the delay of e_i and e^i to $\tau_i=u^m(1-1/u^i)$ and $\tau^i=0$, respectively. The equilibrium inflow rate is u^m.

We consider the instance in which we set $u_1=1$ and $u^i=(m/(m-1))^{i-1}$. Note that this is a feasible choice as it is strictly increasing in i, and therefore, $u_i>0$. We set the total amount of flow to send through the network to $M=u^m$.

Figure 2. An illustration of the tight instance.

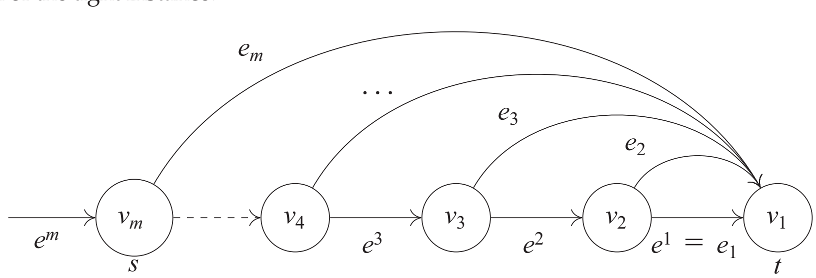

The following lemmata show that the price of anarchy for this instance tends to $e/(e-1)$ for $m \to \infty$.

Lemma 3. *The completion time of the equilibrium is* $T_{EQ} = u^m$.

Proof. Because $\tau^i = 0$ for all i, the total delay of the straight path is zero. Therefore, in the first phase of the equilibrium, all particles take the straight path, and we get

$$x'_e = \begin{cases} 0 & \text{for } e = e_i \quad i = 2, \ldots, m \\ u^m & \text{for } e = e^i \quad i = 1, \ldots, m \end{cases} \quad \text{and} \quad \ell'_{v_i} = \frac{u^m}{u^i} \quad \text{for } i = 1, \ldots, m.$$

This yields

$$\ell'_{v_1} - \ell'_{v_i} = \left(\frac{u^m}{u^1} - \frac{u^m}{u^i}\right) = \tau_i \text{ for all } i,$$

and therefore, the first phase lasts until time $\theta = 1$, when all paths enter the dynamic shortest path network.

Because the equilibrium inflow rate is u^m, we have $\hat{\theta} = M/u^m = 1$. Therefore,

$$T_{EQ} = \ell_t(\hat{\theta}) = \ell_t(1) = 1 + \tau_m = 1 + u^m\left(1 - \frac{1}{u^m}\right) = u^m. \quad \square$$

Lemma 4. *The completion time of the optimum flow is* $T_{OPT} = u^m(1 - (m-1)/(mu^m))$.

Proof. See Figure 3 for an illustration of $f_t^-(\theta)$, the inflow rate into t of the optimum flow as a function of θ. Note that, because M equals the area under this curve,

$$M = \sum_{i=1}^{m-1} (\tau_{i+1} - \tau_i)u^i + (T_{OPT} - \tau_m)u^m. \tag{12}$$

Now observe that, for all $i = 1, \ldots, m-1$,

$$(\tau_{i+1} - \tau_i)u^i = u^m\left(\frac{1}{u^i} - \frac{1}{u^{i+1}}\right)u^i = u^m\left(1 - \frac{u^i}{u^{i+1}}\right) = u^m\left(1 - \left(\frac{m}{m-1}\right)^{-1}\right) = \frac{u^m}{m}.$$

Replacing this in Equation (12), we conclude that

$$T_{OPT} = \frac{M}{u^m} - \frac{m-1}{m} + \tau_m = \frac{1}{m} + u^m\left(1 - \frac{1}{u^m}\right) = u^m\left(1 - \frac{m-1}{mu^m}\right). \quad \square$$

Lemma 5. *The price of anarchy of* $\left(1 - \left(1 - \frac{1}{m_G}\right)^{m_G}\right)^{-1}$ *is tight.*

Proof. From Lemmas 3 and 4, we see that the price of anarchy of this instance equals

$$\frac{T_{EQ}}{T_{OPT}} = \frac{u^m}{u^m\left(1 - \frac{m-1}{mu^m}\right)} = \frac{1}{1 - \frac{m-1}{mu^m}} = \frac{1}{1 - \left(\frac{m-1}{m}\right)^m}.$$

The proof is finished by noting that $m = m_G$. $\square$

It follows that the bound of $e/(e-1)$ is tight in the limit as well.

Figure 3. The inflow rate into t of the optimum flow for the tight instance.

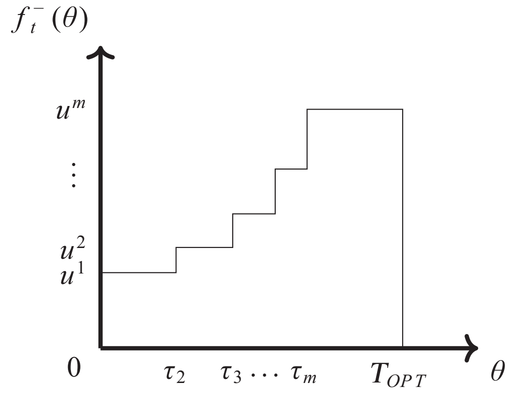

4. Monotonicity Conjecture

In this section, we discuss how the monotonicity conjecture would establish that our main result holds in general and discuss some difficulties in proving it that appears already in series-parallel graphs.

Conjecture 1: Consider a network G and two fixed inflow rates $u_1 < u_2$ with their corresponding dynamic equilibria in G and their corresponding makespans T^1_{EQ} and T^2_{EQ} for routing M units of flow. Then, $T^1_{EQ} \geq T^2_{EQ}$.

The following lemma demonstrates how the universal bound on the PoA would follow from Conjecture 1.

