Long Term Behavior of Dynamic Equilibria in Fluid Queuing Networks

Roberto Cominetti[1], José Correa[2], and Neil Olver[3,4](✉)

[1] Facultad de Ingeniería y Ciencias, Universidad Adolfo Ibáñez, Santiago, Chile
[2] Facultad de Ingeniería Industrial, Universidad de Chile, Santiago, Chile
[3] Department of Econometrics and Operations Research, Vrije Universiteit Amsterdam, Amsterdam, Netherlands
n.olver@vu.nl
[4] CWI, Amsterdam, Netherlands

Abstract. 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.

1 Introduction

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

We warmly thank Schloss Dagstuhl for its hospitality during the Seminar 15412 *"Dynamic Traffic Models in Transportation Science"* at which this research started. We also express our sincere gratitude to Vincent Acary, Umang Bhaskar, and Martin Skutella for enlightening discussions. This work was partially supported by *Núcleo Milenio Información y Coordinación en Redes* (ICM-FIC RC130003), an NWO Veni grant, and an NWO TOP grant.

F. Eisenbrand and J. Koenemann (Eds.): IPCO 2017, LNCS 10328, pp. 161–172, 2017.
DOI: 10.1007/978-3-319-59250-3_14

delay $\tau_e \geq 0$ (see Fig. 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.

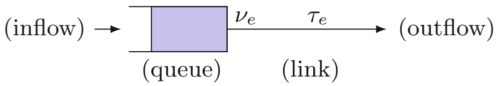

Fig. 1. An arc in the fluid queuing network.

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 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., [2–4,6,7,9–12]), their existence has only been proved recently by Zhu and Marcotte [13] though in a somewhat different setting, and by Meunier and Wagner [8] 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 [5] 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. [1], who proved existence and uniqueness, and provided a constructive proof for the existence of 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, so that no simple monotonicity arguments can be used to study the long term behavior.

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 or whether the dynamic equilibrium may exhibit Zeno-like oscillations in which queues alternate infinitely often over a finite time interval. In such a case the computation by integration would not yield a finite procedure. While this seems very unlikely, we have not been able to prove that it will never happen.

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 Sect. 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 Sect. 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 Sect. 5 we discuss some counterexamples that rule out some natural properties that one might expect to hold in a dynamic equilibrium, and we state some related open questions.

2 Dynamic Equilibria in Fluid Queing 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 [5] and Cominetti et al. [1].

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 $f_e^+ : [0, \infty) \to [0, \infty)$ 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 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

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

These dynamics uniquely determine the queue lengths $z_e(\theta)$ as well as the arc outflows (Fig. 2)

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

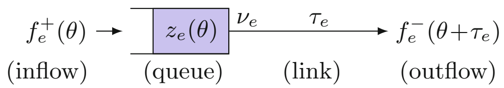

Fig. 2. Dynamics of an arc in the queuing network.

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 $v \in V\backslash\{t\}$, namely for a.e. $\theta \geq 0$

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

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

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

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

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

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

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

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

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

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

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 a.e. $\xi \geq 0$ 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

$$F_e^+(\theta) = \int_0^\theta f_e^+(z)\, dz$$
$$F_e^-(\theta) = \int_0^\theta f_e^-(z)\, dz.$$

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

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

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

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

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,

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

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,

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

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

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

where

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

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

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

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 [1] for the existence and uniqueness of NTFR's and to [5] 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'.

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. □

Theorem 2. *Consider a steady state with queues $z_e^* \geq 0$ and let d_v^* be the minimum travel times with arc times $\tau_e^* = \tau_e + q_e^*$ where $q_e^* = z_e^*/\nu_e$. Let (ℓ', x') with $\ell'_v = 1$ be a corresponding* NTFR *and denote by $\mathcal{F}_0$ the set of st-flows satisfying* (10)*. Then x' and (d^*, q^*) are optimal solutions for the following pair of dual linear programs:*

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

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

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 by (13) the arc e is active, and hence $d_w^* = d_v^* + \tau_e + q_e^*$. And if $q_e^* > 0$, 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. □

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 Example 3 in Sect. 5.

