

Prediction Equilibrium for Dynamic Network Flows

Lukas Graf

*University of Passau
94032 Passau, Germany*

LUKAS.GRAF@UNI-PASSAU.DE

Tobias Harks

*University of Passau
94032 Passau, Germany*

TOBIAS.HARKS@UNI-PASSAU.DE

Kostas Kollias

*Google
Mountain View, CA 94043, USA*

KOSTASKOLLIAS@GOOGLE.COM

Michael Markl

*University of Passau
94032 Passau, Germany*

MICHAEL.MARKL@UNI-PASSAU.DE

Editor: Dan Alistarh

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 define *dynamic prediction equilibrium (DPE)* in which no agent can at any point during their journey improve their predicted travel time by switching to a different route. 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 then proceed to derive properties of the predictors that ensure a dynamic prediction equilibrium exists. Additionally, we define ε -*approximate DPE* wherein no agent can improve their predicted travel time by more than ε and provide further conditions of the predictors under which such an approximate equilibrium can be computed. Finally, we complement our theoretical analysis by an experimental study, in which we systematically compare the induced average travel times of different predictors, including two machine-learning based models trained on data gained from previously computed approximate equilibrium flows, both on synthetic and real world road networks.

Keywords: Dynamic Traffic Assignment, Dynamic Flow, Prediction Equilibrium, Machine-Learning, Regression Predictors

1. Introduction

Modelling and optimizing traffic flows 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 and thus lead to a complex and self-referential system.

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 book of Ford and Fulkerson, 1962, or the more recent surveys of Friesz and Han, 2019, Peeta and Ziliaskopoulos, 2001, and Skutella, 2008). A fundamental base model describing the dynamic flow propagation process is the so-called *deterministic queuing model* due to Vickrey (1969). Here, a directed graph $G = (V, E)$ is given, where edges $e \in E$ are associated with a queue with a positive rate capacity ν_e and a physical transit time τ_e . 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 et al., 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 et al., 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 have been studied extensively in the transportation science literature, see Friesz et al. (1993); Han et al. (2013a,b,c); Meunier and Wagner (2010); Zhu and Marcotte (2000), whereas the works by Koch and Skutella (2011); Cominetti et al. (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, 1996, § VII-IX), (Boyce et al., 1995; Ran et al., 1993), 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 other travelers are usually unknown. For the Vickrey model, Graf et al. (2020) established the existence of instantaneous dynamic equilibria and derived further structural properties. A key difference between dynamic equilibria (in the full information model) and instantaneous dynamic equilibria is the possibility of *cyclic behavior* in the latter. More specifically, Graf et al.

(2020) describe 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.

1.1 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 flow over time. In our model, agents use 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 also allows the use of predictions based on historical data or the flow evolution learned en route. The model also includes the case of finitely many *classes* of agents similar to the concept described by Dafermos (1972), for which agents of the same class use the same prediction function, but the prediction function may vary across the different classes. Technically, this is solved by subdividing the traffic commodities by their different classes and assigning a prediction function to each (sub-)commodity.

Existence Results. As our main theoretical contribution, we define this model formally and derive conditions for the queue predictors leading to the existence of dynamic (prediction) 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 fixed constant $\alpha > 0$ which leads to the existence on the whole \mathbb{R} after a countably infinite number of extensions. The extension step itself is based on a formulation using infinite dimensional variational inequalities in the edge-flow space. If the predictor satisfies a continuity condition, only depends on past information, and respects a First-In-First-Out (FIFO) rule, the extension is possible and leads to existence of an equilibrium.

While this approach is in line with previous existence proofs using variational inequalities as put forth in the seminal papers by Friesz et al. (1993); Han et al. (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 by Graf et al. (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 approach uses an extension-methodology, which does not rely on the complete spatio-temporal unfolding of flow. We demonstrate the applicability of our main result by showing that it applies not only to several simple variants of predictors but also to predictors that implement an arbitrary continuous (machine-learned) transformation.

Computing Approximate Dynamic Prediction Equilibria. We show that under mild conditions on the predictors, an approximate version of a dynamic prediction equilibrium can be computed in finite time by introducing an extension-based algorithm. This algorithm computes a flow in which agents recompute their predicted shortest path to the sink with the latest predictions in a periodic interval. If the predictors behave according to

a Lipschitz condition, we show that these flow instances approximate dynamic prediction flows.

Computational Study. On the experimental side, we conduct experiments on small synthetic networks, on the Sioux Falls network from LeBlanc et al. (1975), and on a larger real road network of Anaheim obtained from Transportation Networks for Research Core Team (2016). For various predictors, we study how the particles’ choice of the predictor affects their average travel time. For this purpose, we also train linear regression and simple neural network models, for use as one of our predictors.

Comparison to Conference Version. An extended abstract of this paper was published in Graf et al. (2022b). In comparison, this work not only includes fully worked-out proofs of all theorems, but also generalizes the theoretical model of equilibria by making the predictors depend on the historical flow instead of its corresponding queue functions. Appropriate and thus weaker conditions for the existence of equilibria are worked out such that the original result occurs as a special case. This allows us to embed arbitrary continuous ML methods as opposed to only simple linear regression models. A detailed description of the algorithm for computing approximate equilibria, and the proof of its correctness and termination are added. Finally, we revised the computational study: We implement a new neural network based predictor and evaluate all predictors in a larger experiment setup based on real world origin-destination demand pairs.

1.2 Related Work

The idea of using real-time information and traffic predictions en route and subsequently change the route is not new and has been proposed under varying names such as ATIS (advanced traveller information systems), see (Chorus et al., 2005; Watling, 1994; Yang, 1998) for an overview. 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 (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 (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. Yu et al. (2018) propose a spatio-temporal graph convolutional network which model the temporal dependency, whereas Wu et al. (2019) model 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), introduced by Velickovic et al. (2018), have also been used in the context

of traffic predictions by 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 (2016) considered a mathematical approach incorporating traffic predictions in a 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 et al. (2004); Hamdouch et al. (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 et al. (2015); Cao et al. (2017); Correa et al. (2019); Graf and Harks (2020).

2. 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 , each with a commodity-specific source node $s_i \in V$ and a commodity-specific sink node $t_i \in V \setminus \{s_i\}$. We assume that there is at least one s_i - t_i path for each $i \in I$ and we denote the set of nodes and edges lying on an s_i - t_i -path as V_i and E_i , respectively. The infinitesimal agents of every commodity $i \in I$ enter the network according to a locally integrable network inflow rate function $u_i \in \mathcal{R} := \{g \in L^1_{\text{loc}}(\mathbb{R}, \mathbb{R}_{\geq 0}) \mid g|_{(-\infty, 0)} \stackrel{a.e.}{=} 0\}$.

A *flow over time* is a tuple $f = (f^+, f^-)$, where $f^+, f^- \in \mathcal{R}^{I \times E}$ model 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 $\theta \in \mathbb{R}$. The *queue length* of edge e at time θ is given by

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

where $F^+_{i,e}(\theta) := \int_0^\theta f^+_{i,e}(z) dz$ and $F^-_{i,e}(\theta) := \int_0^\theta f^-_{i,e}(z) dz$ denote the *cumulative (edge) inflow* and *cumulative (edge) outflow*. For 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; their cumulative variants are denoted as F^+_e and F^-_e , respectively.

A *deterministic* flow over time f satisfies the following conditions (2) and (3). We assume that the queue operates at capacity which can be modeled by requiring

$$f^-_e(\theta) = \begin{cases} \nu_e, & \text{if } q_e(\theta - \tau_e) > 0, \\ \min\{f^+_e(\theta - \tau_e), \nu_e\}, & \text{else,} \end{cases} \quad (2)$$

for all $e \in E$, $\theta \in \mathbb{R}$. Moreover, we want the flow to follow a FIFO principle on the queues, which can be formalized by

$$f_{i,e}^-(\theta) = \begin{cases} f_e^-(\theta) \cdot \frac{f_{i,e}^+(\vartheta)}{f_e^+(\vartheta)}, & \text{if } f_e^+(\vartheta) > 0, \\ 0, & \text{else,} \end{cases} \quad (3)$$

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 while $\frac{q_e(\vartheta)}{\nu_e}$ is the current waiting time to be spent in the queue of edge e . Consequently, constraint (3) 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.

A *feasible flow* is a deterministic flow that also fulfills the *flow conservation constraints* (4) and (5). These are modeled for commodity i and node $v \neq t_i$ as

$$\sum_{e \in \delta_v^+} f_{i,e}^+(\theta) - \sum_{e \in \delta_v^-} f_{i,e}^-(\theta) = \begin{cases} u_i(\theta), & \text{if } v = s_i, \\ 0, & \text{if } v \neq s_i, \end{cases} \quad (4)$$

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 t_i we require

$$\sum_{e \in \delta_{t_i}^+} f_{i,e}^+(\theta) - \sum_{e \in \delta_{t_i}^-} f_{i,e}^-(\theta) \leq 0 \quad \text{for all } \theta \in \mathbb{R}_{\geq 0}. \quad (5)$$

Cominetti et al. (2015) analysed the queue dynamics of deterministic flows f : They show that the queue length is given by $q_e(\theta) = \max_{\xi \leq \theta} \int_{\xi}^{\theta} f_e^+ - \nu_e \, d\lambda$ for any $\theta \in \mathbb{R}$, $e \in E$. Thus, queues are independent of the edge outflow rates f_e^- . Moreover, it can be shown that given a set of inflow rates $f^+ = (f_{i,e}^+)_{i,e} \in \mathcal{R}^{I \times E}$ there exists a unique (up to a Lebesgue-null set) family of outflow rates $f^- = (f_{i,e}^-)_{i,e} \in \mathcal{R}^{I \times E}$, such that $f := (f^+, f^-)$ is deterministic; we call f the *deterministic flow with inflow rates* f^+ . If the inflow rates of any two deterministic flows f and f' into an edge e match almost everywhere up to some time H , the corresponding queue length q_e and exit time $T_e(\theta) := \theta + \frac{q_e(\theta)}{\nu_e} + \tau_e$ of both flows match up to time H , and the outflow rates out of edge e match almost everywhere up to time $T_e(H)$.

2.1 Instantaneous Dynamic Equilibrium

For an instantaneous dynamic equilibrium (IDE) as defined by Graf et al. (2020) we assume that, whenever an agent arrives at an intermediate node $v \in V_i$ 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 $c_e(\theta) := \frac{q_e(\theta)}{\nu_e} + \tau_e$. 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

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

We say that edge $e = vw \in E_i$ 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 $i \in I, \theta \in \mathbb{R}$ and $e \in E$ it satisfies*

$$f_{i,e}^+(\theta) > 0 \implies e \in E_i(\theta). \quad (7)$$

2.2 Dynamic Nash Equilibrium

In contrast, in the full information model we assume that agents have complete knowledge of the entire (future) evolution of the flow over time. If an agent enters an edge $e = vw$ at 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 et al., 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 IDE. So, for any $i \in I, v \in V$ and $\theta \in \mathbb{R}$ 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 $\theta \in \mathbb{R}$ it holds that*

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

3. 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 relax this behavioral assumption by introducing a model wherein every commodity $i \in I$ maintains a predictor $\hat{q}_{i,e}$ for every edge $e \in E$. The value $\hat{q}_{i,e}(\theta, \bar{\theta}, f)$ is then the queue length at time θ on edge e as predicted by commodity i at time $\bar{\theta}$ using the historical flow over time f . Formally, a predictor $\hat{q}_{i,e}$ has the following signature:

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

In general, such a predictor can depend in any arbitrary way on the entire input data including, in particular, the future evolution of the flow after the prediction time $\bar{\theta}$. However, for our theoretical results we require the predictors to behave in a slightly more restricted way. The first property we require is that the predictors do not use (and, therefore, do not need) any information on the future evolution of the queues.

Notation 3 For a vector of functions $g = (g_i)_{i \in [d]}$ with $g_i : \mathbb{R} \rightarrow \mathbb{R}$ and some $H \in \mathbb{R}$, we write $g_{\leq H} := (g_i|_{(-\infty, H]})_{i \in [d]}$ for the restriction of the functions to $(-\infty, H]$.

For two vectors of functions $g = (g_i)_{i \in [d]}$, $h = (h_i)_{i \in [d]}$ with $g_i, h_i : X \rightarrow \mathbb{R}$, $X \subseteq \mathbb{R}$, we write $g \stackrel{a.e.}{c.w.} h$ if $g_i \stackrel{a.e.}{c.w.} h_i$ holds for all $i \in [d]$.

