Convergence of a Packet Routing Model to Flows Over Time

LEON SERING, Technische Universität Berlin, Germany
LAURA VARGAS KOCH, RWTH Aachen University, Germany
THERESA ZIEMKE, Technische Universität Berlin, Germany

The mathematical approaches for modeling dynamic traffic can roughly be divided into two categories: discrete *packet routing models* and continuous *flow over time models*. Despite very vital research activities on models in both categories, the connection between these approaches was poorly understood so far. In this work we build this connection by specifying a (competitive) packet routing model, which is discrete in terms of flow and time, and by proving its convergence to the intensively studied model of flows over time with deterministic queuing. More precisely, we prove that the limit of the convergence process, when decreasing the packet size and time step length in the packet routing model, constitutes a flow over time with multiple commodities. In addition, we show that the convergence result implies the existence of approximate equilibria in the competitive version of the packet routing model. This is of significant interest as exact pure Nash equilibria, similar to almost all other competitive models, cannot be guaranteed in the multi-commodity setting.

Moreover, the introduced *packet routing model with deterministic queuing* is very application-oriented as it is based on the network loading module of the agent-based transport simulation MATSim. As the present work is the first mathematical formalization of this simulation, it provides a theoretical foundation and an environment for provable mathematical statements for MATSim.

CCS Concepts: • **Theory of computation** → **Algorithmic game theory**; **Network games**; **Network flows**; *Exact and approximate computation of equilibria*; *Routing and network design problems*; • **Mathematics of computing** → **Network flows**.

Additional Key Words and Phrases: flow over time, packet routing, multi-commodity, dynamic equilibrium, convergence of discrete to continuous, discretization error, approximate equilibrium, dynamic traffic assignment

ACM Reference Format:
Leon Sering, Laura Vargas Koch, and Theresa Ziemke. 2021. Convergence of a Packet Routing Model to Flows Over Time. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC '21), July 18–23, 2021, Budapest, Hungary.* ACM, New York, NY, USA, 20 pages. https://doi.org/10.1145/3465456.3467626

Related Version: A full version of this paper including all proofs is available at http://arxiv.org/abs/2105.13202.

1 INTRODUCTION

Modeling traffic is an essential, but difficult task for which many different approaches have been developed during the last decades. Some, mainly older approaches model traffic as time-independent *static* flows [21, 22, 28]. These models are relatively simple and well-understood, e.g. regarding existence, uniqueness, efficiency, and structure of equilibria, but they can only model aggregated traffic flows. In contrast to that, there are time-dependent *dynamic* approaches, which lead to more realistic, but also more complex models. There exists various mathematical approaches as

EC '21, July 18–23, 2021, Budapest, Hungary

ACM ISBN 978-1-4503-8554-1/21/07…$15.00
https://doi.org/10.1145/3465456.3467626

well as simulation tools in this category. Some consider discrete time steps [11, 13, 23, 29], while others model continuous time [12, 17, 24]. Some model traffic as discrete travelers (vehicles/packets) [11–13, 23, 24, 29], while others model it as continuous splittable flows [9, 10, 17]. In this paper we consider two of these approaches: on the one hand, the intensively studied flow over time model with deterministic queuing [17], which considers continuous time and continuous flow, and on the other hand, a packet routing model considering discrete travelers and discrete time steps. The latter is based on the transport simulation MATSim (Multi-agent transport simulation, see [13]) and it is very similar to the model studied by Werth et al. [29].

The flows over time model yields exact user equilibria, called *dynamic equilibria* or *Nash flows over time* [17]. They are guaranteed to exist [3, 4] and their structure has been studied intensively [5, 7, 15, 16, 20, 27]. Even though most of the work has been done for single-origin-destination networks, the existence results and further structural insights have been transferred to the multi-terminal setting as well [4, 26]; see also [25] for an extended analysis on all facets of Nash flows over time. Even more general, Meunier and Wagner [18] introduce a broader class of dynamic congestion games and show the existence of dynamic equilibria.

On the application side, the large-scale, agent-based transport simulation MATSim is used to model real-world traffic scenarios for transport planning purposes. To achieve this in reasonable time, MATSim simplifies real-world traffic by discretizing time and aggregating vehicles. A co-evolutionary algorithm is used to compute approximate user equilibria. Despite the different perspectives of the two approaches, experiments by Ziemke et al. [30] indicate that Nash flows over time are the limits of the convergence processes when decreasing the vehicle size and time step length in the simulation coherently. This motivates the question whether the convergence of the flow models can be verified and proven mathematically.

1.1 Our Contribution

To answer this question, we introduce a multi-commodity packet routing model, which is basically a mathematical formalization of MATSim's network loading module. This *packet routing model with deterministic queuing* can be seen as a refinement of the existing packet model studied by Werth et al. [29].

The main contribution of this work consists of the proof that for a fixed commodity-wise flow supply the packet flow in the discrete model converges to the corresponding flow in the continuous model, when the time step length and packet size go to zero. The key idea is to identify particles with packets and to show convergence of their travel times. Note that this result focuses on the network loading aspect of the models and does not consider any game theoretical aspects yet, i.e., the routes for the commodities are fixed. Up to our knowledge, the presented work is the first to prove convergence of a discrete to a continuous model in the dynamic setting.

Additionally, we show that, as a consequence of the convergence proof, the competitive version of the packet routing model possesses approximate user equilibria. This is particularly surprising since – up to our knowledge – the existence of (approximate) pure Nash equilibria in a multi-commodity setting has not been proven for other similar competitive packet routing models. Note that *exact* multi-commodity pure Nash equilibria are not guaranteed to exist in most packet routing models; this is also the case for the model presented here.

Our results are of significant interest as they are the first step towards the transfer of established structural results on dynamic equilibria from flows over time to competitive packet routing games. On the one hand, this provides the first building block for a mathematical foundation of MATSim, on the other hand, the convergence of the network loading module of the widely used transport simulation strengthens the relevance of dynamic equilibria in the flows over time model.

1.2 Further Related Work

The connection between discrete and continuous models in dynamic transportation settings has been studied before. Otsubo and Rapoport [19] consider the Vickrey model and focus on experimental comparisons and Cantarella and Watling [1] compare discrete and continuous time models by presenting mathematical tools and applications and unify both in a single model.

In static routing games convergence has been shown very recently by Cominetti et al. [6]. They prove convergence of equilibria in routing games from two perspectives, on the one hand, when player sizes decreases, and on the other hand, when the probability with which players participate in the game decreases.

For flows over time a connection between discrete and continuous optimization problems was established by Fleischer and Tardos [8] in 1998. They transferred optimization algorithms for discrete-time flow models to a flow model with continuous time.

The competitive packet routing model presented in this paper is similar to existing packet routing models with first in first out (FIFO) rules; see [2, 14, 23, 29]. However, each model has a different focus. Werth et al. [29] consider complexity questions and the price of anarchy for several different objective functions: arrival times and bottleneck costs as players' objectives and sum of completion times as well as makespan as social cost. While Ismailli [14] analyzes the computational complexity of computing a best response and a Nash equilibirum, Scarsini et al. [23] analyze the price of anarchy of a packet routing model with periodic inflow rates. Finally, Cao et al. [2] analyze the existence of subgame perfect equilibria.

2 PACKET ROUTING MODEL WITH DETERMINISTIC QUEUING

In this section, we formally introduce our packet routing model, which is motivated by the network loading model of MATSim. Packets with individual origins and destinations travel in discrete time steps through a network, where arcs have a transit time and a capacity. After traversing an arc, packets might be delayed at bottlenecks, where they have to wait with other packets in a point queue. In order to enter the next arc, packets from different incoming arcs are merged as smoothly as possible, similar to the zipper method from real traffic. The length of a time step and the packet size are given by some discretization parameters (α, β), which are crucial to show convergence later on. Note that our model is similar to the model of Werth et al. [29] with the only difference that we consider the zipper method (with arc priorities as tie-breaking) for the merging process, while Werth et al. directly break ties according to arc priorities without using the zipper-like merging.

2.1 Preliminaries

Units. Since we consider convergence later on, we keep the time and packets size flexible. As help for orientation, we introduce some units for time and for the size of the packets. We fix two base units, namely sec for a base unit of time, and vol for the base unit of packet volume. Note, however, that the used variables are all considered to be unitless and we use these base units only for intuition.

Network. A network is given by a directed graph $G = (V, E)$, where every arc is equipped with a transit time $\tau_e > 0$ and a capacity $\nu_e > 0$. The transit time denotes how long it takes to traverse an arc, measured in secs. The capacity restricts the flow rate that is allowed to leave the arc. Hence, it is measured in flow volume per time unit, i.e., vol/sec.

Discretization. For our packet routing model we specify two discretization parameters $\alpha, \beta > 0$, one for the time step length and one for the packet sizes. In other words, one time step has a duration of α sec and each packet has a volume of β vol. Since the point in time corresponding to a time step

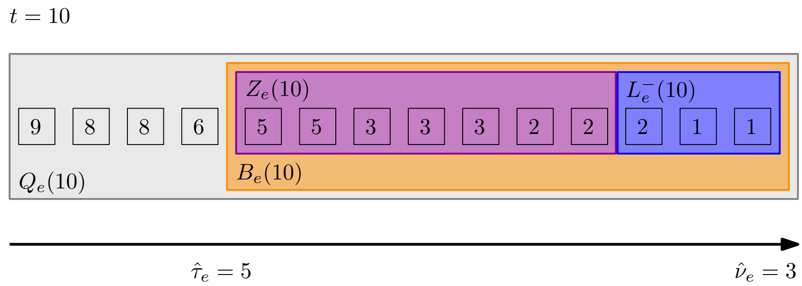

Fig. 1. Example of an arc queue Q_e at time step $t = 10$. Different parts are highlighted: Packets in the buffer B_e, leaving packets L_e^- and waiting packets Z_e. Within each packet the arc entrance time step is denoted.