Lemma 6. *Suppose Conjecture* 1 *holds. Then, in any instance,*

$$T_{EQ} \leq \left(1-\left(1-\frac{1}{m_G}\right)^{m_G}\right)^{-1} \cdot T_{OPT} \leq \frac{\mathrm{e}}{\mathrm{e}-1} \cdot T_{OPT}.$$

Proof. Let a graph G and a total amount of flow M be given and consider a dynamic equilibrium with inflow rate $u_{EQ} > u_{OPT}$. By Conjecture 1, we know that $T_{EQ} \leq T'_{EQ}$, where T'_{EQ} is the makespan of the dynamic equilibrium in G with inflow rate $u_{EQ}' = u_{OPT}$. Because the network itself does not change, neither does the makespan of the optimal flow. Then, using Theorem 2, we obtain

$$T_{EQ} \leq T'_{EQ} \leq \left(1-\left(1-\frac{1}{m_G}\right)^{m_G}\right)^{-1} \cdot T_{OPT} \leq \frac{\mathrm{e}}{\mathrm{e}-1} \cdot T_{OPT}.$$

Even though there are good reasons to believe that Conjecture 1 is true, proving it would require a deeper understanding of the evolution of dynamic equilibria. Consider the following argument that illustrates where the difficulty resides. If, for a network, the number of phases is finite regardless of the inflow rate, the continuity of the derivatives of the equilibrium in terms of the inflow rate implies that we can partition $\mathbb{R}_{\geq 0}$ into intervals $[0, u_1), [u_1, u_2), \ldots$ such that, for all inflow rates $u \in [u_i, u_{i+1})$, the sequence of phases of the dynamic equilibrium, that is, the sequence of configurations (E'_θ, E^*_θ), is the same. Thus, if we can prove the monotonicity for a pair of inflow rates within any of these intervals, we conclude the monotonicity in general again by continuity.

Now, denote by $\boldsymbol{\ell}^u$ the labels of the dynamic equilibrium with inflow rate u. For a fixed total amount of flow M, the last particle arrives at the sink t at time $\ell^u_t(M/u)$, so Conjecture 1 states that $\ell^u_t(M/u)$ is a nonincreasing function of u. Denoting by $0 = \theta^u_0, \theta^u_1, \ldots, \theta^u_J = M/u$ the times of the phase transitions for the inflow rate u, we have that

$$\begin{aligned} \ell^u_t(M/u) &= \int_0^{M/u} \ell^{u\prime}_t(\theta)\mathrm{d}\theta \\ &= \sum_{j=0}^{J-1} (\theta^u_{j+1} - \theta^u_j) \cdot \ell^{u\prime}_t(\theta^u_j) \\ &= \sum_{j=0}^{J-1} (u\theta^u_{j+1} - u\theta^u_j) \cdot \frac{\ell^{u\prime}_t(\theta^u_j)}{u}. \end{aligned} \tag{13}$$

Here, the term $\ell^{u\prime}_t(\theta^u_j)/u$ is exactly the inverse of the *outflow* of the network, that is, the inflow into the sink t. At a time θ, this quantity depends only on u and the configuration (E'_θ, E^*_θ) and is known to be a decreasing function of u (Kaiser [16, lemma 18]. Thus, if we look within any given phase, the outflow of the network is, in fact, larger when the inflow is larger. However, Equation (13) also depends on the length of each phase, which, in turn, depends on the labels and not only their instantaneous derivatives.

In the next section, we prove the monotonicity conjecture for a special class of networks that we call series composition of parallel paths. Essentially, we complete the previous argument by proving that, for this class of networks, $\ell^{u\prime}_t(\theta)/u$ is decreasing in θ and that, for each phase j, $u\theta^u_j$ is decreasing in u. These two properties strongly rely on the fact that, in these networks, all phase transitions correspond to the activation of new edges and none to queues depleting. A natural extension would be to prove that parallel composition preserves monotonicity, but this does not seem to be as straightforward. The main reason for this stems from the fact that small graphs in the class of series-parallel graphs can already exhibit quite intricate equilibrium behavior depending on the inflow rate.

In order to illustrate this, consider the graph known as the pulse gadget in Figure 4(a). It consists of three nodes, s, v, and t, and four edges $e_1 = (s, t)$, $e_2 = (s, v)$, $e_3 = (v, t)$, and $e_4 = (v, t)$. The edge capacities are $1/3$, $3/4$, $1/3$, and 1, respectively, and their transit times are 2, 0, 0, and 2. There are three s,t paths, namely, $P_1 = (e_1)$, $P_2 = (e_2, e_3)$, and $P_3 = (e_2, e_4)$.

Figure 4. An example illustrating the intricacies involved in proving the monotonicity conjecture.