Definition 4 A predictor $\hat{q}_{i,e}$ is called *causal*, if it fulfills the condition

$$f_{\leq \bar{\theta}} \stackrel{a.e.}{c.w.} f'_{\leq \bar{\theta}} \implies \hat{q}_{i,e}(\cdot, \bar{\theta}, f) = \hat{q}_{i,e}(\cdot, \bar{\theta}, f')$$

for all $\bar{\theta} \in \mathbb{R}$ and deterministic flows $f, f' \in (\mathcal{R} \times \mathcal{R})^{I \times E}$.

The next 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 $\hat{c}_{i,e}(\theta, \bar{\theta}, f) := \tau_e + \frac{\hat{q}_{i,e}(\theta, \bar{\theta}, f)}{\nu_e}$ and the predicted exit time of edge e is given by $\hat{T}_{i,e}(\theta, \bar{\theta}, f) := \theta + \hat{c}_{i,e}(\theta, \bar{\theta}, f)$. 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 , deterministic flow f and prediction time $\bar{\theta}$ the predicted exit time $\hat{T}_{i,e}(\cdot, \bar{\theta}, f)$ is non-decreasing.

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

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

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 $v \in V_i$ can reach t_i is given by

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

If all predictors respect FIFO, the label functions defined in (11) satisfy the equations

$$\hat{\ell}_{i,v}(\theta, \bar{\theta}, f) = \begin{cases} \theta, & \text{if } v = t_i, \\ \min_{vw \in \delta_v^+} \hat{\ell}_{i,w}(\hat{T}_{i,vw}(\theta, \bar{\theta}, f), \bar{\theta}, f), & \text{if } v \neq t_i. \end{cases} \quad (12)$$

Moreover, the *predicted delay for taking an edge* $e = vw \in E_i$ at time θ is then defined as

$$\hat{\Delta}_{i,e}(\theta, \bar{\theta}, f) := \hat{\ell}_{i,w}(\hat{T}_{i,e}(\theta, \bar{\theta}, f), \bar{\theta}, f) - \hat{\ell}_{i,v}(\theta, \bar{\theta}, f).$$

We say that an edge $e = vw \in E_i$ is $\bar{\theta}$ -estimated active for commodity i at time θ , if $\hat{\Delta}_{i,e}(\theta, \bar{\theta}, f) = 0$ holds true. Furthermore, let us denote the set of $\bar{\theta}$ -estimated active edges for commodity i at time θ by $\hat{E}_i(\theta, \bar{\theta}, f)$.

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 feasible 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

$$f_{i,e}^+(\theta) > 0 \implies e \in \hat{E}_i(\theta, \theta, f). \quad (13)$$

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

In other words, in a DPE, particles only follow predicted shortest paths, and particles predict a shortest path (according to their commodity's predictor) when they enter the network and every time they arrive at an intermediate node $v \neq t_i$.

4. Existence of Dynamic Prediction Equilibria

In this section we show that for causal predictors that respect FIFO and a suitable continuity property there always exists a dynamic prediction equilibrium. We will also give several examples of such predictors, including one for which any DPE corresponds to an IDE and vice versa.

4.1 Existence of DPE Using a Variational Inequality

To show the existence of DPE we make use of a result by Brézis (1968, Theorem 24) guaranteeing the existence of solutions to certain variational inequalities on a reflexive Banach space X . For this, we call a mapping $\mathcal{A} : K \rightarrow Y$ from $K \subseteq X$ to the normed space Y *sequentially weak-strong continuous* if \mathcal{A} maps every sequence $(u_i)_i$ in X that *converges weakly* to $u \in X$, i.e., $\lim_{i \rightarrow \infty} \langle g, u_i \rangle = \langle g, u \rangle$ for all $g \in X'$, to a strongly convergent sequence, i.e., $\lim_{i \rightarrow \infty} \|A(u_i) - A(u)\|_Y = 0$. Here, $\langle \cdot, \cdot \rangle$ denotes the canonical pairing between X and its continuous dual space X' . Cominetti et al. (2015, Section 5) have found the following useful special case of Brézis's theorem.

Theorem 7 Let X be a reflexive Banach space and let $\mathcal{A} : K \rightarrow X'$ be a sequentially weak-strong continuous map defined on a non-empty, closed, bounded and convex set $K \subseteq X$. Then there exists a solution $u \in X$ to the variational inequality

$$\forall v \in X : \quad \langle \mathcal{A}(u), v - u \rangle \geq 0. \quad (14)$$

Here, we apply this theorem on the reflexive Banach space $X = L^p(S)^d$ with S a compact set in \mathbb{R}^2 and $1 < p < \infty$. Its continuous dual space is isomorphic to $X' = L^q(S)^d$, where q is the conjugate of p with $1/p + 1/q = 1$. The canonical pairing is then given by $\langle g, f \rangle := \sum_{i \in [d]} \int_S g_i \cdot f_i \, d\lambda$ for $f \in X$ and $g \in X'$.

This theorem can be used to build up a dynamic prediction 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. A sufficient regularity condition on the predictors to apply the theorem is depicted in the following.

Definition 8 A predictor $\hat{q}_{i,e}$ is p -continuous for $1 < p < \infty$, if $\hat{q}_{i,e}(\cdot, \cdot, \cdot, f)$ is continuous for every deterministic flow f with $f^+ \in L_{\text{loc}}^p(\mathbb{R}, \mathbb{R}_{\geq 0})^{I \times E}$ and

$$L^p([0, M], \mathbb{R}_{\geq 0})^{I \times E} \rightarrow C([0, M] \times D, \mathbb{R}_{\geq 0}), \quad f^+ \mapsto \hat{q}_{i,e}(\cdot, \cdot, \cdot, f)$$

is sequentially weak-strong continuous for every $M > 0$ and compact interval D . Here, f denotes the deterministic flow with inflow rates $\mathbb{1}_{[0,M]} \cdot f^+$.

Now, we formally introduce the concept of partial dynamic prediction equilibria.

Definition 9 *A deterministic flow f is feasible up to time H if f fulfills the flow conservation conditions (4) and (5) for $\theta \leq H$. We call f a partial dynamic prediction flow with respect to a set of predictors \hat{q} up to time H if f is feasible up to time H and $f_{i,e}^+(\theta) > 0$ implies $e \in \hat{E}_i(\theta, \theta, f)$ for all $\theta \leq H, e \in E, i \in I$.*

We now show that the time horizon H of any partial dynamic prediction flow can be extended. We employ a proof-technique similar to the one used in (Graf et al., 2020, Lemma 5.6) for the extension property of IDE 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 10 *Let u_i be network inflow rates in $\mathcal{R} \cap L_{\text{loc}}^p(\mathbb{R}, \mathbb{R}_{\geq 0})$ and $\hat{q} = (\hat{q}_{i,e})_{i,e}$ be a set of p -continuous and causal predictors that respect FIFO. We can extend any partial dynamic prediction flow f with respect to \hat{q} up to time H to a dynamic prediction flow up to time $H + \tau_{\min}$ with $\tau_{\min} := \min_{e \in E} \tau_e$.*

The main idea is to first define a set K of all possible extensions of f . We then define a mapping $\mathcal{A} : K \rightarrow L^q(D)^{I \times E}$ associating with each possible extension g a function which for every commodity i , edge e and time θ returns the predicted delay $\hat{\Delta}_{i,e}(\theta, \theta, g)$. The p -continuity of the predictors implies that \mathcal{A} is sequentially weak-strong continuous such that Theorem 7 gives a solution to the variational inequality (14). This solution is an extension of f which satisfies the properties of a partial dynamic prediction flow up to time $H + \tau_{\min}$.

In this section, we equip the space of continuous functions $C(S)$ on a compact set $S \subseteq \mathbb{R}^d$ with the topology induced by the uniform norm $\|g\|_{\infty} := \max_{x \in S} g(x)$.

Lemma 11 *If the predictors $(\hat{q}_{i,e})_{e \in E}$ are p -continuous, the mappings*

- (i) $f^+ \mapsto \hat{T}_{i,e}(\cdot, \cdot, f)$ for all $e \in E$,
- (ii) $f^+ \mapsto \hat{\ell}_i^P(\cdot, \cdot, f)$ for all simple paths P ,
- (iii) $f^+ \mapsto \hat{\ell}_{i,v}(\cdot, \cdot, f)$ for all nodes $v \in V_i$, and
- (iv) $f^+ \mapsto \hat{\Delta}_{i,e}(\cdot, \cdot, f)$ for all $e \in E_i$.

are sequentially weak-strong continuous from $L^p([0, M], \mathbb{R}_{\geq 0})^{I \times E}$ to $C([0, M] \times D)$ for $M > 0$ and compact intervals D . Here, f denotes the deterministic flow with inflow rates $\mathbb{1}_{[0,M]} \cdot f^+$.

Proof (i). The sequential weak-strong continuity of $f^+ \mapsto \hat{T}_{i,e}(\cdot, \cdot, f)$ directly follows from the p -continuity of $\hat{q}_{i,e}$.

(ii). For $f^+ \mapsto \hat{\ell}_i^P(\cdot, \cdot, f)$ we use induction on the length of the simple path P . In the base case, P is an empty path such that $f^+ \mapsto \hat{\ell}_i^P(\cdot, \cdot, f)$ is a constant (and thus continuous) map. Assume that P consists of some path P' and an additional final edge e

and let $(f^{+,k})_{k \in \mathbb{N}}$ be weakly converging to some f^+ . Using the triangle inequality, the term $\|\hat{\ell}_i^P(\cdot, \cdot, \cdot, f^k) - \hat{\ell}_i^P(\cdot, \cdot, \cdot, f)\|_\infty$ can be upper bounded by

$$\begin{aligned} \alpha^k + \beta^k &:= \max_{\substack{\theta \in [0, M] \\ \bar{\theta} \in D}} \left| \hat{T}_{i,e} \left(\hat{\ell}_i^{P'}(\theta, \bar{\theta}, f^k), \bar{\theta}, f^k \right) - \hat{T}_{i,e} \left(\hat{\ell}_i^{P'}(\theta, \bar{\theta}, f), \bar{\theta}, f \right) \right| \\ &\quad + \max_{\substack{\theta \in [0, M] \\ \bar{\theta} \in D}} \left| \hat{T}_{i,e} \left(\hat{\ell}_i^{P'}(\theta, \bar{\theta}, f^k), \bar{\theta}, f \right) - \hat{T}_{i,e} \left(\hat{\ell}_i^{P'}(\theta, \bar{\theta}, f), \bar{\theta}, f \right) \right|. \end{aligned}$$

By induction there exists an $N \in \mathbb{N}$ such that

$$\|\hat{\ell}_i^{P'}(\cdot, \cdot, \cdot, f^k) - \hat{\ell}_i^{P'}(\cdot, \cdot, \cdot, f)\|_\infty < 1$$

holds for all $k \geq N$. With $M' := \|\hat{\ell}_i^{P'}(\cdot, \cdot, \cdot, f)\|_{[0, M] \times D} + 1$, the value $\hat{\ell}_i^{P'}(\theta, \bar{\theta}, f^k)$ is contained in $[0, M']$ for all $k \geq N$, $\theta \in [0, M]$, and $\bar{\theta} \in D$. Moreover, $\mathbb{1}_{[0, M]} \cdot f^{+,k}$ converges weakly to $\mathbb{1}_{[0, M]} \cdot f^+$ in $L^p([0, M'])^{I \times E}$. Using (i) for M' we conclude $\alpha^k \rightarrow 0$. The term β_k vanishes because of the uniform continuity of $\hat{T}_{i,e}(\cdot, \cdot, \cdot, f)$ on the compact set $[0, M'] \times D$.

(iii). Let $(f^{+,k})_{k \in \mathbb{N}}$ weakly converge to f^+ in $L^p([0, M], \mathbb{R}_{\geq 0})^{I \times E}$. By (ii), the sequence $(\hat{\ell}_i^P(\cdot, \cdot, \cdot, f^k))_k$ uniformly converges to $\hat{\ell}_i^P(\cdot, \cdot, \cdot, f)$ on $C([0, M] \times D, \mathbb{R}_{\geq 0})$ for any simple v - t_i -path P . The minimum of these functions also converges uniformly.