$t \in \mathbb{N}_0$ is given by αt, the set of all these points in time is given by $\Theta_\alpha := \alpha\mathbb{N}_0 = \{ 0, \alpha, 2\alpha, \dots \}$. In addition, we use rounding brackets in order to round to the previous or next element of Θ_α:

$$\lfloor x \rfloor_\alpha := \max \{ \theta \in \Theta_\alpha \mid \theta \leq x \} \quad \text{and} \quad \lceil x \rceil_\alpha := \min \{ \theta \in \Theta_\alpha \mid \theta \geq x \} .$$

Analogously, this notation is also used for β, i.e.,

$$\lfloor x \rfloor_\beta := \max \{ \phi \in \beta\mathbb{N}_0 \mid \phi \leq x \} \quad \text{and} \quad \lceil x \rceil_\beta := \min \{ \phi \in \beta\mathbb{N}_0 \mid \phi \geq x \} .$$

Discretized network. Given a network and discretization parameters we normalize the arc properties as follows. The number of time steps required for a packet to traverse an arc is denoted by $\hat{\tau}_e := \frac{\lceil \tau_e \rceil_\alpha}{\alpha} \in \mathbb{N}$. The minimum number of packets that are allowed to leave an arc in each time step is given by $\hat{\nu}_e := \frac{\nu_e \cdot \alpha}{\beta}$. Note that $\hat{\nu}_e$ does not need to be an integer. The fractional remainder is transferred to the next time step, as described later.

Packets. Given a network G and discretization parameters α and β, we consider n packets of volume β with a path P_i connecting an origin-destination pair $(o_i, d_i) \in V^2$ and a release time $r_i \in \Theta_\alpha$ for all $i \in N := \{ 1, 2, \dots, n \}$. The release time step is hence given by $\hat{r}_i := \frac{r_i}{\alpha} \in \mathbb{N}_0$.

2.2 Network Loading

Given a network, discretization parameters and a set of packets it is necessary to determine the movement of the packets through the network. We do this by specifying the arc dynamics and the node transitions algorithmically, which combined determines the position of every packet over time. In every discrete time step, we first consider the dynamic of all arcs. This primarily means to determine the packets that leave each arc. Afterwards, the node transition is executed for all nodes, moving the packets to the next arc on their respective path. This alternating process is repeated for every time step until every packet has reached its destination. Note that this happens in finite time, as packets travel along simple paths. Since all transit times are strictly positive, a packet can only enter one arc per time step. This justifies processing all arc dynamics first while computing the node transitions afterwards in each time step.

Arc queue. We consider every arc e of the network to be a packet queue Q_e denoted as an ordered finite sequence $Q_e(t)$ of packets for every time step $t \in \mathbb{N}_0$. In addition, packets are labeled with the time step at which they have entered e. Note that at any time step the ordering of all packets that are currently on this arc respects the arc entrance times. In other words, the queues operate by the first in first out (FIFO) principle. How ties are broken for particles that enter an arc at the same time step is specified by the node transition algorithm given below. In the following we denote different parts of the arc queue, which are also illustrated in Figure 1.

Buffer. At every time step t the *buffer* $B_e(t)$ is defined as the sequence of packets on arc e, that have been in Q_e long enough time to possibly leave e at time step t. More formally, we define $B_e(t)$ to be the maximal sequence of packets that satisfy

(1) $B_e(t)$ is a suffix of $Q_e(t)$,
(2) the arc entrance time step of all packets in $B_e(t)$ is at most $t - \hat{\tau}_e$.

Current capacities. Not necessarily all packets in the buffer will be able to leave the arc at time step t, as their volume might exceed the current arc capacity. Since the discretized capacity $\hat{\nu}_e$ can be non-integer and surely the number of packets leaving the arc has to be integer, the remainder is passed on to the next time step. For this reason we introduce the current packet capacity $\hat{\nu}_e(t)$. It is initialized by $\hat{\nu}_e(0) := \hat{\nu}_e$ and then defined iteratively by

$$\hat{\nu}_e(t) := \begin{cases} \hat{\nu}_e & \text{if } |B_e(t-1)| \leq \hat{\nu}_e(t-1), \\ \hat{\nu}_e + \hat{\nu}_e(t-1) - \lfloor \hat{\nu}_e(t-1) \rfloor & \text{else.} \end{cases} \tag{1}$$

Note that $\lfloor \hat{\nu}_e(t-1) \rfloor$ describes the number of packets that will be allowed to leave arc e at time step $t-1$ as we specify in the next paragraph. Hence, $\hat{\nu}_e(t-1) - \lfloor \hat{\nu}_e(t-1) \rfloor$ is the unused remaining capacity of the last time step, in the case that the number of buffering packets in $B_e(t-1)$ exceeds the previous capacity $\hat{\nu}_e(t-1)$.

Leaving packets. At every time step the algorithm determines at every arc e a list of leaving packets $L_e^-(t)$, which is an ordered list of packets in the buffer that will leave the arc in this time step. Formally, $L_e^-(t)$ is the maximal sequence of packets that satisfies the following two conditions:

(1) $L_e^-(t)$ is a suffix of $B_e(t)$,
(2) the number of packets in $L_e^-(t)$ does not exceed the current capacity $\hat{\nu}_e(t)$.

In other words, these are the first $\lfloor \hat{\nu}_e(t) \rfloor$ packets in line which have spent enough time on e to traverse it. Clearly, this $L_e^-(t)$ is unique and it can be determined algorithmically by considering one packet after the other in $B_e(t)$.

Waiting packets. The waiting packets $Z_e(t)$ at time step t are exactly the packets in the buffer $B_e(t)$ that are not leaving in time step t, i.e., that are not contained in $L_e^-(t)$. Following the naming of the flow over time model we refer to $Z_e(t)$ as *(waiting) queue*.

Released packets. We have an additional list $L_v^-(t)$ for each node v, which contains the packets that are released into the network through v at time step t. In other words,

$$L_v^-(t) := \{ i \in N \mid o_i = v \text{ and } \hat{r}_i = t \}.$$

Here, the ordering is given by the packet index.

Node transition. In each time step every leaving packet is moved from its current arc to its subsequent arc according its path. During this process the ordering of the packets is preserved. Whenever packets from multiple arcs enter into the same arc they are merged similar to the zipper method in real traffic. This means that packets are merged separately on each arc proportional to the number of packets sharing the same previous arc; see Figure 2 for an example. The precise procedure can be found in Algorithm 1 and is explained in the following.

The algorithm considers a node v at a fixed time step t and describes how packets traverse this node. The basic idea is to consider each outgoing arc e individually. All incoming arcs e' that contain at least one leaving packet that wants to continue its journey on arc e are collected in the set A. Here, packets that enter the network at node v at the current time step are treated as if they would be leaving packets of an additional arc. The packets should merge as smoothly as possible. Hence, in order to obtain a good ordering of packets entering e, we define a priority counter for all

ALGORITHM 1: Node transition at node v at a fixed time step t (the parameter t is omitted here).

```
Input: leaving packets L⁻_{e'} of all incoming arcs e' ∈ δ⁻_v;
       list of packets L⁻_v that want to enter the network at v at the current time step.
Output: sorted list L⁺_e of packets entering arc e for all outgoing arcs e ∈ δ⁺_v.
L⁺_e ← () for all e ∈ δ⁺_v
for each e ∈ δ⁺_v do
    // initialization of incoming arcs:
    for each e' ∈ δ⁻_v do
        Y_{e'} ← ordered sublist of packets i from L⁻_{e'} with (e', e) on path P_i
        if |Y_{e'}| ≥ 1 then
            y_{e'} ← |Y_{e'}|   // number of packets going from e' to e
            a_{e'} ← 1/y_{e'}    // priority counter (small means higher priority)
        end
    end
    A ← { e' ∈ δ⁻_v | |Y_{e'}| ≥ 1 }   // arcs that contain packets moving into e
    // initialization of released packets:
    Y_v ← ordered sublist of packets i from L⁻_v where e is the first arc of path P_i
    if |Y_v| ≥ 1 then
        y_v ← |Y_v|   // number of packets released into e
        a_v ← 1/y_v    // priority counter
        A ← A ∪ { v }   // v is treated as an incoming arc
    end
    // transition:
    while A ≠ ∅ do
        e* ← argmin_{e'∈A} a_{e'}   // ties are broken arbitrarily (but deterministically)
        L⁺_e.append(Y_{e*}.pop())   // remove first packet from Y_{e*} and add it to the end of L⁺_e
        a_{e*} ← a_{e*} + 1/y_{e*}
        if |Y_{e*}| = 0 then
            A ← A \ { e* }
        end
    end
end
return (L⁺_e)_{e∈δ⁺_v}
```

incoming arcs e' that operates proportionally to the number of packets $y_{e'}$ transiting from e' to e in the current time step. Initially, these counters are set to $\frac{1}{y_{e'}}$ and the incoming arc with the lowest counter (highest priority) can send the next packet. Afterwards, the lowest counter is increased by $\frac{1}{y_{e'}}$ and if e' does not contain any further leaving packets it is removed from A. A visualization of the priority counter for the merging process of the top outgoing arc in the example of Figure 2 is given in Figure 3. Note that in the end all leaving packets are transferred onto their next arc and that this merging procedure is mainly important for the ordering on the outgoing arcs.

