

On the Price of Anarchy for flows over time

Citation for published version (APA):

Correa, J., Cristi, A., & Oosterwijk, T. (2019). On the Price of Anarchy for flows over time. In *EC '19 Proceedings of the 2019 ACM Conference on Economics and Computation* (pp. 559-577). ACM New York. <https://doi.org/10.1145/3328526.3329593>

Document status and date:

Published: 01/07/2019

DOI:

[10.1145/3328526.3329593](https://doi.org/10.1145/3328526.3329593)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

On the Price of Anarchy for flows over time

JOSÉ CORREA and ANDRÉS CRISTI, Universidad de Chile, Chile
TIM OOSTERWIJK, Maastricht University, Netherlands

Dynamic network flows, or network flows over time, constitute an important model for real-world situations where 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 exactly $e/(e-1) \approx 1.582$. This significantly extends a result by Bhaskar, Fleischer, and Anshelevich (SODA 2011). Furthermore, our methods allow to determine that the Price of Anarchy in parallel-link 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 $e/(e-1)$.

CCS Concepts: • **Theory of computation** → **Algorithmic game theory**; **Quality of equilibria**; **Network games**; *Routing and network design problems*; • **Mathematics of computing** → *Network flows*; • **Networks** → *Network dynamics*.

Additional Key Words and Phrases: flows over time, price of anarchy, dynamic equilibrium

ACM Reference Format:

José Correa, Andrés Cristi, and Tim Oosterwijk. 2019. On the Price of Anarchy for flows over time. In *ACM EC '19: ACM Conference on Economics and Computation (EC '19), June 24–28, 2019, Phoenix, AZ, USA*. ACM, New York, NY, USA, 19 pages. <https://doi.org/10.1145/3328526.3329593>

1 INTRODUCTION

In the study of traffic in networks it is often crucial to take the underlying dynamical 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 where 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 non-negative delay τ_e and a capacity per time unit v_e . A continuous stream of particles is injected at a source $s \in V$, at constant rate u_0 , and travels towards 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 v_e and, after leaving the queue, move along the edge to reach its head after τ_e time units. The

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC '19, June 24–28, 2019, Phoenix, AZ, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6792-9/19/06...\$15.00

<https://doi.org/10.1145/3328526.3329593>

discrete version of the problem was initially studied from an optimization perspective. Indeed, Ford and Fulkerson [10, 11] considered a fluid queuing model and designed 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 timespan. Shortly after, Gale [12] showed the existence of a flow pattern that achieves this optimum simultaneously for all time horizons. These results were extended to continuous time by Anderson and Philpott [1] and Fleischer and Tardos [9]. 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 will be reached.

This *dynamic equilibrium* model was initially considered, in a very simple network, by Vickrey [30], and shortly after in the transportation science community [31]. Since then it has attracted much attention as a showcase model to understand the surprising behavior of dynamic routing games [23, 24]. In the last decade there have been significant efforts in understanding the structure and computational properties of dynamic equilibria in the fluid queuing model [3–8, 14, 16, 17, 19, 22, 25–27]. Meunier and Wagner [22] proved, 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 [19] gave an elegant characterization of the derivatives (w.r.t. time) of a dynamic equilibrium and thus proposed an algorithm to construct a dynamic equilibrium by concatenating static flows. Using this characterization, Cominetti, Correa and Larré [7] gave a constructive proof of existence of equilibria and proved they are essentially unique. 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. [5] established this property (on a slightly different atomic model that does not influence the result) for series-parallel networks, while Correa et al [8] established 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 3). 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 to the optimal solution. Our main result, which can be seen as an improvement upon a result of Bhaskar, Fleischer and Anshelevich [3, 4], establishes that if the monotonicity conjecture

holds for the dynamic equilibrium, then the Price of Anarchy, defined as the worst case ratio of the quality of an equilibrium to that of an optimal solution, is exactly $e/(e - 1)$.

1.1 The Price of Anarchy

The usual way of quantifying the inefficiency of selfish behavior is the Price of Anarchy (PoA). It is defined as the worst possible ratio between the quality of an optimal solution and the quality of an equilibrium [20]. In the context of fluid queuing networks there are two natural and related goals which 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 [12], 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 towards t and the time it takes in a dynamic equilibrium.

The first to study the PoA in this context were Koch and Skutella [18, 19] who proved 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-PoA and the makespan-PoA are 1. Interestingly, it has long been conjectured that the makespan-PoA is bounded by a small constant [29]. 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, Fleischer and Anshelevich [3, 4] studied this question from a mechanism design perspective and found 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 since reducing the capacities too much may also block good routes significantly, increasing the makespan of a dynamic equilibrium. More precisely, Bhaskar et al. 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 is to establish that the same bound still holds by doing this only for the inflow, i.e., leaving all capacities unchanged but only reducing the inflow.

Example 1. Consider the network in Fig. 1, where an edge e is labelled (v_e, τ_e) , and let the total flow M to be sent through the network be 2 and let $u_0 = 1$. We claim that the optimal flow will send $\frac{1}{2}$ units of flow both along the path (s, u, t) and along the path (s, v, t) until time 2. Therefore, the makespan of the optimal flow is 3. On the other hand, the equilibrium will first send 1 unit of flow along the path (s, u, v, t) from time 0 until time 1. Then, from time 1 to 2, it will send $\frac{1}{2}$ units of flow both along the path (s, u, v, t) and along the path (s, v, t) . Since a particle originating at s at time 2 will encounter a queue time of 1 on edges (s, u) and (v, t) , the makespan of the equilibrium

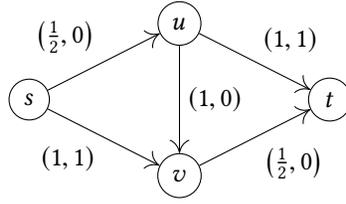


Fig. 1. An illustration of the network of Example 1.

is 4, and hence, the PoA of this instance is $\frac{4}{3}$. Note that if we set $v_{uv} = 0$ (as in [4]), the equilibrium in the modified network will do exactly the same as the optimal flow, and the new PoA is 1.

Finally, Cominetti, Correa and Olver [8] 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 v_e . This result readily implies that the Price of Anarchy converges to 1 as the amount of flow to be routed grows to infinity.

