Algorithms for flows over time with scheduling costs[*]

Dario Frascaria[†] and Neil Olver[‡]

December 3, 2019

Abstract

Flows over time have received substantial attention from both an optimization and (more recently) a game-theoretic perspective. In this model, each arc has an associated delay for traversing the arc, and a bound on the rate of flow entering the arc; flows are time-varying. We consider a setting which is very standard within the transportation economic literature, but has received little attention from an algorithmic perspective. The flow consists of users who are able to choose their route but also their departure time, and who desire to arrive at their destination at a particular time, incurring a *scheduling cost* if they arrive earlier or later. The total cost of a user is then a combination of the time they spend commuting, and the scheduling cost they incur. We present a combinatorial algorithm for the natural optimization problem, that of minimizing the average total cost of all users (i.e., maximizing the social welfare). Based on this, we also show how to set tolls so that this optimal flow is induced as an equilibrium of the underlying game.

1 Introduction

The study of *flows over time* is a classical one in combinatorial optimization; it began already with the work of Ford and Fulkerson [10] in the 50s. It is a natural extension of static flows, which associates a single numerical value, representing a total quantity or rate of flow on the arc. In a flow over time, a second value associated with each arc represents the time it takes for flow to traverse it; the flow is then described by a function on each arc, representing the rate of flow entering the arc as a function of time.

Classical optimization problems involving static flows have natural analogs in the flow over time setting (see the surveys [17, 24]). For example (restricting the discussion to single commodity flows), the *maximum flow over time* problem asks to send as much flow as possible, departing from the source starting from time 0 and arriving to the sink by a given time horizon T; this can be solved in polynomial time [10, 11, 9]. A *quickest flow* asks, conversely, for the shortest time horizon necessary to send a given amount of flow. Of particular importance for us is the notion of an *earliest arrival flow*: this has the very strong property that simultaneously for all $T' \leq T$, the amount of flow arriving by time T' is as large as possible [12]. Such a flow can also be characterized as minimizing the average arrival time [15]. Earliest arrival flows can be "complicated", in that they can require exponential space (in the input size) to describe [29], and determining the average

arrival time of an earliest arrival flow is NP-hard [8]. But they can be constructed in time strongly polynomial in the sum of the input and output size [3].

Another important aspect of many settings were flow-over-time models are applicable—such as traffic—involves game theoretic considerations. In traffic settings, the flow is made up of a large number of individuals making their own routing choices, and aiming to maximize their own utility rather than the overall social welfare (e.g., average journey time). *Dynamic equilibria*, which is the flow over time equivalent of Wardrop equilibria for static flows, are key objects of study. Existence, uniqueness, structural and algorithmic issues, and much more have been receiving increasing recent interest from the optimization community [4, 5, 6, 7, 16, 22, 23].

Traffic, being such a relevant and important topic, has received attention from many different communities, each with their own perspective. Within the transportation economic literature, modelling other aspects of user choice besides route choice has been considered particularly important. A very standard setting, motivated by morning rush-hour traffic, is the following [26, 2].Users are able to choose not only their route, but also their *departure time*. They are then concerned not only with their journey time, but also their *arrival time* at the destination. This is captured in a *scheduling cost function* which we will denote by ρ: a user arriving at time θ will experience a scheduling cost of $\rho(\theta)$. The total disutility of a user is then the sum of their scheduling cost and their journey time (scaled by some factor $\alpha > 0$ representing their value for time spent commuting). A very standard choice of ρ is

$$\rho(\theta) = \begin{cases} -\beta\theta & \text{if } \theta \leq 0 \\ \gamma\theta & \text{if } \theta > 0 \end{cases}, \tag{1}$$

where $\beta < \alpha < \gamma$ (it is very bad to be late, but time spent in the office early is better than time spent in traffic). We will allow general scheduling cost functions, though for most of the paper we will focus on strongly unimodal cost functions; these are the most relevant, and this avoids some distracting technical details.

Two very natural questions can be posed at this point. The first is a purely optimization question, with no attention paid to the decentralized nature of traffic.

Problem 1. How can one compute a flow over time minimizing the average total cost paid by users, i.e., maximizing the social welfare?

From now on, we will call a solution to this problem simply an *optimal flow*.

It is well understood that users will typically not coordinate their actions to induce a flow that minimizes total disutility. There is a huge body of literature (particularly in the setting of static flows [20]) investigating this phenomenon. In the traffic setting, the relevance of an optimal flow represented by an answer to this question comes primarily via the possibility of *pricing*. By putting appropriate tolls on roads, we can influence the behaviour of users and the resulting dynamic equilibrium. Thus:

Problem 2. How can one set tolls (possibly time-varying) on the arcs of a given instance so that an optimal flow is obtained in dynamic equilibrium?

One subtlety is that since dynamic equilibria need not be precisely unique, there is a distinction between tolls that induce an optimal flow as *an* equilibrium, compared to tolls for which *all* dynamic equilibria are optimal. (This is called *weak* and *strong* enforcement by Harks [14] in a general pricing setting.) We will return to this subtlety shortly.

Questions like these are of great interest to transportation economists. However, most work in that community has focused on obtaining a fine-grained understanding of very restricted topologies (such as a single link, or multiple parallel links); see [25] for a survey.

Both of these question (for general network topologies) were considered by Yang and Meng [28] in a discrete time setting, by exploiting the notion of *time-expanded graphs.* This is a standard tool in the area of flows over time; discrete versions all of the optimization questions concerning flows over time mentioned earlier can (in a sense) be dealt with in this way. A node v in the graph is expanded to a collection (v, i) of nodes, for $i \in \mathbb{Z}$ in a suitable interval, and an arc vw of delay τ_{vw} becomes a collection of arcs $((v, i), (w, i + \tau_{vw}))$ (this assumes a scaling so that τ_{vw} is a length in multiples of the chosen discrete timesteps). Scheduling costs are encoded by appropriately setting arc costs from (t, i) to a supersink t' for each i, and the problem can be solved by a minimum cost static flow computation. A primary disadvantage of this approach (and in the use of time-expanded graphs more generally) is that the running time of the algorithm depends polynomially on the number of time steps, which can be very large. Further, it cannot be used to exactly solve the continuous time version (our interest in this paper); by discretizing time, it can be used to approximate it, but the size of the time-expanded graph is inversely proportional to the step size of the discretization. In the same work [28], the authors also observe that in the discrete setting, an answer to the second question can be obtained from the time-expanded graph as well. Taking the LP describing the minimum cost flow problem on the time-expanded graph, the optimal dual solution to this LP provides the necessary tolls to enforce (weakly) an optimal flow. (This is no big surprise—dual variables can frequently be interpreted as prices.)

An assumption on ρ. Suppose we consider ρ in the standard form given in (1), but with $\beta > \alpha$. This means that commuting is considered to be less unpleasant than arriving early. A user arriving earlier than time 0 at the sink would be better off "waiting" at the sink before leaving, in order to pay a scheduling cost of 0. Whether waiting in this way is allowed or not depends on the precise way one specifies the model, but it is most natural (and convenient) to allow this. If we do so, then it is clear that a scheduling cost function ρ can be replaced by

$$\hat{\rho}(\theta) := \min_{\xi \geq \theta} \rho(\xi) + \alpha(\xi - \theta)$$

without changing the optimal flow (except there is no longer any incentive to wait at the sink, and we need not even allow it). Then $\theta \to \hat{\rho}(\theta) + \alpha\theta$ is nondecreasing. From now on, we always assume that ρ satisfies this; we will call it the *growth bound* on ρ.