Network loading. The loading of the network can be defined algorithmically as follows. For every time step we first consider all arcs and determine the leaving packets. Next, all leaving packets are moved to the next arc of their path according to the node transition. Finally, the queues and the current capacities are updated. This is described formally in Algorithm 2.

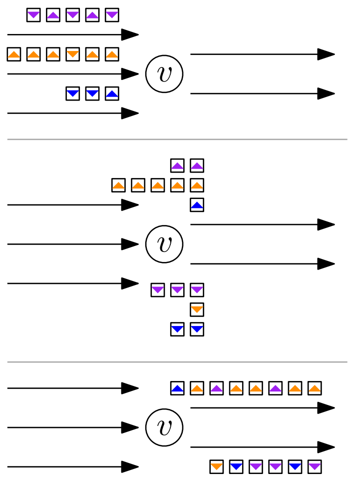

Fig. 2. The node transition in three steps: leaving packets of each incoming arc are depicted on the left. The small arrows indicate onto which outgoing arc the packets will be moved. As a first step the packets get grouped by desired outgoing arcs (middle). Finally, the packets are merged according to the priority procedure (right side). Here, ties are broken from top incoming arc to bottom incoming arc.

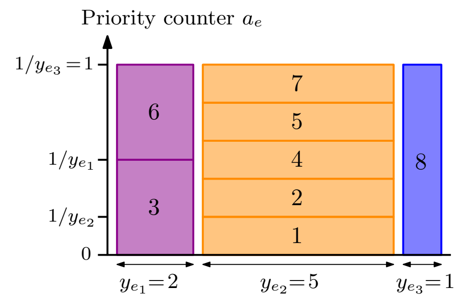

Fig. 3. Visualization of the merging process with priority counter. We consider the merging process of the upper outgoing arc of Figure 2. For each incoming arc we draw the packets as a stack of rectangles each of area 1 with a width of y_{e_i}. The height is then describing the increment of the priority counter. Whenever a packet is selected by the algorithm, the height of the upper arc of the rectangle corresponding to the packet equals the priority counter. As packets with low priority counter are selected first, we can order the packets from the bottom to the top (always considering the upper edge of the rectangle). In this examples ties are broken in favour of the arc with lower index.

ALGORITHM 2: Network loading.

Input: discretized network $(V, E, (\hat{\nu}_e)_{e\in E}, (\hat{\tau}_e)_{e\in E})$; set of packets N with $(o_i, d_i, P_i, \hat{r}_i)_{i\in N}$.
Output: arrival time steps $(\hat{c}_i)_{i\in N}$.
$t \leftarrow 0$
$\hat{c}_i \leftarrow \infty$ for all $i \in N$
$Q_e(0) \leftarrow ()$ for all $e \in E$
$\hat{\nu}_e(0) \leftarrow \hat{\nu}_e$ for all $e \in E$

while $\exists\, i \in N$ *with* $\hat{c}_i = \infty$ **do**
  determine $L_e^-(t)$ as defined in paragraph *leaving packets* for all $e \in E$
  for *each* $i \in N$ *with* $i \in L_e^-(t)$ *and* $e \in \delta_{d_i}^-$ **do**
    $\hat{c}_i \leftarrow t$ // packet i has reached its destination
  end
  for *each* $v \in V$ **do**
    $L_v^-(t) \leftarrow$ list of all packets $i \in N$ with $\hat{r}_i = t$ and $o_i = v$ ordered by packet index
    $(L_e^+(t))_{e\in\delta_v^+} \leftarrow \mathrm{NodeTransition}((L_e^-(t))_{e\in\delta_v^-}, L_v^-(t))$
  end
  for *each* $e \in E$ **do**
    set arc entrance time step to t for each packet in $L_e^+(t)$
    $Q_e(t+1) \leftarrow Q_e(t) - L_e^-(t) + L_e^+(t)$ // remove leaving and append entering packets
    update the current capacity $\hat{\nu}_e(t+1)$ according to Equation (1)
  end
  $t \leftarrow t+1$
end
return $(\hat{c}_i)_{i\in N}$

2.3 Competitive Packet Routing Game

In this section we consider the model from a game-theoretical perspective, which means that every packet is considered to be a player who aims at arriving as early as possible at her destination.

Strategies and costs. We are given a discretized network $(V, E, (\hat{\nu}_e)_{e\in E}, (\hat{\tau}_e)_{e\in E})$ and a set of packets N, each packet $i \in N$ equipped with an (o_i, d_i)-pair and a release time $\hat{r}_i$. Note that we assume that d_i is reachable from o_i for all $i \in N$. By considering the packets as players and denoting the set of all simple o_i-d_i-paths as the strategy set of player i, we obtain a competitive packet routing game.

For every path profile $\pi = (P_i)_{i\in N}$ the network loading given by Algorithm 2 determines the arrival time steps $(\hat{c}_i(\pi))_{i\in N}$. The arrival time given by $c_i(\pi) := \alpha \cdot \hat{c}_i(\pi)$ is the cost of player i, which she aims to minimize.

Equilibria. We consider pure Nash equilibria in this model. These are states in which no player can arrive strictly earlier by unilaterally changing her path. This equilibrium concept has been considered in related models before [2, 29]. With the arguments from [29] it follows that in single-o-d-games, i.e., when all packets share the same origin-destination-pair, a pure Nash equilibrium always exists. The idea is that players decide on a shortest path one-by-one in their release order. Since the tie-breaking mechanism at the node transition is known in advance, each player can choose a shortest path on which she can never be overtaken, and hence, on which she cannot be delayed by subsequent players. This yields a pure Nash equilibrium, as changing the path does not interfere with preceding players.

For multi-commodity games pure Nash equilibria do not exist in general. This follows from the example presented by Ismaili [14], which we adapt to our model here.

Proposition 2.1. *There are competitive packet routing games in our model for which no pure Nash equilibrium exists.*

Consider a game with six players played on the discretized network depicted in Figure 4 with $\alpha = \beta = 1$. The two main players, pursuer 1 going from o_P to d_P and evader 2 going from o_E to d_E, are the only players who can choose a path. The purpose of the remaining four players is to transfer the information between the main players. By giving the four long arcs priority in the merging, we obtain a matching pennies game between the pursuer and the evader, which does not possess a pure Nash equilibrium. For a detailed proof we refer to the full version of this paper.

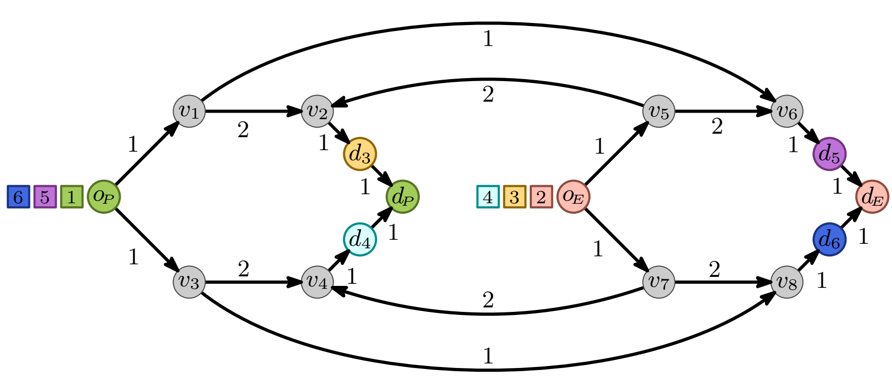

Fig. 4. Network of a game without pure Nash equilibrium. They are six players, indicated by different colors. The transit times are depicted beside the arcs and all arcs have a capacity of 1.

ε-Equilibria. Since the existence of pure Nash equilibria cannot be guaranteed, we consider approximate pure Nash equilibria. In these states players choose nearly optimal strategies, i.e., they cannot improve by more than ε by changing their paths.

Definition 2.2. For $\varepsilon > 0$ an ε-equilibrium is a strategy profile $\pi = (P_i)_{i \in N}$ such that for every player $i \in N$ and every simple o_i-d_i-path P_i' it holds that

$$c_i(\pi) \leq c_i(P_{-i}, P_i') + \varepsilon.$$

In the example of *Proposition* 2.1 every state is a 1-equilibrium but there do not exist any ε-equilibria for $\varepsilon < 1$. Thus, for a given game and a given ε it is not obvious whether there exists an ε-equilibrium or not. In Section 6, however, we prove that for any network (with given supply rates, which are defined later on) and any given $\varepsilon > 0$ we can find a discretization such that an ε-equilibrium exists.

