

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

Long term behavior of dynamic equilibria in fluid queuing networks

Roberto Cominetti

Facultad de Ingeniería y Ciencias, Universidad Adolfo Ibáñez, roberto.cominetti@uai.cl

José Correa

Departamento de Ingeniería Industrial, Universidad de Chile, correa@uchile.cl

Neil Olver

Department of Mathematics, London School of Economics and Political Science, N.Olver@lse.ac.uk

A fluid queuing network constitutes one of the simplest models in which to study flow dynamics over a network. In this model we have a single source-sink pair and each link has a per-time-unit capacity and a transit time. A dynamic equilibrium (or equilibrium flow over time) is a flow pattern over time such that no flow particle has incentives to unilaterally change its path. Although the model has been around for almost fifty years, only recently results regarding existence and characterization of equilibria have been obtained. In particular the long term behavior remains poorly understood. Our main result in this paper is to show that, under a natural (and obviously necessary) condition on the queuing capacity, a dynamic equilibrium reaches a steady state (after which queue lengths remain constant) in finite time. Previously, it was not even known that queue lengths would remain bounded. The proof is based on the analysis of a rather non-obvious potential function that turns out to be monotone along the evolution of the equilibrium. Furthermore, we show that the steady state is characterized as an optimal solution of a certain linear program. When this program has a unique solution, which occurs generically, the long term behavior is completely predictable. On the contrary, if the linear program has multiple solutions the steady state is more difficult to identify as it depends on the whole temporal evolution of the equilibrium.

Key words: Flows over time, Dynamic Equilibria, Steady State

History:

1. Introduction

The theory of flows over time provide a natural and convenient model to describe the dynamics of a continuous stream of particles traveling from a source to a sink in a network, such as urban or Internet traffic. Probably the most basic model for the propagation of flow

is the so-called *fluid-queue model* in which each arc in the network consists of a fluid queue with an arc-specific capacity followed by a link with constant delay. Thus, the time to traverse an edge is composed of a flow-dependent waiting time in the queue plus a constant travel time after leaving the queue. This model was initially studied in the framework of optimization. Ford and Fulkerson (1958, 1962) considered a fluid queue model in a discrete time setting and designed an algorithm to compute a flow over time carrying the maximum possible flow from the source s to the sink t in a given timespan. Gale (1959) then 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 (1994) and Fleischer and Tardos (1998). We refer to Skutella (2009) for an excellent survey. However, when network flows suffer from a lack of coordination among the participating agents, it is natural to take a game theoretic approach. As first described by Vickrey (1969) for a simple *bottleneck model*, in a dynamic network routing game each infinitesimal particle is interpreted as a player that seeks to complete its journey in the least possible time. Players are forward-looking and anticipate the congestion and queuing delays induced by others upon arrival to any edge in the network. Equilibrium occurs when each particle travels along a shortest path.

More formally, a *fluid queuing network* is a directed graph $G = (V, E)$ where each arc $e \in E$ consists of a fluid queue with capacity $\nu_e > 0$ followed by a link with constant delay $\tau_e \geq 0$ (see Figure 1). A constant inflow rate $u_0 > 0$ enters the network at a fixed source $s \in V$ and travels towards a terminal node $t \in V$. A *dynamic equilibrium* models the temporal evolution of the flows in the network. Loosely speaking, it consists of a flow pattern in which every particle travels along a shortest path, accounting for the fact that travel times depend on the instant at which a particle enters the network as well as the state of the queues that will be encountered along its path by the time at which they are reached. Intuitively, if the queues are initially empty, the equilibrium should start by sending all the flow along shortest paths considering only the free-flow delays τ_e . These paths are likely to become overloaded so that queues will grow on some of its edges and at some point in time new paths will become competitive and will be incorporated into the equilibrium. These new paths may in turn build queues so that even longer paths may come into play. Hence one might expect that the equilibrium proceeds in phases in which the paths used by the equilibrium remain stable. However, it is unclear if the number of

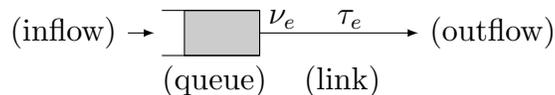


Figure 1 An arc in the fluid queuing network.

such phases is finite and whether the equilibrium will eventually reach a steady state in which the queues and travel times stabilize.

Although dynamic equilibria have been around for almost fifty years (see, e.g., (Ford and Fulkerson 1958, Friesz et al. 1993, Gale 1959, Merchant and Nemhauser 1978a,b, Peeta and Ziliaskopoulos 2001, Ran and Boyce 1996, Vickrey 1969, Xu et al. 1999)), their existence has only been proved recently by Zhu and Marcotte (2000) though in a somewhat different setting, and by Meunier and Wagner (2010) who gave the first existence result for a model that covers the case of fluid queuing networks. These proofs, however, rely heavily on functional analysis techniques and provide little intuition on the combinatorial structure of dynamic equilibria, their characterization, or feasible approaches to compute them. Substantial progress was recently achieved by Koch and Skutella (2011) by introducing the concept of *thin flows with resetting* that characterize the time derivatives of a dynamic equilibrium, and which provide in turn a method to compute an equilibrium by integration. A slightly refined notion of *normalized thin flows with resetting* was considered by Cominetti et al. (2015), who proved existence and uniqueness, and provided a constructive proof for the existence of a dynamic equilibrium.

In recent work, further extensions and variants of the model have been studied. In particular, Sering and Skutella (2018) extend some of the known results about dynamic equilibria to the case in which there are multiple sources and terminal nodes. However, the multi-commodity case is largely open. Furthermore, Sering and Vargas Koch (2019) consider spill back effects to model the fact that in practice the queues cannot grow arbitrarily large and that their effect propagates back in the network. Graf and Harks (2019) consider a related model in which particles behave myopically and make routing decisions based on the current state of the network, without anticipating its evolution. Finally, we mention the work of Cao et al. (2017) who considered an atomic model and established that in series-parallel networks queues remain bounded in a dynamic equilibrium.

In this paper we focus on the long term behavior of dynamic equilibria in fluid queuing networks. Clearly if the inflow u_0 is very large compared to the queuing capacities, the

queues will grow without bound, and no steady state can be expected. More precisely, let $\delta^+(S)$ be an st -cut with minimum queuing capacity $\bar{\nu} = \sum_{e \in \delta^+(S)} \nu_e$; if there are multiple options, choose S (containing s) to be setwise minimal. If $u_0 > \bar{\nu}$ all the arcs in $\delta^+(S)$ will grow unbounded queues, whereas for $u_0 \leq \bar{\nu}$, it is natural to expect that the equilibrium should eventually reach a steady state, where queue lengths remain constant. This was not known—in fact, it was not even known that queue lengths remain bounded!