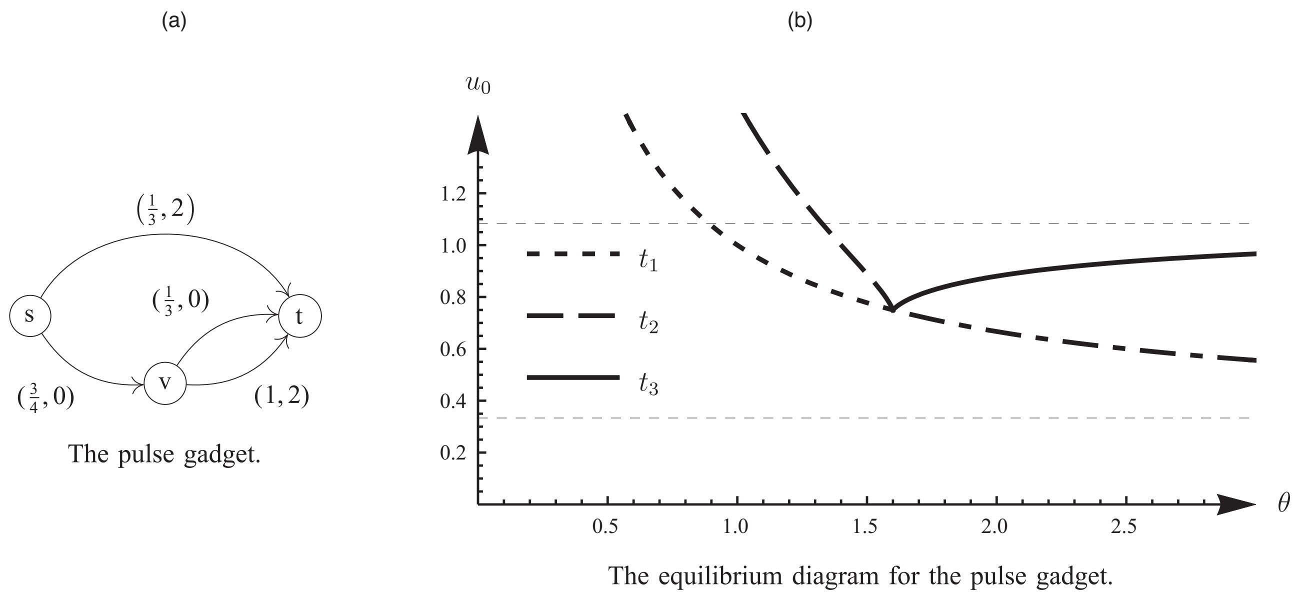

In this rather uncomplicated series-parallel graph, the phases of the dynamic equilibrium depend on the inflow rate in a nontrivial way. In the first phase, all flow is sent along path P_2. It turns out that three events can occur that cause a new phase of the equilibrium. First, it is possible that edge e_1 enters the shortest path network. Second, edge e_4 might enter the shortest path network. Finally, a queue that has grown on edge e_2 might deplete. Let us define t_1, t_2, and t_3 as the time at which each of these events happen, respectively, as a function of u_0. As an example, for an inflow rate of $u_0 = 1$, it holds that $t_1 = 1$, $t_2 = 7/5$, and $t_3 = 4$. The dependence of t_1, t_2, and t_3 on u_0 has been visualized in the diagram in Figure 4(b). It shows for every value of u_0 the value of t_1, t_2, and t_3 if these events happen at all for this given inflow rate.

As one can observe in this diagram, the length and existence of the phases depends in a nontrivial way on u_0. For some values of u_0, some events might never occur. For $u_0 \leq 1/3$, the thin flow never changes. For $1/3 < u_0 \leq 3/4$, edges e_1 and e_4 enter the shortest path network at the same time. For $3/4 < u_0 < 13/12$, they enter sequentially, and moreover, the edge on e_2 depletes. Finally, for $u_0 \geq 13/12$, they enter sequentially, but the queue on e_2 never vanishes.

Expressing the length of the phases depending on u_0 for a general graph is hard. Especially for inflow rates close to the bifurcation points in such an equilibrium diagram, it is not clear how to prove the monotonicity conjecture. These diagrams show the intricate subtleties one has to resolve in order to prove, for example, that parallel composition of two graphs preserves monotonicity.

5. Series Concatenation of Parallel-Paths Networks

In this section, we prove the monotonicity for series concatenation of parallel-paths networks. We say a graph is a parallel-paths network if all s,t paths are edge-disjoint. We say G is a series concatenation of parallel-paths networks if it results from taking parallel-paths networks $G_1, G_2, \ldots, G_{m-1}$ and identifying the sink of G_i with the source of G_{i+1} for each $i \leq m-2$. We denote by s_i the resulting node and by s_m the sink of G_{m-1}. We prove the following theorem.

Theorem 3. *Conjecture* 1 *holds when G is a series concatenation of parallel paths.*

We start by showing that paths whose internal nodes have exactly one incoming and one outgoing edge can be contracted to a single edge without changing the dynamic equilibrium. Therefore, we can focus on series concatenation of parallel-links networks, that is, in which all nodes of G are in $\{s_1, s_2, \ldots, s_m\}$.

Lemma 7. *Let $G_1 = (V_1, E_1)$ be a graph with a vertex $v \in V_1$ with exactly one incoming edge $e_1 = (u, v) \in E_1$ and exactly one outgoing edge $e_2 = (v, w) \in E_1$. Let $(f_e)_{e \in E_1}$ be a flow over time in G_1. Let $G_2 = (V_2, E_2)$ be the graph obtained by contracting e_1 and e_2 into one edge e_3, that is, removing the vertex v and the edges e_1 and e_2 and adding the new edge $e_3 = (u, w)$. Let $\nu_{e_3} = \min\{\nu_{e_1}, \nu_{e_2}\}$ and $\tau_{e_3} = \tau_{e_1} + \tau_{e_2}$. Let $(g_e)_{e \in E_2}$ be a flow over time in G_2. Then, if $g^+_{e_3}(\theta) = f^+_{e_1}(\theta)$ for all $\theta \geq 0$, then $g^-_{e_3}(\theta + \tau_{e_3}) = f^-_{e_2}(\theta + \tau_{e_1} + \tau_{e_2})$ for all $\theta \geq 0$.*

Proof. The result follows from using the flow conservation constraint in v and concatenating the definition of the outflows and the queue dynamics.

Consider first the case in which $\nu_{e_1} \le \nu_{e_2}$. By definition, we have that $f^-_{e_1}(\theta) \le \nu_{e_1}$ for all $\theta \ge 0$. By the flow conservation constraint in v, we have that $f^-_{e_1}(\theta) = f^+_{e_2}(\theta)$ for all $\theta \ge 0$. Thus, the inflow into e_2 is always smaller than its capacity, so $z_{e_2}(\theta) = 0$ for all $\theta \ge 0$. Hence, $f^-_{e_2}(\theta + \tau_{e_1} + \tau_{e_2}) = f^+_{e_2}(\theta + \tau_{e_1}) = f^-_{e_1}(\theta + \tau_{e_1})$. From the definition of $f^-_{e_1}$, we obtain that

$$f^-_{e_2}(\theta + \tau_{e_1} + \tau_{e_2}) = \begin{cases} \nu_{e_1} & \text{if } z_{e_1}(\theta) > 0, \\ \min\{f^+_{e_1}(\theta), \nu_{e_1}\} & \text{if } z_{e_1}(\theta) = 0. \end{cases}$$

Note that, if $\nu_{e_1} \le \nu_{e_2}$, then $\nu_{e_3} = \nu_{e_1}$, and thus, the queue z_{e_1} satisfies the same dynamics as z_{e_3}. Thus, $z_{e_1}(\theta) = z_{e_3}(\theta)$ for all θ, and then, $g^-_{e_3}(\theta + \tau_{e_3}) = f^-_{e_2}(\theta + \tau_{e_1} + \tau_{e_2})$, which concludes the proof for this case.

