

Equilibria in Dynamic Selfish Routing

Elliot Anshelevich* Satish Ukkusuri†

May 2009

Abstract

In both transportation and communication networks we are faced with “selfish flows”, where every agent sending flow over the network desires to get it to its destination as soon as possible. Such flows have been well studied in time-invariant networks in the last few years. A key observation that must be taken into account in defining and studying selfish flow, however, is that a flow can take a non-negligible amount of time to travel across the network from the source to destination, and that network states like traffic load and congestion can vary during this period. Such flows are called dynamic flows (a.k.a. flows over time). This variation in network state as the flow progresses through the network results in the fundamentally different and significantly more complex nature of dynamic flow equilibria, as compared to those defined in static network settings.

In this paper, we study equilibria in dynamic flows, and prove various bounds about their quality, as well as give algorithms on how to compute them. In general, we show that unlike in static flows, Nash equilibria may not exist, and the price of anarchy can be extremely high. If the system obeys FIFO (first-in first-out), however, we show the existence and how to compute an equilibrium for all single-source single-sink networks. In addition, we prove a set of much stronger results about price of anarchy and stability in the case where the delay on an edge is flow-independent.

1 Introduction

The concepts of “selfish flow” and “selfish routing” in time-invariant networks have been thoroughly explored over the last few years (see, for example, [17, 18, 21]). These models apply to networks that involve routing by a large number of independent self-interested agents, such as transportation networks, Internet routing, and peer-to-peer file sharing systems. In all these systems, individual agents using the network (vehicles on highways, Internet peer-to-peer clients, etc.) can be expected to be somewhat “selfish”, and may only be interested in optimizing their own performance metric when making routing decisions. Understanding equilibria in such networks is crucial in order to measure system efficiency and performance, and to design suitable mechanisms that improve the properties of the system.

Unlike in the usual selfish routing models, however, in many networked systems the state is a function of both the number of agents and *the time when the agents use the network*. In addition, a flow (or equivalently, the agents/entities that constitute the flow) can *take a non-negligible amount of time to travel across the network from the source to destination, and network state like traffic load and congestion can vary during this period*. For instance, in hurricane evacuations, the number of people evacuating in the transportation network is a time-dependent process which can vary considerably during the time it takes to move from homes to safe shelters. In the context of the Internet, network congestion levels (and thereby, download speeds) can vary significantly within a few hours, which is the typical time-scale for downloads of large video files in peer-to-peer systems.

We study routing done by self-interested agents in the context of such flows, which we call “dynamic flow” (or equivalently, “flow over time”). The behavior of dynamic flows is very different from static (time-invariant) flows, as was seen before in numerous papers studying dynamic flows from a centralized

*Dept. of Computer Science, Rensselaer Polytechnic Institute, Troy, NY 12180. eanshel@cs.rpi.edu.

†Dept. of Civil and Environmental Engineering, Rensselaer Polytechnic Institute, Troy, NY 12180. ukkuss@rpi.edu

point of view [9, 11, 13] (i.e., without self-interested agents). To further illustrate the difference between static and dynamic flow, consider that in static selfish flow, the congestion of a link is a function of all the users that use this link, aggregated over time. This is problematic, however, in time varying flows. In our model, when carefully choosing their best route, user i does not consider the congestion on link e based on the total number of users that traverse e , or even on the users that are currently using e , but instead considers the congestion on e that will take place when i reaches e . For a detailed description of our model, see Section 2.

In this paper we study equilibria in dynamic flows, and prove various bounds about their quality, as well as give algorithms on how to compute them. We concentrate on Nash equilibria in dynamic flow models, and show that their properties can be quite different from the ones in static models. We call such equilibria “dynamic equilibria” to differentiate them from Nash equilibria in static selfish flow. The notion of dynamic Nash equilibrium in this context refers to a system state where each agent is not better off by deviating from its chosen solution (which can consist of the agent’s route, start and waiting times, etc.). As this paper illustrates, the variation in network state as the flow progresses through the network results in the fundamentally different and significantly more complex nature of dynamic flow equilibria problems, as compared to those defined in static network settings.

Background and Related Work While both dynamic flows (a.k.a. flows over time) and selfish flows have been studied extensively, few outside the transportation community [14] have yet attempted to study flows that are *selfish* and *take time to traverse a link*. Below we outline some of the related work on these topics.

Static Selfish Flows: Traditional computer science research concerning routing of self-interested agents in networks has mostly focused on *static* flows [3, 17]. In a static routing game, we are given a graph $G = (V, E)$ of links, with latency functions $\ell_e(x)$ for each edge e (these functions are usually assumed to be non-decreasing and convex). We are also given a source s and sink t (or possibly many sources and sinks), along with some amount of flow that desires to be routed from s to t . This flow is sometimes considered to be *non-atomic* (composed of an infinite number of users, each controlling an infinitesimal amount of flow) or *atomic* (composed of a finite number of users, each controlling a discrete amount of flow). In this paper, we focus on non-atomic flow, although many of our results can be adopted to the atomic case. Each user is able to choose the route along which its flow proceeds, and since the users are self-interested, they will choose the route that has minimum congestion. Specifically, if x_e is the total amount of flow on edge e , then the congestion or delay on path P is defined to be $\sum_{e \in P} \ell_e(x_e)$. When considering the solution quality for such flows, the most common measure is the total congestion (which also corresponds to the social welfare), measured by $\sum_{e \in E} x_e \ell_e(x_e)$. The centralized optimal solution is considered to be the flow that minimizes this value (and therefore maximizes social welfare). Many properties are known about such static selfish flows, such as results about the prices of anarchy [4, 17] and stability [5], and the fact that a Nash equilibrium exists and (for certain classes of games) is unique.

Flows over time: The single major difference between the unsolved problems we address and most traditional network flow research is the notion of time passing while flow (i.e., cars/packets) move through the network. This type of flow is often called “flow over time” or “dynamic flow”, and is applicable not only to all kinds of transportation networks, but also a wide range of data communication networks. Specifically, a flow over time means that flow on different links in the network can change over time, and more importantly, a flow requires a certain amount of *transit time* to travel through each link, this amount possibly dependent on the current congestion. Unlike in static flow systems, the amount of flow or congestion on a link changes at every time step, as some flow enters and some flow leaves this link, adding an extra temporal component. There has been some excellent research in flows over time using various optimization techniques (for surveys, see [1, 11, 15, 16]). The work in [8, 12] is especially relevant to ours, since these papers consider traffic delays that are flow-dependent, as we do in Sections 3 and 4. All of this research, however, does not take into account the strategic nature of agents in the system. One of the main goals of this paper is to combine the techniques used for analyzing dynamic flows with the ones used for static selfish flows [6].