(iv). This follows with the same techniques applied in (ii). \blacksquare

With this lemma we can now turn back to the proof of the extension-lemma:

Proof of Lemma 10 Let f be a partial dynamic prediction flow up to time H with respect to \hat{q} . The outflow rate on an edge e of any deterministic flow h whose inflow rates coincide with f up to time H are already uniquely determined up to time $T_e(H) \geq H + \tau_e$. Therefore, the rates $b_{i,v}^-(\theta) := \sum_{e \in \delta_v^-} h_{i,e}^-(\theta) + \mathbb{1}_{v=s_i} u_i(\theta)$ on the interval $D := [H, H + \tau_{\min}]$ are independent of the edge inflow rates $h_{i,e}^+|_D$ on D .

We now want to apply Brézis's theorem in the form of Theorem 7 to find suitable inflow rates $h_{i,e}^+|_D$. For that, we define the set $K \subseteq L^p(D)^{I \times E}$ as follows:

$$K := \left\{ g \in L^p(D, \mathbb{R}_{\geq 0})^{I \times E} \left| \begin{array}{l} \forall i \in I, v \in V \setminus \{t_i\} : \sum_{e \in \delta_v^+} g_{i,e} \stackrel{a.e.}{=} b_{i,v}^- \\ \forall i \in I : \sum_{e \in \delta_{t_i}^+} g_{i,e} \leq b_{i,t_i}^- \end{array} \right. \right\}.$$

The elements of K are the possible inflow rates for the interval D that fulfill the flow conservation constraints. More specifically, for any $g \in K$, let \bar{g} denote the unique deterministic flow with inflow rates

$$\bar{g}_{i,e}^+(\theta) := \begin{cases} g_{i,e}(\theta), & \text{if } \theta \in D, \\ f_{i,e}^+(\theta), & \text{otherwise.} \end{cases}$$

Claim 12 For each $g \in K$, \bar{g} is a partial dynamic prediction flow up to time H and feasible up to time $H + \tau_{\min}$.

Proof As f is feasible up to time H and $g \in K$, it follows that \bar{g} is feasible up to time $H + \tau_{\min}$. Because all predictors are causal, the equilibrium property (13) is transferred from f over to g for $\theta < H$. \blacksquare

Claim 13 *The set K is nonempty, closed, bounded and convex.*

Proof By $u_i \in L^p_{\text{loc}}(\mathbb{R}, \mathbb{R}_{\geq 0})$ and $f_{i,e}^-(\theta) \leq \nu_e$, it follows that $b_{i,v}^-$ is p -integrable on D for all $v \in V$. Then, $g \in K$ implies $\|g_{i,e}\|_{L^p(D)} \leq \|b_{i,v}^-\|_{L^p(D)}$ for all $i \in I$ and $e = vw \in E$. Therefore, K is bounded in $L^p(D)^{I \times E}$. To see that K is nonempty, observe that, in a partial dynamic prediction flow, particles of commodity i can only arrive at a node v if t_i is reachable from v . Hence, for all $v \in V \setminus \{t_i\}$ with $b_{i,v}^-(\theta) > 0$ we select an arbitrary edge $e \in \delta_v^+$ and set $g_{i,e} = b_{i,v}^-$. For all other edges, we set $g_{i,e} = 0$. This implies $g \in K$. It can be easily verified that K is convex and closed. \blacksquare

By Claim 12, \bar{g} is a dynamic prediction flow up to time H for every $g \in K$. Therefore, we are looking for some $g \in K$ such that $\bar{g}_{i,e}^+(\theta) > 0$ implies $e \in \hat{E}_i(\theta, \theta, \bar{g})$ for $\theta \in D$. For an edge $e = vw \in E_i$ lying on a directed s_i - t_i -path, the statement $e \in \hat{E}_i(\theta, \theta, \bar{g})$ is equivalent to $\hat{\Delta}_{i,e}(\theta, \theta, \bar{g}) \leq 0$. Using this observation, we define $\mathcal{A} : K \rightarrow L^q(D)^{I \times E}$ with $1/p + 1/q = 1$ as the predicted delay operator when using an edge $e = vw$:

$$\mathcal{A}(g)_{i,e}(\theta) := \begin{cases} \hat{\Delta}_{i,e}(\theta, \theta, \bar{g}), & \text{if } e \in E_i, \\ 1, & \text{otherwise.} \end{cases}$$

We note that for edges $e \notin E_i$ that are irrelevant to a commodity i , we simply set the predicted delay to 1; however, any strictly positive constant would work. By the p -continuity of \hat{q} , the function $\mathcal{A}(g)_{i,e}$ is continuous as a function on \mathbb{R} . Thus, $\mathcal{A}(g)$ is indeed contained in $L^q(D)^{I \times E}$ as a function on D .

Claim 14 *The map \mathcal{A} is non-negative and sequentially weak-strong continuous.*

Proof The non-negativity of \mathcal{A} results from the predictors respecting FIFO.

To show that \mathcal{A} is sequentially weak-strong continuous, let $(g^k)_{k \in \mathbb{N}}$ be a sequence in $L^p(D)^{I \times E}$ that converges weakly to g^+ . Then, $\bar{g}^{k,+}|_{[0,M]}$ also converges weakly to $\bar{g}^+|_{[0,M]}$ in $L^p([0,M])^{I \times E}$ with $M := H + \tau_{\min}$. By Lemma 11 and the causality of the predictors, the sequence $\hat{\Delta}_{i,e}(\cdot, \cdot, \bar{g}^k)$ converges to $\hat{\Delta}_{i,e}(\cdot, \cdot, \bar{g})$ with respect to the uniform norm on $C([0,M] \times D)$ for any $e \in E_i$. With this we conclude that $\mathcal{A}(g^k)$ converges to $\mathcal{A}(g)$ in $C(D)^{I \times E}$, and hence also in $L^q(D)^{I \times E}$. \blacksquare

The variational inequality as given in Theorem 7 now admits a solution $g \in K$ such that for all $h \in K$ we have $\langle \mathcal{A}(g), h - g \rangle \geq 0$. We assume that the equilibrium property for \bar{g} does not hold almost everywhere on D , implying that the union of all sets

$$\Phi_{i,e} := \left\{ \theta \in D \mid \bar{g}_{i,e}^+(\theta) > 0, \mathcal{A}(g)_{i,e}(\theta) > 0 \right\}$$

has positive measure. We now construct flow rates $h \in K$ that lead to the contradiction $\langle \mathcal{A}(g), h - g \rangle < 0$. To do this, we observe that by the continuity of the maps $(\theta, \bar{\theta}) \mapsto \ell_{i,w}(\theta, \bar{\theta}, \bar{g})$ and $\theta \mapsto \hat{T}_{i,e}(\theta, \theta, \bar{g})$ the set of times $\Theta_{i,e} := \{\theta \in \mathbb{R} \mid e \in \hat{E}_{i,e}(\theta, \theta, \bar{g})\}$, at which an edge e is active, is closed and thus measurable. We define

$$h_{i,e}(\theta) := \frac{b_{i,v}^-(\theta)}{|\delta_v^+ \cap \hat{E}_{i,e}(\theta, \theta, \bar{g})|}$$

for $\theta \in \Theta_{i,e}$ and $h_{i,e}(\theta) := 0$ otherwise. As the set $\delta_v^+ \cap \hat{E}_i(\theta, \theta, \bar{g})$ is non-empty for any $\theta \in D$ and $v \in V_i \setminus \{t\}$, and as $b_{i,v}^-(\theta)$ can only be positive for nodes in V_i , h is an element of K .

By definition, $h_{i,e}(\theta)$ is only positive if e is active at time θ (w.r.t. \bar{g}) which implies $\mathcal{A}(g)_{i,e}(\theta) = 0$. Therefore, we conclude

$$\langle \mathcal{A}(g), h - g \rangle = \sum_{i,e} \int_D \mathcal{A}(g)_{i,e} \cdot h_{i,e} \, d\lambda - \langle \mathcal{A}(g), g \rangle = - \sum_{i,e} \int_D \mathcal{A}(g)_{i,e} \cdot g_{i,e} \, d\lambda.$$

Let $i \in I$ and $e \in E$ such that $\Phi_{i,e}$ has positive measure. Because both $g_{i,e}$ and $\mathcal{A}(g)_{i,e}$ are positive on $\Phi_{i,e}$, this implies $\int_{\Phi_{i,e}} \mathcal{A}(g)_{i,e} \cdot g_{i,e} \, d\lambda > 0$. The non-negativity of $\mathcal{A}(g)$ and g now implies

$$\langle \mathcal{A}(g), h - g \rangle \leq - \int_{\Phi_{i,e}} \mathcal{A}(g)_{i,e} \cdot g_{i,e} \, d\lambda < 0,$$

a contradiction. ■

With this, we can show the existence of dynamic prediction flows for causal and p -continuous predictors respecting FIFO.

Starting with the zero-flow with horizon 0 and iteratively extending it using Lemma 10, we obtain a dynamic prediction flow, as shown in the following theorem. Note, that this result allows different commodities to use different sets of predictors.

Theorem 15 *For any network with a finite set of commodities, each associated with a network inflow rate in $L_{\text{loc}}^p(\mathbb{R}, \mathbb{R}_{\geq 0})$ and causal and p -continuous predictors $\hat{q}_{i,e}$ that respect FIFO, there exists a dynamic prediction flow with respect to \hat{q} .*

Proof Let $f^0 \in (\mathcal{R} \times \mathcal{R})^{I \times E}$ with $f_{i,e}^{0,+}, f_{i,e}^{0,-} := 0$ for all $i \in I, e \in E$. Then f^0 is a partial dynamic prediction flow up to time 0. Recursively define f^{n+1} as the dynamic prediction flow up to time $\tau_{\min} \cdot (n+1)$ given through the extension of f^n as per Lemma 10. It follows that $f_{\leq k \cdot \tau_{\min}}^n \stackrel{\text{a.e.}}{\underset{\text{c.w.}}{=}} f_{\leq k \cdot \tau_{\min}}^k$ whenever $n \geq k$.

We define the flow f^∞ by $f_{i,e}^{\infty,+}(\theta) := f_{i,e}^{n_\theta,+}(\theta)$ and $f_{i,e}^{\infty,-}(\theta) := f_{i,e}^{n_\theta,-}(\theta)$ where $n_\theta := \lceil \theta / \tau_{\min} \rceil$. Then f^∞ is a dynamic prediction flow with respect to \hat{q} up to time ∞ : The feasibility and the equilibrium properties are checked by using $f_{\leq n \cdot \tau_{\min}}^\infty \stackrel{\text{a.e.}}{\underset{\text{c.w.}}{=}} f_{\leq n \cdot \tau_{\min}}^n$, the fact that f^n fulfills the desired properties for all $n \in \mathbb{N}$, and the causality of the predictors. ■

Example 1 *To see why we require the predictors to be continuous, consider the non-continuous predictor*

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

Using this predictor in a network consisting of only a single source-sink pair connected by two parallel edges e_1 and e_2 (see Figure 1) 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. Assume a dynamic prediction flow f exists. Then, clearly, during the time




Figure 1: A network where the use of a non-continuous predictor can result in a situation where no dynamic prediction equilibrium exists. Edges are labeled with (τ_e, ν_e) .

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), and thus $q_{e_1}^f(1) = 1$. Let $O \subseteq [1, \infty]$ denote the open set containing all times $\theta \geq 1$ with $q_{e_1}^f(\theta) < 1$, and let $O = \bigcup_{j \in J} (a_j, b_j)$ be a partition of O into at most countably many disjoint, non-empty, open intervals. For every $j \in J$, we have $a_j \notin O$, and thus $q_{e_1}^f(a_j) = 1$ holds by the continuity of $q_{e_1}^f$. During (a_j, b_j) , e_1 is the only active edge, and thus, for any $\theta \in (a_j, b_j)$, it follows $1 > q_{e_1}^f(\theta) = q_{e_1}^f(a_j) + (\theta - a_j) > 1$. Thus, O is the empty set. This means, starting from time 1 all particles would choose edge e_2 , implying $q_{e_1}^f(\theta) < 1$ for all $\theta > 1$; a contradiction.

4.2 A Sufficient Regularity Condition

To make the previous theorem more accessible, we give a sufficient, less abstract regularity condition for p -continuity.

