Machine-Learned Prediction Equilibrium for Dynamic Traffic Assignment

Lukas Graf[1], Tobias Harks[1], Kostas Kollias[2], Michael Markl[1]

[1]University of Augsburg
[2]Google
{lukas.graf, tobias.harks}@math.uni-augsburg.de, kostaskollias@google.com, michael.markl@student.uni-augsburg.de

Abstract

We study a dynamic traffic assignment model, where agents base their instantaneous routing decisions on real-time delay predictions. We formulate a mathematically concise model and derive properties of the predictors that ensure a *dynamic prediction equilibrium* exists. We demonstrate the versatility of our framework by showing that it subsumes the well-known full information and instantaneous information models, in addition to admitting further realistic predictors as special cases. We complement our theoretical analysis by an experimental study, in which we systematically compare the induced average travel times of different predictors, including a machine-learning model trained on data gained from previously computed equilibrium flows, both on a synthetic and a real road network.

Introduction

Understanding and optimizing traffic networks is a significant effort that impacts billions of people living in urban areas, with key challenges including managing congestion and carbon emissions. These phenomena are heavily impacted by individual driver routing decisions, which are often influenced by ML-based predictions for the delays of road segments (see, for instance, (Jiang and Luo 2021) for an overview of convolutional and graph neural network based approaches). One key aspect that is not well understood, is that these routing decisions, in turn, influence the forecasting models by changing the underlying signature of traffic flows.

In this paper, we address this interplay focusing on the popular dynamic traffic assignment (DTA) framework, on which there has been substantial work over the past decades (see the classical book of Ford and Fulkerson (Ford and Fulkerson 1962) or the more recent surveys of Friesz et al. (Friesz and Han 2019), Peeta and Ziliaskopoulos (Peeta and Ziliaskopoulos 2001) and Skutella (Skutella 2008)). A fundamental base model describing the dynamic flow propagation process is the so-called *deterministic queuing model* due to Vickrey (Vickrey 1969). Here, a directed graph $G = (V, E)$ is given, where edges $e \in E$ are associated with a queue with positive rate capacity $\nu_e \in \mathbb{R}_{>0}$ and a physical transit time $\tau_e \in \mathbb{R}_{>0}$. If the total inflow into an edge $e = vw \in E$ exceeds the rate capacity ν_e, a queue builds up and agents need to wait in the queue before they are forwarded along the edge. The total travel time along e is thus composed of the waiting time spent in the queue plus the physical transit time τ_e. The Vickrey model is arguably one of the most important traffic models (see Li, Huang and Yang (Li, Huang, and Yang 2020) for an up to date research overview of the past 50 years), and yet, it is mathematically quite challenging to analyze (see Friesz et al. (Han, Friesz, and Yao 2013b) for a discussion of the inherent complexities).

Given a physical flow propagation model, the routing and traffic prediction algorithms are usually subsumed under a *behavioral model* of agents in order to solve a DTA model. The behavior of agents is modelled based on their informational assumption which in turn defines their respective utility function. Most works in the DTA literature on the Vickrey model can roughly be classified into two main informational categories: the *full information model* and the *instantaneous information model*. In the full information model, an agent is able to exactly forecast future travel times along a chosen path effectively anticipating the whole spatio-temporal flow evolution over the network. This assumption has been justified by letting travelers learn good routes over several trips and a dynamic equilibrium then corresponds to an attractor of an underlying learning dynamic. Existence and computation of dynamic equilibria in the full information model has been studied extensively in the transportation science literature, see (Friesz et al. 1993; Han, Friesz, and Yao 2013a,b,c; Meunier and Wagner 2010; Zhu and Marcotte 2000), whereas the works in (Koch and Skutella 2011; Cominetti, Correa, and Larré 2015) allow a direct combinatorial characterization of dynamic equilibria leading to existence and uniqueness results in the realm of the Vickrey bottleneck model. While certainly relevant and key for the entire development of the research in DTA, this concept may not accurately reflect the behavioral changes caused by the wide-spread use of navigation devices and resulting real-time decisions by agents.

In the instantaneous route choice model, agents are informed in real-time about the current traffic situations and, if beneficial, reroute instantaneously no matter how good or bad that route was in hindsight, see Ran and Boyce (Ran and Boyce 1996, § VII-IX), Boyce, Ran and LeBlanc (Boyce, Ran, and LeBlanc 1995; Ran, Boyce, and LeBlanc 1993), Friesz et al. (Friesz et al. 1989). Indeed it seems more realistic that the information available to a navigation device is rather instantaneous and certainly not complete, that is, congestion information is available only as an aggregate (estimated waiting times for road traversal) but the individual routes and/or source and destination of travelers are usually unknown. For the Vickrey model, Graf, Harks and Sering (Graf, Harks, and Sering 2020) established the existence of instantaneous dynamic equilibria and derived further structural properties. One key property that differentiates dynamic equilibria (in the full information model) from instantaneous dynamic equilibria is the possibility of *cyclic behavior* in the latter. More specifically, (Graf, Harks, and Sering 2020) show that there are instances with only two origin-destination pairs and a finite flow volume in which *any* instantaneous dynamic equilibrium cycles forever. This can never happen in the full information model as an agent plays a best-response given the collective decisions of all other agents, thus, any cycle only increases the travel time.

Our Contribution

We propose a new DTA formulation within the Vickrey model that is based on *predicted* travel times. Since the physical transit times are known a priori, the only unknown is the precise evolution of the queues over time. In our model, every agent is associated with a *queue prediction function* that provides for any future point in time a prediction of queues. This model includes as special cases the full information model and the instantaneous information model but it allows to use predictions based on historical data or the queueing evolution learned en route. Besides these special cases, our model allows other queue prediction functions and even includes the case of finitely many *classes* of agents that may use different predictors.

As our main theoretical contribution, we define this model formally and derive conditions for the queue predictors leading to the existence of dynamic equilibria. The main approach is based on an *extension property* of partial equilibrium flows, that is, we show that any equilibrium flow up to some time $\theta \geq 0$ can be extended to time $\theta + \alpha$ for some $\alpha > 0$ which leads to the existence on the whole $\mathbb{R}$ using Zorns' lemma. The extension step itself is based on a formulation using infinite dimensional variational inequalities in the edge-flow space and whenever the predictor satisfies a natural continuity condition, only depends on past information and the predicted arrival times are non-decreasing, the extension is possible.