Dynamic Equilibrium in Transportation Networks: Due to the inherent nature of time-varying flow in transportation networks, dynamic flow (referred to as *Dynamic Traffic Assignment*) models are heavily used in transportation network planning, operations, and evaluating real-time systems. These are typically classified into two categories: analytical models based on mathematical programming formulations, and simulation-based heuristic models. Extensive work has been performed for both types of approaches and an overview of this literature can be found in [14, 20]. A key limitation of the previously developed models is that even under simplified assumptions, the models lack rigorous theoretical results on the existence, uniqueness, and algorithms to compute dynamic flow equilibria. While simulation models allow us to compute a “solution”, it is difficult to guarantee that it is indeed a dynamic equilibrium. Simply put, our work provides a much needed theoretical foundation to the dynamic strategic flow concept, while building on the work done by transportation researchers. Portions of the results in this paper were presented at DTA 2008 (International Symposium on Dynamic Traffic Assignment).

2 Model and Our Results

In this section we describe our general model of strategic flow over time. In our model we are given a road or computer network, represented by a directed graph $G = (V, E)$. We also have some flow demands, with source-sink pairs $\{s_i, t_i\}$ such that a unit of flow desires to move from s_i to t_i . This flow is selfish and non-atomic, so a player in this case corresponds to an infinitesimal amount of flow, and the goal of every player is simply to reach its destination in the least amount of time possible. The strategies of the players consist of picking a path from their source to their destination.

Everything we defined so far is exactly the same as in the usual selfish routing model. The main difference is the congestion function. Every link has a function $d_e(x, H_e^t)$, which determines how long it takes x units of flow that enter edge e at time t to traverse this edge. This amount can depend not just on the amount of flow x , but also on the “history” of the edge usage. Specifically, define H_e^t to be the set of amounts of flow using e before time t , and how long ago it entered e before time t . For example, if 2 units of flow entered e one time step before t , and $1/3$ units of flow entered e two timesteps before t , then $H_e^t = \{(2, 1), (1/3, 2)\}$. This is an extremely general way to model flow over time. The fact that the delay d_e depends on the history vector H_e means that, for example, the delay for x flow entering at time t can depend on the total flow that is currently on the edge, or on the time until all previous flow leaves the edge, etc. We assume that the functions d_e are monotone increasing, i.e., that increasing the amount of flow in H_e or in x can only increase delays. Given these congestion functions and the behavior of other users, each user chooses a path from s_i to t_i that would minimize its delay.

What makes our model drastically different from static selfish flow is that it takes *time* for a user to traverse a link. As we show in Section 4, the properties of our model are also very different from the static flow model. For example, equilibria may not exist in dynamic flows, and unlike for static non-atomic flows, they may not be unique. This results partially from the fact that unlike static flows, self-interested dynamic flows do not form a congestion game [3], and modeling dynamic flow games simply as repeated static flow games can result in great inaccuracies.

Flow-Dependent Delays obeying FIFO We are especially interested in dynamic flows where the FIFO (first-in first-out) property is obeyed. This is certainly true for most communication networks, where packets are forwarded according to their arrival order, and is a common assumption for transportation networks as well, since FIFO is largely enforced in time of high congestion, and is enforced on average always. Because of this, in Section 3 we consider the general model described above, but with the assumption that for all e ,

$$d_e(x, H_e^t) + t \geq d_e(x', H_e^{t'}) + t' \tag{1}$$

where x enters e at time t , and x' enters e at time $t' < t$ (so that $(x', t - t') \in H_e^t$). This condition simply states that if flow x' enters an edge before x , then it also leaves before x . For a discussion of some concrete models that fit into this framework, see Section 3.1.

In Section 4 we show that for delay functions satisfying FIFO, it is still possible that no equilibria exist, and that the price of anarchy is unbounded. On the positive side, however, we show in Section 3

that for single-source single-sink networks obeying FIFO, a Nash equilibrium always exists, and can be computed efficiently.

Flow-Independent Delays In Section 5 we focus on a very special case of flow independent delay functions. In this model, the travel time on a link $e \in E$ does not depend on the amount of flow on that link, so the delay d_e is constant. Concurrently with our work, [10] showed many interesting properties of a continuous version of the model with flow-independent delays. While the flow-dependent delay model above has extremely different properties from static flow, in the case of flow-independent delays there is a close relationship between static and dynamic flows [13].

We exploit this relationship to give price of stability bounds for dynamic flows with flow-independent delays. Specifically, we show that there always exists a Nash equilibrium that is as good as the centralized optimum. We also consider the case when the players’ cost functions do not simply correspond to travel time. For example, a player’s cost might be a function of both travel time and the congestion on the links taken. We show that if these cost functions are linear, then we can compute an equilibrium that is at most $4/3$ more expensive than the centralized optimum (and in general it will be α more expensive, where α is a factor dependent on the class of cost functions in [17, 18]).

Model Extensions and Waiting Times In Section 6, we mention several extensions of our model including the addition of “waiting” at intermediate nodes (referred to as “intermediate storage” in [13]).

3 Flow-Dependent Delays Obeying First-In First-Out (FIFO)

An important goal in modeling dynamic flows is to understand the desirable properties of a good dynamic equilibrium model. In this section we assume that the delays on the links satisfy the first-in first-out (FIFO) property. The FIFO property can be defined as that any unit of flow A entering a link e before some flow unit B also exits the link before B ; in other words overtaking is not allowed¹. It is easy to show that simple inflow [8] or exit flow models violate the FIFO condition. An inflow model is a model where the travel time for x flow entering a link at time t depends only on the size of flow x . In other words, the delay function is of the form $d_e(x)$, and does not rely on H_e . To see that inflow models violate FIFO, consider the case when there is a sharp increase in traffic flow at time t followed by a rapid decrease in flow at time $t + 1$. Then FIFO will be violated since the time taken for a flow entering at time t is much higher than the flow entering at time $t + 1$. Several authors in the transportation community have tried to overcome this limitation by assuming that the number of vehicles on a link at time t is a function of both inflow and exit flow [2, 7]. While these models work well for most practical traffic flows, FIFO can still be violated in such models.

In this section we assume that all delay functions satisfy the FIFO property, i.e., that Inequality 1 always holds. There is a range of models which satisfy this condition, including point queue models and other models that provide approximate “positions” (based on microscopic traffic simulations) of the flow currently on the link (i.e., how far along the link it has traveled so far) or/and a function of speed of the flow. We describe some of these models in Section 3.1.

Even if the FIFO property holds, dynamic equilibria can be extremely different from static equilibria. As we show in Section 4, in multi-source multi-sink networks, equilibria may not exist, and the price of anarchy can be unbounded. This is in contrast with static flows, where an (essentially unique) equilibrium always exists, and can be computed efficiently [17].