Our main goal in this paper is to show that both these properties do indeed hold: more precisely, when $u_0 \leq \bar{\nu}$, the dynamic equilibrium reaches a steady state in *finite* time. At first glance, these convergence properties might seem “obvious”, and it might seem surprising that they are at all difficult to prove. We will present some examples that illustrate why this is not the case. For instance, it may occur that the flow across the cut $\delta^+(S)$ may temporarily exceed its capacity $\bar{\nu}$ by an arbitrarily large factor, forcing the queues to grow very large. This phenomenon may occur since the inflow u_0 entering the network at different points in time may experience different delays and eventually superpose at $\delta^+(S)$ which gets an inflow larger than u_0 . In other cases some queues may grow during a period of time after which they reduce to zero and then grow again later on. In fact, we give a construction that shows that this can happen an exponential (in the input size) number of times during the evolution! Along the way to our main result, we provide a characterization of the steady state as an optimal solution of a certain linear programming problem and we discuss when this problem has a unique solution. Despite the fact that convergence to a steady state occurs in finite time, it remains as an open question whether this state is attained after finitely many phases.

The paper is structured as follows. Section 2 reviews the model of fluid queuing networks, including the precise definition of dynamic equilibrium and the main results known so far. Then, in Section 3 we discuss the notion of steady state and provide a characterization in terms of a linear program. Inspired by the objective function of this linear program, in Section 4 we introduce a potential function and we prove that it is a Lyapunov function for the dynamics. This potential turns out to be piecewise linear in time with finitely many possible slopes. We then prove that the potential remains bounded so that there is a finite time at which its slope is zero, and we show that in that case the system has reached a steady state. Further, we provide an explicit pseudopolynomial bound on the convergence time. Finally, in Section 5 we discuss the interesting (and perhaps surprising) examples alluded to earlier, as well as some remaining open questions.

2. Dynamic equilibria in fluid queuing networks

In this section we recall the definition of dynamic equilibria in fluid queuing networks, and we briefly review the known results on their existence, characterization, and computation. The results are stated without proofs for which we refer to Koch and Skutella (2011) and Cominetti et al. (2015).

2.1. The model

Consider a fluid queuing network $G = (V, E)$ with arc capacities ν_e and delays τ_e . The network dynamics are described in terms of the inflow rates $f_e^+(\theta)$ that enter each arc $e \in E$ at time θ , where $f_e^+ : [0, \infty) \rightarrow [0, \infty)$ is measurable.

Arc dynamics. If the inflow $f_e^+(\theta)$ exceeds ν_e a queue $z_e(\theta)$ will grow at the entrance of the arc. The queues are assumed to operate at capacity, that is to say, when $z_e(\theta) > 0$ the

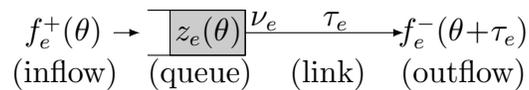


Figure 2 Dynamics of an arc in the queuing network.

flow is released at rate ν_e , whereas when the queue is empty the outflow is the minimum between $f_e^+(\theta)$ and the capacity ν_e . Hence the queue evolves from its initial state $z_e(0)=0$ according to

$$\dot{z}_e(\theta) = \begin{cases} f_e^+(\theta) - \nu_e & \text{if } z_e(\theta) > 0 \\ [f_e^+(\theta) - \nu_e]_+ & \text{if } z_e(\theta) = 0. \end{cases} \quad (1)$$

These dynamics uniquely determine the queue lengths $z_e(\theta)$ as well as the arc outflows

$$f_e^-(\theta + \tau_e) = \begin{cases} \nu_e & \text{if } z_e(\theta) > 0 \\ \min\{f_e^+(\theta), \nu_e\} & \text{if } z_e(\theta) = 0. \end{cases} \quad (2)$$

Flow conservation. A *flow over time* is a family $(f_e^+)_{e \in E}$ of arc inflows such that flow is conserved at every node $v \in V \setminus \{t\}$, namely for a.e. $\theta \geq 0$

$$\sum_{e \in \delta^+(v)} f_e^+(\theta) - \sum_{e \in \delta^-(v)} f_e^-(\theta) = \begin{cases} u_0 & \text{if } v = s \\ 0 & \text{if } v \neq s, t. \end{cases} \quad (3)$$

Dynamic shortest paths. A particle entering an arc e at time θ experiences a queuing delay $z_e(\theta)/\nu_e$ plus a free-flow delay τ_e to traverse the arc after leaving the queue, so that it will exit the arc at time

$$T_e(\theta) = \theta + \frac{z_e(\theta)}{\nu_e} + \tau_e. \quad (4)$$

Consider a particle entering the source node s at time θ . If this particle follows a path $p = e_1 e_2 \cdots e_k$, it will reach the end of the path at time

$$T_p(\theta) = T_{e_k} \circ \cdots \circ T_{e_2} \circ T_{e_1}(\theta). \quad (5)$$

Denoting \mathcal{P}_v the set of all sv -paths, the minimal time at which node v can be reached is

$$\ell_v(\theta) = \min_{p \in \mathcal{P}_v} T_p(\theta). \quad (6)$$

The paths attaining these minima are called *dynamic shortest paths*. The arcs in these paths are said to be *active* at time θ and we denote them by E'_θ . Observe that $\ell_v(\theta)$ can also be defined through the dynamic Bellman's equations

$$\begin{cases} \ell_s(\theta) = \theta \\ \ell_w(\theta) = \min_{e=vw \in E} T_e(\ell_v(\theta)) \end{cases} \quad (7)$$

so that $e = vw$ is active precisely if $\ell_w(\theta) = T_e(\ell_v(\theta))$.

Dynamic equilibrium. A dynamic equilibrium is a flow pattern that uses only dynamic shortest paths. More precisely, let $\Theta_e = \{\theta : e \in E'_\theta\}$ be the set of *entrance times* θ at which the arc e is active, and $\Xi_e = \ell_v(\Theta_e)$ the set of *local times* $\xi = \ell_v(\theta)$ at which e will be active. A flow over time $(f_e^+)_{e \in E}$ is called a *dynamic equilibrium* iff for almost every $\xi \geq 0$ we have $f_e^+(\xi) > 0 \Rightarrow \xi \in \Xi_e$.

2.2. Characterization of dynamic equilibria

Since the inflows $f_e^+(\cdot)$ are measurable the same holds for $f_e^-(\cdot)$ and we may define the *cumulative inflows* and *cumulative outflows* as

$$\begin{aligned} F_e^+(\theta) &= \int_0^\theta f_e^+(z) dz \\ F_e^-(\theta) &= \int_0^\theta f_e^-(z) dz. \end{aligned}$$

These cumulative flows allow to express the queues as $z_e(\theta) = F_e^+(\theta) - F_e^-(\theta + \tau_e)$. It turns out that a dynamic equilibrium can be equivalently characterized by the fact that for each arc $e = vw \in E$ we have

$$F_e^+(\ell_v(\theta)) = F_e^-(\ell_w(\theta)) \quad \forall \theta \geq 0. \quad (8)$$

In this case, the functions $x_e(\theta) \triangleq F_e^+(\ell_v(\theta))$ are static flows with

$$\sum_{e \in \delta^+(v)} x_e(\theta) - \sum_{e \in \delta^-(v)} x_e(\theta) = \begin{cases} u_0 \theta & \text{if } v = s \\ -u_0 \theta & \text{if } v = t \\ 0 & \text{if } v \neq s, t. \end{cases} \quad (9)$$

2.3. Derivatives of a dynamic equilibrium

The labels $\ell_v(\theta)$ and the static flows $x_e(\theta)$ are nondecreasing functions which are also absolutely continuous so that they can be reconstructed from their derivatives by integration.¹ Moreover, from these functions one can recover the equilibrium inflows $f_e^+(\cdot)$ using the relation $x'_e(\theta) = f_e^+(\ell_v(\theta))\ell'_v(\theta)$. Hence, finding a dynamic equilibrium reduces essentially to computing the derivatives $\ell'_v(\theta), x'_e(\theta)$.