Our results. We give a combinatorial algorithm to compute an optimal flow. Similarly to the case of earliest arrival flows, this flow can be necessarily complicated, and involves a description length that is exponential in the input size.

The algorithm is also similar to that for computing an earliest arrival flow. It is based on the (possibly exponentially sized) path decomposition of a minimum cost flow into *successive shortest paths.* In particular, suppose we choose the scheduling cost function to be

$$\rho(\theta) = \begin{cases} -\alpha\theta & \text{if } \theta \leq 0 \\ \infty & \text{if } \theta > 0 \end{cases}. \tag{2}$$

Then the disutility a user experiences is precisely described by how much before time 0 they depart; all users must arrive by time 0 to ensure finite cost. This is precisely the reversal (both in time and direction of all arcs) of an earliest arrival flow, from the sink to the source. Our algorithm will be the same as the earliest arrival flow in this case. This also shows that it may be the case that all optimal solutions to Problem 1 require exponential size (as a function of the input encoding length), since this is the case for earliest arrival flows.

Despite the close relation to earliest arrival flows, the proof of optimality of our algorithm is rather different. A key reason for this is the following. As mentioned, earliest arrival flows have the strong property that the amount of flow arriving before a given deadline T' is the maximum possible, *simultaneously for all choices of T'* (up to some maximum depending on the total amount of flow being sent). This implies that an earliest arrival flow certainly minimizes the average arrival time amongst all possible flows [15], but is a substantially stronger property. A natural analog of this stronger property in our setting would be to ask for a flow for which, simultaneously for any given cost horizon $C' \leq C$, the amount of flow consisting of agents experiencing disutility at most C' is as large as possible. Unfortunately, in general no such flow exists. The example is too involved to discuss here, but it relates to some questions on the behaviour of dynamic equilibria in this model that are investigated in a parallel manuscript.

Since the proofs for earliest arrival flows [12, 19, 27, 3] show this stronger property which does not generalize, we take a different approach. Our proof is based on duality (of an infinite dimensional LP, though we do not require any technical results on such LPs). The main technical challenge in our work comes from determining the correct ansatz for the dual solution, as well as exploiting properties of the residual networks obtained from the successive shortest paths algorithm in precisely the right way to demonstrate certain complementary slackness conditions. As was the case with the time-expanded graph approach, the optimal dual solution immediately provides us with the tolls. However, we obtain an explicit formula for the optimal tolls, in terms of the successive shortest paths of the graph (see Section 3). This may be useful in obtaining a better structural understanding of optimal tolls, beyond just their computation. We also remark that a corollary of our result is that there is always an optimal solution without waiting (except at the source).

Consider for a moment the model where users cannot choose their departure time, but instead are released from the source at a fixed rate u_0, and simply wish to reach the destination as early as possible. This is the game-theoretic model that has received the most attention from the flow-over-time perspective [4, 6, 7, 16, 22]. Our construction of optimal tolls is applicable to this model as well. Reverse all arcs, as well as the role of the source and sink (thus making s the new sink), and also introduce a replacement sink s' and arc ss' of capacity u_0 in the original instance . Then by choosing ρ as described in (2), the optimal flow is an earliest arrival flow, and the tolls we construct will induce it in the original instance (after appropriate time reversal).

We now return to the subtlety alluded to earlier: the distinction between strongly enforcing an optimal flow, and only weakly enforcing it.
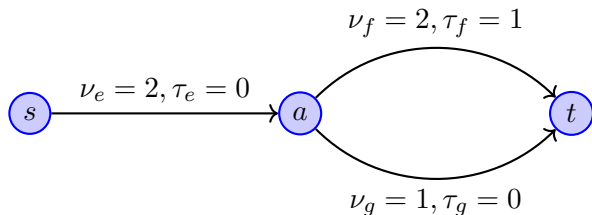

Figure 1: An instance where time-varying arc tolls cannot enforce that *all* equilibria are optimal flows.

Consider the simple instance in Figure 1. Suppose that the outflow of arc a is larger than 1 for some period in the optimum flow, due to the choice of scheduling cost function. In this period, one unit of flow would take the bottom arc g, and the rest will be routed on f. Since the total cost (including tolls) of all users is the same in a tolled dynamic equilibrium, a toll of cost equivalent to

a unit delay on arc g is needed in this period to induce the optimal flow. But then it will also be an equilibrium to send *all* flow in this period along f.

To strongly enforce an optimal flow, we need more flexible tolls. One way that we can do it is by "tolling lanes". If we are allowed to dynamically divide up the capacity of an arc into "lanes" (say a "fast lane" and a "slow lane"), and then separately set time-varying tolls on each lane, then we *can* strongly enforce any optimal flow. We discuss this further in Section 5. We are not aware of settings where this phenomenon has been previously observed, and it would be interesting to explore this further in a more applied context.

Outline of the paper. We introduce some basic notation and notions, as well as a formally define our model, in Section 2. In Section 3, we describe our algorithm, and show that it returns a feasible flow over time; we restrict ourselves to the most relevant case of a strictly unimodal scheduling cost function. In Section 4 we show optimality of this algorithm, and in Section 5 we derive optimal tolls from this analysis. Finally, in Section 6 we remark on the technical changes needed to the algorithm to handle general scheduling cost functions.

2 Model and preliminaries

The notation $(z)^+$ is used to denote the nonnegative part of z, i.e., $(z)^+ = \max\{z, 0\}$. Given $v : X \to \mathbb{R}$ and $A \subseteq X$, we will use the shorthand notation $v(A) := \sum_{a \in A} v(a)$. We will not distinguish between a map $v : X \to \mathbb{R}$ and a vector in $\mathbb{R}^X$, and so the notation v_a and $v(a)$ is interchangeable. All graphs considered will be directed. We assume all graphs to be simple, and that there are no digons (i.e., there are no pairs $v, w \in V$ so that vw and wv are both arcs). This is for notational convenience only—this restriction can easily be lifted.

We begin with some basic notions and results about static flows and flows over time. For further details regarding static flows, we refer the reader to the book by Ahuja, Magnanti and Orlin [1]. For more about flows over time, we suggest the surveys by Skutella [24] and Köhler et al. [17].

Static flows. Let $G = (V, E)$ be a directed graph, with source node $s \in V$ and sink node $t \in V$. Each arc $e \in E$ has a *capacity* ν_e and a *delay* τ_e (both nonnegative). We use $\delta^+(v)$ to denote the set of arcs in E with tail v, and $\delta^-(v)$ the set of arcs with head v.

Consider some $f : E \to \mathbb{R}_+$ (which we will equivalently view as a vector in $\mathbb{R}_+^E$). For $v \in V$, we define the *net flow* at v (denoted ∇f_v) to be the quantity

$$\nabla f_v := f(\delta^-(v)) - f(\delta^+(v)) = \sum_{e \in \delta^-(v)} f_e - \sum_{e \in \delta^+(v)} f_e.$$

We say that f is a *(static) s-t-flow of value Q* if

(i) $\nabla f_v = 0$ for all $v \in V \setminus \{s, t\}$, with $\nabla f_t = -\nabla f_s = Q$; and

(ii) $f_e \leq \nu_e$ for all $e \in E$.

