Continuity, Uniqueness and Long-Term Behavior of Nash Flows Over Time*

Neil Olver[1,2], Leon Sering[3], and Laura Vargas Koch[3,4]

[1]Department of Mathematics, London School of Economics and Political Science
[2]Centrum Wiskunde & Informatica, Amsterdam
[3]Department of Mathematics, ETH Zürich
[4]School of Business and Economics, RWTH Aachen University

November 15, 2021

Abstract

We consider a dynamic model of traffic that has received a lot of attention in the past few years. Users control infinitesimal flow particles aiming to travel from a source to destination as quickly as possible. Flow patterns vary over time, and congestion effects are modeled via queues, which form whenever the inflow into a link exceeds its capacity. Despite lots of interest, some very basic questions remain open in this model. We resolve a number of them:

- We show *uniqueness* of journey times in equilibria.
- We show *continuity* of equilibria: small perturbations to the instance or to the traffic situation at some moment cannot lead to wildly different equilibrium evolutions.
- We demonstrate that, assuming constant inflow into the network at the source, equilibria always settle down into a "steady state" in which the behavior extends forever in a linear fashion.

One of our main conceptual contributions is to show that the answer to the first two questions, on uniqueness and continuity, are intimately connected to the third. Our result also shows very clearly that resolving uniqueness and continuity, despite initial appearances, cannot be resolved by analytic techniques, but are related to very combinatorial aspects of the model. To resolve the third question, we substantially extend the approach of [CCO21], who show a steady-state result in the regime where the input flow rate is smaller than the network capacity.

1 Introduction

Motivated especially by congestion in transportation networks and communication networks, the study of routing games has received a huge amount of attention. Most of this work concerns *static* models; that is, the model posits a constant, unchanging demand, and a solution is represented by some kind of flow. Congestion effects are modeled via a relationship between the amount of traffic using a particular link, and the resulting delay experienced. Many variants have been considered: nonatomic games (where each individual player controls an infinitesimal amount of flow), atomic games with a finite number of players, multi-commodity and single-commodity settings, different choices of congestion functions, and much more. Much is understood about equilibrium behavior: conditions for existence and uniqueness; bounds on the *price of anarchy* and

*N.O. is partially supported by NWO Vidi grant 016.Vidi.189.087.

various related notions; the phenomenon of Braess's paradox; tolling to improve equilibrium efficiency; and so on (see [Rou05] for a survey of the area).

While these static models can be a good approximation in many situations, this is not always the case. More recently, there has been a lot of interest in models that are explicitly *dynamic*—that is, time-varying. A canonical situation to consider is morning rush-hour traffic; clearly there are substantial variations of traffic behavior and congestion over time. Another motivating setting is the routing of packets in communication networks, which traverse through a network of limited bandwidth over time, and are processed in queues at the nodes of the network.

In the static case, models typically allow for some relationship between traffic density and delay to be posited (for example, a linear relationship, or something more refined based on empirical data). This is substantially more difficult to do in the dynamic setting. It is rather crucial to maintain a *first-in-first-out* property for flow on an arc; overtaking is questionable from a modeling perspective, and also introduces various pathologies. This property is unfortunately easily to violate; for example, if one attempts to specify the delay a user experiences as a function of the inflow rate into a link at the moment of entry, a sharp decrease in inflow will lead to overtaking. Well-behaved models in this generality require detailed modeling of traffic along links (as opposed to describing the traffic on a link via a single time-varying value); an example of a model taking this approach is the LWR, or kinematic wave, model [LW55]. These models are very challenging to analyze, even for a single link. Fortunately, there is one model that is both natural, and quite relevant to real-world traffic behavior (as well as communication network settings). This is the *fluid queueing model*, also known as the *deterministic queueing model* or the *Vickrey bottleneck model* [Vic69], and it is this model that will concern us.

In the fluid queueing model, each link has a *capacity* and a *transit time*. If the inflow rate into the link always remains below its capacity, then the time taken to traverse the link is constant, as given by the transit time. However, if the inflow rate exceeds the link capacity for some period, a queue grows on the entrance of the link. The delay experienced by a user is then equal to the transit time, plus whatever time is spent waiting in the queue. As long as there is a queue present it will empty at rate given by the link capacity; depending on whether the inflow rate is smaller or larger than the capacity, this queue will decrease or increase size. (See Figure 1; the notation will be fully described in Section 2.) Note that the queues are considered to be *vertical*, meaning they can hold an unlimited amount of flow (*horizontal* or *spatial* queues have been considered in [SVK19, Ser20]).

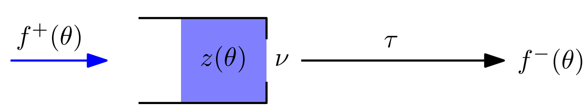

Figure 1: The dynamic of an arc at snapshot time θ. The inflow $f^+(\theta)$ and outflow $f^-(\theta)$ describe the flow entering or leaving the arc at time θ. The amount of flow in the queue $z_e(\theta)$ can leave the queue with rate ν and afterwards traverses the arc, which takes τ units of time.

Consider a network of such links, and suppose that all traffic in the network has the same origin and destination. Clearly this is a restrictive assumption, but the single-commodity setting is already very challenging. Starting at time 0, users are released into the system at some constant rate; each user controls an infinitesimal amount of flow. Their joint choices yield a *flow over time*, as first introduced by Ford and Fulkerson in the fifties [FF56]. But here, each user is self-interested and wants to find a *quickest* route to the destination. This will depend on the choices of the other users, since these will impact the queue lengths. A Nash equilibrium in this model is then a joint choice of routes for all users, such that all users are satisfied in hindsight with their choices: no user can switch to an alternative route and arrive at the destination at a strictly earlier time. Note that this means that when making a decision, it is not the queueing delays *now* that matter, but the queues that the user would actually experience upon reaching an arc, which may be different.

An alternative notion of "instantaneous dynamic equilibria" [GHS20] has also been considered, but will not concern us here.

A lot of attention has also been paid to (essentially) discrete versions of this model, where individual indivisible packets must be routed through a network. This has been studied from the perspective of algorithms and optimization (e.g., [LMR94, LMR99]) and in various variants for self-interested strategic users. While e.g., [CCCW17, Ism17, SST18, SVKZ21] examine the discrete version of the deterministic queueing model, Hoefer et al. [HMRT11] (see also [KM15]) introduce *temporal congestion games*, where arcs represent machines that must process a packet (job) before it can continue onto a next machine and [HPS⁺18, SSVK18] examine different rules to handle the interaction of the packets. The latter two variants do not have a first-in-first-out requirement, and instead the emphasis is on obtaining good scheduling rules.

We return to discussing the fluid queueing model. There has been some substantial progress in understanding equilibria in this model. Koch and Skutella [KS11] (see also [Koc12]) were the first to study it using tools from combinatorial optimization. They uncovered an intricate structure in the time derivatives of equilibria, which they could precisely describe as a solution to a certain nonlinear system (which they called *thin flow with resetting*, and which we will encounter later). Assuming solutions to this system always exist, they showed how an equilibrium could be constructed by essentially integrating; the resulting equilibrium has the property that all queue lengths are piecewise linear. They call a maximal interval of departure times on which queues are affine a *phase*. (Note that there is no guarantee that there are a finite number of phases, even within a compact interval; that is, the algorithm is not known to be finite).

Existence was shown with precision by Cominetti, Correa and Larré [CCL15]. They showed that the nonlinear system of Koch and Skutella always has a solution (in fact a unique solution), and hence that the integration procedure always succeeds in constructing an equilibrium. Surprisingly, it is an open question as to whether the thin flow equations can be efficiently solved; in other words, whether there is a polynomial-time algorithm to compute the next phase, given that the behavior for all previous phases has been computed.

A good upper bound on the price of anarchy (suitably defined—it is necessary to consider average arrival time rather than average journey time as the cost function to have any hope of positive results) also remains an open question. A constant $\frac{e}{e-1}$ bound is conjectured, and some partial progress has been made [CCO19, BFA15]. Cominetti, Correa and Olver [CCO21] answered a perhaps even more basic question about equilibrium efficiency: if the inflow at the source is not larger than the network capacity (that is, the minimum capacity of a cut separating s and t), do queue lengths and hence journey times remain bounded in an equilibrium? They give an answer in the affirmative; this gives at least some sense in which equilibria are well-behaved and at least not disastrously inefficient.

In this paper, we answer some significant open problems in the model.

Uniqueness. Are equilibria unique in this model? After existence, this is among the first questions to ask about equilibria in any model.

Some care is needed in phrasing the question. Imagine an inflow rate of 1 at the source, and two parallel links between s and t, both of capacity 1 or larger. Then flow can be split arbitrarily (and in an arbitrarily time-varying way) between the two links, and the result is an equilibrium. So there is no uniqueness at the level of flows. Instead, the right question concerns the *journey times* of users in equilibrium. All flow particles leaving at time θ incur the same (smallest possible) journey time in an equilibrium. The correct uniqueness question is not about the routes chosen in equilibrium, but the costs (journey times) experienced. Is this unique?

[CCL15] give a partial answer to this question. They show that if one restricts to "right differentiable" solutions (the precise meaning of this we postpone until later), then equilibria are unique with respect to journey times. First impressions might be that such a restriction is of a technical nature, and that this *essentially* shows uniqueness, but this is misleading. The precise reason for this requires more technical preliminaries, and we postpone this until Section 3. For now, we remark that there is no a priori reason why equilibria not respecting this right-differentiability requirement could not exist, nor anything "unphysical" about them.

We prove that this cannot happen, and that equilibria are indeed unique, without any assumptions. To further motivate why this uniqueness result is important, we will shortly discuss how it relates to the next question we resolve.

Continuity. If an equilibrium is disturbed in some way, does the disturbed equilibrium remain close in some sense to the undisturbed equilibrium, or could it veer off wildly in another direction? Such a property is *crucial* for the model to have any bearing on reality. Clearly, real traffic situations will not *exactly* match up with equilibria in the model, even under the most optimistic assumptions. Individual cars or packets are not really infinitesimally small; this is an approximation. Given two routes that have slightly different but almost-equal journey times, a user may not notice or be sufficiently concerned, and pick the slightly longer one. Actual travel delays on a link could vary slightly due to all sorts of factors. All of this is of course obvious—there is no expectation that the abstract, simplified model would capture all these real-world aspects. Nonetheless, the hope is that the model is a good one, in the sense that it captures qualitatively important aspects of the real situation, *and* that the closer some more complex (artificial or real-world) situation matches the conditions of the model, the closer the behavior of the model would match the more complex setting. If slightly perturbing the situation at some moment can lead to completely different equilibrium behavior, then any conclusions drawn from the model must be treated with extreme skepticism. Does the model really have anything to tell us, in this case?

It might seem that continuity should be a straightforward property to show; or at least, that it should be a matter of proficiency with analytic techniques. Similarly to the uniqueness question, this turns out to not be the case at all, and for very similar reasons. Again, we hold the explanation of this to the technical overview.

We are able to show continuity of equilibria with respect to a wide range of natural perturbations. For instance, if we change the capacity and/or transit times of some arcs in the network by very small amounts, the resulting equilibrium will not change too much.

We have already mentioned that the fluid queueing model is an approximation of reality where "small" users are replaced by infinitesimal ones. It is assumed that this replacement does not make much difference. To actually justify this, one would need to show some form of convergence of equilibria for discrete packet model to dynamic equilibria in the fluid queueing model. No such rigorous justification has been made so far, however: empirical evidence has been provided [ZSV+21], and rigorous results have been obtained for the *static* model [CSSSM21]. As some first steps, the convergence of the underlying packet model to the fluid queueing model has been proven (without considering equilibrium behavior) [SVKZ21]. We anticipate our work being a main ingredient in proving convergence of equilibria; it should be clear that without continuity, there is little hope for such a convergence result. We leave this for followup work.

It should seem natural, given the above discussion, that uniqueness and continuity are related. The following question, however, appears entirely unrelated. Surprisingly, this is not the case: one of our main conceptual contributions is to show how uniqueness and continuity follow from resolving it.

Long-term behavior. As already mentioned, [CCO21] prove that if inflow at the source is constant (starting from some initial time) and not larger than the capacity of the network, as measured by the minimum capacity of an s-t-cut, then queues remain bounded in an equilibrium. Their proof actually shows something stronger: Namely, they show that as long as this inflow condition is satisfied, after some (instance dependent) time the equilibrium will reach a "steady state", in which all queues remain constant from this time forward.

This raises a question. What happens if the inflow rate is larger than the minimum s-t-cut capacity? Clearly, queues can no longer stay bounded. But one may still ask whether the evolution eventually reaches a steady state—which no longer means a situation where no queues change (clearly impossible), but a situation where queues change linearly, forever into the future.

Just like the [CCO21] result, this can be viewed as a positive statement about the efficiency of equilibria. It turns out that there is only one possible rate vector at which queues can grow in a steady state. So if steady state is reached, it means that over a sufficiently long time horizon, the equilibrium does use the network in the most efficient way possible; queues only grow because they must, due to bottlenecks in the network. Further, this queue behavior at steady state can be efficiently characterized.

While this result is already interesting in its own right, a very slight generalization of this also turns out to be the main technical ingredient in our proof of uniqueness and continuity. In fact, it is *precisely* what one needs; there can be no way of showing uniqueness or continuity without showing this result about long-term behavior. It is in this sense that we mean that our result is fundamentally combinatorial, rather than analytic; our proof of convergence to steady state is, like the one of [CCO21], based on a non-obvious potential related to a primal-dual program that characterizes possible steady-state situations.

2 Model and preliminaries

An instance is described by a directed graph $G = (V, E)$ with a source s and sink t, where each arc is equipped with a transit time $\tau_e \geq 0$ and a capacity $\nu_e > 0$. At s we have a constant *network inflow rate* of $u_0 > 0$, which begins a time 0. We may assume that every node in G is both reachable from s, and can reach t. For technical convenience, we will follow previous works and assume that G has no directed cycle consisting of 0-length arcs.

