University of Augsburg

Institute for Mathematics
Discrete Mathematics, Optimization, and Operations Research

Master's Thesis

Prediction Equilibria in Dynamic Traffic Assignment

Michael Markl

Supervisor:	Prof. Dr. Tobias Harks
Second Reviewer:	Prof. Dr. Mirjam Dür
Matriculation Number:	1474802
Submission Date:	March 07, 2022
Revision:	June 21, 2022 (#497dc5)

Contents

1 Introduction

The prevailing behavior of traffic participants on today's streets highly differs from that only a few decades ago. In the past, road users based their routing decisions mainly on physical maps. The information derived from these maps is an estimate on the transit time of each street under normal, congestion-free conditions. Since the early 2000s agents' routing decisions rely more and more upon the navigation systems built into their vehicles. With the advent of the Traffic Message Channel (TMC), these systems were not only able to compute a shortest path to the destination under normal conditions, but they could also react to the current congestion status reported via radio signal.

In the past years, however, these navigation systems have experienced another stage of improvement: Services like Google Maps not only offer real-time analyses of the current traffic load using data gathered through the Internet, but their method of finding the best routes to the destination also reflects a predicted future evolution of the network's traffic load based both on historical and real-time data. This enables users, for example, to avoid the usual commuting traffic that is likely to boost travel times in rush hour on certain streets.

Routing games play a significant role in the analysis and optimization of transportation and communication networks. Traditionally, traffic load in these networks was modeled using static flows. Modern techniques for modeling traffic networks capture the evolution of traffic over time, resulting in dynamic flows, which are also called flows over time. In this thesis the agents' behavior is subject to the well-known Vickrey's fluid queuing model. This means that agents are modeled as infinitesimally small particles in a continuous flow setting. Here, a network is modeled as a directed graph where each edge has a certain transit time and a capacity rate. The transit time is the time a particle needs to travel from the one end of the edge to the other, whereas the capacity rate describes the maximum rate at which vehicles can enter the edge. If this rate is exceeded, a queue starts to build up in front of the edge, and particles have to wait for an additional queuing time until they can traverse the edge.

A game-theoretical structure is added to this physical model to reflect the agents' behavior and to define states of equilibrium. In such an equilibrium, all agents, modeled as infinitesimal particles, travel along "shortest paths" to their destination – or rather what they think are shortest paths. Here, the various models of equilibria analyzed in the past differ mainly in the informational access of the agents.

In the classical model, the so-called *full information model*, agents know the exact future evolution of the traffic load on every edge. Therefore, they can determine actual shortest paths when beginning their travel at their source node. Such a flow is called a *dynamic (Nash) equilibrium (DE) flow.*

Of course, assuming the full information model is quite unrealistic. In an effort to reflect the adaptive route choices of navigation systems using TMC, a new model was considered: The *instantaneous information model.* Here, agents base their routing decision on the current congestion of all edges that can be retrieved instantaneously. This means, whenever a particle arrives at an intermediate node, it will retrieve the current waiting times of all edges, and calculate a shortest path according to these instantaneous travel times (assuming the reported queue sizes will remain constant). A corresponding flow is then called *instantaneous dynamic equilibrium (IDE) flow*.

This thesis analyzes yet another model which takes the latest innovations of navigation systems into account. Traffic participants are equipped with a means of forecasting the traffic based on the historical and real-time evolution of the traffic. This forecast might not be accurate, however agents will base their shortest path calculations on the predicted traffic evolution using this forecast method. If all agents travel along these predicted shortest paths, a so-called *dynamic prediction equilibrium (DPE)* arises. This model turns out to be a generalization of the two models mentioned above.

The thesis is organized in the following chapters. First, several properties of First-In-First-Out (FIFO) ordered cost functions and algorithms for determining shortest paths in these kinds of networks in Chapter 2 are discussed. This is particularly interesting for Vickrey's fluid queuing model which induces a FIFO-ordered cost function itself. Dynamic flows that adhere to this physical model are then formally introduced in Chapter 3. Moreover, the deterministic nature of this model which has often been implicitly used in literature, is formally proven. Chapter 4 focuses on the theoretical analysis of dynamic prediction equilibria. After providing an intuition of the underlying behavioral assumptions, the specific notion of equilibrium is formally defined. The chapter continues with a discussion of an unconstructive existence proof of DPE under mild assumptions on the predictors. This existence result is obtained by solving a variational inequality in the L^p function space. After presenting a set of example predictors, the model is compared with the dynamic Nash equilibrium and the instantaneous dynamic equilibrium flows, and it is shown that the proposed model generalizes both of them. Finally, Chapter 5 presents an algorithm for computing so-called *approximate dynamic prediction equilibria (ε-DPE)*. The created simulation is run both on synthetic and real-world traffic networks and a comparison of the performance of the considered predictors is deduced.

2 Computing Dynamic Shortest Paths in FIFO-Networks

Time-dependent cost functions that follow the First-In-First-Out (FIFO) order play an important role in Vickrey's fluid queuing model and in the dynamic prediction equilibria introduced later: Once we equip every agent with a method of predicting the traffic, we expect these predictions to induce dynamic cost functions that follow the FIFO order.

In this chapter – as well as in the rest of this thesis – we only work on directed graphs $G = (V, E)$ with a finite number of nodes V and a finite number of edges E. Although we allow parallel edges, we often write $e = vw \in E$ for a directed edge e from node v to node w. For a node v, we denote the set of *outgoing edges* of v as $\delta_v^+ := \{vw \in E\}$ and the set of *incoming edges* as $\delta_v^- := \{uv \in E\}$. For a subset of nodes $V' \subseteq V$ we denote the set of edges leaving V' by $\delta^+(V') := \{vw \in E \mid v \in V' \not\ni w\}$ and those entering V' by $\delta^-(V') := \{vw \in E \mid v \notin V' \ni w\}$. Moreover, we use the common notation $[n] := \{1, \dots, n\}$ for the first n natural numbers for $n \in \mathbb{N}$.

Throughout this chapter we assume that all nodes can reach a specific sink node $t \in V$. Given FIFO-ordered cost functions, we aim to find dynamic shortest paths to this sink node as well as so-called *active edges* for a specific point in time θ. These are edges that lie on a shortest path from a node v to the sink t.

This chapter gives a brief overview of some properties of FIFO-ordered cost functions and introduces basic dynamic shortest paths algorithms. The focus is to pave the way for the computation of approximated dynamic prediction equilibria as performed in Chapter 5. A more in-depth analysis of the shortest-path problem in FIFO networks has been worked out in [8]. A complexity analysis of the shortest path problem for piecewise linear cost functions was carried out in [10].

We begin by defining the FIFO order for dynamic cost functions in Section 2.1. After discussing some basic properties of FIFO-ordered cost functions in Section 2.2, we focus on the duality between earliest arrival and latest departure times in Section 2.3 leading to a convenient characterization of active edges. In Section 2.4 a dynamic variant of Dijkstra's algorithm for FIFO-ordered cost functions is analyzed and used in Section 2.5 to compute the set of active outgoing edges of a single node. Section 2.6 wraps up the chapter with the discussion of the Bellman-Ford algorithm which is capable of computing the earliest arrival times at the sink as functions over time.

2.1 Definition of the FIFO Order

This section defines the notions of shortest paths and active edges for time-dependent, FIFO-ordered cost functions.

For a time-dependent (also called dynamic) cost function $c : \mathbb{R} \to \mathbb{R}_{\geq 0}^E$, the *exit time of an edge* $e \in E$ when entering it at time θ is given by $T_e(\theta) := \theta + c_e(\theta)$ with $T_e : \mathbb{R} \to \mathbb{R}$.

Definition 2.1.1. A time-dependent cost function $c : \mathbb{R} \to \mathbb{R}^E_{\geq 0}$ is *FIFO-ordered*, if for all edges $e \in E$ the function T_e is monotonically increasing. It is *strongly FIFO-ordered* if T_e is strictly increasing for all $e \in E$.

In other words, a cost function is FIFO-ordered if particles that enter the edge at a later point in time than others also arrive at a later point in time at the target node of an edge.

Let $c : \mathbb{R} \to \mathbb{R}^E_{\geq 0}$ be a time-dependent cost function. The exit time T_P of a finite path $P = e_1 \cdots e_k$ is given by the concatenation of the edges' exit times as $T_P := T_{e_k} \circ \cdots \circ T_{e_1}$. A path is called *simple* if it does not contain a cycle, and we denote the set of all simple v-w-paths as $\mathcal{P}_{v,w}$ for any two nodes $v, w \in V$. The *earliest arrival time* at t when starting at time θ in v is then given by $l_{v,t}(\theta) := \min_{P \in \mathcal{P}_{v,t}} T_P(\theta)$. A path that attains this minimum is called a *shortest v-t-path at time θ*.

We call an edge $e = vw \in E$ *active at time θ* if the condition $l_{v,t}(\theta) = l_{w,t}(T_e(\theta))$ holds true. All active edges at time θ are collected in the set $E(\theta)$. Here, the underlying idea is that we call an edge active if it lies on a shortest v-t-path. It should be noted, however, that an active edge does not necessarily lie on a *simple* shortest path.

2.2 Properties of the FIFO Order

In the definition of the earliest arrival time, cyclic paths were ignored. One characteristic property of FIFO-ordered cost functions is that cycles do not have a beneficial effect on the arrival time of a path, thus allowing us to only consider simple paths.

Proposition 2.2.1. *For a FIFO-ordered cost function $c : \mathbb{R} \to \mathbb{R}^E_{\geq 0}$, removing a cycle in any v-t-path P does not increase the path's exit time.*

More specifically, if $P = P_1\, C\, P_2$ is the concatenation of a path P_1, a cycle C and another path P_2, then $T_P(\theta) \geq (T_{P_2} \circ T_{P_1})(\theta)$ holds for all $\theta \in \mathbb{R}$.

Proof. The statement is a direct consequence of the monotonicity of T_{P_2} and the fact that $T_C(\theta) \geq \theta$ holds for all $\theta \in \mathbb{R}$. $\square$

The next observation shows that the FIFO order enables us to find a representation of the earliest arrival functions $(l_{v,t})_v$ based on the earliest arrival functions of neighboring nodes. This observation is also used to describe the dynamic Bellman-Ford algorithm discussed in Section 2.6.

Proposition 2.2.2. *Let $c : \mathbb{R} \to \mathbb{R}^E_{>0}$ be a FIFO-ordered cost function. The vector $(l_{v,t})_{v \in V}$ of functions is the pointwise maximal solution of the following system of equations in the function-valued variables $(\tilde{l}_v : \mathbb{R} \to \mathbb{R})_{v \in V}$:*

$$\tilde{l}_v(\theta) = \begin{cases} \theta, & \text{if } v = t, \\ \min_{e = vw \in \delta_v^+} \tilde{l}_w\left(T_e(\theta)\right), & \text{otherwise.} \end{cases}$$

Proof. We first prove that $(l_{v,t})_{v \in V}$ is a solution of the system. For $v = t$ we have $l_{v,t}(\theta) = \theta$ for all $\theta \in \mathbb{R}$. For $v \neq t$ let $e = vw \in \delta_v^+$ and let P be a (simple) shortest w-t-path at time $T_e(\theta)$ and $P' := eP$ be the concatenation of e and P.

If P' contains a cycle, then it was introduced by e and v occurs a second time in P'. By removing this cycle we obtain a simple v-t-path Q with

$$l_{v,t}(\theta) \leq T_Q(\theta) \leq T_{P'}(\theta) = l_{w,t}(T_e(\theta))$$

using Proposition 2.2.1. If P' does not contain a cycle, we analogously infer

$$l_{v,t}(\theta) \leq T_{P'}(\theta) = l_{w,t}(T_e(\theta)).$$

Hence, $l_{v,t}(\theta)$ is a lower bound on $\{l_{w,t}(T_e(\theta)) \mid e = vw \in \delta_v^+\}$. Let $P = e_1 \cdots e_k$ be a shortest v-t-path at time θ and let $e_1 = vw$ and $P' = e_2 \cdots e_k$. Furthermore, let Q be a shortest w-t-path at time $T_{e_1}(\theta)$. Assuming $T_P(\theta) > l_{w,t}(T_{e_1}(\theta))$ implies

$$T_P(\theta) > T_Q(T_{e_1}(\theta)),$$

which means that removing any cycle in the path $e_1 Q$ would yield a strictly shorter path than P; a contradiction.

We now have to show that $(l_{v,t})_{v\in V}$ is the pointwise maximal solution. This means that for any other solution $(\tilde{l}_v)_{v\in V}$ of the system of equations, $\tilde{l}_v(\theta) \leq l_v(\theta)$ must hold for all $v \in V$ and $\theta \in \mathbb{R}$. Let $P = e_1 \cdots e_k$ be a shortest v-t-path at time θ with $v_0 = v, v_k = t$ and $e_i = v_{i-1}v_i$ for $i \in [k]$. Then, applying the system of equations yields

$$\tilde{l}_v(\theta) \leq \tilde{l}_{v_1}(T_{e_1}(\theta)) \leq \tilde{l}_{v_2,t}(T_{e_1 e_2}(\theta)) \leq \cdots \leq \tilde{l}_t(T_P(\theta)) = T_P(\theta) = l_{v,t}(\theta).$$

□

The following example illustrates why the above system of equations above does not always have a unique solution. The method used works in any cyclic graph: We build a self-confirming cycle proposing an earlier arrival time than actually possible.

Example 2.2.3. For a counterexample of the uniqueness, we consider a network consisting of three nodes $V = \{s, v, t\}$, three edges $E = \{st, sv, vs\}$ and an arbitrary FIFO-ordered cost function $c : \mathbb{R} \to \mathbb{R}^E_{\geq 0}$. The graph is depicted below:

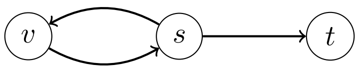

We artificially bound the arrival time of s at 0 using $\tilde{l}_s(\theta) := \min\{T_{st}(\theta), 0\}$. The arrival times of v and t are defined as expected with $\tilde{l}_t(\theta) := \theta$ and $\tilde{l}_v(\theta) := \tilde{l}_s(T_{vs}(\theta))$. Obviously, the system of equations is satisfied for nodes v and t. For node s, the equation reads

$$\tilde{l}_s(\theta) = \min\left\{\tilde{l}_t(T_{st}(\theta)), \tilde{l}_v(T_{sv}(\theta))\right\}.$$

By inserting the definitions of $\tilde{l}_v$ and $\tilde{l}_t$, the right-hand side equates to

$$\min\left\{T_{st}(\theta), \tilde{l}_s(T_{vs}(T_{sv}(\theta)))\right\} = \min\left\{T_{st}(\theta), T_{st}\Big(T_{vs}(T_{sv}(\theta))\Big), 0\right\}.$$

Proposition 2.2.1 implies the inequality $T_{st}(\theta) \leq T_{st}(T_{vs}(T_{sv}(\theta)))$, so that the right-hand side reduces to $\min\{T_{st}(\theta), 0\} = \tilde{l}_s(\theta)$.

This solution $(\tilde{l}_w)_{w\in V}$ is different to the actual earliest arrival times $(l_w)_{w\in V}$: The earliest arrival time when starting in s at any positive time $\theta > 0$ is given by

$$l_s(\theta) = T_{st}(\theta) \geq \theta > 0 = \tilde{l}_s(\theta).$$

2.3 Duality of Arrival and Departure Times

Sometimes it is useful not to work with the earliest arrival time, but with the latest possible departure time instead. To enable this switch, we define a kind of inverse of a monotonically increasing function.

Although this technique was developed independently of past studies, a similar method has been introduced in [8].

Definition 2.3.1. We define the function space

$$\mathcal{F} := \left\{ f : \mathbb{R} \to \mathbb{R} \mid\ f \text{ is increasing and } \lim_{|x|\to\infty} |f(x)| = \infty \right\}.$$

The *reversal of $f \in \mathcal{F}$* is defined as

$$f^{\leftarrow} : \mathbb{R} \to \mathbb{R}, \quad \theta \mapsto \sup \{\xi \in \mathbb{R} \,|\, f(\xi) \leq \theta\}.$$

We can interpret the reversal of f in the following way: If $f(\theta)$ is the earliest arrival time when departing at time θ, then $f^{\leftarrow}(\theta)$ is the latest departure time for arriving before or at time θ.

Proposition 2.3.2. *For $f, g \in \mathcal{F}$ the following statements are true:*

(i) It holds that $f^{\leftarrow} \in \mathcal{F}$, i.e. $f^{\leftarrow}$ is increasing and $\lim_{|x|\to\infty} |f(x)| = \infty$.

(ii) If f is continuous, then $f^{\leftarrow}(\theta) = \max\{\xi \in \mathbb{R} \mid f(\xi) = \theta\}$ *holds for all $\theta \in \mathbb{R}$ and $f^{\leftarrow}$ is strictly increasing. Moreover, we have $f \circ f^{\leftarrow} = \mathbf{id}_{\mathbb{R}}$.*

(iii) If f is continuous and strictly increasing, then $f^{\leftarrow}$ is the inverse of f.

(iv) It holds that $(g \circ f)^{\leftarrow} = f^{\leftarrow} \circ g^{\leftarrow}$ and $(\min\{f, g\})^{\leftarrow} = \max\{f^{\leftarrow}, g^{\leftarrow}\}$.

(v) If $f \geq \mathbf{id}_{\mathbb{R}}$ holds pointwise, then so does $f^{\leftarrow} \leq \mathbf{id}_{\mathbb{R}}$.

Proof. (i). Let $\theta_1, \theta_2 \in \mathbb{R}$ with $\theta_1 < \theta_2$. Then

$$f^{\leftarrow}(\theta_1) = \sup\{\xi \in \mathbb{R} \mid f(\xi) \leq \theta_1\} \leq \sup\{\xi \in \mathbb{R} \mid f(\xi) \leq \theta_2\} = f^{\leftarrow}(\theta_2) \tag{2.1}$$

implies the monotonicity. From $\lim_{|\theta|\to\infty} |f(\theta)|$ and the monotonicity of f we conclude

$$\lim_{|\theta|\to\infty} |f^{\leftarrow}(\theta)| = \lim_{|\theta|\to\infty} |\sup\{\xi \in \mathbb{R} \mid f(\xi) \leq \theta\}| = \infty.$$

(ii). We note that $\{\xi \in \mathbb{R} \mid f(\xi) = \theta\}$ is non-empty, closed and bounded from above by the condition $\lim_{|\theta|\to\infty} |f(\theta)| = \infty$ and the continuity and monotonicity of f. The reversal $f^{\leftarrow}$ is strictly increasing as the inequality (2.1) is strict for continuous f.

(iii). Let f be continuous and strictly increasing. Then it holds that

$$f^{\leftarrow}(\theta) = \max\{\xi \mid f(\xi) = \theta\} = f^{-1}(\theta).$$

(iv). By definition, the first statement evaluated at time θ becomes

$$l \coloneqq \sup\{\xi \mid g(f(\xi)) \leq \theta\} = \sup\{\xi_f \mid f(\xi_f) \leq \sup\{\xi_g \mid g(\xi_g) \leq \theta\}\} \eqqcolon r.$$

Let $\xi_f \in \mathbb{R}$ fulfill $f(\xi_f) \leq g^{\leftarrow}(\theta)$. Then for all ξ_g with $g(\xi_g) \leq \theta$ we have $f(\xi_f) \leq \xi_g$. Using the monotonicity of g and f we infer $g(f(\xi_f)) \leq g(\xi_g) \leq \theta$ which implies $l \geq r$. To see that $r \geq l$ holds, any $\xi \in \mathbb{R}$ with $g(f(\xi)) \leq \theta$ fulfills $f(\xi) \leq g^{\leftarrow}(\theta)$.

The statement $(\min\{f, g\})^{\leftarrow} = \max\{f^{\leftarrow}, g^{\leftarrow}\}$ evaluated in θ is equivalent to

$$l \coloneqq \sup\{\xi \mid \min\{f(\xi), g(\xi)\} \leq \theta\} = \max\{\sup\{\xi \mid f(\xi) \leq \theta\}, \sup\{\xi \mid g(\xi) \leq \theta\}\} \eqqcolon r.$$

Let $\xi \in \mathbb{R}$ fulfill $\min\{f(\xi), g(\xi)\} \leq \theta$ and, without loss of generality, assume $f(\xi) \leq g(\xi)$. Then it follows $f(\xi) \leq \theta$ and therefore $\xi \leq f^{\leftarrow}(\theta) \leq r$, implying $l \leq r$. In order to show $l \geq r$, assume $f^{\leftarrow}(\theta) \leq g^{\leftarrow}(\theta)$ without loss of generality, so that $r = g^{\leftarrow}(\theta)$. Any $\xi \in \mathbb{R}$ with $g(\xi) \leq \theta$ fulfills $\min\{f(\xi), g(\xi)\} \leq g(\xi) \leq \theta$ and hence $l \geq r$ holds true.

(v). For $f \geq \mathbf{id}_R$ we conclude

$$f^{\leftarrow}(\theta) = \sup\{\xi \in \mathbb{R} \mid f(\xi) \leq \theta\} \leq \sup\{\xi \in \mathbb{R} \mid \xi \leq \theta\} = \theta$$

for all $\theta \in \mathbb{R}$. □

Corollary 2.3.3. *For a FIFO-ordered cost $c : \mathbb{R} \to \mathbb{R}^E_{\geq 0}$ with $\lim_{\theta \to -\infty} T_e(\theta) = -\infty$, the following statements hold true:*

(i) For all edges $e \in E$ we have $T_e \in \mathcal{F}$.

(ii) For any path $P = e_1 \cdots e_k$ it holds that $T_P^{\leftarrow} = T_{e_1}^{\leftarrow} \circ \cdots \circ T_{e_k}^{\leftarrow}$.

(iii) For any node $v \in V$ it holds that $l_{v,t}^{\leftarrow} = \max_{P \in \mathcal{P}_{v,t}} T_P^{\leftarrow}$.

The reversal of a function can be utilized for a characterization of active edges in the case of continuous cost function:

Lemma 2.3.4. *Let $c : \mathbb{R} \to \mathbb{R}^E_{\geq 0}$ be a continuous cost function following the FIFO order with $\lim_{\theta \to -\infty} T_e(\theta) = -\infty$ for all $e \in E$ and let $t \in V$ be a sink node. Then an edge $e = vw$ is active at time θ if and only if $T_e(\theta) \leq l_{w,t}^{\leftarrow}(l_{v,t}(\theta))$.*

Proof. Let edge e be active at time θ, i.e. $l_{w,t}(T_e(\theta)) \leq l_{v,t}(\theta)$. By definition of the reversal this already implies $l_{w,t}^{\leftarrow}(l_{v,t}(\theta)) \geq T_e(\theta)$.

If on the other hand $T_e(\theta) \leq l_{w,t}^{\leftarrow}(l_{v,t}(\theta))$ holds, then the monotonicity of $l_{w,t}$ implies

$$l_{w,t}(T_e(\theta)) \leq l_{w,t}\Big(l_{w,t}^{\leftarrow}(l_{v,t}(\theta))\Big).$$

By the continuity of $l_{w,t}$ the claim now follows from Proposition 2.3.2 (ii) which states that $l_{w,t} \circ l_{w,t}^{\leftarrow} = \mathbf{id}_{\mathbb{R}}$. □

2.4 The Dynamic Dijkstra Algorithm

The first algorithm we discuss is a simple modification of Dijkstra's algorithm to determine the earliest arrival times $(l_{s,w}(\theta))_{w \in V'}$ at all nodes $w \in V$ that are reachable from s. Adjusting Dijkstra's algorithm for static edge costs to our setting yields the Dynamic Dijkstra Algorithm as depicted in Algorithm 1.

Algorithm 1 The Dynamic Dijkstra Algorithm

```
def dynamic_dijkstra(
  theta: float, source: Node, costs: Dict[Edge, Callable[[float], float]]
) -> Dict[Node, float]:
  arrival: Dict[Node, float] = {}
  queue: PriorityQueue[Node] = PriorityQueue({ source: theta })
  while len(queue) > 0:
    v, xi = queue.pop_min()
    arrival[v] = xi
    for e in v.outgoing_edges:
      w = e.node_to
      if w in arrival:
        continue
      relaxation = arrival[v] + costs[e](arrival[v])
      if w not in queue:
        queue.push(w, relaxation)
      elif relaxation < queue.key_of(w):
        queue.decrease_key(w, relaxation)
  return arrival
```

A priority queue, consisting of items together with associated priority keys, operates at the core of the algorithm. The procedure requires the queue to support the operations `push(item, key)`, `pop_min()`, `decrease_key(item, new_key)`, `item in queue`, and `key_of(item)`. The operation `push(item, key)` adds the item `item` with priority `key` to the queue, `pop_min()` returns the item-key-pair with the minimum key and removes it from the queue, `decrease_key(item, new_key)` replaces the priority key associated to the item `item` with `new_key`, `item in queue` returns whether `item` is contained in the queue and the operation and `key_of(item)` returns the key of `item` in the queue.

The idea of Algorithm 1 is to *visit* a node v only once its earliest arrival time $l_{s,v}(\theta)$ is determined. We say a node w has been *discovered* if a node v of an incoming edge vw has been visited. Here, all visited nodes together with their earliest arrival are recorded in the dictionary `arrival`. The priority queue `queue` keeps track of all unvisited, discovered nodes and associates each of them with a *currently suspected earliest arrival time* as its priority key.

In the main loop of the procedure, we retrieve an unvisited, discovered node v with the smallest priority key ξ amongst all unvisited, discovered nodes. Later, we will prove $\xi = l_{s,v}(\theta)$. Then, for each outgoing edge $e = vw$ we realize the edge's cost at time ξ and update the currently suspected arrival time of the target node w.

The following proposition states, that the Dynamic Dijkstra Algorithm is correct.

Proposition 2.4.1. *Given FIFO-ordered cost function* $c : \mathbb{R} \to \mathbb{R}^E_{\geq 0}$*, the Dynamic Dijkstra Algorithm initiated on* $s \in V$ *and* $\theta \in \mathbb{R}$ *computes the vector* $(l_{s,w}(\theta))_{w \in V'}$*, where* V' *is the set of nodes reachable from* s*.*

Proof. As an invariant of the main loop, we prove that all key-value-pairs of the dictionary `arrivals` are of the form $(w, l_{s,w}(\theta))$. In the beginning this is clearly true as `arrival` is initially empty. After the first iteration of the loop, the only key-value-pair in `arrivals` is $(s, \theta) = (s, l_{s,s}(\theta))$. Assume the loop invariant holds before entering the body of the loop again at a later time and let v be the popped node. As nodes are added to the queue once at most, we have $v \neq s$. Let u be the node in whose loop iteration v was added to the queue or in whose loop iteration the key of v in the queue was decreased, whatever happened the most recently.

Because the invariant implies that the value of u in `arrival` equals $l_{s,u}(\theta)$, the value of v in `arrival` equals $l_{s,u}(\theta) + c_{uv}(l_{s,u}(\theta)) = T_{uv}(l_{s,u}(\theta)) \geq l_{s,v}(\theta)$. Assume this inequality is strict and let P be a shortest s-v-path at time θ. If all nodes of P were available in `arrival`, then the last node before v in P would have set the key of v in its iteration to $l_{s,v}(\theta)$. Let u be the first node in P that is not available in `arrival`. Because u cannot be the source s, the predecessor u' of u in P must have set the key of u to at most $T_{u'u}(l_{s,u'}(\theta)) \leq T_P(\theta)$. As $T_{uv}(l_{s,u}(\theta)) > l_{s,v}(\theta) = l_P(\theta)$ the key of u in the queue was smaller than the key of v, so the priority queue would have popped u before v. $\square$

A simple binary min-heap together with a lookup table was implemented to support the operations of the queue efficiently. With this data structure, the worst case running time is logarithmic in the number of queue items for the operations `push(item, key)`, `pop()` and `decrease_key(item, new_key)` and constant for the operations `min_key()` and `contains(item)`. Thus, the Dynamic Dijkstra Algorithm terminates with a runtime of $\mathcal{O}((|V| + |E|) \cdot \log |V| \cdot \mathcal{T}_c)$ where $\mathcal{T}_c$ denotes the time it takes to evaluate the cost $c_e(\theta)$ for some $e \in E$ and $\theta \in \mathbb{R}$.

2.5 Computing Active Outgoing Edges

Given a FIFO-ordered cost function $c : \mathbb{R} \to \mathbb{R}^E_{\geq 0}$, nodes $s, t \in V$ and a time $\theta \in \mathbb{R}$, we want to compute the set of active outgoing edges $E(\theta) \cap \delta^+_s$ of s, i.e. the edges $e = sw$ with

$$l_{s,t}(\theta) = l_{w,t}(T_e(\theta)).$$

Unfortunately, we cannot determine all active edges from the arrival times $(l_{s,w}(\theta))_{w \in V'}$ obtained by a simple run of the Dynamic Dijkstra Algorithm: In a static scenario, the idea is to backtrack all shortest paths by searching edges $e = vw$ backwards starting from t for which equality holds in $l_{s,w}(\theta) \leq T_e(l_{s,v}(\theta))$. This approach is described in Algorithm 2, which aims to return the set $E(\theta) \cap \delta^+_s$. The vector $(l_{s,w}(\theta))_{w \in V'}$ as returned by the Dynamic Dijkstra Algorithm is passed to the function via the parameter `arrivals`. During the procedure, we only enqueue a node u in `queue`, if there exists a u-t-path in which each edge $e = vw$ fulfills $l_{s,w}(\theta) = T_e(l_{s,v}(\theta))$. Once we discover an edge from s to such an enqueued node that also fulfills the equality, we add it to the set of active outgoing edges of s.

Algorithm 2 Backtracking Shortest Paths

```
def backtrack_shortest_paths(
  source: Node, sink: Node, arrival: Dict[Node, float],
  costs: Dict[Edge, Callable[[float], float]]
) -> Set[Edge]:
  active_edges = set()
  queue: List[Node] = [sink]
  discovered: Set[Node] = {sink}
  while len(queue) > 0:
    w = queue.pop()
    for e in w.incoming_edges:
      v = e.node_from
      if v not in arrivals \
         or arrivals[v] + costs[e](arrivals[v]) > arrivals[w]:
        continue
      if v == source:
        active_edges.add(e)
      if v not in discovered:
        queue.append(v)
        discovered.add(v)
  return active_edges
```

The following proposition proves that paths found using the described approach are in fact shortest paths. Moreover, it shows that for strongly FIFO-ordered costs we find all shortest paths, which implies that Algorithm 2 returns the whole set $E(\theta)\cap\delta_s^+$ for this type of cost functions.

Proposition 2.5.1. *Let $P = e_1 \cdots e_k$ be a path with $e_i = v_{i-1}v_i$ and $v_0 = s, v_k = t$ and let $c : \mathbb{R} \to \mathbb{R}^E_{\geq 0}$ be a cost function following the FIFO order. Then*

$$(\forall i \in [k] : T_{e_i}(l_{s,v_{i-1}}(\theta)) = l_{s,v_i}(\theta)) \implies T_P(\theta) = l_{s,t}(\theta).$$

If c is strongly FIFO-ordered, the statements are equivalent.

Proof. Assume $T_{e_i}(l_{s,v_{i-1}}(\theta)) = l_{s,v_i}(\theta)$ holds for all i. Then we have

$$T_P(\theta) = T_P(l_{s,s}(\theta)) = T_{e_2\cdots e_{k-1}}(l_{s,v_1}(\theta)) = \cdots = l_{s,v_k}(\theta) = l_{s,t}(\theta).$$

Pretend that c is strongly FIFO-ordered and let P fulfill $T_P(\theta) = l_{s,t}(\theta)$. We assume there is some $i \in [k]$ with $T_{e_i}(l_{s,v_{i-1}}(\theta)) > l_{s,v_i}(\theta)$. Let P' be a shortest s-v_i-path at time θ. We extend P' with $e_{i+1}\cdots e_k$ and obtain an s-t-path P''. The strong monotonicity of all T_e yields the contradiction

$$l_{s,t}(\theta) \leq T_{P''}(\theta) = T_{e_{i+1}\cdots e_k}(l_{s,v_i}(\theta)) < T_{e_{i+1}\cdots e_k}(T_{e_1\cdots e_i}(\theta)) = T_P(\theta). \qquad \square$$

For general FIFO-ordered costs however, not all subpaths of shortest paths are again shortest paths. This means, there might be edges $e = vw$ that do not fulfill the equality

$l_{s,w}(\theta) = T_e(l_{s,v}(\theta))$ but still lie on a shortest s-t-path at time θ. This might be the case if there exists a bottleneck edge e' closer to t together with an interval on which $T_{e'}$ is constant. An example of this can be seen in Figure 2.1.