1.2 Our results

As mentioned above, our main result is to improve upon the result of Bhaskar, Fleischer and Anshelevich [3, 4] and show that the PoA is upper bounded by $e/(e-1)$ 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 since it potentially makes further progress (e.g. on multi-commodity 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 . The simplest algorithm to compute this quantity is that of Ford and Fulkerson [11] 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 $u_{OPT} = |f|$. As in the dynamic equilibrium particles are selfish, its inflow rate u_{EQ} always equals u_0 .

Similarly, letting T_{EQ} be the time it takes for the equilibrium to route the M units of flow,¹ our main result, which we prove in Section 3, can be stated in the following terms.

Theorem 2. *If $u_{EQ} = u_{OPT}$, then $T_{EQ} \leq \frac{e}{e-1} \cdot T_{OPT}$ and this is tight.*

Note that $u_{EQ} = u_0 \geq u_{OPT}$, therefore the missing case left by Theorem 2 is when $u_{EQ} > u_{OPT}$. Intuitively, this case should be easier. Indeed, for the theorem to hold in general it is enough to prove the monotonicity conjecture, which basically states that by decreasing the inflow, the makespan T_{EQ} increases. To formalize this conjecture, denote the makespan of the equilibrium with inflow rate u_i by T_{EQ}^i .

Conjecture 3 (Monotonicity conjecture). *Consider a network G and two fixed inflow rates $u_1 < u_2$ with their corresponding dynamic equilibria in G . Then $T_{EQ}^1 \geq T_{EQ}^2$.*

For the special and simpler case of parallel-link networks, it turns out that the monotonicity conjecture holds (cf. Lemma 16), immediately implying, by Theorem 2, that the PoA is bounded for these networks. Furthermore, for parallel-link networks we are able to obtain an improved bound

¹In Section 2.2 we denote by $\hat{\theta} = M/u_0$ the time this last particle enters s and by $\ell_t(\hat{\theta})$ the time at which it reaches t . Since the dynamic equilibrium satisfies FIFO, it is clear that $\ell_t(\hat{\theta}) = T_{EQ}$.

by refining the analysis. In fact, the argument can be extended to parallel-path networks. Indeed, we prove the following result.

Theorem 4. *In parallel-link networks, $T_{EQ} \leq \frac{4}{3} \cdot T_{OPT}$ and this is tight.*

The proofs proceed 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, in both the general case and the parallel-link case, the formula can be used to upper bound the sum of the queues at equilibrium by an appropriately small constant times $u_{EQ} \cdot T_{EQ}$.

1.3 Further related literature

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

Hoefer et al. [15] 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 [16] considers a similar atomic model with the 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 has been some recent efforts to carry over the results to more general multi-commodity networks. In particular, Garrido [13] was able to extend some of the results for dynamic equilibria to the case of multiple sinks, while Sering and Skutella [26] do it for the much more involved multi-source multi-sink case. However, we are still lacking a good understanding of the general multi-commodity case.

As mentioned earlier, the issue of bounded queues was studied by Cao et al [5] who proved 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 [8] 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, Larson, and Steskal [21] study new types of Braess's paradox appearing in the dynamic equilibrium.

Some very recent work considers other aspects of the problem. In particular, Sering and Vargas Koch [27] consider spillback effects, which is the study how an a priori bound on the amount of flow that can be waiting on a queue affects the equilibrium behavior. Graf and Harks [14] 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, Schröder, and Tomala [25] 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 [19] and that of Cominetti et al [7], 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 [19] by parametric flow techniques, and in general it can be written as a non-linear complementarity problem [6, 18]. Very recently, Kaiser [17] noted that the problem is actually a linear complementarity problem and that it can be solved efficiently in series-parallel networks.

2 THE MODEL

Let $G = (V, E)$ be a directed graph, where each edge $e \in E$ has a positive capacity v_e and a non-negative 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 , where 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 $f_e^+ : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ be the function associated with an edge $e \in E$ that maps a non-negative time θ to the inflow rate into e at time θ . In case the inflow rate $f_e^+(\theta)$ exceeds the edge capacity v_e , a queue will grow at the tail of the edge at rate $f_e^+(\theta) - v_e$. The queue mass at time θ is denoted by $z_e(\theta)$, and if $f_e^+(\theta) < v_e$, the queue will deplete at a rate equal to $f_e^+(\theta) - v_e$, until the inflow rate changes again or until $z_e = 0$. Therefore, a particle that enters edge e at time θ will wait in the queue $z_e(\theta)/v_e$ units of time and subsequently travel across the edge, taking time τ_e . Hence, this particle has link exit time

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

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

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

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

$$\frac{dz_e(\theta)}{d\theta} = \begin{cases} f_e^+(\theta) - v_e & \text{if } z_e(\theta) > 0, \\ \max\{f_e^+(\theta) - v_e, 0\} & \text{if } z_e(\theta) = 0, \end{cases} \quad (1)$$

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 $v \in V \setminus \{t\}$ and for almost all $\theta \geq 0$.

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

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

2.1 Optimal flows over time

In a directed graph $G = (V, E)$ with edge capacities v_e and source and sink $s, t \in V$, a *static flow* is a function $f : E \rightarrow \mathbb{R}_{\geq 0}$ of flow values f_e that satisfies $f_e \leq v_e$ 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 f is denoted $|f| = \sum_{e \in \delta^+(s)} f_e$. Since we have an inflow of u_0 in our model, we restrict the size of the flow to be at most this quantity, i.e., $|f| \leq u_0$.³ If G is acyclic and \mathcal{P} denotes the set of all s, t -paths, a static flow f can be decomposed into path flows $(f_p)_{p \in \mathcal{P}}$ such that $f_e = \sum_{p \in \mathcal{P}: e \in p} f_p$ [2].

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

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

In the *maximum flow over time* problem [10, 11] 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{f} that solves the following linear program [11].

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

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{f} is called its *underlying static flow*. We define the flow rate or inflow of this temporally repeated flow as $|f|$.