Given an *s-t*-flow f, its *residual network* $G^f = (V, E^f)$ is defined by

$$E^f = \{vw : vw \in E \text{ and } f_{vw} < \nu_{vw}\} \cup \{vw : wv \in E \text{ and } f_{wv} > 0\}.$$

Call arcs in $E^f \cap E$ *forward arcs* and arcs in $E^f \setminus E$ *backwards arcs*. The *residual capacity* ν_e^f of an arc $e \in E^f$ is then $\nu_{vw}^f = \nu_{vw} - f_{vw}$ for vw a forward arc, and $\nu_{vw}^f = f_{wv}$ for vw a backwards arc. We also define $\tau_{vw} = -\tau_{wv}$ for all backwards arcs vw.

5

Given a subset $F \subseteq E$, we use $\chi(P)$ to denote the characteristic vector of F. In particular, if P is a path from v to w, then $\chi(P)$ is a unit flow from v to w.

We make the definitions $\overleftarrow{E} := \{wv : vw \in E\}$ and $\overleftrightarrow{E} := E \cup \overleftarrow{E}$. Given $f, g \in \mathbb{R}_+^E$, we define $f + g$ in the obvious way, and also define $f - g \in \mathbb{R}_+^{\overleftrightarrow{E}}$, by interpreting a negative value on vw instead as a positive value on wv.

Given a choice of value Q, a *minimum cost flow* is an s-t-flow f^* minimizing $\sum_{e \in E} f_e^* \tau_e$ (amongst all s-t-flows of value Q). An s-t-flow f (of the correct value) is a minimum cost flow if and only if E^f contains no negative cost cycles, i.e., cycles $C \subseteq E^f$ with $\tau(C) < 0$.

Flows over time. Consider some $f : E \times \mathbb{R} \to \mathbb{R}_+$. We will generally write $f_e(\theta)$ rather than $f(e, \theta)$. Define the *net flow into v at time θ* by

$$\nabla f_v(\theta) := \sum_{e \in \delta^-(v)} f_e(\theta - \tau_e) - \sum_{e \in \delta^+(v)} f_e(\theta).$$

Note that $f_e(\theta)$ represents the flow *entering* arc e at time θ; this flow will exit the arc at time $\theta + \tau_e$ (explaining the asymmetry between the terms for flow entering and flow leaving in the above).

We say that f is a *flow over time of value Q* if the following hold.

(i) f has compact support (i.e., for some K, $f_e(\theta) = 0$ whenever $|\theta| > K$).

(ii) $\int_{-\infty}^{\infty} \nabla f_v(\theta) d\theta = Q(\mathbf{1}_{v=t} - \mathbf{1}_{v=s})$ for all $v \in V$.

(iii) $\int_{-\infty}^{\xi} \nabla f_v(\theta) d\theta \geq 0$ for all $v \in V \setminus \{s, t\}$ and $\xi \in \mathbb{R}$.

(iv) $f_e(\theta) \leq \nu_e$ for all $e \in E$ and $\theta \in \mathbb{R}$.

Note that this definition allows for flow to wait at a node; to disallow this and consider only *flows over time without waiting*, we would replace (iii) with the condition

(iii′) $\nabla f_v(\theta) = 0$ for all $v \in V \setminus \{s, t\}$ and $\theta \in \mathbb{R}$.

We also have a natural notion of a residual network in the flow over time setting. Define, for any flow over time f and $\theta \in \mathbb{R}$,

$$E^f(\theta) = \{vw : vw \in E \text{ and } f_{vw}(\theta) < \nu_{vw}\} \cup \{vw : wv \in E \text{ and } f_{wv}(\theta - \tau_{wv}) > 0\}.$$

Minimizing scheduling cost. We are concerned with the following optimization problem. Given a *scheduling cost function* $\rho : \mathbb{R} \to \mathbb{R}_+$, as well as a value $\alpha > 0$, determine a flow over time f of value Q that minimizes the sum of the *commute cost* $\alpha \sum_{e \in E} \int_{\mathbb{R}} f_e(\theta) d\theta$ and the *scheduling cost* $\int_{\mathbb{R}} \nabla f_t(\theta) \cdot \rho(\theta) d\theta$. As already discussed, we assume that ρ satisfies the growth bound, i.e., that $\theta \to \rho(\theta) + \alpha\theta$ is nondecreasing. This ensures that waiting at t is not needed, which is in fact disallowed by our definition[1], and makes various arguments cleaner. We will also make the assumption that ρ is strongly unimodal[2]. We then assume w.l.o.g. that the minimizer of ρ is at 0, and that $\rho(0) = 0$. For further technical convenience, by adjusting ρ on a set of measure zero we take ρ to be lower semi-continuous.

The unimodal assumption is not necessary; the algorithm and analysis can be extended to essentially general ρ, under some very weak technical conditions. We postpone this discussion to the end of the paper.

---

[1] Were this really needed, one could simply add a dummy arc tt' to a new sink t'.
[2] I.e., (strictly) decreasing until some moment, and then (strictly) increasing.

We also assume that we are able to query $\rho^{-1}(y)$ for a given rational $y > 0$, obtaining a pair of solutions (one positive, one negative) of moderate bit complexity. Alternatively, we can consider the case where ρ is represented as a piecewise linear function.

3   A combinatorial algorithm

In this section we present an algorithm that computes an optimal flow over time, assuming that ρ is strongly unimodal. The proof of optimality is discussed in Section 4.

We begin by recalling the *successive shortest paths (SSP)* algorithm for computing a minimum cost static flow. It is not a polynomial time algorithm, so it is deficient as an algorithm for static flows, but it provides a structure that is relevant for flows over time. This is of course well known from its role in constructing earliest arrival flows, which we will briefly detail.

The SSP algorithm construct a sequence of paths $(P_1, P_2, \ldots)$ and associated amounts $(x_1, x_2, \ldots)$ inductively as follows. Suppose $P_1, \ldots, P_j$ and $x_1, \ldots, x_j$ have been defined. Let $f^{(j)} = \sum_{i=1}^{j} x_i \chi(P_i)$, and let G_j denote the residual graph of $f^{(j)}$. Also let $d_j(v, w)$ denote the length (w.r.t. arc delays τ in G_j) of a shortest path from v to w in G_j (this may be infinite). By construction, G_j will contain no negative cost cycles, so that d_j is computable. If $d_j(s, t) = \infty$, we are done; set $m := j$. Otherwise, define P_{j+1} to be any shortest s-t-path in G_j, and x_{j+1} the minimum capacity in G_j of an arc in P_{j+1}. It can be shown that $\sum_{j=1}^{r} \tilde{x}_j \chi(P_j)$, with $\sum_{j=1}^{r} \tilde{x}_j = Q$ and $\tilde{x}_j = x_j$ for $j < r$, $0 \leq \tilde{x}_r \leq x_r$, is a minimum cost flow of value Q, as long as Q is not larger than the value of a maximum flow.