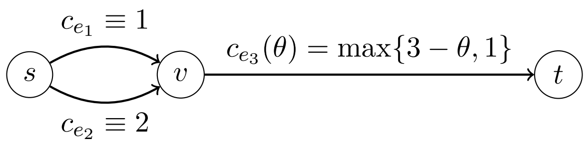

Figure 2.1: Both e_1 and e_2 are active at time 0 with $T_{e_2}(l_{s,s}(0)) > l_{s,v}(0)$.

This leaves us with the question of how to find the rest of the active edges for general FIFO-ordered cost functions. For the rest of this section, we assume that $c : \mathbb{R} \to \mathbb{R}^E_{\geq 0}$ is a continuous, FIFO-ordered cost function with $\lim_{\theta \to -\infty} T_e(\theta) = -\infty$ for all $e \in E$. The idea for finding all active outgoing edges $\delta^+_s \cap E(\theta)$ is described as follows: Once we determined $l_{s,t}(\theta)$ using a first run of the Dynamic Dijkstra Algorithm, we carry out a second run of the Dynamic Dijkstra Algorithm on the reverse graph and the reverse time to compute the latest departure time vector $(l^{\leftarrow}_{w,t}(l_{s,t}(\theta)))_w$. Lemma 2.3.4 then states that an outgoing edge $e = sw \in \delta^+_s$ of s is active if and only if $T_e(\theta) \leq l^{\leftarrow}_{w,t}(l_{s,t}((\theta)))$.

More specifically, the *reverse graph* of $G = (V, E)$ is defined as $G^{\leftarrow} = (V, E^{\leftarrow})$ with the *reversed edges* $E^{\leftarrow} := \{e^{\leftarrow} = wv \mid e = vw \in E\}$. We define a new cost function $\tilde{c}$ on the reverse graph as

$$\tilde{c} : \mathbb{R} \to \mathbb{R}^{E^{\leftarrow}}_{\geq 0}, \quad \tilde{c}_{e^{\leftarrow}}(\theta) := -T^{\leftarrow}_e(-\theta) - \theta.$$

We denote the exit times induced by $\tilde{c}$ by $\tilde{T}_{e^{\leftarrow}}$ and $\tilde{T}_{P^{\leftarrow}}$, where $P^{\leftarrow} = e^{\leftarrow}_k \cdots e^{\leftarrow}_1$ is a path in $G^{\leftarrow}$ for a corresponding $P = e_1 \cdots e_k$ in G. Moreover, the earliest arrival time functions induced by $\tilde{c}$ are denoted as $\tilde{l}_{v,w}$ for $v, w \in V$.

Lemma 2.5.2. *Let c be a continuous, FIFO-ordered cost function $c : \mathbb{R} \to \mathbb{R}^E_{\geq 0}$ fulfilling $\lim_{\theta \to -\infty} T_e(\theta) = -\infty$ for all $e \in E$. Then $\tilde{l}_{v,w}(\theta) = -l^{\leftarrow}_{w,v}(-\theta)$ holds for all $\theta \in \mathbb{R}$.*

Proof. The non-negativity of c implies $T_e \geq \mathbf{id}_{\mathbb{R}}$ and by Proposition 2.3.2 (v) we infer $T^{\leftarrow}_e \leq \mathbf{id}_{\mathbb{R}}$. Therefore, we have $T^{\leftarrow}_e(-\theta) \leq -\theta$ and thus $\tilde{c}_{e^{\leftarrow}}(\theta) \geq 0$ holds for all $\theta \in \mathbb{R}$.

The new edge exit times fulfill $\tilde{T}_e(\theta) = -T^{\leftarrow}_e(-\theta)$ and, by Proposition 2.3.2 (ii), $\tilde{T}_e$ is strictly increasing and $\tilde{c}$ is strongly FIFO-ordered. The exit time of a path $P^{\leftarrow} = e^{\leftarrow}_k \cdots e^{\leftarrow}_1$ when entering at time θ equals

$$\tilde{T}_{P^{\leftarrow}}(\theta) = \tilde{T}_{e^{\leftarrow}_{k-1} \cdots e^{\leftarrow}_1}(-T^{\leftarrow}_{e_k}(-\theta)) = \tilde{T}_{e^{\leftarrow}_{k-2} \cdots e^{\leftarrow}_1}(-T^{\leftarrow}_{e_{k-1}}(T^{\leftarrow}_{e_k}(-\theta)) = \cdots = -T^{\leftarrow}_P(-\theta).$$

Therefore, the earliest arrival times are given as

$$\tilde{l}_{v,w}(\theta) = \min_{P \in \mathcal{P}_{w,v}} \tilde{T}_{P^{\leftarrow}}(\theta) = \min_{P \in \mathcal{P}_{w,v}} -T^{\leftarrow}_P(-\theta) = -\max_{P \in \mathcal{P}_{w,v}} T^{\leftarrow}_P(-\theta) = -l^{\leftarrow}_{w,v}(-\theta).$$

□

With this setup, we execute a second run of the Dynamic Dijkstra Algorithm on the reverse graph $G^{\leftarrow}$ together with the costs $\tilde{c}$, the starting time $-l_{s,t}(\theta)$, and the sink as the start node. The resulting vector is of the form

$$\left(\tilde{l}_{t,w}(-l_{s,t}(\theta))\right)_w = \left(-l^{\leftarrow}_{t,w}(l_{s,t}(\theta))\right)_w ,$$

for which a negation yields the desired vector.

To implement this algorithm, we have to restrict ourselves to cost functions where the reversal of T_e can be evaluated. This is the case, for example, if the cost functions c_e are piecewise linear functions for all $e \in E$. Assuming the reversal of T_e can be determined in $\mathcal{T}_c^{\leftarrow}$ time the running time of the complete procedure is $\mathcal{O}((|V| + |E|) \cdot \log |V| \cdot (\mathcal{T}_c + \mathcal{T}_c^{\leftarrow}))$. We collect our result in the following theorem.

Theorem 2.5.3. *Let c be a continuous, FIFO-ordered cost function $c : \mathbb{R} \to \mathbb{R}^E_{\geq 0}$ fulfilling $\lim_{\theta \to -\infty} T_e(\theta) = -\infty$ for all $e \in E$. Then, the above procedure computes the set $\delta_s^+ \cap E(\theta)$ in $\mathcal{O}((|V| + |E|) \cdot \log |V| \cdot (\mathcal{T}_c + \mathcal{T}_c^{\leftarrow}))$ time.*

2.6 Computing the Earliest Arrival Functions

Often, it is useful not only to compute active edges for some fixed point in time, but to have the earliest arrival functions $(l_{v,t})_{v \in V'}$ at some sink available as functions over time.

A simple but also quite expensive method of computing these functions is a modification of the Bellman-Ford algorithm as described for example in [9]. It uses the representation given in Proposition 2.2.2. This means that we want to find the pointwise maximal solution of the system of equations

$$\tilde{l}_v(\theta) = \begin{cases} \theta, & \text{if } v = t, \\ \min_{e = vw \in \delta_v^+} \tilde{l}_w\left(T_e(\theta)\right), & \text{otherwise.} \end{cases}$$

Hence, the idea is to initialize all functions with $\tilde{l}_v(\theta) \coloneqq \infty$ for $v \neq t$ and $\tilde{l}_t(\theta) \coloneqq \theta$ and decrease the functions (pointwise) using the operation $\tilde{l}_v \coloneqq \min_{vw \in \delta_v^+} \tilde{l}_w \circ T_{vw}$ for all $v \in V \setminus \{t\}$ until no further changes can be made. Then all equations are automatically fulfilled.

More specifically, if some function $\tilde{l}_w$ changes, then all nodes v with an edge vw leading to w might need to be adjusted as well using the operation $\tilde{l}_v \coloneqq \min\{\tilde{l}_v, \tilde{l}_w \circ T_{vw}\}$. Therefore, we need some operations on the class of functions operated on to formulate the algorithm: In order to calculate the edge exit times $T_e = c_e + \mathbf{id}_{\mathbb{R}}$ we need pointwise addition and a representation of the identity function; for updates of the functions the pointwise minimum and the composition of functions have to be implemented. To detect changes we also need to be able to identify whether one function is pointwise smaller or equal to some other function.

In the context of this thesis, all the necessary operations explained above have been created for piecewise linear functions. The resulting procedure is shown in Algorithm 3. Here, `cost + identity` is the piecewise linear function representing the sum of `cost` and `identity`. The function `compose` computes the piecewise linear function representing the composition of the first argument with the second argument. The decision whether the function `arrivals[v]` is element wise smaller or equal to `relaxation` is expressed by the

Algorithm 3 Dynamic Bellman-Ford Algorithm

```
def dynamic_bellman_ford(
  sink: Node, costs: Dict[Edge, PiecewiseLinear], theta: float
) -> Dict[Node, PiecewiseLinear]:
  arrivals: Dict[Node, PiecewiseLinear] = { sink: identity }
  exit_times = {e: cost + identity for (e, cost) in costs.items()}
  changed_nodes = {sink}
  while len(changed_nodes) > 0:
    change_detected = {}
    for w in changed_nodes:
      for e in w.incoming_edges:
        v = e.node_from
        relaxation = compose(arrivals[w], exit_times[e])
        if v not in arrivals.keys():
          change_detected.add(v)
          arrivals[v] = relaxation
        elif not arrivals[v] <= relaxation:
          change_detected.add(v)
          arrivals[v] = minimum(arrivals[v], relaxation)
    changed_nodes = change_detected
  return arrivals
```

term `arrivals[v] <= relaxation`. Finally, the pointwise minimum of two functions is computed using the operation `minimum`.

For a more thorough analysis of the Dynamic Bellman-Ford algorithm and similar algorithms the reader may is referred to [9].

The benefit of calculating the earliest arrival times as functions is that only a single run of the Dynamic Bellman-Ford Algorithm is necessary to compute the set of active edges $E(\theta)$ for all $\theta \in \mathbb{R}$. This can be achieved by simply evaluating $l_{w,t}(T_e(\theta)) \leq l_{v,t}(\theta)$ for all $e = vw \in E$ and $\theta \in \mathbb{R}$. Additionally, the result can be used to calculate average travel times as explained in Section 5.4.

3 Vickrey's Fluid Queuing Model

This chapter discusses the physical model of agents that participate in the traffic network. Here, agents act as infinitesimal particles in a continuous, dynamic flow. The formal model goes back to a work by William Vickrey in [26], and therefore it is often referred to as *Vickrey's Fluid Queuing Model*.

In this work we follow the mathematical notation given by Cominetti et al. in [7] which is similar to that in [12–14, 17]. In Section 3.1 we formally introduce dynamic flows and the corresponding restrictions necessary to establish the deterministic queuing behavior. Before we analyze these so-called feasible dynamic flows we prove some fundamental properties of absolute continuous functions in Section 3.2. These are useful for demonstrating intuitive characteristics of feasible flows in Section 3.3. Finally, in Section 3.4 we formally prove that given a set of edge inflow rates there exists a unique set of outflow rates for any edge such that the capacity constraints of an edge are satisfied. This property was claimed by Cominetti et al. and their insights in [7] were used to obtain the formal proof.

3.1 Fundamental Definitions

We begin by defining the model of dynamic flows, which are also called flows over time. Each edge $e \in E$ of a finite, directed graph $G = (V, E)$ has a positive *rate capacity* $\nu_e > 0$ and a non-negative *transit time* $\tau_e \geq 0$. The rate capacity limits the amount of flow an edge can transfer per time unit. As such, this capacity can be imagined as the width of a conveyor belt. The transit time is the time the conveyor belt needs to transfer particles from its beginning to its end. Throughout the thesis we allow $\tau_e = 0$, however for most results we require strict positive transit times or, in some cases, that the transit time $\sum_{e \in C} \tau_e$ of any directed cycle C is strictly positive.

Moreover, in this work we consider multicommodity flows. That means, we have a finite set I of commodities and each commodity $i \in I$ has a *source node* $s_i \in V$ and a *sink node* $t_i \in V \setminus \{s_i\}$. We require that s_i can reach t_i in G. Furthermore, we define

$$\begin{aligned} V_i &\coloneqq \{v \in V \mid v \text{ lies on a directed path from } s_i \text{ to } t_i \text{ in } G\} \quad \text{and} \\ E_i &\coloneqq \{e \in E \mid e \text{ lies on a directed path from } s_i \text{ to } t_i \text{ in } G\} \end{aligned}$$

as the subset of nodes and edges that are *relevant* to commodity i.

Definition 3.1.1. The space of locally p-integrable functions from an interval I to a measurable set $X \subseteq \mathbb{R}$ is defined as

$$L^p_{\text{loc}}(I, X) \coloneqq \left\{ f \in L^0(I, X) \,\middle|\, \int_a^b |f|^p \,\mathrm{d}\lambda < \infty \text{ for all } a < b \text{ in } I \right\},$$

where $L^0(I, X)$ is the set of equivalence classes of measurable functions that are almost everywhere identical.

The space of *rate functions* consists of all non-negative, locally integrable functions that vanish on $(-\infty, 0)$, and is denoted by

$$\mathcal{R} \coloneqq \left\{ f \in L^1_{\mathrm{loc}}(\mathbb{R}, \mathbb{R}_{\geq 0}) \,\middle|\, f|_{(-\infty,0)} \overset{a.e.}{=} 0 \right\}.$$

Each commodity $i \in I$ has a *network inflow rate* which is given by a rate function $u_i \in \mathcal{R}$. For a time $\theta \in \mathbb{R}$, the value $u_i(\theta)$ denotes the rate at time θ at which particles of commodity i enter the network in the commodity's source s_i.

Definition 3.1.2. A *(dynamic) flow* is a pair $f = (f^+, f^-)$ of families of functions with $f^+_{i,e}, f^-_{i,e} \in \mathcal{R}$ for all $e \in E$ and $i \in I$.

Here, $f^+_{i,e}(\theta)$ is called the *inflow rate* of commodity i into edge e at time θ; $f^-_{i,e}(\theta)$ describes the *outflow rate* of commodity i out of e at time θ. We denote the *total inflow and outflow rates* of an edge e by $f^+_e(\theta) \coloneqq \sum_{i\in I} f^+_{i,e}(\theta)$ and $f^-_e(\theta) \coloneqq \sum_{i\in I} f^-_{i,e}(\theta)$, respectively.

The *cumulative inflow and cumulative outflow* of a commodity i up to time θ are defined as $F^+_{i,e}(\theta) \coloneqq \int_0^\theta f^+_{i,e}(t)\,\mathrm{d}t$ and $F^-_{i,e}(\theta) \coloneqq \int_0^\theta f^-_e(t)\,\mathrm{d}t$, respectively. Similarly, the total cumulative in- and outflow are defined as $F^+_e(\theta) \coloneqq \sum_{i\in I} F_{i,e}(\theta)$ and $F^-_e(\theta) \coloneqq \sum_{i\in I} F^-_{i,e}(\theta)$, respectively. Based on that, the *queue length* of an edge e at time θ is given by $q_e(\theta) \coloneqq F^+_e(\theta) - F^-_e(\theta + \tau_e)$.

Now, $c_e(\theta) \coloneqq \tau_e + q_e(\theta)/\nu_e$ denotes the time-dependent *cost* of traversing an edge e at time θ where $q_e(\theta)/\nu_e$ is often referred to as the *waiting time* of edge e at time θ. Based on that, the *exit time* when entering an edge e at time θ is defined as $T_e(\theta) \coloneqq \theta + c_e(\theta)$.

Notation 3.1.3. If the dynamic flow is not clear from the context, we often write $F^{+,f}_e$, $F^{-,f}_e$, q^f_e, c^f_e, or T^f_e to explicitly state that f is the underlying dynamic flow on which these functions depend on.

This definition already gives a good insight into how particles should behave in a dynamic flow of this type: A particle of commodity i is generated at the source s_i at some time θ and aims to arrive at the sink t_i as early as possible. At the time the agent spawns at s_i, it immediately enters an outgoing edge e of s_i. At this moment, the cost of e is $c_e(\theta) = \tau_e + q_e(\theta)/\nu_e$, which means that the particle has to queue for $q_e(\theta)/\nu_e$ time units before it can traverse the edge in τ_e additional time units. It will thus arrive at v at time $T_e(\theta) = \theta + c_e(\theta)$. Once it arrives at an intermediate node $v \neq t_i$, it will again choose a new outgoing edge in δ^+_v until it arrives at its destination t_i.

Figure 3.1 illustrates the queuing behavior for a simple network. In this example, all edges have a transit time of $\tau_e = 1$ and a capacity of $\nu_e = 1$ except edge s_1v, which has a capacity of $\nu_e = 2$. Here, the set of commodities $I = \{r, g\}$ consists of a red commodity r and a green commodity g that both share the same sink t. The network inflow rates are given as $u_r|_{[0,\infty)} \equiv 2$ and $u_g|_{[0,\infty)} \equiv 1$. For $i \in I$, all particles that spawn at some time θ at the source s_i immediately enter the first edge s_iv which can carry all particles because its capacity matches the network inflow rate. As shown in Figure 3.1(a), the first particles that have appeared in the network reach the center of the first edge at time 0.5. At time 1 these particles reach the node v. As soon as a particle arrives at v, it will immediately enter the next edge vt. Starting from time 1, particles enter the edge vt at a rate of $f^+_{vt}|_{[1,\infty)} = 3$ exceeding the edge's capacity by 2. Therefore, a queue starts to build up in front of vt. At

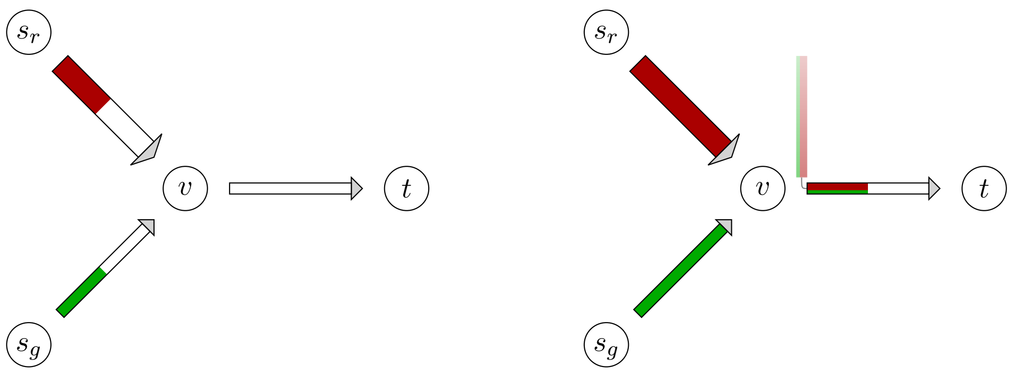

(a) The flow at time 0.5.

(b) The flow at time 1.5.

Figure 3.1: Queuing behavior of particles in feasible dynamic flows.

time 1.5 this queue has a length of $q_{vt}(1.5) = 1$ as shown in Figure 3.1(b). Because the incoming flow of the red commodity makes up $2/3$ of the total inflow rate into edge vt, the red commodity is designated $2/3$ of the capacity of vt.

In order to make dynamic flows follow the physical model as imagined above, we have to make several restrictions. These result in the class of feasible dynamic flows.

Definition 3.1.4. Given a dynamic flow f, we introduce the following constraints:

(F1) The outflow rate on edge e does not exceed its capacity, i.e. $f_e^-(\theta) \leq \nu_e$.

(F2) Edges operate at capacity, i.e.

$$\forall e \in E: \quad f_e^-(\theta) = \begin{cases} \nu_e, & \text{if } q_e(\theta - \tau_e) > 0, \\ f_e^+(\theta - \tau_e), & \text{otherwise.} \end{cases}$$

(F3) Flow traverses an edge in a FIFO-manner, i.e.

$$\forall e \in E: \quad f_{i,e}^-(\theta) = \begin{cases} f_e^-(\theta) \cdot \frac{f_{i,e}^+(\xi)}{f_e^+(\xi)}, & \text{if } f_e^+(\xi) > 0, \\ 0, & \text{otherwise,} \end{cases}$$

where $\xi \coloneqq \min\{\xi \mid T_e(\xi) = \theta\}$ is the earliest possible time for entering edge e in order to exit the edge at time θ.

(F4) Flow is preserved on intermediate nodes, i.e.

$$\forall i \in I, v \in V: \quad \sum_{e \in \delta_v^+} f_{i,e}^+(\theta) - \sum_{e \in \delta_v^-} f_{i,e}^-(\theta) \begin{cases} = u_i(\theta), & \text{for } v = s_i, \\ = 0, & \text{for } v \in V \setminus \{s_i, t_i\}, \\ \leq 0, & \text{for } v = t_i. \end{cases}$$

A flow f is called *feasible up to time* $H \in \mathbb{R} \cup \{\infty\}$ if it fulfills all these properties for almost all $\theta < H$.

A flow f is called *deterministic on edge e up to time $H \in \mathbb{R} \cup \{\infty\}$* if it fulfills properties (F1), (F2) and (F3) for edge e and almost all $\theta < H$. A flow f which is deterministic on all $e \in E$ up to time H, is called a *deterministic flow up to time H*.

For $H = \infty$, we usually omit the term "up to ∞" in the above definitions.

Remark 3.1.5. We defined property (F3) in the same way as in literature such as [14]. However, we note that by Proposition 3.3.1 and Corollary 3.2.9 the set

$$\{\theta \mid \min\{\xi \mid T_e(\xi) = \theta\} < \max\{\xi \mid T_e(\xi) = \theta\}\}$$

has zero measure whenever (F1) is fulfilled. Thus, we might choose any ξ with $T_e(\xi) = \theta$ without any effect on the feasibility definitions. Taking the minimum of such ξ usually simplifies the handling of flows that fulfill (F1) up to some finite horizon H.

3.2 Absolutely Continuous Functions

Before discussing some properties of feasible dynamic flows, it is useful to first introduce so-called *absolutely continuous functions*.

Definition 3.2.1. A function $f : [a, b] \to \mathbb{R}$ is *absolutely continuous* on the interval $[a, b]$, if for all $\varepsilon > 0$ there exists some $\delta > 0$ such that any finite sequence $(a_1, b_1), \ldots, (a_N, b_N)$ of pairwise disjoint, open subintervals of $[a, b]$ fulfills

$$\sum_{i=1}^{N}(b_i - a_i) < \delta \implies \sum_{i=1}^{N}|f(b_i) - f(a_i)| < \varepsilon.$$

The following two theorems as given help to understand why absolutely continuous functions play an important role in the discussion of feasible flows.

Theorem 3.2.2 ([24, Theorem 10 in Section 6.5])**.** *Let $f : [a, b] \to \mathbb{R}$ be absolutely continuous for some $a, b \in \mathbb{R}$, $a < b$. Then f is differentiable almost everywhere on (a, b), and its derivative f' is integrable over $[a, b]$ with*

$$\int_a^b f' \,\mathrm{d}\lambda = f(b) - f(a).$$

The next statement is often referred to as the *Lebesgue differentiation theorem*.

Theorem 3.2.3 ([24, Theorem 11 in Section 6.5])**.** *A function $f : [a, b] \to \mathbb{R}$ is absolutely continuous on $[a, b]$ if and only if there exists an integrable function $g : [a, b] \to \mathbb{R}$ with*

$$f(x) = f(a) + \int_a^x g \,\mathrm{d}\lambda.$$

The immediate results of these theorems in the context of dynamic flows are that F_e^+ and F_e^- are absolutely continuous on any interval $[a, b]$ because the edge inflow and outflow rates are locally integrable. Later, we will see that the same applies to the queue length functions and edge exit time functions as well.

The next proposition gives an equivalent formulation of absolutely continuous functions using possibly infinite sequences of subintervals.

Proposition 3.2.4. *A function $f : [a, b] \to \mathbb{R}$ is absolutely continuous if and only if for all $\varepsilon > 0$ there exists some $\delta > 0$ such that any finite or countably infinite sequence $((a_i, b_i))_{i \in I}$ of pairwise disjoint, open subintervals of $[a, b]$ fulfills*

$$\sum_{i \in I} (b_i - a_i) < \delta \implies \sum_{i \in I} |f(b_i) - f(a_i)| < \varepsilon.$$

Proof. Let f be absolutely continuous and let $\varepsilon > 0$ be arbitrary. We take $\delta > 0$ as the witness fulfilling Definition 3.2.1 for $\varepsilon/2$. We only need to check the property for countably infinite sequences $([a_i, b_i])_{i \in \mathbb{N}}$ of pairwise disjoint, closed subintervals of $[a, b]$ with $\sum_{i \in \mathbb{N}} (b_i - a_i) < \delta$. Here, it holds $\sum_{i=1}^{N} |f(b_i) - f(a_i)| < \varepsilon/2$ for any $N \in \mathbb{N}$. Hence, we deduce $\sum_{i=1}^{\infty} |f(b_i) - f(a_i)| \leq \varepsilon/2 < \varepsilon$. □

In the following statements, λ^* refers to the Lebesgue outer measure.

Proposition 3.2.5. *Let $f : [a, b] \to \mathbb{R}$ be absolutely continuous and $\varepsilon > 0$ arbitrary. For $\delta > 0$ fulfilling the requirements in Proposition 3.2.4, $\lambda^*(A) < \delta$ implies $\lambda^*(F(A)) < \varepsilon$.*

In particular, absolutely continuous functions map Lebesgue-null sets to Lebesgue-null sets. This is known as the Luzin N property.

Proof. By definition of the Lebesgue outer measure there exists an infinite sequence (a_i, b_i) of open intervals with $N \subseteq \bigcup_{i \in \mathbb{N}} (a_i, b_i)$ such that $\sum_{i \in \mathbb{N}} (b_i - a_i) < \delta$. Without loss of generality we can assume that these intervals are pairwise disjoint by taking the union of overlapping intervals.

By continuity, f attains its minimum m_i and maximum M_i on $[a_i, b_i]$ in points $c_i, d_i \in [a_i, b_i]$. We then have $\lambda^*(f((a_i, b_i))) = \lambda([m_i, M_i]) = |f(d_i) - f(c_i)|$, respectively. With $\sum_{i \in \mathbb{N}} |d_i - c_i| \leq \sum_{i \in \mathbb{N}} (b_i - a_i) < \delta$ and Proposition 3.2.4 it follows

$$\lambda^*(f(N)) \leq \lambda^* \left(\bigcup_{i \in \mathbb{N}} f\left((a_i, b_i)\right) \right) \leq \sum_{i \in \mathbb{N}} \lambda^* \left(f((a_i, b_i)) \right) = \sum_{i \in \mathbb{N}} |f(d_i) - f(c_i)| < \varepsilon.$$

□

For the analysis of some properties of feasible dynamic flows, a variant of Sard's theorem will be used. In its original form introduced by Sard in [25], the theorem states that the set of critical values of a differentiable map is a Lebesgue-null set (c.f. [5, Theorem 6.1]). We show that a similar result holds true for absolutely continuous functions.

This result builds upon Vital's Covering Theorem as presented in [24]. We call an interval non-degenerate, if its interior is non-empty.

Definition 3.2.6. A collection $\mathcal{V}$ of closed, bounded, non-degenerate intervals is called a *Vitali-Covering* of a set $A \subseteq \mathbb{R}$ if for all $x \in A$ and $\varepsilon > 0$ there exists an interval $[a, b] \in \mathcal{V}$ with $x \in [a, b]$ and $b - a < \varepsilon$.

Lemma 3.2.7 (Vitali's Covering Theorem [24, Section 6.2])**.** *Let A be a measurable set with $\lambda^*(A) < \infty$ and let $\mathcal{V}$ be a Vitali-Covering of A consisting of closed, bounded intervals. Then there exists a finite or countably infinite disjoint sequence $([a_i, b_i])_{i \in I}$ of intervals in $\mathcal{V}$ with*

$$\lambda^* \left(A \setminus \bigcup_{i \in I} [a_i, b_i] \right) = 0.$$

Theorem 3.2.8 (Sard's theorem)**.** *Let $f : [a,b] \to \mathbb{R}$ be absolutely continuous and let A be the set of points $x \in [a,b]$ in which f is differentiable with $f'(x) = 0$. Then, $f(A)$ is a Lebesgue-null set.*

Proof. For arbitrary $\varepsilon > 0$ we show $\lambda^*(f(A)) < \varepsilon$. We define a Vitali-Covering as the set

$$\mathcal{V} := \left\{ B_h(x) \,\middle|\, x \in A, h > 0, \forall y \in B_h(x) : \, |f(y) - f(x)| < \frac{\varepsilon}{2(b-a)} \cdot |y - x| \right\},$$

where we denote $B_h(x) := [x - h, x + h]$. This is indeed a Vitali-Covering for A as for all $x \in A$ we have

$$\lim_{y \to x} \frac{|f(y) - f(x)|}{|y - x|} = 0$$

and thus for all $\delta > 0$ there exists an interval of length smaller than δ in $\mathcal{V}$ containing x. By Lemma 3.2.7, there is a finite and pairwise disjoint sequence $(B_{h_i}(x_i))_{i \in I}$ of intervals in $\mathcal{V}$ with $\lambda^*(A \setminus \bigcup_{i \in I} B_{h_i}(x_i)) = 0$. We use this together with Proposition 3.2.5 to estimate

$$\begin{aligned} \lambda^*(f(A)) &\leq \lambda^* \left(f \left(A \setminus \bigcup_{i \in I} B_{h_i}(x_i) \right) \cup f \left(\bigcup_{i \in I} B_{h_i}(x_i) \right) \right) \\ &\leq \sum_{i \in I} \lambda^* \left(f \left(B_{h_i}(x_i) \right) \right). \end{aligned}$$

For all $i \in I$ we know that $f(B_{h_i}(x_i))$ is contained in an interval of length $h_i \cdot \varepsilon / (b - a)$. Moreover, as all $B_{h_i}(x_i)$ are pairwise disjoint subsets of $[a, b]$, we have $\sum_{i \in I} 2 \cdot h_i \leq b - a$. This implies

$$\lambda^*(f(A)) \leq \sum_{i \in I} \lambda^*(f(B_{h_i}(x_i))) \leq \sum_{i \in I} \frac{h_i \cdot \varepsilon}{b - a} \leq \frac{\varepsilon \cdot (b - a)}{2 \cdot (b - a)} < \varepsilon.$$

□

We note that Sard's theorem even holds for arbitrary continuous functions. The proof of this is also based on Vitali's Covering Theorem, however it is a bit more technical. We conclude this section with a corollary for absolutely continuous and non-decreasing functions.

Corollary 3.2.9. *Let $f : [a,b] \to \mathbb{R}$ be an absolutely continuous, non-decreasing function. Then the set of points $t \in f([a,b])$ with $\min f^{-1}(\{t\}) < \max f^{-1}(\{t\})$ is a Lebesgue-null set.*

Proof. For every such point t, let ξ_t denote the center point of the closed, non-degenerate interval $I := f^{-1}(\{t\})$. The derivative of f in ξ_t exists and vanishes. By Sard's theorem, the set of the points $f(\xi_t) = t$ is a Lebesgue-null set. □

3.3 Some Properties of Feasible Dynamic Flows

To get a better intuition of the behavior of feasible dynamic flows, we collect a few statements that will prove beneficial later in this thesis. The first observation we make is that particles are processed by the queue in a FIFO order. Therefore, all properties of FIFO-ordered cost functions that were obtained in Chapter 2 hold true for the cost function induced by feasible dynamic flows.

Proposition 3.3.1. *The time-dependent cost function c_e induced by a dynamic flow f fulfilling (F1) for almost all $\theta \in \mathbb{R}$ is FIFO-ordered.*

Proof. We apply property (F1) in the following inequality with $\theta_1 \leq \theta_2$:

$$\begin{aligned} T_e(\theta_2) - T_e(\theta_1) &= \theta_2 - \theta_1 + \frac{q_e(\theta_2) - q_e(\theta_1)}{\nu_e} = \theta_2 - \theta_1 + \frac{\int_{\theta_1}^{\theta_2} f_e^+(t)\,\mathrm{d}t - \int_{\theta_1+\tau_e}^{\theta_2+\tau_e} f_e^-(t)\,\mathrm{d}t}{\nu_e} \\ &\geq \theta_2 - \theta_1 + \frac{\int_{\theta_1}^{\theta_2} f_e^+(t)\,\mathrm{d}t - (\theta_2 - \theta_1)\cdot \nu_e}{\nu_e} \geq 0. \end{aligned}$$

□

The same property (F1) is responsible for the fact that the queue stays strictly positive on any interval $(\theta, \theta + q_e(\theta)/\nu_e)$.