Let θ be a point of differentiability and set $\ell'_v = \ell'_v(\theta) \geq 0$ and $x'_e = x'_e(\theta) \geq 0$. From (9) we see that x' is a static st -flow of size u_0 , namely,

$$\sum_{e \in \delta^+(v)} x'_e - \sum_{e \in \delta^-(v)} x'_e = \begin{cases} u_0 & \text{if } v = s \\ -u_0 & \text{if } v = t \\ 0 & \text{if } v \neq s, t \end{cases} \quad (10)$$

while using (7), (4), (1) and the differentiation rule for a minimum we get

$$\begin{cases} \ell'_s = 1 \\ \ell'_w = \min_{e=vw \in E'_\theta} \rho_e(\ell'_v, x'_e) \end{cases} \quad (11)$$

where

$$\rho_e(\ell'_v, x'_e) = \begin{cases} x'_e/\nu_e & \text{if } e \in E_\theta^* \\ \max\{\ell'_v, x'_e/\nu_e\} & \text{if } e \notin E_\theta^* \end{cases} \quad (12)$$

with E_θ^* the set of arcs $e = vw$ with positive queue $z_e(\ell_v(\theta)) > 0$. In addition to this, the conditions for dynamic equilibria imply $E_\theta^* \subset E'_\theta$ as well as

$$\begin{aligned} (\forall e \in E'_\theta) \quad & x'_e > 0 \Rightarrow \ell'_w = \rho_e(\ell'_v, x'_e) \\ (\forall e \notin E'_\theta) \quad & x'_e = 0. \end{aligned} \quad (13)$$

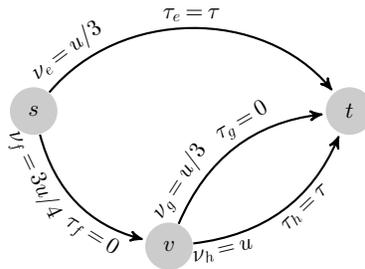
These equations fully characterize the derivatives of a dynamic equilibrium. In fact, for all subsets $E^* \subseteq E' \subseteq E$ the system (10)-(13) admits at least one solution (ℓ', x') and moreover the ℓ' component is unique. These solutions are called *normalized thin flows with resetting* (NTFR) and can be used to reconstruct a dynamic equilibrium by integration, proving the existence of equilibria. We refer to Cominetti et al. (2015) for the existence and uniqueness of NTFR's and to Koch and Skutella (2011) for a description of the integration algorithm and how to find the equilibrium inflows $f_e^+(\cdot)$.

Observe that there are only finitely many options for E^* and E' . Since the corresponding ℓ' is unique, it follows that the functions $\ell_v(\theta)$ will be uniquely defined and piecewise linear with finitely many options for the derivatives. Although the static flows $x_e(\theta)$ are not unique in general, one can still find an equilibrium in which these functions are also piecewise linear by fixing a specific x' in the NTFR for each pair E^*, E' .

2.4. A detailed example

We now work out the details of a small example that already provides some intuition on how the dynamic equilibria behaves. In particular this example exhibits an unexpected property, namely that the flow coming into the sink can be larger than the network inflow.

EXAMPLE 1. Consider the network consisting of the vertices $\{s, v, t\}$ with edges $e = (s, t), f = (s, v), g = (v, t), h = (v, t)$ and inflow $u_0 = u$. Capacities are $\nu_e = u/3, \nu_f = 3u/4, \nu_g = u/3, \nu_h = u$, and delays are $\tau_e = \tau_h = \tau$, and $\tau_f = \tau_g = 0$.



In a dynamic equilibrium for this instance, flow is initially routed through the shortest path fg . Then queues grow in both edges until at time $\tau/2$ the path consisting of e enters the shortest path network. From that point in time the flow splits in equal proportions between paths e and fg implying that a queue starts growing on edge e , the queue of f starts decreasing, while the queue on g continues to increase. Then, at time 0.7τ , the

queue on g has grown enough to make h enter the shortest path network. At this point all edges except h have a queue and the flow starts splitting as follows: $4u/13$ take path e , $4u/13$ take path fg , and $5u/13$ take path fh . Therefore all queues start decreasing until at time 2τ the queues on edges e and f deplete simultaneously. When this happens the shortest path network stays the same but the edges with queue change since only g still has a queue. The new thin flow is thus computed, and the flow starts splitting evenly between the paths e , fg , and fh (each gets flow $u/3$). This last phase constitute a steady state so it lasts forever. More precisely in this instance one can compute the derivative of the distance labels at node t as

$$\ell'_t(\theta) = \begin{cases} 3 & \text{for } \theta \in [0, \tau/2) \\ 3/2 & \text{for } \theta \in [\tau/2, \tau/2 + \tau/5) \\ 12/13 & \text{for } \theta \in [\tau/2 + \tau/5, 2\tau) \\ 1 & \text{for } \theta \in [2\tau, \infty) \end{cases}.$$

Interestingly, the amount of flow arriving at t at time $\ell_t(\theta)$ can readily be computed as $u/\ell'_t(\theta)$. So that if we consider the local time at node t this flow is then

$$f_e^-(\theta) + f_g^-(\theta) + f_h^-(\theta) = \begin{cases} u/3 & \text{for } \theta \in [0, 3\tau/2) \\ 2u/3 & \text{for } \theta \in [3\tau/2, 9\tau/5) \\ 13u/12 & \text{for } \theta \in [9\tau/5, 3\tau) \\ u & \text{for } \theta \in [3\tau, \infty). \end{cases}$$

This brings us to the surprising fact that for some time interval, the flow arriving at the sink is larger than the inflow.

3. Steady states

We say that a dynamic equilibrium attains a *steady state* if for sufficiently large times all the queues are frozen to a constant $z_e(\theta) \equiv z_e^*$. This is clearly equivalent to the fact that the arc travel times become constant equal to $\tau_e^* = \tau_e + q_e^*$ with $q_e^* = z_e^*/\nu_e$ the corresponding queuing times.