While this approach is in line with previous existence proofs using variational inequalities as put forth in the seminal papers by Friesz et al. (Friesz et al. 1993; Han, Friesz, and Yao 2013a,b,c), there are some remarkable differences. The above works rely on the complete spatio-temporal unfolding of the path-inflows over the network which is known as *network loading*. As shown in (Graf, Harks, and Sering 2020), already the simple prediction function given by the constant current queues (which leads to the instantaneous route choice model) leads to dynamic equilibria with cycling behavior (forever) and thus puts a path-based formulation over the entire time horizon out of reach. Our extension approach follows the extension-methodology used in (Graf, Harks, and Sering 2020) for the case of constant prediction functions. The more general model, however, comes with several technical difficulties that we need to address. We demonstrate the applicability of our main result by showing that it applies for instance to a natural linear regularized predictor $\hat{q}_e^{\mathrm{RL}}$. The idea here is to predict the queue growth linearly based on previously observed data on the time interval $[\bar{\theta} - \delta, \bar{\theta}]$. The regularization is necessary to obtain a continuous predictor since the purely linearized predictor $\hat{q}_e^{\mathrm{L}}(\theta; \bar{\theta}; q) := (\theta - \bar{\theta})\partial_- q_e(\bar{\theta})$ may be discontinuous as a function in the variable q.

On the experimental side, we conduct a simulation on a small synthetic network, on the commonly used Sioux Falls network from (LeBlanc, Morlok, and Pierskalla 1975) and on a larger real road network of Tokyo, Japan, obtained from Open Street Maps (OpenStreetMap contributors 2017). We study how the average travel time of vehicles in the network is impacted by the application of various predictors. For this purpose, we also train a linear regression model, for use as one of our predictors.

Related Work

The idea of using real-time information and traffic predictions en route and subsequently change the route is by no means new and has been proposed under varying names such as ATIS (advanced traveller information systems), see (Chorus, Molin, and Wee 2005; Watling 1994; Yang 1998) for an overview. Ben-Akiva et al. (Ben-Akiva et al. 2002) introduced DynaMIT, a simulation-based approach designed to predict future traffic conditions. Other works that also rely on simulation-based models include (Mahmassani 2001). Peeta and Mahmassani (Peeta and Mahmassani 1995) introduced a rolling horizon framework addressing the real-time traffic assignment problem. This approach concatenates for fixed consecutive time-intervals static flow assignments and thus does not comply to our definition of dynamic equilibrium in which at any time (also within stages) equilibrium conditions must hold. Huang and Lam (Huang and Lam 2003) allow for different user classes where each class may use a different travel time prediction. Their model is formulated in discrete time and assumes an acyclic path formulation.

A large body of research has been dedicated to the use of deep learning techniques, in particular *graph neural networks* (GNNs), for predicting street segment delays in road networks. It is impossible to list all relevant work in this section, we instead describe some key papers and point the reader to (Jiang and Luo 2021) for a complete survey. The work in (Li et al. 2018) uses a random walk-based graph diffusion process to create a convolutional operator that captures spatial relations. In (Yu, Yin, and Zhu 2018), the authors propose a spatio-temporal graph convolutional network which model the temporal dependency, whereas (Wu et al. 2019) models the spatial dependency through an adaptive learnable dependency matrix and the temporal dependency with dilated convolution (Oord et al. 2016). Finally, *graph attention networks* (GATs) (Velickovic et al. 2018) have also been used in the context of traffic predictions (Zheng et al. 2020). We note that our work bridges the above areas of dynamic route updates based on real time information and of applying ML models for predicting traffic delay.

Gentile (Gentile 2016) considered a mathematical approach incorporating traffic predictions in a dynamic traffic assignment (DTA) model. He derives the existence of equilibria using a variational inequality approach for the considered DTA model under simplifying assumptions such as an acyclic graph. The VI approach is arc and node-based and for its correctness, the assumption on acyclic (finite) paths is necessary as he uses a telescopic sum of edge travel times in order to arrive at a path-based VI formulation as used in (Friesz et al. 1993). Note that this approach fails in the general setting we consider in this paper. For further references on adaptive route choice models we refer to (Kucharski and Gentile 2019; Marcotte, Nguyen, and Schoeb 2004; Hamdouch, Marcotte, and Nguyen 2004; Unnikrishnan and Waller 2009; Watling and Hazelton 2003).

For works analyzing the inefficiency of dynamic equilibria (within the full or instantaneous information model), we refer to (Bhaskar, Fleischer, and Anshelevich 2015; Cao et al. 2017; Correa, Cristi, and Oosterwijk 2019; Graf and Harks 2020).

The Flow Model

In the following, we describe the Vickrey fluid queuing model that we will use throughout this paper. We consider a finite directed graph $G = (V, E)$ with positive rate capacities $\nu_e \in \mathbb{R}_{>0}$ and positive transit times $\tau_e \in \mathbb{R}_{>0}$ for every edge $e \in E$. There is a finite set of commodities $I = \{1, \ldots, n\}$, each with a commodity-specific source node $s_i \in V$ and a commodity-specific sink node $t_i \in V$. We assume that there is at least one s_i-t_i path for each $i \in I$. The (infinitesimally small) agents of every commodity $i \in I$ enter the network according to a locally integrable, bounded network inflow rate function $u_i : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$.

A *flow over time* is a tuple $f = (f^+, f^-)$, where $f^+, f^- : \mathbb{R}_{\geq 0} \times E \times I \to \mathbb{R}_{\geq 0}$ are locally integrable functions modeling the edge inflow rate $f^+_{i,e}(\theta)$ and edge outflow rate $f^-_{i,e}(\theta)$ of commodity i of an edge $e \in E$ at time $\theta \in \mathbb{R}_{\geq 0}$. The *queue length* of edge e at time θ is given by

$$q_e(\theta) := \sum_{i \in I} F^+_{i,e}(\theta) - \sum_{i \in I} F^-_{i,e}(\theta + \tau_e), \tag{1}$$

for $\theta \in \mathbb{R}_{\geq 0}$, where $F^+_{i,e}(\theta) := \int_0^\theta f^+_{i,e}(z) \,\mathrm{d}z$ and $F^-_{i,e}(\theta) := \int_0^\theta f^-_{i,e}(z) \,\mathrm{d}z$ denote the *cumulative (edge) inflow* and *cumulative (edge) outflow*. We implicitly assume $f^-_{i,e}(\theta) = 0$ for all $\theta \in [0, \tau_e)$, which will ensure together with Constraint (4) (see below) that the queue lengths are always nonnegative. For the sake of simplicity, we denote the aggregated in- and outflow rates for all commodities by $f^+_e := \sum_{i \in I} f^+_{i,e}$ and $f^-_e := \sum_{i \in I} f^-_{i,e}$, respectively.