Proposition 3.3.2. *If a dynamic flow f fulfills (F1) up to time H on an edge e, then q_e is strictly positive on any interval $(\theta, \theta + q_e(\theta)/\nu_e)$ with $T_e(\theta) \leq H$.*

Proof. If $q_e(\theta)$ is non-positive, the interval is empty. For positive $q_e(\theta)$, property (F1) implies

$$q_e(\theta + \delta) = q_e(\theta) + \int_\theta^{\theta+\delta} f_e^+\,\mathrm{d}\lambda - \int_{\theta+\tau_e}^{\theta+\tau_e+\delta} f_e^-\,\mathrm{d}\lambda \geq q_e(\theta) - \delta \cdot \nu_e > 0$$

for all $\delta \in (0, q_e(\theta)/\nu_e)$. □

Once we add property (F2) we can show that the cumulative flow that has entered an edge up to some time θ has left the edge until time $T_e(\theta)$.

Proposition 3.3.3. *If a dynamic flow f fulfills (F1) and (F2) up to some time H on an edge $e \in E$, then $F_e^+(\theta) = F_e^-(T_e(\theta))$ holds for all θ with $T_e(\theta) \leq H$. If (F1) is fulfilled almost everywhere, then $F_e^+(T_e^{\leftarrow}(\xi)) = F_e^-(\theta)$ holds true for $\xi \leq H$.*

Proof. Proposition 3.3.2 and property (F2) imply $\int_\theta^{\theta+q_e(\theta)/\nu_e} f_e^-\,\mathrm{d}\lambda = q_e(\theta)$. We conclude

$$F_e^-(T_e(\theta)) = F_e^-(\theta + \tau_e) + \int_{\theta+\tau_e}^{\theta+\tau_e+q_e(\theta)/\nu_e} f_e^-\,\mathrm{d}\lambda = F_e^+(\theta) + q_e(\theta) - q_e(\theta) = F_e^+(\theta).$$

If (F1) is fulfilled almost everywhere, c_e is FIFO-ordered by Proposition 3.3.1. Thus, we can apply Proposition 2.3.2 (ii) with $\theta = T_e^{\leftarrow}(\xi)$ which yields $F_e^+(T_e^{\leftarrow}(\xi)) = F_e^-(\xi)$. □

Property (F2) also implies that the outflow rate of an edge vanishes up to time τ_e. Therefore, if arbitrary functions $(f_{i,e}^-)_{i\in I} \in L^1_{\mathrm{loc}}(\mathbb{R}, \mathbb{R}_{\geq 0})^I$ fulfill (F2) up to some time $H \geq 0$, the following proposition already implies $f_{i,e}^- \in \mathcal{R}$.

Proposition 3.3.4. *If a dynamic flow f fulfills (F2) on an edge e up to some time $H \in \mathbb{R}_{\geq 0} \cup \{\infty\}$, then $f^-_{i,e}(\theta)$ vanishes for almost all $\theta \leq \min\{\tau_e, H\}$ for all $i \in I$. Moreover, for all $\theta < H - \tau_e$ the queue length attains the form*

$$q_e(\theta) = \int_0^\theta f^+_e(\xi) - f^-_e(\xi + \tau_e)\,\mathrm{d}\lambda(\xi).$$

These statements hold true for arbitrary $(f^+_{i,e})_{i\in I} \in \mathcal{R}^I$ and $(f^-_{i,e})_{i\in I} \in L^1_{\mathrm{loc}}(\mathbb{R}, \mathbb{R}_{\geq 0})^I$ that fulfill (F2) up to time H.

Proof. Because f^+_e vanishes almost everywhere on $(-\infty, 0)$ and f^-_e is non-negative, the queue function q_e is non-positive on $(-\infty, 0)$. (F2) implies that $f^-_e(\theta) = f^+_e(\theta - \tau_e) = 0$ holds for almost all $\theta \leq \min\{\tau_e, H\}$. The non-negativity of all $f^-_{i,e}$ shows $f^-_{i,e}|_{(-\infty,\min\{\tau_e,H\})} \overset{a.e.}{=} 0$ for all $i \in I$.

For $\theta + \tau_e < H$, the queue length can be analyzed using a case distinction. For $\theta \leq 0$ any integral involving f^-_e vanishes in the following calculation due to the above observation and $\theta + \tau_e \leq \min\{\tau_e, H\}$:

$$\begin{aligned} q_e(\theta) &= \int_0^\theta f^+_e\,\mathrm{d}\lambda - \int_0^{\theta+\tau_e} f^-_e\,\mathrm{d}\lambda = 0 = \int_0^\theta f^+_e\,\mathrm{d}\lambda - \int_{\tau_e}^{\theta+\tau_e} f^-_e\,\mathrm{d}\lambda \\ &= \int_0^\theta f^+_e(\xi) - f^-_e(\xi+\tau_e)\,\mathrm{d}\lambda(\xi). \end{aligned}$$

On the other hand, $\theta > 0$ implies $H > \tau_e$ and thus $\int_0^{\tau_e} f^-_e\,\mathrm{d}\lambda = 0$. We infer

$$\begin{aligned} q_e(\theta) &= \int_0^\theta f^+_e(\xi) - f^-_e(\xi+\tau_e)\,\mathrm{d}\lambda(\xi) - \int_0^{\tau_e} f^-_e\,\mathrm{d}\lambda \\ &= \int_0^\theta f^+_e(\xi) - f^-_e(\xi+\tau_e)\,\mathrm{d}\lambda(\xi). \end{aligned}$$

□

The last proposition in this chapter shows that flow cannot occur at arbitrary places in the network.

Proposition 3.3.5. *Let f be a feasible dynamic flow and let $i \in I$ be a commodity in a network in which any cycle C has a strictly positive transit time $\sum_{e\in C} \tau_e > 0$. Let $V' \subseteq V$ be a set of nodes with $s_i \in V'$ for which the inflow rate $f^+_{i,e}$ vanishes almost everywhere for any $e \in \delta^+(V')$. Then $f^+_{i,e}$ vanishes for all $e = vw \in E$ with $w \notin V'$.*

Proof. Let E' denote the set of edges $\{e = vw \in E \mid w \notin V'\}$ and let $\theta \in \mathbb{R}$ be the infimum of all $\xi \in \mathbb{R}$ for which there exists some $e \in E'$ with $F^+_{i,e}(\xi)$. Then there is an edge $e_0 = v_0 w$ with $F^+_{i,e}(\xi) > 0$ for any $\xi > \theta$. By the flow conservation in v_0, there exists an edge $e_1 = v_1 v_0$ with $F^-_{i,e_1}(\xi) > 0$ and hence $F^+_{i,e_1}(\xi - \tau_{e_1}) > 0$ for any $\xi > \theta$. The statement's requirement implies $v_1 \notin V'$, and therefore $e_1 \in E'$ and $\tau_{e_1} = 0$ must hold. Continuing this approach leads to an infinite path of edges within E' each with a zero transit time. As there are only finitely many edges, this sequence must contain a cycle of zero transit time; a contradiction. □

3.4 The Existence and Uniqueness of Deterministic Flows

In this section, we discuss the existence and uniqueness of deterministic edge outflow rates for any given family of inflow rates $(f_{i,e}^+)_{i,e}$.

In a first step, we discuss the unique existence of functions that behave similar to f_e^- in (F1) and (F2). More specifically, let $f \in \mathcal{R}$ be any rate function and let $\nu : \mathbb{R} \to \mathbb{R}_{\geq 0}$ be any locally integrable function, which acts as the time-dependent capacity of an edge. To reflect our model, we will later set ν to be constant. For now, however, we want to find a function $g \in \mathcal{R}$ that fulfills $g(\theta) \leq \nu(\theta)$ and

$$g(\theta) = \begin{cases} \nu(\theta), & \text{if } \int_0^\theta f - g \,\mathrm{d}\lambda > 0, \\ f(\theta), & \text{otherwise}, \end{cases} \tag{3.2}$$

for almost all $\theta \in \mathbb{R}$. Here, $\int_0^\theta f - g \,\mathrm{d}\lambda$ should be interpreted as the length of a queue operating at a certain capacity rate ν with inflow rate f and outflow rate g.

Cominetti et al. have given a representation of the queue length for feasible flows in [7, Section 2.2] that only depends on the edge's inflow and capacity rate. In the following, we formally prove a slightly more general variant of their statement.

Lemma 3.4.1. *Given $f \in \mathcal{R}$ and $g \in L^1_{\mathrm{loc}}(\mathbb{R}, \mathbb{R}_{\geq 0})$ that fulfill $g \leq \nu$ and Equation (3.2) almost everywhere on $(-\infty, H)$ for some $H \in \mathbb{R}_{\geq 0} \cup \{\infty\}$, then for all $\theta < H$ it holds*

$$\int_0^\theta f - g \,\mathrm{d}\lambda = \max_{\xi \leq \theta} \int_\xi^\theta f - \nu \,\mathrm{d}\lambda$$

and $\xi^* := \max\{\xi \leq \theta \mid \int_0^\xi f - g \,\mathrm{d}\lambda = 0\}$ is such a maximizer. Moreover, g vanishes almost everywhere on $(-\infty, 0)$, implying $g \in \mathcal{R}$.

Proof. We first show that $q(\theta) := \int_0^\theta f - g \,\mathrm{d}\lambda$ is never negative for $\theta < H$. Assume the contrary and let $\theta < H$ fulfill $q(\theta) < 0$. If $\theta > 0$, choose $\xi^* := \max\{\xi \leq \theta \,|\, q(\xi) = 0\}$ to be the latest time before θ at which the queue was empty. Because q vanishes at time 0, this maximum ξ^* exists. The continuity of q implies that q is non-positive on $[\xi^*, \theta]$. Applying Equation (3.2) on (ξ^*, θ) yields

$$q(\theta) = \int_0^\theta f - g \,\mathrm{d}\lambda = \int_0^{\xi^*} f - g \,\mathrm{d}\lambda + \int_{\xi^*}^\theta f - f \,\mathrm{d}\lambda = q(\xi^*) = 0,$$

a contradiction. For $\theta < 0$ we have $q(\theta) \leq 0$ as f vanishes on $(-\infty, 0)$. Because of Equation (3.2), $g(\theta) = f(\theta) = 0$ holds for almost all $\theta < 0$, showing $q(\theta) = 0$ for $\theta < 0$.

We now prove the desired equation for all $\theta < H$. For $\theta < 0$ we have $q(\theta) = 0$. Furthermore, the right-hand side equates to $\max_{\xi \leq \theta} \int_\xi^\theta -\nu \,\mathrm{d}\lambda = 0$. For $\theta > 0$ we again define $\xi^* := \max\{\xi \leq \theta \mid q(\xi) = 0\}$ to be the latest time before θ at which the queue is empty. Then, by definition, q is positive on (ξ^*, θ) and hence Equation (3.2) implies $g = \nu$ almost everywhere on (ξ^*, θ). This shows

$$q(\theta) = q(\xi^*) + \int_{\xi^*}^\theta f - \nu \,\mathrm{d}\lambda = \int_{\xi^*}^\theta f - \nu \,\mathrm{d}\lambda.$$

It remains to show that $\int_\xi^\theta f - \nu \,\mathrm{d}\lambda \le q(\theta)$ holds for any other $\xi \in [0,\theta]$. In the case $\xi \le \xi^*$, it follows $g \overset{a.e.}{\le} \nu$ which implies

$$\int_\xi^\theta f - \nu \,\mathrm{d}\lambda \le \int_\xi^\theta f - g \,\mathrm{d}\lambda = q(\theta) - q(\xi) \le q(\theta).$$

For $\xi > \xi^*$ we use Equation (3.2) and $q(\xi^*) = 0$ to get

$$\int_\xi^\theta f - \nu \,\mathrm{d}\lambda = q(\theta) - \int_{\xi^*}^\xi f - \nu \,\mathrm{d}\lambda - q(\xi^*) = q(\theta) - q(\xi) \le q(\theta).$$

□

We transfer this proposition to our scenario to obtain the following representation of the queue length function.

Corollary 3.4.2. *Let a dynamic flow $f = (f^+, f^-)$ fulfill properties (F1) and (F2) on an edge e up to some time $H \in \mathbb{R}_{\ge 0} \cup \{\infty\}$. Then*

$$q_e(\theta) = \max_{\xi \le \theta} \int_0^\xi f_e^+ - \nu_e \,\mathrm{d}\lambda$$

holds for all $\theta \le H - \tau_e$ and $\xi^* := \max\{\xi \le \theta \mid q_e(\xi) = 0\}$ is such a maximizer.

Proof. We define $g_e(\theta) := f_e^-(\theta + \tau_e)$. Proposition 3.3.4 shows $q_e(\theta) = \int_0^\theta f_e^+ - g_e \,\mathrm{d}\lambda$ for all $\theta < H$. Moreover, g_e fulfills $g_e \le \nu$ by (F1) and Equation (3.2) by (F2) almost everywhere on $(-\infty, H - \tau_e)$. Applying Lemma 3.4.1 yields the desired statement. □

Lemma 3.4.3. *For a locally integrable function γ with $\gamma(\theta) \le 0$ for almost all $\theta < 0$, the function $\Gamma(\theta) := \max_{\xi \le \theta} \int_\xi^\theta \gamma \,\mathrm{d}\lambda$ is absolutely continuous on any interval $[\alpha, \beta]$ with $\alpha < \beta$. Moreover, Γ is almost everywhere differentiable with derivative*

$$\Gamma'(\theta) = \begin{cases} \gamma(\theta), & \textit{if } \Gamma(\theta) > 0, \\ 0, & \textit{otherwise,} \end{cases}$$

and $\gamma(\theta) \le 0$ holds for almost all θ with $\Gamma(\theta) = 0$.

Proof. We begin by showing that Γ is absolutely continuous on any interval $[\alpha, \beta]$. For the most part, we follow the proof of the *Fundamental Theorem of Lebesgue Integral Calculus* as explained in [3, Theorem 4.4.1].

As γ is locally integrable, we have $\int_\alpha^\beta |\gamma| \,\mathrm{d}\lambda < \infty$. By [3, Proposition 2.5.8], for any $\varepsilon > 0$ there exists $\delta > 0$ such that for any measurable set A with $\lambda(A) < \delta$ we have $\int_A |\gamma| \,\mathrm{d}\lambda < \varepsilon$. Now, let $((a_i, b_i))_i$ be a finite sequence of pairwise disjoint intervals with $\sum_i (b_i - a_i) < \delta$. We have to show that $\sum_i |\Gamma(b_i) - \Gamma(a_i)| < \varepsilon$.

Claim 3.4.4. *For any $a < b$ we have $\Delta := |\Gamma(b) - \Gamma(a)| \le \int_a^b |\gamma| \,\mathrm{d}\lambda$.*

Proof. Let $\xi_a \leq a$ and $\xi_b \leq b$ both be maximal with $\Gamma(a) = \int_{\xi_a}^a \gamma \, \mathrm{d}\lambda$ and $\Gamma(b) = \int_{\xi_b}^b \gamma \, \mathrm{d}\lambda$. In the case $\Gamma(a) \geq \Gamma(b)$, we utilize the definition of Γ to get

$$\Delta = \int_{\xi_a}^a \gamma \, \mathrm{d}\lambda - \max_{\xi \leq b} \int_{\xi}^b \gamma \, \mathrm{d}\lambda \leq \int_{\xi_a}^a \gamma \, \mathrm{d}\lambda - \int_{\xi_a}^b \gamma \, \mathrm{d}\lambda = \int_a^b -\gamma \, \mathrm{d}\lambda \leq \int_a^b |\gamma| \, \mathrm{d}\lambda.$$

Now, we analyze the case $\Gamma(b) > \Gamma(a)$. If $\xi_b \leq a$ holds true, we can show $\xi_b = \xi_a$: If we assume $\xi_b < \xi_a$, then $\int_{\xi_b}^b \gamma \, \mathrm{d}\lambda > \int_{\xi_a}^b \gamma \, \mathrm{d}\lambda$ holds by the maximality of ξ_b, which implies the contradiction $\int_{\xi_b}^a \gamma \, \mathrm{d}\lambda > \int_{\xi_a}^a \gamma \, \mathrm{d}\lambda$. Assuming $\xi_b > \xi_a$ we get $\int_{\xi_a}^a \gamma \, \mathrm{d}\lambda > \int_{\xi_b}^a \gamma \, \mathrm{d}\lambda$ by the maximality of ξ_a and thus $\int_{\xi_a}^b \gamma \, \mathrm{d}\lambda > \int_{\xi_b}^b \gamma \, \mathrm{d}\lambda$; another contradiction. With this observation, we deduce

$$\Delta = \int_{\xi_b}^b \gamma \, \mathrm{d}\lambda - \int_{\xi_a}^a \gamma \, \mathrm{d}\lambda = \int_a^b \gamma \, \mathrm{d}\lambda \leq \int_a^b |\gamma| \, \mathrm{d}\lambda.$$

Now the case $\Gamma(b) > \Gamma(a)$ with $\xi_b > a$ remains. Here, we have

$$\Delta = \int_{\xi_b}^b \gamma \, \mathrm{d}\lambda - \Gamma(a) \leq \int_{\xi_b}^b \gamma \, \mathrm{d}\lambda \leq \int_a^b |\gamma| \, \mathrm{d}\lambda. \qquad \blacksquare$$

The claim above implies the absolute continuity of Γ, because it shows that

$$\sum_i |\Gamma(b_i) - \Gamma(a_i)| \leq \sum_i \int_{a_i}^{b_i} |\gamma| \, \mathrm{d}\lambda = \int_{\bigcup_i [a_i, b_i]} |\gamma| \, \mathrm{d}\lambda < \varepsilon.$$

Next, we prove that on closed intervals $[a, b]$ with $a < b$ on which Γ is strictly positive, there exists a common $\xi \leq a$ such that $\Gamma(\theta) = \int_{\xi}^{\theta} \gamma \, \mathrm{d}\lambda$ holds for all $\theta \in [a, b]$. Let $\xi_a \in \arg\max_{\xi \leq a} \int_{\xi}^a \gamma \, \mathrm{d}\lambda$. We define

$$\theta := \max \left\{ \theta \in [a, b] \,\middle|\, \forall \theta' \in [a, \theta] : \Gamma(\theta') = \int_{\xi_a}^{\theta'} \gamma \, \mathrm{d}\lambda \right\}.$$

If $\theta = b$, the claim follows directly. If we assume $\theta < b$, there exists a sequence $(\theta_k)_{k \in \mathbb{N}}$ with $\theta_k \in (\theta, b]$, $\lim_{k \to \infty} \theta_k = \theta$ and $\Gamma(\theta_k) \neq \int_{\xi_a}^{\theta_k} \gamma \, \mathrm{d}\lambda$. By definition of Γ we even have $\Gamma(\theta_k) > \int_{\xi_a}^{\theta_k} \gamma \, \mathrm{d}\lambda$. Moreover, this implies $\Gamma(\theta_k) = \max_{\xi \in (\theta, \theta_k]} \int_{\xi}^{\theta_k} \gamma \, \mathrm{d}\lambda$, because for $\xi \leq \theta$ we have

$$\int_{\xi}^{\theta_k} \gamma \, \mathrm{d}\lambda \leq \Gamma(\theta) + \int_{\theta}^{\theta_k} \gamma \, \mathrm{d}\lambda = \int_{\xi_a}^{\theta_k} \gamma \, \mathrm{d}\lambda < \Gamma(\theta_k).$$

Therefore, $\Gamma(\theta_k)$ converges to 0 for $k \to \infty$. Because Γ is continuous we derive the contradiction $\Gamma(\theta) = 0$.

This shows that Γ is almost everywhere differentiable on the set $\{\theta \in \mathbb{R} \,|\, \Gamma(\theta) > 0\}$ with $\Gamma'(\theta) = \gamma(\theta)$. Being absolutely continuous on any interval implies that Γ is almost everywhere differentiable. This means that Γ is also almost everywhere differentiable in $A := \{\theta \mid \Gamma(\theta) = 0\}$. As Γ is never (strictly) negative, the derivative on A can only be 0 if it exists.

It remains to show that $\gamma(\theta) \leq 0$ holds for almost all $\theta \in A$. By the Lebesgue differentiation theorem (Theorem 3.2.3), the function $h(x) := \int_0^x \gamma \, \mathrm{d}\lambda$ is almost everywhere differentiable with derivative γ. For $\theta \in A$ with $\Gamma'(\theta) = 0$ and $h'(\theta) = \gamma(\theta)$ we have

$$0 = \Gamma'(\theta) = \lim_{x \searrow \theta} \frac{\Gamma(x) - \Gamma(\theta)}{x - \theta} \geq \lim_{x \searrow \theta} \frac{\int_\theta^x \gamma \,\mathrm{d}\lambda - \int_\theta^\theta \gamma \,\mathrm{d}\lambda}{x - \theta} = \lim_{x \searrow \theta} \frac{h(x) - h(\theta)}{x - \theta} = \gamma(\theta),$$

which concludes the proof. $\square$

Lemma 3.4.5. *Let $f \in \mathcal{R}$ be an inflow rate and let $\nu \in L^1_{\mathrm{loc}}(\mathbb{R}, \mathbb{R}_{\geq 0})$ be a non-negative locally integrable capacity rate. Then the following statements hold true.*

(i) There exists a rate function $g \in \mathcal{R}$ that fulfills $g \leq \nu$ and Equation (3.2) *almost everywhere.*

(ii) If another $\tilde{g} \in L^1_{\mathrm{loc}}(\mathbb{R}, \mathbb{R}_{\geq 0})$ fulfills $\tilde{g} \leq \nu$ and Equation (3.2) *almost everywhere on $(-\infty, H)$ for some $H \in \mathbb{R}_{\geq 0} \cup \{\infty\}$, then $\tilde{g}|_{(-\infty,H)} \overset{a.e.}{=} g|_{(-\infty,H)}$ holds true.*

This means, there exists a unique $g \in L^1_{\mathrm{loc}}(\mathbb{R}, \mathbb{R}_{\geq 0})$ fulfilling $g \leq \nu$ and Equation (3.2) *almost everywhere. Moreover, g vanishes almost everywhere on $(-\infty, 0)$, implying $g \in \mathcal{R}$.*

Proof. (i). By Lemma 3.4.3, the function $q(\theta) := \max_{\xi \leq \theta} \int_\xi^\theta f - \nu \,\mathrm{d}\lambda$ is almost everywhere differentiable. We define $g(\theta) := f(\theta) - q'(\theta)$ whenever $q'(\theta)$ is defined, and $g(\theta) := 0$ otherwise. Then, we have $q(\theta) = \int_0^\theta f - g \,\mathrm{d}\lambda$ and g automatically fulfills Equation (3.2) almost everywhere.

It remains to show that g is bounded from above by ν almost everywhere. For $q(\theta) > 0$, this is given by Equation (3.2) already. Hence, we can concentrate on θ with $q(\theta) = 0$. However, Lemma 3.4.3 states $q'(\theta) = 0$ and $f(\theta) \leq \nu(\theta)$ for almost all such θ.

(ii). Assume there is a locally integrable function $\tilde{g}$ fulfilling $\tilde{g} \leq \nu$ and Equation (3.2) almost everywhere on $(-\infty, H)$. Applying Lemma 3.4.1 for any $\theta < H$ yields

$$\int_0^\theta f - \tilde{g} \,\mathrm{d}\lambda = \max_{\xi \leq \theta} \int_\xi^\theta f - \nu \,\mathrm{d}\lambda = \int_0^\theta f - g \,\mathrm{d}\lambda.$$

Hence, g and $\tilde{g}$ coincide almost everywhere on $(-\infty, H)$ (cf. [3, Theorem 4.4.2]). $\square$

So far, we have proved that given any total inflow rate function f_e^+ there exists a rate function g fulfilling $g \leq \nu$ and Equation (3.2). In our scenario this shows that the total outflow rate function f_e^- is uniquely determined by the total inflow rate function and properties (F1) and (F2). We will now extend this result to multicommodity inflow and outflow rates in such a way that property (F3) holds too.

Theorem 3.4.6. *Let I be a finite set of commodities and let $(f_{i,e}^+)_{i \in I} \in \mathcal{R}^I$ be a family of inflow rates on an edge e. Then the following statements hold true.*

(i) There exist deterministic outflow rates $(f_{i,e}^-)_{i \in I} \in \mathcal{R}^I$ corresponding to $(f_{i,e}^+)_{i \in I}$.

(ii) If $(h_{i,e}^-)_{i \in I} \in L^1_{\mathrm{loc}}(\mathbb{R}, \mathbb{R}_{\geq 0})^I$ also fulfill (F1), (F2) and (F3) almost everywhere on $(-\infty, H)$ for some $H \in \mathbb{R}_{\geq 0} \cup \{\infty\}$, then $h_{i,e}^-|_{(-\infty,H)} \overset{a.e.}{=} f_{i,e}^-|_{(-\infty,H)}$ holds for all $i \in I$.

We call $(f_{i,e}^-)_{i \in I}$ the (unique) deterministic outflow rates corresponding to *$(f_{i,e}^+)_{i \in I}$. Given inflow rates on all edges $(f^+)_{i \in I, e \in E}$ we call (f^+, f^-) the* (unique) deterministic flow corresponding to *$(f_{i,e}^+)_{i \in I, e \in E}$.*

Proof. (i). Let $f_e^+ \coloneqq \sum_{i \in I} f_{i,e}^+ \in \mathcal{R}$ denote the total rate function, for which Lemma 3.4.5 yields the unique existence of a function $g_e \in \mathcal{R}$ fulfilling $g_e \leq \nu_e$ and Equation (3.2) almost everywhere. We define

$$f_{i,e}^-(\theta) \coloneqq \begin{cases} g_e(\theta - \tau_e) \cdot \frac{f_{i,e}^+(\xi_\theta)}{f_e^+(\xi_\theta)}, & \text{if } f_e^+(\xi_\theta) > 0, \\ 0, & \text{otherwise,} \end{cases}$$

where $\xi_\theta \coloneqq \max\{\xi \mid \Theta(\xi) \coloneqq \xi + \tau_e + \int_0^\xi f_e^+ - g_e \,\mathrm{d}\lambda / \nu_e = \theta\}$. We aim to show

$$f_e^-(\theta) \coloneqq \sum_{i \in I} f_{i,e}^-(\theta) = g_e(\theta - \tau_e)$$

for almost all $\theta \in \mathbb{R}$. Once this is proven, the queue $q_e(\theta)$ coincides with $\int_0^\theta f_e^+ - g_e \,\mathrm{d}\lambda$ (and therefore T_e coincides with Θ), which can be seen by applying Lemma 3.4.1:

$$\begin{aligned} q_e(\theta) \coloneqq \int_0^\theta f_e^+ \,\mathrm{d}\lambda - \int_0^{\theta+\tau_e} f_e^- \,\mathrm{d}\lambda &= -\int_0^{-\tau_e} f_e^+ - g_e \,\mathrm{d}\lambda + \int_0^\theta f_e^+ - g_e \,\mathrm{d}\lambda \\ &= -\max_{\xi \leq -\tau_e} \int_0^\xi f_e^+ - \nu_e \,\mathrm{d}\lambda + \int_0^\theta f_e^+ - g_e \,\mathrm{d}\lambda = \int_0^\theta f_e^+ - g_e \,\mathrm{d}\lambda. \end{aligned}$$

Therefore, the defined flow rates fulfill properties (F1), (F2) and (F3) through $g_e \leq \nu_e$ and Equation (3.2) almost everywhere.

If $f_e^+(\xi_\theta)$ is positive, then the claim follows directly from the definition of $f_{i,e}^-(\theta)$. Therefore, we only analyze θ with $f_e^+(\xi_\theta) = 0$.

The function Θ is non-decreasing and absolutely continuous on any interval $[a, b]$. We define the Lebesgue-null set

$$A \coloneqq \left\{\xi \in \mathbb{R} \;\middle|\; \Theta'(\xi) \neq 1 + (f_e^+(\xi) - g_e(\xi))/\nu_e \text{ or Equation (3.2) does not hold in } \xi\right\}$$

and the set of points where the derivative of Θ vanishes

$$B \coloneqq \{\xi \in \mathbb{R} \mid \Theta \text{ is differentiable in } \xi \text{ with } \Theta'(\xi) = 0\}.$$

By the Luzin N property described in Proposition 3.2.5 all absolute continuous functions map null sets to null sets, which implies $\lambda(\Theta(A)) = 0$. Moreover, Sard's theorem (Theorem 3.2.8) states that $\Theta(B)$ is also a Lebesgue-null set. Therefore, $\Theta(A) \cup \Theta(B)$ is a Lebesgue-null set, and we take $\theta \in \Theta(A)^c \cap \Theta(B)^c$. Since $\Theta(\xi_\theta) = \theta$, we know $\xi_\theta \notin A \cup B$. The monotonicity of Θ together with $\xi_\theta \notin A \cup B$ implies

$$0 < \Theta'(\xi_\theta) = 1 + \frac{f_e^+(\xi_\theta) - g_e(\xi_\theta)}{\nu_e},$$

which is equivalent to $g_e(\xi_\theta) < f_e^+(\xi_\theta) + \nu_e = \nu_e$. By Equation (3.2) this can only be the case, if $\int_0^{\xi_\theta} f_e^+ - g_e \,\mathrm{d}\lambda = 0$ and $g_e(\xi_\theta) = f_e^+(\xi_\theta) = 0$ hold true. The observation $\theta = \Theta(\xi_\theta) = \xi_\theta + \tau_e$ concludes the proof of $\sum_{i \in I} f_{i,e}^-(\theta) = 0 = g_e(\theta - \tau_e)$.

(ii). Let $(h_{i,e}^-)_{i \in I} \in L^1_{\mathrm{loc}}(\mathbb{R}, \mathbb{R}_{\geq 0})$ fulfill properties (F1), (F2) and (F3) in the place of $(f_{i,e}^-)_{i \in I}$ almost everywhere on $(-\infty, H)$ for some $H \in \mathbb{R}_{\geq 0} \cup \{\infty\}$. We want to show that

$h^-_{i,e}$ coincides almost everywhere on $(-\infty, H)$ with $f^-_{i,e}$ for all $i \in I$. The following claim is the essential result for this proof:

Claim 3.4.7. *The functions h^-_e and f^-_e coincide almost everywhere on $(-\infty, H)$.*

Proof. Proposition 3.3.4 states that $q^h_e(\theta) = \int_0^\theta f^+_e - \tilde{g}_e \,\mathrm{d}\lambda$ holds for all $\theta < H - \tau_e$ with $\tilde{g}_e(\theta) := h^-_e(\theta + \tau_e) = \sum_{i \in I} h^-_{i,e}(\theta + \tau_e)$. Now, (F2) implies that $\tilde{g}_e$ fulfills Equation (3.2) almost everywhere on $(-\infty, H - \tau_e)$. The same holds true for g_e with $g_e(\theta) = f^-_e(\theta + \tau_e)$. (F1) implies that both functions are bounded from above by ν_e on $(-\infty, H - \tau_e)$. Lemma 3.4.5 states that any two functions with these properties are identical almost everywhere on $(-\infty, H - \tau_e)$; therefore h^-_e and f^-_e coincide almost everywhere on $(-\infty, H)$. ■

By the above claim, q^h_e and T^h_e also coincide with q^f_e and T^f_e on $(-\infty, H - \tau_e)$, respectively. Property (F3) now implies $h^-_{i,e}(\theta) = f^-_{i,e}(\theta)$ for almost all $\theta < H$ and all $i \in I$. □

As the final property in this chapter we show that, if the inflow rates of two deterministic flows match up to some time H, then their outflow rates match up to time $T_e(H)$, too.

Lemma 3.4.8. *Let (f^+, f^-), (g^+, g^-) be two flows that are deterministic on an edge $e \in E$. If their inflow rates on e coincide up to some time $H \in \mathbb{R}_{\geq 0}$, i.e.*

$$\forall i \in I: \quad f^+_{i,e}|_{(-\infty,H)} \overset{a.e.}{=} g^+_{i,e}|_{(-\infty,H)},$$

then the corresponding queue functions q^f_e and q^g_e and hence the exit times T^f_e and T^g_e coincide on $(-\infty, H]$. Moreover, their outflow rates on e coincide up to time $T_e(H)$, i.e.

$$\forall i \in I: \quad f^-_{i,e}|_{(-\infty,T_e(H))} \overset{a.e.}{=} g^-_{i,e}|_{(-\infty,T_e(H))}.$$

Proof. By Corollary 3.4.2, the queue functions $q^f_e|_{(-\infty,H]}$ and $q^g_e|_{(-\infty,H]}$ only depend on $f^+_e|_{(-\infty,H)}$ and $g^+_e|_{(-\infty,H)}$, respectively. Therefore, both queue functions and both exit times are identical up to time H.