3 FLOWS OVER TIME

In order to be able to analyze the connection between the packet routing model and flows over time in the following sections, we briefly introduce the multi-commodity flow over time model which is studied in [4, 10, 25, 26]. This model extends the network flow concept by a continuous time component. Each infinitesimally small flow particle needs a fixed time to traverse an arc before it reaches a bottleneck given by the capacity rate of the arc. Whenever the flow rate exceeds this capacity a queue builds up, which operates at capacity rate.

3.1 Model

A flow over time *instance* consists of a network $(V, E, (\nu_e)_{e \in E}, (\tau_e)_{e \in E})$ together with a finite set of commodities J, where each commodity $j \in J$ is equipped with a set of particles $M_j = [0, m_j]$, $m_j \in \mathbb{R}_{>0}$, an origin-destination-pair (o_j, d_j) and an integrable and bounded *supply rate* function $u_j \colon [0, \infty) \to [0, \infty)$. The inflow rate function u_j must have bounded support and must satisfy $\int_0^\infty u_j(\xi)\,\mathrm{d}\xi = m_j$.

A multi-commodity flow over time is represented by a set of functions $f = (f^+_{j,e}, f^-_{j,e})_{j \in J, e \in E}$ that are bounded and locally integrable. These functions describe the rate with which the flow of commodity j enters and leaves arc e at every point in time. Since flow does not enter the network before time 0, we set $f^+_{j,e}(\theta) = f^-_{j,e}(\theta) = 0$ for $\theta < 0$. We denote the *cumulative flow* by

$$F^+_{j,e}(\theta) := \int_0^\theta f^+_{j,e}(\xi)\,\mathrm{d}\xi \quad \text{and} \quad F^-_{j,e}(\theta) := \int_0^\theta f^-_{j,e}(\xi)\,\mathrm{d}\xi.$$

Furthermore, we denote the *total (cumulative) flow*, which is the sum over all commodities, by

$$f^+_e(\theta) := \sum_{j \in J} f^+_{j,e}(\theta), \quad f^-_e(\theta) := \sum_{j \in J} f^-_{j,e}(\theta), \quad F^+_e(\theta) := \sum_{j \in J} F^+_{j,e}(\theta) \quad \text{and} \quad F^-_e(\theta) := \sum_{j \in J} F^-_{j,e}(\theta).$$

In order to describe the flow dynamics in the deterministic queuing model, several conditions need to be fulfilled for almost all $\theta \in [0, \infty)$. First of all, we need commodity-wise flow conservation at every node at every point in time:

$$\sum_{e \in \delta^+_v} f^+_{j,e}(\theta) - \sum_{e \in \delta^-_v} f^-_{j,e}(\theta) \begin{cases} = 0 & \text{for all } v \in V \setminus \{o_j, d_j\}, \\ = u_j(\theta) & \text{for } v = o_j, \\ \leq 0 & \text{for } v = d_j. \end{cases} \tag{2}$$

Let the total flow volume in the queue at time θ be denoted by $z_e(\theta)$. For particles that enter at time θ we obtain a *waiting time* $q_e(\theta)$ and an *exit time* $T_e(\theta)$. Formally, they are given by

$$z_e(\theta) := F_e^+(\theta - \tau_e) - F_e^-(\theta), \qquad q_e(\theta) := \frac{z_e(\theta + \tau_e)}{\nu_e} \quad \text{and} \quad T_e(\theta) := \theta + \tau_e + q_e(\theta).$$

As the queue operates at capacity rate whenever we have a positive queue, we obtain the following condition on the total outflow rate:

$$f_e^-(\theta) = \begin{cases} \nu_e & \text{if } z_e(\theta) > 0, \\ \min\left\{ f_e^+(\theta - \tau_e), \nu_e \right\} & \text{if } z_e(\theta) = 0. \end{cases} \tag{3}$$

Finally, we require that the outflow rate of a commodity j at time $T_e(\vartheta)$ corresponds proportionally to its fractional part of the total inflow rate at the entrance time ϑ:

$$f_{j,e}^-(\theta) = \begin{cases} f_e^-(\theta) \cdot \frac{f_{j,e}^+(\vartheta)}{f_e^+(\vartheta)} & \text{if } f_e^+(\vartheta) > 0, \\ 0 & \text{else,} \end{cases} \tag{4}$$

where $\vartheta = \min\{ \xi \le \theta \mid T_e(\xi) = \theta \}$. This ensures that the proportions of the commodities are preserved during the arc traversal. Additionally, this condition guarantees FIFO since flow of a commodity cannot overtake flow of other commodities.

If Conditions (2) to (4) are met, we call f a *feasible flow over time.*

Arrival times. A given feasible flow over time f uniquely determines the node arrival times of a particle $\phi \in M_j$ along a simple o_j-d_j-path P. Note that particle ϕ enters the source o_j at time

$$\ell_{j,o_j}^P(\phi) := \min\left\{ \theta \ge 0 \,\middle|\, \int_0^\theta u_j(\xi)\,\mathrm{d}\xi = \phi \right\}.$$

For an arc $e = uv$ of path P the *arrival time* at node v of particle ϕ is given by

$$\ell_{j,v}^P(\phi) := T_e\left(\ell_{j,u}^P(\phi)\right).$$

3.2 Multi-Commodity Nash Flows Over Time

In order to consider dynamic equilibria in this setting we first need to define the earliest point in time a particle can arrive at a certain node. This depends on a given feasible flow over time f (in particular on the preceding particles), but for clearance we omit f in the following notation.

Earliest arrival times. For every commodity $j \in J$ and particle $\phi \in M_j$ the earliest arrival time of ϕ at some node v is given by taking the minimum arrival time at node v over all simple o_j-d_j-paths $\mathcal{P}_j$. Formally,

$$\ell_{j,v}(\phi) = \min_{P \in \mathcal{P}_j} \ell_{j,v}^P(\phi).$$

Current shortest paths network. For every commodity $j \in J$ and every particle $\phi \in M_j$ we define the current shortest paths network $E'_{j,\phi}$ by all arcs that are on a shortest o_j-d_j-path, i.e.,

$$E'_{j,\phi} := \left\{ e \in E \,\middle|\, \text{there exists an } P \in \mathcal{P}_j \text{ with } e \in P \text{ and } \ell_{j,d_j}^P(\phi) = \ell_{j,d_j}(\phi) \right\}.$$

Finally, this enables us to define a dynamic equilibrium, which is a Nash equilibrium with infinitely many players. It is a feasible flow over time, in which (almost) every particle corresponding to a player uses a quickest o_j-d_j-path.

Definition 3.1. A feasible flow over time is a *multi-commodity Nash flow over time* if for all arcs $e = uv \in E$, all commodities $j \in J$ and almost all point in time $\theta \in [0, \infty)$ it holds that

$$f^+_{j,e}(\theta) > 0 \quad \Rightarrow \quad \theta \in \left\{ \ell_{j,u}(\phi) \;\middle|\; \phi \in M_j \text{ such that } e \in E'_{j,\phi} \right\}.$$

For all networks and all supply rate functions $u_j \in L^p$ with bounded support such a Nash flow over time exists [4]. Since this is the case in our model the existence of a Nash flow over time is always guaranteed.

4 COUPLING PARTICLES AND PACKETS

In the following we connect a flow over time instance with discretization parameters to obtain a packet routing instance. Additionally, we transfer the definitions and notations from flows over time to the packet routing model. Overall, this section serves as foundation for the convergence proof in Section 5.

The setting is as follows. We are given a network G and a finite set of commodities J, each with an integrable and bounded supply rate function $u_j : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ that has bounded support. Let $m_j := \int_0^\infty u_j(\xi)\,\mathrm{d}\xi$ and $\bar{u}_j := \max_{\theta \in \mathbb{R}_{\geq 0}} u_j(\theta)$. As we focus on the network loading problem we assume without loss of generality that each commodity only uses a single fixed simple path P_j.

For given discretization parameters (α, β) we define for each commodity a set of discrete packets with suitable release times and set up a tool box of properties that connects discrete packet routings with flows over time.

Notation. Let $f = (f^+_{j,e}, f^-_{j,e})_{j \in J, e \in E}$ be a feasible multi-commodity flow over time with functions $F^+_{j,e}$, $F^-_{j,e}$, z_e, q_e, T_e and $\ell_{j,v}$ as described in Section 3.

Packets. In order to obtain a packet flow, we consider packets $N_j := \{ 1, 2, \ldots, \lfloor m_j/\beta \rfloor \}$ with path P_j for every commodity $j \in J$. The release time of packet $i \in N_j$ is given by the release time of the last particle corresponding to i, i.e.,

$$r_i := \min \left\{ \theta \in \Theta_\alpha \;\middle|\; \int_0^\theta u_j(\xi)\,\mathrm{d}\xi \geq i \cdot \beta \right\}.$$

Note that the total volume of all packets that commodity j is sending is given by $\beta \cdot |N_j| = \lfloor m_j \rfloor_\beta$.