Definition 16 A predictor $\hat{q}_{i,e}$ depends continuously on the cumulative inflow, total cumulative outflow and queue lengths, if there exists a continuous map

$$\gamma_{i,e} : \mathbb{R} \times \mathbb{R} \times C(\mathbb{R}, \mathbb{R}_{\geq 0})^{(I \times E) + E + E} \rightarrow \mathbb{R}_{\geq 0}$$

such that $\hat{q}_{i,e}(\theta, \bar{\theta}, f) = \gamma_{i,e}(\theta, \bar{\theta}, F_{I \times E}^{+,f}, F_E^{-,f}, q^f)$ for all $\theta, \bar{\theta} \in \mathbb{R}$ and deterministic flows f . Here, $F_{I \times E}^{+,f} := (F_{i,e}^{+,f})_{i \in I, e \in E}$, $F_E^{-,f} := (F_e^{-,f})_{e \in E}$, and $q^f := (q_e^f)_{e \in E}$ denote the cumulative edge in- and outflow functions, and queue length functions with respect to f .

Lemma 17 A predictor $\hat{q}_{i,e}$ that depends continuously on the cumulative inflow, total cumulative outflow and queue lengths, is p -continuous for any p with $1 < p < \infty$.

Proof We show that $f^+ \mapsto \hat{q}_{i,e}(\cdot, \cdot, f)$ is sequentially weak-strong continuous from $L^p([0, M] \times D, \mathbb{R}_{\geq 0})^{I \times E}$ to $C([0, M] \times D)$ for all M and compact intervals D .

Cominetti et al. (2015, Lemma 16) showed that the maps $f_{i,e}^+ \mapsto F_{i,e}^{+,f}$ and $f_e^+ \mapsto q_e^f$ are sequentially weak-strong continuous from $L^p([0, M], \mathbb{R}_{\geq 0})$ to $C([0, M], \mathbb{R}_{\geq 0})$ for any $M > 0$. Thus, the same holds for $f_e^+ \mapsto (\theta \mapsto F_e^{+,f}(\theta - \tau_e) - q_e^f(\theta - \tau_e)) = F_e^{-,f}$.

Let $(f^{+,k})_{k \in \mathbb{N}}$ be a sequence converging weakly to f^+ in $L^p([0, M])^{I \times E}$. Then, the sequence $(f_e^{+,k})_{k \in \mathbb{N}}$ converges weakly to f_e^+ and $(f_{i,e}^{+,k})_{k \in \mathbb{N}}$ converges weakly to $f_{i,e}^+$ for every $e \in E$ and $i \in I$. By Cominetti et al. (2015), the sequences $(F_{i,e}^{+,k})_{k \in \mathbb{N}}$, $(F_e^{-,k})_{k \in \mathbb{N}}$ and $(q_e^k)_{k \in \mathbb{N}}$ converge strongly to $F_{i,e}^+$, F_e^- and q_e , respectively, in $C([0, M])$ for all $e \in E$ and $i \in I$. This implies that these sequences also converge in $C(\mathbb{R}, \mathbb{R}_{\geq 0})$.

We write $g^k := (F_{I \times E}^{+,k}, F_E^{-,k}, q_E^k)$ for $k \in \mathbb{N}$ and $g := (F_{I \times E}^+, F_E^-, q_E)$. The sequence g^k converges to g in $\mathcal{G} := C(\mathbb{R}, \mathbb{R}_{\geq 0})^{(I \times E) + E + E}$. Let $\varepsilon > 0$ be arbitrary. By the compactness of $[0, M] \times D$ and the continuity of $\gamma_{i,e}$, there exists some δ such that for all $(\theta, \bar{\theta}), (\theta', \bar{\theta}') \in [0, M] \times D$ and $g' \in \mathcal{G}$, the condition $\|g - g'\|_\infty < \delta$ implies $|\gamma_{i,e}(\theta, \bar{\theta}, g) - \gamma_{i,e}(\theta', \bar{\theta}', g')| < \varepsilon/2$. There is $N \in \mathbb{N}$ with $\|g^k - g\|_\infty < \delta$ for all $k \geq N$. For arbitrary $\theta \in [0, M]$, $\bar{\theta} \in D$ and $k \geq N$ we infer

$$|\gamma_{i,e}(\theta, \bar{\theta}, g^k) - \gamma_{i,e}(\theta, \bar{\theta}, g)| < \varepsilon/2,$$

which implies the desired result $\|\gamma_{i,e}(\cdot, \cdot, g^k) - \gamma_{i,e}(\cdot, \cdot, g)\|_{C([0,M] \times D)} < \varepsilon$. \blacksquare

4.3 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.,

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

This predictor is trivially continuous, causal, and respects FIFO. In the resulting dynamic prediction flow, particles just always follow physically shortest paths.

The *constant predictor* predicts that all queues will stay constant:

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

This leads to the mentioned special case of IDE flows. Since the constant predictor clearly is causal, depends continuously on the queue lengths and respects FIFO, we can apply Theorem 15 and, thus, reprove the existence of IDE flows shown in Graf et al. (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

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

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 $\partial_- q_e^f(\bar{\theta})$ might be discontinuous.

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

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

Using Lemma 17 it is easy to show that \hat{q}^{RL} respects FIFO, is causal and p -continuous for any $p > 0$. It, thus, induces the existence of a DPE.

The regularized linear predictor takes two samples of the past queue lengths (at times $\bar{\theta}$ and $\bar{\theta} - \delta$) and uses this information to predict future queue lengths up to the prediction horizon by a pw. linear function. We generalize this idea by taking a fixed, finite set of observations of the past flow and employ a continuous transformation on these values to

predict the future queues: Let k_p be the number of past samples, k_f the number of future samples, and $\delta > 0$ be some step size. The arguments of the function $\gamma_{i,e}$ from Definition 16 are mapped to \mathbb{R}^d using an observation map σ . This map returns k_p samples of the queue length $q_{e'}^f$ and the edge load $L_{e'}^f := F_{e'}^{+,f} - F_{e'}^{-,f}$ at the past times $\bar{\theta} - (i-1)\delta$ for every $i \in [k_p]$ and $e' \in E$. Then, a continuous (machine-learned) transformation $\phi_{i,e}$ is applied resulting in interpolation points $\hat{q}_{i,e}^{\text{ML,raw}}(\bar{\theta} + j\delta, \bar{\theta}, f) \geq 0$ for $j \in [k_f]$. These two maps are depicted in the following:

$$\begin{array}{c}
 \theta, \bar{\theta}, \\
 F_{I \times E}^{+,f}, F_E^{-,f}, q^f
 \end{array}
 \xrightarrow{\sigma}
 \begin{array}{c}
 \bar{\theta}, \\
 \left(q_{e'}^f(\bar{\theta} - (j-1)\delta) \right)_{\substack{e' \in E \\ j \in [k_p]}}, \\
 \left(L_{e'}^f(\bar{\theta} - (j-1)\delta) \right)_{\substack{e' \in E \\ j \in [k_p]}}
 \end{array}
 \xrightarrow{\phi_{i,e}}
 \left(\hat{q}_{i,e}^{\text{ML,raw}}(\bar{\theta} + j\delta, \bar{\theta}, f) \right)_{j \in [k_f]}$$

Finally, these points are linearly interpolated (in the first argument) to determine the value $\hat{q}_{i,e}^{\text{ML,raw}}(\theta, \bar{\theta}, f)$. This results in a causal, and p -continuous predictor. However, $\hat{q}_{i,e}^{\text{ML,raw}}$ could violate FIFO. To prevent that, we apply the following post-processing: We derive new interpolation points with $\hat{q}_{i,e}^{\text{ML}}(\bar{\theta}, \bar{\theta}, f) = q_e^f(\bar{\theta})$ and

$$\hat{q}_{i,e}^{\text{ML}}(\bar{\theta} + j\delta, \bar{\theta}, f) := \max \left\{ \begin{array}{l} \hat{q}_{i,e}^{\text{ML,raw}}(\bar{\theta} + j\delta, \bar{\theta}, f), \\ \hat{q}_{i,e}^{\text{ML}}(\bar{\theta} + (j-1)\delta, \bar{\theta}, f) - \nu_e \delta \end{array} \right\}$$

for all $j \in [k_f]$. We extend the predictor constantly outside the interpolation points. We call the resulting predictor $\hat{q}_{i,e}^{\text{ML}}$ the *machine-learned predictor*.

Proposition 18 *If all maps $\phi_{i,e}$ are continuous and non-negative, the predictor \hat{q}^{ML} is causal, respects FIFO, and depends continuously on the queue length functions. It, thus, induces the existence of a dynamic prediction equilibrium.*

This leaves the question of how to choose the transformation maps in order to retrieve a good predictor. In our experimental study, we use both a linear regression model and a simple neural network and train these models on previously computed flows. We provide more details in Section 6.3.

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

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

This predictor clearly is not causal 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 Nash equilibria for which existence has been proven in Cominetti et al. (2015).

5. Approximate Dynamic Prediction Equilibria

We augment the introduction of DPE by an analysis of the computation of an approximate variant of DPE. We recall that in a DPE, flow is only sent along predicted shortest paths. More specifically, flow enters an edge only if the predicted delay $\hat{\Delta}_{i,e}(\theta, \theta, f)$ vanishes. This inspires the following definition of approximate DPE, in which agents enter an edge only if the predicted delay is small enough.

Definition 19 A pair (\hat{q}, f) of predictors \hat{q} and a flow over time f is called an ε -approximate DPE up to time $H \in \mathbb{R}_{>0} \cup \{\infty\}$, if f is feasible up to time H and for all $i \in I$, $e \in E$, and $\theta < H$ it holds that

$$f_{i,e}^+(\theta) > 0 \implies e \in E_i \wedge \hat{\Delta}_{i,e}(\theta, \theta, f) \leq \varepsilon.$$

The goal of this section is to develop an algorithm for computing ε -approximate DPE provided that all predictors are causal, respect FIFO, fulfill a regularity condition, and are computable as piecewise linear functions. We start off with an algorithm that computes so-called δ -routed DPE in which route updates happen only at discrete time steps of size $\delta > 0$. After an outline of the algorithm in Section 5.1, Section 5.2 points out how to compute outflow rates and queue lengths given piecewise constant inflow rates. In Section 5.3 we prove the correctness and termination of the algorithm. Finally, Section 5.4 shows that by imposing a Lipschitz-condition on the predictors, any δ -routed equilibrium flow is also an ε -approximate DPE.

5.1 Extension-Based Approximation Algorithm

As agents in a DPE behave like infinitesimal particles in a continuous flow, the routes chosen may change indefinitely often. For computational analyses it is infeasible to calculate shortest paths this often. Hence, we discretize the points in time at which routes are updated.

More specifically, shortest paths are recalculated every δ time units for some $\delta > 0$. This means that particles arriving at an intermediate node v at some time θ will use a predicted shortest path computed at time

$$\vartheta_\delta(\theta) := \left\lfloor \frac{\theta}{\delta} \right\rfloor \cdot \delta,$$

leading to the following variant of DPE whose computation we will focus on for the moment.

Definition 20 Let f be a dynamic flow, $\hat{q} = (\hat{q}_{i,e})_{i,e}$ a set of predictors and $\delta > 0$. The pair (f, \hat{q}) is called a δ -routed DPE up to time $H \in \mathbb{R} \cup \{\infty\}$ if f is feasible up to time H and for all $e \in E$, $i \in I$ and almost all $\theta < H$ it holds that

$$f_{i,e}^+(\theta) > 0 \implies e \in \hat{E}_i(\vartheta_\delta(\theta), \vartheta_\delta(\theta), f).$$

We require that network inflow rates are piecewise (*pw.*) constant with finitely many jumps. This enables us to find a δ -routed DPE with pw. constant edge inflow and outflow rates. As a result, cumulative edge inflow and outflow functions, queue length and exit time functions of edges and paths, as well as earliest arrival time functions are pw. linear. Consequently, we assume that the predicted queue length functions $\hat{q}_{i,e}(\cdot, \bar{\theta}, f)$ of all predictors are pw. linear.

The predictors in our simulation must be causal and respect FIFO. The first requirement is necessary, as we build our flow in an extension based approach. The second allows us to use shortest path algorithms for FIFO cost functions.

In the following we give an overview on the computation of a δ -routed DPE. Initially, we begin with a “zero”-flow f with $f^+, f^- \equiv 0$ as our δ -routed DPE flow up to time $H = 0$. We call H the *flow horizon* of f . We aim to compute a δ -routed DPE flow up to

some $H_{\text{comp}} \in \mathbb{R}$. The algorithm consists of two different phases that need to be repeated multiple times: A *routing phase* and a *distribution phase*.