Assume now that $\nu_{e_1} > \nu_{e_2}$. From the definition of $f^-_{e_1}$ and $f^-_{e_2}$ and the fact that $f^+_{e_2}(\theta) = f^-_{e_1}(\theta)$, we obtain that

$$f^-_{e_2}(\theta + \tau_{e_1} + \tau_{e_2}) = \begin{cases} \nu_{e_2} & \text{if } z_{e_2}(\theta + \tau_{e_1}) > 0, \\ \min\{\nu_{e_1}, \nu_{e_2}\} & \text{if } z_{e_2}(\theta + \tau_{e_1}) = 0 \text{ and } z_{e_1}(\theta) > 0, \\ \min\{\min\{f^+_{e_1}(\theta), \nu_{e_1}\}, \nu_{e_2}\} & \text{if } z_{e_2}(\theta + \tau_{e_1}) = 0 \text{ and } z_{e_1}(\theta) = 0. \end{cases}$$

Now, if we define $\bar{z}_{e_3}(\theta) = z_{e_1}(\theta) + z_{e_2}(\theta + \tau_{e_1})$, we can simplify the preceding to

$$f^-_{e_2}(\theta + \tau_{e_1} + \tau_{e_2}) = \begin{cases} \nu_{e_2} & \text{if } \bar{z}_{e_3}(\theta) > 0, \\ \min\{f^+_{e_1}(\theta), \nu_{e_2}\} & \text{if } \bar{z}_{e_3}(\theta) = 0. \end{cases}$$

To complete the proof, we need to show that the $\bar{z}_{e_3}(\theta)$ we defined in G_1 has the same dynamics as $z_{e_3}(\theta)$ in G_2. From its definition, we see that the dynamics of $\bar{z}_{e_3}$ are

$$\frac{\mathrm{d}\bar{z}_{e_3}(\theta)}{\mathrm{d}\theta} = \begin{cases} f^+_{e_1}(\theta) - \nu_{e_1} + \nu_{e_1} - \nu_{e_2} & \text{if } z_{e_1}(\theta) > 0, z_{e_2}(\theta + \tau_{e_1}) > 0, \\ f^+_{e_1}(\theta) - \nu_{e_1} + [\nu_{e_1} - \nu_{e_2}]_+ & \text{if } z_{e_1}(\theta) > 0, z_{e_2}(\theta + \tau_{e_1}) = 0, \\ [f^+_{e_1}(\theta) - \nu_{e_1}]_+ + \min\left\{f^+_{e_1}(\theta), \nu_{e_1}\right\} - \nu_{e_2} & \text{if } z_{e_1}(\theta) = 0, z_{e_2}(\theta + \tau_{e_1}) > 0, \\ [f^+_{e_1}(\theta) - \nu_{e_1}]_+ + \left[\min\left\{f^+_{e_1}(\theta), \nu_{e_1}\right\} - \nu_{e_2}\right]_+ & \text{if } z_{e_1}(\theta) = 0, z_{e_2}(\theta + \tau_{e_1}) = 0, \end{cases}$$

where $[\cdot]_+ = \max\{\cdot, 0\}$. Because $\nu_{e_1} > \nu_{e_2}$, the first two cases give the same expression $f^+_{e_1}(\theta) - \nu_{e_2}$. In the latter two cases, a case distinction between $f^+_{e_1}(\theta) \ge \nu_{e_1}$ and $f^+_{e_1}(\theta) < \nu_{e_1}$ reveals that we can simplify the preceding to

$$\frac{\mathrm{d}\bar{z}_{e_3}(\theta)}{\mathrm{d}\theta} = \begin{cases} f^+_{e_1}(\theta) - \nu_{e_2} & \text{if } \bar{z}_{e_3}(\theta) > 0, \\ [f^+_{e_1}(\theta) - \nu_{e_2}]_+ & \text{if } \bar{z}_{e_3}(\theta) = 0. \end{cases}$$

Because z_{e_3} satisfies the exact same dynamics, we obtain that $\bar{z}_{e_3}(\theta) = z_{e_3}(\theta)$ for all θ if the inflow pattern is the same. Hence, we conclude that, also in this case, $g^-_{e_3}(\theta + \tau_{e_3}) = f^-_{e_2}(\theta + \tau_{e_1} + \tau_{e_2})$. □

Assume $s_1, s_2, \ldots, s_m$ are ordered by increasing distance to s. Because of Lemma 7, to study dynamic equilibria in G, we can assume without loss of generality (w.l.o.g.) that G is, in fact, a series concatenation of parallel-links networks, that is, that all edges in G are of the form (s_i, s_{i+1}) for some $1 \le i < m$. We also can assume w.l.o.g. that all pairs of parallel edges have different delays.[7] We say that a pair of nodes s_i, s_{i+1} and all edges of the form (s_i, s_{i+1}) form a *component* of G. We denote by $E'_{i,\theta}$ and $E^*_{i,\theta}$ the set of active edges, respectively, the set of edges with positive queue in this component at time θ. As explained in the previous section, to prove the monotonicity in a graph, we can restrict ourselves to the case in which the change in the inflow is so small that the sequence of phases is the same. We prove some facts about dynamic equilibria in G that we later use to prove Theorem 3 and conclude this section.

Lemma 8. *Consider the dynamic equilibrium with inflow u in G. For all $\theta \ge 0$ and all $1 \le i < m$, it holds that*

$$\ell'_{s_{i+1}}(\theta) = \max\left\{\ell'_{s_i}(\theta), \frac{u}{\sum_{e \in E'_{i,\theta}} \nu_e}\right\}, \tag{14}$$

and $\ell'_{s_1}(\theta) = 1$. Moreover, the set $E'_{i,\theta}$ is inclusion-wise increasing in θ, all phase transitions occur because a new edge becomes active, and there is a finite number of phases.