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 . This can be found with a binary search. First guess a time T and solve the previous linear program. 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 which we denote by T_{OPT} . Hence, there is a quickest flow that is a temporally repeated flow. We will also refer to a quickest flow as an optimal flow over time \hat{f} . Finally, throughout the paper we will refer to the inflow or flow rate of this quickest flow over time \hat{f} as its size $|\hat{f}|$ and we will denote it by u_{OPT} . 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 [12], which justifies the binary search procedure above. Even though earliest arrival flows may not be temporally repeated flows, there is always one that is a *generalized temporally repeated flow*. We refer the interested reader to the survey by Skutella [28].

2.2 Equilibrium flows

In our definitions we follow the refined notion of dynamic equilibria from [7]. 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 . 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 $E^*_\theta = \{e = (u, v) \in E : z_e(\ell_u(\theta)) > 0\}$. Cominetti et al. [7] proved that we can equivalently write

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

so it is immediate that $E^*_\theta \subseteq E'_\theta$.

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.

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 $F_e^+(\ell_v(\theta)) = F_e^-(\ell_w(\theta))$ [7]. It will be convenient to define the *cumulative flow* induced by an equilibrium 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 Eq. (2) over the interval $[0, \ell_v(\theta)]$ yields that the cumulative flow $x(\theta)$ is a static s, t -flow of value $u_0\theta$ for every $\theta \in \mathbb{R}_{\geq 0}$. Now define

$$x'_e(\theta) = \frac{dx_e(\theta)}{d\theta} = f_e^+(\ell_v(\theta))\ell'_v(\theta),$$

where

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

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

Theorem 5 ([7, 19]). *Consider a dynamic equilibrium 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 $x'(\theta)$ satisfies:*

$$\ell'_w(\theta) \leq \ell'_v(\theta) \quad \forall e = (v, w) \in E'_\theta \setminus E_\theta^* : x'_{vw}(\theta) = 0, \quad (4)$$

$$\ell'_w(\theta) = \max \left\{ \ell'_v, \frac{x'_e(\theta)}{v_e} \right\} \quad \forall e = (v, w) \in E'_\theta \setminus E_\theta^* : x'_{vw}(\theta) > 0, \quad (5)$$

$$\ell'_w(\theta) = \frac{x'_e(\theta)}{v_e} \quad \forall e = (v, w) \in E_\theta^*. \quad (6)$$

Moreover, it turns out that for a given pair (E'_θ, E_θ^*) , there always exists a pair (ℓ', x') that satisfies conditions (4), (5) and (6), and such that $x' = (x'_e(\theta))_{e \in E}$ is a static s, t -flow of value u_0 with support in E'_θ . Furthermore the ℓ'_v labels are unique [7].

Therefore, the derivatives $\ell' = (\ell'_v(\theta))_{v \in V}$ 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 [7].⁴ The dynamic equilibrium thus consists of a sequence of phases, where 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 ℓ' labels, and therefore the length of each phase can be computed, integrating the derivatives, with the α -extension algorithm of Koch and Skutella [19].

To be more precise, fix a time θ and let (ℓ', x') be a solution to the conditions (4), (5) and (6). Then for the pair (E'_θ, E_θ^*) there exists an $\alpha > 0$ such that if one integrates the ℓ' labels, all inactive edges remain inactive and positive queues remain positive. In other words, for all $\Delta \in [0, \alpha]$,

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

⁴Assuming right-continuity of the ℓ' -labels of the dynamic equilibrium, or that there is no Zeno-type behavior.

Note that if (4) holds with strict inequality for an edge $e = (v, w)$, integrating ℓ' will make e inactive immediately, i.e., 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, (ℓ', x') is still a solution at time $\theta + \Delta$ because there are no conditions on inactive edges. Also, if $\ell'_v < x'_e/v_e$ for an edge e in condition (5), a queue will start to grow immediately after θ . This does not pose a problem either, since in this case $\ell'_w = x'_e/v_e$, so e also satisfies condition (6). As a result, the derivatives in $[0, \alpha]$ are constant and equal to ℓ' , 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—i.e., that the sequence defined by the α -extension algorithm does not have accumulation points—we can enumerate all the phases as $0, 1, 2, \dots$ 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 ℓ' labels, and the static flow x' , remain constant. For ease of exposition we make this assumption in this paper, noting however that it is only used in the proof of Proposition 7, which in turn is used in Claim 11. Nevertheless, our main results hold even without this non-Zeno-type behavior assumption. In the full version of the paper we present a similar derivation (that implicitly includes a weaker version of Claim 11) that does not use Proposition 7 and that is enough to establish our main theorems.

3 THE PRICE OF ANARCHY

In this section we present our main result. 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 , i.e., $\hat{\theta} = M/u_{EQ}$. Thus, since dynamic equilibria satisfy FIFO [19], 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. *If $u_{EQ} = u_{OPT}$, then $T_{EQ} \leq \frac{e}{e-1} \cdot T_{OPT}$ and this is tight.*

To prove Theorem 2, we first establish the following three auxiliary propositions. Later, we build upon these results to upper bound $T_{EQ} - T_{OPT}$ in terms of the sum of the queues at equilibrium and find a formula to evaluate this quantity. The first proposition states that if flow is sent along an edge, the derivatives of the distance labels of both its vertices are positive.

Proposition 6. *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 $\ell'_w(\theta) = x'_e/v_e > 0$ by thin flow condition (6). If $e \in E'_\theta \setminus E_\theta^*$, then $\ell'_w(\theta) = \max\{\ell'_v(\theta), x'_e/v_e\} \geq x'_e/v_e > 0$ by thin flow condition (5). The claim is proved for vertex w . Now, if $v = s$ then the result follows immediately since $\ell'_s(\theta) = 1$. On the other hand if $v \neq s$, because $x'(\theta)$ 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 $\ell'_v(\theta) > 0$. \square

The second proposition states that if an edge has a queue, then it has strictly positive flow, and hence, strictly positive derivatives of its distance labels.