A *routing phase* is run when the δ -routed DPE flow f has been calculated up to some multiple $H = k \cdot \delta$ of δ . All routes that have been calculated up to time H are invalidated, and new routes are determined. This is done in the following two steps:

- (R1) Gather predictions $\hat{q}_{i,e}(\cdot, H, f)$ as pw. linear functions for all commodities $i \in I$ and edges $e \in E$.
- (R2) Compute the set of active outgoing edges $\delta_v^+ \cap \hat{E}_i(H, H, f)$ for all nodes $v \in V_i$ and commodities $i \in I$.

In the *distribution phase* the dynamic flow gets extended. Assume we have computed a δ -routed DPE flow f up to flow horizon H and want to extend H up to some time $H' > H$. To ensure flow conservation, we distribute the *node inflow rate* $b_{i,v}^-(\theta)$ of $v \in V_i$ to outgoing edges. A distribution phase consists of the following steps.

- (D1) Update f with the deterministic flow with respect to the updated inflow rates

$$f_{i,e}^+|_{[H,\infty)} := \begin{cases} \frac{b_{i,v}^-(H)}{|\delta_v^+ \cap \hat{E}_i(\vartheta_\delta(H), \vartheta_\delta(H), f)|}, & \text{if } e \in \hat{E}_i(\vartheta_\delta(H), \vartheta_\delta(H), f), \\ 0, & \text{otherwise,} \end{cases}$$

for all commodities $i \in I$ and edges $e = vw \in E$.

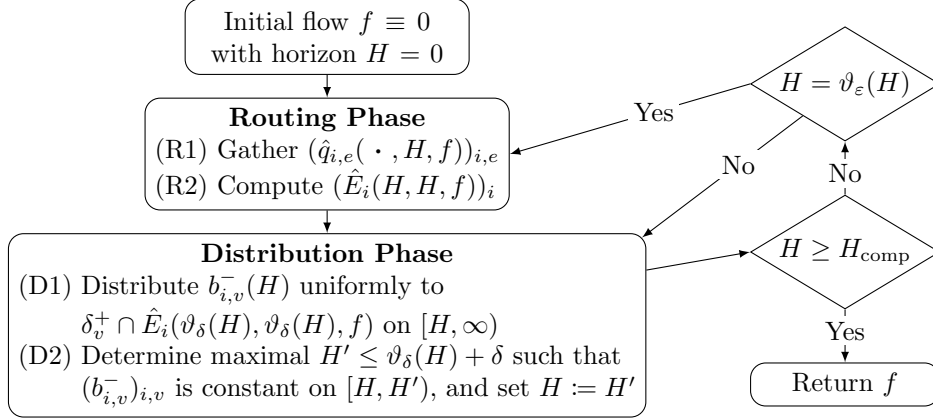
- (D2) Determine the maximal $H' \leq \vartheta_\delta(H) + \delta$ such that $b_{i,v}^-$ is constant on $[H, H')$ (according to the updated deterministic flow) for every node v and commodity $i \in I$, and set $H := H'$.

The distribution phase is executed multiple times until the flow horizon H reaches the next multiple of δ ; then a new routing phase is initiated. Once the flow horizon H reaches the computation target H_{comp} we terminate the process and return the computed flow f . A schematic overview of this algorithm is depicted in Figure 2.

REMARKS ON THE IMPLEMENTATION

There are three major aspects to consider when implementing the algorithm. Firstly, while it is cheap to compute the predicted queue length functions $\hat{q}_{i,e}(\cdot, H, f)$ for the different causal predictors introduced in Section 4.3, it is computationally expensive to compute the set of active edges $\hat{E}_i(H, H, f)$. Secondly, a difficulty in the distribution phase lies in computing the outflow rates when extending the inflow rates with a new constant as in Step (D1). Thirdly, it is unclear how to find the maximal time H' such that $b_{i,v}^-$ is constant on $[H, H')$ quickly. It is a good idea to solve the latter two problems in combination, as any change to an edge outflow rate $f_{i,vw}^-$ (induced by changing edge inflow rates) results in a change of the target node's inflow rate $b_{i,w}^-$. The solution to these two problems is discussed in Sections 5.2 and 5.3.

Computing the active edges $\hat{E}_i(H, H, f)$ can be solved by a dynamic variant of the Bellman-Ford algorithm covered by Ding et al. (2008). However, because of its poor running


 Figure 2: A schematic overview of the computation of a δ -routed DPE.

time, it is beneficial to use a Dijkstra-based algorithm. With this approach we have to calculate the active outgoing edges for each node and commodity with a separate run of the algorithm. Therefore, we defer this calculation until we notice a node v in the distribution phase with positive inflow rate $b_{i,v}^-$ and only calculate the active outgoing edges of v on demand. In most cases, only a small subset of nodes will experience inflow of any given commodity; in our experiments this proved to be much faster than the Dynamic Bellman-Ford algorithm.

Nevertheless, for predictors \hat{q}_i that are constant in the first argument, i.e. if $\hat{q}_{i,e}(\cdot, \bar{\theta}, f)$ is constant for all $e \in E$, $\bar{\theta} \in \mathbb{R}$ and deterministic flows f , a single run of a simple static version of Dijkstra's algorithm suffices to compute the set of all active edges. For predictors of this type, e.g., the Zero-predictor and the constant predictor, this approach is the preferred method to compute $\hat{E}_i(H, H, f)$.

An implementation of this algorithm using Python is publicly available in Graf et al. (2022a).

5.2 Outflow Rates of Piecewise Constant Inflow Rates

The idea of the introduced algorithm is to utilize the simple structure of flows over time generated by pw. constant inflow rates. This is demonstrated by the following theorem.

Theorem 21 *Let f be a deterministic flow, let $(g_{i,e})_{i \in I} \in \mathbb{R}_{\geq 0}^I$ be a set of new constant inflow rates into an edge e beginning from time $H \in \mathbb{R}$, and let $g_e := \sum_{i \in I} g_{i,e}$. Let $h = (h^-, h^+)$ denote the deterministic flow with inflow rates*

$$h_{i,e}^+(\theta) := \begin{cases} f_{i,e}^+(\theta), & \text{for } \theta < H, \\ g_{i,e}, & \text{for } \theta \geq H, \end{cases}$$

Then $h_{i,e}^-$ is given by $h_{i,e}^-|_{(-\infty, T_e^f(H))} \stackrel{\text{a.e.}}{=} f_{i,e}^-|_{(-\infty, T_e^f(H))}$ and by the following cases:

(C1): $g_e = 0$. Then $h_{i,e}^-|_{[T_e^f(H), \infty)} \stackrel{\text{a.e.}}{=} 0$.

(C2): $g_e > 0 \wedge (q_e^f(H) = 0 \vee g_e \geq \nu_e)$. Then $h_{i,e}^-|_{[T_e^f(H), \infty)} \stackrel{\text{a.e.}}{=} \min\{\nu_e, g_e\} \cdot \frac{g_{i,e}}{g_e}$.

(C3): $g_e \in (0, \nu_e) \wedge q_e^f(H) > 0$. Then, with $T_{\text{depl}} := H + \frac{q_e^f(H)}{\nu_e - g_e}$,

$$h_{i,e}^-|_{[T_e^f(H), T_{\text{depl}} + \tau_e]} \stackrel{\text{a.e.}}{\equiv} \nu_e \cdot \frac{g_{i,e}}{g_e} \quad \text{and} \quad h_{i,e}^-|_{[T_{\text{depl}} + \tau_e, \infty)} \stackrel{\text{a.e.}}{\equiv} g_{i,e}.$$

Proof As the inflow rates of f and h into edge e match up to time H , the queue and exit time functions of edge e coincide up to time H , and the outflow rate functions match up to time $T_e(H)$. For $\theta > T_e(H)$, we handle all three cases individually and utilize the observation

$$q_e^h(\theta - \tau_e) = q_e(H) + g_e \cdot (\theta - \tau_e - H) - \int_{H + \tau_e}^{\theta} h_e^- d\lambda.$$

(C1). From $g_e = 0$ and (2) it follows $q_e^h(\theta - \tau_e) \leq q_e(H) - \int_{H + \tau_e}^{T_e(H)} h_e^- d\lambda = 0$ for all $\theta > T_e(H)$. Applying (2) again yields $h_{i,e}^-(\theta) = 0$ for $\theta > T_e(H)$.

(C2). We show $h_e^-(\theta) = \min\{\nu_e, g_e\}$ for almost all $\theta > T_e(H)$. Then, since any ξ fulfilling $T_e^h(\xi) = \theta$ is greater than H , property (3) leads to the conclusion $h_{i,e}^-(\theta) = \min\{\nu_e, g_e\} \cdot \frac{g_{i,e}}{g_e}$ for almost all $\theta > T_e(H)$.

In the case $g_e > \nu_e$, for any $\theta > T_e(H)$, the queue length fulfills $q_e^h(\theta - \tau_e) > 0$. Therefore $h_e^-(\theta) = \nu_e$ holds for almost all $\theta > T_e(H)$. For $g_e = \nu_e$ and $q_e(H) > 0$, the same reasoning applies. Now assume $g_e \in (0, \nu_e)$ and $q_e(H) = 0$. For any $\theta > T_e(H)$, we have

$$q_e^h(\theta - \tau_e) = \max_{\xi \in [H, \theta - \tau_e]} \int_{\xi}^{\theta - \tau_e} h_e^+ - \nu_e d\lambda = 0.$$

Constraint (2) yields $h_e^-(\theta) = g_e$ for almost all $\theta > T_e(H)$.

(C3). It suffices to prove $h_e^-|_{(T_e(H), T_{\text{depl}} + \tau_e)} \stackrel{\text{a.e.}}{\equiv} \nu_e$ and $h_e^-|_{(T_{\text{depl}} + \tau_e, \infty)} \stackrel{\text{a.e.}}{\equiv} g_e$. For $\theta \in (T_e(H), T_{\text{depl}} + \tau_e)$ we have

$$q_e^h(\theta - \tau_e) \geq q_e(H) - (\nu_e - g_e) \cdot (\theta - \tau_e - H) > 0,$$

implying $h_e^-(\theta) = \nu_e$ for almost all of these θ . Therefore, T_{depl} fulfills $q_e^h(T_{\text{depl}}) = q_e^h(H) - (\nu_e - g_e) \cdot (T_{\text{depl}} - H) = 0$. For $\theta > T_{\text{depl}} + \tau_e$, we compute

$$q_e^h(\theta - \tau_e) = \max_{\xi \in [T_{\text{depl}}, \theta - \tau_e]} \int_{\xi}^{\theta - \tau_e} h_e^+ - \nu_e d\lambda = 0,$$

implying $h_e^-(\theta) = g_e$. ■

Remark 22 *The proof also describes the evolution of the queue length starting at time H . In the cases (C1) and (C3) we have*

$$q_e^h(\theta) = \begin{cases} q_e^f(H) - (\theta - H) \cdot (\nu_e - g_e), & \text{for } \theta \in [H, T_{\text{depl}}], \\ 0, & \text{for } \theta \geq T_{\text{depl}}, \end{cases}$$

where $T_{\text{depl}} := H + \frac{q_e^f(H)}{\nu_e - g_e}$ denotes the depletion time. For (C2), no queue depletion occurs and for $\theta \geq H$ we have

$$q_e(\theta) = q_e(H) + (\theta - H) \cdot \max\{g_e - \nu_e, 0\}.$$

5.3 Correctness and Termination

The previous section allows us to calculate the deterministic outflow rates for the case that new constant inflow rates are assigned to an edge as in Step (D1): In the implementation, we maintain pw. constant functions $f_{i,e}^+$ and $f_{i,e}^-$ for all $i \in I$ and $e \in E$, as well as pw. linear functions q_e for all $e \in E$. Before and after each distribution phase, these functions represent a deterministic flow that is feasible up to time H . In Step (D1) the inflow rates are then updated starting from time H and the outflow rates as well as the queue length functions are updated according to Theorem 21 and Remark 22 starting from times $T_e(H)$ and H , respectively, while preserving the deterministic property of f .

To detect the maximum H' such that $b_{i,v}^-$ stays constant on $[H, H')$ for all $i \in I, v \in V$ in Step (D2), we employ a priority queue that contains all events of the form

(E1) “The outflow rate $f_{i,e}^-$ changes at time θ ” for all $i \in I, e \in E$ and $\theta > H$, and

(E2) “The network inflow rate u_i changes at time θ ” for all $i \in I$ and $\theta > H$.