To construct an earliest arrival flow of value Q and time horizon T, we (informally) send flow at rate x_j along path P_j for the time interval $[0, T - \tau(P_j)]$, for each $j \in [m]$ (if $\tau(P_j) > T$, we send no flow along the path). By this, we mean that for each $e = vw \in P_j$, we increase by x_j the value of $f_e(\theta)$ for $\theta \in [d_{j-1}(s, v), T - d_{j-1}(v, t)]$ (or if e is a backwards arc, we instead decrease $f_{wv}(\theta)$). An argument is needed to show that this defines a valid flow, since we must not violate the capacity constraints, and moreover, P_j may contain reverse arcs not present in G.

We are now ready to describe our algorithm for minimizing the disutility, which is a natural variation on the earliest arrival flow algorithm. It is also constructed from the successive shortest paths, but using a *cost horizon* rather than a *time horizon*. For now, consider C to be a given value (it will be the "cost horizon"). For each $j \in [m]$ with $\alpha d_{j-1}(s, t) \leq C$, we send flow at rate x_j along path P_j for the time interval $[a_j, b_j]$ chosen maximally so that

$$\rho(\xi + d_{j-1}(s, t)) \leq C - \alpha d_{j-1}(s, t) \quad \text{for all} \quad \xi \in [a_j, b_j].$$

(If ρ is continuous, then of course $\rho(a_j + d_{j-1}(s, t)) = \rho(b_j + d_{j-1}(s, t)) = C - \alpha d_{j-1}(s, t)$.) Note that a user leaving at time a_j or b_j and using path P_j, without waiting at any moment, incurs disutility C; whereas a user leaving at some time $\theta \in (a_j, b_j)$ and using path P_j will incur a strictly smaller total cost.

As we will shortly argue, this results in a feasible flow over time f. Given this, its value will be $\sum_{j=1}^{m} x_j(b_j - a_j)$. It is easy to see that this value changes continuously and monotonically with C (here we use the strong unimodality). Thus a bisection search can be used to determine the correct choice of C for a given value Q. Alternatively, bisection search can be avoided by using Megiddo's parametric search technique [18]; this will ensure a strongly polynomial running time, if queries to ρ^{-1} are considered to be of unit cost.

This requires only oracle access to ρ^{-1} (and reasonable control on the bit complexity of $\rho^{-1}(y)$ in terms of y). If ρ is piecewise linear with not too many breakpoints (as is the case, in particular, for the "standard" β/γ choice generally used in the transportation economics literature), a third

option presents itself. One can generate the entire parametric curve of C as a function of Q, which will also be piecewise linear, from which the correct choice of C can readily be determined. This is fairly straightforward, and we omit the details.

Feasibility. In the following we show that the flow resulting flow f is a feasible flow over time. Given a vertex $v \in V$, a time $\theta \in \mathbb{R}$ and $j \in [m]$, let

$$c_j(v, \theta) = \alpha d_{j-1}(s, t) + \rho(\theta + d_{j-1}(v, t)).$$

If $v \in P_j$ then $c_j(v, \theta)$ is the travel cost of a user that utilizes path P_j and passes through node v at time θ; there does not seem to be a simple interpretation if $v \notin P_j$ however. Now define

$$J(v, \theta) = \max\{j \in [m] : c_j(v, \theta) \leq C\}, \tag{3}$$

with the convention that the maximum over the empty set is 0. The motivation for this definition comes from the following theorem, which completely characterizes f. (If preferred, one could even think of this theorem as providing the definition of f.)

Theorem 1. $f_{vw}(\theta) = f_{vw}^{(J(v,\theta))}$ for any $vw \in E$ and $\theta \in \mathbb{R}$.

Corollary 2. f is a feasible flow over time.

Proof. By the way that we constructed f, it has value Q and satisfies flow conservation. Only nonnegativity and the capacity constraint remain, which follows from the theorem. $\square$

Before proving Theorem 1, we need the following lemma.

Lemma 3. $c_j(v, \theta)$ is nondecreasing with j for any $\theta \in \mathbb{R}$.

Proof. Consider any $j \in [m-1]$; we show that $c_{j+1}(v, \theta) \geq c_j(v, \theta)$. Suppose Q is a shortest v-t-path in G_{j-1}, so $\tau(Q) = d_{j-1}(v, t)$. Consider the unit v-t flow $g = \chi(P_{j+1}) - \chi(P_j) + \chi(Q)$ in $\overleftrightarrow{E}$. Now observe that the support of g is contained in G_j: P_{j+1} and $\overleftarrow{P_j}$ are certainly contained in G_j; and if $e \in Q \cap (E_{j-1} \setminus E_j)$, then $e \in P_j$. Since G_j contains no negative cost cycles, the cost of g is at least that of a shortest v-t-path in G_j, and so $d_j(v, t) \leq \tau(P_{j+1}) - \tau(P_j) + \tau(Q)$. Finally, we can conclude

$$\begin{aligned}
\alpha d_j(s, t) + \rho(\theta + d_j(v, t)) &= \alpha d_j(s, t) + \rho(\theta + d_{j-1}(v, t)) - \rho(\theta + d_{j-1}(v, t)) + \rho(\theta + d_j(v, t)) \\
&\geq \alpha d_j(s, t) + \rho(\theta + d_{j-1}(v, t)) - \alpha(d_j(v, t) - d_{j-1}(v, t)) \\
&\geq \alpha d_{j-1}(s, t) + \rho(\theta + d_{j-1}(v, t)),
\end{aligned}$$

where the first inequality follows from the growth assumption, using $d_j(v, t) \geq d_{j-1}(v, t)$. $\square$

Proof of Theorem 1. Fix some $vw \in E$ and $\theta \in \mathbb{R}$. Consider now any P_j (with $\alpha \tau(P_j) \leq C$, so that it is used for a nontrivial interval), with $vw \in P_j$. Since P_j is a shortest path in G_{j-1}, if we send flow along this path starting from some time ξ, it will arrive at v at time $\xi + d_{j-1}(s, v)$. Considering the definition of the interval $[a_j, b_j]$, we see that P_j contributes flow to vw at time θ if $c_j(v, \theta) \leq C$. By Lemma 3, this occurs precisely if $j \leq J(v, \theta)$.

Considering in similar fashion paths P_j with $wv \in P_j$ (and noting that $J(w, \theta + \tau_{vw}) = J(v, \theta)$), we determine that

$$f_{vw}(\theta) = \sum_{\substack{j:vw\in P_j \\ j\leq J(v,\theta)}} x_j - \sum_{\substack{j:wv\in P_j \\ j\leq J(v,\theta)}} x_j = f_{vw}^{(J(v,\theta))}. \qquad \square$$

8

4 Optimality

In this section, we show that our proposed algorithm does return an optimal flow.