LEMMA 1. *A dynamic equilibrium attains a steady state iff there exists some $\theta^* \geq 0$ such that $\ell'_v(\theta) = 1$ for every node $v \in V$ and all $\theta \geq \theta^*$.*

Proof. In a steady state we clearly have $\ell_v(\theta) = \theta + d_v^*$ where d_v^* is the minimum travel time from s to v with arc times τ_e^* , so that $\ell'_v(\theta) = 1$. Conversely, if all these derivatives

are equal to 1 then $\ell_v(\theta) = \theta + d_v^*$ for some constant d_v^* and $\theta \geq \theta^*$. Moreover, an arc $e = vw$ with nonempty queue must be active so that $\ell_w(\theta) = T_e(\ell_v(\theta))$ which yields

$$z_e(\theta + d_v^*) = z_e(\ell_v(\theta)) = \nu_e(\ell_w(\theta) - \ell_v(\theta) - \tau_e) = \nu_e(d_w^* - d_v^* - \tau_e)$$

which shows that all queues eventually become constant. \square

THEOREM 1. *Consider a steady state with queues $z_e^* \geq 0$ and let d_v^* be the minimum travel time from s to v under arc travel times $\tau_e^* = \tau_e + q_e^*$, where $q_e^* = z_e^*/\nu_e$. Let (ℓ', x') with $\ell'_v = 1$ for all $v \in V$ be a corresponding NTFR and denote by \mathcal{F}_0 the set of st-flows of value u_0 . Then x' and (d^*, q^*) are optimal solutions to the following pair of dual linear programs:*

$$\begin{aligned} \min_{y'} \quad & \sum_{e \in E} \tau_e y'_e \\ \text{s.t.} \quad & y' \in \mathcal{F}_0 \\ & 0 \leq y'_e \leq \nu_e \quad \forall e \in E, \end{aligned} \tag{P}$$

$$\begin{aligned} \max_{d, q} \quad & u_0 d_t - \sum_{e \in E} \nu_e q_e \\ \text{s.t.} \quad & d_s = 0 \\ & d_w \leq d_v + \tau_e + q_e \quad \forall e = vw \in E \\ & q_e \geq 0 \quad \forall e \in E. \end{aligned} \tag{D}$$

Proof. Clearly (d^*, q^*) is feasible for (D). Also (10) gives $x' \in \mathcal{F}_0$, while (13) implies that if $x'_e > 0$ then $1 = \rho_e(1, x'_e)$. This implies that $x'_e \leq \nu_e$, so x' is feasible for (P). If $x'_e > 0$ then certainly the arc is active (formally, by (13)) and hence $d_w^* = d_v^* + \tau_e + q_e^*$. And if $q_e^* > 0$, also implying that e is active, then (11) implies that $1 \leq \rho_e(1, x'_e) = x'_e/\nu_e$, which yields $x'_e = \nu_e$. This proves that x' and (d^*, q^*) are complementary solutions, and hence are optimal for (P) and (D) respectively. \square

According to this result, if a dynamic equilibrium eventually settles to a steady state then the corresponding queue lengths must be optimal for (D). Generically (after perturbing capacities) this linear program has a unique solution in which case the steady state is fully characterized. Otherwise, if (D) has multiple solutions it is not evident which queue lengths will be obtained in steady state. Note that even if the min cost flow for (P) is unique, this does not mean that only one steady state situation is possible because there may be flexibility in the queue lengths. For instance, if $u_0 = 1$ and the network has a single

link from s to t of unit capacity, if we create a queue of some length at time 0 this queue will remain in the steady state solution. This point will be further discussed in Section 5.3.

Remark. It is not difficult to show that when we start with initial conditions $z_e(0) = z_e^*$ where $z_e^* = \nu_e q_e^*$ with q^* optimal for (D), then the dynamic equilibrium is already at a steady state and the queues remain constant.

4. Convergence to a steady state

In this section we prove that a steady state exists and that it is actually reached in finite time. To this end we introduce a Lyapunov potential function that increases along the evolution of the dynamic equilibrium. The potential function is inspired from the previous dual program and is given by

$$\Phi(\theta) := u_0(\ell_t(\theta) - \ell_s(\theta)) - \sum_{e \in E} z_e(\ell_v(\theta)).$$

Remark. The potential is the difference between the total travel time experienced by users leaving at time θ , and the total queue volumes, as seen by users leaving at time θ . We are not aware of a more insightful interpretation of it. In fact, none of the more “natural” quantities we tried as candidate potential functions (total delay, time spent queueing, total delay excluding queueing delays, ...) are monotone.

THEOREM 2. $\Phi'(\theta)$ is nonnegative for all θ and strictly positive unless the dynamic equilibrium has reached a steady state.

Proof. The queues can be expressed as $z_e(\ell_v(\theta)) = \nu_e [\ell_w(\theta) - \ell_v(\theta) - \tau_e]_+$, and therefore $\Phi(\theta) = u_0(\ell_t(\theta) - \ell_s(\theta)) - \sum_{e \in E} \nu_e [\ell_w(\theta) - \ell_v(\theta) - \tau_e]_+$. To take the derivative we recall that E'_θ is the set of active edges, i.e., those for which $\ell_w(\theta) - \ell_v(\theta) \geq \tau_e$, while in E_θ^* the inequality is strict. Using the derivative of a max function and taking a NTFR (ℓ', x') at time θ , we thus obtain

$$\Phi'(\theta) = u_0(\ell'_t - \ell'_s) - \sum_{e \in E'_\theta \setminus E_\theta^*} \nu_e [\ell'_w - \ell'_v]_+ - \sum_{e \in E_\theta^*} \nu_e (\ell'_w - \ell'_v).$$

Notice that the dependency of the τ_e 's in the previous derivative is somewhat hidden in the set of active edges E'_θ . Now, for $e \in E'_\theta \setminus E_\theta^*$ we have $\ell'_w \leq \rho_e(\ell'_v, x'_e) = \ell'_v$ if $x'_e = 0$ and $\ell'_w = \rho_e(\ell'_v, x'_e) \geq \ell'_v$ if $x'_e > 0$, so that letting $E_\theta^+ = E_\theta^* \cup \{e \in E'_\theta \setminus E_\theta^* : x'_e > 0\}$ we may write

$$\Phi'(\theta) = u_0(\ell'_t - \ell'_s) - \sum_{e \in E_\theta^+} \nu_e (\ell'_w - \ell'_v).$$

Let us introduce a return arc ts with capacity $\nu_{ts} = u_0$ and flow $x'_{ts} = u_0$ so that x' is a circulation. Let $E_\theta^r = E_\theta^+ \cup \{ts\}$ and for each $e = vw \in E_\theta^r$ define the function

$$H_e(z) = \begin{cases} 1 & \text{if } \ell'_v \leq z < \ell'_w \\ -1 & \text{if } \ell'_w \leq z < \ell'_v \\ 0 & \text{otherwise.} \end{cases}$$

Then the derivative $\Phi'(\theta)$ can be expressed as

$$\Phi'(\theta) = - \int_0^\infty \sum_{e \in E_\theta^r} \nu_e H_e(z) dz.$$

For the remainder of the proof, let $\delta^+(S)$ (respectively $\delta^-(S)$) denote the arcs in E_θ^r leaving (respectively entering) S . Let $V_z = \{v : \ell'_v \leq z\}$ and consider an arc $e = vw \in E_\theta^+$. If $e \in \delta^+(V_z)$ then $\ell'_v \leq z < \ell'_w$ and therefore $\ell'_w = x'_e / \nu_e$. Similarly, if $e \in \delta^-(V_z)$ then $\ell'_w \leq z < \ell'_v$ which implies $e \in E_\theta^*$ and again $\ell'_w = x'_e / \nu_e$. Hence $x'_e = \nu_e \ell'_w$ for all $e \in E_\theta^+ \cap \delta(V_z)$. This equality also holds for the return arc ts , while in the remaining arcs $x'_e = 0$. Hence

$$\sum_{e \in \delta^+(V_z)} \nu_e z \leq \sum_{e=vw \in \delta^+(V_z)} \nu_e \ell'_w = \sum_{e \in \delta^+(V_z)} x'_e = \sum_{e \in \delta^-(V_z)} x'_e = \sum_{e=vw \in \delta^-(V_z)} \nu_e \ell'_w \leq \sum_{e \in \delta^-(V_z)} \nu_e z \quad (14)$$

with strict inequality if $\delta^+(V_z)$ is nonempty. It follows that for all $z > 0$ we have

$$\sum_{e \in E_\theta^r} \nu_e H_e(z) = \sum_{e \in \delta^+(V_z)} \nu_e - \sum_{e \in \delta^-(V_z)} \nu_e \leq 0$$

and therefore $\Phi'(\theta) \geq 0$ with strict inequality unless $\delta^+(V_z)$ is empty for almost all $z \geq 0$. Since for $\delta^+(V_z)$ to be empty we need that it either contains all vertices in V or none of them, we have that $\Phi'(\theta) = 0$ if and only if all ℓ'_v are equal, and hence (since $\ell'_s = 1$) all equal to 1. By Lemma 1, this exactly characterizes a steady state. \square