By Proposition 3.3.2 both queue functions q^f_e and q^g_e are positive on $(H, H + q_e(H)/\nu_e)$. Hence, for $\theta \in (H + \tau_e, T_e(H))$, the queues of both flows are non-empty at time $\theta - \tau_e$ implying $g^-_e(\theta) = \nu_e = f^-_e(\theta)$ for almost all $\theta \in (H + \tau_e, T_e(H))$ according to property (F2). For almost all $\theta < H + \tau_e$, property (F2) implies

$$g^-_e(\theta) = \begin{Bmatrix} \nu_e, & \text{if } q_e(\theta - \tau_e) > 0, \\ g^+_e(\theta - \tau_e) = f^+_e(\theta - \tau_e), & \text{otherwise,} \end{Bmatrix} = f^-_e(\theta).$$

By the Luzin N property of T_e, the set of points $\theta < T_e(H)$ that fulfill the condition

$$\exists\, \xi \leq H: \; T_e(\xi) = \theta \,\wedge\, \Big(f^-_e(\theta) \neq g^-_e(\theta) \,\vee\, \exists i \in I: \; f^+_{i,e}(\xi) \neq g^+_{i,e}(\xi) \Big)$$

is a Lebesgue-null set. For almost all other $\theta < T_e(H)$ we define $\xi := \min\{\xi \leq H \mid T_e(\xi) = \theta\}$, and we conclude using property (F3)

$$g^-_{i,e}(\theta) = \begin{Bmatrix} g^-_e(\theta) \cdot \frac{g^+_{i,e}(\xi)}{g^+_e(\xi)}, & \text{if } g^+_e(\xi) > 0, \\ 0, & \text{otherwise,} \end{Bmatrix} = \begin{Bmatrix} f^-_e(\theta) \cdot \frac{f^+_{i,e}(\xi)}{f^+_e(\xi)}, & \text{if } f^+_e(\xi) > 0, \\ 0, & \text{otherwise,} \end{Bmatrix} = f^-_{i,e}(\theta). \quad \square$$

4 Dynamic Prediction Equilibria

In this chapter, a game theoretical structure is added to the flows introduced in the previous section. We describe the behavior of the infinitesimally small particles by interpreting them as agents belonging to one of finitely many commodities. An agent shares the same origin and destination nodes with all other agents of his commodity. Moreover, all agents of a commodity use the same prediction method which can be invoked at any time $\bar{\theta}$ to retrieve a (not necessarily accurate) traffic forecast of all edges. This prediction method receives the historical data of the total traffic flow up to time $\bar{\theta}$ as an input.

We impose an important assumption on the agents' behavior: The agents act as if the forecasts retrieved by their prediction methods accurately describe the future traffic (although it may not). However, whenever there is a new routing decision to take, a new prediction is triggered. Therefore, when an agent arrives at an intermediate node at some time θ, and they have to choose which of the outgoing edges to enter, they will retrieve a new forecast at time $\bar{\theta} = \theta$ and calculate a shortest path to their destination according to the new traffic forecast. Then, they decide to enter the first edge on this shortest path. If all agents adapt to this behavior, a *dynamic prediction equilibrium (DPE)* emerges. DPE were first introduced in [13]. Here, we use the same concepts, although with a more general class of predictors.

In the first section of this chapter, the behavioral model as well as the notions of a predictor and a dynamic prediction equilibrium are introduced. To illustrate the definition, an example of such an equilibrium is analyzed in Section 4.2. Section 4.3 provides a recapitulation of some mathematical tools necessary to prove the existence of dynamic prediction equilibria. The existence of DPE is then proven in Section 4.4 under quite mild assumptions on the predictors. To apply the existence theorem more easily, Section 4.5 shows sufficient conditions for the regularity of predictors. Section 4.6 introduces several example predictors, that were considered in this thesis, and their compatibility with the existence theorem is analyzed. Finally, a comparison with other popular models of equilibria is carried out in Section 4.7.

4.1 Definition

Predicting future travel times in Vickrey's fluid queuing model comes done to predicting the queue length functions q_e for all $e \in E$. This inspires the following definition of a *predictor*.

Definition 4.1.1. A *predictor* of a commodity $i \in I$ is a function

$$\hat{q}_i : \mathbb{R} \times \mathbb{R} \times \mathcal{F} \to \mathbb{R}^E_{\geq 0}, \quad (\theta, \bar{\theta}, f) \mapsto \left(\hat{q}_{i,e}(\theta, \bar{\theta}, f)\right)_{e \in E},$$

where $\mathcal{F} := (\mathcal{R} \times \mathcal{R})^{I \times E}$ denotes the set of all dynamic flows. We call $\hat{q}_{i,e}(\theta, \bar{\theta}, f)$ the *predicted queue length* of edge e at time θ as predicted at time $\bar{\theta}$ with the historical flow f.

Based on that, we define the *predicted dynamic cost* $\hat{c}_{i,e}(\theta, \bar{\theta}, f) := \hat{q}_{i,e}(\theta, \bar{\theta}, f)/\nu_e$ at time θ as predicted at time $\bar{\theta}$. Furthermore, $\hat{T}_{i,e}(\theta, \bar{\theta}, f) := \theta + \hat{c}_{i,e}(\theta, \bar{\theta}, f)$ denotes the *predicted exit time* of edge e when e at time θ as predicted at time $\bar{\theta}$. Given a path $P = e_1 \cdots e_k$ we define the *predicted exit time* when entering path P at time θ as predicted at time $\bar{\theta}$ as

$$\hat{T}_{i,P}(\theta, \bar{\theta}, f) := \left(\hat{T}_{e_k}(\,\cdot\,, \bar{\theta}, f) \circ \cdots \circ \hat{T}_{e_1}(\,\cdot\,, \bar{\theta}, f)\right)(\theta).$$

As in Chapter 2, let $\mathcal{P}_{v,w}$ denote the set of all simple paths from v to w. We define $\hat{l}_{i,v}(\theta, \bar{\theta}, f) := \min_{P \in \mathcal{P}_{v,t_i}} \hat{T}_{i,P}(\theta, \bar{\theta}, f)$ as the *earliest predicted arrival time* at the destination t_i when starting in $v \in V_i$ at time θ as predicted at time $\bar{\theta}$. A path that attains this minimum is called a *predicted shortest v-t_i-path* at time θ as predicted at time $\bar{\theta}$.

We can only restrict ourselves to simple paths in the above definition, if waiting at a node, or traveling through a short cycle, is never helpful. This can be realized by assuming the FIFO-compatibility of predictors:

Definition 4.1.2. A predictor $\hat{q}_i$ is *FIFO-compatible* if for all $\bar{\theta} \in \mathbb{R}$ and for all deterministic dynamic flows $f \in \mathcal{F}$ the predicted dynamic costs $(\hat{c}_{i,e}(\,\cdot\,, \bar{\theta}, f))_{e \in E}$ are FIFO-ordered.

It is worth noting that a predictor $\hat{q}_i$ is FIFO-compatible if and only if

$$\hat{q}_{i,e}(\theta_2, \bar{\theta}, f) - \hat{q}_{i,e}(\theta_1, \bar{\theta}, f) \geq (\theta_2 - \theta_1) \cdot (-\nu_e)$$

holds for any $\theta_1 < \theta_2$, prediction time $\bar{\theta}$ and deterministic flow f. For a FIFO-compatible predictor the dynamic variant of the triangle inequality in Proposition 2.2.2 states that

$$\hat{l}_{i,v}(\theta, \bar{\theta}, f) \leq \hat{l}_{i,w}(\hat{T}_{vw}(\theta, \bar{\theta}, f), \bar{\theta}, f)$$

holds for all $\theta, \bar{\theta} \in \mathbb{R}$ and feasible flows $f \in \mathcal{F}$.

Definition 4.1.3. Given a predictor $\hat{q}_i$, the *predicted delay when entering edge $e = vw \in E_i$* at time θ as predicted at time $\bar{\theta}$ using the historical flow data f is given as

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

An edge $e = vw \in E_i$ is called *active* for commodity $i \in I$ at time $\theta \in \mathbb{R}$ as predicted at time $\bar{\theta}$, if the predicted delay $\hat{\Delta}_{i,e}(\theta, \bar{\theta}, f) = 0$ vanishes. All edges for commodity i that are active at time θ as predicted at time $\bar{\theta}$ are collected in the set $\hat{E}_i(\theta, \bar{\theta}, f)$.

In other words, an edge is active at some time θ if it lies on a shortest v-t_i-path at time θ as predicted at time $\bar{\theta}$. With these fundamental definitions we are able to define the state of equilibrium in which the behavior of agents matches the prior description.

Definition 4.1.4. A pair $(\hat{q}, f)$ of a set of predictors $\hat{q} = (\hat{q}_i)_{i \in I}$ and a dynamic flow f is a *(partial) dynamic prediction equilibrium 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

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

For $H = \infty$ the pair $(\hat{q}, f)$ is called a *dynamic prediction equilibrium (DPE)*. Moreover, we call f a *dynamic prediction flow with respect to the predictor $\hat{q}$*.

A dynamic prediction equilibrium can be interpreted as follows: If agents of a commodity i enter an edge $e = vw$ at some time θ, expressed by $f^+_{i,e}(\theta) > 0$, then this edge must lie on a shortest v-t_i-path at time θ as predicted at time θ.

4.2 Example of a Dynamic Prediction Equilibrium

Before we analyze the properties of dynamic prediction equilibria, we first take a look at an example of such an equilibrium to get a better intuition of these mathematical objects. For that we introduce a first, simple predictor that demonstrates the capabilities of the model.

Definition 4.2.1. Given some prediction horizon $H \in \mathbb{R}_{>0} \cup \{\infty\}$, the *linear predictor* $\hat{q}^{\mathrm{L}}_{i,e}$ is defined as

$$\hat{q}^{\mathrm{L}}_{i,e}(\theta, \bar{\theta}, q) \coloneqq \Big(q_e(\bar{\theta}) + \partial_- q_e(\bar{\theta}) \cdot \min\{\theta - \bar{\theta}, H\} \Big)^+,$$

where $(x)^+ \coloneqq \max\{x, 0\}$ denotes the positive part of $x \in \mathbb{R}$ and $\partial_- q_e$ denotes the left-side derivative of q_e.

In other words, the linear predictor states that the queue length grows with the current derivative up to some prediction horizon H. Starting from H it assumes a constant course.

For the following example, we will assume that no prediction horizon is given, i.e. $H = \infty$. The example network is shown in Figure 4.1(a). All edges e except edge ut have a capacity of $\nu_e = 2$ and edge ut has a capacity of 1. Moreover, all edges e except vw have a transit time of $\tau_e = 2$. Edge vw has a transit time of $\tau_{vw} = 2$.

There is a single commodity with a constant network inflow rate of $u|_{[0,\infty)} \equiv 2$, source node s and sink node t. As the edge ut is the only edge with a capacity smaller than 2, it is the only edge that might build up a queue. This implies that for any edge e other than ut, the predicted queue length can be determined as $\hat{q}^L_e(\theta, \bar{\theta}, f) = 0$.

Particles starting at the source node s have to decide between the upper path sut and the lower path $svwt$. This decision is based on the predicted exit times of both paths. Because no queues will ever be predicted in the lower path, the predicted exit time simplifies to $\hat{T}_{svwt}(\theta, \bar{\theta}, f) = \theta + 4$. For the upper path, particles starting at time θ first travel through the first edge with a transit time of 1, and then might have to enqueue at edge ut with a predicted queue length $\hat{q}^L_{ut}(\theta + 1, \bar{\theta}, f)$. Thus, the predicted exit time of the upper path can be computed as $\hat{T}_{sut}(\theta, \bar{\theta}, f) = \theta + 2 + \hat{q}^L_{ut}(\theta + 1, \bar{\theta}, f)$.

This means, for particles starting at s at time $\bar{\theta}$ the lower path is the predicted shortest path whenever the predicted queue length of ut at time $\bar{\theta} + 1$ exceeds the value 2. If that value is smaller than 2, the upper path is the preferred alternative.

For particles starting at times $\bar{\theta} \leq 1$, the predicted queue length of ut at time $\bar{\theta} + 1$ vanishes, and therefore all particles start streaming into the upper path which is shown in Figure 4.1(b).

Starting from time $\bar{\theta} = 1$, a queue starts to build up and thus its derivative jumps to 1. However, for any $\bar{\theta} \in [1, 2]$, the predicted queue length at time $\bar{\theta} + 1$ does still not exceed the value 2. Hence, the upper path remains the preferred path such that it is the only path used up until time 2. This can be seen in Figure 4.1(c).

Figure 4.1(c) also shows that, starting from time $\bar{\theta} = 2$, the predicted queue length exceeds the value 2 making the lower path the preferred path for the first time. Therefore, particles now start flowing into the edge sv instead of su. This behavior continues as long

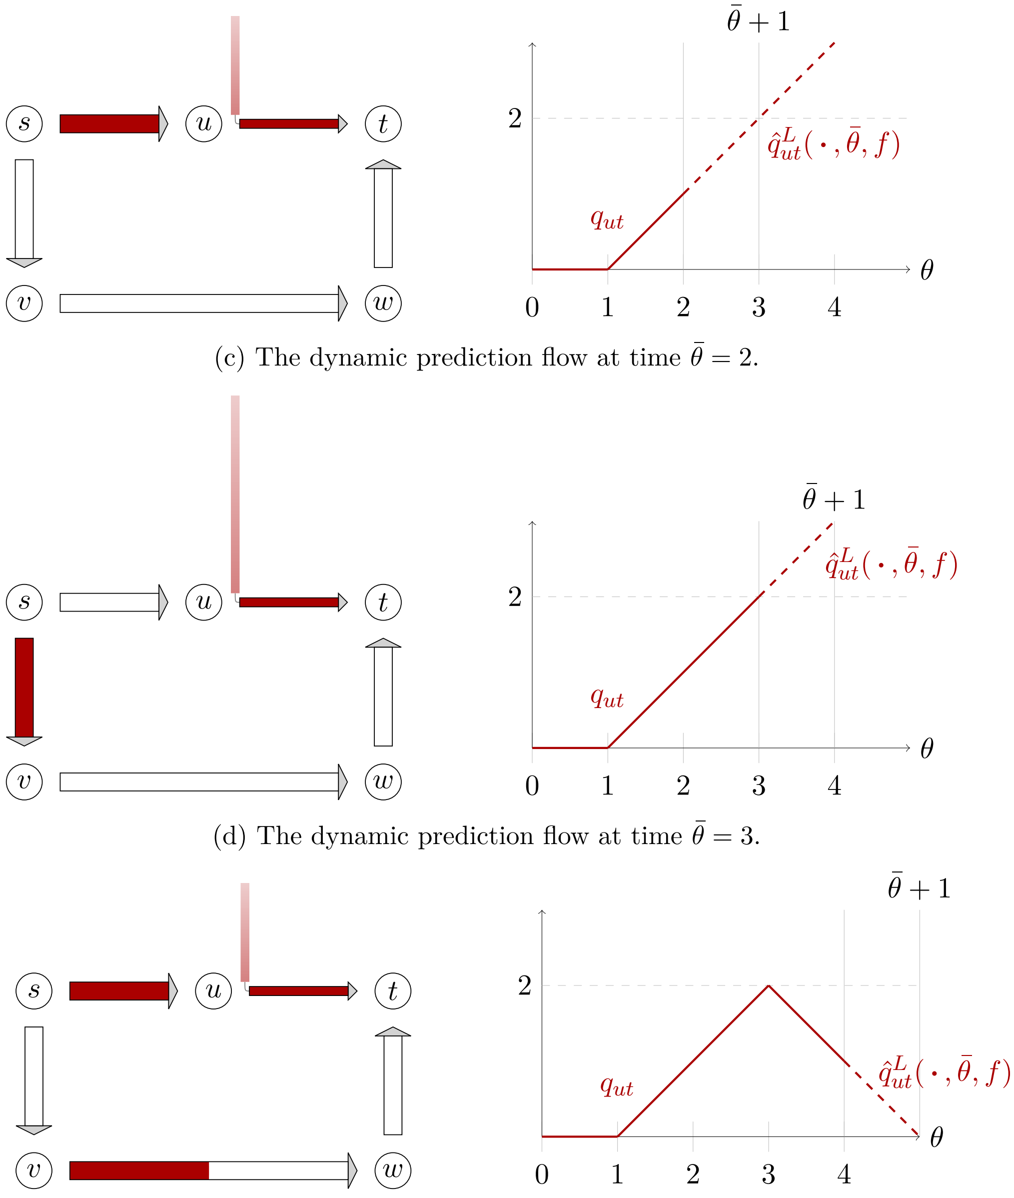

(a) The dynamic prediction flow at time $\bar{\theta} = 0$.

(b) The dynamic prediction flow at time $\bar{\theta} = 1$.

(c) The dynamic prediction flow at time $\bar{\theta} = 2$.

(d) The dynamic prediction flow at time $\bar{\theta} = 3$.

(e) The dynamic prediction flow at time $\bar{\theta} = 4$.

Figure 4.1: The evolution of a dynamic prediction equilibrium using the linear predictor.

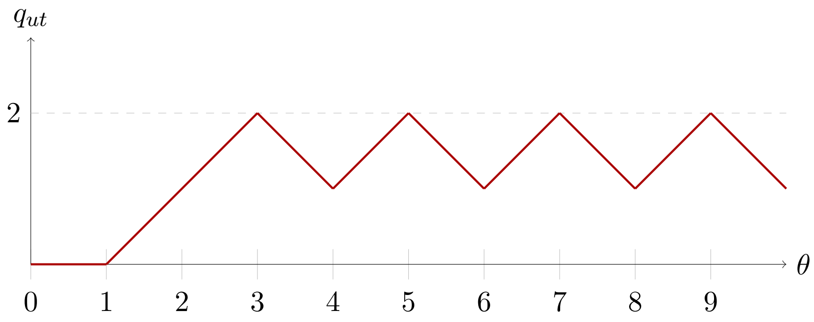

Figure 4.2: The evolution of the queue length q_{ut} of a dynamic prediction flow up to time $\theta = 10$.

as particles still arrive at node u: until $\bar{\theta} = 3$. Figure 4.1(d) shows the equilibrium flow at time $\bar{\theta} = 3$.

Once no more particles arrive at node u, the edge can process its queue and the derivative of the queue length becomes -1. At this point, the predicted queue length at time $\bar{\theta} + 1$ jumps from value 3 to the value 1. Hence, particles start to choose the upper path again as long as no particles arrive at node u as depicted in Figure 4.1(e).

At time $\bar{\theta} = 4$ particles arrive at node u, the queue length of ut will increase, and particles will start to choose the lower path again. This switching behavior between the upper and the lower path continues forever accordingly. Figure 4.2 shows the evolution of the queue up to time $\theta = 10$.

4.3 Fundamentals for the Existence Theorem

Before we may show that under certain conditions a dynamic prediction equilibrium always exists, we first need to introduce various mathematical tools.

The extended norm and its topology

The following sections involve the space of real-valued continuous functions $C(\mathbb{R}^d, \mathbb{R})$ on $\mathbb{R}^d$ in various settings. Unbounded functions on $\mathbb{R}^d$ hinder us from installing the uniform norm on these function spaces. However, as we still want to work with the same topology induced by the uniform norm, we "extend" the notion of a norm to allow the value ∞:

Definition 4.3.1. Given a vector space X on $\mathbb{R}$, a function $\|\cdot\| : X \to \mathbb{R}_{\geq 0} \cup \{\infty\}$ is called an *extended norm on X*, if it fulfills the following conditions:

(i) $\|x\| = 0 \Leftrightarrow x = 0$ for all $x \in X$,

(ii) $\|\alpha \cdot x\| = |\alpha| \cdot \|x\|$ for all $\alpha \in \mathbb{R} \setminus \{0\}$,

(iii) $\|x + y\| \leq \|x\| + \|y\|$ for all $x, y \in X$.

An extended norm $\|\cdot\|$ induces a topology on X in the same way a usual norm would: A set $B \subseteq X$ is called $\|\cdot\|$-*open* if for every $x \in B$ there exists some $\varepsilon > 0$ such that

$U_\varepsilon(x) := \{y \in X \mid \|x-y\| < \varepsilon\} \subseteq B$. Moreover, limits of sequences in this topology behave in accordance to the extended norm. For a general topology X with open sets τ, we say a sequence $(x_n)_{n\in N}$ *converges to* $x \in X$ if for every open set $U \in \tau$ containing x there is some $N \in \mathbb{N}$ with $x_n \in U$ for all $n \geq N$.

Proposition 4.3.2. *Let X be a vector space equipped with the topology induced by an extended norm $\|\cdot\|$. Given a sequence $(x_n)_{n\in\mathbb{N}}$ with $x_n \in X$ and a point $x \in X$, the following statements are equivalent:*

(i) The sequence $(x_n)_{n\in\mathbb{N}}$ converges to x.

(ii) For $n \to \infty$, $\|x_n - x\|$ converges to 0.

Proof. Assume (i). Then, for every $\varepsilon > 0$, there is some $N \in \mathbb{N}$ with $x_n \in U_\varepsilon(x)$ for all $n \geq N$. Therefore, $\|x_n - x\| \leq \varepsilon$ for all $n \geq N$, which implies that the sequence converges to 0.

Assume (ii) and let U be a $\|\cdot\|$-open set containing x. Then there exists some $\varepsilon > 0$ with $U_\varepsilon(x) \subseteq U$. Furthermore, there exists $N \geq 0$ such that $\|x_n - x\| < \varepsilon$ and thus $x_n \in U$ for all $n \geq N$. □

Finally, we note that a vector space X equipped with an extended norm $\|\cdot\|$ is locally metrizable, i.e. for every point $x \in X$ there is a neighborhood U of x such that U is metrizable: We select the neighborhood $U := \{y \in X \mid \|x-y\| < \infty\}$ together with the metric $d_U(y,z) := \|y-z\|$. Every locally metrizable vector space is first-countable, implying that X is a sequential vector space. This means that all properties of its topology can be expressed using sequences.

Definition 4.3.3. A topological space X with open sets τ is called a *sequential space*, if it fulfills the so-called *universal property of sequential spaces*:

For every topological space Y, a function $f : X \to Y$ is continuous if and only if it is *sequentially continuous*, i.e. if for every sequence $(x_n)_{n\in\mathbb{N}}$ converging to some $x \in X$, the sequence $(f(x_n))_{n\in N}$ converges to $f(x)$.

In this general setting a function $f : X \to Y$ is called *continuous*, if for every open set U in Y, the inverse image $f^{-1}(U)$ is an open set in X.

Proposition 4.3.4. *A topological space induced by an extended norm is sequential.*

Proof. Let Y be a topological space and $f : X \to Y$ continuous and let $(x_n)_{n\in\mathbb{N}}$ be a sequence in X converging to some $x \in X$. Let U be any open set in Y containing $f(x)$. Then $f^{-1}(U)$ is an open set in X containing x. Hence, there is some $N \in \mathbb{N}$ with $x_n \in f^{-1}(U)$ for all $n \geq N$. Moreover, $f(x_n) \in U$ for all of these $n \geq N$ implying that $(f(x_n))_{n\in\mathbb{N}}$ converges to $f(x)$ in Y, and f is sequentially continuous.

Now, assume that f is sequentially continuous. We aim to show that $f^{-1}(U)$ is open in X for any open set U in Y. Assume, $f^{-1}(U)$ is not open. This means that there exists some $x \in f^{-1}(U)$ such that $U_\varepsilon(x) \nsubseteq f^{-1}(U)$ holds for every $\varepsilon > 0$. Therefore, for every $n \in \mathbb{N}$ there exists x_n with $\|x - x_n\| < 1/n$ and $x_n \notin f^{-1}(U)$. The sequence $(x_n)_{n\in\mathbb{N}}$ converges to x in X. As f is sequentially continuous, the sequence $(f(x_n))_{n\in\mathbb{N}}$ converges to $f(x)$ in Y. Because U is open, there exists $N \in \mathbb{N}$ with $f(x_n) \in U$, or equivalently $x_n \in f^{-1}(U)$, for all $n \geq N$; a contradiction. □

Now that we have build up enough trust in the topology induced by an extended norm and in its behavior, we introduce the extended uniform norm on the continuous function spaces.

Definition 4.3.5. The *extended (uniform) norm* on the space $C(\mathbb{R}^n, \mathbb{R})$ is defined as

$$\|f\|_\infty := \sup_{x \in \mathbb{R}^n} |f(x)| .$$

For $k \in \mathbb{N}$ the *extended (uniform) norm* on the vector space $C(\mathbb{R}^n, \mathbb{R}^m)$ is given as

$$\|(f_1, \dots, f_m)\|_\infty := \max_{i \in [m]} \|f_i\|_\infty .$$

The space $L^p(S)^d$ and its continuous dual space

Showing the existence of DPE involves the existence of variational inequalities studied in functional analysis. In the following, we introduce some fundamental definitions and theorems. We essentially follow the introduction by Royden in [24, Section 8.1].

For a measurable set $S \subseteq \mathbb{R}$ and some $p \geq 1$, we define the function space

$$\mathcal{L}^p(S) := \left\{ f : S \to \mathbb{R} \,\middle|\, \|f\|_p := \left(\int_S |f|^p \,\mathrm{d}\lambda \right)^{1/p} < \infty \right\}.$$

Identifying any two functions $f, g \in \mathcal{L}^p(S)$ that fulfill $f \overset{a.e.}{=} g$ yields the Banach space

$$L^p(S) := \mathcal{L}^p(S) / \overset{a.e.}{=},$$

together with its norm $\|\cdot\|_p$.

The *continuous dual space* X' of a normed vector space X over $\mathbb{R}$ is the set of all bounded linear functionals $f : X \to \mathbb{R}$. Here, a linear functional f is called *bounded* if $\|f\|_{X'} := \sup_{x \in X} |f(x)|$ is finite. Moreover, $\|\cdot\|_{X'}$ defines a norm on X'. The *canonical pairing* between a Banach space X and its continuous dual refers to $\langle f, x \rangle := f(x)$ for $f \in X'$, $x \in X$. The term *continuous* in *continuous dual space* stems from the following observation.

Proposition 4.3.6 ([24, Theorem 1 in Section 13.2])**.** *A linear mapping between two normed spaces is continuous if and only if it is bounded.*

The Riesz Representation Theorem now leads to the following insight into the dual space of $L^p(S)$:

Theorem 4.3.7 (Riesz Representation Theorem for $L^p(S)$ [24, Section 8.1])**.** *Let S be a measurable set, $1 \leq p < \infty$ and let q be the conjugate of p with $1/p + 1/q = 1$. For each $g \in L^q(S)$ the linear functional $T_g(f) := \int_S g \cdot f \,\mathrm{d}\lambda$ is bounded and hence $T_g \in L^p(E)'$. Moreover, for each $T \in L^p(S)'$ there exists a unique $g \in L^q(S)$ with*

$$T = T_g \quad \textit{and} \quad \|T\|_{L^p(S)'} = \|g\|_q .$$

In other words, the continuous dual space of $L^p(S)$ is isomorphic to $L^q(S)$ where q is the conjugate with $1/p + 1/q = 1$. The canonical pairing between the two spaces is given as $\langle g, f \rangle := \int_S g \cdot f \,\mathrm{d}\lambda < \infty$ for $g \in L^q(S)$ and $f \in L^p(S)$.

For vectors of p-integrable functions living in $L^p(S)^d := (L^p(S))^d$ with $d \in \mathbb{N}$ it is easy to show that the continuous dual space is $L^q(S)^d$ with $1/p + 1/q = 1$ together with the pairing

$$\langle f, g \rangle = \sum_{i=1}^{d} \int_S g_i \cdot f_i \, \mathrm{d}\lambda$$

for $f \in L^p(S)^d$ and $g \in L^q(S)^d$ and the norm $\|g\|_{L^q(S)^d} = \sum_{i=1}^d \|g_i\|_q$.

The Weak Topology and Sequential Weak-Strong Continuity

As in [24, Section 11.4] we define the *weak topology* for a non-empty set X and a collection of mappings $\mathcal{F} = (f_i : X \to Y_i)_{i \in I}$ where each Y_i is a topological space, as the weakest topology that makes all mappings f_i continuous, i.e. the smallest topology containing the open sets $\{f_i^{-1}(U) \mid i \in I,\ U \text{ open in } Y_i\}$.

For a normed vector space X, the *weak topology of* X refers to the weakest topology that makes all mappings $\langle f, \cdot \rangle = f \in X'$ continuous. Here, whenever we refer to the weak topology, it will be stated explicitly. In any other case, we use the "strong" topology induced by the norm of X. We say a sequence $(x_i)_{i \in \mathbb{N}}$ in a normed space X *converges weakly* to some $x \in X$, if the sequence converges to x in the weak topology. The following proposition, as given in [24, Section 11.4], characterizes weak convergence of sequences using the mappings of the continuous dual space:

Proposition 4.3.8. *A sequence $(x_i)_{i \in \mathbb{N}}$ in a normed space X converges weakly to $x \in X$ if and only if the pairing $\langle f, x_i \rangle$ converges to $\langle f, x \rangle$ for every $f \in X'$.*

Using these preliminaries we can define sequential weak-strong continuity:

Definition 4.3.9. A mapping $f : X \to Y$ from a normed space X into another normed space Y is called *sequentially weak-strong continuous* if $\lim_{i \to \infty} \|f(x_i) - f(x)\|_Y = 0$ holds true for every sequence $(x_i)_{i \in I}$ converging weakly to some x in X.

Reflexive Banach Spaces

The following definitions are based on [24, Section 14]. For a normed vector space X, the linear mapping $J : X \to X''$ with $J(x) = (f \mapsto f(x))$ is called the *natural embedding* of X into its *bidual space* $X'' := (X')'$. A normed vector space X is called *reflexive* if the natural embedding is surjective.

For proving the existence of DPE, we are interested in the space $L^p(S)$ for a measurable set $S \subseteq \mathbb{R}$. Here, [24, Theorem 8 in Section 19.4] proves that $L^p(S)$ is reflexive for all $1 < p < \infty$. This result can easily be extended to the case $L^p(S)^d$ for $d \in \mathbb{N}$.

Solutions to Variational Inequalities

In this subsection, we formulate the theorem necessary to show the existence of DPE. The underlying theory involves functional analysis and solutions of variational inequalities. The main statement this section builds upon is a generalization of the results by Browder and Minty [6, 20]. This generalization was first introduced by Haïm Brézis in [4, Théorème 24].

Lions has shown a variant of Brézis's very general theorem in [15, Théorème 8.1 in Chapitre 2]. In contrast to Brézis, Lions uses the rather basic concepts introduced in the first part of this section to define pseudo-monotone operators in [15, Définition 2.1].

Definition 4.3.10. Given a reflexive Banach space X we call a mapping $\mathcal{A} : K \to X'$ *pseudo-monotone*, if $\mathcal{A}$ is bounded and if for every sequence $(x_i)_{i\in\mathbb{N}}$ converging weakly to x in X fulfilling

$$\limsup_{i\in\mathbb{N}} \langle \mathcal{A}(x_i), x_i - x\rangle \leq 0$$

it holds that

$$\forall y \in K : \quad \liminf_{i\in\mathbb{N}} \langle \mathcal{A}(x_i), x_i - y\rangle \geq \langle \mathcal{A}(x), x - y\rangle .$$

Theorem 4.3.11 ([15, Théorème 8.1 in Chapitre 2])**.** *Let X be a reflexive Banach space, let K be a non-empty, closed, and convex subset $K \subseteq X$, and let $\mathcal{A} : K \to X'$ be a pseudo-monotone mapping. Then for any $f \in X'$ there exists a solution $u \in K$ to the following variational inequality:*

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

Cominetti et al. formulated a useful special case of this theorem for the existence of dynamic Nash equilibrium flows in [7]. This corollary is presented in the following:

Corollary 4.3.12 ([7, Section 5])**.** *Let X be a reflexive Banach space and let $\mathcal{A} : K \to X'$ 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*

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

4.4 Existence of Dynamic Prediction Equilibria

In this section, the existence of dynamic prediction equilibria under certain conditions is proven. However, an equilibrium might not exist for general predictors. In the following example the discontinuity of the predictor is exploited to show that there is no dynamic prediction flow with respect to this predictor.

Example 4.4.1. We consider the network given in Figure 4.3 with a single commodity with source s, sink t and network inflow rate $u|_{[0,\infty)} \equiv 2$. Each edge of the network has transit time $\tau_e = 1$ and capacity $\nu_e = 2$ except edge st which has a capacity of $\nu_{st} = 1$. We define the non-continuous predictor

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

and we assume there exists a dynamic prediction flow with respect to $\hat{q}$.