4.1 Duality-based certificates of optimality

We can write the problem we are interested in as a (doubly) infinite linear program as follows:

$$\min \quad \int_{-\infty}^{\infty} \rho(\theta) \nabla f_t(\theta) d\theta + \alpha \sum_{e \in E} \tau_e \int_{-\infty}^{\infty} f_e(\theta) d\theta + \alpha \sum_{v \in V \setminus \{s,t\}} \int_{-\infty}^{\infty} z_v(\theta) d\theta$$

$$\text{s.t.} \quad \int_{-\infty}^{\infty} \nabla f_s(\theta) d\theta = -Q$$

$$\int_{-\infty}^{\infty} \nabla f_t(\theta) d\theta = Q \tag{4}$$

$$\int_{-\infty}^{\theta} \nabla f_v(\xi) d\xi = z_v(\theta) \qquad \forall v \in V \setminus \{s,t\}, \theta \in \mathbb{R}$$

$$f_e(\theta) \leq \nu_e \qquad \forall e \in E, \theta \in \mathbb{R}$$

$$z, f \geq 0$$

Here, $z_v(\theta)$ represents the amount of flow waiting at node v at time θ (which must always be nonnegative). The travel cost is captured on a per-arc basis, including waiting time as well.

The following theorem provides a certificate of optimality of a feasible solution to (4).

Theorem 4. *Let f be a flow over time with value Q, and suppose that $\pi : V \times \mathbb{R} \to \mathbb{R}$ satisfies the following, for some choice of C:*

(i) $\theta \to \pi_v(\theta) - \alpha\theta$ is nonincreasing.

(ii) $\pi_w(\theta + \tau_{vw}) \leq \pi_v(\theta) + \alpha\tau_{vw}$ for all $\theta \in \mathbb{R}, vw \in E^f(\theta)$.

(iii) $\pi_s(\theta) = 0$ for all $\theta \in \mathbb{R}$.

(iv) $\pi_t(\theta) = (C - \rho(\theta))^+$ for all $\theta \in \mathbb{R}$, and $\nabla f_t(\theta) = 0$ whenever $\rho(\theta) > C$.

Then f is an optimal solution.

Essentially, $\pi_v(\theta)$ are dual variables, and the assumptions of the theorem are that f and π satisfy the complementary slackness conditions. There are many extensions of LP duality theory to infinite dimensional settings, e.g., [13, 21]; however the situation is subtle, since strong duality and even weak duality can fail [21]. We prefer to avoid technicalities and derive it directly.

Proof of Theorem 4. We will need the following technical lemma (obvious via integrating by parts in the case that h is also absolutely continuous).

Claim 5. *Let $h : \mathbb{R} \to \mathbb{R}$ be a nonincreasing function, and $z : \mathbb{R} \to \mathbb{R}_+$ be an absolutely continuous nonnegative function with compact support. Then $\int_{\mathbb{R}} h(\theta) z'(\theta) d\theta \geq 0$.*

Proof. Since $\int_{\mathbb{R}} z'(\theta) d\theta = 0$, we may assume by shifting if necessary that h is nonnegative on the support of z. Let μ be a measure so that $\mu([\theta, \infty)) = h(\theta)$ for almost every θ in the support of z. We certainly have that for any θ,

$$\int_{-\infty}^{\theta} z'(\xi) d\xi = \left[z(\xi) \right]_{-\infty}^{\theta} = z(\theta) - 0 \geq 0.$$

9

Thus

$$\int_{\mathbb{R}} \int_{\mathbb{R}} \mathbf{1}_{\xi \leq \theta} z'(\xi) d\xi d\mu(\theta) \geq 0,$$

from which we obtain the result by applying Fubini's theorem. $\qquad\square$

Define, for each $vw \in E$,

$$\delta_{vw}(\theta) = (\pi_w(\theta + \tau_{vw}) - \pi_v(\theta) - \alpha\tau_{vw})^+.$$

Now let g, z be any feasible solution to (4) with compact support. Consider any $v \in V \setminus \{s, t\}$, and observe that

$$\int_{\mathbb{R}} (\pi_v(\theta)\nabla g_v(\theta) + \alpha z_v(\theta)) d\theta = \int_{\mathbb{R}} ((\pi_v(\theta) - \alpha\theta)\nabla g_v(\theta)) d\theta \; + \; \left[\alpha\theta z_v(\theta)\right]_{-\infty}^{\infty} \geq 0, \qquad (5)$$

by the above claim, exploiting property (i).

We have

$$\text{cost}(g) = \int_{\mathbb{R}} \rho(\theta)\nabla g_t(\theta) d\theta + \alpha \sum_{e \in E} \int_{\mathbb{R}} \tau_e g_e(\theta) d\theta + \alpha \sum_{v \in V \setminus \{s,t\}} \int_{\mathbb{R}} z_v(\theta) d\theta$$

$$\overset{(*)}{\geq} \int_{\mathbb{R}} (C - \pi_t(\theta))\nabla g_t(\theta) d\theta + \alpha \sum_{vw \in E} \int_{\mathbb{R}} \big(\pi_w(\theta + \tau_{vw}) - \pi_v(\theta) - \delta_{vw}(\theta)\big) g_e(\theta) d\theta$$

$$+ \alpha \sum_{v \in V \setminus \{s,t\}} \int_{\mathbb{R}} z_v(\theta) d\theta$$

$$\overset{(**)}{=} CQ + \sum_{v \in V \setminus \{s,t\}} \int_{\mathbb{R}} [\pi_v(\theta)\nabla g_v(\theta) + \alpha z_v(\theta)] d\theta - \sum_{e \in E} \delta_e(\theta) g_e(\theta)$$

$$\overset{(***)}{\geq} CQ - \sum_{e \in E} \int_{\mathbb{R}} \delta_e(\theta)\nu_e d\theta.$$

In the above, $(*)$ holds by the definitions of π_t and δ_e; $(**)$ follows by recombining the $g_e(\theta)$ terms and recalling that $\pi_s \equiv 0$ and that g has value Q; and $(***)$ follows from (5), and the inequalities $\delta_e(\theta) \geq 0$ and $g_e(\theta) \leq \nu_e$ that hold for all $e \in E$ and $\theta \in \mathbb{R}$.

Finally, observe that all of the inequalities in the above hold with equality if $g = f$. Property (ii) implies that if $f_{vw}(\theta) > 0$ (so that $wv \in E^f(\theta)$), then $\delta_{vw}(\theta) = \pi_w(\theta + \tau_{vw}) - \pi_v(\theta) - \alpha\tau_{vw}$, yielding equality in $(*)$. It also implies that if $f_{vw}(\theta) < \nu_{vw}$ (so that $vw \in E^f(\theta)$) then $\delta_{vw}(\theta) = 0$, yielding equality in $(***)$. $\qquad\square$

As is often the case, the optimal dual solution also provides us the prescription for tolls to induce the optimum flow. We delay this discussion to Section 5.

4.2 The dual prescription

We now give a certificate of optimality $\pi : V \times \mathbb{R} \to \mathbb{R}$ for (4) that satisfies the conditions of Theorem 4. Given a vertex $v \in V$ and a time $\theta \in \mathbb{R}$ let

$$\pi_v(\theta) = \max\{\pi'_v(\theta), \bar{\pi}_v(\theta), 0\}$$

where

$$\pi'_v(\theta) = -\alpha d_{J(v,\theta)}(v,s),$$
$$\bar{\pi}_v(\theta) = C - \alpha d_{J(v,\theta)}(v,t) - \rho(\theta + d_{J(v,\theta)}(v,t)).$$

Notice that $\pi_s(\theta) = 0$ and $\pi_t(\theta) = \max\{C - \rho(\theta), 0\}$ for all $\theta \in \mathbb{R}$ and thus conditions (iii) and (iv) of Theorem 4 hold. For the remaining conditions, we begin with some basic facts about distance labels associated with successive shortest paths (statements of a similar flavour can be found in Ahuja et al. [1], for example).