THEOREM 3. *Let $\bar{\nu} = \sum_{e \in C} \nu_e$ be the minimal queuing capacity among all st -cuts C . If $u_0 \leq \bar{\nu}$ then the dynamic equilibrium attains a steady state in finite time.*

Proof. From Theorem 2 it follows that there is some $\kappa > 0$ such that $\Phi'(\theta) \geq \kappa$ for every phase other than the steady state. This is simply because the thin flow depends only on the current shortest path network E'_θ and the set of queuing edges E_θ^* , and so there are only finitely many possible derivatives.

Thus, in order to prove that a steady state is reached in finite time it suffices to show that $\Phi(\theta)$ remains bounded. To this end we note that the condition $u_0 \leq \bar{\nu}$ implies that (P) is feasible and hence it has a finite optimal value α . The conclusion then follows by noting that the point (d, q) with $d_v = \ell_v(\theta) - \ell_s(\theta)$ and $q_e = z_e(\ell_v(\theta))/\nu_e$ is feasible for the dual (D) so that $\Phi(\theta) \leq \alpha$. \square

Given that convergence to a steady state does happen in finite time, it is natural to ask for explicit bounds. It is easy to see that a polynomial bound (in the input size encoding) is impossible; simply consider a network consisting of two parallel links, one with capacity $1 - 2^{-L}$ and length zero, the other with capacity 1 and length 1. The first phase, where all traffic takes the shorter edge, lasts until time $2^L - 1$. However, we can give a *pseudopolynomial* bound on the convergence time (and hence, queue lengths). The following results shows this bound. We present it in a slightly more general setting that allows for rational inflow and capacities and arbitrary initial queues since we will need it in this form in Section 5.2.

THEOREM 4. *Consider an instance for which all arc capacities ν_e as well as the inflow u_0 are multiples of $1/K$, $K \in \mathbb{Z}_+$. We allow for an arbitrary initial state at time 0 with possibly nonempty queues. Let $M = \sum_{e \in E} \nu_e$ and $T = \sum_{e \in E} (\tau_e + q_e(0))$. Then assuming the dynamic equilibrium attains a steady state, it is reached by time $2KMT$, and moreover, the waiting time in any queue never exceeds $2u_0K^2MT$.*

Proof. We first remark that it suffices to prove the result for $K = 1$, i.e., integer capacities and inflow. This follows by considering the instance where all arc capacities and the inflow have been scaled by a factor K ; this scaling has no impact on the time to reach steady state, nor the queue lengths. So assume $K = 1$ for the remainder.

We use the same notions defined in the proof of Theorem 2. Assume that a steady state is attained; thus (P) has a finite objective value, and this is at most $\sum_{e \in E} \tau_e$. Thus $\Phi(\theta) \leq \sum_{e \in E} \tau_e$ for all θ . Initially,

$$\Phi(0) = \ell_t(0) - \ell_s(0) - \sum_{e \in E} \nu_e q_e(0) \geq -M \sum_{e \in E} q_e(0).$$

So $\Phi(\theta) - \Phi(0) \leq T$ for all θ . Also clearly $u_0 \leq M$, since we assume that there is a steady state.

Consider some time θ which is not in the steady state phase; so $\Phi'(\theta) > 0$ by Theorem 2. Our first goal will be to show that $\Phi'(\theta) \geq 1/(2M)$; this clearly implies the bound on the time to reach steady state.

Let θ be any time before steady state is reached. Let $z_1 = \min_{v \in V} \ell'_v(\theta)$ and $z_2 = \max_{v \in V} \ell'_v(\theta)$. From the proof of Theorem 2, we have $\sum_{e \in E_\theta^r} \nu_e H_e(z) \leq -1$ for any $z \in (z_1, z_2)$ (it is strictly negative and integral). And certainly $H_e(z) = 0$ for all $e \in E_\theta^r$ and $z \notin [z_1, z_2]$.

Thus

$$\Phi'(\theta) = - \int_{z_1}^{z_2} \sum_{e \in E_\theta^r} \nu_e H_e(z) dz \geq z_2 - z_1.$$

To bound this, choose an arbitrary z for which $z_1 < z < z_2$ (recall that $z_1 < z_2$ since we have not reach steady state). We have, following the lines of (14),

$$z_1 \sum_{e \in \delta^-(V_z)} \nu_e \leq \sum_{e=vw \in \delta^-(V_z)} \nu_e \ell'_w = \sum_{e \in \delta^-(V_z)} x'_e = \sum_{e \in \delta^+(V_z)} x'_e = \sum_{e=vw \in \delta^+(V_z)} \nu_e \ell'_w \leq z_2 \sum_{e \in \delta^+(V_z)} \nu_e.$$

Thus

$$z_2 - z_1 \geq z_2 \left(1 - \frac{\sum_{e \in \delta^+(V_z)} \nu_e}{\sum_{e \in \delta^-(V_z)} \nu_e} \right) \geq \frac{1}{2M},$$

since $\sum_{e \in \delta^-(V_z)} \nu_e \leq M + u_0 \leq 2M$, and $z_2 \geq \ell'_s(\theta) = 1$.

Our next goal is to bound the queue lengths. The only extra ingredient we need is a bound on the speed at which a queue can grow. There is always a NTFR which does not route flow along cycles (see Koch 2012, Theorem 6.64), so that $x'_e \leq u_0$ for all $e \in E$. Therefore, $\max_v \ell'_v = \max_e x'_e / \nu_e \leq u_0$. Hence, for all times θ before steady state is reached,

$$\ell_t(\theta) - \ell_s(\theta) \leq \ell_t(0) - \ell_s(0) + (u_0 - 1)\theta \leq T + 2(u_0 - 1)MT.$$

This implies the same bound on all queue waiting times. \square

5. Some constructions and conjectures

While we have settled the finite-time convergence to a steady state, there are a number of questions about dynamic equilibria that remain open. In this section we provide some constructions exhibiting somehow surprising behavior. First we show that in a dynamic equilibrium the flow across a cut can be arbitrarily larger than the inflow. Then we build an instance for which the dynamic equilibria has exponentially many phases. We wrap up the section by discussing the possibility of characterizing the steady state queues, and some conjectures regarding more general steady state results.