A *feasible* flow over time satisfies the following conditions (2), (3), (4), and (5). The *flow conservation constraints* are modeled for a commodity $i \in I$ and all nodes $v \neq t_i$ as

$$\sum_{e \in \delta^+_v} f^+_{i,e}(\theta) - \sum_{e \in \delta^-_v} f^-_{i,e}(\theta) = \begin{cases} u_i(\theta) & \text{if } v = s_i, \\ 0 & \text{if } v \neq s_i, \end{cases} \tag{2}$$

for $\theta \in \mathbb{R}_{\geq 0}$ where $\delta^+_v := \{\, vu \in E \,\}$ and $\delta^-_v := \{\, uv \in E \,\}$ are the sets of outgoing edges from v and incoming edges into v, respectively. For the sink node t_i of commodity i we require

$$\sum_{e \in \delta^+_{t_i}} f^+_{i,e}(\theta) - \sum_{e \in \delta^-_{t_i}} f^-_{i,e}(\theta) \leq 0 \quad \text{for all } \theta \in \mathbb{R}_{\geq 0}. \tag{3}$$

We assume that the queue operates at capacity which can be modeled by requiring

$$f^-_e(\theta + \tau_e) = \begin{cases} \nu_e & \text{if } q_e(\theta) > 0, \\ \min\{\, f^+_e(\theta), \nu_e \,\} & \text{else}, \end{cases} \tag{4}$$

for all $e \in E, \theta \in \mathbb{R}_{\geq 0}$.

Finally, we want the flow to follow a strict FIFO principle on the queues, which can be formalized by

$$f^-_{i,e}(\theta) = \begin{cases} f^-_e(\theta) \cdot \frac{f^+_{i,e}(\vartheta)}{f^+_e(\vartheta)} & \text{if } f^+_e(\vartheta) > 0, \\ 0 & \text{else}, \end{cases} \tag{5}$$

where $\vartheta := \min\{\, \vartheta \leq \theta \mid \vartheta + \tau_e + \frac{q_e(\vartheta)}{\nu_e} = \theta \,\}$ is the earliest point in time a particle can enter edge e and leave at time θ and $\frac{q_e(\vartheta)}{\nu_e}$ is the current waiting time to be spent in the queue of edge e. Consequently, constraint (5) ensures that the share of commodity i of the aggregated outflow rate at any time equals the share of commodity i of the aggregated inflow rate at the time the particles entered the edge.

Instantaneous Dynamic Equilibrium

In an instantaneous dynamic equilibrium (IDE) as defined in (Graf, Harks, and Sering 2020) we assume that, whenever an agent arrives at an intermediate node v at time θ, she is given the information about the current queue length $q_e(\theta)$ and transit time τ_e of all edges $e \in E$, and, based on this information, she computes a shortest v-t_i path and enters the first edge on this path. We define the *instantaneous travel time* of an edge e at time θ as $c_e(\theta) := \tau_e + \frac{q_e(\theta)}{\nu_e}$. With this we can define commodity-specific node labels $\ell_{i,v}(\theta)$ corresponding to current earliest arrival times when travelling from v to the sink t_i at time θ by

$$\ell_{i,v}(\theta) := \begin{cases} \theta & \text{for } v = t_i, \\ \min_{e = vw \in E} \ell_{i,w}(\theta) + c_e(\theta) & \text{else}. \end{cases} \tag{6}$$

We say that edge $e = vw$ is *active* for $i \in I$ at time θ, if $\ell_{i,v}(\theta) = \ell_{i,w}(\theta) + c_e(\theta)$ and we denote the set of active edges for commodity i by $E_i(\theta) \subseteq E$.

Definition 1. A feasible flow over time f is an *instantaneous dynamic equilibrium (IDE)*, if for all $i \in I, \theta \in \mathbb{R}_{\geq 0}$ and $e \in E$ it satisfies

$$f^+_{i,e}(\theta) > 0 \implies e \in E_i(\theta). \tag{7}$$

Dynamic Nash Equilibrium

In contrast, in the full information model we assume that agents have complete knowledge of the entire (future) evolution of the dynamic flow. If an agent enters an edge $e = vw$ at time time θ, the travel time is $c_e(\theta) := \tau_e + \frac{q_e(\theta)}{\nu_e}$ and the exit time of edge e is given by $T_e(\theta) := \theta + c_e(\theta)$. In this setting it is common (cf. (Cominetti, Correa, and Larré 2015)) to define the node labels in such a way as to denote the earliest possible arrival time at each node (starting from the commodity's source node). Here, however, we will instead use an equivalent definition more in line with the node labels for IDEs. So, for any $i \in I, v \in V$ and $\theta \in \mathbb{R}_{\geq 0}$ we define a node label $\ell_{i,v}(\theta)$ denoting the earliest possible arrival time at node t_i for a particle starting at time θ at node v by setting

$$\ell_{i,v}(\theta) := \begin{cases} \theta & \text{for } v = t_i, \\ \min_{e=vw\in\delta_v^+} \ell_{i,w}(T_e(\theta)) & \text{else.} \end{cases} \quad (8)$$

We, again, say that an edge $e = vw$ is *active* for commodity $i \in I$ at time θ, if it holds that $\ell_{i,v}(\theta) = \ell_{i,w}(T_e(\theta))$ and denote by $E_i(\theta)$ the set of active edges for commodity i at time θ.

Definition 2. A feasible flow over time f is a *dynamic equilibrium (DE)*, if for all $e \in E, i \in I$ and $\theta \geq 0$ it holds that

$$f^+_{i,e}(\theta) > 0 \implies e \in E_i(\theta). \quad (9)$$

Dynamic Prediction Equilibria

IDE is a short-sighted behavioral concept assuming that agents at time $\bar{\theta}$ predict the future evolution of queue sizes according to the constant function $q_e(\theta) = q_e(\bar{\theta})$ for all $\theta \geq \bar{\theta}$. In the following we will relax this behavioral assumption by introducing a model wherein every commodity $i \in I$ maintains a predictor $\hat{q}_{i,e}$ for every edge $e \in E$. For a given flow over time and any two times $\theta \geq \bar{\theta}$ the value $\hat{q}_{i,e}(\theta; \bar{\theta}; q)$ is then the queue length at time θ on edge e as predicted by commodity i at time $\bar{\theta}$. Formally, a predictor $\hat{q}_{i,e}$ has the following signature:

$$\hat{q}_{i,e} : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \times \mathrm{C}\left(\mathbb{R}_{\geq 0}, \mathbb{R}_{\geq 0}\right)^E \to \mathbb{R}_{\geq 0}$$

In general such a predictor can depend in any arbitrary way on the entire input data including, in particular, the future evolution of the queue lengths after the prediction time $\bar{\theta}$. However, for our theoretical results we require the predictors to behave in a slightly more restricted way. First we want the predictors to depend continuously on query time, prediction time and the observed queue lengths.