Proof. By definition, we have that $\ell'_{s_1}(\theta) = 1$ for all θ. Recall that

$$E'_\theta = \{e = (u,v) \in E : \ell_v(\theta) \geq \ell_u(\theta) + \tau_e\}, \text{ and}$$
$$E^*_\theta = \{e = (u,v) \in E : \ell_v(\theta) > \ell_u(\theta) + \tau_e\}.$$

This immediately implies that, for every component i and all times θ, $|E'_{i,\theta} \setminus E^*_{i,\theta}| \leq 1$, and that if $\ell'_{s_{i+1}}(\theta) \geq \ell'_{s_i}(\theta)$, which is a consequence of Equation (14), then edges do not leave E^*_θ, and a phase transition can occur only when a new edge enters E'_θ. If $\ell'_{s_{i+1}}(\theta) \geq \ell'_{s_i}(\theta)$ for all i and all θ, then we immediately obtain also that E'_θ is inclusion-wise increasing, and therefore, the number of phases is finite.

At time $\theta = 0$, there is only one edge per component that is active, the one with smallest delay, and no edge is resetting. This means that there is only one s,t path, so $x'_e(0) = u$ for all $e \in E'_0$. Therefore, Equation (6) immediately implies Equation (14). Before the first phase transition occurs, the set E'_θ remains unchanged, and the active edge in a component i such that $\ell'_{s_{i+1}}(0) - \ell_{s_i}{}'(0) > 0$ enters E^*_θ immediately after $\theta = 0$.

Consider now $\theta > 0$, and suppose that, for a component, it holds that

$$\ell'_{s_i}(\theta) \leq \frac{u}{\sum_{e \in E^*_{i,\theta}} \nu_e}. \tag{15}$$

It is easy to check that, if this is true, then $\ell'_{s_{i+1}} = \max\left\{\ell'_{s_i}, u/\left(\sum_{e \in E'_{i,\theta}} \nu_e\right)\right\}$, $x'_e = \ell'_{s_{i+1}} \cdot \nu_e$ for $e \in E^*_{i,\theta}$, and $x'_e = u - \sum_{e' \in E_{i,\theta}} x'_{e'}$ for at most one edge $e \in E'_{i,\theta} \setminus E^*_{i,\theta}$, satisfies the thin flow conditions in Equations (5)–(7) for the edges of the ith component.

Now, inductively, assume that a phase transition occurs at time $\theta > 0$, and assume that Equation (15) has held for all $\theta' < \theta$ and all components. By the previous argument, we also have that Equation (14) has held for all $\theta' < \theta$ and all components. Thus, for all times $\theta' < \theta$ and all components $i < m$, $\ell'_{s_{i+1}}(\theta') \geq \ell'_{s_i}(\theta')$, and therefore, the phase transition at time θ was caused by an edge becoming active. This, in particular, means that the set E^*_θ is the same as immediately before θ. Let i^* be the component for which a new edge became active. In components $j < i^*$, the sets $E'_{j,\theta}$ and $E^*_{j,\theta}$ are the same as immediately before θ, so the thin flow remains unchanged. Because of this, in component i^*, $\ell'_{s_{i^*}}(\theta)$ does not change, so Equation (15) still holds, and therefore, $\ell'_{s_{i^*+1}}(\theta) = \max\left\{\ell'_{s_{i^*}}(\theta), u/\left(\sum_{e \in E'_{i^*,\theta}} \nu_e\right)\right\}$. Because $E'_{i^*,\theta}$ grew by one element, $\ell'_{s_{i^*+1}}(\theta)$ decreased or remained constant. Consider now a component $j > i^*$, and assume inductively that $\ell'_{s_j}(\theta)$ decreased or remained constant. Clearly Equation (15) still holds, and then also Equation (14) holds in component j, and $\ell'_{s_{j+1}}(\theta)$ decreased or remained constant.

Right after the phase transition, it can happen that, in a component i, the edge in $E'_{i,\theta} \setminus E^*_{i,\theta}$ enters E^*_θ. This happens if $\ell'_{s_{i+1}}(\theta) > \ell'_{s_i}(\theta)$. But note that, in this case, by Equation (14), $u/\left(\sum_{e \in E'_{i,\theta}} \nu_e\right) > \ell'_{s_i}(\theta)$, and therefore, Equation (15) continues to hold. □

Lemma 9. *In a dynamic equilibrium in G, it holds that, for all* $1 \leq i \leq m$, $\ell'_{s_i}(\theta)$ *is a nonincreasing function of* θ. □

Proof. This is a direct consequence of inductively applying Lemma 8. For $i = 1$, $\ell'_{s_1}(\theta) = 1$ is nonincreasing in θ. Now, assume $\ell'_{s_i}(\theta)$ is nonincreasing in θ. Because $E'_{i,\theta}$ is inclusion-wise increasing in θ, $u/\left(\sum_{e \in E'_{i,\theta}} \nu_e\right)$ is nonincreasing in θ. Therefore, by Equation (14), $\ell'_{s_{i+1}}(\theta)$ is also nonincreasing in θ. □

Lemma 10. *Consider a dynamic equilibrium in G and a component i of G. If* $E'_{i,\theta}$ *is constant for* $\theta \in [\theta_1, \theta_2]$, *then* $\ell'_{s_{i+1}}(\theta) - \ell'_{s_i}(\theta)$ *is a nondecreasing function in the same interval.*

Proof. By Lemma 8,

$$\ell'_{s_{i+1}}(\theta) - \ell'_{s_i}(\theta) = \max\left\{\ell'_{s_i}(\theta), \frac{u}{\sum_{e \in E'_{i,\theta}} \nu_e}\right\} - \ell'_{s_i}(\theta),$$

which is a decreasing function of $\ell'_{s_i}(\theta)$. Because $u/\left(\sum_{e \in E'_{i,\theta}} \nu_e\right)$ is constant and, by Lemma 9, $\ell'_{s_i}(\theta)$ is nonincreasing, we conclude the statement of the lemma. □