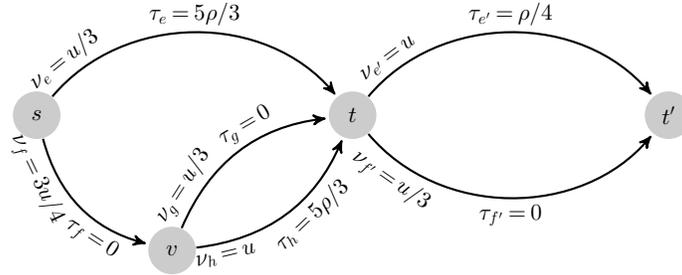


Figure 3 The modified instance.

5.1. Flow across a cut

As mentioned in the introduction a first conjecture would be that, similarly to what happens for static flows, the flow across any cut is always bounded by the inflow. This would provide a way to estimate the queues and to prove their boundedness. Unfortunately the property fails in a dynamic equilibrium. The reason for this is that flow entering the network at different times may experience different delays in such a way that they later superpose across an intermediate cut. In Example 1 we constructed one such instance, where the peak outflow rate was a factor $13/12$ larger than the inflow rate. We will now show how to construct instances in which the outflow is arbitrarily larger than the inflow.

The base of the construction, given in Figure 3, consists on appending to the instance in Example 1 (slightly rescaled for convenience) an extra sink t' and two arcs, e' and f' , connecting t with t' . Since the outflow in Example 1 (with $\tau = 5\rho/6$) is

$$\begin{cases} u/3 & \text{for } \theta \in [0, 5\rho/4) \\ 2u/3 & \text{for } \theta \in [5\rho/4, 3\rho/2) \\ 13u/12 & \text{for } \theta \in [3\rho/2, 5\rho/2) \\ u & \text{for } \theta \in [5\rho/2, \infty), \end{cases}$$

the outflow in the new instance will be

$$\begin{cases} u/3 & \text{for } \theta \in [0, 7\rho/4) \\ 13u/12 & \text{for } \theta \in [7\rho/4, 11\rho/4) \\ u & \text{for } \theta \in [11\rho/4, \infty). \end{cases}$$

Indeed, while the flow leaving t is $u/3$, it will directly take arc f' and therefore immediately reach t' . Then the flow leaving t increases to $2u/3$ and it will continue to choose arc f' , though a queue will start to build up on that arc. At time $3\rho/2$ (considering the local time

at t) the queue in f' will be of size $(u/3) \cdot (\rho/4)$, implying a queuing time of $\rho/4$. Since this is exactly the delay of arc e' , at this point e' enters the shortest path network. From this point the flow leaving t increases to $13u/12$ and thus the flow splits: $u/3$ flow units take arc f' while the remaining $9u/12$ take arc e' . The latter pattern stays until time $5\rho/2$ when the flow leaving t changes to u and thus it splits as $u/3$ taking arc f' and $2u/3$ taking arc e' , at which point steady state has been reached.

Some observations are in order:

- The length of the “pulse” in this construction is exactly ρ , and so this can be made as large as required.
- The full pulse is produced as long as the inflow is u for a period of $7\rho/4$; if it were to decrease or otherwise vary after this, it would only interfere with the final steady state phase.
- All arcs in the gadget have capacity at least $u/3$, meaning that no queues will form if the inflow is bounded by $u/3$. This means that the instance would still produce a pulse if the inflow was at most $u/3$ for some initial period, and then equal to u for a period (of length at least $7\rho/4$).

Building on these observations, we now construct an instance where the “pulse” is arbitrarily larger than the inflow. More precisely, we will construct a gadget $\text{PULSE}(u, k, q)$, for any $k \in \mathbb{N}$ and $u, q \in \mathbb{R}_+$, with the following properties (we define $\lambda := 13/12$ for convenience).

- (i) There is some value $\alpha_k = O((5/3)^k)$ such that, assuming a constant inflow rate of u in the interval $[0, (5/3)^k \rho]$, the inflow rate into the sink is at most $u/3$ for $\theta < \rho \alpha_k$ and exactly $\lambda^k \cdot u$ for $\theta \in [\rho \alpha_k, \rho \alpha_k + \rho]$ (where θ is the local time at the sink).
- (ii) All arcs have capacity at least $u/3$.
- (iii) The gadget has $6k$ arcs, with all free-flow delays being multiples of $3^{-k} \rho$ and bounded by $O((5/3)^k \rho)$, and all arc capacities being multiples of $12^{-k} u$ and bounded by $\lambda^k u$.

We have already seen in Figure 3 the construction of a $\text{PULSE}(u, 1, \rho)$ gadget; all the required properties clearly hold. We construct a $\text{PULSE}(u, k, \rho)$ gadget for $k \geq 1$ by chaining together in series a $\text{PULSE}(u, k-1, \frac{5}{3}\rho)$ gadget (call it G) followed by a $\text{PULSE}(\lambda^{k-1}u, 1, \rho)$ gadget (call it H). Properties (i)–(iii) are then mostly easy to verify inductively. An important observation is that since the outflow of G before its pulse is at most $u/3$, which

is smaller than the smallest capacity arc in H , this initial flow does not cause any disturbance. The bound on α_k follows easily by making the stronger inductive claim that $\alpha_k = \frac{21}{8}((5/3)^k - 1)$. Now $\alpha_1 = 7/4$, so this is correct for $k = 1$. For $k \geq 2$, the flow into H increases to $\lambda^{k-1}u$ at time $\rho\alpha_{k-1} \cdot (5/3)$, meaning that the outflow increases to $\lambda^k u$ at time $\rho\alpha_{k-1}(5/3) + (7/4)\rho = \rho\alpha_k$, and this the inductive claim holds.

5.2. Instances with an exponential number of phases

A natural hope would be that the number of phases is always polynomial in the input size (ideally as measured by the number of arcs in the instance, but failing that, as measured by the total encoding length of the instance). Unfortunately, as we show next, this is not the case and indeed the number of phases of a dynamic equilibrium may be exponential even in relatively simple series-parallel networks. This may help explaining why it is so notoriously difficult to practically compute dynamic equilibria in real-world networks (Wagner 2012, Friez and Han 2018).

The PULSE gadget of the previous section will be a first key ingredient. The next step of our construction is to use it to build a “damping” gadget. For any $k \in \mathbb{Z}_+$, $\rho \in \mathbb{R}_+$ we construct a gadget $\text{DAMPER}(k, \rho)$ with the following properties (recall that $\lambda := 13/12$).

- (i) There are values $Q = e^{O(k)}\rho$ and $\theta_1 + 2\rho < \theta_2 = e^{O(k)}\rho$ so that the following holds. Given an inflow of rate 1 in the interval $[0, Q)$, the gadget produces an outflow that is precisely 1 in the intervals $[\theta_1, \theta_1 + \rho)$ and $[\theta_2, \theta_2 + \rho)$, and precisely λ^{-k} in the interval $[\theta_2 - \rho, \theta_2)$.
- (ii) All arc capacities are multiples of 12^{-k} between $\lambda^{-k}/3$ and 1.
- (iii) Assuming an inflow rate that is always bounded by 1, the sum of queueing delays and free-flow delays within the gadget can never exceed $e^{O(k)}\rho$.
- (iv) The construction has $O(k)$ arcs, with total encoding length $O(k \cdot |\rho|)$ (where $|\rho|$ denotes the encoding length of ρ).