We use the notation $\delta^-(v)$ and $\delta^+(v)$ to denote the set of incoming and outgoing arcs at v, respectively, and similarly $\delta^-(S)$ and $\delta^+(S)$ for arcs entering or leaving a set S. We define $[z]^+ := \max\{0, z\}$.

A *flow over time* is given by a family of locally-integrable functions $f_e^+ : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ and $f_e^- : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ that describe the inflow and outflow rate at each arc $e \in E$ at every point in time. The *cumulative* inflows and outflows are given by the absolutely continuous functions $F_e^+(\xi) := \int_0^\xi f_e^+(\xi')\,\mathrm{d}\xi'$ and $F_e^-(\xi) := \int_0^\xi f_e^-(\xi')\,\mathrm{d}\xi'$, respectively. A flow particle in the queue of link e at time ξ must have entered the link at time ξ or before, and will not leave the link until after time $\xi + \tau_e$. Thus the *queue volume* on a link e at time ξ, denoted by $z_e(\xi)$, is given by $z_e(\xi) := F_e^+(\xi) - F_e^-(\xi + \tau_e)$. Since queues always empty at maximum rate, the amount of time spent waiting in a queue is simply the queue volume upon arrival divided by the arc capacity. We call a flow over time *feasible* if it satisfies flow conservation at every node $v \in V$ for every point in time ξ:

$$\sum_{e \in \delta^+(v)} f_e^+(\xi) - \sum_{e \in \delta^-(v)} f_e^-(\xi) \begin{cases} = 0 & \text{for } v \in V \setminus \{ s, t \}, \\ = u_0 & \text{for } v = s, \\ \leq 0 & \text{for } v = t, \end{cases}$$

and if the queues empty at a rate given by the capacity:

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

The *earliest arrival time* at w of a particle starting at s at time $\theta \in \mathbb{R}_{\geq 0}$ is given by

$$\ell_w(\theta) := \begin{cases} \theta & \text{if } w = s, \\ \min\limits_{e = vw \in \delta^-(w)} \ell_v(\theta) + \frac{z_e(\ell_v(\theta))}{\nu_e} + \tau_e & \text{otherwise.} \end{cases}$$

We will often refer to the collection of earliest arrival times for a fixed θ as a *labeling*. Given earliest arrival times ℓ, we define the associated *queueing delays* (or *queueing delay functions*) by $q_e(\theta) \coloneqq z_e(\ell_v(\theta))/\nu_e$ for each $e = vw \in E$. Note that $q_e(\theta)$ is the queueing delay on arc e for a flow particle departing s at time θ and taking a shortest path to e. For any earliest arrival times ℓ, we define the *active* arcs $E'_{\ell(\theta)}$ and the *resetting* arcs $E^*_{\ell(\theta)}$ for some time θ as $E'_{\ell(\theta)} \coloneqq \{ e = vw \in E \mid \ell_w(\theta) = \ell_v(\theta) + q_e(\theta) + \tau_e \}$ and $E^*_{\ell(\theta)} \coloneqq \{ e = vw \in E \mid q_e(\theta) > 0 \}$, respectively. We will write E'_θ and E^*_θ as shorthand, if the choice of ℓ is clear. So E'_θ consists of all arcs that lie on a shortest path from s to some node in the network, from the perspective of a user that departs at time θ; and E^*_θ is the set of arcs where such a user would find a queue if they enter the arc as early as possible.

An *equilibrium* (also referred to as a *dynamic equilibrium* or a *Nash flow over time*) is a feasible flow over time in which almost all flow particles travel along a shortest path from s to t, i.e., only along active arcs. In this case the active and resetting arcs are characterized by

$$E'_{\ell(\theta)} = \{ e = vw \in E \mid \ell_w(\theta) \geq \ell_v(\theta) + \tau_e \} \quad \text{and} \quad E^*_{\ell(\theta)} = \{ e \in E \mid \ell_w(\theta) > \ell_v(\theta) + \tau_e \}. \tag{1}$$

For an arc $e = vw \in E^*_\theta$, the delay experienced by such a user departing at time θ is given by $q_e(\theta) = \ell_w(\theta) - \ell_v(\theta) - \tau_e$.

As proven in [CCL15], a feasible flow over time is an equilibrium if and only if $F_e^+(\ell_v(\theta)) = F_e^-(\ell_w(\theta))$ for all arcs $e = vw$ and all departure times θ. Define $x_e(\theta) \coloneqq F_e^+(\ell_v(\theta))$ for all e and θ. For an equilibrium flow, the derivative of x_e at time θ, which exists almost everywhere, can be interpreted as a flow describing what proportion of flow particles departing at time θ use arc e; x' is an s-t-flow of value u_0 almost everywhere. It has been shown [CCL15, KS11] that a flow over time is an equilibrium if and only if the resulting pair (x, ℓ) satisfies the following for almost every θ: setting $x' = x'(\theta)$, $\ell' = \ell'(\theta)$, $E' = E'_\theta$ and $E^* = E^*_\theta$,

$$x' \text{ is a static } s\text{-}t \text{ flow of value } u_0, \tag{TF-1}$$

$$\ell'_s = 1, \tag{TF-2}$$

$$\ell'_w = \min_{e = vw \in E'} \rho_e(\ell'_v, x'_e) \qquad \text{for all } w \in V \setminus \{ s \}, \tag{TF-3}$$

$$\ell'_w = \rho_e(\ell'_v, x'_e) \qquad \text{for all } e = vw \in E' \text{ with } x'_e > 0, \tag{TF-4}$$

$$\text{where} \qquad \rho_e(\ell'_v, x'_e) \coloneqq \begin{cases} \frac{x'_e}{\nu_e} & \text{if } e = vw \in E^*, \\ \max\{\ell'_v, \frac{x'_e}{\nu_e}\} & \text{if } e = vw \in E' \setminus E^*. \end{cases}$$

These are called the *thin flow* conditions (sometimes *normalized thin flows* or *thin flows with resetting*). Note that the conditions are fully determined by the pair E' and E^*. As long as $E^* \subseteq E'$, E' is acyclic and each node v is reachable from s, they always have a solution, and ℓ' is uniquely determined [CCL15]. We will sometimes call this unique ℓ' (leaving out x') the *thin flow direction*. For the case $E^* = \emptyset$, there always exists a unique solution [Koc12], even if E' contains cycles, and it can be computed efficiently.

We will take this viewpoint throughout, and always think of an equilibrium as a pair (x, ℓ) satisfying these conditions. Such a pair fully determines the flow over time (f^+, f^-) for the equilibrium; for example, the queue volume is determined by $\frac{1}{\nu_e} z_e(\ell_v(\theta)) = \max\{0, \ell_w(\theta) - \ell_v(\theta) - \tau_e\}$ for all $e = vw$ and $\theta \in \mathbb{R}_{\geq 0}$. Note that from this perspective, given (x, ℓ) satisfying the thin flow conditions for all times $\theta \in [0, T]$, extending the equilibrium to later departure times only requires knowing $\ell(T)$; nothing about the earlier (with respect to departure time) behavior of the equilibrium is needed. Just ℓ alone captures all the truly important information; one could easily verify whether a given curve ℓ can be extended to an equilibrium, since given $\ell'(\theta)$, it is easy to determine whether a matching $x'(\theta)$ exists so that $(x'(\theta), \ell'(\theta))$ satisfy the thin flow conditions. We call such an ℓ an *equilibrium trajectory*.

Note that for any configuration (E', E^*), the solution (x', ℓ') to the thin flow equations satisfy $\ell'_v \leq \kappa$, where $\kappa \coloneqq u_0 / \min_{e \in E} \nu_e$. Thus any equilibrium trajectory is κ-Lipschitz.

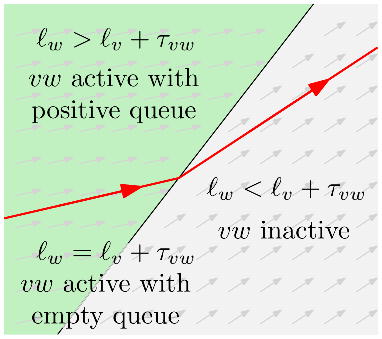

Figure 2: Dynamic equilibria can be seen as trajectories in $\mathbb{R}^V$ that follow a piecewise-constant vector field.

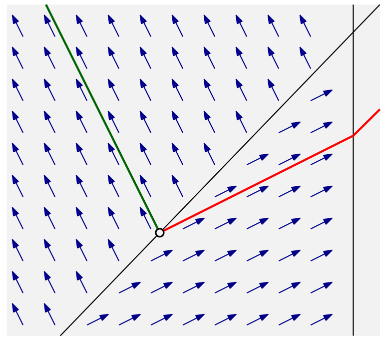

Figure 3: The simplest situation which would produce non-uniqueness (if it were possible).

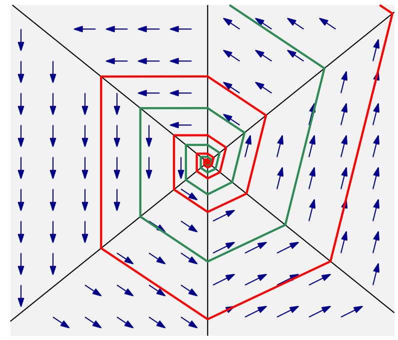

Figure 4: A more subtle potential situation which would produce non-uniqueness; we show that this cannot occur.

It has been shown [KS11, CCL15] that in an equilibrium, ℓ is a piecewise linear function; each linear segment is referred to as a *phase* of the equilibrium. However, it is not known whether the number of phases is necessarily finite, even within a compact interval.

We call a pair (E', E^*) with $E^* \subseteq E' \subseteq E$ a *valid configuration* if (i) for every node $v \in V$ there is an s-v-path in E', (ii) every arc $e \in E^*$ lies on an s-t-path in E', and (iii) either E' is acyclic or $E^* = \emptyset$. Furthermore, we call a vector $\ell^\Diamond \in \mathbb{R}^V$ *feasible* if $(E'_{\ell^\Diamond}, E^*_{\ell^\Diamond})$ is a valid configuration. Requirement (i) is equivalent to labels denoting the earliest arrival time starting from s; requirement (ii) ensures that flow in queues travels on active s-t-paths, which is the definition of dynamic equilibria. We define $\Omega \subseteq \mathbb{R}^V$ to be the subset of feasible labelings. It is easy to see that any equilibrium trajectory starting from some $\ell^\circ \in \Omega$ remains in Ω for all time.

3 Technical overview

Equilibrium trajectories. We first further develop the view of equilibria as trajectories whose derivatives are controlled by a vector field.

Let $X : \Omega \to \mathbb{R}^V$ be the vector field for which $X(\ell^\circ)$ is the solution to the thin flow equation for $(E'_{\ell^\circ}, E^*_{\ell^\circ})$, for all $\ell^\circ \in \Omega$. So if ℓ is an equilibrium trajectory, $\ell'(\theta) = X(\ell(\theta))$ for almost every θ. Now since $X(\ell^\circ)$ depends only on E'_{ℓ° and $E^*_{\ell^\circ}$, we can deduce that X is piecewise constant. The regions on which X is constant also have a very simple structure, given by (1): each arc $e = vw$ divides Ω into two open halfspaces separated by the hyperplane $\{\ell^\circ \in \Omega : \ell^\circ_w - \ell^\circ_v = \tau_e\}$; a region is determined by the choice of sign (positive, negative or zero) for each link (not all combinations necessarily yield a region); see Figure 2.

Uniqueness and continuity of equilibrium trajectories. We begin our discussion by considering uniqueness of equilibrium trajectories.

To understand why uniqueness is a strong property, and nontrivial to prove, let us consider some other piecewise constant vector fields, restricting ourselves to two dimensions. First, consider the situation shown in Figure 3. Non-uniqueness of trajectories is quite obvious in this case, as we have two regions with a common boundary and the vector fields in each region pointing away from this boundary. Previous work does in fact rule this out for the vector field of interest to us—it can be shown to be a consequence of the uniqueness of solutions to the thin flow conditions.

A more difficult situation is shown in Figure 4, where the trajectory (informally) spirals outwards. Again, non-uniqueness is rather clear; two possible trajectories starting from the origin are shown. This possibility is *not* excluded by previous work. In particular, if a trajectory of this type existed, it would provide an example of an equilibrium (x, ℓ) for which ℓ is not right-differentiable. By contrast, the procedure of [CCL15] and Koch and Skutella [KS11] produces an equilibrium whose labeling is right-differentiable, and so this would immediately be an example of non-uniqueness.

Theorem 3.1. *For any $\ell^\circ \in \Omega$, there is a unique equilibrium trajectory ℓ with $\ell(0) = \ell^\circ$.*

Before discussing how we prove this result, we first discuss continuity, which turns out to be intimately related. For now, we are referring to continuity of the equilibrium trajectory, as a function of the initial feasible labeling. This is very natural from the perspective of equilibrium trajectories, but care is needed in interpreting this; for example, slightly modifying a queue or queues at some moment in the evolution of an equilibrium does not immediately correspond to a change in the labeling. We will discuss other more "interpretable" continuity results at the end of this section.

Let $\mathcal{L}$ denote the Banach space of κ-Lipschitz functions from $[0, \infty)$ to $\mathbb{R}$ imbued with the uniform norm $\|g\|_\infty = \sup_{\theta \geq 0} g(\theta)$. (Recall that $\kappa = u_0 / \min_e \nu_e$ is a Lipschitz constant for all equilibrium trajectories.) By $\mathcal{X}$ we denote the normed space $\mathcal{L}^V$ imbued with the norm $\|\ell\| := \max_{v \in V} \|\ell_v\|_\infty$. Our continuity result is the following.