On the interval $[0, 1)$ only edge st is active, as its queue length $q_{st}^f(\theta)$ must be smaller than 1 and thus $\hat{T}_{st}(\theta, \theta, f) < 2 \leq \hat{T}_{svt}(\theta, \theta, f)$ holds for $\theta \in [0, 1)$. Hence, f must fulfill $f^+_{st}|_{[0,1]} \overset{a.e.}{\equiv} 2$, $f^+_{sv}|_{[0,1]} \overset{a.e.}{\equiv} 0 \overset{a.e.}{\equiv} f^+_{vt}|_{[0,1]}$ and $q^f_{st}(1) = 1$.

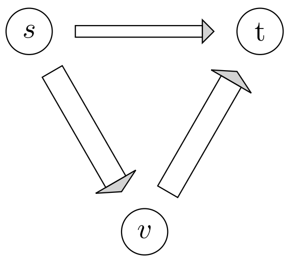

(a) The network at time $\theta = 0$.

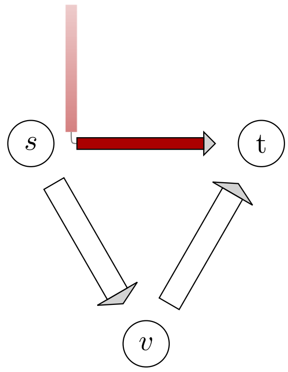

(b) Any dynamic prediction flow at time $\theta = 1$.

Figure 4.3: The network considered in Example 4.4.1.

For $\theta \geq 1$ the queue $q^f_{st}(\theta)$ must be at least of length 1: The set $O := \{\theta \mid q^f_{st}(\theta) < 1\}$ is open and at times in O only edge st is active, implying $f^+_{st}|_O \overset{a.e.}{\equiv} 2$ and $f^+_{e_2}|_O \overset{a.e.}{\equiv} 0$. We partition O into an at most countably infinite union $O = \bigcup_{j \in J}(a_j, b_j)$ of disjoint, open intervals. Let (a_j, b_j) be one of these intervals with $a_j \geq 1$. By continuity, we have $q^f_{st}(a_j) = 1$ and there exists $\theta \in (a_j, b_j)$ with $q^f_{st}(\xi) > 0$ for all $\xi \in (a_j, \theta]$. Property (F2) implies

$$q^f_{st}(\theta) = q^f_{st}(a_j) + \int_{a_j}^{\theta} f^+_{st} \,\mathrm{d}\lambda - \int_{a_j+\tau_{st}}^{\theta+\tau_{st}} f^-_{st} \,\mathrm{d}\lambda = 1 + (\theta - a_j) \cdot (2 - 1) > 1.$$

This contradicts $\theta \in O$ and thus $O \cap [1, \infty)$ is empty. Hence, edge sv is the only active edge from time 1 onwards. This, however, implies that only edge sv is used starting from time 1 such that the queue at edge sv depletes at time 2; again a contradiction.

Next, we discuss several properties of predictors that ensure the existence of a dynamic prediction equilibrium. As we have seen above, the predictors must fulfill some regularity conditions. The following condition will ensure existence of dynamic prediction flows by using the existence of solutions to the variational inequality of Brézis.

Definition 4.4.2. A predictor $\hat{q}_i$ is called *p-continuous* for some $p \geq 1$, if it fulfills the following two conditions:

(i) For every deterministic flow f with $f^+ \in L^p_{\mathrm{loc}}(\mathbb{R}, \mathbb{R}_{\geq 0})^{I \times E}$ the function $\hat{q}_{i,e}(\,\cdot\,, \,\cdot\,, f)$ is continuous, i.e. $\hat{q}_{i,e}(\,\cdot\,, \,\cdot\,, f) \in C(\mathbb{R}^2, \mathbb{R}_{\geq 0})$, for every $e \in E$.

(ii) For every edge $e \in E$ and $M > 0$ and any compact interval D, the mapping

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

where f denotes the deterministic flow with inflow rates $\mathbf{1}_{[0,M]} \cdot f^+$, is sequentially weak-strong continuous.

The final property we require is that predictors cannot use the future evolution of the flow in their predictions. This means, we want to be able to make predictions solely based on the past evolution of the dynamic flow. We formalize this notion in the following definition.

Notation 4.4.3. For a vector of functions $g = (g_i)_{i\in[d]}$ with $g_i : \mathbb{R} \to \mathbb{R}$ 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 $g_i, h_i : X \to \mathbb{R}$, $X \subseteq \mathbb{R}$, we write $g \underset{c.w.}{\overset{a.e.}{=}} h$ if $g_i \overset{a.e.}{=} h_i$ holds for all $i \in [d]$.

Definition 4.4.4. A predictor $\hat{q}_i$ is called *oblivious*, if it fulfills the condition

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

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

The following theorem shows that we can extend any dynamic prediction flow by some strictly positive amount. This will suffice to prove the existence of dynamic prediction equilibria in Theorem 4.4.13.

Theorem 4.4.5. *Let I be a finite set of commodities with network inflow rates $u_i \in \mathcal{R}$ with $u_i \in L^p_{\mathrm{loc}}(\mathbb{R}, \mathbb{R}_{\geq 0})$. Let $\hat{q} = (\hat{q}_i)_{i\in I}$ be a set of oblivious, p-continuous and FIFO-compatible predictors with $p > 1$ and assume τ_e is strictly positive for all $e \in E$.*

Then we can extend the horizon of any locally p-integrable partial dynamic prediction flow. More specifically, given a dynamic prediction flow f up to time H with $f^+ \in L^p_{\mathrm{loc}}(\mathbb{R},\mathbb{R})^{I\times E}$, there exists a dynamic prediction flow h up to time $H + \alpha$ with $\alpha := \min_{e\in E} \tau_e$ such that $h_{\leq H} \underset{c.w.}{\overset{a.e.}{=}} f_{\leq H}$ and $h^+ \in L^p_{\mathrm{loc}}(\mathbb{R},\mathbb{R})^{I\times E}$ hold.

The proof of the existence theorem is based on the existence of a solution to the variational inequality as given in Corollary 4.3.12. The predicted delay function $\hat{\Delta}_{i,e}(\theta, \bar{\theta}, f)$ plays an essential role in formulating the variational inequality. Therefore, we first make sure that the mapping $f^+ \mapsto \hat{\Delta}_{i,e}(\,\cdot\,,\,\cdot\,, f)$ is sequentially weak-strong continuous. On the way, we prove the same sequential weak-strong continuity for predicted exit times of edges, paths and for earliest arrival times.

Proposition 4.4.6. *If $\hat{q}_i$ is a p-continuous predictor, the mapping*

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

is sequentially weak-strong continuous for every $e \in E$, $M > 0$ and compact interval D.

Proof. Let $(f^{+,k})_{k\in\mathbb{N}}$ be a sequence in $L^p(D)^{I\times E}$ converging weakly to some f^+. We use the p-continuity of $\hat{q}_i$ to conclude

$$\left\|\hat{T}_{i,e}(\,\cdot\,,\,\cdot\,,f^k) - \hat{T}_{i,e}(\,\cdot\,,\,\cdot\,,f)\right\|_\infty = \frac{1}{\nu_e} \cdot \left\|\hat{q}_{i,e}(\,\cdot\,,\,\cdot\,,f^k) - \hat{q}_{i,e}(\,\cdot\,,\,\cdot\,,f)\right\|_\infty \xrightarrow[k\to\infty]{} 0.$$

□

The sequential weak-strong continuity of the exit time of an edge constitutes the base case for the sequential weak-strong continuity of the exit time of a finite path.

Proposition 4.4.7. *If $\hat{q}_i$ is a p-continuous predictor, the mapping*

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

is sequentially weak-strong continuous for any finite path P, $M>0$ and compact interval D.

Proof. We use a simple induction on the length of the path P. In the base case, P is an empty path such that $\hat{T}_{i,P}(\,\cdot\,,\,\cdot\,,f) = ((\theta,\bar{\theta}) \mapsto \theta)$ for any deterministic flow f, and $f^+ \mapsto \hat{T}_{i,P}(\,\cdot\,,\,\cdot\,,f)$ is a constant (and thus continuous) map.

Now, 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^+. We want to show that

$$\begin{aligned}
&\left\|\hat{T}_{i,P}(\,\cdot\,,\,\cdot\,,f^k) - \hat{T}_{i,P}(\,\cdot\,,\,\cdot\,,f)\right\|_\infty \\
&= \sup_{\substack{\theta\in[0,M]\\ \bar{\theta}\in D}} \left|\hat{T}_{i,e}\left(\hat{T}_{i,P'}(\theta,\bar{\theta},f^k),\bar{\theta},f^k\right) - \hat{T}_{i,e}\left(\hat{T}_{i,P'}(\theta,\bar{\theta},f),\bar{\theta},f\right)\right|
\end{aligned}$$

converges to 0 for $k\to\infty$. We split up this term using the triangle inequality to get

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

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

$$\left\|\hat{T}_{i,P'}(\,\cdot\,,\,\cdot\,,f^k) - \hat{T}_{i,P'}(\,\cdot\,,\,\cdot\,,f)\right\|_\infty < 1$$

holds for all $k\geq N$. We define $M' = \max_{\theta\in[0,M],\bar{\theta}\in D} \hat{T}_{i,P'}(\theta,\bar{\theta},f)+1$. Then for all $\theta\in[0,M]$ and $\bar{\theta}\in D$, the value $\hat{T}_{i,P'}(\theta,\bar{\theta},f^k)$ is contained in $[0,M']$ for all $k\geq N$. Moreover, $\mathbf{1}_{[0,M]}\cdot f^{+,k}$ converges weakly to $\mathbf{1}_{[0,M]}\cdot f^+$ in $L^p([0,M'])^{I\times E}$. Using Proposition 4.4.6 we conclude

$$\alpha^k \leq \left\|\hat{T}_{i,e}(\,\cdot\,,\,\cdot\,,f^k) - \hat{T}_{i,e}(\,\cdot\,,\,\cdot\,,f)\right\|_{C([0,M']\times D)} \xrightarrow[k\to\infty]{} 0.$$

The term β_k converges to 0 because of the uniform continuity of $\hat{T}_{i,e}(\,\cdot\,,\,\cdot\,,f)$ on the compact set $[0,M']\times D$. □

Proposition 4.4.8. *If $\hat{q}_i$ is a p-continuous predictor, the mapping*

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

is sequentially weak-strong continuous for any node $v\in V_i$, $M>0$ and compact interval D.

Proof. Let $(f^{+,k})_{k\in\mathbb{N}}$ weakly converge to f^+ in $L^p([0,M],\mathbb{R}_{\geq 0})^{I\times E}$. By Proposition 4.4.7, the sequence $(\hat{T}_{i,P}(\,\cdot\,,\,\cdot\,,f^k))_k$ uniformly converges to $\hat{T}_{i,P}(\,\cdot\,,\,\cdot\,,f)$ on $C([0,M]\times D,\mathbb{R}_{\geq 0})$ for any simple v-t_i-path P. The minimum of these functions then also converges uniformly. ☐

Now, we are ready to establish the sequential weak-strong continuity for the predicted delay function:

Proposition 4.4.9. *If $\hat{q}_i$ is a p-continuous predictor, the mapping*

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

is sequentially weak-strong continuous for any $e\in E_i$, $M>0$ and compact interval D.

Proof. Let $(f^{+,k})_{k\in\mathbb{N}}$ weakly converge to f^+ in $L^p([0,M],\mathbb{R}_{\geq 0})^{I\times E}$. For $e=vw$, we show

$$\left\|\hat{\Delta}_{i,e}(\,\cdot\,,\,\cdot\,,f^k)-\hat{\Delta}_{i,e}(\,\cdot\,,\,\cdot\,,f)\right\|_\infty \xrightarrow[k\to\infty]{} 0.$$

In a first step, we split up the sum again using the triangle inequality to obtain

$$\begin{aligned}
&\left\|\hat{\Delta}_{i,e}(\,\cdot\,,\,\cdot\,,f^k)-\hat{\Delta}_{i,e}(\,\cdot\,,\,\cdot\,,f)\right\|_\infty \\
&\leq \left\|\hat{l}_{i,v}(\,\cdot\,,\,\cdot\,,f^k)-\hat{l}_{i,v}(\,\cdot\,,\,\cdot\,,f)\right\|_\infty \\
&\quad + \sup_{\substack{\theta\in[0,M]\\ \bar\theta\in D}} \left|\hat{l}_{i,w}\left(\hat{T}_{i,e}(\theta,\bar\theta,f^k),\bar\theta,f^k\right)-\hat{l}_{i,w}\left(\hat{T}_{i,e}(\theta,\bar\theta,f),\bar\theta,f\right).\right|
\end{aligned}$$

The first term converges to 0 by Proposition 4.4.8. The second term is handled analogously as in the proof of Proposition 4.4.7: We apply the triangle inequality once again to get

$$\begin{aligned}
&\sup_{\substack{\theta\in[0,M]\\ \bar\theta\in D}} \left|\hat{l}_{i,w}\left(\hat{T}_{i,e}(\theta,\bar\theta,f^k),\bar\theta,f^k\right)-\hat{l}_{i,w}\left(\hat{T}_{i,e}(\theta,\bar\theta,f),\bar\theta,f\right)\right| \\
&\leq \overbrace{\sup_{\substack{\theta\in[0,M]\\ \bar\theta\in D}} \left|\hat{l}_{i,w}\left(\hat{T}_{i,e}(\theta,\bar\theta,f^k),\bar\theta,f^k\right)-\hat{l}_{i,w}\left(\hat{T}_{i,e}(\theta,\bar\theta,f^k),\bar\theta,f\right)\right|}^{=:\alpha^k} \\
&\quad + \underbrace{\sup_{\substack{\theta\in[0,M]\\ \bar\theta\in D}} \left|\hat{l}_{i,w}\left(\hat{T}_{i,e}(\theta,\bar\theta,f^k),\bar\theta,f\right)-\hat{l}_{i,w}\left(\hat{T}_{i,e}(\theta,\bar\theta,f),\bar\theta,f\right)\right|}_{=:\beta^k}.
\end{aligned}$$

By Proposition 4.4.6 there exists an $N\in\mathbb{N}$ with $\|\hat{T}_{i,e}(\,\cdot\,,\,\cdot\,,f^k)-\hat{T}_{i,e}(\,\cdot\,,\,\cdot\,,f)\|_\infty<1$ for all $k\geq N$ and define $M'=\max_{\theta\in[0,M],\bar\theta\in D}\hat{T}_{i,e}(\theta,\bar\theta,f)+1$. Then $\hat{T}_{i,e}(\theta,\bar\theta,f^k)\in[0,M']$ for any $\theta\in[0,M],\bar\theta\in D$ for all $k\geq N$. Using Proposition 4.4.8 we conclude

$$\alpha^k \leq \left\|\hat{l}_{i,w}(\,\cdot\,,\,\cdot\,,f^k)-\hat{l}_{i,w}(\,\cdot\,,\,\cdot\,,f)\right\|_{C([0,M']\times D)} \xrightarrow[k\to\infty]{} 0.$$

The term β_k converges to 0 because of the uniform continuity of $\hat{l}_{i,w}(\,\cdot\,,\,\cdot\,,f)$ on the compact set $[0,M']\times D$. ☐

Proof of Theorem 4.4.5. Let f be a partial dynamic prediction flow up to time H with respect to $\hat{q}$. Without loss of generality, we may assume that $(f^-_{i,e})_{i,e}$ are the deterministic outflow rates with respect to the inflow rates $(f^+_{i,e})_{i,e}$ as this does not change the outflow rates up to time H as per Theorem 3.4.6. Lemma 3.4.8 states that the outflow rate on an edge e of any dynamic flow h whose inflow rates coincide with f up to time H are already uniquely determined up to time $T_e(H)$. Therefore, also the rates $b^-_{i,v}(\theta) \coloneqq \sum_{e\in\delta^-_v} h^-_{i,e}(\theta) + \mathbf{1}_{v=s_i} u_i(\theta)$ on the interval $D \coloneqq [H, H+\alpha]$ are independent of the edge inflow rates $h^+_{i,e}|_D$ on D.

We now want to apply Brézis' theorem in the form of Corollary 4.3.12 to find suitable inflow rates $h^+_{i,e}|_D$. For that, we define the set $K \subseteq L^p(D)$ as follows:

$$K \coloneqq \left\{ g \in L^p(D, \mathbb{R}_{\geq 0})^{I\times E} \;\middle|\; \begin{array}{cccc} \forall i \in I, v \in V \setminus \{t_i\}: & \sum_{e\in\delta^+_v} g_{i,e} & \overset{a.e.}{=} & b^-_{i,v}, \\ \forall i \in I: & \sum_{e\in\delta^+_{t_i}} g_{i,e} & \overset{a.e.}{\leq} & b^-_{i,t_i} \end{array} \right\}.$$

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

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

Claim 4.4.10. *For each $g \in K$, $\bar{g}$ is a feasible flow up to time $H+\alpha$ and $(\hat{q}, \bar{g})$ is a DPE up to time H.*

Proof. The properties (F1), (F2) and (F3) are fulfilled for almost all $\theta \in \mathbb{R}$ since $\bar{g}$ is a deterministic flow. For the feasibility of $\bar{g}$, it remains to show that flow is conserved for almost all $\theta < H+\alpha$. By Lemma 3.4.8 the outflow rates of $\bar{g}$ coincide with the outflow rates of f for all edges $e \in E$ on $(-\infty, T_e(H))$. Hence, for almost all $\theta < H$ property (F4) is fulfilled due to the feasibility of f up to time H. For $\theta \in D$ the constraint is directly implied by the first two conditions of $g \in K$.

Because all predictors are oblivious, the equilibrium property

$$\bar{g}^+_{i,e}(\theta) > 0 \implies e \in \hat{E}_i(\theta, \theta, \bar{g})$$

is transferred from f for almost all $\theta < H$. ■

Claim 4.4.11. *The set K is nonempty, closed, bounded and convex.*

Proof. We note that $b^-_{i,v}$ is p-integrable on D for all $v \in V$ because of the assumption $u_i \in L^p_{\text{loc}}(\mathbb{R}, \mathbb{R}_{\geq 0})$ and $f^-_{i,e}(\theta) \leq \nu_e$ for almost all θ. By the constraints of $g \in K$ we have $\|g_{i,e}\|_{L^p(D)} \leq \|b^-_{i,v}\|_{L^p(D)}$ for all $i \in I$ and $e = vw \in E$. This implies

$$\begin{aligned} \|g\|_{L^p(D)^{I\times E}} &= \sum_{i,e} \|g_{i,e}\|_{L^p(D)} = \sum_{i,v} \sum_{e\in\delta^+_v} \|g_{i,e}\|_{L^p(D)} \\ &= \sum_{i,v} \left\| \sum_{e\in\delta^+_v} g_{i,e} \right\|_{L^p(D)} \leq \sum_{i,v} \left\| b^-_{i,v} \right\|_{L^p(D)} < \infty. \end{aligned}$$

Therefore, K is bounded in $L^p(D)^{I\times E}$. To understand that K is nonempty, we 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 simply 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 $g \in K$. The convexity of K can be verified easily. By checking that the constraints of K also hold at limit points of converging sequences of K, we conclude that K is closed. ■

By Claim 4.4.10, $\bar{g}$ 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

$$\bar{g}^+_{i,e}(\theta) > 0 \implies e \in \hat{E}_i(\theta, \theta, \bar{g})$$

is also fulfilled for almost all $\theta \in D$. For an edge $e = vw \in E_i$ lying on a directed s_i-t_i-path, the statement $e \in \hat{E}_i(\theta, \theta, \bar{g})$ can be reformulated as $\hat{\Delta}_{i,e}(\theta, \theta, \bar{g}) \leq 0$. Using this observation, we define our operator $\mathcal{A} : L^p(D)^{I\times E} \to L^q(D)^{I\times E}$ 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 here. By the p-continuity of $\hat{q}$, the function $\mathcal{A}(g)_{i,e}$ is continuous as a function on $\mathbb{R}$, too. Thus, $\mathcal{A}(g)$ is indeed contained in $L^q(D)^{I\times E}$ as a function on D.

Claim 4.4.12. *The map $\mathcal{A}$ is non-negative and sequentially weak-strong continuous.*

Proof. The non-negativity of $\mathcal{A}$ results from the FIFO-compatibility of the predictor.

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 $L^p([0,M])^{I\times D}$ with $M := H + \alpha$. By Proposition 4.4.9 and the obliviousness 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 $e \in E_i$. We use this information to compute the convergence on the space $C(D)^{I\times E}$:

$$\begin{aligned} \left\| \mathcal{A}(g^k) - \mathcal{A}(g) \right\|_\infty &= \max_{\substack{i\in I \\ e\in E_i}} \sup_{\theta\in D} \left| \hat{\Delta}_{i,e}(\theta,\theta,\bar{g}^k) - \hat{\Delta}_{i,e}(\theta,\theta,\bar{g}) \right| \\ &\leq \max_{\substack{i\in I \\ e\in E_i}} \left\| \hat{\Delta}_{i,e}(\,\cdot\,,\,\cdot\,,\bar{g}^k) - \hat{\Delta}_{i,e}(\,\cdot\,,\,\cdot\,,\bar{g}) \right\| \xrightarrow[k\to\infty]{} 0. \end{aligned}$$

This also implies convergence on $L^p(D)^{I\times E}$: For arbitrary ε there exists $N \in \mathbb{N}$ with $\|\mathcal{A}(g^k) - \mathcal{A}(g)\|_\infty < \varepsilon/\alpha$ for all $k \geq N$. This implies

$$\left\| \mathcal{A}_{i,e}(g^k) - \mathcal{A}_{i,e}(g) \right\|_{L^q(D)} \leq \alpha \cdot \left\| \mathcal{A}_{i,e}(g^k) - \mathcal{A}_{i,e}(g) \right\|_\infty < \varepsilon.$$

■

The variational inequality as given in Corollary 4.3.12 now states that there exists some $g \in K$ 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 following set has positive measure:

$$\Phi \coloneqq \bigcup_{i,e} \Phi_{i,e} \coloneqq \bigcup_{i,e} \left\{ \theta \in D \,\middle|\, \bar{g}^+_{i,e}(\theta) > 0 \text{ and } \mathcal{A}(g)_{i,e}(\theta) > 0 \right\}.$$

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 $(\theta, \bar{\theta}) \mapsto \hat{l}_{i,w}(\theta, \bar{\theta}, \bar{g})$ and $\theta \mapsto \hat{T}_{i,e}(\theta, \theta, \bar{g})$, the set of times at which an edge e is active, denoted by

$$\Theta_{i,e} \coloneqq \left\{ \theta \in \mathbb{R} \,\middle|\, e \in \hat{E}_{i,e}(\theta, \theta, \bar{g}) \right\},$$

is closed and thus measurable. We now define

$$h_{i,e} : D \to \mathbb{R}, \quad \theta \mapsto \begin{cases} \frac{b^-_{i,v}(\theta)}{\left|\delta^+_v \cap \hat{E}_i(\theta,\theta,\bar{g})\right|}, & \text{if } \theta \in \Theta_{i,e}, \\ 0, & \text{otherwise,} \end{cases}$$

for all $i \in I$, $e = vw \in E$ and $\theta \in D$. This function is measurable, because the function $\theta \mapsto \left|\delta^+_v \cap \hat{E}_i(\theta, \theta, \bar{g})\right|$ can be expressed as $\sum_{e \in \delta^+_v} \mathbf{1}_{\Theta_{i,e}}$. The boundedness of $b^-_{i,v}(\theta)$ implies $h \in L^2(D)^{I \times E}$. As the set $\delta^+_v \cap \hat{E}_i(\theta, \theta, \bar{g})$ is non-empty for any $\theta \in D$ and $v \in V_i \setminus \{t\}$, it follows that

$$\sum_{e \in \delta^+_v} h_{i,e}(\theta) = \sum_{e \in \delta^+_v \cap \hat{E}_i(\theta,\theta,\bar{g})} \frac{b^-_{i,v}(\theta)}{\left|\delta^+_v \cap \hat{E}_i(\theta, \theta, \bar{g})\right|} \begin{cases} = b^-_{i,v}(\theta), & \text{if } v \in V \setminus \{t\}, \\ \leq b^-_{i,v}(\theta), & \text{otherwise,} \end{cases}$$

because $b^-_{i,v}(\theta)$ can only be positive for $v \in V_i$. Therefore, h is an element of K.

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

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

There exist $i \in I$ 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 $\int_{\Phi_{i,e}} \mathcal{A}(g)_{i,e} \cdot g_{i,e} \,\mathrm{d}\lambda > 0$. The non-negativity of $\mathcal{A}(g)$ and g now implies

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

which contradicts that g is a solution to the variational inequality. □

This theorem provides us with the possibility to extend any dynamic prediction flow by some positive amount of time α. This is the key observation for the existence of dynamic prediction equilibria:

Theorem 4.4.13. *Let I be a finite set of commodities with network inflow rates $u \in \mathcal{R}^I$ with $u_i \in L^p_{\text{loc}}(\mathbb{R}, \mathbb{R}_{\geq 0})$, let $\hat{q} = (\hat{q}_i)_{i \in I}$ be a set of oblivious, p-continuous and FIFO-compatible predictors with $p > 1$ and assume τ_e is strictly positive for all $e \in E$. Then there exists a dynamic prediction flow with respect to $\hat{q}$.*

Proof. Let $\alpha := \min_{e \in E} \tau_e > 0$ and let $f^0 \in (\mathcal{R} \times \mathcal{R})^{I \times E}$ with $f^{0,+}_{i,e}, f^{0,-}_{i,e} :\equiv 0$ for all $i \in I$, $e \in E$. Then f^0 is a dynamic prediction flow with respect to $\hat{q}$ up to time 0. Recursively define f^{n+1} as the dynamic prediction flow with respect to $\hat{q}$ up to time $\alpha \cdot (n+1)$ given through the extension of f^n using Theorem 4.4.5. By the extension procedure, we have $f^n_{\leq k \cdot \alpha} \overset{a.e.}{\underset{c.w.}{=}} f^k_{\leq k \cdot \alpha}$ whenever $n \geq k$. We define the flow f^∞ using

$$f^{\infty,+}_{i,e}(\theta) := f^{n,+}_{i,e}(\theta) \quad \text{and} \quad f^{\infty,-}_{i,e}(\theta) := f^{n,-}_{i,e}(\theta) \quad \text{with } n = \lceil \theta / \alpha \rceil$$

for all $\theta \geq 0$. Then f^∞ is a dynamic prediction flow with respect to $\hat{q}$ up to time ∞: The feasibility and the equilibrium properties are checked for almost all $\theta \in \mathbb{R}$ by using $f^\infty_{\leq n \cdot \alpha} \overset{a.e.}{\underset{c.w.}{=}} f^n_{\leq n \cdot \alpha}$, the fact that f^n fulfills the desired properties for all $n \in \mathbb{N}$, and the obliviousness of the predictors. ☐

4.5 Sufficient Conditions for p-Continuity of Predictors

This section introduces sufficient conditions for the p-continuity of predictors. More specifically, we want to show that every predictor that depends continuously on the past queues, edge loads and cumulative inflow functions is p-continuous. The exact notion of these types of predictors is as follows:

Definition 4.5.1. A predictor $\hat{q}_i$ *depends continuously on the cumulative inflow, total cumulative outflow functions and queue length functions*, if there exist mappings

$$\gamma_{i,e} : \mathbb{R} \times \mathbb{R} \times C(\mathbb{R}, \mathbb{R}_{\geq 0})^{(I \times E) + 2 \cdot E} \to \mathbb{R}_{\geq 0}$$

for each $e \in E$ with the following two properties:

(i) For all $\theta, \bar{\theta} \in \mathbb{R}$ and for all deterministic flows $f \in (\mathcal{R} \times \mathcal{R})^{I \times E}$, it holds that

$$\hat{q}_{i,e}(\theta, \bar{\theta}, f) = \gamma_{i,e}\left(\theta, \bar{\theta}, F^{+,f}_{I \times E}, F^{-,f}_{E}, q^f\right),$$

where $F^{+,f}_{I \times E} := (F^{+,f}_{i,e})_{i \in I, e \in E}$ and $F^{+,f}_{E} := (F^{+,f}_{e})_{e \in E}$ denote the cumulative edge in- and outflow functions with respect to f.

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

The proof of the desired result involves the so-called *Arzelà-Ascoli Theorem.* We first define the notions of uniform boundedness and equicontinuity of a set of functions, before presenting the theorem.

Definition 4.5.2. Let (X, d) be a compact metric space and let F be a set of continuous functions on X, i.e. $F \subseteq C(X)$.

The set F is *uniformly bounded*, if there is some $M \in \mathbb{R}$ with $|f(x)| \leq M$ for all $f \in F$ and $x \in X$.

The set F is *equicontinuous at some point* $x \in X$, if for each $\varepsilon > 0$ there exists a $\delta > 0$ such that for all $y \in X$ with $d(x, y) < \delta$ it follows $|f(y) - f(x)| < \varepsilon$ for all $f \in F$. The set F is *equicontinuous on* X, if it is equicontinuous at all points $x \in X$.

Theorem 4.5.3 (Arzelà-Ascoli Theorem [24, Section 10.1])**.** *Let X be a compact metric space and $(f_n)_{n \in \mathbb{N}}$ a uniformly bounded, equicontinuous sequence of functions $f_n \in C(X)$. Then $(f_n)_n$ has a subsequence that converges uniformly on X to a continuous function $f \in C(X)$.*

Corollary 4.5.4. *Let $a, b \in \mathbb{R}$ with $a < b$. If a set $F \subseteq C([a, b])$ of continuous functions on $[a, b]$ is uniformly bounded and all $f \in F$ fulfill the* Hölder condition *of a fixed order $\alpha > 0$ and a fixed constant $K > 0$, i.e.*

$$\forall f \in F,\ x, y \in [a, b] : \quad |f(y) - f(x)| \leq K \cdot |y - x|^{\alpha},$$

then F is relatively compact in $C([a, b])$.

Proof. We apply the Arzelà-Ascoli Theorem to the compact metric space $X = [a, b]$. For relative compactness in the normed vector space $(C([a, b]), \|\cdot\|_{\infty})$ it suffices to show that any sequence in F has a convergent subsequence in $(C([a, b]), \|\cdot\|_{\infty})$, i.e. a uniformly convergent subsequence. The equicontinuity of F now immediately follows from the Hölder condition by choosing $\delta = (\varepsilon/K)^{1/\alpha}$. $\square$

Next, we discuss several properties of operators $T : X \to Y$ where X and Y are Banach spaces over $\mathbb{R}$. We call T a *linear operator*, if it supports addition and multiplication with a scalar, i.e. $T(x + y) = T(x) + T(y)$ and $T(c \cdot x) = c \cdot T(x)$ holds for all $x, y \in X, c \in \mathbb{R}$. A linear operator $T : X \to Y$ is *bounded*, if it maps bounded subsets of X to bounded subsets of Y; T is called *compact*, if it maps bounded subsets B of X to relatively compact subsets of Y, which means that $\overline{T(B)} \subseteq Y$ is compact in Y. This already implies that all compact operators are bounded.

Corollary 4.5.5. *For any $a, b \in \mathbb{R}$ with $a < b$ and $p > 1$, the integration map*

$$I : L^p([a, b]) \to C([a, b]), \quad f \mapsto \left(t \mapsto \int_a^t f \,\mathrm{d}\lambda \right)$$

is a compact linear operator.

Proof. The linearity of I is clear. It remains to show that $I(F)$ is relatively compact in $C([a, b])$ for a bounded subset F of $L^p([a, b])$. For that, it suffices to show that Corollary 4.5.4 can be applied to $I(F)$. In fact, for any $I(f) \in I(F)$ and any $x, y \in [a, b]$ with $x < y$, the Hölder inequality yields

$$|I(f)(y) - I(f)(x)| = \left| \int_x^y f \,\mathrm{d}\lambda \right| = \left\| f \cdot \mathbf{1}_{(x,y)} \right\|_{L^1} \leq \|f\|_{L^p} \cdot \left\| \mathbf{1}_{(x,y)} \right\|_{L^q} = \|f\|_{L^p} \cdot |y - x|^{1/q},$$

where q is the Hölder conjugate of p fulfilling $1/p + 1/q = 1$. As F is bounded, $\|f\|_{L^p}$ can be estimated independently of f, which shows that F fulfills the required Hölder condition of fixed power $\alpha = 1/q$ and a fixed constant $K = \sup_{f \in F} \|f\|_{L^p}$.