For single-source single-sink networks, however, dynamic flows behave much better in the presence of FIFO. Below we assume that there is a single unit of flow starting from a node s and with destination t . If more than a single unit of flow is present, we can always scale the flow so that there is only a single unit: the important thing is that flow only leaves s at one moment in time, since otherwise this essentially becomes a multi-source problem. For the single-source single-sink case, we prove that there exists a Nash equilibrium for this general class of dynamic models, and provide an algorithm to compute the equilibrium

¹For a precise definition, see Inequality 1.

solution. To the best of our knowledge, this is the first theoretical result on the existence/computation of Nash equilibria for this general class of models. Moreover, in Proposition 3.2, we prove that if *strict* FIFO is satisfied (Inequality 1 is made strict), then the equilibrium is unique.

Theorem 3.1 *If delay functions obey FIFO, when a single (splittable) unit of flow desires to get from a source s to a sink t , a Nash equilibrium always exists and can be computed efficiently.*

PROOF: We will form an equilibrium solution as follows. Consider x units of flow to be the first flow to enter link e . How long would it take this flow to traverse e ? This time could be arbitrary, but it will be a function of only the current flow on the link, x , since there is no flow ahead of it. Thus, we let $g_e(x)$ be the delay that x amount of traffic would encounter on link e if there were no one else ahead of it on e . In other words, $g_e(x) = d_e(x, \emptyset)$.

Compute a static equilibrium flow $\{f_e\}$ on graph G with latency functions $g_e(x)$. To form a dynamic equilibrium, we now transform $\{f_e\}$ into a dynamic flow as follows. If some amount of flow x takes the path P in the static flow, send the same amount of flow on the same path in the dynamic flow. We claim that this dynamic flow F is a (dynamic) Nash equilibrium.

In general, in a dynamic flow it is possible that different units of flow reach a node at different times. We will first show that this does not happen in the dynamic flow F . That is, we will show that for every node v , all the flow of F that reaches v arrives at it simultaneously. In the static equilibrium $\{f_e\}$, the latency on any path with positive flow from s to any node v is the same (since otherwise there would be incentive for flow to switch its path). Call this delay d_v . Let G_F be the subgraph of G only containing edges with positive flow f_e . Notice that we can form an ordering of nodes in nondecreasing order of d_v that is also a topological sort of G_F . This is because if there were flow through a node v to node w , then it must be that $d_w \geq d_v$. We do not have to worry about cycles of delay 0, as we can assume that G_F is acyclic (otherwise we would be able to remove the cycle, and still have a static equilibrium). Without loss of generality, rename the nodes so that this order is $s = v_1, v_2, \dots, v_n = t$, and set $d_i = d_{v_i}$. We now prove inductively on this ordering that all flow in the dynamic solution F that arrives at a node v arrives there simultaneously at time d_v .

The base case is trivial, since all the flow arrives at node s at time $d_s = 0$. Now, consider a node v_i , and suppose that for all nodes v_j with $j < i$, we have that the only time when flow in the dynamic solution F arrives at v_j is at time d_j . Now consider a link (v_j, v_i) with flow on it. In the dynamic flow F , there is only one time that flow uses this link, since there is only one time when flow reaches node v_j , and the flow of F enters its next edge as soon as it arrives at node v_j . This means that all the flow enters this link at the same time d_j , so the flow that traverses $e = (v_j, v_i)$ experiences a delay of exactly $g_e(f_e)$. By our definition of d_i , we know that $d_i = d_j + g_e(f_e)$. Therefore, all of the flow of F that arrives at v_i arrives there at the same time, specifically at time d_i .

We have shown that in the dynamic flow F , all the flow arriving at a node arrives there simultaneously. In particular, this is true for the destination node t , so all of the flow paths have the same delay (call it d_{\max}), and no one would want to switch to a different flow path of F . To show that F is a dynamic Nash equilibrium, all that is left to show is that there does not exist a path P that is not used by the flow F , with a delay $d(P)$ smaller than d_{\max} . Suppose to the contrary that such a path P exists. Let v be the first node on this path that can be reached in time strictly less than d_v on this path P (we know such a node exists, since t is one such node). Let w be the previous node before v in P that is in G_F (that is, that has some flow of F going through it). We know such a node exists since s is a such a node. By our choice of v , we know that the time it takes to reach node w along path P in the presence of the dynamic flow F is at least d_w . Consider the subpath $P(w, v)$ of P from w to v .

If $P(w, v)$ does not include any edges of G_F , consider the latency along $P(w, v)$ in the static equilibrium $\{f_e\}$. This latency is $\sum_{e \in P(w, v)} g_e(0)$, since there is no flow traversing the edges of $P(w, v)$. Since $g_e(x)$ is defined as the delay for x amount of flow to traverse edge e if not affected by any other flow, this means that $g_e(0)$ is the smallest possible time that it would take to traverse edge e under any circumstances in our dynamic model (assuming that the presence of other flow only increases the travel time). Because of this, we know that the delay for traversing $P(w, v)$ in our dynamic model would be at least the latency of $P(w, v)$ in the static equilibrium $\{f_e\}$. But the delay of any path to v in the static equilibrium is at

least d_v . Therefore, the time necessary to get to v along path P in the presence of the dynamic flow F is at least $d_w + \sum_{e \in P(w,v)} g_e(0) \geq d_v$. This gives us a contradiction, since we chose v as a node that can be reached earlier than d_v .

The path $P(w, v)$ cannot contain more than one edge of G_F by our choice of w , but it might consist of a single edge (w, v) of G_F . By the FIFO property, any flow that enters (w, v) after time d_w must reach node v after d_v , since we have flow on the edge (w, v) leaving at d_w and reaching v at time d_v . Therefore, the time necessary to get to v along path P in the presence of the dynamic flow F is once again at least d_v . This once again yields a contradiction, and so we have shown that the needed time to get to any node v is at least d_v . In particular, for all paths P from s to t , the delay $d(P)$ on path P is at least d_{\max} , and so F is a dynamic Nash equilibrium. ■

The proof of this theorem essentially shows that equilibria in our model correspond exactly to static equilibria in the appropriate static network. Since static equilibria are efficiently computable (especially for linear congestion functions), this theorem provides us with an algorithm for computing dynamic equilibria as well. In addition to the existence of Nash equilibria, we can also show that dynamic equilibria are unique in the presence of strict FIFO.