Theorem 3.2. *Let $\Psi : \Omega \to \mathcal{X}$ be the map that takes $\ell^\circ \in \Omega$ to the unique equilibrium trajectory ℓ satisfying $\ell(0) = \ell^\circ$. Then Ψ is a continuous map.*

We remark that the continuity of Ψ is a statement over the entire evolution of a trajectory—a rather strong form of continuity, and in particular, stronger than convergence of the trajectory at any fixed time θ. Our result says that if we look at the equilibrium trajectory ℓ with $\ell(0) = \ell^\circ$, and then look at the equilibrium trajectory $\tilde{\ell}$ with $\tilde{\ell}(0)$ equal to a slight perturbation of ℓ°, then ℓ and $\tilde{\ell}$ stay close forever, rather than possibly drifting apart very slowly.

While continuity with respect to initial labels is very natural given our view on equilibria, which is very centred on the label functions, care is needed in interpreting this result. We will later discuss a number of other continuity results that are somewhat more "physical" and interpretable in terms of the actual dynamics of flows over time.

The connection to long-term behavior. It is convenient to make this precise by extending the vector field in the neighborhood of ℓ° in a conic fashion. Consider again the structure of the vector field X; its regions are determined by the hyperplanes $\{\ell^\diamondsuit : \ell^\diamondsuit_w - \ell^\diamondsuit_v = \tau_e\}$ for $e = vw \in E$. Call an arc $e = vw$ *tight* (with respect to ℓ°) if $\ell^\circ_w - \ell^\circ_v = \tau_e$; so the hyperplane of arc e passes through ℓ° precisely if it is tight. To obtain a vector field whose regions are all cones around ℓ°, we essentially ignore the hyperplanes defined by non-tight arcs—while still remembering on which side of these hyperplanes ℓ° lies. All arcs that are inactive we will completely ignore; so consider the subnetwork $\hat{G} = (V, E'_{\ell^\circ})$. But all arcs in $E^*_{\ell^\circ}$ are special: not only are they resetting at ℓ°, they will *always* be resetting. This can be interpreted as allowing the queues on these arcs to become negative. We will denote these arcs by E^∞, and call them *free arcs*. We can think of this as defining an entirely new "local instance" defined by the network $\hat{G}$ and free arcs E^∞ (and initial point $\ell(0) = \ell^\circ$). The notion of a valid configuration and hence feasible labeling should be adjusted in the obvious way for this local instance: $(\hat{E}', \hat{E}^*)$ is a valid configuration if $E^\infty \subseteq \hat{E}^* \subseteq \hat{E}' \subseteq E'_{\ell^\circ}$. The feasible set thus changes; denote it by $\hat{\Omega}$.

Since we only consider hyperplanes passing through ℓ°, we get an explicit conic self-similarity: the vector field X satisfies $X(\ell^\circ + \alpha d) = X(\ell^\circ + d)$ for any $d \in \mathbb{R}^V$ with $\ell^\circ + d \in \hat{\Omega}$ and $\alpha > 0$.

We are now ready to make the connection between uniqueness/continuity and long-term behavior of equilibrium trajectories in this conic setting. Let (y, λ) be the solution of the thin flow problem at ℓ°; so $\lambda = X(\ell^\circ)$. Suppose that we could show that for some given initial condition $\hat{\ell}^\circ$, any equilibrium trajectory ℓ

starting from $\hat{\ell}^\circ$ eventually, after some finite amount of time T, satisfies $\ell'(\theta) = \lambda$ for all $\theta \geq T$. Here, T may certainly depend on the choice of $\hat{\ell}^\circ$. But now we can exploit the conic symmetry. Consider an equilibrium trajectory $\ell^{(\varepsilon)}$ starting from $\ell^\circ + \varepsilon(\hat{\ell}^\circ - \ell^\circ)$. Then it is quite easy to see that the trajectory ℓ defined by $\ell(\theta) := (\ell^{(\varepsilon)}(\varepsilon\theta) - \ell^\circ)/\varepsilon + \ell^\circ$ is an equilibrium trajectory starting from $\hat{\ell}^\circ$. Thus $\ell^{(\varepsilon)}(\theta)$ will reach steady state at time εT. So as we move the initial condition closer and closer to ℓ°, any equilibrium trajectory looks more and more like the trajectory ℓ^* given by $\ell^*(\theta) = \ell^\circ + \theta\lambda$. The maximum distance between $\ell^{(\varepsilon)}(\varepsilon)$ and $\ell^*(\varepsilon)$ can be controlled by the initial distance between the initial condition and ℓ°.

This shows uniqueness at ℓ°. Let ℓ be any equilibrium trajectory starting from ℓ°. For any $\varepsilon > 0$, we can think of ℓ on $[\varepsilon, \infty)$ as an equilibrium trajectory starting from $\ell(\varepsilon)$. Since $\ell(\varepsilon) \to \ell^\circ$ as $\varepsilon \to 0$, it follows from the above that $\sup_{\theta \geq 0}[\ell(\theta + \varepsilon) - \ell^*(\theta)]$ converges to 0 as $\varepsilon \to 0$. So $\ell = \ell^*$.

This also shows continuity locally around ℓ°: equilibrium trajectories starting from small perturbations of ℓ° remain close to the equilibrium trajectory starting from ℓ°. Essentially, this shows continuity of the trajectories over a single phase, and it is not too difficult to deduce continuity over the entire equilibrium from this.

Equilibria reach steady state. To prove Theorem 3.1 and Theorem 3.2, the main remaining ingredient is to show that equilibria do always reach steady state. We prove the following theorem.

Theorem 3.3. *Let $G = (V, E)$ be a given network, with $E^\infty \subseteq E$ also given such that (E, E^∞) is a valid configuration. Let ℓ° be any feasible labeling for G with $E^*_{\ell^\circ} \supseteq E^\infty$, and let λ be the solution to the thin flow equations for configuration (E, E^∞). Then for any equilibrium trajectory ℓ with $\ell(0) = \ell^\circ$, there exists some T such that $\ell'_v(\theta) = \lambda_v$ for all $v \in V$ and $\theta \geq T$.*

One can give explicit bounds on T in terms of the instance and ℓ°; we do so in Section 4.

For uniqueness and continuity, the theorem is applied with $E^\infty = E^*_{\ell^\circ}$. The choice $E^\infty = \emptyset$ is of independent interest: it corresponds to a real equilibrium trajectory with no artificial free arcs. As mentioned earlier, [CCO21] proved the following:

Theorem 3.4 ([CCO21])**.** *Consider an instance satisfying the following* inflow condition*: the inflow rate u_0 does not exceed the minimum capacity of an s-t-cut in the network $G = (V, E)$. Then for any feasible initial condition ℓ°, after some finite time all queues remain constant; that is, there exists T such that $\ell'_v(\theta) = 1$ for all $v \in V$ and $\theta \geq T$.*

(Their theorem was stated for ℓ° corresponding to the empty network, but their approach extends directly to arbitrary initial conditions.) Our result for $E^\infty = \emptyset$ thus removes the inflow condition, while modifying the notion of "steady state" appropriately.

Proving Theorem 3.3 constitutes the bulk of the technical work in this paper. Our approach has its genesis in the proof of Theorem 3.4 by [CCO21], but generalizing their result is by no means straightforward. It is not the introduction of free arcs that cause difficulty, but rather violation of the inflow condition. We will now give a high-level sketch of our approach, highlighting the main new difficulties and novelties compared to [CCO21]. The detailed proof can be found in Section 4.

Let us first summarize the approach taken by [CCO21] under the inflow condition and $E^\infty = \emptyset$. They first pose and answer the following question (paraphrased): *which choices of initial condition ℓ° have the property that $\ell(\theta) = \ell^\circ + \lambda\theta$ is immediately an equilibrium trajectory?* (In their case, the thin flow direction λ is the all-ones vector.) They show that the answer is provided in full by considering the following primal-dual

LP:

$$\begin{aligned} \text{minimize} \quad & \sum_{e \in E} \tau_e f_e \\ \text{s.t.} \quad & f \text{ is an } s\text{-}t\text{-flow of value } u_0 \\ & f_e \leq \nu_e \quad \text{for all } e \in E \\ & f_e \geq 0 \quad \text{for all } e \in E \end{aligned} \qquad \begin{aligned} \text{maximize} \quad & u_0(d_t - d_s) - \sum_{e \in E} \nu_e p_e \\ \text{s.t.} \quad & d_w - d_v - p_e \leq \tau_e \quad \text{for all } e = vw \in E \\ & p_e \geq 0 \quad \text{for all } e \in E \end{aligned} \tag{2}$$

Dual optima (d^*, p^*) are in one-to-one correspondence with steady-state initial conditions: d^* represent a possible initial labeling (with no requirement that $d^*_s = 0$ — d^* can be shifted arbitrarily), and p^*_e the queue length on arc e. Primal optima, on the other hand, are in one-to-one correspondence with equilibria departure flows. In other words: (x, ℓ) in which $\ell_v(\theta) = d^*_v + \theta$ for some dual optimum d^*, and $x'(\theta)$ is a primal optimum for almost every θ, is an equilibrium; and all steady-state equilibria are of this form.

It is by no means immediately apparent why answering this question regarding the characterization of steady states is helpful in proving convergence to steady state. The key novelty in [CCO21] is that the dual LP provides us with the correct potential function. Namely, they define, given an equilibrium (x, ℓ) with corresponding queueing delays q, the potential function

$$\Phi(\theta) := u_0[\ell_t(\theta) - \ell_s(\theta)] - \sum_{e=vw \in E} \nu_e q_e(\theta). \tag{3}$$

Then $\Phi(\theta)$ is the objective value of the feasible dual solution given by $d_v = \ell_v(\theta)$ for all $v \in V$, and $p_e = q_e(\theta)$ for all $e \in E$ (feasibility being a consequence of feasibility of the labeling $\ell(\theta)$). The inflow condition ensures that the primal LP is feasible, and hence that the dual optimum has finite value; thus $\Phi(\theta)$ is bounded. Moreover, $\Phi(\theta)$ turns out to be monotone—in fact, strictly monotone with slope bounded away from zero, until the point that steady state is reached. The proof involves rewriting the derivative of Φ, namely (for almost every θ)

$$\Phi'(\theta) = u_0[\ell'_t(\theta) - \ell'_s(\theta)] - \sum_{e \in E^*_\theta} \nu_e q'_e(\theta),$$

as an integral over a family of cuts, after which the inequality $\Phi'(\theta) \geq 0$ follows from the thin flow equations. This shows convergence to steady state in finite time.

To generalize this result, we follow the same basic plan, but each stage presents new (and in some cases significant) additional challenges. We'll only consider $E^\infty = \emptyset$ in this discussion, since as mentioned, this is not the major difficulty. Characterizing steady state solutions is not much more difficult; replacing ν_e with $\hat{\nu}_e := \nu_e \lambda_w$ for all $e = vw \in E$ in (2) (where λ is as defined in Theorem 3.3) does the job, in fact. One can then attempt to define a potential based on the dual objective value in the same way, which would yield

$$\Phi(\theta) = u_0[\ell_t(\theta) - \ell_s(\theta)] - \sum_{e=vw \in E} \lambda_w \nu_e q_e(\theta).$$

This candidate potential is bounded, by feasibility of the primal (which is not difficult to show). But unfortunately, it is *not* monotone![1]

It turns out that while a primal-dual LP characterizing the steady state is still the key to produce the correct potential, the situation is much more subtle. The obvious generalization of (2), with ν_e replaced by $\hat{\nu}_e$ and no other changes, is not the correct one. Rather, one must observe that there is a larger class of candidate LPs, from which a choice must be carefully made. Let $y \in \mathbb{R}^V_{\geq 0}$ be such that (y, λ) is a thin flow for configuration $(E, \emptyset)$. A first observation is that for arcs $e = vw$ with $\lambda_w > \lambda_v$, we may enforce the constraint

[1] An explicit counterexample can be found in [Fra21].

$f_e = \hat{\nu}_e$, and for arcs $e = vw$ with $\lambda_w < \lambda_v$, we may enforce the constraint $f_e = 0$, without changing the feasible set. This comes from observing that if one looks at any set S of the form $S = \{v : \lambda_v \leq t\}$, for some $t \in \mathbb{R}_{\geq 0}$, and this set is nontrivial (neither empty nor V), then $y_e = \lambda_w \nu_e$ for all $e = vw \in \delta^+(S)$ and $y_e = 0$ for all $e \in \delta^-(S)$ by the thin flow conditions.

A followup observation, and the more crucial one, is that we have substantial flexibility in the primal objective function. First, the coefficient in the objective for arcs $e = vw$ with $\lambda_w \neq \lambda_v$ has no impact on the set of optimal solutions, since the flows on these arcs are fixed. Further, we can arbitrarily rescale (by positive values) the coefficients within each level set of λ; this again has no impact, because the flows entering and leaving the level set are determined. In particular, we may replace the objective with $\sum_{e \in E} \hat{\tau}_e f_e$, where $\hat{\tau}_e \coloneqq \tau_e/\lambda_w$, which turns out to be the correct choice.

Even with the (in hindsight) correct choice of primal-dual LP, obtaining the correct potential is not straightforward. The dual no longer points out a precise choice, and to obtain monotonicity, a very careful consideration of different types of arcs is needed. The sign of $\lambda_w - \lambda_v$ for an arc $e = vw$ plays a crucial role. Arcs with $\lambda_w \leq \lambda_v$ contribute in the same way to the potential as in (3); a contribution of $-\nu_e q_e(\theta)$. Arcs with $\lambda_w > \lambda_v$, however, behave differently. For example, they contribute both when the arcs is active and when it is inactive (though in different ways).