Lemma 11. *Let* $u_0 < u_1$ *be two inflow rates in G such that, up until a total flow M has entered the network, the sequence of phases of the dynamic equilibrium, that is, the sequence of configurations* (E', E^*), *is the same for both inflows. For* $j \geq 1$, *denote by $\theta_j^{u_0}$ and $\theta_j^{u_1}$ the time of the jth phase transition when the inflow is u_0 and when the inflow is u_1, respectively. Then, for all $j \geq 1$, $u_1\theta_j^{u_1} \leq u_0\theta_j^{u_0}$.*

Proof. For an inflow $u \in [u_0, u_1]$, denote by $\boldsymbol{\ell}^u$ the corresponding time labels. In a given phase, we know, by Lemma 8, that, for each component i, there is a component $i' < i$ such that $\ell_{s_i}^u{}'(\theta) = u/\left(\sum_{e\in E'_{i',\theta}} \nu_e\right)$, or $\ell_{s_i}^u{}' = \ell_{s_1}^u{}' = 1$, and also that $\ell_{s_{i+1}}^u{}' - \ell_{s_i}^u{}' \geq 0$. Therefore, $(\ell_{s_{i+1}}^u{}' - \ell_{s_i}^u{}')/u$ is increasing in u for any fixed $\theta > 0$.

Now, inductively, consider the jth phase transition and assume that all previous phase transitions $j' < j$ satisfy that $u\theta_{j'}^u$ is nonincreasing in u. This transition happens when a new edge e becomes active. Let i be the component containing e, and let e'' be the last edge that became active in i before e. Let j'' be the phase transition when that happened. The transition times satisfy

$$\begin{aligned}\tau_e - \tau_{e''} &= \int_{\theta_{j''}^u}^{\theta_j^u} \ell_{s_{i+1}}^u{}'(\theta) - \ell_{s_i}^u{}'(\theta)\mathrm{d}\theta \\ &= \sum_{j''\leq j'<j} (\theta_{j'+1}^u - \theta_{j'}^u)(\ell_{s_{i+1}}^u{}'(\theta_{j'}^u) - \ell_{s_i}^u{}'(\theta_{j'}^u)) = \sum_{j''\leq j'<j} (u\theta_{j'+1}^u - u\theta_{j'}^u)(\ell_{s_{i+1}}^u{}'(\theta_{j'}^u) - \ell_{s_i}^u{}'(\theta_{j'}^u))/u\,.\end{aligned}$$

But, in this time interval, $E'_{i,\theta}$ stays constant, so by Lemma 10, $\ell_{s_{i+1}}^u{}'(\theta) - \ell_{s_i}^u{}'(\theta)$ is nondecreasing in θ. As we already proved, $(\ell_{s_{i+1}}^u{}'(\theta_{j'}^u) - \ell_{s_i}^u{}'(\theta_{j'}^u))/u$ is increasing in u, so if $u\theta_{j''}^u$ increased with u, then the right-hand side would become strictly larger than $\tau_e - \tau_{e''}$, causing a contradiction. □

We are now ready to prove Theorem 3.

Proof of Theorem 3. By Lemma 8, there are finitely many phases, so it suffices to prove the monotonicity for any pair of inflow rates $u_0 < u_1$ such that the sequence of phases of the dynamic equilibrium is the same. Now, if we send a total amount of flow of M and we denote by $\boldsymbol{\ell}^u$ the time labels for an inflow of $u \in [u_0, u_1]$, the last particle arrives at the sink t at time $\ell_t^u(M/u)$. Denoting by θ_j^u the time of the jth phase transition, for which we take the last to be $\theta_J^u = M/u$, recall from Equation (13) that

$$\ell_t^u(M/u) = \sum_{j=0}^{J-1} (u\theta_{j+1}^u - u\theta_j^u) \cdot \frac{\ell_t^u{}'(\theta_j^u)}{u}\,.$$

By Lemma 8, we have that $\ell_t^{u\,\prime}(\theta)/u$ is decreasing in u,and, by Lemma 9, that it is decreasing in θ. Finally, by Lemma 11, $u\theta_j^u$ is nonincreasing in u for every j, including $j = J$, as $u\theta_J^u = uM/u = M$. These facts combined imply that $\ell_t^u(M/u)$ is decreasing in u. □

Note that, by Lemma 7, a parallel-path network behaves exactly as a parallel-link network. Therefore, Theorem 2 implies that the price of anarchy for parallel-path networks is at most $(1-(1-1/2)^2)^{-1} = 4/3$. Furthermore, this bound is tight because Figure 2 provides a tight example for a two-link network.

6. Discussion

In this section, we provide some computational evidence to support the monotonicity conjecture, show alternative conjectures that would imply a universal bound on the price of anarchy, and extend our results to the total delay price of anarchy.

6.1. Computational Experiments

Monotonicity is an important basic property of the fluid queuing model that seems intuitive yet appears to be notoriously hard to prove. Although a formal proof still eludes us, we provide computational experiments to support the conjecture.

We consider a small graph outside of the class of graphs for which we prove the monotonicity conjecture in Theorem 3, namely, the graph for Braess' paradox visualized in Figure 1. This graph consists of the nodes s, u, v, and t connected by five edges $e_1 = (s, u)$, $e_2 = (s, v)$, $e_3 = (u, v)$, $e_4 = (u, t)$, and $e_5 = (v, t)$. We check the monotonicity in this graph computationally for every possible combination of the values $\nu_j \in \{0.5, 1\}$ and $\tau_j \in \{0, 1\}$, where ν_j and τ_j are the capacity respectively the delay of edge e_j for $j = 1, \ldots, 5$. This covers many different scenarios and equilibrium flows with bottlenecks at different locations in the graph. We check the conjecture for each of these 2^{10} different underlying graphs for the values $M \in \{2, 5\}$. For each combination of parameters, we compare the makespan for the values $u_{EQ} = 1$, $u_{EQ} = 1.01$, and $u_{EQ} = 1$, which are denoted by T_{EQ}^1, $T_{EQ}^{1.01}$, and $T_{EQ}^{1.1}$, respectively. Everything was computed using the Nash Flow Computation Tool developed by Max Zimmer [31].