To achieve this, whenever we update an edge’s outflow rate $f_{i,e}^-$ in Step (D1), we generate corresponding events of type (E1) and possibly remove existing events, that are rendered invalid by the update. As we require $\tau_e > 0$, the added events happen after H . The events of type (E2) are enqueued in the initialization of the algorithm.

A change in a node inflow rate $b_{i,v}^-$ requires that either the network inflow rate has changed (and $v = s_i$) or that an outflow rate of an incoming edge has changed. Thus, the minimum time θ of all events in the priority queue is a lower bound to the maximum time H' such that $b_{i,v}^-|_{[H,H']}$ is constant for all $i \in I$ and $v \in V$. If there is only a single event in the queue at time θ , then $b_{i,v}^-$ changes at time θ , and we have $H' = \theta$. If multiple events occur at time θ and their changes to $b_{i,v}^-$ balance out for all $i \in I$ and $v \in V$, we can simply ignore these events. Once H' is determined, we remove all events with time $\theta \leq H'$ from the priority queue.

We now show the correctness and termination of the proposed algorithm.

Theorem 23 *If all predictors are causal, computable as pw. linear functions and respect FIFO, then the algorithm described in Section 5.1 returns a δ -routed DPE up to time $H_{\text{comp}} \in \mathbb{R}$ in finite time.*

Proof If the algorithm terminates, the returned flow is feasible up to time H_{comp} . The fact that at any time θ we extend the flow along edges in $\hat{E}_i(\vartheta_\delta(\theta), \vartheta_\delta(\theta), f)$ implies that (\hat{q}, f) is a δ -routed DPE up to time H_{comp} .

To ensure that the algorithm terminates, we note that a single change in the node inflow rate $b_{i,v}^-$ at some time θ can cause finitely many events of type (E1) that all happen no earlier than at time $\theta + \tau_{\min} > \theta$. Assume, we have computed a flow up to time H . Then, all events that occur up to time $H + \tau_{\min}/2$ have already been enqueued. After processing these finitely many events, we will have computed a flow up to time $H + \tau_{\min}/2$. This suffices to assess that the algorithm terminates. \blacksquare

Remark 24 *A simple analysis on the number of events shows that the number of distribution phases is in $\mathcal{O}(P \cdot (1 + 2d_{\max}^+ \cdot |I|)^{H_{\text{comp}}/\tau_{\min} + 2})$ where P is the total number of jumps of any network inflow rate u_i , and $d_{\max}^+ := \max_{v \in V} |\delta_v^+|$ denotes the maximum out-degree.*

5.4 Computing Approximate Dynamic Prediction Equilibria

Finally, we show that if the predictors fulfill a Lipschitz condition, any δ -routed DPE is also an ε -approximate DPE. The exact notion of the Lipschitz condition is the following.

Definition 25 *A predictor $\hat{q}_{i,e}$ is called L -Lipschitz with $L \in [0, \infty)$ if for every feasible flow f it holds*

$$\forall (\theta, \bar{\theta}), (\theta', \bar{\theta}') \in \mathbb{R}^2 : \quad |\hat{q}_{i,e}(\theta, \bar{\theta}, f) - \hat{q}_{i,e}(\theta', \bar{\theta}', f)| \leq L \cdot \|(\theta, \bar{\theta}) - (\theta', \bar{\theta}')\|_\infty.$$

Theorem 26 *If (\hat{q}, f) is a δ -routed DPE and all predictors are L -Lipschitz, then (\hat{q}, f) is an ε -approximated DPE with $\varepsilon := (L/\nu_{\min} + 2) \cdot (L/\nu_{\min} + 1)^{|V|-1} \cdot \delta$ and $\nu_{\min} := \min_{e \in E} \nu_e$.*

Proof Using the triangle inequality and the Lipschitz-property of $\hat{q}_{i,e}$, we deduce that the function $\hat{T}_{i,e}(\cdot, \cdot, \cdot, f)$ is $(L/\nu_e + 1)$ -Lipschitz. Hence, for any simple path $P = e_1 \cdots e_k$, the function $\hat{T}_{i,P}(\cdot, \cdot, \cdot, f)$ is Lipschitz-continuous with constant $(L/\nu_{\min} + 1)^k \leq (L/\nu_{\min} + 1)^{|V|-1}$. The pointwise minimum $\hat{\ell}_{i,v}(\cdot, \cdot, \cdot, f)$ of finitely many such functions is $(L/\nu_{\min} + 1)^{|V|-1}$ -Lipschitz. By applying the triangle inequality once more, we see that $\hat{\Delta}_{i,e}(\cdot, \cdot, \cdot, f)$ is Lipschitz-continuous with constant $(L/\nu_{\min} + 2) \cdot (L/\nu_{\min} + 1)^{|V|-1}$. Now, the equality $\hat{\Delta}_{i,e}(\lfloor \theta/\delta \rfloor \cdot \delta, \lfloor \bar{\theta}/\delta \rfloor \cdot \delta, f) = 0$ implies $\hat{\Delta}_{i,e}(\theta, \bar{\theta}, f) \leq (L/\nu_{\min} + 2) \cdot (L/\nu_{\min} + 1)^{|V|-1} \cdot \delta = \varepsilon$. ■

Remark 27 *Note, that the approximation guarantee of Theorem 26 degrades exponentially with the number of nodes in a network. There exist networks in which also the achieved approximation quality ε degrades exponentially: Assume that the network consists of a chain P of edges from node s to node t and the predictor for each edge e of the chain is given by $\hat{q}_{i,e}(\theta, \bar{\theta}, f) := \theta \cdot L$, and let $\tau_e = 0$ and $\nu_e = 1$. Then, the predicted exit time of this chain P is $\hat{T}_{i,P}(\theta, \bar{\theta}, f) = \theta \cdot (L + 1)^{|V|-1}$. We add an additional edge e^* from s to t with $\hat{q}_{i,e^*}(\theta, \bar{\theta}, f) := 0$ and $\tau_{e^*} := 0$. Now assume a network inflow rate of $u := \mathbf{1}_{[0,1]}$. A flow f which sends all particles along the chain is a δ -routed DPE with $\delta = 1$. However, the minimal ε such that f is an ε -approximated DPE is $\varepsilon = (L + 1)^{|V|-1}$.*

Note, that this does not imply that other methods of finding ε -approximated DPE must also suffer from this lower bound.

We shortly discuss whether the introduced predictors are Lipschitz. Clearly, the Zero-Predictor $\hat{q}_{i,e}^Z$ is 0-Lipschitz. However, the linear predictor with any $H > 0$ and the regularized linear predictor with $H = \infty$ are not Lipschitz. On the positive side, the Lipschitz-condition holds for several predictors if all network inflow rates are bounded with $\|u_i\|_\infty < \infty$. Then, the queue length function q_e^f of a feasible flow f is L_{q_e} -Lipschitz with

$$L_{q_e} := \max\{\nu_e, \sum_{i \in I} \mathbf{1}_{v=s_i} \|u_i\|_\infty + \sum_{e' \in \delta_v^-} \nu_{e'}\},$$

where v is the tail of edge e . This directly implies that the constant predictor $\hat{q}_{i,e}^C$ and the perfect predictor $\hat{q}_{i,e}^P$ are L_{q_e} -Lipschitz as well. For $H < \infty$, the regularized linear predictor $\hat{q}_{i,e}^{\text{RL}}$ is L -Lipschitz with $L = 2L_{q_e}(2 + H/\delta)$ where δ is the size of the regularization window. The ML predictors also fulfill a Lipschitz-condition under an additional Lipschitz-constraint on the map $\phi_{i,e}$.

6. Computational Study

In the following we compare the different predictors introduced in Section 4.3 by measuring their performance in synthetic and real-world road networks. As ML-based predictors we introduce both a linear regression based predictor \hat{q}^{LR} and a more advanced neural network based predictor \hat{q}^{NN} .

6.1 Experiment Setup

We run an experiment for each considered network. The experiment consists of three phases: In the first phase, we generate training data by computing δ_r -routed DPE in which all commodities use the constant predictor. Secondly, this data is used to train both the linear regression and the neural network based predictors. In the third step, we evaluate the performance of all predictors by computing δ_r -routed DPE in which all predictors are used side-by-side.

More specifically, we are given a network in which all commodities $i \in I$ use the constant predictor \hat{q}^{C} . Each commodity has a positive constant network inflow rate up to some common time h , after which all network inflow vanishes. We sample this constant network inflow rate according to a normal distribution $\bar{u}_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$ and set $u_i(\theta) := \bar{u}_i$ for $\theta \in [0, h]$ and $u_i(\theta) = 0$ otherwise. We compute several of these δ_r -routed DPE according to Section 5.1 with a fixed reroute-interval δ_r and randomly chosen network inflow rates, and use the resulting flows as training data.

For the evaluation step, we randomly pick one commodity $j \in I$ as the so-called *focused commodity*. For each introduced, causal predictor, we add an additional commodity $i \in \{\hat{q}^{\text{Z}}, \hat{q}^{\text{C}}, \hat{q}^{\text{L}}, \hat{q}^{\text{RL}}, \hat{q}^{\text{LR}}, \hat{q}^{\text{NN}}\}$ using the corresponding predictor. These extra commodities allow us to measure the predictors' performance: They have the same source and sink as the focused commodity, and a very small constant inflow rate \bar{u}_i up to time h . This minimizes their influence on the flow of the original commodities.

With this setup, we compute a δ_r -routed DPE. To assess the performance, we monitor the average travel time of particles of the additional commodities: The network outflow rate of a commodity i is given by $o_i(\theta) := \sum_{e \in \delta_i^-} f_{i,e}^-(\theta) - \sum_{e \in \delta_i^+} f_{i,e}^+(\theta)$. 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 :

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

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 $T_{i,\text{OPT}}^{\text{avg}} := \int_0^h \min\{H, l_{i,s}(\theta)\} - \theta \, d\theta / h$. The *slowdown* is defined as $T_i^{\text{avg}} / T_{i,\text{OPT}}^{\text{avg}} - 1$.

We also monitor the accuracy of the predictor in hindsight: Let δ_s denote the fixed step size of two consecutive interpolation points of the machine-learned predictor. The *mean absolute error* of predictor \hat{q}_i is defined as

$$\text{MAE}_i := \frac{\sum_{\bar{\theta} \in \bar{\Theta}} \sum_{j \in [k_f]} \sum_{e \in E} |\hat{q}_{i,e}(\bar{\theta} + j\delta_s, \bar{\theta}, f) - q_e(\bar{\theta} + j\delta_s)|}{|\bar{\Theta}| \cdot k_f \cdot |E|},$$




Figure 3: A network with source s and sink t . Edges are labeled with (τ_e, ν_e) .

where $\bar{\Theta} := \{j\delta_r \mid j \in \mathbb{N}_0, j\delta_r \leq H\}$ is the set of times at which routes were computed.

We compute a total of N δ_r -routed DPE with randomly chosen commodities and network inflow rates, and aggregate the results to obtain reliable performance metrics.

6.2 Data

We conduct our experiment on several networks. The first is a warm-up synthetic graph with 4 nodes and 5 edges as presented in Figure 3; it has a single commodity with a manually chosen demand μ_i . The second network called the *Nguyen* network was taken from Nguyen (1984) augmented with capacities and edge transit times by Han et al. (2019). Appropriate demand values were generated by hand. The other graphs are real world road networks collected in (Transportation Networks for Research Core Team, 2016). Thereof, we use the networks of the cities Sioux Falls and Anaheim. These datasets come with edge attributes travel time τ_e and capacity ν_e . Moreover, they include static origin-destination demand values μ_i . We add versions of these networks with a single manually chosen commodity and a corresponding demand μ_i . All manually chosen parameters can be found in Graf et al. (2022a). Details of each network are depicted in Table 1.

Table 1: Attributes of the considered networks

Network	$ E $	$ V $	$ I $	ν_{avg}	τ_{avg}	μ_{avg}	σ_{avg}
Synthetic	5	4	1	1.4	1.4	4	0.5
Nguyen	19	13	4	50	2.2	100	50
Sioux Falls I	75	24	1	10,247	4.13	8,000	2,412
Sioux Falls II	75	24	528	10,247	4.13	682.95	603
Anaheim I	914	416	1	6,030	0.88	8,000	900
Anaheim II	914	416	1.406	6,030	0.88	74.5	900

Furthermore, we use the following parameters for computing the δ_r -DPE in all networks: The network inflow rates vanish at time $h = 12$ and the computation horizon is set to $H_{\text{comp}} = 60$ which is large enough for all particles to reach the sink. We set the reroute interval to $\delta_r = 1/8$. The predictors \hat{q}^L and \hat{q}^{RL} use a prediction horizon of $H = 20$, and \hat{q}^{RL} has a regularization window of $\delta = 1$. The ML predictors \hat{q}^{ML} have a step size of $\delta = 1$ with $k_p = 20$ past and $k_f = 20$ future samples.