Proposition 3.2 *With the same assumptions as in Theorem 3.1, if strict FIFO is satisfied (i.e., Inequality 1 is satisfied without equality), then the dynamic Nash equilibrium is unique.*

PROOF: We want to show that if the model obeys *strict* FIFO, then the dynamic equilibrium corresponding to the static one in the proof of Theorem 3.1 is the only existing Nash equilibrium. To see this, consider a dynamic equilibrium in which flow reaches some node v at two distinct times τ_1 and τ_2 , with $\tau_1 < \tau_2$. Let P_i be the flow path from v to t of the flow reaching v at time τ_i , for $i = 1, 2$. Since this is an equilibrium for nonatomic flow, then all flow must reach the destination at the same time, so the delay on P_1 is strictly greater than the delay on P_2 . In this case, every player representing an infinitesimal amount of the flow arriving at time τ_1 would have incentive to switch its path to P_2 . Since it would leave on the same path strictly before the flow arriving at τ_2 , it would also arrive strictly before it at the destination. Therefore, this solution could not be an equilibrium. Since all flow arriving at any node v arrives there simultaneously, we have that no edge is ever used by a dynamic flow equilibrium at two different times. Let f_e be the amount of flow traversing edge e in the dynamic equilibrium. We claim that $\{f_e\}$ is a static equilibrium under latency functions $g_e(x)$ (as defined in the proof of Theorem 3.1).

To prove this, let d'_v be the time that flow reaches node v in this dynamic equilibrium. Let G_F be as above, and let $s = v_1, v_2, \dots, v_n = t$ be a topological sort of the nodes in G_F , ordered in nondecreasing order of d'_v . As before, this is possible since a cycle in G_F would have a time delay of 0, and we could simply remove such a cycle and still have a dynamic equilibrium. Using similar arguments to the ones in the proof of Theorem 3.1, we can show that the latency of all flow paths of $\{f_e\}$ to v in the static problem is exactly d'_v , and the latency of paths not carrying flow is at least d'_v . In particular, this is true for the destination node t , and so $\{f_e\}$ is a static equilibrium. Notice that the above argument about uniqueness only holds if the delay functions are continuous, as otherwise we could have dynamic equilibria where not all flow reaches the sink at the same time. By the uniqueness of static equilibria in uncapacitated graphs [17], we conclude that the dynamic equilibrium is unique as well. ■

More accurately, the statement of uniqueness for static equilibria does not mean that the actual flows are unique, but that the cost of the equilibrium is unique. In the case of dynamic flows, this means that for all equilibria, there is a unique time when all flow must arrive at the destination.

3.1 Concrete Models Satisfying FIFO

There are many realistic flow models that satisfy the FIFO property. Unfortunately, the delay functions d_e for most of the models cannot be written in closed form and do not possess a general analytical solution. However, these models can be solved using numerical methods using various discrete time forward or backward differencing methods for solving first-order differential equations. To illustrate the types of functions d_e that obey FIFO, we present a specific model below. For another, simpler set of functions, see the examples in Section 4.

Variations of the Point Queue Model Consider the following form of what is known in the transportation community as the point queue model. Suppose that a flow x enters the tail of a link e at time t , and assume that if there were no flow on the link at time t , then x will traverse the link in a fixed travel time which can be a constant c_e or, to more accurately account for congestion effects, an expression such as $c_e x$. Now suppose instead that some flow y were already present on the link at time t . In this model, we will think of flow x as traveling on edge e with constant velocity V_x , picked so that it never passes any flow that is in front of it, which is also traveling at constant velocity. Specifically, let $\tau(y)$ be the time at which the last unit of the y flow exits the link e . If y entered e at time t' , then it is traveling at velocity $V_y = 1/(\tau(y) - t')$. Since it is now time t , the flow y has traversed $\frac{t-t'}{\tau(y)-t'}$ of the edge. If we now set $V_x = 1/(c_e x + \tau(y))$, then the flow x will never overtake any flow in front of it, guaranteeing the FIFO property. To see this, notice that

$$\frac{\tau - t'}{\tau(y) - t'} > \frac{\tau - t}{c_e x + \tau(y)}$$

for all $\tau \geq t$. This is true because $\tau c_e x + \tau t' + t \tau(y) > t' c_e x + t t' + t' \tau(y)$, which holds because $\tau \geq t > t'$. Since $\frac{\tau - t}{c_e x + \tau(y)}$ is exactly the position of the x flow at time τ if it is traveling at velocity V_x , then the flow x always lags behind flow y . The travel time for flow x in this case will be $d_e(x, H_e^t) = c_e x + \tau(y)$.

In fact, any travel times that result in the x flow exiting the edge after $\tau(y)$ would satisfy FIFO. For example, we could instead say that x exits the edge at time $\max(c_e x + t, \tau(y))$. This would correspond to the x flow traveling along the link at a top speed of $1/c_e x$, but being unable to pass any flow that is in front of it. In general, any model where flow x has a top speed for traveling on an edge, with its actual speed dependent on the presence or on the density of the flow ahead of x , fits into our framework. FIFO is satisfied in such a model as long as the speed constraints are such that flow never passes anyone ahead of it.

4 Lower Bounds and Examples with No Equilibrium

In the previous section we saw that for the single-source single-sink case of dynamic flow satisfying FIFO, a unique Nash equilibrium exists, and can be computed efficiently. Unfortunately, as we show below, this equilibrium can be much worse than the optimal solution, leading to large prices of anarchy and stability (which are the same in this case, since the equilibrium is unique). We also give an example with multiple sources where an equilibrium does not exist at all, and an example with multiple equilibria.

Example with unbounded Price of Anarchy In the proof of Theorem 3.1, we established that the equilibria of our dynamic model correspond to static equilibria in the same graph G with cost functions $g_e(x)$. In this case $g_e(x)$ is the time it would take x amount of flow to traverse edge e if there were no flow ahead of it on e . The price of anarchy in our model can be very different from the price of anarchy for static flows, however, since while the cost of the equilibrium remains the same, the optimal solution in the dynamic model can be very different from the optimal solution for static flow. This leads to the cost of the optimal solution being very small compared to the cost of the equilibrium.




Figure 1: An example with unbounded Price of Anarchy. The edges are labeled with values $d_e(x, \emptyset)$.

To illustrate this, consider the example in Figure 1, where we have k parallel paths leading from s to

v , followed by a path of 100 edges leading from v to t . Suppose the parallel paths have constant delay, with the i 'th path having delay $i - 1$. That is, if e is the i 'th parallel edge, then set $d_e(x, H_e^t) = i - 1$. The 100 edges in the bottleneck path each have a delay such that if x amount of flow enters an edge, with no other flow currently on this edge, then it will take x time units for this flow to leave the edge. Any such delay functions that obey FIFO will give us the desired example. To be fully concrete, however, we will give a specific example of d_e , defined recursively. First, we set $d_e(x, \emptyset) = x$. In general, if the last amount of flow that entered e before t was an amount x' at time t' , then the latest element of H_e^t is $(x', t - t')$. In such a case, we define $d_e(x, H_e^t) = \max\{d_e(x', H_e^{t'}) + t' - t, 0\} + x$. These functions satisfy the desired condition, since if there is no other flow on edge e at time t , then $d_e(x', H_e^{t'}) + t' \leq t$, and so $d_e(x, H_e^t) = x$. It is also easy to show that these functions satisfy FIFO.