Moreover, $I(F)$ is uniformly bounded by $K \cdot (b-a)^{1/q}$ as for any $t \in [a,b]$ and $f \in F$ the Hölder inequality implies

$$I(f)(t) \leq \int_a^b |f| \,\mathrm{d}\lambda = \left\| f \cdot \mathbf{1}_{[a,b]} \right\|_{L^1} \leq \|f\|_{L^p} \cdot \left\| \mathbf{1}_{[a,b]} \right\|_{L^q} \leq K \cdot (b-a)^{1/q}.$$

□

Remark 4.5.6. The statement of Corollary 4.5.5 does not hold for $p = 1$. In fact, for $[a,b] = [0,1]$, the sequence $f_n = n \cdot \mathbf{1}_{[0,1/n]}$ is bounded by $\|f_n\|_{L^1} = 1$. However, $F_n(x) := \int_0^x f_n \,\mathrm{d}\lambda = nx \cdot \max\{x, 1/n\}$ converges pointwise to $F(0) := 0$ and $F(x) = 1$ for $x > 0$. Therefore, no subsequence of $(F_n)_{n \in \mathbb{N}}$ converges uniformly in $C([0,1])$, so that the integration operator is not compact from $L^1([0,1])$ to $C([0,1])$.

Proposition 4.5.7. *Every compact linear operator $T : L^p([a,b]) \to C([a,b])$ is sequentially weak-strong continuous for any $p > 1$.*

Proof. This statement is a standard result in the study of compact operators and is shown e.g. in [2, Lemma 8.2]. □

Corollary 4.5.8. *The mappings $f_{i,e}^+ \mapsto F_{i,e}^+$ and $f_e^+ \mapsto F_e^+$ are sequentially weak-strong continuous from $L^p([0,M])$ to $C([0,M])$ for any $M > 0$ and $p > 1$.*

The following property of Vickrey's fluid queuing system has been shown by Cominetti et al. in [7, Lemma 4]. The authors used this property to prove the existence of dynamic Nash equilibrium flows (c.f. Section 4.7.1).

Proposition 4.5.9. *The mapping $f_e^+ \mapsto q_e^f$ is sequentially weak-strong continuous from $L^p([0,M])$ to $C([0,M])$ for any $M > 0$ and $p > 1$.*

Proof. We use the representation of the queue length function given in Corollary 3.4.2 as

$$q_e^f(\theta) := \max_{\xi \leq \theta} \int_\xi^\theta f_{i,e}^+ - \nu_e \,\mathrm{d}\lambda.$$

Now, for any weakly converging sequence $f_e^{+,k}$ in $L^p([0,M])$ we use the equality

$$q_e^f(\theta) = \int_0^\theta f_e^+ - \nu_e \,\mathrm{d}\lambda - \min_{\xi \leq \theta} \int_0^\xi f_e^+ - \nu_e \,\mathrm{d}\lambda = F_e^+(\theta) - \theta \cdot \nu_e - \min_{\xi \leq \theta} \int_0^\xi f_e^+ - \nu_e \,\mathrm{d}\lambda$$

to compute

$$\begin{aligned} \left\| q_e^{f^k} - q_e^f \right\|_\infty &= \sup_{\theta \in [0,M]} \left| \max_{\xi \leq \theta} \int_\xi^\theta f_e^{+,k} - \nu_e \,\mathrm{d}\lambda - \max_{\xi \leq \theta} \int_\xi^\theta f_e^+ - \nu_e \,\mathrm{d}\lambda \right| \\ &\leq \left\| F_e^{+,k} - F_e^+ \right\|_\infty + \sup_{\theta \in [0,M]} \left| \min_{\xi \leq \theta} \int_0^\xi f_e^{+,k} - \nu_e \,\mathrm{d}\lambda - \min_{\xi \leq \theta} \int_0^\xi f_e^+ - \nu_e \,\mathrm{d}\lambda \right|. \end{aligned}$$

The first term converges by Corollary 4.5.8 to 0. Moreover, $\gamma^k : \xi \mapsto \int_0^\xi f_e^{+,k} - \nu_e \, \mathrm{d}\lambda$ converges strongly to $\gamma : \xi \mapsto \int_0^\xi f_e^{+} - \nu_e \, \mathrm{d}\lambda$ in $C([0, M])$.

For $\theta \in \mathbb{R}$ we denote the minimizers as $\xi_{k,\theta} \in \arg\min_{\xi \leq \theta} \gamma^k(\xi)$ and $\xi_\theta \in \arg\min_{\xi \leq \theta} \gamma(\xi)$. Then we compute

$$\begin{aligned}
\left|\min_{\xi\leq\theta} \gamma^k(\xi) - \min_{\xi\leq\theta} \gamma(\xi)\right| &= \left|\gamma^k(\xi_{k,\theta}) - \gamma(\xi_\theta)\right| \\
&\leq \left|\gamma^k(\xi_{k,\theta}) - \gamma(\xi_{k,\theta})\right| + \left|\gamma(\xi_{k,\theta}) - \gamma(\xi_\theta)\right| \\
&\leq \left\|\gamma^k - \gamma\right\|_\infty + \gamma(\xi_{k,\theta}) - \gamma(\xi_\theta) \\
&\leq 3 \cdot \left\|\gamma^k - \gamma\right\|_\infty + \gamma^k(\xi_{k,\theta}) - \gamma^k(\xi_\theta) \leq 3 \cdot \left\|\gamma^k - \gamma\right\|_\infty,
\end{aligned}$$

which goes to 0 as k approaches ∞. The observation

$$\sup_{\theta\in[0,M]} \left|\min_{\xi\leq\theta} \int_0^\xi f_e^{+,k} - \nu_e \, \mathrm{d}\lambda - \min_{\xi\leq\theta} \int_0^\xi f_e^{+} - \nu_e \, \mathrm{d}\lambda\right| = \left\|\gamma^k - \gamma\right\|_\infty$$

concludes the proof of the statement. □

Proposition 4.5.10. *The mapping $f_e^+ \mapsto F_e^-$ is sequentially weak-strong continuous from $L^p([0, M])$ to $C([0, M])$ for any $M > 0$ and $p > 1$.*

Proof. We use the following representation of the total cumulative outflow in terms of the total cumulative inflow and the queue length:

$$F_e^-(\theta) = F_e^+(\theta - \tau_e) - q_e(\theta - \tau_e).$$

Now, the sequential weak-strong continuity is due to Corollary 4.5.8 and Proposition 4.5.9 together with the fact that F_e^- vanishes on $[0, \tau_e]$. □

Lemma 4.5.11. *A predictor $\hat{q}_{i,e}$ that depends continuously on the cumulative inflow, total cumulative outflow and queue length functions is a p-continuous predictor for any $p > 1$.*

Proof. By the continuity of $\gamma_{i,e}$ it is trivial to conclude that $\hat{q}_{i,e}(\,\cdot\,, \,\cdot\,, f) \in C(\mathbb{R}^2, \mathbb{R}_{\geq 0})$ holds for all $e \in E$ and deterministic flows f. It remains to show that the map $f^+ \mapsto \hat{q}_{i,e}(\,\cdot\,, \,\cdot\,, f)$ is sequentially weak-strong continuous from $L^p([0, M], \mathbb{R}_{\geq 0})^{I\times E}$ to $C([0, M] \times D, \mathbb{R}_{\geq 0})$ for every $M > 0$ and compact interval D.

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 Corollary 4.5.8, Proposition 4.5.10 and Proposition 4.5.9, 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 $C(\mathbb{R}, \mathbb{R}_{\geq 0})$ with respect to the extended uniform norm.

We write $g^k := (F_{I\times E}^{+,k}, F_E^{-,k}, q_E^k)$ for all $k \in \mathbb{N}$ and $g := (F_{I\times E}^{+}, F_E^{-}, q_E)$ and conclude that g^k converges to g in $\mathcal{G} := C(\mathbb{R}, \mathbb{R}_{\geq 0})^{(I\times E)+2\cdot E}$ with respect to the extended uniform norm.

We want to show that

$$\left\|\hat{q}_{i,e}(\,\cdot\,,\,\cdot\,,f^k) - \hat{q}_{i,e}(\,\cdot\,,\,\cdot\,,f)\right\|_{C([0,M]\times D)} = \left\|\gamma_{i,e}(\,\cdot\,,\,\cdot\,,g^k) - \gamma_{i,e}(\,\cdot\,,\,\cdot\,,g)\right\|_{C([0,M]\times D)}$$

vanishes as k approaches ∞.

Let $\varepsilon > 0$ be arbitrary. Because $\gamma_{i,e}$ is continuous, for any pair $(\theta, \bar{\theta}) \in [0, M] \times D$ there exists some $\delta_{(\theta,\bar{\theta})} > 0$ such that whenever $\|(\theta, \bar{\theta}, g) - (\theta', \bar{\theta}', g')\|_\infty < \delta_{(\theta,\bar{\theta})}$ for some $g' \in \mathcal{G}$ it holds that

$$\left|\gamma_{i,e}(\theta, \bar{\theta}, g) - \gamma_{i,e}(\theta', \bar{\theta}', g')\right| < \frac{\varepsilon}{2}.$$

The compact set $[0, M] \times D$ is covered by

$$C := \left\{B_{\delta_{(\theta,\bar{\theta})}}(\theta, \bar{\theta}) \,\middle|\, \theta \in [0, M], \bar{\theta} \in D\right\}$$

where $B_r(x) := \{y \in [0, M] \times D \mid \|y - x\|_\infty < r\}$ denotes the open ball in $[0, M] \times D$ around x with radius r. By the compactness of $[0, M] \times D$ there is a subcover using finitely many of these open balls. Let δ be the minimum radius of these finitely many balls. Then we get some $N \in \mathbb{N}$ 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

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

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

4.6 Applied Predictors

In this section we introduce several types of predictors. We start with quite unsophisticated predictors and improve them gradually.

Definition 4.6.1. The *Zero-Predictor* $\hat{q}_i^{\mathrm{Z}}$ predicts that no queues occur, i.e.

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

for all $e \in E$, $\theta, \bar{\theta} \in \mathbb{R}$ and dynamic flows f.

If a commodity $i \in I$ uses the Zero-Predictor, the induced edge costs are constant with $\hat{c}_{i,e}^{\mathrm{Z}}(\,\cdot\,,\,\cdot\,,\,\cdot\,) \equiv \tau_e$ for all $e \in E$. Moreover, predicted shortest paths will always remain the same no matter when the prediction will be taken. One can imagine that agents of that commodity only use a physical map with recorded static edge costs as the information their routing decision is based upon.

Definition 4.6.2. The *constant predictor* $\hat{q}_i^{\mathrm{C}}$ predicts that the queue lengths remain constant from the time of the prediction onwards, i.e.

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

for all $e \in E$, $\theta, \bar{\theta} \in \mathbb{R}$ and dynamic flows f.

Agents using the constant predictor take a snapshot of the queue lengths of all edges at time $\bar\theta$ and assume this current situation for the future. Here, the dynamic edge costs are constant in the argument θ with $\hat c^{\mathrm{C}}_{i,e}(\,\cdot\,,\bar\theta,f) \equiv \tau_e + q^f_e(\bar\theta)/\nu_e$ for all edges $e \in E$. As depicted in Section 4.7.2, this predictor is deeply connected with so-called instantaneous dynamic equilibria (IDE) as introduced in [14, Definition 2.1].

Proposition 4.6.3. *The Zero-Predictor and the constant predictor are oblivious, FIFO-compatible and p-continuous for any $p > 1$.*

Proof. For the Zero-Predictor the statement is trivially fulfilled.

The constant predictor is oblivious, as the queue at time $\bar\theta$ only depends on the total inflow rate up to time $\bar\theta$ as stated in Lemma 3.4.8. The FIFO-compatibility is again fulfilled trivially. The p-continuity results from the predictor's continuous dependence on the queue length: For any sequence of real values $(\bar\theta^k)_{k\in\mathbb{N}}$ converging to some $\bar\theta$ and for every sequence of queue length functions q^k_e converging in $C(\mathbb{R},\mathbb{R}_{\geq 0})$ with respect to the extended uniform norm to some q_e, it holds

$$\left|q^k_e(\bar\theta^k) - q_e(\bar\theta)\right| \leq \left|q^k_e(\bar\theta^k) - q_e(\bar\theta^k)\right| + \left|q_e(\bar\theta^k) - q_e(\bar\theta)\right| \xrightarrow[k\to\infty]{} 0.$$

□

The *linear predictor* $\hat q^{\mathrm{L}}_{i,e}$ as given in Definition 4.2.1 for which an example dynamic prediction flow is depicted in Section 4.2, is a bit more involved. It is defined as

$$\hat q^{\mathrm{L}}_e(\theta,\bar\theta,f) := \left(q^f_e(\bar\theta) + \partial_- q^f_e(\bar\theta)\cdot \min\{\theta-\bar\theta, H\}\right)^+,$$

where $(x)^+ := \max\{x,0\}$ denote the positive part of a real value x, and $H \in \mathbb{R}_{>0}\cup\{\infty\}$ is some prediction horizon. However, as we have seen in Section 4.2 the predictions are not continuous in the argument $\bar\theta$ whenever the gradient $\partial_- q^f_e$ jumps. Therefore, $\hat q^{\mathrm{L}}_{i,e}$ is not p-continuous and hence does not qualify for Theorem 4.4.13, although it is oblivious and FIFO-compatible. However, it remains an open question whether there exists an example network in which there is no dynamic prediction flow with respect to the linear predictor.

A simple regularization technique can be applied to the linear predictor in order to make it p-continuous:

Definition 4.6.4. The *regularized linear predictor* with a window size $\delta > 0$ and a prediction horizon $H \in \mathbb{R}_{>0}\cup\{\infty\}$ is defined as

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

The regularized linear predictor is obtained from the linear predictor by taking the rolling average of the gradient of the queue length in a rolling window with δ time units into the past. This leads to the left-sided difference quotient of q^f_e.

Proposition 4.6.5. *The regularized linear predictor is oblivious, FIFO-compatible and p-continuous for any $p > 1$.*

Proof. It is clear, that $\hat{q}^{\mathrm{RL}}$ is oblivious. The fact that the dynamic cost c_e follows the FIFO order implies that the regularized linear predictor is FIFO-compatible. It remains to show that $\hat{q}^{\mathrm{RL}}$ depends continuously on the queue length function. In the proof for the constant predictor, we have seen that the map $(\bar{\theta}, q) \mapsto q(\bar{\theta})$ is continuous from $\mathbb{R} \times C(\mathbb{R}, \mathbb{R}_{\geq 0})$ to $\mathbb{R}$. The same holds true for the map $(\bar{\theta}, q) \mapsto q(\bar{\theta} - \delta)$ for any $\delta > 0$. Moreover, $(\theta, \bar{\theta}) \mapsto \min\{\theta - \bar{\theta}, H\}$ is continuous. Then, $\hat{q}^{\mathrm{RL}}$ depends continuously on the queue function, as the corresponding map is simply the sum and product of continuous functions. □

Taking a closer look at the regularized linear predictor, we notice that it is just the linear interpolation of two past samples of the queue length function (given H is finite). Now, we enhance this method by allowing the predictor to take a multitude of past samples and predict multiple samples of the future evolution of the queue length. This way we establish a machine-learning based predictor that learns the linear coefficients connecting past queue lengths with the predicted future queue length samples.

More specifically, for an edge $e = vw \in E$ the new predictor uses k_p samples of the past queue lengths at times $\bar{\theta} - \delta \cdot (i-1)$ for $i \in [k_p]$ to predict k_f samples of the future queue at times $\bar{\theta} + \delta \cdot j$ for $j \in [k_f]$. Moreover, the predictor does not only use the past samples of the concerned edge, but also of neighboring edges. Here, the set of neighboring edges is given by $N(e) = \delta_v^- \cup \{e\} \cup \delta_w^+$. This allows the predictor to be aware of flow that occurs in the neighborhood of an edge.

For each edge e, the predictor learns coefficient matrices $W^{e',e} \in \mathbb{R}^{k_p \times k_f}$ for $e' \in N(e)$ as well as biases $\beta \in \mathbb{R}^{k_f}$ to compute interpolation points

$$\hat{q}_e^{\mathrm{ML,raw}}\left(\bar{\theta} + j \cdot \delta, \bar{\theta}, f\right) := \left(\sum_{e' \in N(e)} \sum_{i \in [k_p]} w_{i,j}^{e',e} \cdot q_{e'}^f\left(\bar{\theta} - (i-1) \cdot \delta\right) + \beta_j \right)^+$$

for $j \in [k_f]$. The idea is to linearly interpolate these points which would result in a p-continuous predictor for $p > 1$, as it depends continuously on the queue length functions. However, this raw machine-learned predictor is not guaranteed to be FIFO-compatible. To achieve this, we apply some post-processing: We derive the new interpolation points $\hat{q}_e^{\mathrm{ML}}(\bar{\theta}, \bar{\theta}, f) = q_e^f(\bar{\theta})$ and

$$\hat{q}_e^{\mathrm{ML}}\left(\bar{\theta} + j \cdot \delta, \bar{\theta}, f\right) := \max \left\{ \begin{array}{c} \hat{q}_e^{\mathrm{ML,raw}}\left(\bar{\theta} + j \cdot \delta, \bar{\theta}, f\right), \\ \hat{q}_e^{\mathrm{ML}}\left(\bar{\theta} + (j-1) \cdot \delta, \bar{\theta}, f\right) - \nu_e \cdot \delta \end{array} \right\}$$

for all $j \in [k_f]$ and linearly interpolate these points.

Definition 4.6.6. The predictor $\hat{q}_{i,e}^{\mathrm{ML}}$ is called the *linear regression predictor.*

After applying the post-processing the resulting predictor still depends continuously on the queue length functions which implies the following proposition.

Proposition 4.6.7. *The linear regression predictor is oblivious, FIFO-compatible and p-continuous for any $p > 1$.*

More technical details of the learning procedure used for the linear regression predictor are described in Section 5.4.2.

As the last predictor in this section, we introduce the so-called *perfect predictor*.

Definition 4.6.8. The *perfect predictor* $\hat{q}^{\mathrm{P}}_{i,e}$ predicts the future queue length correctly with

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

It is clear from the definition, that this predictor is FIFO-compatible and depends continuously on the queue length function q^f_e. It is therefore p-continuous for any $p > 1$. However, it is also clear that it is a non-oblivious predictor, and as such it does not fulfill the constraints of Theorem 4.4.13. Nevertheless, the flow in a dynamic prediction equilibrium, in which all commodities use the perfect predictor, is the same as a dynamic Nash equilibrium flow (as analyzed in Section 4.7.1). In this scenario, particles can predict the future evolution of the traffic exactly. Hence, they can choose a perfect shortest path already when starting at the source node.

4.7 Comparison with Existing Forms of Equilibria

Multiple models of equilibrium involving Vickrey's fluid queuing model have been studied in the past. In this section we compare two popular models with dynamic prediction equilibria. As they use the same physical model as introduced in Chapter 3, they differ only in the behavioral model. More specifically, the difference between the models is that the agents base their routing decisions on different levels of detail in the information given to them.

4.7.1 Dynamic Nash Equilibrium Flows

In a so-called *dynamic (Nash) equilibrium* all agents have full information about the future traffic. Using this information, agents in such a dynamic flow can choose a perfect path when starting at their source, because they know the correct future evolution of all queues. This means that they choose an s_i-t_i-path which will turn out to be a perfect choice also in hindsight. In the following, we give a brief formal definition of such an equilibrium based on the definition by Cominetti et al. in [7]. It is weaker than a similar definition given by Koch and Skutella in [17].

Given a feasible dynamic flow f, we use the notation of Chapter 2 with respect to the dynamic cost function $c_e(\theta) := \tau_e + q_e(\theta)/\nu_e$. This means that the *exit time of a finite path* $P = (e_1, \dots, e_k)$ when entering at time θ is given by the concatenation of the exit times of its edges as $T_P(\theta) := (T_{e_k} \circ \dots \circ T_{e_1})(\theta)$. For two nodes $v, w \in V$, we denote the earliest arrival at w when starting at node v at time θ as $l_{v,t}(\theta) := \min_{P \in \mathcal{P}_{v,w}} T_P(\theta)$. For a commodity $i \in I$, we call an edge $e \in E_i$, that lies on a directed s_i-t_i-path, *active at time* θ if $T_e(l_{s_i,v}(\theta)) = l_{s_i,w}(\theta)$ holds true. The set of active times of an edge $e \in E$ is collected in the set $\Theta_{i,e}$.

Definition 4.7.1. A feasible flow f is called a *dynamic (Nash) equilibrium (DE) flow*, if for every $e = vw \in E$, $i \in I$ and for almost all $\theta \in \mathbb{R}$ the following implication holds:

$$f^+_{i,e}(\theta) > 0 \implies \theta \in l_{s_i,v}(\Theta_{i,e}).$$

As explained previously, in dynamic Nash equilibrium flows, agents choose perfect paths. In our framework this would mean that the predictors of all agents are perfect. Therefore, we want to show equivalence between the two concepts. This is achieved with the following theorem.

Theorem 4.7.2. *Let I be a finite set of commodities and f a feasible dynamic flow in a network in which any cycle has a strictly positive transit time. Then f is a dynamic Nash equilibrium flow if and only if $(\hat{q}^{\mathrm{P}}, f)$ is a dynamic prediction equilibrium.*

On the way of proving this statement, we will find a novel characterization of dynamic Nash equilibrium flows. We start by inspecting the perfect predictor more closely:

Proposition 4.7.3. *Let I be a finite set of commodities and f a feasible dynamic flow. The pair $(\hat{q}^{\mathrm{P}}, f)$ is a DPE if and only if for all commodities $i \in I$, $e \in E$ and almost all $\theta \in \mathbb{R}$ it holds that*

$$f^+_{i,e}(\theta) > 0 \implies e \in E_i \wedge l_{w,t_i}(T_e(\theta)) = l_{v,t_i}(\theta).$$

Proof. As the perfect predictor always predicts the queue correctly, i.e. $\hat{q}^{\mathrm{P}}_{i,e}(\theta, \bar{\theta}, f) = q_e(\theta)$, the predicted exit time of edges e and paths P coincide with T_e and T_P. Thus, for $e \in E_i$, $l_{w,t_i}(T_e(\theta)) = l_{v,t_i}(\theta)$ is equivalent to $e \in \hat{E}_i(\theta, \theta, f)$. □

Comparing the above representation of a dynamic prediction flow with respect to the perfect predictor with the definition of dynamic Nash equilibrium flows, the most prominent difference is the perspective of when edges are thought of as *active*. In dynamic Nash equilibrium flows, an active edge $e = vw$ lies on a shortest s_i-w-path, whereas following the DPE notion a (predicted) active edge lies on a shortest v-t_i-path.

To realize that these two notions of equilibria are indeed equivalent, we need a series of intermediate results. For a feasible flow f and an edge $e = vw$ reachable from s_i we define the cumulative in- and outflow up to the earliest arrival time when starting at time θ in s_i as $x^+_{i,e}(\theta) \coloneqq F^+_{i,e}(l_{s_i,v}(\theta))$ and $x^-_{i,e}(\theta) \coloneqq F^-_{i,e}(l_{s_i,w}(\theta))$. Moreover, we introduce the cumulative balance functions $B_{i,v}(\theta) \coloneqq \sum_{e \in \delta^+_v} F^-_{i,e}(\theta) - \sum_{e \in \delta^-_v} F^+_{i,e}(\theta)$ for all $v \in V$ and $i \in I$. The value $b_{i,v}(\theta) \coloneqq B_{i,v}(l_{s_i,v}(\theta))$ denotes the cumulative balance at the earliest arrival when starting at time θ in s_i for all nodes $v \in V$ that are reachable from s_i. In [19, Theorem 3.1.11] the following theorem based on techniques introduced in [17] was shown:

Theorem 4.7.4. *Given a feasible flow f in a network in which any cycle has positive transit time, the following statements are equivalent:*

(i) The flow f is a dynamic Nash equilibrium flow.

(ii) For all $i \in I$ and $e \in E$ reachable from s_i it holds $x^+_{i,e} = x^-_{i,e}$.

(iii) For all $i \in I$ it holds $b_{i,s_i} = -b_{i,t_i}$.

Before we show the equivalence, we will prove a similar characterization of dynamic prediction flows that use the perfect predictor. Analogously to the above theorem, we define $\hat{x}^+_{i,e}(\theta) \coloneqq F^+_{i,e}(l^{\leftarrow}_{v,t_i}(\theta))$ and $\hat{x}^-_{i,e}(\theta) \coloneqq F^-_{i,e}(l^{\leftarrow}_{w,t_i}(\theta))$ as the cumulative in- and outflow up to the latest departure time for arriving at time θ at the sink t_i for all $e \in E_i$ (utilizing the notation described in Section 2.3). Correspondingly, $\hat{b}_{i,v}(\theta) \coloneqq B_{i,v}(l^{\leftarrow}_{v,t_i}(\theta))$ describes the cumulative balance of node v at the latest departure time to arrive at t_i at time θ for all $v \in V_i$.

Theorem 4.7.5. *Given a feasible flow f in a network in which any cycle has positive transit time, the following statements are equivalent:*

(i) The pair $(\hat{q}^{\mathrm{P}}, f)$ is a dynamic prediction equilibrium.

(ii) For all commodities $i \in I$, $e \in E$ and almost all $\theta \in \mathbb{R}$ it holds that

$$f^{+}_{i,e}(\theta) > 0 \implies e \in E_i \wedge l_{w,t_i}(T_e(\theta)) = l_{v,t_i}(\theta).$$

(iii) For all $i \in I$ it holds $\hat{x}^{+}_{i,e} = \hat{x}^{-}_{i,e}$ for $e \in E_i$, and $f^{+}_{i,e} \overset{a.e.}{\equiv} 0$ for $e \notin E_i$.

(iv) For all $i \in I$ it holds $\hat{b}_{i,s_i} = -\hat{b}_{i,t_i}$.

Before we prove the above theorem, we begin with a simple observation on the connection between $\hat{x}^{+}_{i,e}$ and $\hat{x}^{-}_{i,e}$.

Proposition 4.7.6. *For a feasible flow f, $\hat{x}^{+}_{i,e}(\theta) \geq \hat{x}^{-}_{i,e}(\theta)$ holds true for all $i \in I$, $e \in E_i$, and $\theta \in \mathbb{R}$.*

Proof. For any edge $e = vw \in E_i$ and $\theta \in \mathbb{R}$ we have $T^{\leftarrow}_e(l^{\leftarrow}_{w,t_i}(\theta)) \leq l^{\leftarrow}_{v,t_i}(\theta)$. Using Proposition 3.3.3 it follows

$$\hat{x}^{-}_{i,e}(\theta) = F^{-}_{i,e}(l^{\leftarrow}_{w,t_i}(\theta)) = F^{+}_{i,e}(T^{\leftarrow}_e(l^{\leftarrow}_{w,t_i}(\theta))) \leq F^{+}_{i,e}(l^{\leftarrow}_{v,t_i}(\theta)) = \hat{x}^{+}_{i,e}(\theta).$$

□

Proof of Theorem 4.7.5. The equivalence of (i) and (ii) is due to Proposition 4.7.3.

Next, we show $(ii) \Leftrightarrow (iii)$. Let f be a feasible flow and let

$$\hat{\Theta}_{i,e} \coloneqq \{\theta \in \mathbb{R} \,|\, e \in E_i \wedge l_{w,t_i}(T_e(\theta)) = l_{v,t_i}(\theta)\}$$

denote the set of times an edge is active for commodity i as predicted by the perfect predictor. For an edge $e = vw \in E_i$ we define

$$\omega_{i,e}(\theta) \coloneqq \max\{\omega \leq \theta \,|\, l_{w,t_i}(T_e(\omega)) = l_{v,t_i}(\theta)\}.$$

Claim 4.7.7. *The complement of $\hat{\Theta}_{i,e}$ is denoted by $\hat{\Theta}^c_{i,e}$. Then, for all $e \in E_i$ it holds that*

$$\hat{\Theta}^c_{i,e} = \bigcup_{\theta \in \mathbb{R}} (\omega_{i,e}(\theta), \theta).$$

Proof. Assume $\theta \in \hat{\Theta}^c_{i,e}$, or equivalently $l_{w,t_i}(T_e(\theta)) > l_{v,t_i}(\theta)$. By the continuity of l_{w,t_i}, T_e and l_{v,t_i}, there exists some $\varepsilon > 0$ such that $l_{w,t_i}(T_e(\theta')) > l_{v,t_i}(\theta+\varepsilon)$ holds for all $\theta' \in [\theta, \theta+\varepsilon]$. This implies $\omega_{i,e}(\theta+\varepsilon) < \theta$ and thus $\theta \in (\omega_{i,e}(\theta+\varepsilon), \theta+\varepsilon)$.

On the other hand, let θ be contained in an interval $(\omega_{i,e}(\xi), \xi)$ for some $\xi \in \mathbb{R}$. Then $l_{w,t_i}(T_e(\theta)) > l_{v,t_i}(\xi) \geq l_{v,t_i}(\theta)$ holds by definition of $\omega_{i,e}$ and the monotonicity of l_{v,t_i} implying $\theta \in \hat{\Theta}^c_{i,e}$. ■

Cominetti et al. showed in [7, Lemma 8] that given a possibly uncountable family $\{(a_j, b_j)\}_{j \in J}$ of intervals, a function $g \in L^1_{\mathrm{loc}}(\mathbb{R}, \mathbb{R}_{\geq 0})$ vanishes on (a_j, b_j) for every $j \in J$

if and only if it vanishes on the union $\bigcup_{j\in J}(a_j, b_j)$ (c.f. [19, Lemma 3.1.7]). Applying this statement together with the observation above proves the following claim.

Claim 4.7.8. *The pair $(\hat{q}^{\mathrm{P}}, f)$ is a DPE if and only if $f^+_{i,e}$ vanishes on all intervals $(\omega_{i,e}(\theta), \theta)$ for $e \in E_i$ and on all of $\mathbb{R}$ for edges $e \notin E_i$.*

For any $e \in E_i$ and $\theta \in \mathbb{R}$ it holds that

$$l_{w,t_i}\left(T_e(\omega_{i,e}(l^{\leftarrow}_{v,t_i}(\theta)))\right) = l_{v,t_i}(l^{\leftarrow}_{v,t_i}(\theta)) = \theta.$$

Therefore, we can infer $l^{\leftarrow}_{w,t_i}(\theta) \geq T_e(\omega_{i,e}(l^{\leftarrow}_{v,t_i}(\theta)))$ which implies

$$F^+_{i,e}(\omega_{i,e}(l^{\leftarrow}_{v,t}(\theta)) = F^-_{i,e}\left(T_e(\omega_{i,e}(l^{\leftarrow}_{v,t_i}(\theta)))\right) \leq F^-_{i,e}(l^{\leftarrow}_{w,t_i}(\theta)) = \hat{x}^-_e(\theta).$$

If $(\hat{q}^{\mathrm{P}}, f)$ is a dynamic prediction equilibrium, $F_{i,e}(l^{\leftarrow}_{v,t_i}(\theta)) = F^+_{i,e}(\omega_{i,e}(l^{\leftarrow}_{v,t_i}(\theta)))$ holds because $f^-_{i,e}$ vanishes on $(\omega_{i,e}(l^{\leftarrow}_{v,t_i}(\theta)), l^{\leftarrow}_{v,t_i}(\theta))$. Thus, it holds that $\hat{x}^+_{i,e}(\theta) \leq \hat{x}^-_{i,e}(\theta)$ which combined with Proposition 4.7.6 shows (iii).

Assume (iii) is fulfilled and let $\theta \in \mathbb{R}$ with $\omega_{i,e}(\theta) < \theta$. Then, by the definition of $\omega_{i,e}$, it follows $\omega_{i,e}(\theta) = T^{\leftarrow}_e(l^{\leftarrow}_{w,t_i}(l_{v,t_i}(\theta)))$. Using this observation we deduce

$$\begin{aligned} F^+_{i,e}(\omega_{i,e}(\theta)) &= F^-_{i,e}(l^{\leftarrow}_{w,t_i}(l_{v,t_i})) = \hat{x}^-_{i,e}(l_{v,t_i}(\theta)) \\ &= \hat{x}^+_{i,e}(l_{v,t_i}(\theta)) = F^+_{i,e}(l^{\leftarrow}_{v,t_i}(l_{v,t_i}(\theta)) \geq F^+_{i,e}(\theta), \end{aligned}$$

where the last inequality stems from $l^{\leftarrow}_{v,t_i}(l_{v,t_i}(\theta)) \geq \theta$. As $\omega_{i,e}(\theta)$ is smaller than θ, this shows that $f^+_{i,e}$ vanishes on $(\omega_{i,e}(\theta), \theta)$.