Proposition 7. *In the dynamic equilibrium, for all $\theta \geq 0$ such that $e = (v, w) \in E_\theta^*$, we have $x'_e(\theta) > 0$, and hence, both $\ell'_v(\theta) > 0$ and $\ell'_w(\theta) > 0$.*

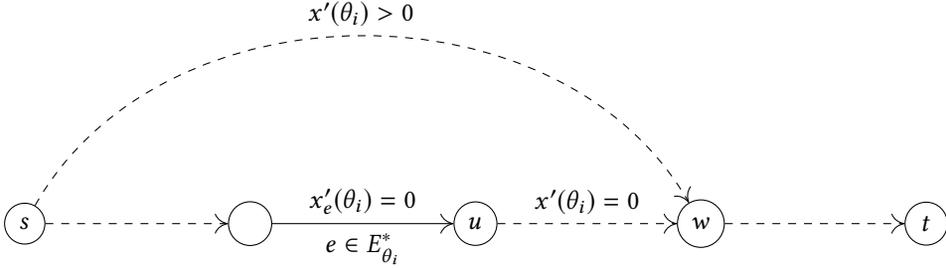


Fig. 2. An illustration to clarify the proof of Proposition 7.

PROOF. We prove the proposition by contradiction. We refer to Fig. 2 for an illustration of the proof.

Consider the earliest phase i of the dynamic equilibrium such that there is an edge e satisfying that for all times θ in phase i , $e \in E_{\theta}^*$ but $x'_e(\theta) = 0$.⁵ Because of our assumption that i is the earliest such phase, we can assume that $x'_e(\theta') > 0$ for θ' in phase $i - 1$. Indeed, otherwise $x'_e(\theta') = 0$ and the queue in e did not increase, so $e \in E_{\theta'}^*$, contradicting the assumption that i is the earliest.

Let θ_i be the starting time of phase i . Since $x'_e(\theta') > 0$ for θ' in phase $i - 1$, there exists some active⁶ s, t -path p containing e . Of course, p is still active by time θ_i .

Without loss of generality we can assume that e is the last edge (in the order induced by p) with the property of having a queue but not carrying flow in phase i , i.e. $e \in E_{\theta_i}^*$ and $x'_e(\theta_i) = 0$. Let $u \in V$ be the head of e . Since $x'_e(\theta_i) = 0$ and $e \in E_{\theta_i}^*$, thin flow condition (6) imposes that $\ell'_u(\theta_i) = 0$. Moreover, all vertices v coming after u along p that do not receive any flow would still satisfy $\ell'_v(\theta_i) = 0$. Thus consider w , the first vertex after e along p receiving some flow (possibly $w = t$). Note that this vertex does not receive flow from p as e does not carry flow. The latter implies that $\ell'_w(\theta_i) = 0$ by the thin flow conditions (4) or (6). On the other hand, as w receives flow from some edge outside p , the thin flow conditions (5) or (6) imply that $\ell'_w(\theta_i) > 0$. A contradiction.

Therefore, $x'_e(\theta_i) > 0$ and by Proposition 6, also $\ell'_v(\theta_i) > 0$ and $\ell'_w(\theta_i) > 0$. This extends to any θ in phase i because the derivatives are constant in a phase. \square

Finally, we will use the following proposition.

Proposition 8. $\frac{\log(x)}{x} \leq \frac{1}{e}$ for all $x > 0$.

PROOF. Consider the function $f(x) = \log(x)/x$ for $x > 0$. Then $f'(x) = (1 - \log(x))/x^2$, and clearly $f'(x) > 0$ for $x \in (0, e)$ and $f'(x) < 0$ for $x \in (e, +\infty)$. So $x = e$ is the global maximum of f and the lemma follows since $f(e) = 1/e$. \square

We are now ready to state the two main lemmata that together form the heart of the proof of Theorem 2. The following lemma relates the completion time of the optimal flow and the equilibrium flow. It assumes the inflow rate of the optimum flow and the equilibrium flow are equal.

Lemma 9. If $u_{OPT} = u_{EQ}$, the completion time of the optimal flow T_{OPT} and of the equilibrium T_{EQ} are related as follows.

$$T_{EQ} - T_{OPT} \leq \frac{1}{u_{EQ}} \sum_{e=(v,w) \in E} z_e(\ell'_v(\hat{\theta})). \quad (7)$$

⁵Note that x' is constant within a phase.

⁶By active we mean that all edges in p are in $E_{\theta'}$.

PROOF. Consider a path decomposition \mathcal{P} of the optimal flow. From the 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,

$$u_{OPT}T_{OPT} - u_{EQ} \hat{\theta} = \sum_{p \in \mathcal{P}} \hat{f}_p \tau_p. \quad (8)$$

We will rewrite the right-hand side as follows. Note that for an edge $e = (v, w) \in E$, $\ell_w(\theta) \leq \ell_v(\theta) + z_e(\ell_v(\theta))/v_e + \tau_e$ and hence $\tau_e \geq \ell_w(\theta) - \ell_v(\theta) - z_e(\ell_v(\theta))/v_e$. Consider a path p , summing over all edges $e \in p$ gives $\tau_p \geq \ell_t(\theta) - \ell_s(\theta) - \sum_{e \in p} z_e(\ell_v(\theta))/v_e$. Applying this inequality for $\theta = \hat{\theta}$ to Eq. (8) and using that $\ell_t(\hat{\theta}) = T_{EQ}$ yields

$$u_{OPT}T_{OPT} - u_{EQ} \hat{\theta} \geq \sum_{p \in \mathcal{P}} \hat{f}_p \left(T_{EQ} - \hat{\theta} - \sum_{e=(v,w) \in p} \frac{z_e(\ell_v(\hat{\theta}))}{v_e} \right). \quad (9)$$

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

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

and hence,

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

The equality follows by summing over all edges instead of all paths, and the last inequality is implied by $\hat{f}_e \leq v_e$. The result follows from our assumption that $u_{OPT} = u_{EQ}$. \square

To complete the proof of Theorem 2, it remains to bound the sum in the right-hand side of Eq. (7). The following lemma does exactly this and does not rely on any assumption.

Lemma 10. *In the dynamic equilibrium, for all $\theta \geq 0$,*

$$\sum_{e=(v,w) \in E} z_e(\ell_v(\theta)) \leq \frac{u_{EQ}}{e} (\ell_t(\theta) - \ell_t(0)).$$

We prove this lemma using two more technical claims.