Proposition 4.1 *The functions above satisfy FIFO, and the price of anarchy in this example is $\Omega(n)$, with n being the number of nodes.*

PROOF: To show that these functions satisfy FIFO, we proceed by induction. If we assume that all flow before time t' obeyed FIFO, then we only need to prove that the time when x flow leaves e is at least the time when x' flow leaves e , with x and x' defined as above. In other words, we need to show that $t + \max\{d_e(x', H_e^{t'}) + t' - t, 0\} + x > d_e(x', H_e^{t'}) + t'$. In the case that $d_e(x', H_e^{t'}) + t' - t < 0$, this reduces to $t + x > d_e(x', H_e^{t'}) + t'$, which is true since $x > 0$. In the other case where $d_e(x', H_e^{t'}) + t' - t \geq 0$, this reduces to $d_e(x', H_e^{t'}) + t' + x > d_e(x', H_e^{t'}) + t'$, which is once again true since $x > 0$.

Now consider the solutions for the example in Figure 1 when k units of flow must get from s to t . First consider the solution where all the flow takes the path of delay 0 to v , and then travels together on the bottleneck path. All the flow will get to t at time $100k$, so the total delay in this solution is $100k^2$. This solution is a Nash equilibrium. To see this, consider a possible deviation for some infinitesimal unit of flow. The only possible deviations take a path to v with delay greater than 0, and so this flow arrives at v after all the other flow has started on the bottleneck path. By the FIFO property, the flow that has deviated will arrive at t only after everyone else has, and so this is not an improving deviation. Therefore, this solution is a Nash equilibrium. Another way to see this is to notice that with the congestion functions $d_e(x, \emptyset)$ which are pictured in Figure 1, this solution corresponds to the static equilibrium solution, and so it forms a dynamic equilibrium by Theorem 3.1.

On the other hand, the optimum centralized solution can incur a much smaller total delay. Consider the solution where 1 unit of flow takes each of the k parallel paths. The first unit of flow arrives at t at time 100. The second unit of flow is not affected by the first, since by the time it reaches one of the bottleneck edges, the first unit of flow has already left it, and so its delay on each edge is 1. Therefore, the second unit of flow arrives at t at time $1 + 100$. In general, the i 'th unit of flow would arrive at t at time $i + 100$. The total delay is therefore at most $k(k + 1)/2 + 100k$, which is in fact the optimal solution. For a large enough k , this gives us a price of anarchy approaching 200. Instead of using 100, we can use any constant n , and so the price of anarchy can be arbitrarily high in this FIFO model.

Notice that instead of using k units of flow, we can scale down the flow and the delays on the parallel paths, and obtain the same results using only a single unit of flow. In addition, this price of anarchy bound for the example in Figure 1 holds even if our objective function is the maximum delay instead of the total delay. This is because in the optimal solution, the maximum delay is $k + 100$, while in the Nash equilibrium the maximum delay is $100k$. ■

For multi-source, Nash Equilibria may not exist We have shown above that the price of anarchy can be very high in most FIFO models, since the optimal dynamic solution can “stagger” the flow by breaking it up into small pieces, while an equilibrium must keep all the flow together so all of it arrives at the same time. For the multi-source version of our FIFO model, the situation is even worse, since a Nash equilibrium may not exist at all. While below we present a multi-source multi-sink example with no equilibrium, there also exist such examples with only a single sink node (but multiple source nodes). We give one such example in the Appendix.

Consider the example in Figure 2, with functions $d_e(x, \emptyset)$ as shown in the figure. We define the functions $d_e(x, H_e^t)$ as in the previous example, so that FIFO is satisfied, and so that if x units of flow

enters e at time t , and no flow is present on e at that time, then it takes $d_e(x, \emptyset)$ time for this x flow to traverse the link. In the previous example, this value $d_e(x, \emptyset)$ was equal to x , but now it can vary between edges as specified in the figure. For example, for the edge from v to t_2 , if x flow enters this edge, with no other flow currently on this edge, then it will take $100x$ time units for this flow to leave the edge. If an edge is not labeled with a function, then we say that the delay on this link is always 0.




Figure 2: A multi-source multi-sink example with no Nash Equilibrium. The edges are labeled with values $d_e(x, \emptyset)$.

There are three demands d_1, d_2, d_3 with sources (s_i, t_i) for $i = 1, 2, 3$, and with flow sizes of 2, 100, and 10 respectively. All these demands leave their sources at the same time.

Proposition 4.2 *The example pictured in Figure 2 has no pure dynamic Nash equilibrium.*

PROOF: We will first describe the basic idea of why this example has no equilibrium, and then prove it rigorously. Notice that d_1 and d_3 have a choice of exactly two paths in reaching their sinks, and d_2 has only one path from s_2 to t_2 , and so has no choice at all. The basic idea is as follows. If more than 1 unit of flow of d_1 chooses the lower path, then it does not affect d_2 (since it arrives at their shared edges after d_2), and so d_2 proceeds on its path without any interference, and reaches node v at the same time as the flow of d_3 . This greatly increases the travel time on the lower path of d_3 , causing d_3 to choose the upper path. This in turn causes the lower path of d_1 to become slower, causing flow to shift to its upper path. This causes the amount of flow from d_1 using its lower path to become less than 1, which delays the arrival time of d_2 at node v . This means that the lower path of d_3 is fast again, causing d_3 to shift flow from its upper path, which makes the lower path of d_1 cheaper again. This cyclical behavior can continue forever, and in fact we will show that there is no possible equilibrium solution.

We now prove that this network has no Nash equilibrium. Let y_1 be the amount of flow of d_1 that takes the lower path of d_1 , and y_3 the amount of flow of d_3 that takes the lower path of d_3 . We will consider 3 cases: $y_1 = 0$, $0 < y_1 \leq 1$, and $y_1 > 1$.

If $y_1 = 0$, consider the options available to flow of d_1 . The upper path always has a delay of 150, while a small amount of flow (say $\varepsilon < 1/2$) on the lower path would get to all of its edges before any other flow could (i.e., before time 1 to the second edge of its path, and before time 2 to the fourth edge), and so would incur a delay of $22\varepsilon + 1$, which is smaller than 150. Therefore, the flow of d_1 would desire to deviate to its lower path, and so this is not an equilibrium.

If $y_1 > 1$, then d_2 is not affected by this flow since y_1 reaches the edge labeled with delay $x + 1$ after the flow of d_2 entered it. This means that the 100 units of flow of d_2 reach node v at time 102, since no other flow interferes with their travel before node v . Consider now the options available to flow of demand d_3 . The lower path has a delay of $102 + 100(100 + y_3) = 10000 + 102 + 100y_3$, since y_3 reaches v at exactly the same time as the large amount of flow from s_2 . This makes this path undesirable compared to the upper path with delay $2 + 20(10 - y_3) + 1103 \leq 1305$, and so to have a stable solution, it must be that $y_3 = 0$. The delay of the upper path is 1305 because y_1 reaches the edge labeled $20x$ after at least $2y_1 + 1 > 3$ units of time, and so the flow of d_3 is not affected by it. Now let us consider the options available to flow of d_1 . Since $y_3 = 0$ and $y_1 > 1$, then y_1 arrives at the edge labeled $20x$ after the flow of

d_3 , and because of the FIFO property, it can only reach t_1 after time $20(10 - y_3) = 200$. This is longer than the delay of 150 on the upper path of d_1 , and so this cannot be an equilibrium.