In- and outflow functions. In order to describe the corresponding packet routing, the *inflow rate* $g^+_{j,e}(\theta)$, for $\theta \in \Theta_\alpha$, denotes the combined volume of packets of commodity j that enter arc e in time step θ/α divided by α (i.e., we obtain a rate in vol/sec). The *outflow rate* $g^-_{j,e}(\theta)$ is defined analogously. Formally, we have

$$g^+_{j,e}(\theta) := \frac{\beta}{\alpha} \cdot \left|\{ i \in N_j \mid i \in L^+_e(\theta/\alpha) \}\right| \qquad \text{and} \qquad g^-_{j,e}(\theta) := \frac{\beta}{\alpha} \cdot \left|\{ i \in N_j \mid i \in L^-_e(\theta/\alpha) \}\right|.$$

The *cumulative inflow* $G^+_{j,e}(\theta)$ and the *cumulative outflow* $G^-_{j,e}(\theta)$ denote the combined volume of packets of commodity j that have entered/left e up to time $\theta \in \Theta_\alpha$ (including time step $t = \theta/\alpha$). They are defined as

$$G^+_{j,e}(\theta) := \sum_{\xi \in \Theta_\alpha, \xi \leq \theta} g^+_{j,e}(\xi) \cdot \alpha \qquad \text{and} \qquad G^-_{j,e}(\theta) := \sum_{\xi \in \Theta_\alpha, \xi \leq \theta} g^-_{j,e}(\xi) \cdot \alpha.$$

The set of all commodities that use an arc e is denoted by $J_e := \{ j \in J \mid e \in P_j \}$ and the set of all commodities which pass node v is denoted by $J_v := \{ j \in J \mid v \text{ lies on } P_j \}$. For each arc e the *total inflow* function is denoted by $g^+_e := \sum_{j \in J_e} g^+_{j,e}$. Analogously, we define the *total outflow* function g^-_e

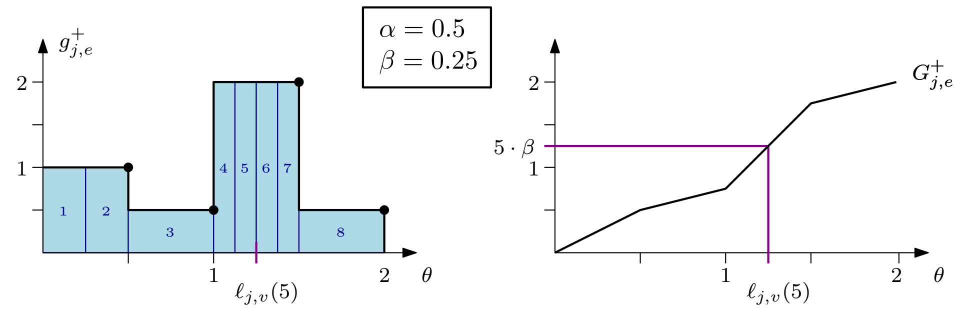

Fig. 5. To exemplary illustrate the refined arrival times, suppose we have $\alpha = 0.5$, $\beta = 0.25$ and eight packets $i = 1, 2, 3, 4, 5, 6, 7, 8$ of commodity j. Suppose packet 1 and 2 enter arc $e = vw$ at time 0.5 packet 3 at time 1, packet 4, 5, 6, 7 at time 1.5 and packet 8 at time 2. As depicted on the left side the packets are represented by a rectangle with area β that are equally distributed under the graph of $g^+_{j,e}$ in the time interval of length α right before the entrance time. The refined entrance time $\ell^D_{j,v}$ of packet i is then given by the right edge of the rectangle of i. This can also be seen by considering the cumulative inflow $G^+_{j,e}$ (right side of the figure). Hence, packet 5 has a refined entrance time of $\ell^D_{j,v}(5) = 1.25$. It enters the arc at time step $\lceil \ell^D_{j,v}(5) \rceil_\alpha = 1.5$ on position $k = (1.25 - (1.5 - 0.5))/0.5 \cdot 4 = 2$.

and the *total cumulative flows* functions G^+_e and G^-_e. Naturally, all these functions map from Θ_α to $\mathbb{R}_{\geq 0}$. We extend this notation to $\mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ as follows.

Refined time. In order to describe the arrival times and cumulative in- and outflows for the packet model we use continuous piece-wise linear functions. The idea is to consider packets as a continuous flow that is equally distributed along the time frame of one time step; see left side of Figure 5. The in- and outflow functions $g^+_{j,e}, g^-_{j,e} : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ are extended to be constant during each time step. In other words, we extend these functions by setting $g^+_{j,e}(\theta) := g^+_{j,e}(\lceil \theta \rceil_\alpha)$ and $g^-_{j,e}(\theta) := g^-_{j,e}(\lceil \theta \rceil_\alpha)$ for all $\theta \in \mathbb{R}_{\geq 0}$. With this the cumulative in- and outflow functions are extended by

$$G^+_{j,e}(\theta) := \int_0^\theta g^+_{j,e}(\xi)\, \mathrm{d}\xi \qquad \text{and} \qquad G^-_{j,e}(\theta) := \int_0^\theta g^-_{j,e}(\xi)\, \mathrm{d}\xi.$$

The *refined arrival time* of packet $i \in N_j$ at node v is defined as

$$\ell^D_{j,v}(i) := \min \left\{ \theta \in \mathbb{R}_{\geq 0} \;\middle|\; G^-_{j,uv}(\theta) \geq i \cdot \beta \right\} = \min \left\{ \theta \in \mathbb{R}_{\geq 0} \;\middle|\; G^+_{j,vw}(\theta) \geq i \cdot \beta \right\}$$

for $uv, vw \in P_j$. Note that for $v = o_j$ the arrival time $\ell^D_{j,v}(i)$ is defined by $G^+_{o_j w}$ with $o_j w \in P_j$, and analogously, for $v = d_j$ it is defined by $G^-_{ud_j}$ with $ud_j \in P_j$. The superscript D indicates that we consider the discrete packet model. We also call $\ell^D_{j,v}(i)$ the *refined entrance time* into vw.

As all packets of a commodity j take the same path and the arc dynamics follow the FIFO principle we have that $\ell^D_{j,v}(i) < \ell^D_{j,v}(i')$ for $i < i'$. Note that packet i is processed in the node transition of v at time $\lceil \ell^D_{j,v}(i) \rceil_\alpha$. The refined arrival time also contains the information about the position $k \in \{1, 2, \dots\}$ of packet i among all packets of commodity j that enter arc vw at this time. More precisely, the position k is obtained by

$$k = \frac{1}{\alpha} \cdot \left(\ell^D_{j,v}(i) - (\overbrace{\lceil \ell^D_{j,v}(i) \rceil_\alpha}^{\text{time step}} - \alpha) \right) \cdot \overbrace{g^+_{j,vw}(\ell_{j,v}(i)) \cdot \frac{\alpha}{\beta}}^{\#\text{packets of } j \text{ in } L^+_{vw}}.$$

A visual representation of the refined arrival times can be found in Figure 5.

Queue sizes, waiting times and exit times. We define the commodity-specific queue size $z^D_{j,e}$, waiting times $q^D_{j,e}$ and exit times $T^D_{j,e}$ for every commodity j corresponding to the flow over time model.

For a commodity j the *commodity-specific queue size* $z^D_{j,e}$ is given for all $\theta \in \mathbb{R}_{\geq 0}$ by

$$z^D_{j,e}(\theta + \lceil \tau_e \rceil_\alpha) := G^+_{j,e}(\theta) - G^-_{j,e}(\theta + \lceil \tau_e \rceil_\alpha).$$

Note that for $\theta \in \Theta_\alpha$ the queue size $z^D_{j,e}(\theta)$ describes the total volume of packets in $Z_e(\theta/\alpha) \cap N_j$. These are exactly the packets of commodity j that have spend at least $\lceil \tau_e \rceil_\alpha$ time on e but are not selected to leave at time step θ/α. The *total queue size* (or simply *queue size*) is given by $z^D_e := \sum_{j \in J_e} z^D_{j,e}$.

The *waiting time* function $q^D_{j,e}$ is obtained by setting

$$q^D_{j,e}(\theta) := \min \left\{ q \geq 0 \;\middle|\; \int_{\theta + \lceil \tau_e \rceil_\alpha}^{\theta + \lceil \tau_e \rceil_\alpha + q} g^-_{j,e}(\xi)\, \mathrm{d}\xi \geq z^D_{j,e}(\theta + \lceil \tau_e \rceil_\alpha) \right\}$$

for $\theta \in \mathbb{R}_{\geq 0}$. It describes the waiting time of packets entering arc e at refined time θ. With this we obtain the *exit time* function

$$T^D_{j,e}(\theta) := \theta + \lceil \tau_e \rceil_\alpha + q^D_{j,e}(\theta).$$

That this function does indeed describe the exit time of a packet that entered e at refined time θ is shown by the following lemma.