Claim 11. In the dynamic equilibrium, for all $\theta \geq 0$,

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

PROOF OF CLAIM 11. Unless indicated otherwise, by an edge e we mean an edge $e = (v, w)$. We begin by writing the queue length in terms of its derivative by using Eq. (1).

$$\sum_{e=(v,w) \in E} z_e(\ell_v(\theta)) = \sum_{e \in E} \int_0^\theta \frac{dz_e(\ell_v(\xi))}{d\xi} \mathbb{1}_{z_e(\ell_v(\xi)) > 0} d\xi = \int_0^\theta \sum_{e \in E'_\xi} \frac{dz_e(\ell_v(\xi))}{d\xi} d\xi. \quad (10)$$

Denoting the flow underlying the dynamic equilibrium by f , for $e \in E_\xi^*$ we have $z'_e(\xi) = f_e^+(\xi) - v_e$ and $\ell'_w(\xi) = x'_e(\xi)/v_e$. Then, using Proposition 7, for edges $e \in E_\xi^*$ we can write

$$\begin{aligned} \frac{dz_e(\ell_v(\xi))}{d\xi} &= z'_e(\ell_v(\xi))\ell'_v(\xi) \\ &= f_e^+(\ell_v(\xi))\ell'_v(\xi) - v_e\ell'_v(\xi) \\ &= x'_e(\xi) - v_e\ell'_v(\xi) \\ &= x'_e(\xi) \left(1 - \frac{v_e\ell'_v(\xi)}{x'_e(\xi)}\right) \\ &= x'_e(\xi) \left(1 - \frac{\ell'_v(\xi)}{\ell'_w(\xi)}\right). \end{aligned}$$

Plugging the above expression into Eq. (10) yields

$$\sum_{e \in E} z_e(\ell_v(\theta)) = \int_0^\theta \sum_{e \in E_\xi^*} x'_e(\xi) \left(1 - \frac{\ell'_v(\xi)}{\ell'_w(\xi)}\right) d\xi.$$

Note that for edges $e = (v, w) \in E'_\xi \setminus E_\xi^*$, we have $\ell_w(\xi) = \ell_v(\xi) + z_e(\ell_v(\theta))/v_e + \tau_e$, so $\ell'_w(\xi) = \ell'_v(\xi) + z'_e(\ell_v(\theta))\ell'_v(\theta)/v_e = \ell'_v(\xi) + 0$ almost everywhere. Using Proposition 6, this gives $1 - \ell'_v(\xi)/\ell'_w(\xi) = 0$ almost everywhere, and hence

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

The next claim bounds the integral on the right-hand side.

Claim 12. In the dynamic equilibrium, for almost all $\xi \geq 0$,

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

PROOF OF CLAIM 12. Consider a path decomposition \mathcal{P} of the dynamic equilibrium f at time ξ . Then we can rewrite the left-hand side as

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

Because of Proposition 6, we can apply the fact that $1 - x \leq \log(1/x)$ for $x = \ell'_v(\xi)/\ell'_w(\xi) > 0$. Therefore,

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

Now rewrite $\log\left(\frac{\ell'_w(\xi)}{\ell'_v(\xi)}\right)$ as $\log(\ell'_w(\xi)) - \log(\ell'_v(\xi))$. This gives a telescopic sum for every path $p \in \mathcal{P}$, so $\sum_{e \in p} \log\left(\frac{\ell'_w(\xi)}{\ell'_v(\xi)}\right) = \log(\ell'_t(\xi)) - \log(\ell'_s(\xi))$. Using that $\ell'_s(\xi) = 1$ and that $x'(\xi)$ is a flow of size u_{EQ} , this yields

$$\sum_{p \in \mathcal{P}} x'_p(\xi) \sum_{e \in p} \log\left(\frac{\ell'_w(\xi)}{\ell'_v(\xi)}\right) \leq \sum_{p \in \mathcal{P}} x'_p(\xi) \log(\ell'_t(\xi)) = u_{EQ} \log(\ell'_t(\xi)). \quad \square$$

We now show how Lemma 10 follows from these claims.

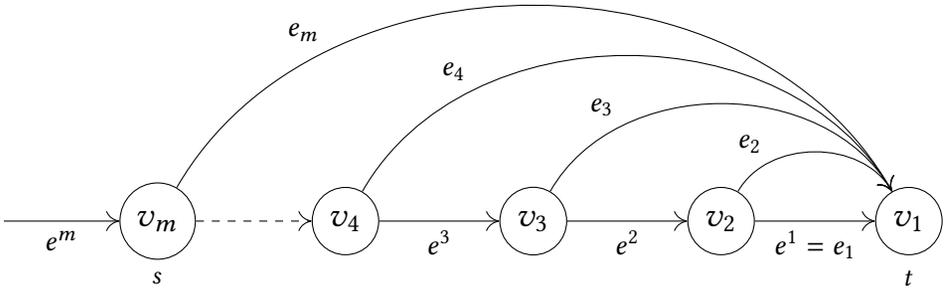


Fig. 3. An illustration of the tight instance.

PROOF OF LEMMA 10. From Claim 11 and Claim 12 we see that

$$\sum_{e=(v,w) \in E} z_e(\ell_v(\theta)) \leq u_{EQ} \int_0^\theta \log(\ell'_t(\xi)) d\xi. \quad (11)$$

Now substitute the monotonic function $\xi = \ell_t^{-1}(\varphi)$, giving $d\xi = \frac{d\varphi}{\ell'_t(\ell_t^{-1}(\varphi))}$. Then Eq. (11) becomes

$$\sum_{e \in E} z_e(\ell_v(\theta)) \leq u_{EQ} \int_{\ell_t(0)}^{\ell_t(\theta)} \frac{\log(\ell'_t(\ell_t^{-1}(\varphi)))}{\ell'_t(\ell_t^{-1}(\varphi))} d\varphi.$$

We now apply Proposition 8 with $x = \ell'_t(\ell_t^{-1}(\varphi)) > 0$ to see that

$$\sum_{e \in E} z_e(\ell_v(\theta)) \leq u_{EQ} \int_{\ell_t(0)}^{\ell_t(\theta)} \frac{1}{e} d\varphi = \frac{u_{EQ}}{e} (\ell_t(\theta) - \ell_t(0)). \quad \square$$

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

PROOF OF THEOREM 2. Applying Lemma 10 for $\theta = \hat{\theta}$ to the bound of Lemma 9 yields

$$T_{EQ} - T_{OPT} \leq \frac{1}{e} (\ell_t(\hat{\theta}) - \ell_t(0)) \leq \frac{1}{e} \ell_t(\hat{\theta}).$$

The result follows by writing $T_{EQ} = \ell_t(\hat{\theta})$ and rearranging this inequality to $\frac{T_{EQ}}{T_{OPT}} \leq \frac{e}{e-1}$. The fact that this is tight follows from Lemma 15 in the next subsection. \square

3.1 Tightness

Consider the family of instances described in [18, Section 7.4], where it is proved that the Price of Anarchy of these instances is at most $\frac{e}{e-1}$. We will prove that for a given choice of the edge capacities, the Price of Anarchy of these instances tends to $\frac{e}{e-1}$ in the limit, thereby proving tightness.