Definition 3. We call a predictor $\hat{q}_{i,e}$ *continuous*, if the mapping

$$\hat{q}_{i,e} : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \times C(\mathbb{R}_{\geq 0}, \mathbb{R}_{\geq 0})^E \to \mathbb{R}_{\geq 0}$$

is continuous from the product topology, where all $C(\mathbb{R}_{\geq 0}, \mathbb{R}_{\geq 0})$ are equipped with the topology induced by the extended uniform norm, to the standard topology on $\mathbb{R}_{\geq 0}$.

The second property, which is also important for implementing the predictors, is that the predictors do not use (and, therefore, do not need) any information on the future evolution of the queues.

Definition 4. A predictor $\hat{q}_{i,e}$ is called *oblivious*, if the following condition holds

$$\forall \theta, \bar{\theta}, q, q' : \quad q_{\leq \bar{\theta}} = q'_{\leq \bar{\theta}} \implies \hat{q}_{i,e}(\theta; \bar{\theta}; q) = \hat{q}_{i,e}(\theta; \bar{\theta}; q'),$$

where $q_{\leq \bar{\theta}}$ denotes the restriction of the function $q : \mathbb{R}_{\geq 0} \times E \to \mathbb{R}_{\geq 0}$ to $[0, \bar{\theta}] \times E$.

The final property ensures that at any point in time there are shortest paths with respect to the predicted queue lengths that are cycle free. However, before we can formally define this property, we need some additional notation. If an agent of commodity $i \in I$ enters an edge $e = vw$ at time θ, the predicted travel time estimated at time $\bar{\theta}$ is given by $\hat{c}_{i,e}(\theta; \bar{\theta}; q) := \tau_e + \frac{\hat{q}_{i,e}(\theta; \bar{\theta}; q)}{\nu_e}$ and the predicted exit time of edge e is given by $\hat{T}_{i,e}(\theta; \bar{\theta}; q) := \theta + \hat{c}_{i,e}(\theta; \bar{\theta}; q)$. We call these times *$\bar{\theta}$-estimated* to emphasize that these values are predictions made at time $\bar{\theta}$.

Definition 5. A predictor $\hat{q}_{i,e}$ *respects FIFO* if for any edge e, queue lengths functions q and prediction time $\bar{\theta}$ the predicted exit time $\hat{T}_{i,e}(\,\cdot\,; \bar{\theta}; q)$ is a monotonically non-decreasing function.

This now allows us to describe how agents determine routes according to the predicted queues. At time $\bar{\theta}$ an agent of commodity $i \in I$ predicts that if she enters a path $P = (e_1, \dots, e_k)$ at time θ she will arrive at the endpoint of P at time

$$\hat{\ell}_i^P(\,\cdot\,; \bar{\theta}; q) := \hat{T}_{i,e_k}(\,\cdot\,; \bar{\theta}; q) \circ \dots \circ \hat{T}_{i,e_1}(\,\cdot\,; \bar{\theta}; q). \quad (10)$$

Denoting the (finite) set of all simple v-t_i paths by $\mathcal{P}_{i,v}$, the earliest $\bar{\theta}$-estimated time at which an agent starting at time θ from node v can reach t_i is given by

$$\hat{\ell}_{i,v}(\theta; \bar{\theta}; q) := \min_{P \in \mathcal{P}_{i,v}} \hat{\ell}_i^P(\theta; \bar{\theta}; q), \quad (11)$$

where the minimum over an empty set is infinity. The label functions defined in (11) satisfy the following equations:

$$\hat{\ell}_{i,v}(\theta; \bar{\theta}; q) = \begin{cases} \theta & \text{if } v = t_i, \\ \min\limits_{vw \in \delta_v^+} \hat{\ell}_{i,w}(\hat{T}_{i,vw}(\theta; \bar{\theta}; q); \bar{\theta}; q) & \text{if } v \neq t_i. \end{cases}$$

We say that an edge $e = vw$ is *$\bar{\theta}$-estimated active* for commodity i at time θ, if $\hat{\ell}_{i,v}(\theta; \bar{\theta}; q) = \hat{\ell}_{i,w}(\hat{T}_{i,e}(\theta; \bar{\theta}; q); \bar{\theta}; q)$ holds true. Furthermore, let us denote the set of $\bar{\theta}$-estimated active edges for commodity i at time θ by $\hat{E}_i(\theta; \bar{\theta}; q)$.

Definition 6. A pair $(\hat{q}, f)$ of a set of predictors $\hat{q} = (\hat{q}_{i,e})_{i \in I, e \in E}$ and a flow over time f is a *dynamic prediction equilibrium (DPE)*, if for all $e \in E, i \in I$ and $\theta \geq 0$ it holds that

$$f^+_{i,e}(\theta) > 0 \implies e \in \hat{E}_i(\theta; \theta; q).$$

We then also call the flow f a *dynamic prediction flow with respect to the predictor $\hat{q}$*.

Existence of Dynamic Prediction Equilibria

In this section we show that for oblivious and continuous predictors that respect FIFO there always exist dynamic prediction equilibria. We will also give several examples of such predictors, including one inducing IDEs as corresponding equilibria.

Existence of DPE Using a Variational Inequality

To show the existence of DPE we make use of a result by Brézis (Brézis 1968, Theorem 24) guaranteeing the existence of solutions to certain variational inequalities.

Theorem 7. *Let $[a, b) \subseteq \mathbb{R}_{\geq 0}$ be some interval, $d \in \mathbb{N}$, $K \subseteq L^2([a, b))^d$ a nonempty, closed, convex and bounded set and $\mathcal{A} : K \to L^2([a, b))^d$ a weak-strong-continuous mapping. Then there exists a point $g^* \in K$ such that*

$$\langle \mathcal{A}(g^*), g - g^* \rangle \geq 0 \text{ for all } g \in K. \tag{12}$$

This theorem can be used to build up a dynamic predicted flow with respect to a given set of predictors by iteratively extending so-called *partial dynamic prediction flows* which fulfill the equilibrium property up to some time horizon. First, we formally introduce these flows:

Definition 8. A *partial flow up to time* ϕ is a tuple $f = (f^+, f^-)$ of locally integrable functions $f^+, f^- : \mathbb{R}_{\geq 0} \times E \times I \to \mathbb{R}_{\geq 0}$ fulfilling conditions (2), (3) and (4) for $\theta \leq \phi$. We call f a *partial dynamic prediction flow with respect to a set of oblivious predictors $\hat{q}$ up to time ϕ*, if $f^+_{i,e}(\theta) > 0$ implies $e \in \hat{E}_i(\theta; \theta; q)$ for all $\theta \leq \phi, e \in E, i \in I$.