Once the precisely correct potential is in place, a cut-based argument along the lines of the [CCO21] proof, with some additional ingredients, shows monotonicity of the potential. Boundedness is straightforward, by comparison with the objective value of the primal-dual program, which is guaranteed to be feasible. One new technical complication is that even once the potential reaches its final value (given by the optimal objective value of the primal-dual program), the label derivatives $\ell'(\theta)$ need not be constant. Rather, the rate of change of *queue lengths* remains constant; some "unimportant" phases involving nodes which are no longer used can still occur.

Continuity revisited. We have so far discussed a notion of continuity that is very natural from our perspective of equilibria as trajectories ℓ. But from the perspective of Nash flows over time, this notion perhaps does not lead itself to immediate interpretation. There are in fact rather a large number of different questions about continuity that one could ask.

We will not attempt to exhaustively consider all possible or interesting notions of continuity in this work. Instead, we will discuss a few results that demonstrate that our "primordial" continuity result, with minimal effort, implies other forms of continuity.

Consider some instance, and the corresponding equilibrium trajectory ℓ starting (for simplicity) from the empty network at time 0. Now suppose we perturb the transit times slightly, leading to a new equilibrium trajectory $\hat{\ell}$ on the perturbed instance. We show that $\hat{\ell}$ remains similar to ℓ.

Theorem 3.5. *Let $G = (V, E)$ be a given network, with fixed capacities ν_e and inflow rate u_0, but variable transit times τ_e. Let $\ell^{(\tau)} \in \mathcal{X}$ be the trajectory corresponding to transit time vector τ, starting from an empty network. Then $\tau \mapsto \ell^{(\tau)}$ is continuous.*

To see how this follows from our main continuity results, consider a single phase (continuity for the entire trajectory then follows by pasting this together appropriately, as with our main continuity result). The steady state direction λ does not depend on τ directly (only through the configuration that defines the phase). The optimal solution and objective value to the primal-dual LP does change, but the optimal objective value, and hence the value that the potential Φ takes upon reaching steady state, is a continuous function of τ. The theorem then follows easily from the monotonicity of Φ.

Similarly, we can consider perturbing the capacities ν_e and/or the inflow rate u_0. Again we get continuity, though necessarily of a slightly weaker form: we must restrict ourselves to a compact interval. This is because the steady state direction λ for a phase *does* change as ν or u_0 changes—but only continuously. This means that once the final steady state of the overall equilibrium is reached, the trajectories may slowly diverge.

Theorem 3.6. *Let $G = (V, E)$, be a given network, with fixed transit times τ_e but variable inflow rate u_0 and capacities ν_e. Let $\ell^{(\nu,u_0)} \in \mathcal{X}$ be the trajectory corresponding to setting capacities to ν and inflow rate to u_0, starting from an empty network. Then $(\nu, u_0) \mapsto \ell^{(\nu,u_0)}(\theta)$ is continuous for any fixed $\theta \in \mathbb{R}_{\geq 0}$.*

As a last example, we consider perturbations of network properties (such as transit time) that occur during the evolution of an equilibrium. There are subtleties here: one must decide at what point knowledge of a perturbation becomes available to a user. We will stick to the most basic choice here: users are aware, from the moment they depart, what perturbations will occur that will impact their arrival time.

Consider an instance, and suppose that the capacity of some arc $e = vw$ changes slightly at time ξ. We show that the resulting equilibrium is close to the equilibrium in which the arc is not perturbed. Of course, continuity with respect to multiple perturbations involving one or more arcs follows immediately. Other types of perturbations (e.g., of transit times) can also be dealt with, though we don't discuss this here. Perturbations of this form are related to dynamic equilibria in time-varying networks, where many fundamental results still hold (see [PS20, Ser20]).

4 Convergence to steady state

In this section, we will prove Theorem 3.3.

We are given a network $G = (V, E)$, as well as $E^\infty \subseteq E$ so that (E, E^∞) is a valid configuration. *Just for this section, it will be very convenient to make the following modification.* We add a free arc ts to the network with capacity $\nu_{ts} := u_0$ and (just for specificity) transit time 0. In the thin flow conditions, (TF-1) should then be replaced by the requirement that x' is a circulation; condition (TF-4) for arc $ts \in E^\infty$ then says that $x'_{ts} = u_0 \ell'_s = u_0$, as needed. The reader may wish to imagine that there is some very large queue on the arc ts, from which the flow emanating from the source s is originating.

Let $\lambda \in \mathbb{R}^V$ be the thin flow direction for configuration (E, E^∞); λ is unique, due to [Koc12, Theorem 6.36] in the case of $E^\infty = \{ ts \}$, or to [CCL15, Theorem 4] in the case that $E \setminus \{ts\}$ is acyclic. Also let

$$\sigma_e := \begin{cases} \lambda_w - \lambda_v & \text{for } e = vw \in E^\infty, \\ [\lambda_w - \lambda_v]^+ & \text{for } e = vw \in E \setminus E^\infty. \end{cases}$$

That is, $\sigma_e = q'_e(\theta)$, where q_e is the queueing delay for e corresponding to the trajectory $\ell(\theta) := \ell^\circ + \lambda\theta$ (for some suitable ℓ°).

Definition 4.1. Given an equilibrium (x, ℓ) with corresponding queueing delays q, and θ that is a point of differentiability of ℓ, we say that the equilibrium is *moving in a steady-state direction at time θ* if $q'(\theta) = \sigma$.

We will sometimes say (somewhat informally) that $\ell'(\theta)$ is a steady-state direction; note however that the definition depends on q' and not just ℓ'—the active and resetting arc sets are important.

Let (x, ℓ) be any equilibrium, and let q be the corresponding queueing delays. Our main task will be showing that there is some time T (depending only on the instance) such that for almost every $\theta \geq T$, the equilibrium is moving in a steady-state direction at time θ. It will then follow immediately that we have reached a steady state by time T, in the sense that all queues grow linearly from time T onwards. This is not quite enough for the conclusion of Theorem 3.3, but it is not hard to show that there exists some $T' \geq T$ such that $\ell'(\theta) = \lambda$ for all $\theta \geq T'$; essentially, after time T, label derivatives only change at "unimportant" nodes that are never again on a shortest path from s to t (and so never see any more flow).

We will later need the following lemma, which says that arcs with a positive queue have a nonzero inflow rate.

Lemma 4.2. *Given an equilibrium (x, ℓ), then $\ell'_v(\theta) > 0$ for all $v \in V$ whenever the derivative exists. In particular, for every resetting arc $e \in E^*_\theta$ we have $x'_e(\theta) > 0$.*

Proof. Fix θ and let $\ell' = \ell'(\theta)$ and $x' = x'(\theta)$. Defining $V_0 := \{ v \in V \mid \ell'_v = 0 \}$ it holds by (TF-4) that $\sum_{e\in\delta^-(v)} x'_e = 0$ for all $v \in V_0$. This immediately imply $t \notin V_0$ as $\sum_{e\in\delta^-(t)} x'_e = u_0 > 0$.

Assume for contradiction that $V_0 \neq \emptyset$ and let v_0 be the first node in V_0 according to a topological ordering with respect to E'_θ. Clearly $v_0 \neq s$ as $\ell'_s = 1$. Due to (TF-3) there has to be an incoming resetting arc $e_0 = uv_0$ (with $x'_{e_0} = 0$). As $\ell(\theta)$ is a feasible labeling there has to be an active v_0-t-path. In particular there has to be an active arc $e = vw$ with $v \in V_0$ and $w \notin V_0$. But since $x'_e = 0$ (flow conservation at v) this is a contradiction to (TF-3). □

Characterizing conditions yielding steady-state directions. The first step will be to give a characterization of possible labels which yield, as a solution to the resulting thin flow conditions, a steady-state direction.

It will be useful to categorize non-free arcs according to their status with respect to λ. Let

$$E^{>} := \{ e = vw \in E \setminus E^{\infty} \mid \lambda_w > \lambda_v \},$$
$$E^{<} := \{ e = vw \in E \setminus E^{\infty} \mid \lambda_w < \lambda_v \},$$
$$E^{=} := \{ e = vw \in E \setminus E^{\infty} \mid \lambda_w = \lambda_v \}.$$

Define $\hat{\nu}_e := \nu_e \cdot \lambda_w$ and $\hat{\tau}_e := \tau_e/\lambda_w$ for all $e = vw \in E$. Before we proceed, we provide useful alternative characterizations of steady-state directions.

Lemma 4.3. *Let (x, ℓ) be an equilibrium, θ a point of differentiability of ℓ and $S = \{v \in V : \ell'_v(\theta) = \lambda_v\}$. Then the following statements are equivalent:*

1. *(x, ℓ) is moving in a steady-state direction at time θ.*
2. *$x'_e(\theta) = 0$ for all $e \in \delta^+(S)$.*
3. *$x'_e(\theta) = 0$ for all $e \notin E[S]$, $E^{>} \subseteq E^*_\theta \subseteq E[S]$, and $E^*_\theta \cap E^{<} = \emptyset$.*

Proof. We omit the dependence on θ in this proof where it is clear.

(1. ⇒ 2.) Let $S_1 = \{v \in V : \ell'_v < \lambda_v\}$. Any arc $e = vw$ entering S_1 satisfies $\ell'_w - \ell'_v < \lambda_w - \lambda_v$. So the only possibility for an arc $e = vw \in \delta^-(S_1)$ that is active is that $\lambda_w - \lambda_v \leq 0$ and $e \notin E^*_\theta$, so that $q'_e = \sigma_e = 0$. But in this case, e is non-resetting and has $\ell'_w < \ell'_v$, and so the thin flow conditions imply that $x'_e = 0$. So there can be no arcs entering S_1 with positive flow; since x' is a circulation, $x'(\delta^+(S_1)) = x'(\delta^-(S_1)) = 0$.

An identical argument for $S_1^* = \{v \in V : \ell'_v \leq \lambda_v\}$ shows that $x'(\delta^-(S_1^*)) = 0$ (notice that still $\ell'_w - \ell'_v < \lambda_w - \lambda_v$ for $vw \in \delta^+(S_1^*)$). Thus since $S = S_1^* \setminus S_1$, $x'(\delta^+(S)) = 0$.

(3. ⇒ 1.) Since $q_e(\theta) = \ell_w(\theta) - \ell_v(\theta) - \tau_e$ for active arcs $e = vw$ (and $q_e(\theta) = 0$ otherwise) we obtain that $q'_e = \ell'_w - \ell'_v$ for resetting arcs, $q'_e = [\ell'_w - \ell'_v]^+$ for active but non-resetting arcs and $q'_e = 0$ for inactive arcs.

Since $E^*_\theta \subseteq E[S]$ and $x'_e = 0$ for all $e \notin E[S]$, $q'_e = 0$ for all $e \notin S$. And since $E^{\infty} \cup E^{>} \subseteq E[S]$, $\sigma_e = 0$ for all $e \notin S$. So we may focus on arcs in $E[S]$. For arcs $e \in E^{<}$, $q'_e = 0 = \sigma_e$ since e is not resetting and $\lambda_w < \lambda_v$. For arcs $e \in E^{>}$, e is active and $\lambda_w > \lambda_v$, and so $q'_e = \ell'_w - \ell'_v = \lambda_w - \lambda_v = \sigma_e$. For arcs $e \in E^{=} \cap E[S]$, since $\ell'_w = \lambda_w = \lambda_v = \ell'_v$ we have that $q'_e = \sigma_e = 0$, irrespective of whether the arc is active, resetting or inactive. So ℓ' is a steady-state direction.

(2. ⇒ 3.) Since there are no directed cycles of arcs with $\tau_e = 0$, there can be no directed cycles of active arcs when ts is excluded. Since x' is supported on active arcs and is a circulation, and $s \in S$, we deduce that $x'_e = 0$ for all $e \notin E[S]$. By Lemma 4.2, $x'_e > 0$ for all $e \in E^*_\theta$, and so $E^*_\theta \subseteq E[S]$.

Consider any $z > 0$, and let $Q = \{v \in S : \lambda_v \leq z\}$. Let y be any circulation so that (y, λ) satisfies the thin flow conditions for (E, E^∞). For any $e = vw \in E^\infty$, $y_e = \lambda_w \nu_e = \ell'_w \nu_e = x'_e$, using that $E^\infty \subseteq E^*_\theta \subseteq E[S]$ as already noted. For an arc $e = vw \in \delta^-(Q) \setminus E^\infty$, we have $e \in E^<$, and so $y_e = 0$. So we can conclude that $y(\delta^-(Q)) \leq x'(\delta^-(Q))$. Next, consider an arc $e = vw \in \delta^+(Q)$. If $e \in E^\infty$, then $x'_e = y_e$, and otherwise, $e \in E^>$. Then $y_e = \lambda_w \nu_e$, but since in this case $w \in S$, we also have by the thin flow conditions tell us that $x'_e \leq \ell'_w \nu_e = \lambda_w \nu_e$. So $x'(\delta^+(Q)) \leq y(\delta^+(Q))$.