For completeness, we describe the family of instances here, and they are illustrated in Fig. 3. Fix parameters $m \in \mathbb{N}$ and $\alpha > 0$. Denote the capacity of edge e_i and e^i respectively by u_i and $u^i = \sum_{k=1}^i u_k$. Set the delay of e_i and e^i to $\tau_i = \alpha u^m \left(\frac{1}{u_{OPT}} - \frac{1}{u^i} \right)$ and $\tau^i = 0$, respectively. The equilibrium inflow rate is u^m .

We consider the instance where we set $u_1 = 1$ and $u^i = \left(\frac{m-1}{m-2} \right)^{i-1}$. 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 $M = \alpha u^m$.

The following lemmata show that the Price of Anarchy for this instance tends to $\frac{e}{e-1}$ for $m \rightarrow \infty$.

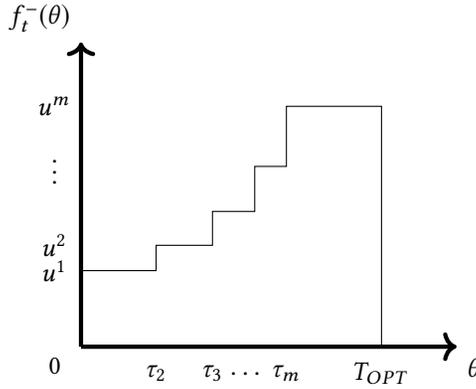


Fig. 4. The inflow rate into t of the optimum flow for the tight instance.

Lemma 13. *The completion time of the equilibrium is $T_{EQ} = \alpha u^m$.*

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

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

This yields

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

and therefore the first phase lasts until time $\theta = \alpha$, when all paths enter the dynamic shortest path network. Since the equilibrium inflow rate is u^m , we have $\hat{\theta} = \frac{M}{u^m} = \alpha$. Therefore,

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

Lemma 14. *The completion time of the optimum flow is $T_{OPT} = \tau_m$.*

PROOF. See Fig. 4 for an illustration of $f_t^-(\theta)$, the inflow rate into t of the optimum flow as a function of θ . Note that since M equals the area under this curve,

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

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

$$(\tau_{i+1} - \tau_i) u^i = \alpha u^m \left(\frac{1}{u^i} - \frac{1}{u^{i+1}} \right) u^i = \alpha u^m \left(1 - \frac{u^i}{u^{i+1}} \right) = \alpha u^m \left(1 - \left(\frac{m-1}{m-2} \right)^{-1} \right) = \frac{\alpha u^m}{m-1}.$$

We conclude that at time τ_m , a total amount of flow of $(m-1) \frac{\alpha u^m}{m-1} = \alpha u^m$ has arrived at t . Since $M = \alpha u^m$, therefore $T_{OPT} = \tau_m$. \square

Lemma 15. *The Price of Anarchy of $\frac{c}{c-1}$ is tight.*

PROOF. From Lemma 13 and Lemma 14 we see that the Price of Anarchy of this instance equals

$$\frac{T_{EQ}}{T_{OPT}} = \frac{\alpha u^m}{\alpha u^m \left(1 - \frac{1}{u^m}\right)} = \frac{u^m}{u^m - 1} = \frac{\left(\frac{m-1}{m-2}\right)^{m-1}}{\left(\frac{m-1}{m-2}\right)^{m-1} - 1},$$

which tends to $\frac{e}{e-1}$ for $m \rightarrow \infty$. \square

4 DISCUSSION

In this section we strengthen our result for parallel-link networks, and we discuss the monotonicity conjecture and related techniques to prove a universal bound on the PoA.

4.1 Parallel-link networks

In this section we strengthen our result for the class of parallel-link networks by proving Theorem 4. First of all, we prove Conjecture 3 for the class of parallel-link networks.

Lemma 16. *Consider a parallel-link network G and two fixed inflow rates $u_1 < u_2$ with their corresponding dynamic equilibria in G . Then $T_{EQ}^1 \geq T_{EQ}^2$.*

PROOF. Denote the edges of G by e_0, \dots, e_m and denote the capacity and delay of edge e_j by v_j and τ_j . Without loss of generality, order the edges such that $\tau_0 < \dots < \tau_m$ (if two or more edges have the same delay, merge them into one edge with the sum of their capacities). For brevity, denote $E_i = \{e_j : j \leq i\}$ and $v^i = \sum_{e \in E_i} v_e = \sum_{j=0}^i v_j$.

First we describe the dynamic equilibrium in a parallel-link network with a generic inflow rate u_0 . Active edges $e = (v, w) \in E$ at time θ can be characterized by the inequality $\ell_w(\theta) - \ell_v(\theta) \geq \tau_e$. Since every edge in a parallel-link network connects s to t , the set of active edges at any time has the form E_i for some i . Moreover, edges will enter the set of active edges sequentially and grow a queue, until the sum of their capacities surpasses the inflow rate. More precisely, if we denote the sets of active edges respectively of edges with positive queue in phase i by E'_i and E_i^* , we see that $E'_i = E_i^* = E_i$, for every phase i such that $v^i < u_0$. Denote by k the first phase where $v^k \geq u_0$. In phase k , the queue lengths will remain constant, and in particular, no queue will accumulate at edge e_k . Consequently, $E'_k = E_k$ and $E_k^* = E_{k-1}$ (if M is sufficiently large to reach this phase).