The construction of this gadget is shown in Figure 4 (the precise value of τ_f is discussed below). Initially, no flow uses arc f ; a queue grows on e , which sends the correct inflow of λ^{-k} into the $\text{PULSE}(\lambda^{-k}, k, \rho)$ gadget. After some time, the pulse gadget generates a pulse of size exactly 1, for a period of length ρ . We set θ_1 to the time that this pulse begins (as measured at t); by Property (i) of the PULSE gadget, $\theta_1 = O((5/3)^k \rho)$. Since all free-flow delays in the gadget are bounded by $O((5/3)^k \rho)$, and all capacities are multiples of 12^{-k}

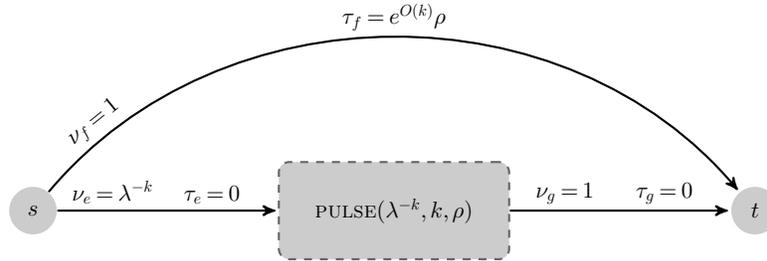


Figure 4 Construction of a $\text{DAMPER}(k, \rho)$ gadget.

bounded by 1, Theorem 4 yields a bound of $e^{O(k)}\rho$, both on the time to reach steady state, and also on the total delay within the pulse gadget. Once the pulse gadget has reached steady state, its outflow remains at λ^{-k} ; this will be the outflow of the DAMPER gadget as well, as long as f has not joined the shortest path network.

We now see how to choose τ_f ; large enough that f only joins the shortest path network after the pulse gadget has been sending outflow λ^{-k} for at least ρ amount of time. Since the queue on arc e grows at rate $\lambda^k - 1$, and the delay within the pulse gadget itself is at most $e^{O(k)}\rho$, it is clear that we can choose $\tau_f = e^{O(k)}\rho$. Once f does join the shortest path network, the entire DAMPER gadget reaches steady state, and the flow into t increases to 1. This determines θ_2 , and from this we can fix Q , ensuring that both are bounded by $e^{O(k)}\rho$. Property (iii) also follows immediately from the choice of τ_f .

Now we come to the construction of the gadget $\text{EXPONENTIAL}(d)$, which for any $d \in \mathbb{Z}_+$ will have size quadratic in d , and at least 2^d phases. The construction is recursive. In the following, C will denote a constant chosen large enough in relation to the hidden implicit constants in the definition of the DAMPER gadget; the precise requirements on C will become clear. We construct $\text{EXPONENTIAL}(1)$ by simply taking two parallel arcs, one of capacity $1/3$ and length 0, and the other of capacity $2/3$ and length 1; this clearly has two phases. To construct $\text{EXPONENTIAL}(d)$ for $d \geq 2$, take a $\text{DAMPER}(15d, C^{(d-1)^2})$ gadget (call it G), and follow this in series by an $\text{EXPONENTIAL}(d-1)$ gadget (call it H).

The idea behind this construction is that because the outflow λ^{-15d} from the damper gadget G during the damped period is smaller than the minimum arc capacity (of at least $\lambda^{-15(d-1)}/3$) in H , queues within the gadget will decrease. The length of the damped phase has been chosen to be long enough that all queues in H empty out completely (this is the only really delicate aspect of this construction). The two high outflow periods of G last

long enough that H runs (inductively) through 2^{d-1} phases during both periods, giving a total of at least 2^d phases.

The following properties about $\text{EXPONENTIAL}(d)$ are then straightforward to confirm inductively, exploiting also the properties of the DAMPER gadget.

- (i) Given a constant inflow of 1 in the interval $[0, C^{d^2})$, the gadget goes through 2^d phases.
- (ii) All arc capacities are multiples of 12^{-15d} between $\lambda^{-15d}/3$ and 1.
- (iii) Assuming an inflow rate that is always bounded by 1, the sum of queueing delays and free-flow delays within the gadget can never exceed C^{d^2-d} .
- (iv) The gadget consists of less than $(10d)^2$ arcs, and has encoding length $O(d^4)$.

Property (iii) requires that C is chosen large enough that the bound on the total delays in G given by Property (iii) is smaller than $C^{k/15-2}\rho$, which for $k = 15d$ and $\rho = C^{(d-1)^2}$ yields a bound of C^{d^2-d-1} . Inductively the total delays in H sum to at most $C^{(d-1)^2-(d-1)} = C^{d^2-3d}$. Overall, we obtain a bound of at most $C^{d^2-d-1} + C^{d^2-3d} \leq C^{d^2-d}$.

Let us now see that all queues in H do empty out during the damped period. By Properties (ii) and (iii), Theorem 4 tells us that the time to reach steady state from the beginning of the damped period is at most

$$2C^{(d-1)^2-(d-1)} \cdot (10d)^2 \cdot 12^{15d} \leq C^{d^2-3d+2} \cdot C^{d-1} = C^{(d-1)^2},$$

which is the length of the damped phase. (We assume in the above that C is large enough that $2(10d)^2 \cdot 12^{15d} \leq C^{d-1}$.) Since the inflow into H in the damped period, namely $\lambda^{-15d} < \lambda^{-15(d-1)}/3$, is lower than the capacity of any arc in H , the steady state necessarily has no queues.

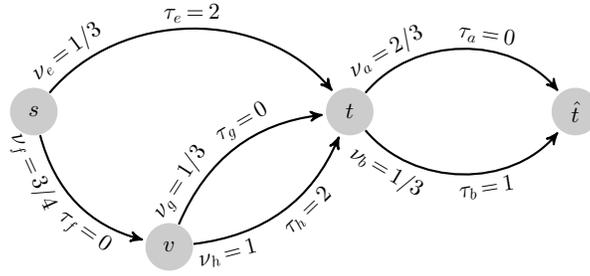
While we know that the number of phases may be very large, it is natural to expect that there are only a *finite* number of phases. While we conjecture that this is true, it is not ruled out by our result. Our result does show that if the length of all phases in the evolution is bounded away from zero, then there can only be a finite number. It is not ruled out, however, that an infinite sequence of phases occurs in a finite amount of time. This is the same issue discussed in Cominetti et al. (2015). The issue is significant; it is the one obstacle to showing uniqueness (in an appropriate sense) of dynamic equilibria.

If such a result could be shown, an even stronger conjecture would be that the number of phases is pseudopolynomially bounded in the input size. This would show that the exponential capacities and free-flow delays in the EXPONENTIAL gadget construction are in fact necessary.

5.3. Steady state queue lengths

Knowing that the dynamic equilibrium always reaches a steady state, a natural question is whether steady state queues can be characterized without having to compute the full equilibrium evolution. While we already observe that this is the case when the dual problem (D) has a unique solution, which occurs generically, the following example suggests that this is likely not possible in general.