Lemma 4.1. *If packet $i \in N_j$ enters $e = uv$ at refined time $\ell^D_{j,u}(i) \in \mathbb{R}_{\geq 0}$, it leaves e at refined time*

$$\ell^D_{j,v}(i) = T^D_{j,e}(\ell^D_{j,u}(i)).$$

Proof. Let $\theta := \ell^D_{j,u}(i)$. By definition of $q^D_{j,e}$ together with the fact that $G^+_{j,e}(\theta) = i\beta$ we have

$$\begin{aligned}
T^D_{j,e}(\theta) &= \theta + \lceil \tau_e \rceil_\alpha + \min \left\{ q \geq 0 \;\middle|\; \int_{\theta + \lceil \tau_e \rceil_\alpha}^{\theta + \lceil \tau_e \rceil_\alpha + q} g^-_{j,e}(\xi)\, \mathrm{d}\xi \geq z^D_{j,e}(\theta + \lceil \tau_e \rceil_\alpha) \right\} \\
&= \min \left\{ T \geq \theta + \lceil \tau_e \rceil_\alpha \;\middle|\; \int_{\theta + \lceil \tau_e \rceil_\alpha}^{T} g^-_{j,e}(\xi)\, \mathrm{d}\xi \geq G^+_{j,e}(\theta) - G^-_{j,e}(\theta + \lceil \tau_e \rceil_\alpha) \right\} \\
&= \min \left\{ T \geq 0 \;\middle|\; G^-_{j,e}(T) \geq G^+_{j,e}(\theta) \right\} \\
&= \min \left\{ T \geq 0 \;\middle|\; G^-_{j,e}(T) \geq i\beta \right\} \\
&= \ell^D_{j,v}(i).
\end{aligned}$$

□

5 CONVERGENCE

In this section we finally show that every feasible flow over time is the limit of the sequence of packet routings when decreasing the packet size and time step length coherently. Observe that if α went quicker to zero than β, packets would get isolated, i.e., the number of time steps between two successive packets would increase (since the current capacity $\hat{\nu}_e$ would go to zero). To avoid this and, even stronger, to guarantee that the number of packets leaving a queue per time step becomes large, we require that β goes quicker to zero than α. More precisely, we only consider sequences of convergence parameters (α, β) with $\alpha \to 0$ and $\frac{\beta}{\alpha} \to 0$ from now on.

As key component for convergence we prove the following theorem.

Theorem 5.1. *Let $(V, E, (\nu_e)_{e\in E}, (\tau_e)_{e\in E})$ be a network and J be a finite set of commodities, each $j \in J$ equipped with a simple path P_j and an integrable and bounded supply rate function u_j with bounded support. Consider the packet routing induced by discretization parameters (α, β) such that α and $\frac{\beta}{\alpha}$ are sufficiently small. Given a node v, it holds for every packet $i \in N_j$ of commodity $j \in J_v$ and the corresponding particle $\phi := i\beta \in M_j$ that*

$$\left|\ell_{j,v}(\phi) - \ell^D_{j,v}(i)\right| \leq \sqrt{\alpha} \cdot C^K \qquad \text{for } K \in \mathbb{N} \text{ such that } \ell_{j,v}(\phi) < K \cdot \tau^* \text{ or } \ell^D_{j,v}(i) < K \cdot \tau^*.$$

Here, $C > 1$ is a constant that only depends on the instance and $\tau^* := \min_{e\in E} \tau_e/2 > 0$.

To be more precise, α is *sufficiently small* if $\sqrt{\alpha}$ is smaller than $\tau^* C^{-H/\tau^*}$, where H is the time when the last particle left the network. For β/α to be *sufficiently small* it has to be smaller than 1 and small enough such that there are at least two packets allowed to leave e at each time step, i.e., $\hat{\nu}_e = \nu_e \cdot \alpha/\beta \geq 2$ (including the upper bound on the supply rates).

Note that this theorem implies a convergence rate of $\sqrt{\alpha}$ for the arrival times of a given packet-particle-pair. Though, the constant might grow exponentially in the packet/particle index.

The roadmap of the proof goes as follows. We do an induction on K, i.e., in each induction step we extend the validity of the statement by τ^* time units. This enables us to consider each arc individually. For the induction step we introduce a hypothetical flow over time that matches the inflow rates of the packet model but follows the flow over time dynamics. With this additional ingredient we are able to show that the exit times, and hence the arrival times, of the packets and corresponding particles do not differ too much.

For an important simplification, note that packets of a commodity j that enter into the network at a node v behave like packets coming from an incoming arc with capacity $\bar{u}_j$ (cf. Algorithm 1). For the sake of clarity we do not introduce additional notation but instead implicitly assume that δ^-_v includes these auxiliary arcs with $\nu_e := \bar{u}_j$.

Furthermore, we define the *rate bound* κ_e for all arcs $e = uv$ by

$$\kappa_e := \max\left\{\sum_{e'\in\delta^-_u} (\nu_{e'} + 1), \nu_e + 1\right\}.$$

This is an instance constant that only depends on the arc capacities, including the upper bounds of the supply rates. Since the number of leaving packets of each predecessor arc e' (including the auxiliary arcs for newly released packets) is upper bounded by $\hat{\nu}_{e'} + 1$ in every time step, we have that $g^-_{e'}(\theta) \leq \nu_{e'} + \beta/\alpha$. Since β/α is assumed to be smaller than 1, κ_e does indeed bound the maximal inflow into e as well as the maximal outflow of e for both the packet model as well as the flow over time model.

We structure the proof with the help of three lemmas, which we present in the following. Due to the technical nature of the proofs and the limited space we only present the key ideas here. The complete proofs can be found in the full version.

We start with the following lemma showing that the waiting time in the packet model can be approximated by the queue size divided by the capacity.

Lemma 5.2. *For all arcs $e = uv$, all commodities $j \in J_e$ and every $\theta \in \mathbb{R}_{\geq 0}$ the waiting time at arc e in the packet model satisfies*

$$\left|q^D_{j,e}(\theta) - \frac{z^D_e(\theta + \lceil\tau_e\rceil_\alpha)}{\nu_e}\right| \leq 2\alpha + \frac{\alpha\kappa_e}{\nu_e} + \frac{\beta}{\nu_e}.$$

The key observation of the proof is that the current capacities in average over the time steps are very similar to ν_e. Combined with several deviations caused by rounding in time and flow volume, this leads to the stated bound.

Hypothetical flow over time. For a fixed arc $e = uv$ we introduce the *hypothetical flow over time* by setting $h^+_{j,e} := g^+_{j,e}$ and $h^-_{j,e} \colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ according to (3) and (4). In other words, we take the inflow rates of the discrete packet model, but determine the outflow rates according to the flow over time model. We define the cumulative flows as before by $H^+_{j,e}(\theta) := \int_0^\theta h^+_{j,e}(\xi)\,\mathrm{d}\xi$ and $H^-_{j,e}(\theta) := \int_0^\theta h^-_{j,e}(\xi)\,\mathrm{d}\xi$. Clearly, we have $H^+_{j,e} = G^+_{j,e}$, but in general we cannot expect that $H^-_{j,e}$ equals $G^-_{j,e}$. Again, we define the *total hypothetical cumulative flow* by $H^+_e := \sum_{j \in J} H^+_{j,e}$ and $H^-_e := \sum_{j \in J} H^-_{j,e}$. The queue size, waiting time and exit time function for the hypothetical inflow rate are denoted by z^H_e, q^H_e and T^H_e. We can even define hypothetical arrival times at u and v for all $\phi \in M_j$ by

$$k_{j,u}(\phi) := \min\{\, \theta \in \mathbb{R}_{\geq 0} \mid H^+_{j,e}(\theta) = \phi \,\} \qquad \text{and} \qquad k_{j,v}(\phi) := \min\{\, \theta \in \mathbb{R}_{\geq 0} \mid H^-_{j,e}(\theta) = \phi \,\}.$$

Clearly, we have that $k_{j,u}(i\beta) = \ell^D_{j,u}(i)$ for all $i \in N_j$. Since the hypothetical flow on arc e follows the continuous flow dynamics we have $k_{j,v}(\phi) = k_{j,u}(\phi) + \tau_e + q^H_e(k_{j,u}(\phi))$. Note that the hypothetical flow rates are considered on each arc separately. In particular, we cannot expect that they satisfy flow conservation.

In the next lemma we show that, given some particle ϕ, whenever the difference of the arrival times at u (hypothetical flow compared to continuous flow) is not too large for all particles arriving earlier than ϕ, then this is also true for the difference of the arrival times of particle ϕ at node v.

Lemma 5.3. *Consider an arc $e = uv$, a commodity $j \in J_e$ and a particle $\phi \in M_j$ with $\phi \in \beta\mathbb{N}$. Suppose there exists a $\delta > \alpha$ with*

$$\left|\ell_{j',u}(\varphi) - k_{j',u}(\varphi)\right| \leq \delta$$

for all commodities $j' \in J_e$, and all particles $\varphi \in M_{j'}$ with $\varphi \in \beta\mathbb{N}$ and $\ell_{j',u}(\varphi) \leq \ell_{j,u}(\phi)$ or $k_{j',u}(\varphi) \leq k_{j,u}(\phi)$. Then it holds that

$$\left|\ell_{j,v}(\phi) - k_{j,v}(\phi)\right| \leq 11\frac{\kappa_e}{\nu_e} \cdot \delta.$$

The proof of this lemma works in three steps. First, we bound the difference between the total cumulative inflows $F^+_e(\ell_{j,u}(\phi))$ and $H^+_e(k_{j,u}(\phi))$ by carefully analyzing the flow of all commodities. In step two we show that the total cumulative outflows $F^-_e(\ell_{j,u}(\phi) + \tau_e)$ and $H^-_e(k_{j,u}(\phi) + \tau_e)$ are not too different by combining the flow over time dynamics with the result of step one. This leads to a bounded difference of the queue lengths, and hence, of the waiting times. In the final step, we use this to show that the deviation of the arrival times does not increase too much when going from u to v.

In the next lemma we bound the difference between the hypothetical arrival time $k_{j,v}$ and the refined packet arrival time $\ell^D_{j,v}$.