We will now show that such a partial dynamic prediction flow can always be extended for some additional time interval. We will employ a similar proof-technique to the one used in (Graf, Harks, and Sering 2020, Lemma 5.6) for the proof of the extension property of IDEs flows. However, the analysis is more involved as we allow for a more general functional dependence of the predicted queue lengths on the past flow evolution. This stands in contrast to IDEs where each prediction only depends on the queue lengths of one edge at a single point in time.

Lemma 9. *Let I be a finite set of commodities with locally integrable, bounded network inflow functions u_i and let $\hat{q} = (\hat{q}_{i,e})_{i \in I, e \in E}$ be a set of continuous and oblivious predictors that respect FIFO. We can extend any partial dynamic prediction flow f with respect to $\hat{q}$ up to time ϕ to a dynamic prediction flow up to time $\phi + \alpha$ for any $0 < \alpha < \min_{e \in E} \tau_e$.*

We will only give a brief proof sketch here – the full proof can be found in the full version (Graf et al. 2021a). The main idea is to first define a set K of all possible extensions of the given partial flow. We then define a mapping $\mathcal{A} : K \to L^2(D)^{I \times E}$ associating with each possible extension a function which for every commodity i, edge e and time θ is zero if and only if this edge is active for this commodity at this time. Using the continuity of the predictors we then show that this mapping is weak-strong continuous such that we can apply Theorem 7 to get a solution to the variational inequality (12). Finally, we show that this solution is indeed an extension which also satisfies the properties of a dynamic prediction flow.

With this key-lemma we can now show the existence of dynamic prediction flows for all oblivious and continuous predictors that respect FIFO. Starting with the zero-flow up to time 0 and iteratively applying Lemma 9 gives us a partial dynamic prediction flow up to any finite time horizon. Zorn's lemma then shows the existence of a dynamic prediction flow for all times, thus, proving our main theorem:

Theorem 10. *For any network with finite set of commodities, each associated with a locally integrable, bounded network inflow rate and oblivious and continuous predictors $\hat{q}_{i,e}$ that respect FIFO, there exists a dynamic prediction flow with respect to $\hat{q}$.*

Example 11. To see why we require the predictors to be continuous, consider the non-continuous predictor

$$\hat{q}_e(\theta; \bar{\theta}; q) := \begin{cases} q_e(\bar{\theta}), & \text{if } q_e(\bar{\theta}) < 1 \\ 2, & \text{else.} \end{cases}$$

Using this predictor in a network consisting of only a single source-sink pair connected by two parallel edges e_1 and e_2 can already lead to a situation where no equilibrium flow exists. Let $\nu_{e_1} = 1, \tau_{e_1} = 1, \nu_{e_2} = 2, \tau_{e_2} = 2$ and assume a constant inflow rate of 2 at the source. Then, clearly, during the time interval $[0, 1)$ agents using the above predictor may only enter edge e_1 (as the predicted travel time along edge e_1 is strictly smaller than 2). Beginning with time $\theta = 1$, however, every possible flow split will violate the equilibrium condition, since at that time edge e_1 has a queue length of 1 and, thus, a predicted queue length of 2. On the one hand, sending agents into edge e_1 at a rate of less than 1 for any period of time after $\theta = 1$, leads to an immediate decrease of its queue lengths and, thus, edge e_2 becomes inactive again. If, on the other hand, agents enter edge e_1 at a rate of 1 or more its queue length will remain at least 1 and, therefore, edge e_1 will be inactive.

Application Predictors

We now discuss several predictors and analyze whether the theorem above can be applied. We begin with simple predictors and make them more sophisticated step-by-step.

The *Zero-Predictor* predicts no queues for all times, i.e.

$$\hat{q}^{\mathrm{Z}}_{i,e}(\theta; \bar{\theta}; q) = 0.$$

This predictor is trivially continuous and oblivious and respects FIFO. The resulting dynamic prediction flow is a flow, where particles just always follow physically shortest paths.

The *constant predictor* predicts in a continuous way that all queues will stay constant:

$$\hat{q}^{\mathrm{C}}_{i,e}(\theta; \bar{\theta}; q) = q_e(\bar{\theta}).$$

This leads to the mentioned special case of IDE flows. Since the constant predictor clearly is continuous and oblivious and respects FIFO we can apply Theorem 10 and, thus, reprove the existence of IDE flows shown in (Graf, Harks, and Sering 2020).

The *linear predictor* takes the derivative of the queues and extends them linearly up to some fixed time horizon $H \in \mathbb{R}_{\geq 0} \cup \{\infty\}$. Formally it is defined as

$$\hat{q}^{\mathrm{L}}_{i,e}(\theta;\bar{\theta};q) := \left(q_e(\bar{\theta}) + \partial_- q_e(\bar{\theta}) \cdot \min\{\theta - \bar{\theta}, H\}\right)^+,$$

where $(x)^+ := \max\{x, 0\}$ denotes the positive part of $x \in \mathbb{R}$. The linear predictor is not in general continuous since the partial derivative $\partial_- q_e(\bar{\theta})$ might be discontinuous.

The *regularized linear predictor* solves this by taking a rolling average of the past gradient (with rolling horizon $\delta > 0$) and extend the prediction queues according to this:

$$\hat{q}^{\mathrm{RL}}_{i,e}(\theta;\bar{\theta};q) := \left(q_e(\bar{\theta}) + \frac{q_e(\bar{\theta}) - q_e(\bar{\theta} - \delta)}{\delta} \cdot \min\{\theta - \bar{\theta}, H\}\right)^+$$

Proposition 12. *The regularized linear predictor is oblivious and continuous and respects FIFO. It, thus, induces the existence of a dynamic prediction equilibrium.*

The proof is a direct computation (see (Graf et al. 2021a) for the detailed proof). A different way of understanding the regularized predictor is that it takes two samples from the past queue lengths (at time $\bar{\theta}$ and $\bar{\theta} - \delta$) and uses this information to predict future queue lengths up to the prediction horizon by a linear function. We can generalize this idea by taking more samples of the past queue (possibly also of queues of neighbouring edges) and use these values to find a piecewise linear prediction of the queue length for the future. More precisely, given some sample number k, some step size δ, and a neighbourhood edge set $N(e) \subseteq E$ we choose real numbers $a^{e'}_{i,j}$ for $i = 1, \dots, k$, $j = 1, \dots, H/\delta$ and $e' \in N(e)$. Our predictor is then the piecewise linear function interpolating between the points $(\bar{\theta} + j\delta, (\sum_{e' \in N(e)} \sum_{i=1}^{k} a^{e'}_{i,j} \cdot q_{e'}(\bar{\theta} - i\delta))^+)$ for $j = 1, \dots, H/\delta$. We will denote such a predictor by $\hat{q}^{\mathrm{ML}}$.