6.3 Machine-Learned Predictors

To assess the impact of different ML-based models in our setting, we train simple linear regression predictors as well as more advanced neural network predictors. To obtain training data, we compute a number of δ_r -routed DPE flows. In these flows, all commodities use the constant predictor introduced in Section 4.3. This allows the model to estimate queues when agents follow instantaneously shortest paths. Moreover, the predictors are trained separately for each network so that they can take local congestion effects into account.

The features for our models are the prediction time $\bar{\theta}$, $k_p = 20$ observations of the past queue length and of the past edge load on a subset $E^{\text{in}} = \{e_1^{\text{in}}, \dots, e_n^{\text{in}}\}$ of edges. The output of the model are $k_f = 20$ samples of the predicted queue length of edges in $E^{\text{out}} = \{e_1^{\text{out}}, \dots, e_m^{\text{out}}\} \subseteq E$. We implemented the following two variants: For the first, a single model taking observations of all edges of the network as input is trained, i.e. $E^{\text{in}} = E^{\text{out}} = E$. In the second, a model for each edge e is trained that only uses observations of surrounding edges, i.e. $E^{\text{in}} = N(e)$, $E^{\text{out}} = \{e\}$, where $N(e)$ consists of edges that are no further than 3 jumps away from e on the undirected graph.

The Linear Regression Predictor As a first machine-learning method we employ a linear regression predictor \hat{q}^{LR} . This method learns a weights matrix W and a bias vector b minimizing the error between $WX_i + b$ and Y_i . We found that the Ridge regressor offered by the scikit-learn library (Pedregosa et al., 2011) works well for our uses. This regressor minimizes the linear squared error together with an L_2 regularization term.

The Neural Network Predictor To compare the capabilities of different machine-learning methods with different levels of complexity, we also develop a neural network based predictor \hat{q}^{NN} . The neural network consists of 4 densely connected layers. The first three layers have the same size as the input. The LeakyReLU function acts as the activation function between the dense layers reducing the vanishing gradient problem. An Adam optimizer minimizes the mean absolute error as the loss function. Moreover, we apply an L_2 regularization term to the weights and biases of each layer; this helps to reduce over-fitting.

6.4 Comparison of Predictors

We first take a closer look at the synthetic network shown in Figure 3. Here, a typical DPE, in which the single commodity with source s and sink t uses the constant predictor, shows the following behavior: Initially, the paths st and $svwt$ are the predicted shortest s - t -paths both with a predicted traversal time of 3. While there is a network inflow rate higher than 3, queues at the edges st , sv and wt build up in such a way that both paths st and $svwt$ have roughly the same predicted traversal time. Moreover, during that time, particles arriving at w never leave their original path by taking edge ws .

However, once the network inflow ends, the queues of sv and st start to decrease immediately. As soon as the queue of st is small enough, particles arriving at w start using both the path wst and the direct connection wt . Eventually, all remaining particles arrive at t without changing their route again.

The return of some particles to s from w indicates that a more advanced predictor could anticipate the situation better and send these particles along edge st in the first place. Figure 4 shows the slowdown and the MAE in $N = 20$ evaluation runs in the

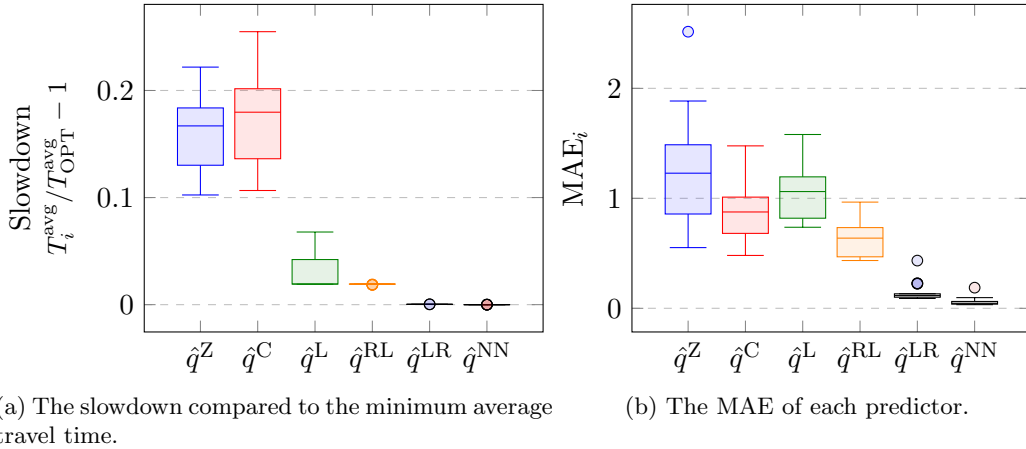


Figure 4: The experiment results of the synthetic network.

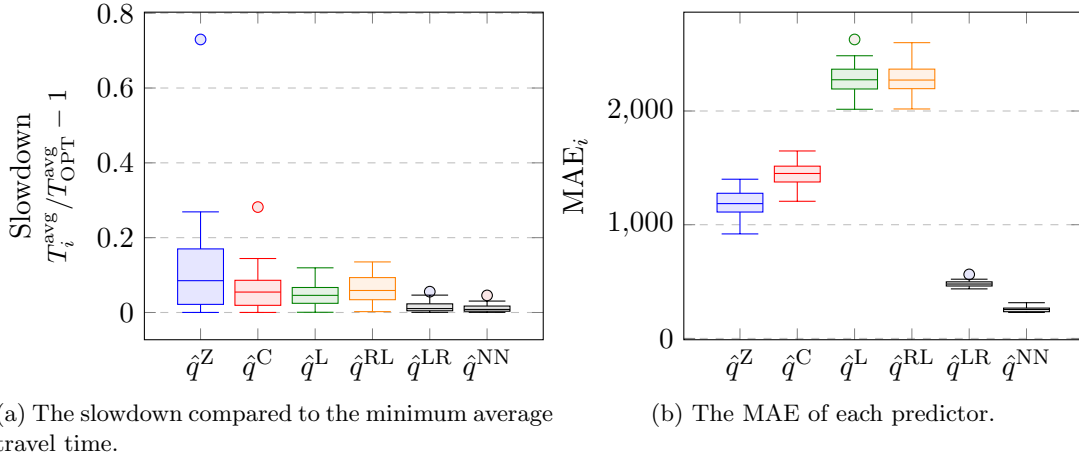


Figure 5: The experiment results of the Anaheim II network.

experiment setup declared above, aggregated in two box plots. Here, we used the machine-learning variants with full-network observations for both the linear regression and the neural network predictors.

The figure shows that the constant predictor actually has the highest slowdown in this scenario. Even the Zero-Predictor, which always sends particles at equal rates along both *st* and *svwt*, has a slightly lower slowdown. While the linear and the regularized linear predictor already show a large improvement compared to the constant predictor and the Zero-Predictor, the two machine learned predictors win against those by another order of magnitude.

In larger networks, these results are less pronounced: While the machine-learned predictors still clearly dominate the others, the difference between the other four predictors is less significant. In most experiments, \hat{q}^L performs the best out of these four. Figure 5 shows the same box plots for the real-world network of Anaheim II.

Table 2: Results of the experiment in the considered networks.

Network	Average Slowdown in %						CPU Time		
	\hat{q}^Z	\hat{q}^C	\hat{q}^L	\hat{q}^{RL}	\hat{q}^{LR}	\hat{q}^{NN}	Train \hat{q}^{LR}	Train \hat{q}^{NN}	Evaluate
Synthetic	16.07	17.36	3.01	1.91	0.03	0.00	6s	7m	6m
Nguyen	2.02	1.21	2.56	1.36	0.27	0.35	14s	10m	10m
Sioux Falls I	7.78	7.16	4.12	4.61	0.73	0.06	9s	5m	46m
Sioux Falls II	5.81	6.18	4.25	4.99	0.26	0.09	1m	1h45m	47m
Anaheim I	22.00	22.11	14.67	16.22	0.50	0.23	6m	1h35m	6h31m
Anaheim II	13.36	6.68	5.11	6.48	1.67	1.25	33m	18h35m	41h49m

The results of the experiment are summarized in Table 2. In the Anaheim II experiment, a total of 163 flows for training were computed to provide a total of 51,019 training samples for the machine-learned predictors. In all other experiments, 500 training flows were computed yielding 198,500 training samples. A training/validation split of 90%/10% was used. Moreover, only for Anaheim I and Anaheim II the machine-learning variants with observations of only neighboring edges (at most three jumps away in the undirected graph) as inputs were used. On average, this corresponds to 1.0% of all edges. For all experiments, $N = 20$ evaluation runs were carried out. The CPU times listed were measured on an Intel[®] Core[™] i7-1165G7 @ 2.80GHz while using all four cores.

 Table 3: The approximation quality $\varepsilon_{\hat{q}_i}$ of each predictor.

Network	Approximation quality $\varepsilon_{\hat{q}_i}$					
	\hat{q}^Z	\hat{q}^C	\hat{q}^L	\hat{q}^{RL}	\hat{q}^{LR}	\hat{q}^{NN}
Synthetic	0.00	0.67	2.94	0.83	0.00	0.50
Nguyen	0.00	0.74	25.52	5.47	0.68	0.63
Sioux Falls I	0.00	0.36	9.52	1.44	0.15	0.17
Sioux Falls II	0.00	0.61	24.23	8.06	0.20	0.41
Anaheim I	0.00	0.58	11.83	6.09	0.51	0.54
Anaheim II	0.00	0.65	6.57	11.25	0.52	0.46

Table 3 shows numerically obtained approximation qualities $\varepsilon_{\hat{q}}$ for each predictor \hat{q} in each network where $\varepsilon_{\hat{q}}$ is the minimal value such that

$$f_{i,e}^+(\theta) > 0 \implies \hat{\Delta}_{i,e}(\theta, \theta, f) \leq \varepsilon_{\hat{q}_i}$$

holds for all computed δ -routed DPE flows f of a particular network instance. As the Zero-predictor never changes its predicted shortest routes, its approximation quality is perfect. On the negative side, the linear and regularized linear predictors have worse approximation qualities. This aligns with our theoretical findings as the linear predictor is not Lipschitz

at all and the Lipschitz constant of the regularized linear predictor is quite high with $L = 2L_{q_e}(2 + H/\delta)$. Furthermore, we can see that for the constant predictor and the machine-learned predictors, a DPE can be approximated well with our method.