Lemma 5.4. *Given an arc $e = uv$, a point in time θ and a commodity $j \in J_e$. For sufficiently small α it holds that*

$$\left|k_{j,v}(\phi) - \ell^D_{j,v}(\lceil \tfrac{\phi}{\beta} \rceil)\right| \leq C_1 \cdot \sqrt{\alpha} + \theta \cdot C_2 \cdot \sqrt{\alpha} \qquad \text{for } \phi \in M_j \text{ with } k_{j,u}(\phi) \leq \theta.$$

Here, C_1 and C_2 are positive constants that only depend on the instance properties κ_e and ν_e.

This lemma is proven by an involved induction over $\phi \in M_j$ in $\sqrt{\alpha} - \beta$ steps. For the induction we merge all commodities into a single commodity, which only exists on e. This causes just a negligible additional error. For the induction step we distinguish three cases. If the discrete and the hypothetical queue are smaller than $\sqrt{\alpha}$ (Case 1) then Lemma 5.2 implies that the arrival times at v are similar. If the hypothetical queue is longer than the discrete queue and longer than $\sqrt{\alpha}$ (Case 2) we consider the particle ϕ' that leaves e at the point in time when ϕ lines up in the queue. The difference between the arrival times of ϕ' and the corresponding packet is bounded by induction. Since the queues in both models operate roughly at the same capacity rate, the difference between the arrival times of ϕ and $\lceil\frac{\phi}{\beta}\rceil$ is small. In the remaining case, the proof idea is similar to Case 2 but with interchanged roles for particles and packets.

With the help of the previous lemmas, we can finally prove Theorem 5.1:

Proof of Theorem 5.1. We prove this theorem by an induction on K. For the base case we consider a point in time $\theta \in [0, \tau^*)$. In this case we only have to consider the nodes o_j for each commodity j since for all other nodes there is no particle $\phi \in M_j$ with $\ell_{j,v}(\phi) \leq \theta$ and also no packet i with $\ell^D_{j,v}(i) \leq \theta$.

For the release times we get

$$\begin{aligned} \ell^D_{j,o_j}(i) = r_i &= \min\left\{ \theta \in \Theta_\alpha \;\middle|\; \int_0^\theta u_j(\xi)\,\mathrm{d}\xi \geq i\beta \right\} \\ &= \min\left\{ \theta \in \Theta_\alpha \;\middle|\; \int_0^\theta u_j(\xi)\,\mathrm{d}\xi \geq \phi \right\} \in [\ell_{j,o_j}(\phi), \ell_{j,o_j}(\phi) + \alpha]. \end{aligned}$$

For the induction step assume the theorem is proven for $K-1$. Consider a packet $i \in N_j$ and a particle $\phi := i\beta \in M_j$ with $\min\{\ell_{j,v}(\phi), \ell^D_{j,v}(i)\} < K\tau^*$. If i is released at node v the statement holds with the considerations of the base case. Otherwise, for the node u with $e = uv \in P_j$ it holds that $\min\{\ell_{j,u}(\phi), \ell^D_{j,u}(i)\} \leq \min\{\ell_{j,v}(\phi), \ell^D_{j,v}(i)\} - \tau_e \leq (K-1)\tau^*$, and therefore, we can apply the induction hypothesis to node u to obtain

$$\max\{\ell_{j,u}(\phi), \ell^D_{j,u}(i)\} \leq \min\{\ell_{j,u}(\phi), \ell^D_{j,u}(i)\} + \frac{\tau_e}{2} \leq \min\{\ell_{j,v}(\phi), \ell^D_{j,v}(i)\} - \frac{\tau_e}{2} \leq (K-1)\tau^*. \quad (5)$$

For the first inequality note that we require $\sqrt{\alpha} < \tau_e/2 \cdot C^{-(K-1)}$.

By introducing the hypothetical flow on arc e, we can apply Lemma 5.3 with $\delta := C^{K-1}\sqrt{\alpha}$. The preconditions of the lemma are satisfied, since we can apply the induction hypothesis to every φ with $\min\{\ell_{j,u}(\varphi), \ell^D_{j,u}(\varphi)\} \leq \max\{\ell_{j,u}(\phi), \ell^D_{j,u}(\phi)\}$ thanks to (5). Note that $k_{j',u}(\phi') = \ell^D_{j',u}(i')$ for all $i' \in N_{j'}$ and $\phi' = i'\beta \in M_{j'}$. Lemma 5.3 gives us

$$\left|\ell_{j,v}(\phi) - k_{j,v}(\phi)\right| \leq 11\frac{\kappa_e}{\nu_e}C^{K-1}\sqrt{\alpha}.$$

For sufficiently small α it holds that $C^{K-1}\sqrt{\alpha} \leq \tau^*$. Hence, again by the induction hypothesis, we obtain

$$\ell^D_{j,u}(i) \leq \min\{\,\ell^D_{j,u}(i), \ell_{j,u}(\phi)\,\} + C^{K-1}\sqrt{\alpha} < (K-1)\tau^* + \tau^* \leq K\tau^*.$$

Applying Lemma 5.4 with $\theta := K\tau^*$ yields

$$\left|k_{j,v}(\phi) - \ell^D_{j,v}(i)\right| \leq C_1\sqrt{\alpha} + K\tau^* C_2\sqrt{\alpha}.$$

Finally, by choosing C large enough we obtain

$$\begin{aligned} \left|\ell_{j,v}(\phi) - \ell^D_{j,v}(i)\right| &\leq \left|\ell_{j,v}(\phi) - k_{j,v}(\phi)\right| + \left|k_{j,v}(\phi) - \ell^D_{j,v}(i)\right| \\ &\leq 11\frac{\kappa_e}{\nu_e}C^{K-1}\sqrt{\alpha} + C_1\sqrt{\alpha} + K\tau^* C_2\sqrt{\alpha} \\ &\leq (C - \sqrt{C})C^{K-1}\sqrt{\alpha} + KC\sqrt{\alpha} \\ &\leq \left(C^K - C^{K-0.5} + KC\right)\sqrt{\alpha} \\ &\leq C^K\sqrt{\alpha}. \end{aligned}$$

The inequalities hold by choosing C large enough such that the following conditions are satisfied

$$C - \sqrt{C} \geq 11 \frac{\kappa_e}{\nu_e}, \qquad C \geq C_1 + C_2 \tau^* \quad \text{and} \quad C^{K-0.5} \geq KC \quad \text{for } K \in \{2, 3, \dots\}.$$

Note that the last condition holds for all $C \geq 4$. □

Since each commodity has a bounded flow mass and the supply rate functions have a bounded support, there exists a point in time when the last particle enters the network. As the particles travel along predefined paths there exists a finite time horizon, i.e., a point in time at which all particles have arrived at their destinations and the network is empty from this point in time onward. By Theorem 5.1 this implies that there is also a finite time horizon for all packet routing instances with $\alpha \leq 1$ simultaneously. With this observation Theorem 5.1 implies convergence of the arrival times, and even stronger, it is possible to show that the flows converge by proving uniform convergence of the cumulative flow rates.

Theorem 5.5. *Let $(V, E, (\nu_e)_{e \in E}, (\tau_e)_{e \in E})$ be a network and J be a finite set of commodities, each $j \in J$ equipped with a simple o_j-d_j-path P_j and an integrable and bounded supply rate function u_j with bounded support. Consider the packet routings induced by a sequence of discretization parameters (α, β) such that α and $\frac{\beta}{\alpha}$ go to zero. Then the following two statements hold:*

1. *For every $v \in V$, $j \in J_v$ and $\phi \in M_j$, the arrival time $\ell^D_{j,v}(\lceil \frac{\phi}{\beta} \rceil)$ of the packet model converges to the arrival time $\ell_{j,v}(\phi)$ of the flow model.*
2. *For all $e = uv \in E$ and all $j \in J_e$ the cumulative flow $G^+_{j,e}$ of the packet model converges uniformly to the cumulative flow $F^+_{j,e}$ of the flow model. Analogously, $G^-_{j,e}$ converges uniformly to $F^-_{j,e}$.*

Proof. For (1) first observe that by the continuity of $\ell_{j,v}$ we have that $\ell_{j,v}(\lceil \phi \rceil_\beta) \longrightarrow \ell_{j,v}(\phi)$ for $\beta \longrightarrow 0$. The statement then follows immediately by Theorem 5.1 with K large enough such that $\ell_{j,v}(\phi) \leq K\tau^*$.

For (2) let $\varepsilon > 0$. By Theorem 5.1 it is possible to choose (α, β) small enough such that

$$\left| \ell^D_{j,u}(i) - \ell_{j,u}(i\beta) \right| \leq \frac{\varepsilon - \beta}{\kappa_e}$$

for every packet $i \in N_j$. This is possible by again choosing K large enough such that at time $K\tau^*$ all particles have left the network.

Given a point in time θ, we consider the particle $\phi := F^+_{j,e}(\theta) \in M_j$. We denote the two packets that are as close as possible to ϕ by $i^- := \lfloor \frac{\phi}{\beta} \rfloor \in N_j$ and $i^+ := \lceil \frac{\phi}{\beta} \rceil \in N_j$ (see below for the case that one of them does not exist).