Proposition 13. *For appropriately chosen $a^{e'}_{i,j}$ the predictor $\hat{q}^{ML}$ is oblivious and continuous and respects FIFO. It, thus, induces the existence of a dynamic prediction equilibrium.*

The predictor $\hat{q}^{\mathrm{ML}}$ always is continuous and oblivious by the same arguments as for the regularized linear predictor. It also respects FIFO, if the numbers $a^{e'}_{i,j}$ are chosen in such a way that the predicted queue-length never decreases faster than by at rate of ν_e. As this function is piece-wise linear we can check easily check this condition for any given set of numbers $a^{e'}_{i,j}$. This leaves the question of how to choose these parameters in order to achieve a good predictor. In our experimental section below, we will use machine learning to learn these by evaluating past data. Consequently, we call the predictor $\hat{q}^{\mathrm{ML}}$ a *linear regression predictor*. We provide more details on the features and data used to train the predictor in the following section.

Finally, the *perfect predictor* predicts the queues exactly as they will evolve, i.e. it satisfies

$$\hat{q}^{\mathrm{P}}_{i,e}(\theta;\bar{\theta};q) := q_e(\theta).$$

This predictor clearly is not oblivious and, thus, we can not apply our existence result here. However, dynamic predicted flows with respect to this predictor do exist as those are just dynamic equilibria for which existence has been proven in (Cominetti, Correa, and Larré 2015).

Computational Study

In the following computational study, we compare the different predictors introduced in the last section with a machine-learning based alternative. To compare the predictors we introduce an extension based algorithm computing an approximation of a DPE for a given set of predictors. We also use this algorithm to generate training data for the machine learning system using the constant predictors, which results in an approximation of IDEs.

As a metric for a predictor's performance we monitor its average travel time in a flow with multiple predictors used side by side: Let i be a commodity with constant net inflow rate up to some time h, i.e. $u_i(\theta) := \bar{u}_i$ for $\theta \leq h$ and $u_i(\theta) := 0$ for $\theta > h$. The net outflow rate of commodity i is given by $o_i(\theta) := \sum_{e \in \delta_t^-} f^-_{i,e}(\theta) - \sum_{e \in \delta_t^+} f^+_{i,e}(\theta)$. Taking the integral of $u_i(\psi) - o_i(\psi)$ over $[0, \phi]$ yields the flow of commodity i inside the network at time ϕ. If we integrate this quantity over $\phi \in [0, H]$ with $H \geq h$, we obtain the *total travel time* of particles of commodity i up to time H:

$$T_i^{\text{total}} := \int_0^H \int_0^\phi u_i(\psi) - o_i(\psi)\, \mathrm{d}\psi\, \mathrm{d}\phi$$

Now, $T_i^{\text{avg}} := T_i^{\text{total}}/(h \cdot \bar{u}_i)$ denotes the *average travel time*. We compare these values against their optimum which can be computed as $T_{i,\text{OPT}}^{\text{avg}} := \int_0^h \min\{H, l_{i,s}(\theta)\} - \theta \,\mathrm{d}\theta / h$.

Extension based simulation

As proposed in our model, each infinitesimal agent updates its route each time after traversing an edge. As our flow is continuous, this would imply that the prediction and therefore also the shortest paths are updated in a continuous manner. For a computational study, this can only be approximated: In our implementation, we assume that shortest paths as predicted at some time $\bar{\theta}$ stay shortest paths for a certain time interval $[\bar{\theta}, \bar{\theta} + \varepsilon)$ and agents compute new shortest paths every ε time units resulting in ε-DPEs:

Definition 14. For $\varepsilon > 0$, a pair $(\hat{q}, f)$ of a set of predictors $\hat{q} = (\hat{q}_{i,e})_{i \in I, e \in E}$ and a flow over time f is an *ε-approximated dynamic prediction equilibrium (ε-DPE)*, if for all $e \in E, i \in I$ and $\theta \geq 0$ it holds that

$$f^+_{i,e}(\theta) > 0 \implies e \in \hat{E}_i(\varepsilon \cdot \lfloor \theta/\varepsilon \rfloor; \varepsilon \cdot \lfloor \theta/\varepsilon \rfloor; q).$$

In our implementation we maintain piece-wise constant inflow and outflow rate functions $f^+_{i,e}$, $f^-_{i,e}$ as well as piece-wise linear queue lengths q_e. We have a sequence of prediction times $\bar{\theta}_k = k \cdot \varepsilon$ at which new predictions are retrieved in the form of piece-wise linear functions $\hat{q}_{i,e}(\,\cdot\,; \bar{\theta}_k; q)$. From these predictions, we derive the time-dependent cost functions $\hat{c}_{i,e}(\,\cdot\,; \bar{\theta}_k; q)$. If we use predictors respecting FIFO, we can use a dynamic variant of the Dijkstra Algorithm to compute the active edges $\hat{E}_i(\bar{\theta}_k; \bar{\theta}_k; q)$.

We then send flow along these active edges until the next prediction time $\bar{\theta}_{k+1}$. This is done in a so-called distribution phase: Let us first assume, that the node inflow $b^-_{i,v}(\theta) := \sum_{e \in \delta_v^-} f^-_{i,e}(\theta)$ is constant on some proper interval $[\phi, \phi + \alpha) \subseteq [\bar{\theta}_k, \bar{\theta}_{k+1})$. The edge inflow rates of

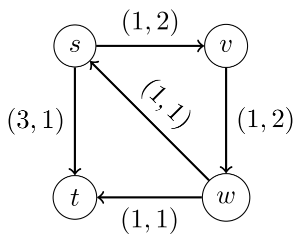

Figure 1: A network with constant inflow at source s. The only sink is node t. Edges are labeled with (τ_e, ν_e).

edges $e \in \delta_v^+$ are then extended on $[\phi, \phi + \alpha)$ by setting $f_{i,e}^+(\theta) \coloneqq b_{v,i}^-(\theta) / \left|\delta_v^+ \cap \hat{E}_i(\bar{\theta}_k; \bar{\theta}_k; q)\right|$ if $e \in \hat{E}_i(\bar{\theta}_k; \bar{\theta}_k; q)$ and $f_{i,e}^+(\theta) \coloneqq 0$ otherwise. The edge outflow rates are then determined using conditions (4) and (5).