Since x' and y are both circulations, we can deduce that $x'_e = y_e$ for $e \in \delta^-(Q) \cup \delta^+(Q)$. Thus arcs $e \in \delta^-(Q) \setminus E^\infty$ do not have a queue (since $x'_e = 0$); and arcs in $\delta^+(Q)$ are active and growing a queue, which means (since θ is a point of differentiability) that such arcs do have a queue. Since this holds for all choices of z, we deduce that all arcs in $E^>$ are resetting, and all arcs in $E^<$ are non-resetting. □

A crucial role will be played by the following minimum cost flow LP:

$$\begin{aligned} \text{minimize} \quad & \sum_{e \in E} \hat{\tau}_e f_e \\ \text{s.t.} \quad & f \ \text{ is a circulation on } G \\ & f_e \leq \hat{\nu}_e \qquad \text{for all } e \in E^= \cup E^< \\ & f_e = \hat{\nu}_e \qquad \text{for all } e \in E^\infty \cup E^> \\ & f_e = 0 \qquad \text{for all } e \in E^< \\ & f \geq 0 \end{aligned} \tag{P}$$

and its dual (which has been slightly massaged),

$$\begin{aligned} \text{maximize} \quad & -\sum_{e \in E} \hat{\nu}_e p_e \\ \text{s.t.} \quad & d_w - d_v - p_e \leq \hat{\tau}_e \qquad \text{for all } e = vw \in E \setminus E^< \\ & p_e \geq 0 \qquad \text{for all } e \in E^= \cup E^< \end{aligned} \tag{D}$$

Note that if (y, λ) is any solution to the thin flow equations for configuration (E, E^∞), then y is a feasible solution to the primal. For the dual, given any assignment of d there is an obvious optimal choice of p that yields a feasible dual solution, namely $p_e = d_w - d_v - \hat{\tau}_e$ for $e = vw \in E \setminus (E^= \cup E^<)$ and $p_e = [d_w - d_v - \hat{\tau}_e]^+$ for $e = vw \in E^= \cup E^<$. As such, we may refer to a given $d \in \mathbb{R}^V$ alone as a solution to the dual, if desired.

Theorem 4.4. *Suppose (x, ℓ) is an equilibrium, and let θ be a point of differentiability of x and ℓ. Let $\tilde{d}_v := \ell_v(\theta)/\lambda_v$ for all $v \in V$. Then $(x'(\theta), \tilde{d})$ is a primal-dual optimal pair if and only if ℓ is moving in a steady-state direction at time θ.*

Proof. First, note that $\tilde{d}$ is well-defined, since Lemma 4.2 implies that $\lambda_v > 0$ for all $v \in V$.

Let q be the queueing delays associated with ℓ. We will omit the dependence on θ (writing x' instead of $x'(\theta)$ for instance) when this is unambiguous. Let $S = \{v \in V : \ell'_v = \lambda_v\}$.

Consider the complementary slackness conditions (with p chosen to be the pointwise minimal solution). They are:

$$\begin{aligned} f_e > 0 \quad &\Rightarrow \quad d_w - d_v \geq \hat{\tau}_e \qquad \text{for all } e \in E \setminus E^<, \\ d_w - d_v > \hat{\tau}_e \quad &\Rightarrow \quad f_e = \hat{\nu}_e \qquad \text{for all } e \in E^=. \end{aligned}$$

Feasible and optimal implies steady-state direction. We begin by showing that if x' is a primal feasible solution and $(x', \tilde{d})$ satisfy the complementary slackness conditions, then ℓ' is a steady-state direction.

We argue that $x'_e = 0$ for any $e = vw \in \delta^+(S)$, and hence that ℓ' is a steady-state direction by Lemma 4.3. Consider any active arc $e = vw \in \delta^+(S)$. Since we are at a point of differentiability, either $e \in E^*_\theta$ or $\ell'_w = \ell'_v$ (or both).

- Suppose $e \in E^*_\theta$. Then $x'_e = \ell'_w \nu_e \notin \{0, \hat{\nu}_e\}$, implying that $e \in E^=$ by primal feasibility. But then we have that $\tilde{d}_w - \tilde{d}_v = \frac{1}{\lambda_w}(\ell_w(\theta) - \ell_v(\theta)) = \frac{1}{\lambda_w}(q_e(\theta) + \tau_e) > \hat{\tau}_e$, which means that we cannot satisfy the second complementary slackness condition. So no such arcs can exist.

- Suppose $\ell'_w = \ell'_v$. If $\lambda_w < \lambda_v$, then $x'_e = 0$ by primal feasibility. If $\lambda_w > \lambda_v$, then primal feasibility implies $x'_e = \lambda_w \nu_e$, so we must have $\lambda_w \leq \ell'_w$ by the thin flow conditions. Since $w \notin S$, $\ell'_w \neq \lambda_w$, so in fact we have $\lambda_w < \ell'_w$. But then $\lambda_w < \ell'_w = \ell'_v = \lambda_v$, a contradiction.

Steady-state direction implies feasible and optimal. Conversely, suppose $\ell'(\theta)$ is a steady-state direction. We need to show that $x'(\theta)$ is feasible for the primal, and that the complementary slackness conditions hold.

We make use of the third characterization of Lemma 4.3. This tells us that $x'_e = 0$ for all $e \notin E[S]$, and $E^> \cup E^\infty \subseteq E^*_\theta \subseteq E[S]$. So feasibility and the complementary slackness condition are satisfied for all arcs outside of $E[S]$.

It remains to consider the situation within S. We begin with primal feasibility. Clearly $0 \leq x'_e \leq \hat{\nu}_e$ for $e \in E[S]$, from the thin flow conditions and $\ell'_v = \lambda_v$ for $v \in S$. For an arc $e = vw \in E^>$, since $E^> \cup E^\infty \subseteq E^*_\theta$ by Lemma 4.3 we must have that $x'_e = \hat{\nu}_e$. And for an arc $e = vw \in E^< \cap E[S]$, $E^< \cap E^*_\theta = \emptyset$ and $\ell'_w < \ell'_v$, implying that $x'_e = 0$.

Now we come to complementary slackness. Arcs $e = vw$ with $x'_e > 0$ satisfy $\ell'_w \geq \ell'_v$ (hence $\lambda_w \geq \lambda_v$) as well as $\ell_w(\theta) \geq \ell_v(\theta) + \tau_e$. Thus

$$\ell_w(\theta)/\lambda_w - \ell_v(\theta)/\lambda_v \geq (\ell_v(\theta) + \tau_e)/\lambda_w - \ell_v(\theta)/\lambda_v \geq \hat{\tau}_e,$$

satisfying the first complementary slackness condition. And if $e = vw \in E^=$ with $\tilde{d}_w - \tilde{d}_v > \hat{\tau}_e$, or equivalently $\ell_w(\theta) - \ell_v(\theta) > \tau_e$, the thin flow conditions imply that $x'_e = \nu_e \ell'_w(\theta) = \hat{\nu}_e$. □

A potential function. For $e \in E \setminus E^\infty$, define $s_e(\theta) := [\tau_e + \ell_v(\theta) - \ell_w(\theta)]^+$, and define $s_e(\theta) := 0$ if $e \in E^\infty$. Then $s_e(\theta) = 0$ if $e \in E'_\theta$, and is otherwise the amount of "slack" by which arc e is inactive.

Let

$$\Phi(\theta) := -\sum_{e=vw \in E^\infty \cup E^>} \hat{\nu}_e \left(\frac{\ell_w(\theta)}{\lambda_w} - \frac{\ell_v(\theta)}{\lambda_v} - \hat{\tau}_e \right) - \sum_{e=vw \in E^*_\theta \setminus (E^\infty \cup E^>)} \nu_e q_e(\theta) - \sum_{e=vw \in E^> \setminus E'_\theta} \nu_e s_e(\theta). \tag{4}$$

Note that if $E^\infty = \{ts\}$ and $\lambda \equiv 1$, then $\Phi(\theta) = -\sum_{e \in E \setminus \{ts\}} \nu_e q_e(\theta) + u_0[\ell_t(\theta) - \ell_s(\theta)]$, exactly matching the potential used in [CCO21]. However, the generalization here is quite far from being the most obvious one; and unlike in [CCO21], it is not clear how to obtain it immediately from the dual objective, even when one has in mind that this is what should motivate the potential. It has been very carefully constructed in order that it is both bounded and monotone, as we will demonstrate.

Boundedness is the easier claim. The optimal objective value of (D), which we denote by OPT, is an explicit bound.

Lemma 4.5. $\Phi(\theta) \leq \mathsf{OPT}$.

Proof. Simply consider the following dual assignment:

$$
\begin{aligned}
d_v &= \tfrac{\ell_v(\theta)}{\lambda_v} && \forall v \in V \\
p_e &= \tfrac{\ell_w(\theta)}{\lambda_w} - \tfrac{\ell_v(\theta)}{\lambda_v} - \hat{\tau}_e && \forall e = vw \in E^\infty \cup (E^> \cap E'_\theta) \\
p_e &= \tfrac{\ell_w(\theta)}{\lambda_w} - \tfrac{\ell_v(\theta)}{\lambda_v} - \hat{\tau}_e + \tfrac{s_e(\theta)}{\lambda_w} && \forall e = vw \in E^> \setminus E'_\theta \\
p_e &= \tfrac{q_e(\theta)}{\lambda_w} && \forall e = vw \in E^*_\theta \setminus (E^\infty \cup E^>) \\
p_e &= 0 && \forall e = vw \in (E^= \cup E^<) \setminus E^*_\theta
\end{aligned}
\tag{5}
$$

It is straightforward to check that this is feasible, and that the resulting dual objective is equal to $\Phi(\theta)$. (Note that the coefficient of s_e and q_e are ν_e, not $\hat{\nu}_e$.) For example, if $e = vw \in E^> \setminus E'_\theta$, then $p_e = \ell_v(\theta)\left(\frac{1}{\lambda_w} - \frac{1}{\lambda_v}\right)$, and so

$$d_w - d_v - p_e = \tfrac{1}{\lambda_w}(\ell_w(\theta) - \ell_v(\theta)) \le \hat{\tau}_e,$$

since $\ell_w(\theta) < \ell_v(\theta) + \tau_e$. □

Now we come to monotonicity. Notice that Φ' takes on one of only finitely many values, since it depends only on the configuration (E'_θ, E^*_θ). The following lemma, combined with Lemma 4.5, thus implies that after some finite time T the equilibrium is moving in a steady-state direction from time T onwards.

Lemma 4.6. *At every point of differentiability θ of Φ, $\Phi'(\theta) \ge 0$, and $\Phi'(\theta) > 0$ if $\Phi(\theta) < \mathsf{OPT}$.*

Proof. Again we omit the parameter θ and just write Φ', ℓ'_v, q'_e and s'_e. We first give a lower bound on Φ'.

Claim 4.7.

$$\Phi' \ge - \sum_{e=vw \in E^*_\theta \cup (E^> \cap E'_\theta)} \hat{\nu}_e \left(\frac{\ell'_w}{\lambda_w} - \frac{\ell'_v}{\lambda_v}\right).$$

Proof. Differentiating Φ, we have

$$\Phi' = - \sum_{e=vw \in E^\infty \cup E^>} \hat{\nu}_e \left(\tfrac{\ell'_w}{\lambda_w} - \tfrac{\ell'_v}{\lambda_v}\right) - \sum_{e=vw \in E^*_\theta \setminus (E^\infty \cup E^>)} \nu_e q'_e - \sum_{e=vw \in E^> \setminus E'_\theta} \nu_e s'_e.$$

We now compare terms. For arcs in $E^=$, E^∞, $E^> \cap E'_\theta$ and $E^< \setminus E^*_\theta$, the contribution is identical. For $e = vw \in E^> \setminus E'_\theta$, we have

$$-\hat{\nu}_e \left(\tfrac{\ell'_w}{\lambda_w} - \tfrac{\ell'_v}{\lambda_v}\right) - \nu_e s'_e = \hat{\nu}_e \ell'_v \left(\tfrac{1}{\lambda_v} - \tfrac{1}{\lambda_w}\right) \ge 0.$$

And for arcs $e = vw \in E^< \cap E^*_\theta$, we have

$$-\nu_e q'_e = -\hat{\nu}_e \left(\tfrac{\ell'_w}{\lambda_w} - \tfrac{\ell'_v}{\lambda_w}\right) \ge -\hat{\nu}_e \left(\tfrac{\ell'_w}{\lambda_w} - \tfrac{\ell'_v}{\lambda_v}\right).$$

■

Let $\tilde{E} := E^*_\theta \cup (E^> \cap E'_\theta)$. For any $z \ge 0$, let

$$S_z := \left\{ v \in V \;\middle|\; \frac{\ell'_v}{\lambda_v} < z \right\}.$$

So

$$\Phi' \ge - \sum_{e=vw \in \tilde{E}} \hat{\nu}_e \cdot \left(\frac{\ell'_w}{\lambda_w} - \frac{\ell'_v}{\lambda_v}\right) = - \int_0^\infty \hat{\nu}(\delta^+(S_z) \cap \tilde{E}) - \hat{\nu}(\delta^-(S_z) \cap \tilde{E})\, \mathrm{d}z. \tag{6}$$

The equality holds since an arc $e = (v, w) \in \tilde{E}$ is in $\delta^+(S_z)$ exactly as long as z fulfills $\frac{\ell'_w}{\lambda_w} \ge z > \frac{\ell'_v}{\lambda_v}$.

Claim 4.8. *If $e = vw \in E$ with $x'_e > 0$ and $\frac{\ell'_w}{\lambda_w} < \frac{\ell'_v}{\lambda_v}$, then $e \in \tilde{E}$.*

Proof. Consider such an arc. Since $x'_e > 0$, either $e \in E^*$ or $e \in E'$ with $\ell'_w \geq \ell'_v$. In the former case, $e \in \tilde{E}$. In the latter case, we must have $\lambda_w > \lambda_v$ (since $\ell'_w/\lambda_w < \ell'_v/\lambda_v$), and since e is flow carrying we have $e \in (E^> \cap E')$ implying $e \in \tilde{E}$. ■

Since x' is a circulation and all flow entering S_z is always counted, it follows that the incoming flow is not smaller than the outgoing flow, i.e., $x'(\delta^-(S_z) \cap \tilde{E}) \geq x'(\delta^+(S_z) \cap \tilde{E})$.