Lastly, we prove the equivalence $(iii) \Leftrightarrow (iv)$. We observe that for any $v \in V_i \setminus \{s_i, t_i\}$ the flow conservation property (F4) implies $\hat{b}_{i,v} \equiv 0$. Moreover, $\hat{b}_{i,v}$ has the representation

$$\hat{b}_{i,v}(\theta) = \sum_{e\in\delta^-_v} \hat{x}^-_{i,e}(\theta) - \sum_{e\in\delta^+_v\cap E_i} \hat{x}^+_{i,e}(\theta) - \sum_{e\in\delta^+_v\setminus E_i} F^+_{i,e}(l^{\leftarrow}_{v,t_i}(\theta).$$

Taking the sum over all $v \in V_i$ yields

$$\hat{b}_{i,s_i}(\theta) + \hat{b}_{i,t_i}(\theta) = \sum_{e\in E_i}\left(\hat{x}^+_{i,e}(\theta) - \hat{x}^-_{i,e}(\theta)\right) + \sum_{v\in V_i, e\in\delta^+_{v_i}\setminus E_i} F^+_{i,e}(l^{\leftarrow}_{v,t_i}(\theta)).$$

From this equation, the implication $(iii) \Rightarrow (iv)$ follows trivially. For $(iv) \Rightarrow (iii)$, assume the expression above equals 0. As all summands are non-negative by Proposition 4.7.6, each summand must vanish. This implies $\hat{x}^+_{i,e} = \hat{x}^-_{i,e}$ for all $e \in E_i$ and $f^+_{i,e} \equiv 0$ for all $e \in \delta^+(V_i)$. By Proposition 3.3.5 this also implies $f^+_{i,e} \equiv 0$ for all $e \notin E_i$. □

Now that we have a characterization on both dynamic Nash equilibrium flows and dynamic prediction flows with respect to the perfect predictor, the idea is to prove equivalence of the two notions using the source's and sink's balance values. The next proposition is the final tool that allows us to compare the corresponding balance values $\hat{b}_{i,v_i}$ and b_{i,v_i} in a meaningful way.

Proposition 4.7.9. *Let $(\hat{q}^{\mathrm{P}}, f)$ be a DPE in a network in which any cycle has positive transit time, and let $i \in I$ and $v \in V_i$. If the function l_{v,t_i} is constant on an interval $[a,b]$ for some $a < b$, then $f^+_{i,e}|_{[a,b]} \equiv 0$ holds for all outgoing edges $e \in \delta^+_v$.*

Proof. For $v = t_i$ the statement follows from the fact that l_{t_i,t_i} is never constant on a non-degenerate interval. For $v \neq t_i$ we assume the contrary: Let l_{v_1,t_i} be constant with value L on an interval $[\theta_0, \theta^*]$ with $\theta_0 < \theta^*$ and assume there is an edge $e_1 = v_1v_2 \in \delta^+_{v_1}$ with $F^+_{i,e_1}(\theta_0) < F^+_{i,e_1}(\theta^*)$. We define $\theta_1 \coloneqq \max\{\theta \leq \theta^* \mid l_{v_2,t_i}(T_{e_1}(\theta)) = L\}$. Because of the DPE property, f^+_{i,e_1} vanishes on (θ_1, θ^*) and hence $F^+_{i,e_1}(\theta_0) < F^+_{i,e_1}(\theta^*) = F^+_{i,e_1}(\theta_1)$. This implies $T_{e_1}(\theta_0) < T_{e_1}(\theta_1)$. Moreover, l_{v_2,t_i} is constant on $[T_{e_1}(\theta_0), T_{e_1}(\theta_1)]$ with value L.

We now construct an infinite sequence of edges $e_j = v_jv_{j+1}$, times θ_j and paths $P_j = (e_1, \ldots, e_j)$ such that the following properties hold for all $j \in \mathbb{N}$:

(i) $F^+_{i,e_j}(T_{P_{j-1}}(\theta_0)) < F^+_{i,e_j}(T_{P_{j-1}}(\theta_j))$,

(ii) l_{v_{j+1},t_i} is constant on $[T_{P_j}(\theta_0), T_{P_j}(\theta_j)]$ with value L,

(iii) $T_{P_j}(\theta_0) < T_{P_j}(\theta_j)$,

(iv) $v_{j+1} \in E_i \setminus \{t_i\}$.

For $j = 1$ the above construction already fulfills the required properties. Assume, we have constructed the sequence up to $j \in \mathbb{N}$. By flow conservation and properties (i) and (iv), there exists an edge $e_{j+1} = v_{j+1}v_{j+2} \in \delta^+_{v_{j+1}}$ with $F^+_{i,e_{j+1}}(T_{P_j}(\theta_0)) < F^+_{i,e_{j+1}}(T_{P_j}(\theta_j))$. We define $\theta_{j+1} \coloneqq \max\{\theta \leq \theta_j \mid l_{v_{j+2},t_i}(T_{P_{j+1}}(\theta)) = L\}$. By property (ii), $f^+_{i,e_{j+1}}$ vanishes on $(T_{P_j}(\theta_{j+1}), T_{P_j}(\theta_j))$ which implies

$$F^+_{i,e_{j+1}}(T_{P_j}(\theta_0)) < F^+_{i,e_{j+1}}(T_{P_j}(\theta_j)) = F^+_{i,e_{j+1}}(T_{P_j}(\theta_{j+1})).$$

Moreover, l_{v_{j+2},t_i} is constant on $[T_{P_{j+1}}(\theta_0), T_{P_{j+1}}(\theta_{j+1})]$ with value L. Because $f^+_{i,e_{j+1}}$ does not vanish almost everywhere on $[T_{P_j}(\theta_0), T_{P_j}(\theta_{j+1})]$, we conclude $T_{P_{j+1}}(\theta_0) < T_{P_{j+1}}(\theta_{j+1})$. Finally, this implies $v_{j+2} \neq t_i$.

Such an infinite sequence cannot exist: Because the transit time of any cycle is positive, the value $T_{P_j(\theta_0)}$ gets arbitrarily large for $k \to \infty$. For $T_{P_j(\theta_0)} > L$, the inequality

$$L = l_{v_{j+1},t_i}(T_{P_j}(\theta_0)) \geq T_{P_j}(\theta_0) > L$$

yields a contradiction. □

Proof of Theorem 4.7.2. By Theorems 4.7.4 and 4.7.5 it suffices to prove the equivalence

$$b_{i,s_i} = -b_{i,t_i} \iff \hat{b}_{i,s_i} = -\hat{b}_{i,t_i}$$

for all commodities $i \in I$. Let us first assume $b_{i,s_i}(\theta) = -b_{i,t_i}(\theta)$ holds for all $\theta \in \mathbb{R}$. Then for all $\theta \in \mathbb{R}$, we compute

$$\begin{aligned}\hat{b}_{i,s_i}(\theta) &= B_{i,s_i}(l^{\leftarrow}_{s_i,t_i}(\theta)) = b_{i,s_i}(l^{\leftarrow}_{s_i,t_i}(\theta)) = -b_{i,t_i}(l^{\leftarrow}_{s_i,t_i}(\theta)) \\ &= -B_{i,t_i}\left(l_{s_i,t_i}(l^{\leftarrow}_{s_i,t_i}(\theta))\right) = -B_{i,t_i}(\theta) = -\hat{b}_{i,t_i}(\theta).\end{aligned}$$

For the other direction, pretend that $\hat{b}_{i,s_i} = -\hat{b}_{i,t_i}$ is fulfilled. For $\theta \in \mathbb{R}$ we have

$$-b_{i,t_i}(\theta) = -B_{i,t_i}(l_{s_i,t_i}(\theta)) = -\hat{b}_{i,t_i}(l_{s_i,t_i}(\theta)) = \hat{b}_{i,s_i}(l_{s_i,t_i}(\theta)) = B_{i,s_i}\left(l^{\leftarrow}_{s_i,t_i}(l_{s_i,t_i}(\theta))\right).$$

We aim to show that this expression equals

$$b_{i,s_i}(\theta) = B_{i,s_i}(l_{s_i,s_i}(\theta)) = B_{i,s_i}(\theta).$$

Let $\theta' := l^{\leftarrow}_{s_i,t_i}(l_{s_i,t_i}(\theta)) \geq \theta$. The function l_{s_i,t_i} is constant on $[\theta, \theta']$ and thus Proposition 4.7.9 states that $f^+_{i,e}$ vanishes on $[\theta, \theta']$ for all $e \in \delta^+_{s_i}$. Therefore, $\sum_{e \in \delta^+_{s_i}} F^+_e$ is constant on $[\theta, \theta']$. The flow conservation property states that $\sum_{e \in \delta^-_{s_i}} F^-_e$ must remain constant on $[\theta, \theta']$, which implies that B_{i,s_i} is constant on $[\theta, \theta']$, concluding our proof. □

We summarize our results in the following theorem.

Theorem 4.7.10. *Given a feasible flow f in a network, in which any cycle has positive transit time, the following statements are equivalent:*

(i) The flow f is a dynamic Nash equilibrium flow.

(ii) The pair $(\hat{q}^{\mathrm{P}}, f)$ is a dynamic prediction equilibrium.

(iii) For all $i \in I$ and $e \in E$ reachable from s_i it holds $x^+_{i,e} = x^-_{i,e}$.

(iv) For all $i \in I$ it holds $\hat{x}^+_{i,e} = \hat{x}^-_{i,e}$ for $e \in E_i$, and $f^+_{i,e} \overset{a.e.}{\equiv} 0$ for $e \notin E_i$.

(v) For all $i \in I$ it holds $b_{i,s_i} = -b_{i,t_i}$.

(vi) For all $i \in I$ it holds $\hat{b}_{i,s_i} = -\hat{b}_{i,t_i}$.

Koch and Skutella showed in [17] that for piecewise right-constant inflow rates, i.e. functions $u_i : \mathbb{R} \to \mathbb{R}$ such that for any θ there is some $\varepsilon > 0$ with $u_i|_{[\theta,\theta+\varepsilon)} \equiv u_i(\theta)$, there always exists a dynamic Nash equilibrium flow. They are not necessarily unique, however, their corresponding arrival functions $l_{s_i,v}$ are.

As stated previously, the perfect predictor is not oblivious and therefore the existence result for DPE cannot be used to obtain existence of dynamic Nash equilibrium flows. However, for inflow rates $u_i \in L^p_{\mathrm{loc}}([0, M], \mathbb{R}_{\geq 0})$ with $1 < p < \infty$, $M \in \mathbb{R}_{>0}$, and strictly positive transit times $\tau_e > 0$, Cominetti et al. have shown existence in [7, Theorem 8] using a similar technique based on the variational equality. This gives rise to the question whether we can drop the obliviousness requirement in the existence theorem for DPE if we restrict ourselves to inflow functions with a bounded support.

4.7.2 Instantaneous Dynamic Equilibrium Flows

Instantaneous Dynamic Equilibrium Flows have been introduced by Graf et al. in [14] to reflect the fact that real-world traffic navigation systems obtain traffic data in real-time. In their model, agents base their routing decisions only on the current traffic load of the edges. In comparison with dynamic Nash equilibrium flows, this is a more realistic scenario, as agents cannot know the future evolution of the traffic.

More specifically, Graf et al. define the so-called *instantaneous edge cost* at time θ as $\bar{c}_e(\theta) := \tau_e + q_e(\theta)/\nu_e$ for a given feasible flow f. The *instantaneous shortest distance* $d_{v,i}$ to the sink t_i when starting at a node $v \in V_i$ at time θ is then defined as

$$d_{i,v}(\theta) = \min_{P \in \mathcal{P}_{v,t_i}} \sum_{e \in P} \bar{c}_e(\theta).$$

Remark 4.7.11. In [14], edge transit times were expected to be strictly positive, and the instantaneous shortest distances were defined as the unique solution to

$$d_{i,v}(\theta) := \begin{cases} 0, & \text{for } v = t_i, \\ \min_{e=vw \in E} d_{i,w}(\theta) + \bar{c}_e(\theta), & \text{for } v \in V_i \setminus \{t_i\}. \end{cases}$$

If there exist cycles C with $\sum_{e \in C} \bar{c}_e(\theta) = 0$ this solution is no longer unique. To perform a more general comparison with dynamic prediction equilibria we here explicitly allow $\tau_e = 0$ and cycles of zero transit time.

An edge $e = vw \in E_i$ is called *instantaneously active at time* θ, if $d_{i,v}(\theta) = d_{i,w}(\theta) + \bar{c}_e(\theta)$ holds true. The set of instantaneously active edges at time θ is denoted as $E_{i,\theta}$.

Definition 4.7.12. A feasible dynamic flow f is called *instantaneous dynamic equilibrium (IDE) flow* if for all $i \in I$, $e \in E$ and $\theta \in \mathbb{R}$ it holds that

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

We already saw that the behavior prescribed in an IDE flow deeply corresponds with the behavior of agents using the constant predictor in a DPE. In both cases, agents use a current snapshot of the traffic and assume these conditions for the future. The following lemma confirms this perception.

Lemma 4.7.13. *A feasible flow f is an instantaneous dynamic equilibrium flow if and only if the pair $(\hat{q}^{\mathrm{C}}, f)$ is a dynamic prediction equilibrium.*

Proof. For any $\theta \in \mathbb{R}$, the predicted cost when using the constant predictor resolves to $\hat{c}_{i,e}(\,\cdot\,, \theta, f) \equiv \bar{c}_e(\theta)$. Given any $\theta, \bar{\theta} \in \mathbb{R}$, this implies

$$\hat{l}_{i,v}(\theta, \bar{\theta}, f) = \min_{P \in \mathcal{P}_{v,t_i}} \hat{T}_{i,P}(\theta, \bar{\theta}, f) = \min_{P \in \mathcal{P}_{v,t_i}} \theta + \sum_{e \in P} \bar{c}_e(\bar{\theta}) = \theta + d_{i,v}(\bar{\theta}).$$

Moreover, the predicted delay equates to

$$\begin{aligned} \hat{\Delta}_{i,e}(\theta, \bar{\theta}, f) &= \hat{l}_{i,w}\left(\hat{T}_{i,e}(\theta, \bar{\theta}, f), \bar{\theta}, f\right) - \hat{l}_{i,v}(\theta, \bar{\theta}, f) \\ &= \theta + \bar{c}_e(\bar{\theta}) + d_{i,w}(\bar{\theta}) - (\theta + d_{i,v}(\bar{\theta})) = c_e(\bar{\theta}) + d_{i,w}(\bar{\theta}) - d_{i,v}(\bar{\theta}). \end{aligned}$$

Hence, the condition $e \in \hat{E}_i(\theta, \theta, f)$ is equivalent to $e \in E_{i,\theta}$, showing equivalence between the two models of equilibrium. ☐

Graf et al. have shown in [14, Theorem 4.6] that if the support of all network inflow rates is bounded, then any IDE flow f terminates. This means that there exists some

$H > 0$ such that $B_{s,i}(\theta) = -B_{t,i}(\theta)$ holds for all $\theta \geq H$ and commodities $i \in I$. Here, $B_{v,i}(\theta) \coloneqq \sum_{e\in\delta_v^+} F_{i,e}^+(\theta) - \sum_{e\in\delta_v^-} F_{i,e}^-(\theta)$ denotes the cumulative flow balance of a node v.

It remains an open question which properties a set of predictors has to fulfill such that a dynamic prediction equilibrium terminates given network inflow rates with bounded support.

In [12, Theorem 3.7] Graf and Harks gave an algorithm that computes IDE flows in single-source networks in finite time. Analogous theorems for DPE strongly depend on the predictors used. However, it might be possible to exploit the same techniques used in [12] to prove the existence of DPE for the linear predictor in single-sink networks, and even formulate a finite-time algorithm for its computation.

5 Computing Approximate Dynamic Prediction Equilibria

In this chapter we discuss the computation of dynamic prediction equilibria. Here, we introduce an extension-based procedure, which requires that all predictors are oblivious and FIFO-compatible. Moreover, we do not discretize the time of the particles' flow, but instead we discretize the points in time at which agents calculate their routes. This gives rise to the definition of an *approximated dynamic prediction equilibrium*.

In the first section, we formally define approximated dynamic prediction equilibria and give an outline of the procedure computing these objects. After that, we study the exact behavior of deterministic edge outflow rate functions given piecewise constant edge inflow rates in Section 5.2. This is used in Section 5.3 to show that the proposed algorithm is correct and that it terminates. Finally, in Section 5.4, an experimental study that compares the performance of different predictors in terms of their average travel times is presented.

5.1 Outline of the Extension-Based Approximation Algorithm

In dynamic prediction equilibria, agents can adjust their route at every intersection. As agents behave like infinitesimal particles in a continuous flow model, these routes may change in a continuous manner. However, for computational analyses it is infeasible to calculate shortest paths in a continuous manner. Therefore, in this computational study, we discretize the points in time at which routes are recalculated.

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

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

leading to the following variant of DPE.

Definition 5.1.1. Let f be a dynamic flow, $\hat{q} = (\hat{q}_{i,e})_{i,e}$ a set of predictors and $\varepsilon > 0$. The pair $(f, \hat{q})$ is called an *ε-approximated dynamic prediction equilibrium (ε-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

$$f^+_{i,e}(\theta) > 0 \implies e \in \hat{E}_i\left(\vartheta_\varepsilon(\theta), \vartheta_\varepsilon(\theta), f\right).$$

In other words, whenever there is positive inflow into an edge, that edge must be active at time $k \cdot \varepsilon$ as predicted at time $k \cdot \varepsilon$ with $k \cdot \varepsilon \leq \theta < (k+1) \cdot \varepsilon$.

The next restriction we require is that network inflow rates are piecewise constant with finitely many jumps. This has the benefit that all edge inflow and outflow rates of a dynamic flow are piecewise constant as well. Cumulative edge inflow and outflow functions, as well as queue length and exit time functions of edges are piecewise linear as a result. The same applies to the exit time functions of paths and the earliest arrival time functions, because the composition and the minimum of piecewise linear functions are in turn piecewise linear. Consequently, we require that the predicted queue length functions $\hat{q}_{i,e}(\,\cdot\,,\bar{\theta},f)$ of all predictors can be expressed as piecewise linear functions with finitely many breakpoints. We specify these classes of functions formally:

Definition 5.1.2. A function $f:\mathbb{R}\to\mathbb{R}$ is called *(right) piecewise constant* if there exists a finite chain $\xi_1<\xi_2<\dots<\xi_k$ with $\xi_i\in\mathbb{R}$ such that f is constant on $(-\infty,\xi_1)$, on $[\xi_i,\xi_{i+1})$ for $i\in[k-1]$ and on $[\xi_k,\infty)$.

Definition 5.1.3. A function $f:\mathbb{R}\to\mathbb{R}$ is called *piecewise linear* if there exist values $\delta_l,\delta_r\in\mathbb{R}$ and a finite chain $\xi_1<\xi_2<\dots<\xi_k$ with $\xi_i\in\mathbb{R}$ such that for all $x\in\mathbb{R}$ it holds

$$f(x)=\begin{cases} f(\xi_1)+(x-\xi_1)\cdot\delta_l, & \text{for } x<\xi_1,\\ f(\xi_i)+(x-\xi_i)\cdot\frac{f(\xi_{i+1})-f(\xi_i)}{\xi_{i+1}-\xi_i}, & \text{for } \xi_i\le x<\xi_{i+1}, i\in[k-1],\\ f(\xi_k)+(x-\xi_k)\cdot\delta_r, & \text{for } x\ge\xi_k.\end{cases}$$

The predictors allowed in our simulation need to be oblivious and FIFO-compatible. The first requirement is necessary, as we build our flow in an extension based approach, so that predictions cannot rely upon the future evolution of the flow. The second requirement allows us to use the algorithms introduced in Chapter 2 to compute shortest paths.

In the following we give an overview on the computation of ε-DPE. Initially, we begin with a “zero”-flow f with $f^+,f^-:\equiv 0$. This is our ε-DPE flow up to time $H=0$, where we call H the *flow horizon* of f. We aim to compute an ε-DPE flow up to time $H_{\text{comp}}\in\mathbb{R}$. 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 whenever the ε-DPE flow f has been calculated up to some multiple $H=k\cdot\varepsilon$ of ε. All routes that have been calculated up to that time 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 piecewise 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 $i\in I$.

The *distribution phase* is responsible for extending the dynamic flow. Assume we have computed an ε-DPE flow f up to time H and want to extend this flow horizon H up to some time $H'>H$. The next routing step based on the flow horizon H is $\vartheta_\varepsilon(H)+\varepsilon$. To follow the conservation constraint (F4), we need to distribute the node inflow rate

$$b_{i,v}^-(\theta):=\sum_{e\in\delta_v^-}f_{i,e}^-(\theta)+\mathbf{1}_{v=s_i}\cdot u_i(\theta)$$

of a node v to outgoing edges for each commodity $i\in I$. A distribution phase then consists of the following two steps.

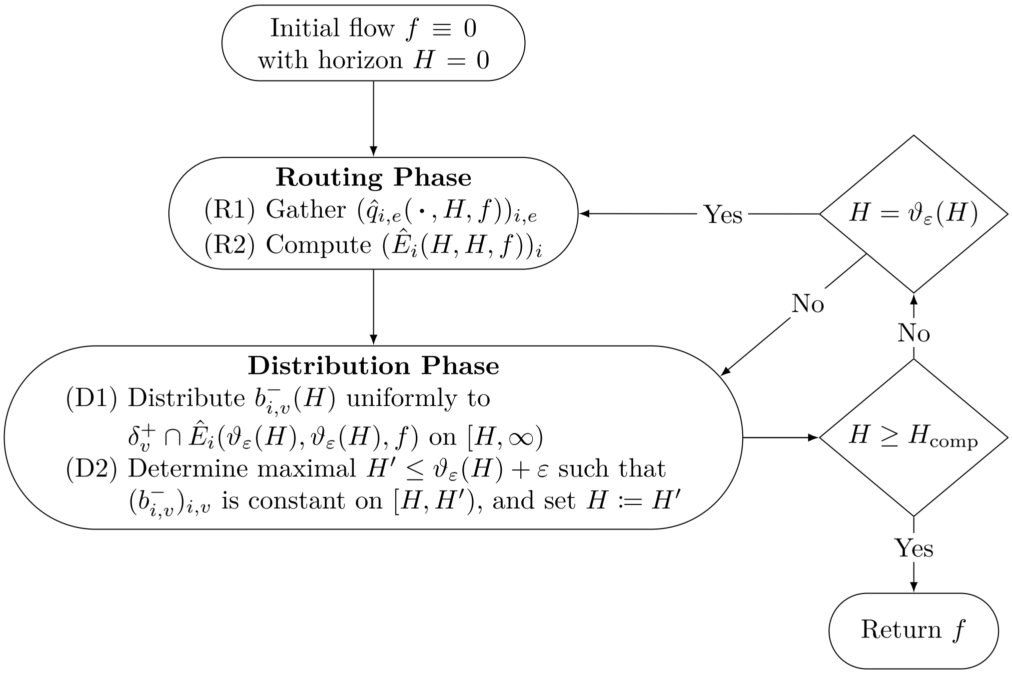

Figure 5.1: A schematic overview of the computation of an ε-DPE.

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

$$f^+_{i,e}|_{[H,\infty)} \coloneqq \begin{cases} \frac{b^-_{i,v}(H)}{\left|\delta^+_v \cap \hat{E}_i(\vartheta_\delta(H), \vartheta_\delta(H), f)\right|}, & \text{if } e \in \hat{E}_i(\vartheta_\delta(H), \vartheta_\delta(H), f), \\ 0, & \text{otherwise,} \end{cases}$$

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 $H \coloneqq H'$.

The distribution phase is executed multiple times until the flow horizon H reaches the next multiple of ε; 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 5.1

Remarks on the Implementation

There are three major considerations that one has to contemplate when implementing this algorithm.

While it is relatively cheap to compute the predicted queue length functions $\hat{q}_{i,e}(\,\cdot\,, H, f)$ for the different predictors introduced in Section 4.6, it is computationally expensive to retrieve 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 done in Step (D1). And 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 $f^-_{i,e}$ (induced by changing edge inflow rates) results in a change of the target node's inflow rate $b^-_{i,w}$ (for $e = vw$). The solution to these two problems is discussed in Sections 5.2 and 5.3. Additionally, because the distribution phase may be run very often, it is essential to do only as much work as necessary. Here, we can reduce the work load in consecutive distribution phases by only considering nodes v whose inflow rates $(b^-_{i,v})_i$ changed at time H.

The problem of computing the set $\hat{E}_i(H, H, f)$ of active edges of a commodity i can be solved by the Dynamic Bellman-Ford algorithm as explained in Section 2.6. However, because of its poor running time, it is more beneficial to use the approach described in Section 2.5 involving the Dynamic Dijkstra algorithm. With this approach we have to calculate the active outgoing edges for each node and commodity with a separate run of the Dynamic Dijkstra Procedure. Therefore, the idea is to delay 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 be experience inflow of a certain commodity; experiments have shown that this approach is 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, this approach is then obviously the preferred method to compute $\hat{E}_i(H, H, f)$.

Adjustments for Computing IDE flows

In Step (D1) we choose to distribute the node inflow rate equally to all active outgoing edges. Of course, this does not ensure, that edges that were active at time $\vartheta(H)$ will remain active throughout the interval (H, H'). In the case of instantaneous dynamic equilibria, that is, if all commodities use the constant predictor, and if we have only a single sink, we can use the more sophisticated so-called *water filling algorithm* (c.f. [14, Online Appendix Section 2]) to distribute the node inflow to the active outgoing edges. This algorithm ensures that edges remain active for some time period.

When computing an ε-DPE we usually deal with arbitrary predictors, and hence use the heuristic to distribute incoming flow uniformly across the active outgoing edges. However, the adjustments necessary to compute exact IDE flows are modest: In the routing phase, we can use a static version of Dijkstra's algorithm to compute the instantaneous active edges $(E_{i,H})$ for some time H. We then compute the maximal H' such that no node inflow rate changes and the set of active edges remains the same on (H, H'). Then, we distribute the node inflow rate according to the water filling algorithm and repeat this procedure until we have reached the targeted flow horizon H_{comp}. The adjusted procedure is depicted in Figure 5.2.

For further analysis of the resulting “natural extension algorithm” the reader is referred to [12]. A key result therein is that there is an implementation of this algorithm that is guaranteed to terminate.

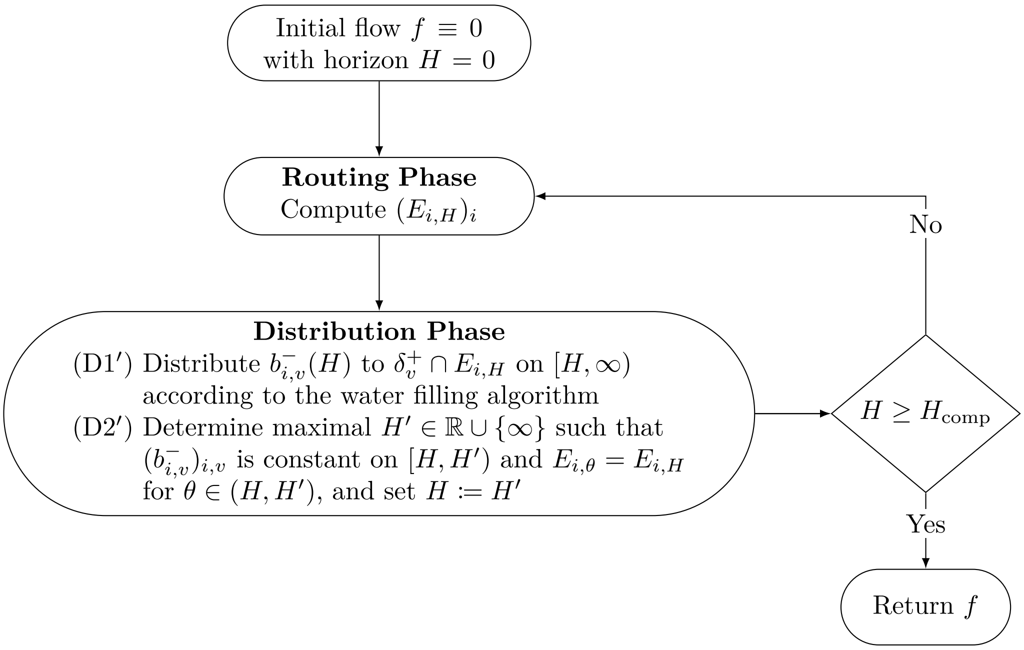

Figure 5.2: A schematic overview of the computation of an exact IDE flow.

5.2 Outflow Rates of Piecewise Constant Inflow Rates

To compute approximate dynamic prediction flows, we first consider a few theoretical observations involving the computation of general, deterministic, dynamic flows. The idea of the introduced algorithm is to utilize the simple structure of dynamic flows generated by piecewise constant inflow rates. This simple structure is demonstrated by the following theorem.

Theorem 5.2.1. *Let f be a dynamic flow that is deterministic on an edge $e \in E$, let $(g_{i,e})_{i \in I} \in \mathbb{R}^I_{\geq 0}$ be a set of new constant inflow rates into e beginning from time $H \in \mathbb{R}$, and let $g_e \coloneqq \sum_{i \in I} g_{i,e}$. The deterministic outflow rates corresponding to the inflow rates*

$$h^+_{i,e}(\theta) \coloneqq \begin{cases} f^+_{i,e}(\theta), & \text{for } \theta < H, \\ g_{i,e}, & \text{for } \theta \geq H, \end{cases}$$

are denoted by $(h^-_{i,e})_i$. Then $h^-_{i,e}$ is given by $h^-_{i,e}|_{(-\infty, T^f_e(H))} \overset{\text{a.e.}}{=} f^-_{i,e}|_{(-\infty, T^f_e(H))}$ and by the following three cases:

Case I: $g_e = 0$. *It holds* $h^-_{i,e}|_{[T^f_e(H), \infty)} \overset{\text{a.e.}}{\equiv} 0$.

Case II: $g_e > 0 \wedge \left(q^f_e(H) = 0 \vee g_e \geq \nu_e\right)$. *It holds* $h^-_{i,e}|_{[T^f_e(H), \infty)} \overset{\text{a.e.}}{\equiv} \min\{\nu_e, g_e\} \cdot \frac{g_{i,e}}{g_e}$.

Case III: $g_e \in (0, \nu_e) \wedge q^f_e(H) > 0$. *For* $T_{\text{depl}} \coloneqq H + \frac{q^f_e(H)}{\nu_e - g_e}$, *it holds that*

$$h^-_{i,e}|_{[T^f_e(H), T_{\text{depl}} + \tau_e)} \overset{\text{a.e.}}{\equiv} \nu_e \cdot \frac{g_{i,e}}{g_e} \quad \text{and} \quad h^-_{i,e}|_{[T_{\text{depl}} + \tau_e, \infty)} \overset{\text{a.e.}}{\equiv} g_{i,e}.$$

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 $T_e(H)$ as per Lemma 3.4.8.

In the following, we use a simple observation on the queue length for any $\theta > T_e(H)$:

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

Case I. From $g_e = 0$ it follows that

$$q_e^h(\theta - \tau_e) \leq q_e(H) - \int_{H+\tau_e}^{T_e(H)} h_e^- \,\mathrm{d}\lambda = 0,$$

for all $\theta > T_e(H)$. Here, the last equality follows from Proposition 3.3.2 and (F2). By using (F2) again, this implies $h_{i,e}^-(\theta) = 0$ for almost all $\theta > T_e(H)$.

Case II. 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 (F3) leads to the conclusion $h_{i,e}^-(\theta) = \min\{\nu_e, g_e\} \cdot \frac{g_{i,e}}{g_e}$ for almost all $\theta > T_e(\theta)$.

Let us first analyze the case $g_e > \nu_e$: For any $\theta > T_e(H)$ the queue length fulfills

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

and therefore $h_e^-(\theta) = \nu_e$ holds for almost all $\theta > T_e(H)$. For $g_e = \nu_e$, the same conclusion holds true.

Now assume $g_e \in (0, \nu_e)$ and $q_e(H) = 0$. By Corollary 3.4.2, for any $\theta > T_e(H)$ we have

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

Applying (F3) yields $h_e^-(\theta) = g_e$ for almost all $\theta > T_e(H)$.