Lemma 6. $d_j(v,s)$ *is nonincreasing with j, for every $v \in V$.*

Proof. We show that $d_{j-1}(v,s) \geq d_j(v,s)$ for any $j \in [m]$. Let Q be a shortest v-s-path in G_{j-1} (if there is no such path, then there is nothing to prove). If $P_j \cap Q = \emptyset$, then $Q \subseteq E_j$, so the claim holds. Otherwise, let xy be the last (i.e., closest to s) edge of Q that is also in P_j, and let Q_y denote the subpath of Q from v to y. Also let P_y denote the subpath of P_j from y to v. Then $C := Q_y + P_y$ is a directed cycle (or collection of cycles, possibly with some edges included twice). Since $C \subseteq E_{j-1}$, $\tau(C) \geq 0$, so $\tau(\overleftarrow{P}_y) \leq \tau(Q_y)$.

Define Q' to be the path obtained by appending the subpath of Q from y to s to $\overleftarrow{P}_y$. Then $Q' \subseteq E_j$, and $\tau(Q') \leq \tau(Q)$, as required. $\square$

Lemma 7. *For all $j \in [m]$ and $v \in V$, $d_{j-1}(v,t) - d_{j-1}(s,t) = d_j(v,s)$.*

Proof. To show that $d_{j-1}(v,t) - d_{j-1}(s,t) \leq d_j(v,s)$, let Q be a shortest v-s-path in G_j, and let Q' be a v-t-path contained in $P_j + Q$ (arcs in opposite directions are cancelled). Then Q' is in G_{j-1}; if e is an arc in Q not in G_{j-1}, then e is the reverse of an arc of P_j, and hence not in $P_j + Q$. So $d_{j-1}(v,t) \leq \tau(Q')$. But since G_{j-1} has no negative cost cycles, $\tau(Q') \leq \tau(Q) + \tau(P_j)$.

To show that $d_{j-1}(v,t) - d_{j-1}(s,t) \geq d_j(v,s)$, let $\bar{Q}$ be a shortest v-t-path in G_{j-1}. Let w be the first (i.e., closest to v) vertex present in both $\bar{Q}$ and P_j (notice that w might be equal to v or t) and let Q be the v-w-path contained in $\bar{Q}$. Then

$$d_{j-1}(v,t) = d_{j-1}(w,t) + \tau(Q). \tag{6}$$

Since $Q \subseteq E_j$, we have that:

$$
\begin{aligned}
d_j(v,s) &\leq d_j(w,s) + \tau(Q) \\
&= d_{j-1}(w,t) - d_{j-1}(s,t) + \tau(Q) &&\text{since } w \in P_j \\
&= d_{j-1}(v,t) - d_{j-1}(s,t) &&\text{by (6).}
\end{aligned}
$$

This concludes the proof. $\square$

Now we are ready to show that π satisfies conditions (i) and (ii) of Theorem 4.

Lemma 8. $\theta \to \pi_v(\theta) - \alpha\theta$ *is nonincreasing.*

Proof. Fix any $\theta \in \mathbb{R}$ and $\epsilon \geq 0$. We show that $\pi_v(\theta) \geq \pi_v(\theta + \epsilon) - \alpha\epsilon$. Let $j := J(v,\theta)$ and $\ell := J(v,\theta + \epsilon)$.

- **Case 1:** $\pi_v(\theta + \epsilon) = -\alpha d_\ell(v,s)$.

  If $\ell \leq j$, then by Lemma 6

  $$\pi_v(\theta) \geq \pi'_v(\theta) = -\alpha d_j(v,s) \geq -\alpha d_\ell(v,s) = \pi_v(\theta + \epsilon).$$

11

So suppose $\ell > j$. By the definition of $J(v, \theta + \epsilon)$, we know that

$$\alpha d_{\ell-1}(s,t) + \rho(\theta + \epsilon + d_{\ell-1}(v,t)) \leq C. \tag{7}$$

As a consequence, we have that:

$$
\begin{aligned}
\pi_v(\theta) &\geq \bar{\pi}_v(\theta) \\
&= C - \alpha d_j(v,t) - \rho(\theta + d_j(v,t)) \\
&\overset{(*)}{\geq} C - \alpha d_j(v,t) - \rho(\theta + \epsilon + d_{\ell-1}(v,t)) - \alpha\left(\epsilon + d_{\ell-1}(v,t) - d_j(v,t)\right) \\
&\geq \alpha d_{\ell-1}(s,t) - \alpha\epsilon - \alpha d_{\ell-1}(v,t) && \text{by (7)} \\
&= -\alpha\epsilon - \alpha d_\ell(v,s) && \text{by Lemma 7} \\
&= \pi_v(\theta + \epsilon) - \alpha\epsilon.
\end{aligned}
$$

Inequality $(*)$ follows from the growth assumption on ρ combined with the fact that $\theta + \epsilon + d_{\ell-1}(v,t) \geq \theta + d_j(v,t)$.

- **Case 2:** $\pi_v(\theta + \epsilon) = C - \alpha d_\ell(v,t) - \rho(\theta + \epsilon + d_\ell(v,t))$.

  If $\ell \geq j$, then:

$$
\begin{aligned}
\pi_v(\theta) &\geq \bar{\pi}_v(\theta) \\
&= C - \alpha d_j(v,t) - \rho(\theta + \epsilon + d_j(v,t)) \\
&= C - \alpha d_j(v,t) - \rho(\theta + \epsilon + d_\ell(v,t)) + \rho(\theta + \epsilon + d_\ell(v,t)) - \rho(\theta + d_j(v,t)) \\
&\geq C - \alpha d_j(v,t) - \rho(\theta + \epsilon + d_\ell(v,t)) - \alpha\left(\epsilon + d_\ell(v,t) - d_j(v,t)\right) \\
&= C - \rho(\theta + \epsilon + d_\ell(v,t)) - \alpha\epsilon - \alpha d_\ell(v,t) \\
&= \pi_v(\theta + \epsilon) - \alpha\epsilon.
\end{aligned}
$$

The second inequality follows again from the growth assumption, this time combined with the inequality $d_\ell(v,t) \geq d_j(v,t)$.

If $\ell < j$, by definition of $J(v, \theta + \epsilon)$ we have that

$$\alpha d_\ell(s,t) + \rho(\theta + \epsilon + d_\ell(v,t)) > C.$$

From this, we obtain

$$
\begin{aligned}
\pi_v(\theta) &\geq \pi'_v(\theta) \\
&= -\alpha d_j(v,s) \\
&> C - \alpha d_\ell(s,t) - \rho(\theta + \epsilon + d_\ell(v,t)) - \alpha d_j(v,s) \\
&\geq C - \alpha d_\ell(s,t) - \rho(\theta + \epsilon + d_\ell(v,t)) - \alpha d_{\ell+1}(v,s) && \text{by Lemma 6} \\
&= C - \alpha d_\ell(v,t) - \rho(\theta + \epsilon + d_\ell(v,t)) && \text{by Lemma 7} \\
&= \pi_v(\theta).
\end{aligned}
$$

$\square$

Lemma 9. If $vw \in E^f(\theta)$, then $\pi_w(\theta + \tau_{vw}) \leq \pi_v(\theta) + \alpha\tau_{vw}$.