For phase i , denote $\rho_i = \ell'_i(\theta)$ and $x_e^i = x'_e(\theta)$ for any $\theta \in (\theta_i, \theta_{i+1})$ and every $e \in E'_i$. From thin flow condition (6) we see that $\ell'_i(\theta) = x'_e(\theta)/v_e$ for $e \in E'_i$. Therefore in phases $i < k$ where $E_i = E'_i$, we obtain $\rho_i = x_e^i/v_e$ for all $e \in E'_i$, and in particular $\rho_i = u_0/v^i$. Since a phase $i < k$ lasts until $\ell_t(\theta) - \ell_s(\theta)$ has grown to value τ_{i+1} , it lasts for $(\tau_{i+1} - \tau_i)v^i/u_0$ units of time, and phase k lasts until time $\hat{\theta}$.

Observe that flow sent in phase $i < k$, i.e., during the interval (θ_i, θ_{i+1}) , arrives at t in the interval $(\theta_i + \tau_i, \theta_{i+1} + \tau_{i+1})$. As a result, the inflow rate at t in this period of time is v^i , which is independent of u_0 . In phase k the inflow rate at t is u_0 and arrives during $(\theta_k, \hat{\theta} + \tau_k)$.

With this dynamic equilibrium structure at hand, we prove the lemma. Since the length of a phase is inversely proportional to the inflow rate, phases last shorter with inflow rate u_2 than with u_1 . Moreover, the inflow rate into t is higher for subsequent phases. Hence, at every point in time the inflow rate into t with inflow rate u_2 is at least the inflow rate into t with inflow rate u_1 . Thus, a total amount of flow of M arrives earlier at t with inflow rate u_2 than with inflow rate u_1 , so $T_{EQ}^2 \leq T_{EQ}^1$. \square

This monotonicity result allows us to state the following universal bound on the Price of Anarchy for parallel-link networks.

Lemma 17. *In parallel-link networks, $T_{EQ} \leq \frac{4}{3} \cdot T_{OPT}$.*

PROOF. From Claim 11, we see that in the dynamic equilibrium for all $\theta \geq 0$,

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

Now note that in parallel-link networks, every path is just an edge from s to t . Hence,

$$\sum_{e \in E'_\xi} x'_e(\xi) \left(1 - \frac{\ell'_v(\xi)}{\ell'_w(\xi)}\right) = u_{EQ} \left(1 - \frac{1}{\ell'_t(\xi)}\right).$$

Now observe that $1 - \frac{1}{x} \leq \frac{1}{4}x$ for any $x > 0$. To see this, denote $f(x) = 1 - \frac{1}{x}$ and $g(x) = \frac{1}{4}x$ and note that $f(2) = g(2)$, while $f'(x) > g'(x)$ for $0 < x < 2$ and vice versa for $x > 2$. Therefore,

$$\sum_{e \in E} z_e(\ell_v(\theta)) \leq u_{EQ} \int_0^\theta \left(\frac{1}{4}\ell'_t(\xi)\right) d\xi = \frac{u_{EQ}}{4} (\ell_t(\theta) - \ell_t(0)).$$

Note that because of Lemma 16, the result of Lemma 9 holds in general without the assumption on the inflow rates (cf. Lemma 19). Invoking Eq. (7), we obtain that $\ell_t(\hat{\theta}) - T_{OPT} \leq \frac{1}{4}\ell_t(\hat{\theta})$, and the bound follows by rearranging the terms. \square

The proof of Theorem 4 is completed by showing tightness of its bound.

Lemma 18. *The Price of Anarchy of $\frac{4}{3}$ for parallel-link networks is tight.*

PROOF. The fact that the bound from Lemma 17 is tight follows from the example depicted in Fig. 5. The network consists of two edges with $v_1 = \frac{1}{2}$, $\tau_1 = 0$, $v_2 = \frac{1}{2}$ and $\tau_2 = 1$. Taking $u_{EQ} = 1$ and $M = 1$, the equilibrium will send all flow along the top edge and the last unit of flow will arrive at time $\ell_t(\hat{\theta}) = 2$. On the other hand, the optimal flow will send $\frac{1}{2}$ along the bottom edge until time $\frac{1}{2}$, and $\frac{1}{2}$ along the top edge until time $T_{OPT} = \frac{3}{2}$. Therefore, the Price of Anarchy of this instance is exactly $\frac{4}{3}$. \square

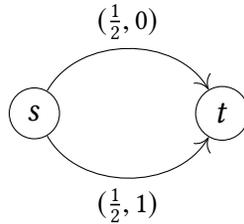


Fig. 5. An illustration of the tight instance for parallel-link networks.

We close this section by noting that the results in Lemmas 16, 17 and 18 also hold for parallel-serial networks, i.e. networks consisting of parallel paths. The reason for this is that a path can be contracted to a single link with capacity equal to the minimum capacity along the path and with delay equals the sum of the delays of the edges in the path. The details of this contraction argument are deferred to the full version of the paper.

4.2 Conjectures

We outline some techniques to prove a universal bound on the PoA. The most natural one would be to prove the monotonicity conjecture as stated in Conjecture 3. This is an important basic property of the fluid queuing model that seems intuitive yet appears to be notoriously hard to prove. The following lemma demonstrates how the universal bound on the PoA would follow from Conjecture 3 and Theorem 2.

Lemma 19. *Suppose Conjecture 3 holds. Then in any instance, $T_{EQ} \leq \frac{c}{c-1} \cdot T_{OPT}$.*

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 3 we know that $T_{EQ} \leq T'_{EQ}$, where T'_{EQ} is the makespan of the dynamic equilibrium in G with inflow rate $u'_{EQ} = u_{OPT}$. Since the network itself does not change, neither does the makespan of the optimal flow. Then, using Theorem 2, we obtain $T_{EQ} \leq T'_{EQ} \leq \frac{c}{c-1} \cdot T_{OPT}$. \square

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 to conjecture (but we are unable to prove it) that for $M > 0$, 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 $0 \leq \theta \leq \hat{\theta}$. The next lemma establishes that if this property holds, then the Price of Anarchy would also be $\frac{c}{c-1}$.