Claim 4.9. *For all $e = vw \in \delta^+(S_z) \cap \tilde{E}$, we have $x'_e = \ell'_w \cdot \nu_e$.*

Proof. If $e \in E^*$ then the claim trivially follows. Otherwise, $e \in E^> \cap \delta^+(S_z)$, i.e., $\lambda_w > \lambda_v$ and $\frac{\ell'_w}{\lambda_w} > \frac{\ell'_v}{\lambda_v}$, which implies that $\ell'_w > \ell'_v$. Since $e \in E'_\theta$, the claim follows from the thin flow conditions. ■

For all active arcs, we certainly have $x'_e \leq \ell'_w \nu_e$. This together with the claims yields

$$\begin{aligned}
z \cdot \hat{\nu}(\delta^-(S_z) \cap \tilde{E}) &\geq \sum_{e=vw \in \delta^-(S_z) \cap \tilde{E}} \frac{\ell'_w}{\lambda_w} \cdot \hat{\nu}_e \\
&= \sum_{e=vw \in \delta^-(S_z) \cap \tilde{E}} \ell'_w \cdot \nu_e \\
&\geq x'(\delta^-(S_z) \cap \tilde{E}) \\
&\geq x'(\delta^+(S_z) \cap \tilde{E}) \\
&= \sum_{e=vw \in \delta^+(S_z) \cap \tilde{E}} \ell'_w \cdot \nu_e \\
&= \sum_{e=vw \in \delta^+(S_z) \cap \tilde{E}} \frac{\ell'_w}{\lambda_w} \cdot \hat{\nu}_e \\
&\geq z \cdot \hat{\nu}(\delta^+(S_z) \cap \tilde{E}).
\end{aligned}$$

This together with (6) shows that $\Phi' \geq 0$.

Finally, suppose that $\Phi(\theta) < \mathsf{OPT}$; we must show that $\Phi'(\theta) > 0$. Let $S = \{v \in V : \ell'_v = \lambda_v\}$ as usual. Since the equilibrium is not moving in a steady-state direction at time θ, we know that that the circulation x' is nonzero across S, by Lemma 4.3. This means that for some nontrivial interval I (either $I = (1-\varepsilon, 1)$ or $I = (1, 1+\varepsilon)$, for some $\varepsilon > 0$), x' crosses S_z for all $z \in I$. For any $z \in I$, choose $e = vw \in \delta^-(S_z)$ with $x'_e > 0$. Either $e \in E^>$ or $\ell'_w > \ell'_v$; either way, $e \in \tilde{E}$. This means that the first inequality in the above displayed chain is strict, for all $z \in I$, and hence Φ' is strictly positive. □

The final phases. We have shown that after finite time, $\ell'(\theta)$ is always a steady-state direction. It remains to show that $\ell'(\theta) = \lambda$ after finite time. For some intuition, the essential reason that $\ell'(\theta)$ can be a steady-state direction differing from λ is the following situation. Take v to be a node where $\ell'_v(\theta) \neq \lambda_v$, chosen with $\ell'_v(\theta)$ minimal amongst all choices; so for any incoming arc uv, $\lambda_u = \ell'_u(\theta)$. It turns out that $\lambda_v < \ell'_v(\theta)$ — the earliest arrival time to v is increasing more quickly than it "should" according to λ_v. Further, there will be an arc wv, currently inactive, with $\ell'_w(\theta) = \lambda_w = \lambda_v < \ell'_v(\theta)$. The slack of this arc is decreasing for as long as $\ell'_v(\theta)$ remains above λ_v, so this situation cannot remain for too long; but once the arc becomes active, $\ell'_v(\theta) = \lambda_w = \lambda_v$. We will now make this picture precise.

Let $(\gamma_v(\theta))_{v \in V}$ be the shortest path labels from s in the network $(V, E \setminus E^<)$ with costs given by the slacks $s_e(\theta) = [\tau_e + \ell_v(\theta) - \ell_w(\theta)]^+$ for $e = vw \in E \setminus E^<$. If $E^< = \emptyset$, clearly $\gamma_v(\theta) = 0$ for all v, as there is always an active s-v path, i.e., a path without slack. Let $\Delta := \max\{\,|\lambda_v - \lambda_w|^{-1} : \lambda_v \neq \lambda_w\,\}$ (or $\Delta := 0$ if the maximum is over an empty set).

Lemma 4.10. *Suppose that $\Phi(T_1) = \mathsf{OPT}$. Then for all $\theta \geq T_2$, where $T_2 := T_1 + \Delta \cdot \max_{v \in V} \gamma_v(T_1)$, we have $\ell'_v(\theta) = \lambda_v$ for all nodes $v \in V$.*

Proof. Let $\theta \geq T_1$ be a point of differentiability and let $S := \{\, v \in V : \ell'_v(\theta) = \lambda_v \,\}$. We need three subsidiary claims.

Claim 4.11. *For any $v \in V$, $\ell'_v(\theta) \geq \lambda_v$.*

Proof. Suppose the claim does not hold; let v be a node where $\ell'_v(\theta) < \lambda_v$, but $\ell'_u(\theta) \geq \lambda_u$ for all $uv \in \delta^-(v) \cap E'_\theta$. Such a node must exist, since (V, E'_θ) is acyclic.

Clearly $v \notin S$. Thus by Lemma 4.3, all arcs entering v are non-resetting, and carry no flow in the thin flow at time θ. Thus by the thin flow conditions (TF-3),

$$\ell'_v(\theta) = \min_{uv \in E'_\theta} \ell'_u(\theta) \geq \min_{uv \in E'_\theta} \lambda_u.$$

But arcs in $\delta^-(v)$ are not in E^∞ nor $E^>$ (again by Lemma 4.3), and so $\min_{uv \in E'_\theta} \lambda_u \geq \lambda_v$, a contradiction. ■

Claim 4.12. *We have the alternative characterization $S = \{v \in V : \gamma_v(\theta) = 0\}$.*

Proof. Suppose $\gamma_v(\theta) = 0$. Then there is a (possibly empty) path P' in $(V, E \setminus E^<)$ of arcs with no slack — that is, active arcs — from s to v. Let u be the last node in P' contained in S, and let P be the subpath of P' from u to v. Since all arcs of P lie in $E \setminus (E^< \cup E[S])$, by Lemma 4.3 all arcs of P lie in $E^=$. Thus $\ell'_v(\theta) = \ell'_u(\theta) = \lambda_u = \lambda_v$, with the first equality coming from the arcs in P being non-resetting and active, the second from $u \in S$, and the third from $P \subseteq E^=$.

Conversely, suppose $\ell'_v(\theta) = \lambda_v$. All arcs in $E^< \cap E[S]$ must be inactive; they cannot be resetting due to Lemma 4.3, and they cannot be active but non-resetting since θ is a point of differentiability. We claim that for every vertex $v \in S$ there is an incoming active arc in $E[S]$. Assume this is not the case; then there exists an active arc $e = uv \in E \setminus E[S] \subseteq E^= \cup E^<$. Since $x'_e = 0$ by Lemma 4.3, this together with Claim 4.11 yields

$$\lambda_u \geq \lambda_v = \ell'_v(\theta) = \ell'_u(\theta) > \lambda_u,$$

a contradiction. Hence, there exists an active path in $E[S]$ from s to v. All arcs on this path have zero slack and lie in $E \setminus E^<$, showing that $\gamma_v(\theta) = 0$. ■

Claim 4.13. *For any $v \in V \setminus S$, $\gamma'_v(\theta) \leq -1/\Delta < 0$ (or if $\Delta = 0$, then $S = V$).*

Proof. Consider any v with $\gamma_v(\theta) > 0$; note that $v \notin S$ by the previous claim. Let P_1 be a u-v path in $E^= \setminus E[S]$ with $u \in S$ and minimal slack, one for which this property remains true for some small interval of time. That is, $\sum_{e \in P_1} s_e(\xi) = \gamma_v(\xi)$ for all $\xi \in [\theta, \theta + \varepsilon)$ for some small $\varepsilon > 0$. Such a path must exist, given that θ is a point of differentiability; we can choose ε so that all slacks are linear in this interval. Let P_2 be a w-v path in $E'_\theta \setminus E[S]$ with $w \in S$. We have that $\lambda_v = \lambda_u = \ell'_u$ and $\lambda_v < \lambda_w = \ell'_w = \ell'_v$. The inequality follows since P_2 must use at least one arc from $E^<$, or otherwise P_2 would be a no slack path in $E \setminus E^<$. (This observation further implies that if $\Delta = 0$, meaning $E^< = \emptyset$, then $S = V$ and we are done.)

As there are no queues on the arcs on P_1 we have

$$\gamma'_v(\theta) = \sum_{e \in P_1} s'_e(\theta) = \sum_{v'w' \in P_1} \ell'_{v'}(\theta) - \ell'_{w'}(\theta) = \ell'_u - \ell'_v = \lambda_v - \lambda_w \leq -1/\Delta. \qquad ■$$

So after $\Delta \cdot \max_{v \in V} \gamma_v(T_1)$ time has passed beyond time T_1, the slack $\gamma_v(\theta)$ has decreased to 0 for every $v \in V$, and hence all nodes are in S. Finally, since $\ell'_v(\theta) = \lambda_v$ for almost all $\theta \geq T_2$, it follows that ℓ_v is differentiable for every $\theta \geq T_2$. □

Altogether, we can now prove Theorem 3.3, and in fact the following more explicit (but somewhat ungainly) statement. First, we define $\eta > 0$ to be a lower bound on $\Phi'(\theta)$ whenever $\Phi(\theta) \neq \mathsf{OPT}$, which exists by Lemma 4.6 and the fact that $\Phi'(\theta)$ takes only finitely many values.

Theorem 4.14. *Let $G = (V, E)$ be a given network, with $E^\infty \subseteq E$ such that (E, E^∞) is a valid configuration. Let ℓ° be any feasible labeling for G and E^∞, and let λ be the solution to the thin flow equations for configuration (E, E^∞). Then for any equilibrium trajectory ℓ with $\ell(0) = \ell^\circ$, it holds that $\ell'_v(\theta) = \lambda_v$ for all $v \in V$ and $\theta \geq T := T_1 + T_2$, where $T_1 := (\mathsf{OPT} - \Phi(0))/\eta$ and $T_2 := \Delta \cdot |E| \left(T_1 \cdot \kappa + \max_{e \in E \setminus E^<} s_e(0)\right)$.*

Proof. By Theorem 3.3 together with Lemma 4.6 the equilibrium trajectory $\ell(\theta)$ follows always a steady state direction from time T_1 onwards. By Lemma 4.10 we have $\ell'(\theta) = \lambda$ for all $\theta \geq T_1 + \Delta \cdot \max_{v \in V} \gamma_v(T_1)$. Finally, we have $T_2 \geq \Delta \cdot \max_{v \in V} \gamma_v(T_1)$ because $\ell'_v(\theta) \leq \kappa$ and thus also $s'_e(\theta) \leq \kappa$. Hence the shortest path labels $\gamma_v(T_1)$ are bounded by the total slack at time T_1, and the slack of each arc e is bounded by $s_e(0) + T_1 \cdot \kappa$. □

5 Uniqueness and continuity

In this section, we prove the uniqueness and continuity of equilibrium trajectories. Our main tool will be the following lemma showing continuity for some small interval.

Lemma 5.1 (Local continuity)**.** *Fix $\ell^\circ \in \Omega$, and let ℓ^* be defined by $\ell^*(\theta) := \ell^\circ + \theta \cdot X(\ell^\circ)$. Then there exists $\varepsilon > 0$ such that the following holds. For any sequence $\ell^{(1)}, \ell^{(2)}, \ldots$ of equilibrium trajectories where $\ell^{(i)}(0) \to \ell^\circ$ as $i \to \infty$, $\ell^{(i)}(\theta) \to \ell^*(\theta)$ as $i \to \infty$ for all $\theta \in [0, \varepsilon]$.*

This lemma will follow from our results on long-term behavior. We first require some preliminaries.

The local network. Fix any initial condition $\ell^\circ \in \Omega$. Then there exists a small $\hat{\theta} > 0$ such that for any equilibrium trajectory ℓ starting from ℓ°, $E^*_{\ell^\circ} \subseteq E^*_{\ell(\theta)}$ and $E'_{\ell(\theta)} \subseteq E'_{\ell^\circ}$ for all $\theta \in [0, \hat{\theta}]$. In other words, arcs with positive queues at the beginning keep a queue during this time and no inactive arc becomes active. As we focus on this configuration we consider the network in which only active arcs are present and in which initially resetting arcs are free arcs. So set $\hat{E} := E'_0$ and $E^\infty := E^*_0$, and define the *local network* to be $\hat{G} = (V, \hat{E})$ with free arcs E^∞. We will use the notation $\hat{\Omega}$ to refer to the set of feasible labelings for the local network. For any $\ell^\circ \in \hat{\Omega}$, we use $\hat{E}'_{\ell^\circ}$ and $\hat{E}^*_{\ell^\circ}$ to refer to the active and resetting arcs for the local network.

Note that the vector field $\hat{X}$ of this local network has a very specific structure: All hyperplanes that separate the different thin flow regions contain the central point ℓ°. In other words, all regions are cones with center ℓ°. This provides us with a scaling invariance, captured in the following lemma.