If $0 < y_1 \leq 1$, then the y_1 flow is still on the edge labeled with delay $x + 1$ when the flow of d_2 enters it. According to our choice of functions $d_e(\cdot)$, this delays d_2 by some amount, and so it takes strictly more than 102 time to reach node v . Consider now the options available to flow of demand d_3 . The lower path has a delay of only $102 + 100y_3$, since it is no longer affected by d_2 . The upper path has a delay of at least 1103, which is greater than the delay of the lower path even if all 10 units of flow are sent on the lower path. Therefore, for this solution to be stable, it must be that $y_3 = 10$. Now let us consider the options available to flow of d_1 . Since $y_3 = 10$, the delay on the lower path of d_1 is at most $1 + (101 + 1) + 20 = 123$. This occurs in the worst case when $y_1 = 1$ and so enters the edge labeled $x + 1$ at exactly the same time as d_2 . However, this delay is still smaller than the delay of 150 on the upper path of d_1 , and so this cannot be an equilibrium, as flow of d_1 would prefer to switch to its lower path. ■

Equilibrium is Not Unique As we proved in Proposition 3.2, if strict FIFO is obeyed, then there exists only a unique dynamic equilibrium. We now present a simple example where equilibria are not unique when FIFO is violated. In Figure 3 we present a simple example where equilibria are not unique. In this example, the delay functions d_e are of a special form: $d_e(x, H_e^t) = d_e(x, \emptyset)$. The delays $d_e(x, H_e^t) = d_e(x, \emptyset)$ are pictured in the figure. These types of delay functions that only depend on the amount of flow entering the link at time t are known as *inflow* delays [8], and have been shown to violate FIFO.




Figure 3: An example with multiple dynamic equilibria. The edges are labeled with values $d_e(x, \emptyset)$.

In Figure 3, there is a single unit of flow desiring to go from s to t . We now show that there are two dynamic Nash equilibria. Let a be the amount of flow that takes the top path. If $a = 2/3$, then all the flow arrives at the middle node at time $2/3$, and then enters the next edge together. Therefore, all the flow arrives at t after time $2/3 + 1 = 5/3$. This is an equilibrium, since all the flow paths have exactly the same delay, and so there is no incentive for a player to switch. For consider the case where $a = 2/5$. In this case, a flow arrives at the middle node at time $4/5$, while $1 - a$ arrives at the middle node at time $3/5$. Since the delays only depend on the amount of flow entering the edge, the a flow arrives at t after time $4/5 + 2/5 = 6/5$, and $1 - a$ flow arrives at t after time $3/5 + 3/5 = 6/5$. Since all the flow paths are of equal delay, this is once again an equilibrium.

5 Flow-Independent Delays

In this section, we consider an important special case of *Flow-Independent* delay functions. Specifically, in the Flow-Independent model, we assume that each link e has a delay d_e , and any flow that enters this link at time τ leaves this link at time $\tau + d_e$. In other words, the delay function is simply $d_e(x, H_e^t) = d_e$, a constant. This is flow-independent in the sense that the delay d_e does not depend on the amount of flow on the link. If edges are uncapacitated, then the best thing for any flow (whether selfish or not) would be to proceed on the fastest path from the source to the sink. Therefore, we assume that each edge also has a capacity c_e so that at every timestep, at most c_e units of flow are allowed to enter e .

In Section 4, we showed that dynamic flow models can behave very differently from static flow. For the flow independent case, however, there is a powerful relationship between static and dynamic flows.

As several papers including [13] pointed out before, looking at a flow over time in the Flow Independent Model is the same as looking at a static flow in the appropriate time-expanded graph. Specifically, form a new graph G' with a set of nodes V^τ in G' for every timestep τ , and add edges $(v^\tau, w^{\tau+d_e})$ with capacity c_e to G' if and only if there is an edge $e = (v, w)$ in G . For every sink node t in G , we also form a special node t^* in G' , and add edges (t^τ, t^*) for every time τ . It is easy to see that any static flow in G' corresponds exactly to a flow over time in G (even with multiple sources and sinks). In fact, [13] gave a PTAS for finding the centralized optimal solution in this model. Notice that in our model we do not allow “waiting” (a.k.a. “intermediate storage”), where flow stays at a vertex instead of immediately leaving it. See the Appendix for further discussion of this topic.

To understand the price of anarchy in the Flow Independent Model, we first consider the quality of Nash equilibria in the static flow model with capacities. It is easy to see that there can exist many equilibria in the presence of capacities, and in fact the price of anarchy can be unbounded [5], since the worst equilibrium can be much more expensive than the centralized optimum. This has led to the investigation of the *price of stability*. Recall that the price of stability is the ratio between the *best* Nash equilibrium and the centralized optimum. The best equilibrium can naturally be viewed as the optimum solution subject to the constraint that the solution be *stable*, with no agent having an incentive to unilaterally deviate from it once it is implemented. For models with a unique equilibrium, the price of stability coincides with the price of anarchy, but in the case of capacitated networks, [5] demonstrated that these can be dramatically different.

Theorem 5.1 *For Flow-Independent delays, the price of stability is 1. In other words, there exists a Nash equilibrium as good as the centralized optimum.*

PROOF: As described in [13], there is a one-to-one correspondence between dynamic flows in G and static flows in G' . Specifically, let F be a dynamic flow in G . We now show how to construct a corresponding static flow in G' . If according to flow F , a unit of flow enters an edge (v, w) at time τ , then we can say that this unit of flow traverses the edge $(v^\tau, w^{\tau+d_e})$ in G' , and vice versa. If a unit of flow arrives at its destination t at time τ in F , then the corresponding unit of flow reaches node t^τ according to the corresponding static flow F' , and then immediately proceeds to node t^* .

It is easy to see that this transformation forms a one-to-one correspondence. In fact, we will show below that the optimal (minimum total delay) dynamic flow OPT in G corresponds to the optimal static flow in G' , and that a flow F is a dynamic equilibrium if and only if the corresponding flow F' is a static equilibrium. Therefore, comparing dynamic equilibria with OPT in G is the same as comparing static equilibria with OPT in G' . This means that we can use existing results about static flows. G' is a static flow graph with constant latency functions, and so OPT is simply a min-cost flow, where the cost of an edge is its latency. Therefore, OPT is a static Nash equilibrium. To see this, suppose to the contrary that in OPT there exists a path with free capacity from some source s to some sink t in G' that is cheaper than some path from s to t with positive flow. Then, by switching a small amount of flow to this cheaper path, we can form a flow that is cheaper than OPT, giving us a contradiction. Since the static optimal solution in G' is a static equilibrium, then the dynamic optimal solution in G is a dynamic equilibrium, and so the price of stability is 1. This does not imply, however, that the price of anarchy is 1, since in capacitated static networks there can be many equilibria, some much more expensive than OPT.