Proof. Let $j := J(v, \theta)$ and $\ell := J(w, \theta + \tau_{vw})$. Note that since $vw \in E^f(\theta)$, Theorem 1 implies that $vw \in E_j$.

- **Case 1:** $\pi_w(\theta + \tau_{vw}) = -\alpha d_\ell(w, s)$.

  If $\ell \le j$, then

  $$\begin{aligned}
  \pi_v(\theta) &\ge -\alpha d_j(v, s) \\
  &\ge -\alpha \tau_{vw} - \alpha d_j(w, s) && \text{since } vw \in E_j \\
  &\ge -\alpha \tau_{vw} - \alpha d_\ell(w, s) && \text{by Lemma 6} \\
  &= \pi_w(\theta + \tau_{vw}) - \alpha \tau_{vw}.
  \end{aligned}$$

  So suppose $\ell > j$. By the definition of $J(w, \theta + \tau_{vw})$ we know that

  $$\alpha d_{\ell-1}(s, t) + \rho(\theta + \tau_{vw} + d_{\ell-1}(w, t)) \le C. \tag{8}$$

  Since $vw \in E_j$ and $d_j(w, t) \le d_{\ell-1}(w, t)$, we also have

  $$\theta + d_j(v, t) \le \theta + \tau_{vw} + d_{\ell-1}(w, t). \tag{9}$$

  Thus

  $$\begin{aligned}
  \pi_v(\theta) &\ge C - \alpha d_j(v, t) - \rho(\theta + d_j(v, t)) \\
  &\ge C - \alpha d_j(v, t) - \rho(\theta + \tau_{vw} + d_{\ell-1}(w, t)) - \alpha \left(\tau_{vw} + d_{\ell-1}(w, t) - d_j(v, t) \right) \\
  &\ge \alpha d_{\ell-1}(s, t) - \alpha \tau_{vw} - \alpha d_{\ell-1}(w, t) && \text{by (8)} \\
  &= -\alpha \tau_{vw} - \alpha d_\ell(w, s) && \text{by Lemma 7} \\
  &= \pi_w(\theta + \tau_{vw}) - \alpha \tau_{vw}
  \end{aligned}$$

  where the second inequality follows from the growth assumption and from (9).

- **Case 2:** $\pi_w(\theta + \tau_{vw}) = C - \alpha d_\ell(w, t) - \rho(\theta + \tau_{vw} + d_\ell(w, t))$.

  If $\ell \ge j$, since $vw \in E_j$ and $d_j(w, t) \le d_\ell(w, t)$, we have that

  $$\theta + d_j(v, t) \le \theta + \tau_{vw} + d_\ell(w, t). \tag{10}$$

  As a consequence, exploiting also the growth assumption, we have

  $$\begin{aligned}
  \pi_v(\theta) &\ge \bar{\pi}_v(\theta) \\
  &= C - \alpha d_j(v, t) - \rho(\theta + d_j(v, t)) \\
  &= C - \alpha d_j(v, t) - \rho(\theta + \tau_{vw} + d_\ell(w, t)) + \rho(\theta + \tau_{vw} + d_\ell(w, t)) - \rho(\theta + d_j(v, t)) \\
  &\ge C - \alpha d_j(v, t) - \rho(\theta + \tau_{vw} + d_\ell(w, t)) - \alpha \left(\tau_{vw} + d_\ell(w, t) - d_j(v, t) \right) \\
  &= C - \rho(\theta + \tau_{vw} + d_\ell(w, t)) - \alpha \tau_{vw} - \alpha d_\ell(w, t) \\
  &= \pi_w(\theta + \tau_{vw}) - \alpha \tau_{vw}.
  \end{aligned}$$

  If $\ell < j$, by definition of $J(w, \theta + \tau_{vw})$ we have that

  $$\alpha d_\ell(s, t) + \rho(\theta + \tau_{vw} + d_\ell(w, t)) > C. \tag{11}$$

  Thus

  $$\begin{aligned}
  \pi_v(\theta) &\ge -\alpha d_j(v, s) \\
  &\ge -\alpha d_j(w, s) - \alpha \tau_{vw} && \text{since } vw \in E_j \\
  &\ge -\alpha d_{\ell+1}(w, s) - \alpha \tau_{vw} && \text{by Lemma 6} \\
  &> C - \alpha d_\ell(s, t) - \rho(\theta + \tau_{vw} + d_\ell(w, t)) - \alpha d_{\ell+1}(w, s) - \alpha \tau_{vw} && \text{by (11)} \\
  &= C - \alpha d_\ell(w, t) - \rho(\theta + \tau_{vw} + d_\ell(w, t)) - \alpha \tau_{vw} && \text{by Lemma 7} \\
  &= \pi_w(\theta + \tau_{vw}) - \alpha \tau_{vw}.
  \end{aligned}$$

□

5   Optimal tolls

Tolls $\delta : E \times \mathbb{R} \to \mathbb{R}_+$ are per-arc, time-varying and nonnegative. The value $\delta_e(\xi)$ represents the toll a user is charged upon entering the link at time ξ.

We have the following theorem.

Theorem 10. *Let (f, π) be an optimal primal-dual solution to* (4) *(as constructed in Section 3 and Section 4) and define, for each $vw \in E$,*

$$\delta_{vw}(\theta) = (\pi_w(\theta + \tau_{vw}) - \pi_v(\theta) - \alpha\tau_{vw})^+.$$

Then f is a dynamic equilibrium under tolls δ.

Of course, to make sense of this theorem we must know what is meant by a dynamic equilibrium under tolls. A precise definition requires introducing the full game-theoretic fluid queueing model (also known as the Vickrey bottleneck model) [26, 16]. Tolls and departure time choice can be introduced into the definition of a dynamic equilibrium discussed in these works. Rather than going this route, we will show that the tolls satisfy a strong property that very clearly ensures the equilibrium property.

We show that the following holds. A user starting from some $v \in V$ at some time $\theta \in \mathbb{R}$ cannot incur a total cost (including scheduling cost, and tolls and commuting cost from this point forward) less than $C - \pi_v(\theta)$. This is even allowing the user to take any link at any time, as if no other users were present in the network. Since the flow represents a solution where all users incur a total cost of precisely C, this must certainly be an equilibrium.

To see this, consider any s-t-path P in E and arrival times θ_v for each $v \in P$ valid for this path; so $\theta_w \geq \theta_v + \tau_{vw}$ for every $vw \in P$. Thus by properties (i) and (ii) of Theorem 4, $\pi_w(\theta_w) - \pi_v(\theta_v) \leq \alpha(\theta_w - \theta_v)$, with equality if $wv \in E^f(\theta_w)$ and $\theta_w = \theta_v + \tau_{vw}$. Then the cost of a user using this route is

$$
\begin{aligned}
\rho(\theta_t) &+ \sum_{e=vw \in P} [\alpha\tau_e + \delta_e(\theta_v)] \\
&\geq \rho(\theta_t) + \pi_t(\theta_t) - \pi_s(\theta_s) \\
&= \rho(\theta_t) + (C - \rho(\theta_t))^+ \\
&\geq C.
\end{aligned}
$$

The inequalities are all tight if for all $vw \in P$, $\theta_w = \theta_v + \tau_{vw}$ and $f_{vw}(\theta_v) > 0$, by the previous observations as well as property (iv). So if the aggregate choices of the users are described by f, all users pay exactly C.