It turns out, computationally, that, for every possible combination of parameters ν_j, τ_j, and M, it holds that $T_{EQ}^{1.1} \le T_{EQ}^{1.01} \le T_{EQ}^{1}$. This suggests that the monotonicity conjecture holds in the Braess graph, which lies beyond the class of series-parallel graphs.

6.2. Alternative Conjectures

Although we strongly believe that the monotonicity conjecture holds, this is not the only way to establish the desired price of anarchy result. Specifically, it seems plausible that the time it takes for a particle to travel from s to t in the dynamic equilibrium (not accounting for the time spent at the source) is less than the makespan of the optimal solution. In other words, $\ell_t(\theta) - \theta < T_{OPT}$ for all $0 \le \theta \le \hat{\theta}$. The next lemma establishes that, if this property holds, then our bound on the price of anarchy remains true.

Lemma 12. *Suppose that* $M > 0$ *and that, for all* $0 \le \theta \le \hat{\theta}$, *we have that* $\ell_t(\theta) - \theta < T_{OPT}$, *and then* $T_{EQ} \cdot \left(1 - \frac{u_{EQ}}{u_{OPT}} \cdot \frac{1}{e}\right) \le T_{EQ} \cdot \left(1 - \frac{u_{EQ}}{u_{OPT}}\left(1 - \frac{1}{m_G}\right)^{m_G}\right) \le T_{OPT}$.

Proof. Consider a network G in which $u_{EQ} > u_{OPT}$ with makespan T_{OPT} for the optimal flow and T_{EQ} for the dynamic equilibrium. Now, consider a network G' that is obtained from G by adding an edge e from s to t with $\nu_e = u_{EQ} - u_{OPT}$ and $\tau_e = T_{OPT} - \varepsilon$.

Because $\tau_e < T_{OPT}$, the optimum flow in G' sends a positive amount of flow along edge e, and it sends flow along the same paths as in G at the same rate for a shorter period of time (provided ε is small enough). This implies that $u_{OPT}' = u_{OPT} + \nu_e = u_{EQ}$, and moreover, $T_{OPT}' < T_{OPT}$. On the other hand, because of our assumption that $\ell_t(\theta) - \theta < T_{OPT}$ and the choice of τ_e, edge e is never active in the dynamic equilibrium (for sufficiently small ε). Therefore, $u_{EQ}' = u_{EQ}$ and $T_{EQ}' = T_{EQ}$, and we can write $T_{EQ}/T_{OPT} \le T_{EQ}'/T_{OPT}'$. The result follows from Theorem 2. □

Finally, we present a third conjecture that would imply the universal bound on the PoA. Indeed, suppose that, for any instance and $\hat{\theta} > 0$, we had that

$$\sum_{e=(v,w)\in E} \hat{f}_e \frac{z_e(\ell_v(\theta))}{\nu_e} < u_{EQ}\hat{\theta}. \tag{16}$$

Then, we could follow the beginning of the proof of Lemma 1, and we can write the following similar to Equation (10).

$$\begin{aligned} u_{OPT}T_{OPT} - u_{EQ}\hat{\theta} &\ge \sum_{p\in\mathcal{P}} \hat{f}_p\left(\ell_t(\hat{\theta}) - \hat{\theta} - \sum_{e\in p} \frac{z_e(\ell_v(\hat{\theta}))}{\nu_e}\right) \\ &= u_{OPT}\ell_t(\hat{\theta}) - u_{OPT}\hat{\theta} - \sum_{e=(v,w)\in E} \hat{f}_e \frac{z_e(\ell_v(\theta))}{\nu_e} \\ &> u_{OPT}\ell_t(\hat{\theta}) - u_{OPT}\hat{\theta} - u_{EQ}\hat{\theta}. \end{aligned}$$

Here, the last inequality is obtained by assuming the conjecture given by Equation (16). Then, by cancelling on both sides and dividing by u_{OPT}, we obtain that $T_{OPT} > \ell_t(\hat{\theta}) - \hat{\theta}$. Because the choice of $\hat{\theta}$ (or that of M) is arbitrary, we obtain that the previous conjecture holds, and therefore, the result follows from Lemma 12.

6.3. Total Delay Price of Anarchy

Our results regarding the price of anarchy with respect to the makespan objective can be extended to bounds on the price of anarchy with respect to the total delay. Given a time horizon T, the total delay is defined as the total travel time of all particles that arrive at the sink t before this time horizon, that is, $\int_0^T \theta f_t^-(\theta)\mathrm{d}\theta$.

Bhaskar et al. [3] prove that the total delay price of anarchy is at most twice the makespan price of anarchy for any temporal routing game. Assuming the monotonicity conjecture, the bounds we obtain in this paper not only hold for the makespan, but, in fact, for every particle. This implies that our bounds would also hold for the total delay price of anarchy with the same factor.

Acknowledgments

The authors thank Dario Frascaria, Marcus Kaiser, Neil Olver, Leon Sering, and Laura Vargas Koch for fruitful discussions as well as Max Zimmer and Leon Sering for their software tool that aided the computational experiments. The authors also thank two anonymous reviewers for helping improving the exposition of the paper. An earlier version of this paper was presented at the 20th ACM conference on Economics and Computation 2019 (Correa et al. [7]).

Endnotes

[1] To be precise, they proved that the distance labels are unique among the class of right-linear equilibria (Cominetti et al. [5, theorem 6]. Note that the equilibrium flow is not unique in general as, in a graph consisting of two identical parallel links with large enough capacity, the flow may split arbitrarily.

[2] We could also model this inflow as a capacity. Indeed, if we add an extra source at which all the flow M resides and add an edge from this extra source to s with capacity u_0, the situation remains unchanged.