To build a feasible flow we have to comply with the flow conservation constraints when extending the flow. As the outflow rate of edges may vary during a single interval $[\bar{\theta}_k, \bar{\theta}_{k+1})$ we can only extend the flow with the above method until some outflow rate changes, after which we start another distribution phase. By choosing $\alpha > 0$ such that $\phi + \alpha$ is the next time an edge outflow rate changes (or $\phi + \alpha = \bar{\theta}_{k+1}$), the flow conservation constraint is satisfied. Hence, there might be a multitude of smaller distribution phases during a single prediction interval. These subsequent distribution phases can be sped up by only updating nodes where edge outflow rates of incoming edges have changed.

The code of our simulations is publicly available in (Graf et al. 2021b).

Data

We conduct our experiments on three graphs. The first is a warm-up synthetic graph with 4 nodes and 5 edges. We present the graph in Figure 1. The second graph is the road map of Sioux Falls as given in (LeBlanc, Morlok, and Pierskalla 1975) which is commonly used in the transport science literature. It comes with edge attributes free-flow travel time τ_e and capacity ν_e. The third graph is the center of Tokyo as obtained from Open Street Maps (OpenStreetMap contributors 2017). This data set the free-flow speed, the length and the numbers of lanes of each road segment e. We compute the transit time τ_e as the product of the free-flow speed and the length of edge e. The capacity ν_e is calculated by multiplying the number of lanes with the free-flow speed. For the latter two networks, commodities are randomly chosen. Details are depicted in Table 1, where $T_{\text{avg}}^{\text{comp}}$ denotes the average computation time for computing a single ε-DPE on a single core of an Intel® Core™ i7-3520M CPU at 2.90GHz.

Network	$\lvert E\rvert$	$\lvert V\rvert$	$\lvert I\rvert$	$T_{\text{avg}}^{\text{comp}}$
Synthetic	5	4	5	$0.33s$
Sioux Falls	75	24	17	$10.92s$
Tokyo	4,803	3,538	40	$343.93s$

Table 1: Attributes of the considered networks

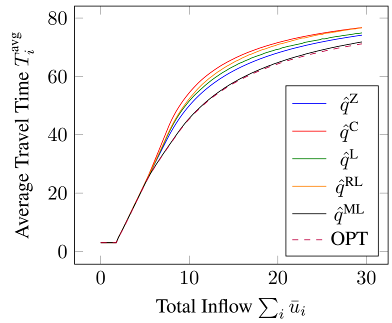

Figure 2: Average travel times of competing predictors in the synthetic network in Figure 1.

The Machine Learned Predictor

To assess the impact of ML-based models in our setting, we train a simple linear regression predictor for each network. To obtain training data for the regression, we run simulations using the proposed extension based algorithm with the simpler constant predictor. This allows the model to estimate the progression of queues when agents follow our behavioral model. The features used to train the model are 10 observations of the past queue length of the edge and of neighboring edges.

Comparison of Predictors

We first take a closer look at the synthetic network shown in Figure 1. Here, we want to analyze how the average travel times of competing predictors evolve while increasing the total network inflow. For each oblivious predictor described above, we add a commodity $i \in \{\hat{q}^{\text{Z}}, \hat{q}^{\text{C}}, \hat{q}^{\text{L}}, \hat{q}^{\text{RL}}, \hat{q}^{\text{ML}}\}$. Each of these commodities has the same source s and sink t and the same constant inflow $\bar{u}_i$ up to time $h = 25$. The outcome of running the simulation with time horizon $H = 100$ for each sampled total inflow in $(0, 30)$ can be seen in Figure 2. The ML based predictor performed best, while notably the Zero-Predictor, who sends flow along paths (s, t) and (s, v, w, t) equally at all times, performs better than the remaining predictors.

For the road-networks of Sioux Falls and Tokyo, we randomly generate inflow rates according to the edge capacities of the network. For each commodity i, we ran the simulation after adding 5 additional commodities with the same source and sink as i – one for each predictor – with a very small constant inflow rate. We monitored their average travel time as a measure of the performance of the different predictors. All other commodities in the network were assigned the constant predictor, such that the resulting queues should behave similar to the training data.

Generally, the Zero-Predictor performs the worst in this scenario; the machine learning based predictor performs similarly well as the remaining predictors. We include more detailed results in (Graf et al. 2021a). We believe it is an interesting future direction to explore more complex learning algorithms and how they interface with the dynamic prediction equilibrium concept as well as understand how different graph topologies impact the various predictors.

References

Ben-Akiva, M.; Bierlaire, M.; Koutsopoulos, H. N.; and Mishalani, R. 2002. *Real Time Simulation of Traffic Demand-Supply Interactions within DynaMIT*, 19–36. Boston, MA: Springer US.

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

Boyce, D. E.; Ran, B.; and LeBlanc, L. J. 1995. Solving an Instantaneous Dynamic User-Optimal Route Choice Model. *Transportation Science*, 29(2): 128–142.

Brézis, H. 1968. Équations et inéquations non linéaires dans les espaces vectoriels en dualité. *Annales de l'Institut Fourier*, 18(1): 115–175.

Cao, Z.; Chen, B.; Chen, X.; and Wang, C. 2017. A Network Game of Dynamic Traffic. In Daskalakis, C.; Babaioff, M.; and Moulin, H., eds., *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*, 695–696. ACM.

Chorus, C. G.; Molin, E. J. E.; and Wee, B. V. 2005. Use and Effects of Advanced Traveller Information Services (ATIS): A Review of the Literature. *Transport Reviews*, 26(2): 127–149.

Cominetti, R.; Correa, J. R.; and Larré, O. 2015. Dynamic Equilibria in Fluid Queueing Networks. *Oper. Res.*, 63(1): 21–34.

Correa, J. R.; Cristi, A.; and Oosterwijk, T. 2019. On the Price of Anarchy for flows over time. In Karlin, A.; Immorlica, N.; and Johari, R., eds., *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019*, 559–577. ACM.

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

Friesz, T. L.; Bernstein, D.; Smith, T. E.; Tobin, R. L.; and Wie, B. W. 1993. A variational inequality formulation of the dynamic network user equilibrium problem. *Oper. Res.*, 41(1): 179–191.

Friesz, T. L.; and Han, K. 2019. The mathematical foundations of dynamic user equilibrium. *Transportation Research Part B: Methodological*, 126: 309 – 328.

Friesz, T. L.; Luque, J.; Tobin, R. L.; and Wie, B. 1989. Dynamic Network Traffic Assignment Considered as a Continuous Time Optimal Control Problem. *Operations Research*, 37(6): 893–901.

Gentile, G. 2016. Solving a Dynamic User Equilibrium model based on splitting rates with Gradient Projection algorithms. *Transportation Research Part B: Methodological*, 92: 120–147.