As already discussed, we cannot in general strongly enforce an optimal flow. The following shows that the "lane tolling" approach suffices to do this.

Theorem 11. *With f, π and δ as in the previous theorem, any dynamic equilibrium g satisfying $g_e(\theta) \leq f_e(\theta)$ for all $e \in E$, $\theta \in \mathbb{R}$ is $g = f$.*

Essentially, being able to dynamically split and separately toll the capacity of a link allows us to easily rule out all other potential equilibria just by using tolls to artificially constrict the capacities (in addition to choosing tolls that weakly enforce the desired flow, which is still needed). Tolling

in this way seems quite distant from what could be imaginable in realistic traffic scenarios. But it does raise the interesting question of whether there is a tolling scheme which can strongly enforce an optimum flow, but which is more restricted (and more plausible) than fully dynamic lane tolling. Another natural question would be to determine if an optimum flow can be strongly enforced using lane tolling only on certain specified edges. We leave these as open questions.

6 General scheduling costs

We now briefly discuss general scheduling costs, satisfying only the growth assumption as well as minor technical conditions. We will not give full details, but just highlight what changes need to be made in the algorithm and analysis.

First, suppose that ρ was unimodal, but not strongly unimodal. This introduces the complication that we cannot uniquely invert ρ on $[0, \infty)$ and $(-\infty, 0]$; and also that ρ^{-1} is discontinuous. If we use the prescription of the interval $[a_j, b_j]$ given in Section 3, namely the maximal interval for which

$$\rho(\xi + d_{j-1}(s,t)) \leq C - \alpha d_{j-1}(s,t) \quad \text{for all} \quad \xi \in [a_j, b_j], \tag{12}$$

we run into the difficulty that Q is no longer a continuous function of C. Thus we may find a choice C_0 so that the resulting amount of flow is strictly less than Q, but the amount of flow corresponding to $C_0 + \epsilon$ is strictly larger than Q, for any positive ϵ.

To remedy this, we can proceed as follows after determining C_0. Let $[a_j^0, b_j^0]$ be the minimal interval so that

$$\rho(\xi + d_{j-1}(s,t)) \geq C_0 - \alpha d_{j-1}(s,t) \quad \text{for all} \quad \xi \in (-\infty, a_j^0] \cup [b_j^0, \infty).$$

Define $[a_j^1, b_j^1]$ as per the previous definition (12). Then for $\delta \in (0, 1)$, define $a_j^\delta = (1 - \delta)a_j^0 + \delta a_j^1$, and similarly for b_j^δ. It is not hard to see that for any choice of δ, using the intervals $[a_j^\delta, b_j^\delta]$ in the original algorithm provides a flow over time with cost horizon C_0; and moreover that the value of the flow depends continuously on δ. The proof of optimality is essentially unaffected. So a further round of bisection or parametric search suffices here.

Next, suppose we drop the requirement that ρ is unimodal (maintaining still the growth assumption). Instead, let us make the mild assumption that for any $z \in \mathbb{R}_+$, $\rho^{-1}([0, z])$ is a union of finitely many compact intervals. Little changes, except that instead of obtaining a single interval $[a_j, b_j]$ for each path P_j, we obtain a finite number of intervals. Theorem 1 remains true as stated.

References

[1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows: Theory, Algorithms, and Applications.* Prentice Hall, 1993.

[2] R. Arnott, A. de Palma, and R. Lindsey. Economics of a bottleneck. *Journal of Urban Economics*, 27(1):111–130, 1990.

[3] N. Baumann and M. Skutella. Solving evacuation problems efficiently–earliest arrival flows with multiple sources. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 399–410, 2006.

[4] U. Bhaskar, L. Fleischer, and E. Anshelevich. A Stackelberg strategy for routing flow over time. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA)*, pages 192–201, 2011.

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

[6] R. Cominetti, J. R. Correa, and N. Olver. Long term behavior of dynamic equilibria in fluid queuing networks. In *Proceedings of the 19th International Conference on Integer Programming and Combinatorial Optimization (IPCO)*, pages 161–172, 2017.

[7] J. R. 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, 2019.

[8] Y. Disser and M. Skutella. The simplex algorithm is NP-mighty. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA)*, pages 858–872, 2015.

[9] L. Fleischer and E. Tardos. Efficient continuous-time dynamic network flow algorithms. *Operations Research Letters*, 23(3):71–80, 1998.

[10] L. R. Ford and D. R. Fulkerson. Constructing maximal dynamic flows from static flows. *Operations Research*, 6(3):419–433, 1958.

[11] L. R. Ford and D. R. Fulkerson. *Flows in Networks*. Princeton University Press, 1962.

[12] D. Gale. Transient flows in networks. *The Michigan Mathematical Journal*, 6(1):59–63, 1959.

[13] R. Grinold. Infinite horizon programs. *Management Science*, 18:157–170, 1971.

[14] T. Harks. Pricing in resource allocation games based on duality gaps. Preprint, arXiv:1907.01976, 2019.

[15] J. J. Jarvis and H. D. Ratliff. Some equivalent objectives for dynamic network flow problems. *Management Science*, 28(1):106–109, 1982.

[16] 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.

[17] E. Köhler, R. H. Möhring, and M. Skutella. Traffic networks and flows over time. In *Algorithmics of Large and Complex Networks - Design, Analysis, and Simulation*, pages 166–196, 2009.

[18] N. Megiddo. Combinatorial optimization with rational objective functions. In *Proceedings of the 10th Annual ACM Symposium on Theory of Computing, (STOC)*, pages 1–12, 1978.

[19] E. Minieka. Maximal, lexicographic, and dynamic network flows. *Operations Research*, 21(2):517–527, 1973.

[20] N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[21] H. E. Romeijn, R. L. Smith, and J. C. Bean. Duality in infinite dimensional linear programming. *Mathematical Programming*, 53:79–97, 1992.

[22] L. Sering and L. V. Koch. Nash flows over time with spillback. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA)*, pages 935–945, 2019.

[23] 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)*, pages 12:1–12:20, 2018.

[24] M. Skutella. An introduction to network flows over time. In *Research Trends in Combinatorial Optimization*, pages 451–482, 2009.

[25] K. A. Small. The bottleneck model: An assessment and interpretation. *Economics of Transportation*, 4(1):110–117, 2015.

[26] W. Vickrey. Congestion theory and transport investment. *American Economic Review*, 59(2):251–60, 1969.

[27] W. L. Wilkinson. An algorithm for universal maximal dynamic flows in a network. *Operations Research*, 19(7):1602–1612, 1971.

[28] H. Yang and Q. Meng. Departure time, route choice and congestion toll in a queuing network with elastic demand. *Transportation Research Part B: Methodological*, 32(4):247–260, 1998.

[29] N. Zadeh. A bad network problem for the simplex method and other minimum cost flow algorithms. *Mathematical Programming*, 5:255–266, 1973.