7. List of Notations

Symbol	Description
$\mathbb{R}, \mathbb{R}_{\geq 0}, \mathbb{R}_{> 0}$	the set of all, non-negative and positive real numbers, respectively
$L_{\text{loc}}^p(I, J)$	the set of locally p -integrable functions $g : I \rightarrow J$ with $I, J \subseteq \mathbb{R}$
$C(A, B)$	the set of continuous functions $g : A \rightarrow B$
\mathcal{R}	the set of $g \in L_{\text{loc}}^1(\mathbb{R}, \mathbb{R}_{\geq 0})$ with $g _{(-\infty, 0)} \stackrel{\text{a.e.}}{=} 0$
$G = (V, E)$	a directed graph with nodes V and edges E
δ_v^-, δ_v^+	the incoming and outgoing edges of node v
$\nu_e \in \mathbb{R}_{> 0}$	the capacity of an edge e
$\tau_e \in \mathbb{R}_{> 0}$	the (free-flow) transit time of an edge e
I	the set of commodities
$s_i, t_i \in V$	the source and sink node of commodity i
$V_i \subseteq V, E_i \subseteq E$	the nodes, edges lying on a s_i - t_i -path
$u_i \in \mathcal{R}$	the network inflow rate of commodity i
$f = (f^+, f^-)$	a flow over time with $f^+, f^- \in \mathcal{R}^{I \times E}$
$f_{i,e}^+(\theta)$	the inflow rate of particles of commodity i into edge e at time θ
$f_{i,e}^-(\theta)$	the outflow rate of particles of commodity i out of edge e at time θ
$F_{i,e}^+(\theta), F_{i,e}^-(\theta)$	the cumulative in- and outflow until time θ
$q_e(\theta)$	the queue length of edge e at time θ
f_e^+, f_e^-	the aggregate in- and outflow rates $f_e^+ := \sum_{i \in I} f_{i,e}^+, f_e^- := \sum_{i \in I} f_{i,e}^-$
$b_{i,v}^-(\theta)$	the inflow rate of particles of commodity i into node v
$T_e(\theta)$	the exit time when entering edge e at time θ , $T_e(\theta) := \theta + q_e(\theta)/\nu_e + \tau_e$
$H \in \mathbb{R}$	a time horizon
$\hat{q}_{i,e}(\theta, \bar{\theta}, f)$	the predicted queue length of edge e at time θ as predicted by commodity i at time $\bar{\theta}$ given the (historical) flow f
$\hat{T}_{i,e}(\theta, \bar{\theta}, f)$	the predicted exit time of edge e when entering at time θ
$\hat{\ell}_i^P(\theta, \bar{\theta}, f)$	the predicted exit time of path P when entering at time θ
$\hat{\ell}_{i,v}(\theta, \bar{\theta}, f)$	the predicted earliest arrival time at t_i when departing at time θ in v
$\hat{\Delta}_{i,e}(\theta, \bar{\theta}, f)$	the predicted delay for taking edge e at time θ
$\hat{E}_i(\theta, \bar{\theta}, f) \subseteq E_i$	the set of $\bar{\theta}$ -estimated active edges at time θ

Funding. The research of the authors was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - HA 8041/1-1 and HA 8041/4-1.

Acknowledgements. We would like to thank the three anonymous reviewers for their constructive feedback during the reviewing process.

References

- Moshe Ben-Akiva, Michel Bierlaire, Haris N. Koutsopoulos, and Rabi Mishalani. *Real Time Simulation of Traffic Demand-Supply Interactions within DynaMIT*, pages 19–36. Springer US, Boston, MA, 2002.
- Umang Bhaskar, Lisa Fleischer, and Elliot Anshelevich. A Stackelberg strategy for routing flow over time. *Games and Economic Behavior*, 92:232–247, 2015.
- David E. Boyce, Bin Ran, and Larry J. LeBlanc. Solving an instantaneous dynamic user-optimal route choice model. *Transportation Science*, 29(2):128–142, 1995.
- Haïm Brézis. équations et inéquations non linéaires dans les espaces vectoriels en dualité. *Annales de l’Institut Fourier*, 18(1):115–175, 1968. doi: 10.5802/aif.280. URL http://www.numdam.org/item/AIF_1968__18_1_115_0.
- Zhigang Cao, Bo Chen, Xujin Chen, and Changjun Wang. A network game of dynamic traffic. In Constantinos Daskalakis, Moshe Babaioff, and Hervé Moulin, editors, *Proceedings of the 2017 ACM Conference on Economics and Computation, EC ’17, Cambridge, MA, USA, June 26-30, 2017*, pages 695–696. ACM, 2017.
- Caspar G. Chorus, Eric J. E. Molin, and Bert Van Wee. Use and Effects of Advanced Traveller Information Services (ATIS): A Review of the Literature. *Transport Reviews*, 26(2):127–149, 2005.
- Roberto Cominetti, José R. Correa, and Omar Larré. Dynamic equilibria in fluid queueing networks. *Oper. Res.*, 63(1):21–34, 2015.
- José R. Correa, Andrés Cristi, and Tim Oosterwijk. On the price of anarchy for flows over time. In Anna Karlin, Nicole Immorlica, and Ramesh Johari, editors, *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019*, pages 559–577. ACM, 2019. URL <https://doi.org/10.1145/3328526.3329593>.
- Stella C. Dafermos. The traffic assignment problem for multiclass-user transportation networks. *Transportation Science*, 6(1):73–87, 1972. ISSN 00411655, 15265447. URL <http://www.jstor.org/stable/25767636>.
- Bolin Ding, Jeffrey Xu Yu, and Lu Qin. Finding time-dependent shortest paths over large graphs. In *Proceedings of the 11th International Conference on Extending Database Technology: Advances in Database Technology, EDBT ’08*, page 205–216, New York, NY, USA, 2008. Association for Computing Machinery. ISBN 9781595939265. doi: 10.1145/1353343.1353371. URL <https://doi.org/10.1145/1353343.1353371>.
- L. R. Ford and D. R. Fulkerson. *Flows in Networks*. Princeton University Press, 1962.
- Terry L. Friesz and Ke Han. The mathematical foundations of dynamic user equilibrium. *Transportation Research Part B: Methodological*, 126:309 – 328, 2019.

- Terry L. Friesz, Javier Luque, Roger L. Tobin, and Byung-Wook Wie. Dynamic network traffic assignment considered as a continuous time optimal control problem. *Operations Research*, 37(6):893–901, 1989.
- Terry L. Friesz, David Bernstein, Tony E. Smith, Roger L. Tobin, and B. W. Wie. A variational inequality formulation of the dynamic network user equilibrium problem. *Oper. Res.*, 41(1):179–191, January 1993.
- Guido Gentile. Solving a dynamic user equilibrium model based on splitting rates with gradient projection algorithms. *Transportation Research Part B: Methodological*, 92:120–147, 2016.
- Lukas Graf and Tobias Harks. The price of anarchy for instantaneous dynamic equilibria. In Xujin Chen, Nikolai Gravin, Martin Hoefer, and Ruta Mehta, editors, *Web and Internet Economics - 16th International Conference, WINE 2020, Beijing, China, December 7-11, 2020, Proceedings*, volume 12495 of *Lecture Notes in Computer Science*, pages 237–251. Springer, 2020. URL https://doi.org/10.1007/978-3-030-64946-3_17.
- Lukas Graf, Tobias Harks, and Leon Sering. Dynamic flows with adaptive route choice. *Math. Program., Ser. B*, 183(1):309–335, 2020.
- Lukas Graf, Tobias Harks, Kostas Kollias, and Michael Markl. Computation of Dynamic Prediction Equilibria. <https://github.com/ArbeitsgruppeTobiasHarks/dynamic-prediction-equilibria/tree/jmlr>, 2022a. Accessed: 2022-09-15.
- Lukas Graf, Tobias Harks, Kostas Kollias, and Michael Markl. Machine-learned prediction equilibrium for dynamic traffic assignment. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(5):5059–5067, Jun. 2022b. doi: 10.1609/aaai.v36i5.20438. URL <https://ojs.aaai.org/index.php/AAAI/article/view/20438>.
- Younes Hamdouch, Patrice Marcotte, and Sang Nguyen. A strategic model for dynamic traffic assignment. *Networks and Spatial Economics*, 4(3):291–315, Sep 2004.
- Ke Han, Terry L. Friesz, and Tao Yao. Existence of simultaneous route and departure choice dynamic user equilibrium. *Transportation Research Part B: Methodological*, 53:17 – 30, 2013a.
- Ke Han, Terry L. Friesz, and Tao Yao. A partial differential equation formulation of Vickrey’s bottleneck model, part i: Methodology and theoretical analysis. *Transportation Research Part B: Methodological*, 49:55 – 74, 2013b.
- Ke Han, Terry L. Friesz, and Tao Yao. A partial differential equation formulation of Vickrey’s bottleneck model, part ii: Numerical analysis and computation. *Transportation Research Part B: Methodological*, 49:75 – 93, 2013c.
- Ke Han, Gabriel Eve, and Terry L. Friesz. Computing dynamic user equilibria on large-scale networks with software implementation. *Networks and Spatial Economics*, 19(3): 869–902, January 2019. doi: 10.1007/s11067-018-9433-y. URL <https://doi.org/10.1007/s11067-018-9433-y>.

- Hai-Jun Huang and William H. K. Lam. 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, 2003.
- Weiwei Jiang and Jiayun Luo. Graph neural network for traffic forecasting: A survey. *CoRR*, abs/2101.11174, 2021. URL <https://arxiv.org/abs/2101.11174>.
- Ronald Koch and Martin Skutella. Nash equilibria and the price of anarchy for flows over time. *Theory Comput. Syst.*, 49(1):71–97, 2011.
- Rafał Kucharski and Guido Gentile. Simulation of rerouting phenomena in Dynamic Traffic Assignment with the Information Comply Model. *Transportation Research Part B: Methodological*, 126(C):414–441, 2019.
- Larry J. LeBlanc, Edward K. Morlok, and William P. Pierskalla. An efficient approach to solving the road network equilibrium traffic assignment problem. *Transportation Research*, 9(5):309–318, 1975. ISSN 0041-1647. doi: [https://doi.org/10.1016/0041-1647\(75\)90030-1](https://doi.org/10.1016/0041-1647(75)90030-1). URL <https://www.sciencedirect.com/science/article/pii/0041164775900301>.
- Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. In *6th International Conference on Learning Representations, ICLR*, 2018.
- Zhi-Chun Li, Hai-Jun Huang, and Hai Yang. Fifty years of the bottleneck model: A bibliometric review and future research directions. *Transportation Research Part B: Methodological*, 139:311 – 342, 2020.
- H. Mahmassani. Dynamic network traffic assignment and simulation methodology for advanced system management applications. *Networks and Spatial Economics*, 1:267–292, 2001.
- Patrice Marcotte, Sang Nguyen, and Alexandre Schoeb. A strategic flow model of traffic assignment in static capacitated networks. *Operations Research*, 52(2):191–212, 2004.
- Frédéric Meunier and Nicolas Wagner. Equilibrium results for dynamic congestion games. *Transportation Science*, 44(4):524–536, 2010. An updated version (2014) is available on Arxiv.
- Sang Nguyen. Estimating origin destination matrices from observed flows. In *Proceedings of a course given at the International Center for Transportation Studies (ICTS) Amalfi, Italy, October 11-16, 1982*, pages 363–380. Elsevier Science Publishers BV, 1984.
- Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau,

- M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Srinivas Peeta and Hani S. Mahmassani. 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, 1995.
- Srinivas Peeta and Athanasios Ziliaskopoulos. Foundations of Dynamic Traffic Assignment: The Past, the Present and the Future. *Networks and Spatial Economics*, 1(3):233–265, 2001.
- B. Ran and D. Boyce. *Dynamic urban transportation network models: theory and implications for intelligent vehicle-highway systems*. Lecture notes in economics and mathematical systems. Springer, Berlin, New York, Paris, 1996.
- Bin Ran, David E. Boyce, and Larry J. LeBlanc. A new class of instantaneous dynamic user-optimal traffic assignment models. *Operations Research*, 41(1):192–202, 1993.
- Martin Skutella. An introduction to network flows over time. In *Research Trends in Combinatorial Optimization, Bonn Workshop on Combinatorial Optimization, November 3-7, 2008, Bonn, Germany*, pages 451–482, 2008.
- Transportation Networks for Research Core Team. Transportation Networks for Research. <https://github.com/bstabler/TransportationNetworks>, March 2016. Accessed: 2022-08-18.
- Avinash Unnikrishnan and Steven Waller. User equilibrium with recourse. *Networks and Spatial Economics*, 9(4):575–593, 2009.
- Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, 2018. URL <https://openreview.net/forum?id=rJXMpikCZ>.
- William S Vickrey. Congestion theory and transport investment. *American Economic Review*, 59(2):251–60, May 1969.
- David Watling. Urban traffic network models and dynamic driver information systems. *Transport Reviews*, 14(3):219–246, 1994.
- David Watling and Martin L. Hazelton. The dynamics and equilibria of day-to-day assignment models. *Networks and Spatial Economics*, 3:349–370, 2003.
- Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, and Chengqi Zhang. Graph wavenet for deep spatial-temporal graph modeling. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*. AAAI Press, 2019.
- Hai Yang. Multiple equilibrium behaviors and advanced traveler information systems with endogenous market penetration. *Transportation Research Part B: Methodological*, 32(3): 205–218, 1998.

Bing Yu, Haoteng Yin, and Zhanxing Zhu. 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.*, pages 3634–3640, 2018.

Chuanpan Zheng, Xiaoliang Fan, Cheng Wang, and Jianzhong Qi. 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*, pages 1234–1241, 2020. URL <https://aaai.org/ojs/index.php/AAAI/article/view/5477>.

Daoli Zhu and Patrice Marcotte. On the existence of solutions to the dynamic user equilibrium problem. *Transportation Science*, 34(4):402–414, 2000.