[3] This makes the situation compatible when adding the extra source in the model.

[4] These labels are well-defined because of Cominetti et al. [5, equation (11)].

[5] In a given instance, (G, ν, τ, u_0).

[6] This is assuming right-continuity of the ℓ'-labels of the dynamic equilibrium or that there is no Zeno-type behavior.

[7] It is easy to see that, if two parallel links have the exact same delay, then they behave as one link with the same delay and capacity equal to the sum of the capacities.

References

[1] Anderson EJ, Philpott AB (1994) Optimisation of flows in networks over time. Kelly FP, ed. *Probability, Statistics and Optimisation* (Wiley, New York), 369–382.
[2] Ahuja RK, Magnanti TL, Orlin JB (1993) *Network Flows: Theory, Algorithms, and Applications* (Prentice-Hall, Upper Saddle River, NJ).
[3] Bhaskar U, Fleischer L, Anshelevich E (2015) A Stackelberg strategy for routing flow over time. *Games Econom. Behav.* 92:232–247.
[4] Cao Z, Chen B, Chen X, Wang C (2017) A network game of dynamic traffic. *Proc. 2017 ACM Conf. Econom. Comput.* (ACM, NY), 695–696.
[5] Cominetti R, Correa J, Larré O (2015) Dynamic equilibria in fluid queuing networks. *Oper. Res.* 63(1):21–34.
[6] Cominetti R, Correa J, Olver N (2017) Long term behavior of dynamic equilibria in fluid queuing networks. Eisenbrand F, Koenemann J, eds. *Proc. 19th Internat. Conf. Integer Programming Combin. Optim.* (Springer, Cham, Switzerland), 161–172.
[7] Correa J, Cristi A, Oosterwijk T (2019) On the price of anarchy for flows over time. *Proc. 2019 ACM Conf. Econom. Comput.* (ACM, NY), 559–577.
[8] Fleischer L, Tardos E (1998) Efficient continuous-time dynamic network flow algorithms. *Oper. Res. Lett.* 23(3–5):71–80.
[9] Ford LR, Fulkerson DR (1958) Constructing maximal dynamic flows from static flows. *Oper. Res.* 6(3):419–433.
[10] Ford LR, Fulkerson DR (1962) *Flows in Networks* (Princeton University Press, Princeton, NJ).
[11] Gale D (1959) Transient flows in networks. *Michigan Math. J.* 6(1):59–63.
[12] Garrido F (2017) Equilibrio dinamico en redes de transporte con multiples origenes y destinos. Unpublished master's thesis, Universidad de Chile, Santiago de Chile.
[13] Graf L, Harks T (2019) Dynamic flows with adaptive route choice. Lodi A, Nagarajan V, eds. *Proc. 20th Internat. Conf. Integer Programming Combin. Optim.* (Springer, Cham, Switzerland) 219–232.
[14] Hoefer M, Mirrokni V, Röglin H, Teng S (2011) Competitive routing over time. *Theoretical Comput. Sci.* 412(39):5420–5432.
[15] Ismaili A (2017) Routing games over time with FIFO policy. Devanur NR, Lu P, eds. *Proc. 13th Internat. Conf. Web Internet Econom.* (Springer, Cham, Switzerland), 266–280.
[16] Kaiser M (2021) Computation of dynamic equilibria in series-parallel networks. *Math. Oper. Res.*, ePub ahead of print September 24, https://doi.org/10.1287/moor.2020.1108.
[17] Koch R (2012) Routing games over time. Unpublished PhD dissertation, Technische Universität Berlin, Germany.
[18] Koch R, Skutella M (2011) Nash equilibria and the price of anarchy for flows over time. *Theory Comput. Systems* 49(1):71–97.
[19] Koutsoupias E, Papadimitriou C (1999) Worst-case equilibria. Meinel C, Tison S, eds. *Proc. 16th Annual Sympos. Theoretical Aspects Comput. Sci.* (Springer-Verlag, Berlin), 404–413.
[20] Macko M, Larson K, Steskal L (2010) Braess's paradox for flows over time. Kontogiannis S, Koutsoupias E, Spirakis PG, eds. *Proc. Third Internat. Conf. Algorithmic Game Theory* (Springer-Verlag, Berlin), 262–275.
[21] Meunier F, Wagner N (2010) Equilibrium results for dynamic congestion games. *Transportation Sci.* 44(4):524–536.
[22] Peeta S, Ziliaskopoulos AK (2001) Foundations of dynamic traffic assignment: The past, the present and the future. *Networks Spat. Econom.* 1(3–4):233–265.
[23] Ran B, Boyce D (1996) *Modeling Dynamic Transportation Networks* (Springer, Berlin).
[24] Scarsini M, Schröder M, Tomala T (2018) Dynamic atomic congestion games with seasonal flows. *Oper. Res.* 66(2):327–339.
[25] Sering L, Skutella M (2018) Multi-source multi-sink Nash flows over time. Borndörfer R, Storandt S, eds. *Proc. 18th Workshop Algorithmic Approaches Transportation Model. Optim. Systems*, vol. 12 (Schloss Dagstuhl, Leibniz-Zentrum für Informatik, Dagstuhl), 1–20.
[26] Sering L, Vargas Koch L (2019) Nash flows over time with spillback. Chan TM, ed. *Proc. 30th Annual ACM-SIAM Sympos. Discrete Algorithms* (SIAM, Philadelphia), 935–945.
[27] Skutella M (2006) Personal communication.
[28] Skutella M (2009) An introduction to network flows over time. Cook W, Lovász L, and Vygen J, eds. *Research Trends in Combinatorial Optimization* (Springer, Berlin), 451–482.
[29] Vickrey WS (1969) Congestion theory and transport investment. *Amer. Econom. Rev.* 59(2):251–260.
[30] Yagar S (1971) Dynamic traffic assignment by individual path minimization and queuing. *Transportation Res.* 5(3):179–196.
[31] Zimmer M (2017) Nash flows computation. GitHub repository. Accessed September 6, 2020, https://github.com/zimmer-m/NashFlowComputation.