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

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

Theorem 3. *$\Phi'(\theta)$ is nonnegative for all θ and strictly positive unless the dynamic equilibrium has reached a steady state.*

Proof. The queues can be expressed as $z_e(\ell_v(\theta)) = \nu_e \left[\ell_w(\theta) - \ell_v(\theta) - \tau_e\right]_+$. Using the derivative of a max function and taking a NTFR (ℓ', x') at time θ, we thus obtain

$$\Phi'(\theta) = u_0(\ell_t' - \ell_s') - \sum_{e\in E_\theta'\setminus E_\theta^*} \nu_e[\ell_w' - \ell_v']_+ - \sum_{e\in E_\theta^*} \nu_e(\ell_w' - \ell_v').$$

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 $\delta(S)$ denote the edges in E_θ^r crossing S (and similarly for $\delta^+(S)$ and $\delta^-(S)$). 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

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

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

$$\sum_{e\in E^r_\theta} \nu_e H_e(z) = \sum_{e\in\delta^+(V_z)} \nu_e - \sum_{e\in\delta^-(V_z)} \nu_e \le 0$$

and therefore $\Phi'(\theta) \ge 0$ with strict inequality unless $\delta^+(V_z)$ is empty for almost all $z \ge 0$. The latter occurs iff all ℓ'_v are equal. Since $\ell'_s = 1$ it follows that $\Phi'(\theta) = 0$ iff $\ell'_v = 1$ for all v which by Lemma 1 characterizes a steady state. □

Theorem 4. *Let $\bar\nu = \sum_{e\in C} \nu_e$ be the minimal queuing capacity among all st-cuts C. If $u_0 \le \bar\nu$ then the dynamic equilibrium attains a steady state in finite time.*

Proof. From Theorem 3 it follows that there is some $\kappa > 0$ such that $\Phi'(\theta) \ge \kappa$ 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 $u_0 \le \bar\nu$ 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 $\Phi(\theta) \le \alpha$. □

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 (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. However, we can give a *pseudopolynomial* bound on the convergence time (and hence, queue lengths).

Theorem 5. *Consider an instance for which $u_0 \in \mathbb{Z}_+$ and $\nu_e \in \mathbb{Z}_+$ for all $e \in E$. Let $M = \sum_{e\in E} \nu_e$ and $T = \sum_{e\in E} \tau_e$. Then assuming the dynamic equilibrium attains a steady state, it is reached by time $O(MT)$, and moreover, the waiting time in any queue never exceeds $O(M^3T)$.*

The argument to bound the convergence time involves showing that the difference between the smallest and largest label derivative is not too small. Combining this with an upper bound on the rate at which any queue can grow yields the second claim. We delay the proof to the full version of the paper.

5 Some Conjectures and Counterexamples

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 discuss some conjectures and provide counterexamples to some of them.

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. The following instance with unit inflow $u_0 = 1$ exhibits an outflow rate of 13/12 during a time interval.

Example 1. Consider the network consisting of the vertices $\{s, v, t\}$ with edges $e_1 = (s, t), e_2 = (s, v), e_3 = (v, t), e_4 = (v, t)$ and inflow $u_0 = u$. Capacities are $\nu_1 = u/3$, $\nu_2 = 3u/4$, $\nu_3 = u/3$, and $\nu_4 = u$, and delays are $\tau_1 = \tau_4 = \tau$, and $\tau_2 = \tau_3 = 0$. In this instance one can compute the derivative of the distance labels at node t as

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

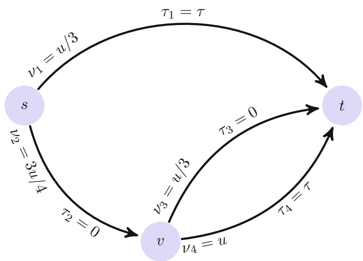

Thus the amount of flow arriving at t at time $\ell_t(\theta)$ can readily be computed as $u/\ell_t'(\theta)$. If we consider the local time at node t this flow is then

$$f_1^-(\theta) + f_3^-(\theta) + f_4^-(\theta) = \begin{cases} u/3 & \text{for } \theta \in [0, 3\tau/2) \\ 2u/3 & \text{for } \theta \in [3\tau/2, 9\tau/5) \\ 13u/12 & \text{for } \theta \in [9\tau/5, 3\tau) \\ u & \text{for } \theta \in [3\tau, \infty). \end{cases}$$

By chaining together slightly modified copies of this instance, one can blow up the maximum outflow to any desired quantity, even with unit inflow. Notice that the length of the "pulse" in the above construction can be made as large as required, by choosing τ appropriately. This pulse can be used to drive a second copy of the construction, with larger u. Figure 3 shows the construction with two copies; there is a phase with outflow $(13/12)^2$. The phases before the pulse of the left gadget only produce a queue on e', which has no impact on the behavior except for essentially shortening e' and h'. We delay the details to the full version of the paper.