Lemma 5.2. *In the local network $\hat{G}$ the following holds for any $\alpha > 0$.*

1. *If $p \in \mathbb{R}^V$ with $\ell^\circ + p \in \hat{\Omega}$ then $\ell^\circ + \alpha p \in \hat{\Omega}$.*
2. *For any equilibrium trajectory ℓ, the trajectory $\ell^{(2)}$ defined by $\ell^{(2)}(\theta) = \alpha \cdot \ell(\frac{\theta}{\alpha}) + (1 - \alpha)\ell^\circ$ is also an equilibrium trajectory.*

Proof. For the first claim, we show that $\hat{E}'_{\ell^\circ+p} = \hat{E}'_{\ell^\circ+\alpha p}$ and $\hat{E}^*_{\ell^\circ+p} = \hat{E}^*_{\ell^\circ+\alpha p}$. For $e \in E^\infty$ this is clear as e is resetting for both initial values. For $e = vw \in \hat{E} \setminus E^\infty$ we have $\ell^\circ_w = \ell^\circ_v + \tau_e$; hence

$$\ell^\circ_w + p_w \geq \ell^\circ_v + p_v + \tau_e \quad \Leftrightarrow \quad \ell^\circ_w + \alpha p_w \geq \ell^\circ_v + \alpha p_v + \tau_e \quad \text{and}$$
$$\ell^\circ_w + p_w > \ell^\circ_v + p_v + \tau_e \quad \Leftrightarrow \quad \ell^\circ_w + \alpha p_w > \ell^\circ_v + \alpha p_v + \tau_e.$$

To show the second claim, let $p := \ell(0) - \ell^\circ$, and note that $\ell^{(2)}(0) = \ell^\circ + \alpha p$. So $\ell^{(2)}(0) \in \hat{\Omega}$ by the above. Furthermore, for almost every point in time the derivative coincides with the vector field:

$$\frac{d}{d\theta}\ell^{(2)}(\theta) = \frac{d}{d\theta}\ell(\theta/\alpha) \overset{a.e.}{=} \hat{X}(\ell(\theta/\alpha)) = \hat{X}\left(1/\alpha\ell^{(2)}(\theta) + (1-1/\alpha)\ell^\circ\right) = X(\ell^{(2)}(\theta)).$$

Note for the last equation that $\frac{1}{\alpha}\ell^{(2)}(\theta) + (1-\frac{1}{\alpha})\ell^\circ$ lies on a line between $\ell^{(2)}(\theta)$ and ℓ°, and hence, the equality follows due to the conic structure of the vector field $\hat{X}$. So $\ell^{(2)}$ is an equilibrium trajectory. □

We are now ready to prove the local continuity lemma.

Proof of Lemma 5.1. Consider the local network $\hat{G}$ and the closed unit ball $B_1(\ell^\circ)$ around ℓ°. Observe that the bound on the time to reach the final phase that Theorem 4.14 provides is clearly continuous in ℓ°. Thus, by compactness of the ball, there is a time T such that for all $\hat{\ell}^\circ \in B_1(\ell^\circ)$, the equilibrium trajectory $\hat{\ell}$ with $\hat{\ell}(0) = \hat{\ell}^\circ$ satisfies $\hat{\ell}'(\theta) = \lambda$ for all $\theta \geq T$.

For a given $\delta > 0$ we obtain with Lemma 5.2 that every equilibrium trajectory $\tilde{\ell}$ in $\hat{G}$ starting at $\tilde{\ell}(0) \in B_\delta(\ell^\circ) \cap \hat{\Omega}$ has reached steady state at time $T \cdot \delta$ (since $\ell^{(2)}(\theta) := 1/\delta \cdot \tilde{\ell}(\delta\theta) + (1-1/\delta)\ell^\circ$ is an equilibrium trajectory starting from $\ell^{(2)}(0) \in B_1(\ell^\circ) \cap \hat{\Omega}$ and hence has reached steady state at time T). By choosing $\delta := \varepsilon/(T \cdot \kappa)$ (recall κ is a Lipschitz constant for any equilibrium trajectory), we have that $\|\ell^*(\theta) - \tilde{\ell}(\theta)\| \leq \varepsilon$, since the steady state is reached within time $T \cdot \delta = \varepsilon/\kappa$. □

Uniqueness. We are now ready to show the uniqueness result stated in Theorem 3.1.

Proof of Theorem 3.1. Assume for contradiction that there are two distinct equilibrium trajectories $\ell^{(1)}$ and $\ell^{(2)}$ starting from ℓ°. Without loss of generality we assume that the trajectories diverge right from the start: otherwise, simply consider the moment in time when the two trajectories diverge, and treat this as the initial condition (shifting times appropriately).

Let $\theta_1, \theta_2, \dots$ be any sequence of strictly positive times converging to zero. For each i, we can consider $\theta \mapsto \ell^{(1)}(\theta + \theta_i)$ as an equilibrium trajectory starting from $\ell^{(1)}(\theta_i)$, and $\theta \mapsto \ell^{(2)}(\theta + \theta_i)$ as an equilibrium trajectory starting from $\ell^{(2)}(\theta_i)$. By Lemma 5.1, there exists an $\varepsilon > 0$ so that $\theta \mapsto \ell^{(1)}(\theta + \theta_i)$ converges to $\theta \mapsto \ell^\circ + (\theta + \theta_i)X(\ell^\circ)$ on the interval $[0, \varepsilon]$. Taking a limit as $i \to \infty$, we find that $\ell^{(1)}$ must be identical to $\theta \mapsto \ell^\circ + \theta X(\ell^\circ)$ on the interval $[0, \varepsilon]$. But the identical argument applies to $\ell^{(2)}$, contradicting our assumption that the trajectories $\ell^{(1)}$ and $\ell^{(2)}$ diverge at time 0. □

Continuity. As defined earlier, let $\Psi\colon \Omega \to \mathcal{X}$ map any feasible initial condition ℓ° to the unique equilibrium trajectory $\ell \in \mathcal{X}$ with $\ell(0) = \ell^\circ$. We wish to prove Theorem 3.2, that Ψ is continuous.

First, we prove a weaker continuity statement. For every $\theta \in \mathbb{R}_{\geq 0}$, define $\Psi_\theta\colon \Omega \to \mathbb{R}^V_{\geq 0}$ by $\Psi_\theta(\ell^\circ) = [\Psi(\ell^\circ)](\theta)$; that is, we map an initial condition to the value of the resulting equilibrium trajectory at a fixed time θ.

Lemma 5.3. *For any $\theta \in \mathbb{R}_{\geq 0}$, Ψ_θ is continuous.*

Proof. Suppose that is not the case; then the following set is bounded:

$$M := \left\{ \vartheta \in \mathbb{R}_{\geq 0} \,\middle|\, \Psi_{\theta'} \text{ is continuous for all } \theta' \in [0, \vartheta] \right\}.$$

Let $\xi := \sup M$ (note that $M \neq \emptyset$ as $0 \in M$). As a first step we show that $\xi \in M$. Recall that all equilibrium trajectories are κ-Lipschitz. This proves that Ψ_ξ is continuous, because for every $\varepsilon > 0$ we can find a $\delta > 0$ such that for all $\theta' \in [\xi - \frac{\varepsilon}{3\kappa}, \xi)$ and $\hat{\ell}^\circ \in B_\delta(\ell^\circ) \cap \Omega$, it holds that

$$\|\Psi_\xi(\ell^\circ) - \Psi_\xi(\hat{\ell}^\circ)\| \leq \|\ell(\xi) - \ell(\theta')\| + \|\Psi_{\theta'}(\ell^\circ) - \Psi_{\theta'}(\hat{\ell}^\circ)\| + \|\hat{\ell}(\theta') - \hat{\ell}(\xi)\| \leq \tfrac{\varepsilon}{3} + \tfrac{\varepsilon}{3} + \tfrac{\varepsilon}{3} \leq \varepsilon.$$

Here, ℓ and $\hat{\ell}$ are the equilibrium trajectories starting with ℓ° and $\hat{\ell}^\circ$, respectively.

We can consider, for every equilibrium trajectory ℓ, $\ell(\xi)$ as the initial condition of the equilibrium trajectory $\theta' \mapsto \ell(\theta' + \xi)$. Applying Lemma 5.1, we obtain that for a small duration $[\xi, \xi + \varepsilon]$ the equilibrium trajectory ℓ depends continuously on the value of $\ell(\xi)$. But since Ψ_ξ is continuous, meaning that $\ell(\xi)$ depends continuously on ℓ°, this shows that $\Psi_{\xi+\varepsilon}$ is also continuous, contradicting our choice of ξ. □

Given an interval $I \subseteq \mathbb{R}_{\geq 0}$, let $\Psi_I(\ell^\circ)$ be the restriction of $\Psi(\ell^\circ)$ to the time interval I, for any $\ell^\circ \in \Omega$. It is a basic fact that a sequence of Lipschitz continuous functions that converge pointwise on a compact interval converge uniformly. Thus Ψ_I is continuous for any compact interval I. To extend to the non-compact interval $\mathbb{R}_{\geq 0}$, we again make use of our steady state result, this time on the original network. Fix some positive δ (say $\delta = 1$). Considering Theorem 4.14, and noting that $\Phi(0)$ depends continuously on the initial conditions, we can deduce that there is a bound T such that every equilibrium trajectory with starting point in $B_\delta(\ell^\circ) \cap \Omega$ has reached steady state by time T. Thus for any $\ell^\Diamond \in B_\delta(\ell^\circ) \cap \Omega$ and $\theta > T$, $\Psi_\theta(\ell^\Diamond) = \Psi_T(\ell^\Diamond) + (\theta - T) \cdot \lambda$, where λ is the steady-state direction for the network G, i.e., λ is the solution to the thin flow equations for configuration $(E, \emptyset)$. Thus $\sup_{\theta \geq 0} \|\Psi_\theta(\ell^\Diamond) - \Psi_\theta(\ell^\circ)\| = \sup_{0 \leq \theta \leq T} \|\Psi_\theta(\ell^\Diamond) - \Psi_\theta(\ell^\circ)\|$, and continuity of Ψ follows.

6 Continuity more generally

Here we show a number of examples of continuity results that can be fairly easily deduced from the continuity of equilibrium trajectories.

Continuity with respect to τ. We first show Theorem 3.5, that equilibria trajectories are continuous with respect to the transit times of the instance. So fix all aspects of the instance apart from the transit time vector τ. We will use Ω^τ to denote the set of feasible labelings for the instance corresponding to τ, and similarly X^τ for the vector field.

An important first observation is that for a fixed configuration (E', E^*), the thin flow equations have no dependence on τ whatsoever.

First, we prove local continuity, in the sense of Lemma 5.1.

Lemma 6.1. *Fix $\tau^\circ \in \mathbb{R}^E_{\geq 0}$ with no directed cycles of zero cost, as well as $\ell^\circ \in \Omega^{\tau^\circ}$. Let ℓ^* be defined by $\ell^*(\theta) = \ell^\circ + \theta \cdot X^{\tau^\circ}(\ell^\circ)$. Then there exists $\varepsilon > 0$ such that the following holds. Consider any sequence $\tau^{(1)}, \tau^{(2)}, \ldots$ of valid transit time vectors converging to τ°, and any corresponding sequence of equilibrium trajectories $\ell^{(1)}, \ell^{(2)}, \ldots$ with $\ell^{(i)}$ being in the instance with transit time vector $\tau^{(i)}$ and $\ell^{(i)}(0) \to \ell^\circ$ as $i \to \infty$. Then $\ell^{(i)}(\theta) \to \ell^*(\theta)$ as $i \to \infty$ for all $\theta \in [0, \varepsilon]$.*

Proof. As in the proof of Lemma 5.1, we pick ε small enough so that we can restrict our attention to the local network at ℓ°. More precisely, suppose $\varepsilon > 0$ is small enough that for all i sufficiently large, the active arcs of $\ell^{(i)}(\theta)$ are in E'_{ℓ° for all $\theta \in [0, \varepsilon]$, and all arcs of $E^*_{\ell^\circ}$ are resetting arcs of $\ell^{(i)}(\theta)$. As before let (λ, y) be a solution to the thin flow equations for the configuration $(E'_{\ell^\circ}, E^*_{\ell^\circ})$ considering transit times τ°. Note that the steady state direction λ is independent of transit times.

Let $\Phi^{(i)}$ be the potential function as defined in (4) for $\ell^{(i)}$ in the network $G^{\tau^{(i)}}$, and let Φ° be the potential for ℓ^* in G^{τ°. Since the potential at time 0 depends continuously on transit times and the initial values, we have $\Phi^{(i)}(0) \to \Phi^\circ(0)$ as $i \to \infty$. These upper bounds converge to the optimal primal-dual value of the original network. The positive lower bound on Φ' when not in a steady state only depends on the configuration (E', E^*) and thus holds for all these potentials simultaneously. Therefore, by Theorem 4.14, the time $T_1^{(i)}$ after which $\ell^{(i)}$ moves in steady state direction goes to 0 as $i \to \infty$. This on the other hand implies that $T_2^{(i)}$, the duration after $T_1^{(i)}$ until $\ell'^{(i)} = \lambda$, goes to 0 as $i \to \infty$ because the initial slack of all these trajectories converges to the slack of ℓ°, which is 0 for all arcs $e \in E'_{\ell^\circ}$. With $T^{(i)} = T_1^{(i)} + T_2^{(i)} \to 0$ we conclude that $\|\ell^*(\theta) - \ell^{(i)}(\theta)\| \leq 2\kappa T^{(i)} + \|\ell^\circ - \ell^{(i)}(0)\| \to 0$ as $i \to \infty$, for any $\theta \in [0, \varepsilon]$. ☐

Given this, the conversion of local continuity to full continuity over the whole trajectory is essentially identical to the argument in Section 5. The extension of local continuity to continuity over any compact interval is the same. To extend this to the entire trajectory, we just again use that the steady-state direction, obtained as the solution to the thin flow equations with configuration $(E, \emptyset)$ is the same. We can obtain a uniform bound on the time to reach this steady state for all choices of τ in some neighborhood, and so the claim follows.

Continuity with respect to ν. Next, we consider perturbing the capacities ν_e and/or the inflow rate u_0. Perturbing the inflow rate can be thought of as perturbing an arc capacity, by inserting a dummy arc $s's$ with capacity u_0, setting s' to be the new source, and choosing the inflow of the new instance to be large; perturbing the capacity of $s's$ is then functionally equivalent to perturbing the inflow of the original instance. So we'll consider only perturbations of ν.