To finish the proof, we must show the correspondence between the optimal and equilibrium solutions of G and G' . First, we show that the total delay of F equals the total latency of F' according to constant latency functions d_e . That is, we define the latency function of $(v^\tau, w^{\tau+d_e})$ to be d_e , and the latency function of (t^τ, t^*) to be 0. In the static flow F' corresponding to a dynamic flow F , we now send flow from s^0 to t^* if F sends flow from s to t . With these latency functions, the latency of a unit of flow of F' is exactly the same as the time it takes to reach its destination according to F . This implies that the optimal (minimum total delay) dynamic flow in G corresponds to the optimal static flow in G' .

Moreover, a flow F is a dynamic equilibrium if and only if the corresponding flow F' is a static equilibrium. F' is not an equilibrium if and only if there exists a path P from some source s to some sink t with free capacity, such that the latency $d(P)$ of this path is smaller than the latency of a path from s to t with positive flow on it. This is because the existence of such a path P guarantees a possible deviation

for some flow of F' . Since the latency of a path in G' exactly equals to the delay of the corresponding path in G , this is exactly the condition needed for F not to be a dynamic Nash equilibrium, thus finishing the proof. Notice, however, that the correspondence between Nash equilibria only holds because we do not allow waiting at intermediate nodes, i.e., a flow entering a node v must immediately leave this node, and cannot remain at this node. For a single source s , Theorem 5.1 still holds if we allow waiting at the source s . The only change needed to the proof is the addition of edges $(s^\tau, s^{\tau+1})$ to G' with latency 1 for all timesteps τ . ■

Because of the correspondence between static flows in G' and dynamic flows in G , we know that comparing dynamic equilibria with OPT in G is the same as comparing static equilibria with OPT in G' . This tells us that the price of anarchy in the Flow-Independent Model is unbounded, using the results of [5].

Using this intuition, we can also generalize Theorem 5.1 as follows. Suppose that the disutility of a player traveling on path P is not simply its delay $\sum_{e \in P} d_e$ (as we assume in the rest of this paper), but is instead a function of how congested its route is. Specifically, for every edge e of G , suppose that there is some function $\delta_e^t(f_e^t)$ that shows the cost to a player using edge e at time t , with this cost dependent on the amount of traffic f_e^t entering edge e at time t . The total cost to a player using path P is then the sum of these costs, and minimizing the total cost to all the players is the same as maximizing social welfare. In other words, we assume that while it still takes a constant amount of time d_e for everyone to traverse edge e (independent of the amount of traffic f_e^t), the disutility $\delta_e^t(f_e^t)$ of a player using edge e actually increases as the amount of traffic increases, modeling the fact that people do not like to drive in heavy traffic, even if it does not cause them to be delayed. This can also model other modes of transportation (such as subway systems), where the travel time may not change with the number of users, but the utility of a user changes greatly (e.g., because the subway car is crowded, or there is no place left to sit).

These new player utilities do not change the possible flows, so the transformation between G and G' is still valid, as well as the correspondence between dynamic flows in G and static flows in G' . If we set the cost of an edge e^t in G' to be a function $\delta_e^t(f_e^t)$ instead of just a constant d_e , then the cost of a flow F' in G' is still the same as the total cost to all players in the corresponding dynamic flow F of G . This means that instead of thinking about dynamic flows in G , we can now think about static flows in G' , and compare the Nash equilibria in G' with OPT. Using existing results about static flows [19], we know that in capacitated graphs, the price of anarchy can be very high, while the price of stability is at most $\alpha(A)$, a value depending on the class of possible latency functions A . For example, if the functions δ_e^t are linear in the amount of traffic, then $\alpha(A)$ equals $4/3$, and for polynomials of degree d and positive coefficients, $\alpha(A) = d/\log d$. In addition, we know that in graphs without capacities, the price of anarchy is at most $\alpha(A)$ [17], and so all equilibria are good compared to the optimal centrally planned solution. The objective function for OPT in G' is now to minimize $\sum_{e^t} \delta_e^t(f_e^t)$. A Nash equilibrium of G now corresponds exactly to a static Nash equilibrium in G' with congestion functions δ_e^t . Therefore, we can now apply the existing results about Nash equilibria in static flows with capacities. This gives us examples where the price of anarchy can be very high, as well as the following result.

Theorem 5.2 *The prices of anarchy and of stability for the Flow-Independent model with player cost functions δ_e are at most those for the static flow model with cost functions δ_e^t . For example, if δ_e^t are all linear functions, then:*

- *The price of stability is at most $4/3$ for capacitated graphs.*
- *The price of anarchy is at most $4/3$ for uncapacitated graphs.*

If instead of linear, the functions δ_e^t were polynomials with degree at most d , then the factor of $4/3$ above can be replaced with $d/\log d$.

PROOF: Using the same arguments as in the proof of Theorem 5.1, we can establish a correspondence between the optimal solutions and the equilibria of dynamic solutions in G and static solutions in G' . Therefore, all the results about comparing static equilibria to OPT still apply to the dynamic solutions.

The results about the price of anarchy above come directly from the classic results of [18]. To prove the results about the price of stability in capacitated networks, we apply the results from [19], which show that there always exists an equilibrium that is close in cost to optimal. In fact, these same results also give bounds on the prices of anarchy and stability for any class of functions, not just linear and polynomial ones. ■

6 Conclusion and Extensions

Several researchers studying flow over time also considered whether the flow is allowed to “wait” at intermediate nodes or not. Such “waiting” (referred to as “intermediate storage” in [13]) makes no sense in static flows, but can become important with the extra temporal component of dynamic flows. Most of our flow-dependent results easily extend to the model where a player strategy is not only a route from s_i to t_i , but also the amount of time they plan to spend waiting at each intermediate node. This is because we can model waiting times by adding self-loops with small constant delay ϵ to every node. Waiting for some multiple of ϵ would correspond to choosing a route that traverses this self-loop several times. It is important to point out, however, that in such a model FIFO might be violated, as flow that has been waiting at some node v until time τ might enter an edge out of v after flow that only reached v at time τ .

Many other extensions of our models are possible. For example, our flow-independent results still hold if the delays depend on the time t , as long as they do not depend on the flow x .