Case III. Similar to the case above, it suffices to prove both $h_e^-|_{(T_e(H), T_{\mathrm{depl}} + \tau_e)} \overset{a.e.}{\equiv} \nu_e$ and $h_e^-|_{(T_{\mathrm{depl}} + \tau_e, \infty)} \overset{a.e.}{\equiv} g_e$. For $\theta \in (T_e(H), T_{\mathrm{depl}} + \tau_e)$ 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. This shows that T_{depl} is in fact chosen such that $q_e^h(T_{\mathrm{depl}}) = q_e^h(H) - (\nu_e - g_e) \cdot (T_{\mathrm{depl}} - H) = 0$. For $\theta > T_{\mathrm{depl}} + \tau_e$ we simply compute

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

implying $h_e^-(\theta) = g_e$. □

Remark 5.2.2. The proof describes not only the structure of the outflow rates beginning at time $T_e(H)$ but also the corresponding evolution of the queue length starting at time H. This means, in the setting of Theorem 5.2.1, the queue q_e^h is given for $\theta \geq H$ as follows.

Cases I and III. Denoting the depletion time by $T_{\text{depl}} := H + q_e^f(H)/(\nu_e - g_e)$, we have

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

Case II. No queue depletion occurs and for all $\theta \geq H$ 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 piecewise constant functions $f_{i,e}^+$ and $f_{i,e}^-$ for all $i \in I$ and $e \in E$, as well as piecewise 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 5.2.1 and Remark 5.2.2 starting from times $T_e(H)$ and H, respectively, while preserving the deterministic property of f.

To quickly 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 already existing events, that are rendered invalid by the update. As we require $\tau_e > 0$ for all $e \in E$, the added events happen at some time $\theta > H$. 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 good 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, because this implies that all $b_{i,v}^-$ stay constant with value $b_{i,v}^-(H)$ for an even longer period than up to time θ. Otherwise, $b_{i,v}^-$ changes at time $\theta = H'$. 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.

Definition 5.3.1. A predictor $\hat{q}_i$ is *computable as piecewise linear functions*, if the function $\hat{q}_{i,e}(\,\cdot\,, \bar{\theta}, f)$ is piecewise linear and there exists an algorithm computing this function for every $e \in E$, $\bar{\theta} \in \mathbb{R}$ and for every deterministic flow f with piecewise constant inflow rates.

Theorem 5.3.2. *If all predictors are FIFO-compatible, oblivious and computable as piecewise linear functions and all τ_e are strictly positive, then the algorithm in Section 5.1 returns an ε-DPE up to time $H_{\text{comp}} \in \mathbb{R}$ in finite time.*

Proof. If the algorithm terminates, the returned flow is both deterministic (up to time ∞) and feasible up to time H_{comp}. The fact that at any time θ we extend the flow along edges in $\hat{E}_i(\vartheta_\varepsilon(\theta), \vartheta_\varepsilon(\theta), f)$ implies that $(\hat{q}, f)$ is an ε-DPE up to time H_{comp}.

To acknowledge that the algorithm terminates, we require that all predictors are computable as piecewise linear functions. This enables us to use the algorithms in Chapter 2 to calculate active edges. Now, 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 later than (or at) time $\theta + \tau_{\min}$ with $\tau_{\min} \coloneqq \min_{e\in E} \tau_e > 0$. Assume, we have computed a flow up to time H. Then the number of 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 is enough to assess that the algorithm terminates.

For a more thorough analysis, let $\mathrm{d}^+_{\max} \coloneqq \max_{v\in V} \lvert\delta^+_v\rvert$ denote the maximum degree of outgoing edges over all nodes and let P_i denote the number of jumps of u_i. We write $P \coloneqq \sum_{i\in I} P_i$. Then, during the interval $[0, \tau_{\min})$, at most $k_0 \coloneqq P$ events can happen. A single event that occurs at time θ can cause at most $2\mathrm{d}^+_{\max} \cdot |I|$ events all of which happen later than $\theta + \tau_{\min}$. Thus, the number of events that occur within the interval $[\tau_{\min}, 2\tau_{\min})$ is bounded from above by $k_1 \coloneqq k_0 \cdot (1 + 2\mathrm{d}^+_{\max} \cdot |I|)$. Similarly, the number of events that happen in the interval $[l\tau_{\min}, (l+1)\tau_{\min})$ is bounded by $k_l \coloneqq k_0 \cdot (1 + 2\mathrm{d}^+_{\max} \cdot |I|)^l$. Utilizing the geometric sum, the number of events that occur up to time $\lceil H_{\text{comp}}/\tau_{\min}\rceil \cdot \tau_{\min} \geq H_{\text{comp}}$ is bounded by

$$\sum_{l=0}^{\left\lceil \frac{H_{\text{comp}}}{\tau_{\min}} \right\rceil} k_l = P \cdot \frac{(1 + 2\mathrm{d}^+_{\max} \cdot |I|)^{\left\lceil \frac{H_{\text{comp}}}{\tau_{\min}} \right\rceil + 1} - 1}{2\mathrm{d}^+_{\max} \cdot |I|},$$

and thus the number of distribution phases is $\mathcal{O}\Big(P \cdot (1 + 2\mathrm{d}^+_{\max} \cdot |I|)^{\frac{H_{\text{comp}}}{\tau_{\min}} + 2}\Big)$. □

5.4 Experimental Study

In this last section we conduct several experiments to demonstrate the capabilities of the proposed algorithm. Most of these results were also described in [13, Section 5]. In the following experiments, we compare the performance of the different predictors in several scenarios. We measure a predictor's performance by the average travel time of particles using that predictor. For that, we use simple network inflow rates: For a commodity i we set $u_i \coloneqq \mathbf{1}_{[0,h)} \cdot \bar{u}_i$ for some constant $\bar{u}_i \in \mathbb{R}_{>0}$ and some (common) inflow horizon $h \in \mathbb{R}_{>0}$.

The outflow rate of a commodity i is denoted as $o_i(\theta) \coloneqq \sum_{e\in\delta^-_{t_i}} f^-_{i,e}(\theta) - \sum_{e\in\delta^+_{t_i}} f^+_{i,e}(\theta)$. Integrating $u_i(\psi) - o_i(\psi)$ over $[0, \phi]$ yields the flow of commodity i that is inside the network at time ϕ. Taking the integral of this quantity over some time period $[0, H_{\text{comp}}]$ with $H_{\text{comp}} \geq h$ gives the *total travel time* of particles of commodity i up to time H_{comp}:

$$T_i^{\text{total}} \coloneqq \int_0^{H_{\text{comp}}} \int_0^{\phi} u_i(\psi) - o_i(\psi) \,\mathrm{d}\psi \,\mathrm{d}\phi.$$

The *average travel time* is defined as $T_i^{\text{avg}} := T_i^{\text{total}}/(h \cdot \bar{u}_i)$.

A particle that follows a predictor blindly has a certain *regret* in comparison to an optimal choice. This regret is defined in the following. For example, a particle using the perfect predictor can choose a path of minimal travel time and thus does not regret its routing decisions. The *average minimum travel time* (resulting from an optimal route choice) of particles up to some time horizon H_{comp} with source-sink-pair (s_i, t_i) can be determined using

$$T_{i,\text{OPT}}^{\text{avg}} := \frac{1}{h} \cdot \int_0^h \min\{H_{\text{comp}}, l_{s_i,t_i}(\theta)\} - \theta \,\mathrm{d}\theta.$$

In this expression, the term $\min\{H_{\text{comp}}, l_{s_i,t_i}(\theta)\} - \theta$ describes the minimum travel time of particles spawning at time θ at the source s_i, limited at the time horizon H_{comp}. Taking the integral of this quantity over $[0, h]$ and dividing by h yields the average minimum travel time. In the conducted experiments, we then measure the *regret* of a commodity i by $T_i^{\text{avg}}/T_{i,\text{OPT}}^{\text{avg}}$. Here, the Dynamic Bellman-Ford algorithm (Section 2.6) is used to determine the piecewise linear function $l_{s_i,t_i}(\,\cdot\,)$ for the dynamic cost function induced by the computed queue lengths.

5.4.1 Data

We consider three different types of networks in our experiments. The first network is a small, synthetic network which was studied in [14]. Its graph is depicted in Figure 5.3.

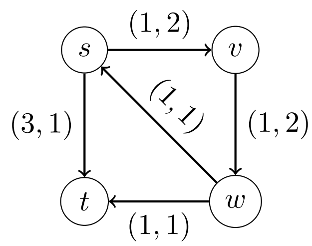

Figure 5.3: A network with source s and sink t. Edges are labeled with (τ_e, ν_e).

Additionally, we use data from two real-world traffic networks. One of them is the Sioux Falls road network as presented in [18] which is a popular example for analyzing models in the transportation science literature. With only 24 nodes, this network is still comparatively small. The publicly available data set contains the (free-flow) travel time and capacity rate of each road segment.

Open Street Maps [21] offers the street networks of almost any region in the world. From this data set, the street network of the center of Tokyo was extracted. It offers the free-flow speed, the length and the number of lanes of each road segment. Here, the transit time τ_e is derived as the product of the length and the free-flow speed of a road segment; the capacity ν_e is computed as the product of the number of lanes and the free-flow speed.

More in-depth information on these networks is displayed in Table 5.1. Here, $T_{\text{comp}}^{\text{avg}}$ is the average computation time for computing an ε-DPE flow up to time H_{comp}. In these calculations, all predictors are used by exactly one commodity, except the constant predictor, which is used by all remaining commodities. The complete experimental study was conducted on a single core of an Intel® Core™ i7-3520M CPU at 2.90GHz.

Network	Synthetic	Sioux Falls	Tokyo
$\lvert E \rvert$	5	75	4,803
$\lvert V \rvert$	4	24	3,538
$\lvert I \rvert$	5	17	40
$[\nu_{\min}, \nu_{\max}]$	$[1, 2]$	$[4823, 25901]$	$[8, 250]$
$[\tau_{\min}, \tau_{\max}]$	$[1, 3]$	$[2, 10]$	$[0.01, 6.6]$
H_{comp}	500	100	100
h	25	25	25
ε	0.25	1	2.5
H for $\hat{q}^{\text{L}}$	10	20	20
δ, H for $\hat{q}^{\text{RL}}$	$5, 10$	$1, 20$	$1, 20$
δ, k_p, k_f for $\hat{q}^{\text{ML}}$	$1, 10, 10$	$1, 20, 20$	$1, 20, 20$
ML model for $\hat{q}^{\text{ML}}$	single	per-edge	single
$T_{\text{comp}}^{\text{avg}}$	$1.65s$	$10.92s$	$343.93s$

Table 5.1: Attributes and configuration of the considered networks.

5.4.2 The Linear Regression Predictor

Before discussing the results of the experimental study, we make a few remarks on the implementation of the linear regression predictor. The other (oblivious) predictors discussed in Section 4.6 are rather easy to implement.

During the experimental study, numerous machine-learning frameworks were considered. These include the *Tensorflow* framework [1], the *Weka* tool [11] and the *Deep Graph Library* [27] together with the *PyTorch* framework [22]. However, more detailed experiments were conducted using the *scikit-learn* framework [23], which supports efficient implementations of linear regression methods.

The linear regression predictor was introduced in Section 4.6. Here, we want to learn coefficient matrices $W^{e'} \in \mathbb{R}^{k_p \times k_f}$ for all edges $e = vw$ and neighboring edges $e' \in N(e)$ as well as biases $\beta \in \mathbb{R}^{k_f}$. Using these values, our raw machine-learned prediction method outputs the interpolation points

$$\hat{q}_{i,e}^{\text{ML,raw}}\left(\bar{\theta} + j \cdot \delta, \bar{\theta}, f\right) := \left(\sum_{e' \in N(e)} \sum_{i \in [k_p]} w_{i,j}^{e'} \cdot q_{e'}^{f}\left(\bar{\theta} - (i-1) \cdot \delta\right) + \beta_j \right)^{+}$$

for all $j \in [k_f]$. Here, we choose $N(e) := \delta_v^+ \cup \{e\} \cup \delta_v^-$.

However, if the number of edges becomes too large – as is the case for the Tokyo instance – we restrict ourselves in learning only a single model for all edges. This means, we only learn a fixed set of coefficients and apply them to all edges $e \in E$. More specifically, let $\mathrm{d}_{\max}^- := \max_{v \in V} \lvert \delta_v^- \rvert$ be the maximum in-degree of any node. We learn matrices $W^{+1}, \dots, W^{+\mathrm{d}_{\max}^+}$, W^0, $W^{-1}, \dots, W^{-\mathrm{d}_{\max}^-}$ and biases $\beta \in \mathbb{R}^{k_f}$. These generalized parameters can now be used to predict the queue lengths of any edge $e = vw$: We (arbitrarily) order the neighboring edges using $\delta_v^- =: \{e^{+1}, \ldots, e^{+|\delta_v^-|}\}$, $e =: e^0$ and $\delta_w^+ =: \{e^{-1}, \ldots, e^{-|\delta_w^+|}\}$ and compute

$$\hat{q}_e^{\mathrm{ML,raw}}\left(\bar{\theta} + j \cdot \delta, \bar{\theta}, f\right) := \left(\sum_{e^l \in N(e)} \sum_{i \in [k_p]} w_{i,j}^l \cdot q_{e^l}^f\left(\bar{\theta} - (i-1) \cdot \delta\right) + \beta_j\right)^+ .$$

The training data fed into the learning procedure stems from previously computed approximated dynamic equilibrium flows: For this, we pick random commodities and network inflow rates and equip every commodity with the constant predictor $\hat{q}_{i,e}^{\mathrm{C}}$. This has two advantages: Firstly, calculating shortest paths when using the constant predictor is efficient. Secondly, the generated dynamic flows allow the machine-learning method to grasp the behavioral model of particles adapting their route choices depending on the queue lengths.

5.4.3 Comparison of Predictors

We have now accomplished all necessary prerequisites for analyzing the results. For each experiment we first describe its exact setup and then inspect the results.

Results for the Synthetic Network

First, we consider the synthetic network as given in Figure 5.3. We set up five commodities, one for each predictor in $\{\hat{q}^{\mathrm{Z}}, \hat{q}^{\mathrm{C}}, \hat{q}^{\mathrm{L}}, \hat{q}^{\mathrm{RL}}, \hat{q}^{\mathrm{ML}}\}$. All commodities have the same network inflow rate $u_i := \mathbf{1}_{[1,h]} \cdot \bar{u}$ for some $\bar{u} \in \mathbb{R}_{>0}$ and $h = 25$. Moreover, every commodity has the same source s and the same sink t. We compare the outcome for varying $\bar{u}$: For each $\bar{u} \in \{0.25, 0.5, \ldots, 5.75, 6\}$ we compute an ε-DPE flow up to time $H_{\mathrm{comp}} = 500$. Furthermore, the rerouting interval is set to $\varepsilon = 0.25$. We note that in this case, the linear regression predictor consists of a single model which was learned on a set of dynamic flows generated in the Tokyo network.

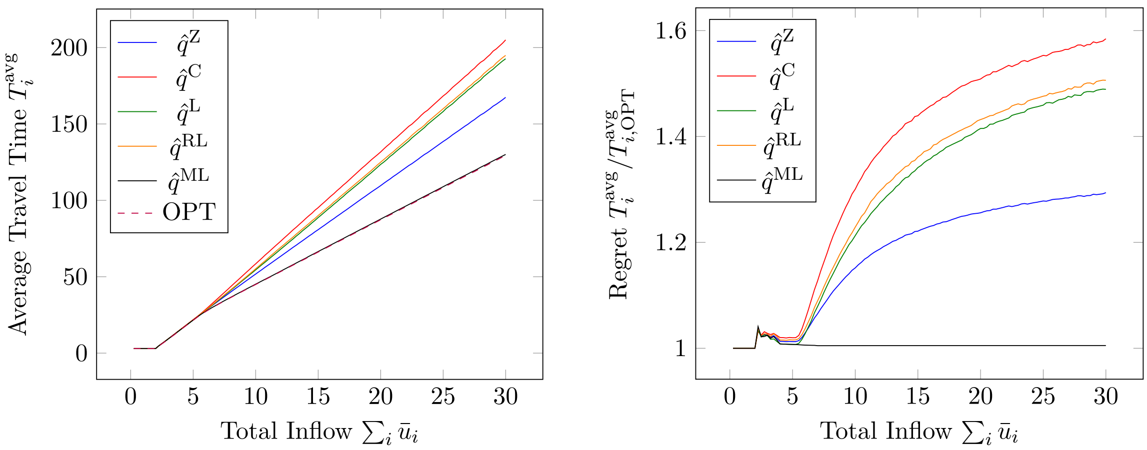

Figure 5.4: Measured average travel times and regrets of competing predictors in the synthetic network in Figure 5.3.

The resulting average travel times of each predictor in each of these runs can be seen in Figure 5.4. It is easy to derive a ranking of the predictors based on the travel times: The linear regression predictor performed best – almost optimally. Notably, the Zero-Predictor, which distributes flow along the paths (s, t) and (s, v, w, t) uniformly at all times, performs better than the remaining predictors. The constant predictor performs worst in this network. This is no surprise, as in [14], this network was deliberately chosen to show that IDE flows admit cyclic behavior. The linear predictor reduces this behavior significantly, as it detects earlier that queues build up at certain edges. The regularized linear predictor dampens this quick recognition of growing queues slightly, but still behaves similar to the linear predictor. It is noteworthy that the performance of the predictors starts deviating at $\sum_i \bar{u}_i = 6$. For an explanation of this deviation, further analysis is required.

Remark 5.4.1. In contrast to [13], the computation horizon H_{comp} was increased from 100 to 500 to ensure that no flow remains in the network after terminating the simulation. This makes the resulting figures much easier to interpret.

Results for the Sioux Falls Network

For the real-world network of Sioux Falls, a set I of 12 commodities is randomly generated together with network inflow rates according to the capacity rates of the edges. We equip all of these commodities with the constant predictor $\hat{q}^{\mathrm{C}}$. For all $i \in I$ we compute an ε-DPE up to time $H_{\text{comp}} = 100$, while calculating new routes every $\varepsilon = 1$ time units.

In the computation of the ε-DPE of commodity i, we generate five more commodities with the common source s_i and sink t_i, and we assign each a different predictor in $\{\hat{q}^{\mathrm{Z}}, \hat{q}^{\mathrm{C}}, \hat{q}^{\mathrm{L}}, \hat{q}^{\mathrm{RL}}, \hat{q}^{\mathrm{ML}}\}$ together with a network inflow rate of $\mathbf{1}_{[0,h]} \cdot 0.125$. This network inflow rate is negligible compared to the network inflow rates of all other commodities, and thus the resulting ε-DPE flows behave just like the flows in the training data, where the constant predictor is used exclusively. For each of these runs, we monitor the regret of the added commodities.

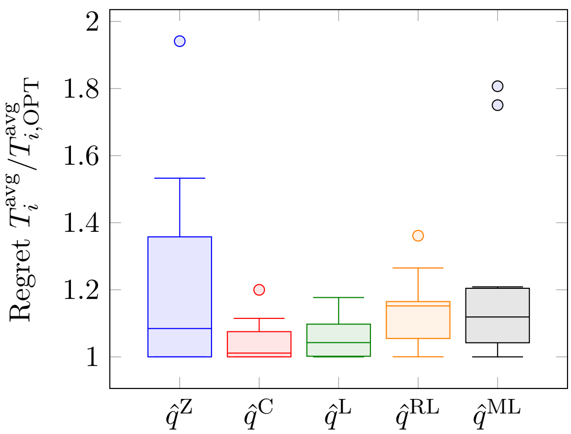

Figure 5.5: Average travel times compared to the minimum average travel time in the Sioux Falls network.

The number of edges of the Sioux Falls network is small enough to learn separate coefficients for each edge. This was done using a 90%/10% split for the training data and test data. Here, the coefficient of determination was above 0.9 for all edges except for six edges but always higher than 0.5. The training data stems from previously computed ε-DPE with randomly generated commodities and network inflow rates as above.

In Figure 5.5 the regret of each predictor is shown in a box plot. As expected, the Zero-predictor performs worst. However, for the remaining predictors there is not a clear winner.

Results for the Tokyo Network

For the Tokyo instance, the same setup was applied as for the Sioux Falls network. However, as the Tokyo network has significantly more edges, it is infeasible to learn a model for each edge separately. Hence, the approach using a single model as discussed in Section 5.4.2 was used. Again, a training and validation split of 90%/10% was applied. Here, we obtained a coefficient of determination of 0.97.

Computing an ε-DPE flow for every of the 35 randomly chosen commodities results in the regrets as shown in the box plot in Figure 5.6. The results are again rather unsatisfying: The Zero-predictor can be identified as the worst predictor, however the performance of the other predictors is mostly indistinguishable.

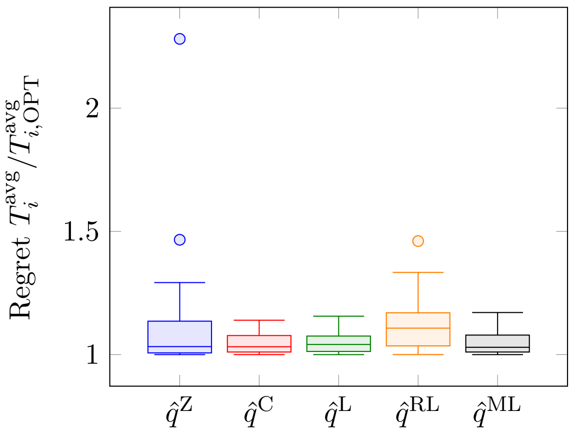

Figure 5.6: Average travel times compared to the minimum average travel time in the Tokyo network.

Concluding Thoughts

On the one hand, the presented results show that the proposed algorithm is capable of computing ε-DPE flows in large-scale real-world traffic networks. On the other hand, there is still plenty of room for analyzing why the linear regression predictor did not perform as well as expected in larger networks. In the following, a few starting points for further research are listed.

I. Add more inputs to the linear regression predictor.

For example, past samples of the edge load $F_{e'}^{+}(\theta) - F_{e'}^{-}(\theta)$ of surrounding edges $e' \in N(e)$ could be added. Whenever a queue occurs, this is due to a lack of capacity on that edge. However, if there was no queue on an incident edge, the predictor has no chance to detect an emerging queue.

II. Increase the computational horizon such that no more flow remains in the network after termination.

This is a relatively simple tweak to the simulation parameters, however, with a possibly large effect, as hinted by Remark 5.4.1.

III. Use manually selected commodities with interesting source and sink pairs.

Often times, randomly selected commodities do not offer an interesting route choice. For example, if any flow starting at the source has to go through a low-capacity edge right at the beginning, any predictor will most likely choose the same path thereafter, because all the flow is stuck at the single bottleneck edge. Using commodities of this type in the simulation makes it harder to compare the performance of the different predictors.

IV. Consider a completely different approach for a machine-learned predictor.

There are plenty of methods for forecasting traffic. Modern routing services often build upon so-called *Graph Neural Networks (GNNs)* where the concepts of deep neural networks are applied to data structured in graphs. A survey on GNNs for traffic forecasts is presented in [16]. It is reasonable to embed one of these modern methods as a new predictor into the computation.

6 Conclusion

In this work, we generalized the equilibrium model introduced in [13] to describe the use of modern routing services. These services are capable of using predictions of the future congestion based on historic and real-time data.

As a first step general FIFO-ordered cost functions were analyzed. Here, the reversal of functions and the duality of earliest arrival and latest departure times led not only to a characterization of active edges, but also turned out to be a convenient notation for proving the equivalence of dynamic Nash equilibrium flows and dynamic prediction flows with respect to the perfect predictor. The characterization of active edges was used to calculate the set of active outgoing edges of a node v using a variant of Dijkstra's algorithm. Moreover, using the Dynamic Bellman-Ford Algorithm it is possible to calculate earliest arrival times as functions over time. This was used in the experimental study to calculate the average minimum travel time to compare the performance of a predictor with an optimal choice.

Before focusing on dynamic prediction equilibria, we introduced Vickrey's Fluid Queuing Model which establishes the physical constraints of the particles' behavior. We formally proved the unique existence of outflow rates given any set of inflow rates such that the resulting flow is deterministic. This is a claim that was often stated in the past literature for which a formal proof has not been published. Conducting the proof required several measure-theoretical results such as Sard's theorem which itself relies on Vitali's Covering Lemma.

Chapter 4 finally introduced dynamic prediction equilibria in its most general form, in which predictors can use the historical flow f in their forecast. In [13], the existence of DPE where oblivious, FIFO-compatible predictors that depend continuously on the queue length functions was proven using the solution of a variational inequality. In this thesis, this result was derived as a special case of the existence theorem that requires a more abstract regularity condition on the predictors in the form of p-continuity. However, there still remain several open questions for the existence of DPE: Do dynamic prediction equilibria still exist if we allow $\tau_e = 0$ for some edges $e \in E$? Can we drop the obliviousness requirement, in the case of network inflow rates with bounded support? Do equilibria still exist if we consider network inflow rates in $L^1_{\text{loc}}(\mathbb{R}, \mathbb{R}_{\geq 0})$?

Besides these general statements we analyzed several predictors to apply the developed theory to. An interesting example is the linear predictor which is incompatible with the existence theorem because of its irregularity at points in which the gradient of the queue length jumps. However, an example network for which the linear predictor does not admit a dynamic prediction flow has not yet been found. For single-sink networks it might even be possible to formulate a procedure for the computation of dynamic prediction flows with respect to the linear predictor similar to the natural extension algorithm proposed in [12] for computing IDE flows.

Furthermore, we proved that dynamic prediction equilibria generalize the popular dynamic Nash equilibrium flows and instantaneous dynamic equilibrium flows by considering dynamic prediction flows with respect to the perfect and the constant predictor, respectively. To show the equivalence to dynamic Nash equilibrium flows, the change in perspective yielded a novel characterization of these equilibria based on the latest departure times.

The last chapter discussed the computation of dynamic prediction equilibria. More precisely, the concept of approximated dynamic prediction equilibria was introduced. This type of equilibrium describes a state in which each agent computes new routes at certain discrete times only. For a set of oblivious, FIFO-compatible predictors that are computable as piecewise linear functions, we introduced an algorithm that computes an ε-DPE flow up to some time horizon H_{comp} and proved its correctness and termination. Finally, an experimental study was conducted on three different networks, two of which are real-world road networks. In the study, the performance of the predictors was compared with respect to the average travel time of particles.

The results show that the proposed simulation is capable of computing ε-DPE in large-scale traffic networks and that the linear regression predictor performs well for the synthetic network. However, future analyses could elaborate on how to improve the linear regression predictor in larger networks. Moreover, other means of current traffic forecast methods could be embedded in the simulation, and it could be investigated whether they render compatible with the proposed existence theorem. Furthermore, a study of the approximation guarantees of an ε-DPE is left for future research.

Bibliography

[1] Martín Abadi et al. "TensorFlow: A System for Large-Scale Machine Learning". In: *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation*. OSDI'16. Savannah, GA, USA: USENIX Association, 2016, pp. 265–283. ISBN: 9781931971331. DOI: 10.5555/3026877.3026899.

[2] Hans Wilhelm Alt. *Lineare Funktionalanalysis*. Springer Berlin Heidelberg, 2012. DOI: 10.1007/978-3-642-22261-0.

[3] Krishna B. Athreya and Soumendra N. Lahiri. *Measure Theory and Probability Theory*. Springer New York, 2006. DOI: 10.1007/978-0-387-35434-7.

[4] Haïm Brézis. "Équations et inéquations non linéaires dans les espaces vectoriels en dualité". French. In: *Annales de l'Institut Fourier* 18.1 (1968), pp. 115–175. DOI: 10.5802/aif.280.

[5] Theodor Bröcker and Klaus Jänich. *Einführung in die Differentialtopologie - Korrigierter Nachdruck*. Wiesbaden: Springer Berlin Heidelberg, 1990. ISBN: 978-3-540-06461-9.

[6] Felix E. Browder. "Nonlinear elliptic boundary value problems". In: *Bulletin of the American Mathematical Society* 69.6 (1963), pp. 862–874. DOI: 10.1090/S0002-9904-1963-11068-X.

[7] Roberto Cominetti, José Correa, and Omar Larré. "Dynamic Equilibria in Fluid Queueing Networks". In: *Operations Research* 63.1 (2015), pp. 21–34. DOI: 10.1287/opre.2015.1348.

[8] Brian Dean. *Shortest Paths in FIFO Time-Dependent Networks: Theory and Algorithms*. Tech. rep. Massachusetts Institute of Technology, Jan. 2004. URL: https://people.cs.clemson.edu/~bcdean/tdsp.pdf.

[9] 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. Nantes, France: Association for Computing Machinery, 2008, pp. 205–216. DOI: 10.1145/1353343.1353371.

[10] Luca Foschini, John Hershberger, and Subhash Suri. "On the Complexity of Time-Dependent Shortest Paths". In: *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*. SODA '11. San Francisco, California: Society for Industrial and Applied Mathematics, 2011, pp. 327–341. DOI: 10.5555/2133036.2133063.

[11] E. Frank et al. "Weka: A machine learning workbench for data mining." In: *Data Mining and Knowledge Discovery Handbook: A Complete Guide for Practitioners and Researchers*. Ed. by O. Maimon and L. Rokach. Berlin: Springer, 2005, pp. 1305–1314. URL: `http://researchcommons.waikato.ac.nz/handle/10289/1497`.

[12] Lukas Graf and Tobias Harks. "A finite time combinatorial algorithm for instantaneous dynamic equilibrium flows". In: *Mathematical Programming* (Feb. 2022). DOI: `10.1007/s10107-022-01772-0`.

[13] Lukas Graf, Tobias Harks, Kostas Kollias, and Michael Markl. *Machine-Learned Prediction Equilibrium for Dynamic Traffic Assignment*. 2021. arXiv: `2109.06713 [cs.GT]`.

[14] Lukas Graf, Tobias Harks, and Leon Sering. "Dynamic flows with adaptive route choice". In: *Mathematical Programming* 183.1-2 (May 2020), pp. 309–335. DOI: `10.1007/s10107-020-01504-2`.

[15] Lions Jacques-Louis. *Quelques méthodes de résolution des problèmes aux limites non linéaires*. French. Études Mathématiques. Paris: Dunod, 1969.

[16] Weiwei Jiang and Jiayun Luo. "Graph Neural Network for Traffic Forecasting: A Survey". In: *CoRR* abs/2101.11174 (2021). arXiv: `2101.11174`.

[17] Ronald Koch and Martin Skutella. "Nash Equilibria and the Price of Anarchy for Flows over Time". In: *Theory of Computing Systems* 49.1 (Nov. 2010), pp. 71–97. DOI: `10.1007/s00224-010-9299-y`.

[18] Larry J. LeBlanc, Edward K. Morlok, and William P. Pierskalla. "An efficient approach to solving the road network equilibrium traffic assignment problem". In: *Transportation Research* 9.5 (1975), pp. 309–318. ISSN: 0041-1647. DOI: `10.1016/0041-1647(75)90030-1`.

[19] Michael Markl. *Berechnung von Nash-Gleichgewichten in dynamischen Flüssen*. Bachelor Thesis. Jan. 2020. URL: `https://michael-markl.de/dynamic-equilibrium-flows/bachelorarbeit.pdf`.

[20] George J. Minty. "On a "Monotonicity" Method for the Solution of Nonlinear Equations in Banach Spaces". In: *Proceedings of the National Academy of Sciences of the United States of America* 50.6 (1963), pp. 1038–1041. ISSN: 00278424. URL: `http://www.jstor.org/stable/71840`.

[21] OpenStreetMap contributors. *Planet dump retrieved from https://planet.osm.org*. `https://www.openstreetmap.org`. 2017.

[22] Adam Paszke et al. "PyTorch: An Imperative Style, High-Performance Deep Learning Library". In: *Advances in Neural Information Processing Systems 32*. Ed. by H. Wallach et al. Curran Associates, Inc., 2019, pp. 8024–8035. URL: `http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf`.

[23] Fabian Pedregosa et al. "Scikit-learn: Machine Learning in Python". In: *Journal of Machine Learning Research* 12.85 (2011), pp. 2825–2830. URL: `http://jmlr.org/papers/v12/pedregosa11a.html`.

[24] Halsey L. Royden and Patrick M. Fitzpatrick. *Real Analysis*. 4th ed. Prentice Hall, 2010. ISBN: 9780131437470.

[25] Arthur Sard. "The measure of the critical values of differentiable maps". In: *Bulletin of the American Mathematical Society* 48.12 (1942), pp. 883–890. DOI: 10.1090/S0002-9904-1942-07811-6.

[26] William S. Vickrey. "Congestion Theory and Transport Investment". In: *The American Economic Review* 59.2 (1969), pp. 251–260. ISSN: 00028282. URL: http://www.jstor.org/stable/1823678.

[27] Minjie Wang et al. "Deep Graph Library: Towards Efficient and Scalable Deep Learning on Graphs". In: *CoRR* abs/1909.01315 (2019). arXiv: 1909.01315.