Graf, L.; and Harks, T. 2020. The Price of Anarchy for Instantaneous Dynamic Equilibria. In Chen, X.; Gravin, N.; Hoefer, M.; and Mehta, R., eds., *Web and Internet Economics - 16th International Conference, WINE 2020, Beijing, China, December 7-11, 2020, Proceedings*, volume 12495 of *Lecture Notes in Computer Science*, 237–251. Springer.

Graf, L.; Harks, T.; Kollias, K.; and Markl, M. 2021a. Machine-Learned Prediction Equilibrium for Dynamic Traffic Assignment. *CoRR*, abs/2109.06713.

Graf, L.; Harks, T.; Kollias, K.; and Markl, M. 2021b. Simulation of Dynamic Prediction Equilibria. https://github.com/schedulaar/dynamic-prediction-equilibria. Accessed: 2021-12-15.

Graf, L.; Harks, T.; and Sering, L. 2020. Dynamic flows with adaptive route choice. *Math. Program., Ser. B*, 183(1): 309–335.

Hamdouch, Y.; Marcotte, P.; and Nguyen, S. 2004. A Strategic Model for Dynamic Traffic Assignment. *Networks and Spatial Economics*, 4(3): 291–315.

Han, K.; Friesz, T. L.; and Yao, T. 2013a. Existence of simultaneous route and departure choice dynamic user equilibrium. *Transportation Research Part B: Methodological*, 53: 17 – 30.

Han, K.; Friesz, T. L.; and Yao, T. 2013b. A partial differential equation formulation of Vickrey's bottleneck model, part I: Methodology and theoretical analysis. *Transportation Research Part B: Methodological*, 49: 55 – 74.

Han, K.; Friesz, T. L.; and Yao, T. 2013c. A partial differential equation formulation of Vickrey's bottleneck model, part II: Numerical analysis and computation. *Transportation Research Part B: Methodological*, 49: 75 – 93.

Huang, H.-J.; and Lam, W. H. K. 2003. A Multi-Class Dynamic User Equilibrium Model for Queuing Networks with Advanced Traveler Information Systems. *Journal of Mathematical Modelling and Algorithms*, 2(4): 349–377.

Jiang, W.; and Luo, J. 2021. Graph Neural Network for Traffic Forecasting: A Survey. *CoRR*, abs/2101.11174.

Koch, R.; and Skutella, M. 2011. Nash Equilibria and the Price of Anarchy for Flows over Time. *Theory Comput. Syst.*, 49(1): 71–97.

Kucharski, R.; and Gentile, G. 2019. Simulation of rerouting phenomena in Dynamic Traffic Assignment with the Information Comply Model. *Transportation Research Part B: Methodological*, 126(C): 414–441.

LeBlanc, L. J.; Morlok, E. K.; and Pierskalla, W. P. 1975. An efficient approach to solving the road network equilibrium traffic assignment problem. *Transportation Research*, 9(5): 309–318.

Li, Y.; Yu, R.; Shahabi, C.; and Liu, Y. 2018. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting. In *6th International Conference on Learning Representations, ICLR*.

Li, Z.-C.; Huang, H.-J.; and Yang, H. 2020. Fifty years of the bottleneck model: A bibliometric review and future research directions. *Transportation Research Part B: Methodological*, 139: 311 – 342.

Mahmassani, H. 2001. Dynamic Network Traffic Assignment and Simulation Methodology for Advanced System Management Applications. *Networks and Spatial Economics*, 1: 267–292.

Marcotte, P.; Nguyen, S.; and Schoeb, A. 2004. A Strategic Flow Model of Traffic Assignment in Static Capacitated Networks. *Operations Research*, 52(2): 191–212.

Meunier, F.; and Wagner, N. 2010. Equilibrium Results for Dynamic Congestion Games. *Transportation Science*, 44(4): 524–536. An updated version (2014) is available on Arxiv.

Oord, A. v. d.; Dieleman, S.; Zen, H.; Simonyan, K.; Vinyals, O.; Graves, A.; Kalchbrenner, N.; Senior, A.; and Kavukcuoglu, K. 2016. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*.

OpenStreetMap contributors. 2017. Planet dump retrieved from https://planet.osm.org . https://www.openstreetmap.org.

Peeta, S.; and Mahmassani, H. S. 1995. Multiple user classes real-time traffic assignment for online operations: A rolling horizon solution framework. *Transportation Research Part C: Emerging Technologies*, 3(2): 83–98.

Peeta, S.; and Ziliaskopoulos, A. 2001. Foundations of Dynamic Traffic Assignment: The Past, the Present and the Future. *Networks and Spatial Economics*, 1(3): 233–265.

Ran, B.; and Boyce, D. 1996. *Dynamic urban transportation network models: theory and implications for intelligent vehicle-highway systems*. Lecture notes in economics and mathematical systems. Berlin, New York, Paris: Springer.

Ran, B.; Boyce, D. E.; and LeBlanc, L. J. 1993. A New Class of Instantaneous Dynamic User-Optimal Traffic Assignment Models. *Operations Research*, 41(1): 192–202.

Skutella, M. 2008. An Introduction to Network Flows over Time. In *Research Trends in Combinatorial Optimization, Bonn Workshop on Combinatorial Optimization, November 3-7, 2008, Bonn, Germany*, 451–482.

Unnikrishnan, A.; and Waller, S. 2009. User Equilibrium with Recourse. *Networks and Spatial Economics*, 9(4): 575–593.

Velickovic, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph Attention Networks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*.

Vickrey, W. S. 1969. Congestion Theory and Transport Investment. *American Economic Review*, 59(2): 251–60.

Watling, D. 1994. Urban traffic network models and dynamic driver information systems. *Transport Reviews*, 14(3): 219–246.

Watling, D.; and Hazelton, M. L. 2003. The Dynamics and Equilibria of Day-to-Day Assignment Models. *Networks and Spatial Economics*, 3: 349–370.

Wu, Z.; Pan, S.; Long, G.; Jiang, J.; and Zhang, C. 2019. Graph WaveNet for Deep Spatial-Temporal Graph Modeling. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*. AAAI Press.

Yang, H. 1998. Multiple equilibrium behaviors and advanced traveler information systems with endogenous market penetration. *Transportation Research Part B: Methodological*, 32(3): 205–218.

Yu, B.; Yin, H.; and Zhu, Z. 2018. Spatio-Temporal Graph Convolutional Networks: A Deep Learning Framework for Traffic Forecasting. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden.*, 3634–3640.

Zheng, C.; Fan, X.; Wang, C.; and Qi, J. 2020. GMAN: A Graph Multi-Attention Network for Traffic Prediction. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, 1234–1241.

Zhu, D.; and Marcotte, P. 2000. On the Existence of Solutions to the Dynamic User Equilibrium Problem. *Transportation Science*, 34(4): 402–414.