The present paper provides important avenues for future research. First, we would like to characterize the equilibrium solutions for the multi-source traffic networks and identify good equilibria. Second, one key reason for obtaining the unbounded price of anarchy for the FIFO model is because of the nature of the optimum solution. The optimum solution can be arbitrarily small since by splitting flow one can improve the optimum solution. One can identify ways to influence the equilibrium solutions to be closer to the optimum, by providing taxes or other incentives. This will help us to obtain a smaller price of anarchy with the influenced equilibrium solution.

Acknowledgements

We thank Lisa Fleischer and Neil Olver for illuminating discussions on modeling flows over time.

References

- [1] J.E. Aronson. A survey of dynamic network flows. *Annals of Operations Research*, 20:1-66, 1989.
- [2] M. Carey. Link travel times ii: properties derived from traffic-flow models. *Networks and Spatial Economics*, 4(4):379–402, 2004.
- [3] G. Christodoulou and E. Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. *ACM Conference on Electronic Commerce (ESA)* 2005.
- [4] R. Cominetti, J. R. Correa, and N. Stier-Moses. Network games with atomic players. *Proc. 33rd ICALP*, pages 525 - 536. Springer, 2006.
- [5] J.R. Correa, A.S.Schulz, and N.E.Stier-Moses. Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29(4):961–976, 2004.
- [6] B. Farzad, N. Olver, and A. Vetta. A priority-based model of routing. *Chicago Journal of Theoretical Computer Science*, Article #1, 2008.
- [7] T. Friesz, J. Luque, R.L. Tobin, and B.W. Wie. Dynamic network traffic assignment considered as a continuous time optimal control problem. *Operations Research*, 37:893–901, 1989.

- [8] A. Hall, K. Langkau, and M. Skutella. An FPTAS for Quickest Multicommodity Flows with Inflow-Dependent Transit Times. *Algorithmica* 47(3), 299–321, 2007.
- [9] Bruce Hoppe and Éva Tardos. *Mathematics of Operations Research*, volume 25, issue 1 (february 2000). pages: 36 – 62.
- [10] R. Koch and M. Skutella. Nash Equilibria and the Price of Anarchy for Flows over Time. Matheon Preprint 524, 2008.
- [11] E. Köhler, R. H. Möhring, and M. Skutella. Traffic Networks and Flows Over Time, in Jrg Kramer (ed.): *DFG Research Center: Mathematics for Key Technologies*, Berliner Mathematische Gesellschaft, 49–70, 2002.
- [12] E. Köhler and M. Skutella. Flows over time with load-dependent transit times, *SIAM Journal on Optimization* 15(4), 1185–1202, 2005.
- [13] L.Fleischer and M. Skutella. Minumum cost flows over time without intermediate storage. *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 66–75, 2003.
- [14] S. Peeta and A. Ziliaskopoulos. Foundations of dynamic traffic assignment: The past, the present, and the future. *Networks and Spatial Economics*, vol. 1, no. 3/4, pp. 233–266, 2001.
- [15] W.B. Powell, P. Jaille, and A. Odoni. Stochastic and dynamic networks and routing. In M.O.Ball, T.L.Magnanti, C.L. Monma, and G.L. Nemhause, editors, *Network Routing*, Ch. 3, pp. 141–295, Vol. 8 of Handbooks in Operations Research and Management Science. North-Holland, Amsterdam, 1995.
- [16] B. Ran and D.E. Boyce. *Modeling Dynamic Transportation Networks*. Springer, Berlin, 1996.
- [17] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.
- [18] T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.
- [19] Nicholas Stier-Moses. *Selfish versus Coordinated Routing in Network Games*. PhD Dissertation, Sloan School of Business, MIT, 2004.
- [20] S.T. Waller and S. Ukkusuri. A linear program for the user optimal dynamic traffic assignment problem. *In Review, Transportation Science*, 2003.
- [21] J.G. Wardrop. Some Theoretical Aspects of Road Traffic Research. *Proceedings of the Institute of Civil Engineers, Part II*, volume 1, pp. 325-378, 1952.

Appendix

A More Lower Bounds and Counterexamples

For multiple sources of flow that is trying to reach a single sink, a Nash equilibrium may not exist, because there can always exist dynamic paths with different delays for getting flow to its destination, leading to incentives for switching to the faster dynamic paths. Consider the example in Figure 4. There are two source nodes s_1 and s_2 , with s_1 having 20 units of flow to get to t , and s_2 having 25 units of flow to get to t . This flow appears at s_1 and s_2 at the same time. Suppose also that the delays $d_e(x)$ for the edges are as pictured in Figure 4. For example, for the edge from s_2 to t , if x flow enters this edge, with no other flow currently on this edge, then it will take x time units for this flow to leave the edge. We do not specify what the delays are if there is already some flow on this edge, except that such delays must obey the FIFO property (or we can define the delays in the same manner as in Section 4).




Figure 4: A multi-source single-sink example with no dynamic equilibrium. The edges are labeled with values $d_e(x, \emptyset)$.

We will now show that this network has no Nash equilibrium. The flow from s_1 has only one possible path, and therefore only one possible strategy. The only difference between solutions is how much flow from s_2 goes directly to t , and how much takes the other path. Let y denote the amount of flow from s_2 that does not take the edge directly to t . We will consider 2 cases: $y < 10$, and $y \geq 10$.

If $y < 10$, then this flow reaches the middle node before the flow from s_1 . The flow from s_1 does not affect it y , so it will reach its destination at time $y + y/2$. For this to be an equilibrium, we need that the $25 - y$ units of flow that go directly to t from s_2 reach there at time $y + y/2$, since otherwise there would be incentive for some of the flow to switch paths. Therefore, it must be that $3y/2 = 25 - y$, implying that $y = 10$, giving us a contradiction. Therefore, if $y < 10$, then some small enough amount of flow that is going directly from s_2 to t will always have an incentive to switch its path.

On the other hand, if $y \geq 10$, then this flow reaches the middle node after (or at the same time) as the flow from s_1 . By the FIFO property, this means that this flow must reach t after (or at the same time) as the flow from s_1 . The earliest that the flow from s_1 could reach t is at time 20. Therefore, in order for the two paths from s_2 to have the same delay, it must be that $25 - y \geq 20$, since otherwise there would be incentive for flow to switch to the direct path from s_2 to t . This implies that $y < 5$, a contradiction.

Therefore, no Nash equilibrium exists in Figure 4. To summarize the above argument, if $y < 10$, then it moves very quickly as it is unaffected by the flow from s_1 , and so the lower path of s_2 is desirable. As soon as y becomes 10, however, it is suddenly affected by the flow from s_1 , and so the direct edge to t suddenly becomes more desirable. The non-existence of Nash equilibria in our (and most other) dynamic flow models is caused by such discontinuities, which never occur in static flow models. In fact, the “intuitive” equilibrium solution for the example in Figure 4 would be a “limit” solution where y approaches 10 from below. In Section 4 we give a more complex example where no such solution exists either.