Define, for any $\nu \in \mathbb{R}^E_{>0}$, G^ν and X^ν for the instance and vector field corresponding to ν. This time, the set of feasible labelings Ω does not depend on ν.

The new ingredient compared to perturbing τ is that the steady-state direction λ^ν *does* depend on ν. However, it does so continuously.

Lemma 6.2. *Fix some valid configuration (E', E^*), and let λ^ν be the thin flow direction for this configuration in G^ν, for any capacity vector ν. Then λ^ν depends continuously on ν in $\mathbb{R}^E_{>0}$.*

Proof. Consider an ordered partition $\mathcal{P} = (V_1, V_2, \ldots, V_k)$ of V, which we can view as an assignment $\pi : V \to \mathbb{N}$ that labels each node with the index of the part it lies in. Associated with this, define the following linear system:

$$\begin{aligned}
\ell'_s &= 1, \\
\ell'_w &= \ell'_v && \text{for all } v, w \text{ where } \pi(w) = \pi(v), \\
x'_e &= \nu_e \ell'_w && \text{for all } e = vw \text{ where } e \in E^* \text{ or } \pi(w) > \pi(v), \\
x'_e &= 0 && \text{for all } e = vw \text{ where } e \notin E^* \text{ and } \pi(w) < \pi(v), \\
\sum_{e \in \delta^+(v)} x'_e - \sum_{e \in \delta^-(v)} x'_e &= \begin{cases} u_0 & \text{if } v = s, \\ 0 & \text{if } v \in V \setminus \{s, t\}. \end{cases}
\end{aligned}$$

This linear system may have multiple solutions. If there exists a solution (x', ℓ') for which $0 \leq x'_e \leq \nu_e \ell'_w$ for all $e = vw$, then it is easy to verify that this satisfies the thin flow conditions; essentially, we have guessed the ordering between label derivatives. Say that $\mathcal{P}$ is a *correct (ordered) partition* if this holds, and let $N[\mathcal{P}]$ denote the set of capacities for which the linear system for this partition is correct. Note that $N[\mathcal{P}]$ is a closed set. Existence of thin flows implies that for every ν, there is at least one ordered partition $\mathcal{P}$ for which $N[\mathcal{P}]$ contains ν (there may be more than one; it is possible that for some partition $(V_1, \dots, V_k)$, we obtain a solution with $\ell'_v = \ell'_w$ for $v \in V_j$, $w \in V_{j+1}$, in which case the partition obtained by merging V_j and V_{j+1} would also be correct). Uniqueness of thin flow directions implies that for any correct partition $\mathcal{P}$, and *any* solution (x', ℓ') to the corresponding system, $\ell' = \lambda^\nu$.

It follows that to show continuity of λ^ν, it suffices to show, for a fixed ordered partition $\mathcal{P}$, continuity of the solution ℓ' of the linear system with respect to ν within $N[\mathcal{P}]$. Substitute out the x' variables, to reduce to a linear system $A^\nu \ell' = b$. The entries of A^ν clearly depend continuously (indeed, linearly) on ν. Further, since ℓ' is uniquely determined in $N[\mathcal{P}]$, A^ν is nonsingular in this set. Hence $\nu \mapsto (A^\nu)^{-1}$ is also continuous in $N[\mathcal{P}]$. □

Since Ω does not depend on ν, we simplify our life and maintain a fixed initial condition (of course, the result could be combined with the results on continuity with respect to initial conditions, if desired). The local continuity statement now becomes much simpler, since we know that the trajectory will move in the thin flow direction for some configuration, and so becomes an immediate consequence of the previous lemma.

Lemma 6.3. *Fix $\nu^\circ \in \mathbb{R}^E_{>0}$, as well as $\ell^\circ \in \Omega$. Let ℓ^* be defined by $\ell^*(\theta) = \ell^\circ + \theta \cdot X^{\nu^\circ}(\ell^\circ)$. Then there exists $\varepsilon > 0$ such that the following holds. Defining ℓ^ν to be the equilibrium trajectory for G^ν with $\ell^\nu(0) = \ell^\circ$, $\ell^\nu(\theta) \to \ell^*(\theta)$ as $\nu \to \nu^\circ$ for all $\theta \in [0, \varepsilon]$.*

Proof. Let λ^ν be the solution to the thin flow equations with configuration $(E'_{\ell^\circ}, E^*_{\ell^\circ})$ in G^ν, for any capacities ν. By Lemma 6.2, λ^ν depends continuously on ν. For any ν, there exists some $\varepsilon(\nu) > 0$ so that $\ell^\nu(\theta) = \ell^\circ + \theta X^\nu(\ell^\circ)$ for $\theta \in [0, \varepsilon(\nu)]$ (this is from considering the equilibrium constructed by the α-extension procedure [KS11, CCL15], which we have shown is the only possible equilibrium). Further, if we restrict our attention to some sufficiently small ball around ν°, we can choose $\varepsilon(\nu) = \varepsilon$ for some fixed $\varepsilon > 0$ for all ν in this ball. The claim thus follows from the previous lemma. □

By Lipschitz continuity, we deduce (precisely as in Section 5) that the equilibrium trajectory is continuous with respect to ν for any compact interval. This time, however, we cannot go beyond this, since now the steady-state direction can depend on ν, meaning that trajectories may slowly diverge as time goes on. So we cannot get uniform convergence of the trajectories for all time.

Variations during an evolution. We can also consider continuity with respect to perturbations during the evolution of the flow. One could delve quite deeply into modeling discussions at this point. However, we restrict our attention *solely* to Nash equilibria, and to deterministic perturbations. This means that users will take all perturbations into account when making their route choice, including ones that occur only after they depart. Note that this is not intended to model, say, unpredictable daily variations in capacity on a daily basis in morning rush-hour traffic.

In this extended abstract, we will only discuss capacity perturbations. Perturbations to transit times can also be handled, with some additional minor complications (in particular, care is required, also in the modeling, when considering decreases in transit times); see [PS20] for a discussion on this topic.

We can restrict our attention to the modification of the capacity of a single arc at a single moment of time; continuity with respect to multiple perturbations follows immediately. So suppose the capacity of arc $e' = v'w'$ (which controls the outflow rate of the arc) changes from $\nu_{e'}$ to $\hat{\nu}_{e'}$ at time $\hat{\xi}$. Let ℓ be an equilibrium trajectory with respect to the unperturbed instance, and q the corresponding queueing delays. Define $\hat{\theta}$ so that $\hat{\xi} = \ell_{v'}(\hat{\theta}) + q_{e'}(\hat{\theta})$. Then ℓ agrees with the equilibrium of the perturbed instance on the interval $[0, \hat{\theta}]$; this follows immediately from the sequential construction of dynamic equilibria (as for example described in [PS20] for time-varying capacities). We can then treat $\ell(\hat{\theta})$ as the initial conditions for an equilibrium trajectory for the instance with the new capacities. The previously stated continuity result immediately applies.

7 Conclusion

As already remarked, our continuity results can be seen as necessary requirements for the fluid queueing model to have any bearing on understanding real traffic networks. Were it the case that continuity did not hold, arbitrarily small deviations from the (simplified) model could yield completely different equilibrium behaviour. While we have shown that this is not the case, there is more work to be done in this direction. For example, one would like to show that discrete versions of the model, where users are atomic and control packets of nonzero size, behave similarly to the continuous model as the packet size goes to zero. Another question concerns the situation where users make approximately, but not necessarily exactly, optimal route choices. Our results and approach should be crucial in resolving these questions, but they do not appear to be immediate corollaries.

An important question that our results touch on concerns whether equilibria have a *finite* number of phases. Our results on long-term behavior imply that after finite time, there are no further phases. They also rule out points of accumulation "from the right": for any given θ, there exists an $\varepsilon > 0$ so that the time interval $(\theta, \theta+\varepsilon)$ lies within a single phase. What our results do not rule out is the possibility of accumulation points "from the left": that is, a moment θ such that there are an infinite number of phases in $(\theta - \varepsilon, \theta)$ for any positive ε. More geometrically, we rule out the outward-spiraling situation of Figure 4, but not an inward-spiraling trajectory. It remains an open question whether our methods can be strengthened to answer this question.

While we have restricted our attention to the setting of a single source and single sink, this is primarily to keep to the setting most considered in the literature, and so that we can directly apply, e.g., results of [CCL15] without comment. However, all our results can be extended to a more general single-commodity setting where all users have the same destination, but not necessarily the same origin; this setting (among others) has been closely studied in [SS18] (see also [Ser20]). Further, while we have restricted our attention to the setting of constant inflow, our uniqueness and continuity results carry over (essentially immediately) to the setting of piecewise-constant time-varying inflows.

Acknowledgments. We are grateful to José Correa, Andrés Cristi, Dario Frascaria, Marcus Kaiser and Tim Oosterwijk for many interesting discussions on topics related to equilibria in flows over time. The first author thanks Dario Frascaria in addition for useful discussions some years prior to this work on the topic of long-term behavior without a bounded inflow assumption.

References

[BFA15] U. Bhaskar, L. Fleischer, and E. Anshelevich. A stackelberg strategy for routing flow over time. *Games and Economic Behavior*, 92:232–247, 2015.

[CCCW17] Z. Cao, B. Chen, X. Chen, and C. Wang. A network game of dynamic traffic. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC*, pages 695–696, New York, NY, USA, 2017. Association for Computing Machinery.

[CCL15] R. Cominetti, J. Correa, and O. Larré. Dynamic equilibria in fluid queueing networks. *Operations Research*, 63(1):21–34, 2015.

[CCO19] J. Correa, A. Cristi, and T. Oosterwijk. On the price of anarchy for flows over time. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC*, pages 559–577, Phoenix, Arizona, USA, 2019. ACM.

[CCO21] Roberto Cominetti, José Correa, and Neil Olver. Long-term behavior of dynamic equilibria in fluid queuing networks. *Operations Research*, published online, 2021.

[CSSSM21] R. Cominetti, M. Scarsini, M. Schröder, and N. Stier-Moses. Approximation and convergence of large atomic congestion games, 2021.

[FF56] D. R. Fulkerson and L. R. Ford. Maximal flow through a network. *Canadian Journal of Mathematics*, 8:399–404, 1956.

[Fra21] Dario Frascaria. *Dynamic traffic equilibria with route and departure time choice*. PhD thesis, Vrije Universiteit Amsterdam, 2021.

[GHS20] L. Graf, T. Harks, and L. Sering. Dynamic flows with adaptive route choice. *Mathematical Programming*, 183(1):309–335, 2020.

[HMRT11] M. Hoefer, V.S. Mirrokni, H. Röglin, and S.-H. Teng. Competitive routing over time. *Theoretical Computer Science*, 412(39):5420–5432, 2011.

[HPS+18] T. Harks, B. Peis, D. Schmand, B. Tauer, and L. Vargas Koch. Competitive packet routing with priority lists. *ACM Transactions on Economics and Computation (TEAC)*, 6(1):1–26, 2018.

[Ism17] A. Ismaili. Routing games over time with fifo policy. In *International Conference on Web and Internet Economics, WINE*, pages 266–280, Bangalore, India, 2017. Springer International Publishing.

[KM15] Janardhan Kulkarni and Vahab S. Mirrokni. Robust price of anarchy bounds via LP and fenchel duality. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 1030–1049, 2015.

[Koc12] R. Koch. *Routing Games over Time*. PhD thesis, Technische Universität Berlin, 2012.

[KS11] R. Koch and M. Skutella. Nash equilibria and the price of anarchy for flows over time. *Theory of Computing Systems*, 49(1):71–97, 2011.

[LMR94] Frank Thomson Leighton, Bruce M. Maggs, and Satish Rao. Packet routing and job-shop scheduling in O(congestion + dilation) steps. *Combinatorica*, 14(2):167–186, 1994.

[LMR99] Frank Thomson Leighton, Bruce M. Maggs, and Andréa W. Richa. Fast algorithms for finding o(congestion + dilation) packet routing schedules. *Combinatorica*, 19(3):375–401, 1999.

[LW55] Michael James Lighthill and Gerald Beresford Whitham. On kinematic waves ii. a theory of traffic flow on long crowded roads. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 229(1178):317–345, 1955.

[PS20] H.M. Pham and L. Sering. Dynamic equilibria in time-varying networks. In *International Symposium on Algorithmic Game Theory, SAGT*, pages 130–145, Augsburg, Germany, 2020. Springer, Springer.

[Rou05] Tim Roughgarden. *Selfish routing and the price of anarchy*. MIT Press, 2005.

[Ser20] L. Sering. *Nash flows over time*. PhD thesis, Technische Universität Berlin, 2020.

[SS18] L. Sering and M. Skutella. Multi-source multi-sink Nash flows over time. In *18th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems, ATMOS*, volume 65, pages 12:1–12:20, Helsinki, Finland, 2018. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

[SST18] M. Scarsini, M. Schröder, and T. Tomala. Dynamic atomic congestion games with seasonal flows. *Operations Research*, 66(2):327–339, 2018.

[SSVK18] R. Scheffler, M. Strehler, and L. Vargas Koch. Nash equilibria in routing games with edge priorities, 2018.

[SVK19] L. Sering and L. Vargas Koch. Nash flows over time with spillback. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 935–945, San Diego, California, USA, 2019. SIAM.

[SVKZ21] L. Sering, L. Vargas Koch, and T. Ziemke. Convergence of a packet routing model to flows over time. In *Proceedings of the 2021 ACM Conference on Economics and Computation, EC*, pages 797–816, Budapest, Hungary, 2021. ACM.

[Vic69] William S Vickrey. Congestion theory and transport investment. *The American Economic Review*, 59(2):251–260, 1969.

[ZSV+21] T. Ziemke, L. Sering, L. Vargas Koch, M. Zimmer, K. Nagel, and M. Skutella. Flows over time as continuous limit of packet-based network simulations. *Transportation Research Procedia*, 52:123–130, 2021.