EXAMPLE 2. Consider the network of Example 1, setting $\tau = 2$ and $u = 1$, with an extra node \hat{t} , which becomes the new sink, and two additional arcs, $a = (t, \hat{t})$ and $b = (t, \hat{t})$. Let $\nu_a = 2/3$, $\nu_b = 1/3$, $\tau_a = 0$, and $\tau_b = 1$. Clearly, up to time $3 + 3/5$ all flow will simply take



arc a and will not queue at t . Therefore we can ignore this initial phase, and the queues that will form at equilibrium in arcs a and b are the same as those that we would have in a network consisting of just nodes t (the source) and \hat{t} (the sink) and inflow

$$u_0(\theta) = \begin{cases} 13/12 & \text{for } \theta \in [0, 2 + 2/5) \\ 1 & \text{for } \theta \in [2 + 2/5, \infty). \end{cases}$$

In this instance all flow will take arc a for time $\theta \in [0, 8/5)$, forming a queue $z_e(8/5) = 2/3$. At this point flow will start splitting between arcs a and b in proportions $2/3, 1/3$, implying that queues will grow on both arcs until time $2 + 2/5$ where the steady state is achieved. The steady state queues will thus be $z_a^* = 32/45$ and $z_b^* = 1/45$. This example shows that the steady state queues are not minimal in any reasonable sense and that, furthermore, slightly changing the instance (e.g. τ_4) will change the steady state queues. Furthermore, if we slightly increase the capacity of arc b , say to $1/3 + \varepsilon$ the steady state queues jump to $z_a^* = 2/3$ and $z_b^* = 0$.

Additionally, one can observe from a slight variant of this instance, namely taking τ large and $\nu_b = 1/3 + \varepsilon$, that queues may grow very large in the transient and then go down to zero at steady state.

5.4. Conjectures on more general steady state results

Suppose that the inflow into the network is not a constant, but a time varying function $u_0(\theta)$. Suppose moreover that $u_0(\theta)$ is always bounded by the min-cut capacity of the network. Of course, there cannot be convergence to a steady state in this setting; but it is natural to expect that queues stay bounded. We conjecture that this is the case. It is not clear how our potential argument can aid in proving this conjecture. In particular, note that the boundedness of Φ alone is not helpful, as this does not imply any bounds on the queue sizes.

Suppose now that the inflow is constant, but *larger* than the min-cut capacity. Then again the evolution can of course not converge to a steady state in the way that we have defined it: queues cannot remain bounded. However, we conjecture that it is still true that after a finite amount of time, the equilibrium settles into a final phase that lasts forever.

Acknowledgments

We thank the organizer of the Dagstuhl Seminar “*Dynamic Traffic Models in Transportation Science*” held in 2015, where this work was initiated. We sincerely thank Vincent Acary, Umang Bhaskar and Martin Skutella for enlightening discussions, and Jonas Isreal for helpful comments on Theorem 4. We also thank the Associate Editor and two reviewers for their comments that helped improve the presentation.

This work was partially supported by ANID through grants FONDECYT 1190043, FONDECYT 1171501, Basal AFB-170001 and Basal AFB-180003; and by the Dutch Science Foundation (NWO) through a TOP grant (614.001.510) and a Vidi grant (016.Vidi.189.087).

Endnotes

1. These derivatives exist almost everywhere and are locally integrable.

References

- E. J. Anderson and A. B. Philpott. Optimisation of flows in networks over time. In F. P. Kelly, editor, *Probability, Statistics and Optimisation*, 369–382. Wiley, New York, 1994.
- Z. Cao, B. Chen, X. Chen, and C. Wang. A network game of dynamic traffic. EC 2017.
- R. Cominetti, J. Correa and O. Larré. Dynamic equilibria in fluid queueing networks. *Operations Research*, 63(1):21–34, 2015.
- L. Fleischer and E. Tardos. Efficient continuous-time dynamic network flow algorithms. *Operations Research Letters*, 23(3-5):71–80, 1998.
- L.R. Ford and D.R. Fulkerson. Constructing maximal dynamic flows from static flows. *Operations Research*, 6:419–433, 1958.

- L. R. Ford and D. R. Fulkerson. *Flows in Networks*. Princeton University Press, 1962.
- T.L. Friesz, D. Bernstein, T.E. Smith, R.L. Tobin, and B.W. Wie. A variational inequality formulation of the dynamic network user equilibrium problem. *Operations Research*, 41(1):179–191, 1993.
- The mathematical foundations of dynamic user equilibrium. T.L. Friesz and K. Han. *Transportation Research Part B: Methodological*, 126:309–328, 2018.
- D. Gale. Transient flows in networks. *Michigan Mathematical Journal*, 6:59–63, 1959.
- L. Graf and T. Harks. Dynamic flows with adaptive route choice. IPCO 2019.
- R. Koch. Routing Games over Time. Doctoral Thesis, TU-Berlin, 2012.
- R. Koch and M. Skutella. Nash equilibria and the price of anarchy for flows over time. *Theory of Computing Systems*, 49:71–97, 2011.
- D.K. Merchant and G.L. Nemhauser. A model and an algorithm for the dynamic traffic assignment problems. *Transportation Science*, 12:183–199, 1978.
- D.K. Merchant and G.L. Nemhauser. Optimality conditions for a dynamic traffic assignment model. *Transportation Science*, 12:200–207, 1978.
- F. Meunier and N. Wagner. Equilibrium results for dynamic congestion games. *Transportation Science*, 44(4):524–536, 2010.
- S. Peeta and A.K. Ziliaskopoulos. Foundations of dynamic traffic assignment: The past, the present and the future. *Networks and Spatial Economics*, 1(3-4):233–265, 2001.
- B. Ran and D. Boyce. *Modeling Dynamic Transportation Networks*. Springer, Berlin, 1996.
- L. Sering and M. Skutella. Multi-source multi-sink Nash flows over time. ATMOS 2018.
- L. Sering and L. Vargas Koch. Nash flows over time with spillback. SODA 2019.
- M. Skutella. An introduction to network flows over time. In W. Cook, L. Lovasz, and J. Vygen, editors, *Research Trends in Combinatorial Optimization*, 451–482, Springer, Berlin, 2009.
- W.S. Vickrey. Congestion theory and transport investment. *American Economic Review*, 59(2):251–260, 1969.
- N. Wagner. The dynamic user equilibrium on a transport network: mathematical properties and economic applications. Doctoral thesis, Université Paris-Est, 2012.
- Y.W. Xu, J.H. Wu, M. Florian, P. Marcotte, and D.L. Zhu. Advances in the continuous dynamic network loading problem. *Transportation Science*, 33(4):341–353, 1999.
- D. Zhu, and P. Marcotte. On the existence of solutions to the dynamic user equilibrium problem. *Transportation Science*, 34(4):402–414, 2000.

Short biographies of authors

José Correa is a Professor in the Department of Industrial Engineering at Universidad de Chile. His research interests are in game theory and mechanism design mostly from an algorithmic viewpoint. Recently he has mostly worked in data driven pricing when users are forward looking and in network games, particularly in dynamic environments.

Roberto Cominetti is a Professor with the Faculty of Engineering and Sciences at Universidad Adolfo Ibáñez, Santiago, Chile. His research interests include convex analysis and game theory, and their applications in transportation networks.

Neil Olver is an Assistant Professor in the Department of Mathematics at the London School of Economics and Political Science; he also holds a visiting appointment at CWI, Amsterdam. He works broadly in combinatorial optimization and adjacent areas, including probability, game theory, and theoretical computer science.