Lemma 20. *Suppose that $M > 0$ and that for all $0 \leq \theta \leq \hat{\theta}$ we have that $\ell_t(\theta) - \theta < T_{OPT}$, then $T_{EQ} \leq \frac{c}{c-1} \cdot T_{OPT}$.*

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

Since $\tau_e < T_{OPT}$, the optimum flow in G' will send a positive amount of flow along edge e , and it will send 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 $u'_{OPT} = u_{OPT} + v_e = u_{EQ}$, 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 will never be active in the dynamic equilibrium (for sufficiently small ε). Therefore $u'_{EQ} = u_{EQ}$ and $T'_{EQ} = T_{EQ}$. Applying Theorem 2 to G' , we can write $T_{EQ}/T_{OPT} \leq T'_{EQ}/T'_{OPT} \leq \varepsilon/(\varepsilon - 1)$. \square

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

$$\sum_{e=(v,w) \in E} \hat{f}_e \frac{z_e(\ell_v(\hat{\theta}))}{v_e} < u_{EQ} \hat{\theta}. \quad (12)$$

Then, we could follow the beginning of the proof of Lemma 9 and we can write the following similar to Equation (9).

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

Here, last inequality is obtained by assuming the conjecture given by Eq. (12). Then by cancelling on both sides and dividing by u_{OPT} we obtain that $T_{OPT} > \ell_t(\hat{\theta}) - \hat{\theta}$. Since 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 20.

We finish by noting that it would actually be enough to show the weaker condition: $\ell_t(\theta) - \theta \leq T_{OPT}$ for all $0 \leq \theta \leq \hat{\theta}$ (without strict inequality). Indeed, this is enough to establish Lemma 20, although the proof requires a more careful and tedious limiting argument which we defer to the full version of the paper.

ACKNOWLEDGMENTS

The authors would like to thank Dario Frascaria, Marcus Kaiser, Neil Olver, Leon Sering and Laura Vargas Koch for fruitful discussions.

This work was partially supported by CONICYT under Grants PCI PII 20150140 and CONICYT-PFCHA/Doctorado Nacional/2018-21180347.

REFERENCES

- [1] Edward J. Anderson and Andy B. Philpott. Optimisation of flows in networks over time. In F. P. Kelly, editor, *Probability, Statistics and Optimisation*, Wiley, New York, 369–382, 1994.
- [2] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice-Hall, NJ, USA, 1993.
- [3] Umang Bhaskar, Lisa Fleischer, and Elliot Anshelevich. A Stackelberg strategy for routing flow over time. SODA 2011.
- [4] Umang Bhaskar, Lisa Fleischer, and Elliot Anshelevich. A Stackelberg strategy for routing flow over time. *Games and Economic Behavior* 92:232–247, 2015.
- [5] Zhigang Cao, Bo Chen, Xujin Chen, and Changjun Wang. A network game of dynamic traffic. EC 2017.
- [6] Roberto Cominetti, José Correa, and Omar Larré. Existence and uniqueness of equilibria for flows over time. ICALP 2011.
- [7] Roberto Cominetti, José Correa, and Omar Larré. Dynamic equilibria in fluid queuing networks. *Operations Research* 63(1):21–34, 2015.
- [8] Roberto Cominetti, José Correa, and Neil Olver. Long term behavior of dynamic equilibria in fluid queuing networks. IPCO 2017.
- [9] Lisa Fleischer and Eva Tardos. Efficient continuous-time dynamic network flow algorithms. *Operations Research Letters*, 23(3-5):71–80, 1998.
- [10] Lester Randolph Ford and Delbert Ray Fulkerson. *Flows in Networks*. Princeton University Press, 1962.
- [11] Lester Randolph Ford and Delbert Ray Fulkerson. Constructing maximal dynamic flows from static flows. *Operations Research*, 6:419–433, 1958.
- [12] David Gale. Transient flows in networks. *Michigan Mathematical Journal*, 6:59–63, 1959.
- [13] Felipe Garrido. Equilibrio dinámico en redes de transporte con múltiples orígenes y destinos. Master’s thesis, Universidad de Chile, 2017.
- [14] Lucas Graf and Tobias Harks. Dynamic flows with adaptive route choice. IPCO 2019.
- [15] Martin Hoefer, Vahab Mirrokni, Heiko Röglin, and Shang-Hua Teng. Competitive routing over time. *Theoretical Computer Science* 412(39):5420–5432, 2011.
- [16] Anisse Ismaili. Routing games over time with FIFO policy. WINE 2017.
- [17] Marcus Kaiser. Personal communication. 2018.
- [18] Ronald Koch. Routing games over time. Ph.D. Dissertation, Technische Universität Berlin, 2012.
- [19] Ronald Koch and Martin Skutella. Nash equilibria and the price of anarchy for flows over time. *Theory of Computing Systems* 49(1):71–97, 2011.
- [20] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. STACS 1999
- [21] Martin Macko, Kate Larson, and Luboš Steskal. Braess’s paradox for flows over time. SAGT 2010.
- [22] Frédéric Meunier and Nicolas Wagner. Equilibrium results for dynamic congestion games. *Transportation Science*, 44(4):524–536, 2010.
- [23] Srinivas Peeta and Athanasios K. Ziliaskopoulos. Foundations of dynamic traffic assignment: The past, the present and the future. *Networks and Spatial Economics* 1(3-4):233–265, 2001.
- [24] Bin Ran and David Boyce. *Modeling dynamic transportation networks*. Springer, Berlin, 1996.

- [25] Marco Scarsini, Marc Schröder, and Tristan Tomala. Dynamic atomic congestion games with seasonal flows. *Operations Research* 66(2):327-339, 2018.
- [26] Leon Sering and Martin Skutella. Multi-source multi-sink Nash flows over time. ATMOS 2018.
- [27] Leon Sering and Laura Vargas Koch. Nash flows over time with spillback. SODA 2019.
- [28] Martin Skutella. An introduction to network flows over time. In W. Cook, L. Lovász, and J. Vygen, editors, *Research Trends in Combinatorial Optimization*, Springer, Berlin, 451–482, 2009.
- [29] Martin Skutella. Personal communication. 2006.
- [30] William S. Vickrey. Congestion theory and transport investment. *American Economic Review*, 59(2):251–260, 1969.
- [31] Samuel Yagar. Dynamic traffic assignment by individual path minimization and queuing. *Transportation Research* 5(3):179–196, 1971.