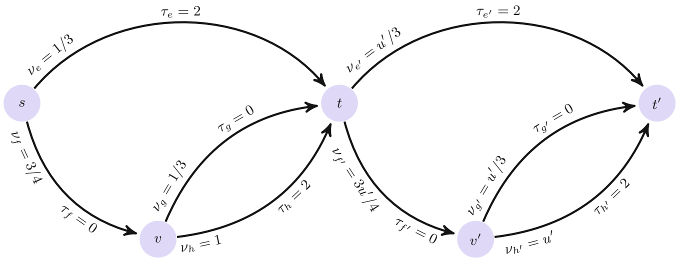

Fig. 3. Creating a larger pulse. Here, $u' = 13/12$.

Even though the previous example shows that intermediate flows can grow very large our main result states that a steady state is actually reached after finite time. This indeed implies that the queues remain bounded along the evolution of a dynamic equilibrium. However this also raises further questions. Indeed it is unclear whether the steady state is attained after finitely many phases of the Koch-Skutella algorithm. It is conceivable that in some situations the phases become shorter and shorter and that infinitely many of them occur in the finite time span before steady state is reached. The next example shows that there may actually be an exponential (in the input size) number of phases.

Example 2. Here we sketch the construction of an instance with an exponential number of phases; we defer the details to the full version of the paper. More precisely, for a given d we construct an instance with $\Omega(2^d)$ phases and $O(d^2)$ nodes. The main idea is to construct a "2-pulse" gadget, based on the "1-pulse" gadget described in Example 1. The outflow rate of this gadget has two, well-separated, periods where the outflow is large; outside of these two periods, the outflow is much smaller. Given a gadget with $\Omega(2^d)$ phases (call it H), we construct one with $\Omega(2^{d+1})$ phases roughly as follows. We begin with the 2-pulse gadget. To the output of this gadget, we attach both a single edge of small capacity and length 0 to the sink; and in parallel, we attach H. In between pulses, all flow uses the short low-capacity edge, and any queues in H decay. During each of the two pulses, flow enters H; inductively, this yields $\Omega(2^d)$ phases each time.

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 3. 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.

References

1. Cominetti, R., Correa, J., Larré, O.: Dynamic equilibria in fluid queueing networks. Oper. Res. **63**(1), 21–34 (2015)
2. Ford, L.R., Fulkerson, D.R.: Constructing maximal dynamic flows from static flows. Oper. Res. **6**, 419–433 (1958)
3. Friesz, T.L., Bernstein, D., Smith, T.E., Tobin, R.L., Wie, B.W.: A variational inequality formulation of the dynamic network user equilibrium problem. Ope. Res. **41**(1), 179–191 (1993)
4. Gale, D.: Transient flows in networks. Mich. Math. J. **6**, 59–63 (1959)
5. Koch, R., Skutella, M.: Nash equilibria and the price of anarchy for flows over time. Theor. Comput. Syst. **49**, 71–97 (2011)
6. Merchant, D.K., Nemhauser, G.L.: A model and an algorithm for the dynamic traffic assignment problems. Transp. Sci. **12**, 183–199 (1978)
7. Merchant, D.K., Nemhauser, G.L.: Optimality conditions for a dynamic traffic assignment model. Transp. Sci. **12**, 200–207 (1978)
8. Meunier, F., Wagner, N.: Equilibrium results for dynamic congestion games. Transp. Sci. **44**(4), 524–536 (2010)
9. Peeta, S., Ziliaskopoulos, A.K.: Foundations of dynamic traffic assignment: the past, the present and the future. Netw. Spat. Econ. **1**(3–4), 233–265 (2001)
10. Ran, B., Boyce, D.: Modeling Dynamic Transportation Networks. Springer, Berlin (1996)
11. Vickrey, W.S.: Congestion theory and transport investment. Am. Econ. Rev. **59**(2), 251–260 (1969)
12. Xu, Y.W., Wu, J.H., Florian, M., Marcotte, P., Zhu, D.L.: Advances in the continuous dynamic network loading problem. Transp. Sci. **33**(4), 341–353 (1999)
13. Zhu, D., Marcotte, P.: On the existence of solutions to the dynamic user equilibrium problem. Transp. Sci. **34**(4), 402–414 (2000)