Note that the arrival time of packet i^- cannot be much later than θ, and analogously, the arrival time of i^+ cannot be much earlier than θ. More precisely,

$$\ell^D_{j,u}(i^-) - \frac{\varepsilon - \beta}{\kappa_e} \leq \ell_{j,u}(i^- \beta) \leq \theta \leq \ell_{j,u}(i^+ \beta) \leq \ell^D_{j,u}(i^+) + \frac{\varepsilon - \beta}{\kappa_e}.$$

Since the slope of $G^+_{j,e}$ is bounded by κ_e and $F^+_{j,e}(\theta) \leq i^- \beta + \beta$, we obtain

$$G^+_{j,e}(\theta) \geq G^+_{j,e}(\ell^D_{j,u}(i^-)) - \frac{\varepsilon - \beta}{\kappa_e} \cdot \kappa_e = i^- \beta - (\varepsilon - \beta) \geq F^+_{j,e}(\theta) - \varepsilon.$$

Analogously, we get an upper bound on $G^+_{j,e}$ by using $F^+_{j,e}(\theta) \geq i^+ \beta - \beta$

$$G^+_{j,e}(\theta) \leq G^+_{j,e}(\ell^D_{j,u}(i^+)) + \frac{\varepsilon - \beta}{\kappa_e} \cdot \kappa_e = i^+ \beta + \varepsilon - \beta \leq F^+_{j,e}(\theta) + \varepsilon.$$

For the case that $i^- = \lfloor \frac{\phi}{\beta} \rfloor = 0 \notin N_j$ we use the trivial lower bound of $G^+_{j,e}(\theta) \geq 0$ and for the case that $i^+ = \lceil \frac{\phi}{\beta} \rceil = |N_j| + 1 \notin N_j$ we use the upper bound $G^+_{j,e}(\theta) \leq |N_j| \cdot \beta \leq \phi + \beta \leq F^+_{j,e}(\theta) + \varepsilon$.

With the exact same arguments it holds that $\left|G^-_{j,e}(\theta) - F^-_{j,e}(\theta)\right| \leq \varepsilon$. □

6 EXISTENCE OF APPROXIMATE PURE NASH EQUILIBRIA

As an application of the convergence result we show the existence of approximate pure Nash equilibria in our competitive packet routing model by using the existence of exact dynamic equilibria for flows over time. This is particularly interesting for the multi-commodity setting, where exact pure Nash equilibria do not exist in general as it was shown in Proposition 2.1.

Theorem 6.1. *Consider a network $(V, E, (\nu_e)_{e \in E}, (\tau_e)_{e \in E})$, a constant $T > 0$ and a set of commodities J, each $j \in J$ equipped with an origin-destination-pair (o_j, d_j) and an integrable and bounded supply rate function $u_j : [0, T] \to \mathbb{R}_{\geq 0}$. For every $\varepsilon > 0$ there are discretization parameters (α, β), such that the corresponding competitive packet routing game possesses an ε-equilibrium.*

Proof. Fix $\varepsilon > 0$. We consider a multi-commodity Nash flow over time f, which exists due to Cominetti et al. [4, Theorem 8]. It can be decomposed into a path-based flow over time $(f_P)_{P \in \mathcal{P}}$, where $\mathcal{P} := \{ P \text{ is an } o_j\text{-}d_j\text{-path of some commodity } j \}$ (cf. [4, Section 2.7]). Additionally, Cominetti et al. have shown in [4, Lemma 7] that the path arrival times $(\ell_{P,d_P})_{P \in \mathcal{P}}$ depend continuously on the path inflow rates $(f_P)_{P \in \mathcal{P}}$. Hence, we can choose $\delta > 0$ small enough such that for all path inflow rates $(\tilde{f}_P)_{P \in \mathcal{P}}$, which differ by at most δ from $(f_P)_{P \in \mathcal{P}}$ (i.e., $\sum_{P \in \mathcal{P}} \|f_P - \tilde{f}_P\|_{L^2} \leq \delta$) the corresponding arrival times $(\tilde{\ell}_{P,d_P})_{P \in \mathcal{P}}$ satisfy

$$\left|\ell_{P,d_P}(\phi) - \tilde{\ell}_{P,d_P}(\phi)\right| \leq \frac{\varepsilon}{4} \quad \text{for all particles } \phi \text{ and all paths } P \in \mathcal{P}. \tag{6}$$

By considering each path P in $\mathcal{P}$ as a separate commodity, we obtain path-based supply rates $(u_P)_{P \in \mathcal{P}} := (f_P)_{P \in \mathcal{P}}$. Combining Theorem 5.1 with a suitable time horizon H provides a pair of discretization parameters (α, β) such that for every packet $i \in N_P$ in the corresponding packet routing it holds that

$$\left|\ell_{P,d_P}(i\beta) - \ell^D_{P,d_P}(i)\right| \leq \frac{\varepsilon}{4} \qquad \text{for all } P \in \mathcal{P}.$$

Additionally, we require β to be small enough such that (6) holds for all $\tilde{f}$ that can be created by shifting a flow volume of at most β within f.

In the following we show that the obtained packet routing constitutes an ε-equilibrium. Assume for contradiction that there is a packet i that would arrive at least ε earlier by switching from its path P to some other o_P-d_P-path P'. By considering the modified supply rates $\tilde{u}$ that result from shifting the β flow volume corresponding to packet i from P to P', we obtain a modified flow over time $\tilde{f}$ and a modified packet routing with arrival times $\tilde{\ell}$ and $\tilde{\ell}^D$. Again with Theorem 5.1 we obtain that

$$\left|\tilde{\ell}_{P',d_P}(i\beta) - \tilde{\ell}^D_{P',d_P}(i)\right| \leq \frac{\varepsilon}{4}.$$

Note, that the constant C in Theorem 5.1 only depends on the instance. By choosing C large enough, it is possible to shift the flow volume without loosing the bound on the arrival times.

Combining all this, we obtain

$$\begin{aligned}\ell_{P,d_P}(i\beta) - \ell_{P',d_{P'}}(i\beta) &\overset{(6)}{\geq} \ell_{P,d_P}(i\beta) - \tilde{\ell}_{P',d_{P'}}(i\beta) - \frac{\varepsilon}{4} \\ &\geq \ell^D_{P,d_P}(i) - \tilde{\ell}^D_{P',d_{P'}}(i) - \frac{3\varepsilon}{4} \\ &\geq \varepsilon - \frac{3\varepsilon}{4} > 0.\end{aligned}$$

But this is a contradiction to the fact that f is a Nash flow over time, since all particles of the flow volume corresponding to packet i use path P even though path P' is strictly faster. □

7 OPEN PROBLEMS

This work is the first that proves a connection between continuous dynamic models and discrete packet routing models. There are several possible directions to proceed with this research.

Our bound on the discretization error increases exponentially in time. Experiments [30] indicate that there might be a time independent bound though. Can a better upper and maybe also a lower bound on the worst case error be shown mathematically? Furthermore, what is the best convergence rate (in α) for a fixed particle? We prove a convergence rate of $\sqrt{\alpha}$, which is clearly not the best possible. However, even rounding the transit time of a single arc causes a deviation of up to $\alpha/2$ showing that a convergence rate cannot be faster than linear in α.

In the single-origin-destination setting, pure Nash equilibria exist for all competitive packet routing games and Nash flows over time are well understood. Given a sequence of these exact packet routing equilibria, whose discretization parameters converge to zero, does the limit exist? If yes, does it constitute a Nash flow over time? What would this mean for the price of anarchy (PoA)? Packet routing models often have a PoA in $\Omega(|N|)$, whereas the PoA of Nash flows over time is conjectured to be $\frac{e}{e-1}$ [7]. This seems to be contradicting at first glance, but in the extreme examples that achieve a PoA of $\Omega(|N|)$ in the packet routing models the networks grow linearly with the number of players. For the convergence process the network is fixed, thus, the PoA might decrease. Further investigation in this regard would certainly be interesting.

In the transport simulation MATSim, traffic assignments are computed by an iterative best response procedure. Naturally, this raises the question whether best response dynamics converge to approximate equilibria in the competitive packet routing model. If this could be proven mathematically, it would further extend the theoretical foundation of MATSim.

ACKNOWLEDGMENTS

We thank Tobias Harks, Kai Nagel and Martin Skutella for fruitful discussions on the topic. We also thank the anonymous reviewers for their very helpful comments and questions.

This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project ID: 390685689) and the Research Training Group 2236 UnRAVeL.

REFERENCES

[1] G.E. Cantarella and D.P. Watling. 2016. Modelling road traffic assignment as a day-to-day dynamic, deterministic process: a unified approach to discrete- and continuous-time models. *EURO Journal on Transportation and Logistics* 5, 1 (2016), 69–98.

[2] Z. Cao, B. Chen, X. Chen, and C. Wang. 2017. A network game of dynamic traffic. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC.* Association for Computing Machinery, New York, NY, USA, 695–696.

[3] R. Cominetti, J. Correa, and O. Larré. 2011. Existence and uniqueness of equilibria for flows over time. In *International Colloquium on Automata, Languages, and Programming, ICALP.* Springer-Verlag, Zurich, Switzerland, 552–563.

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

[5] R. Cominetti, J. Correa, and N. Olver. 2017. Long Term Behavior of Dynamic Equilibria in Fluid Queuing Networks. In *International Conference on Integer Programming and Combinatorial Optimization, IPCO*. Springer, Waterloo, Canada, 161–172.

[6] R. Cominetti, M. Scarsini, M. Schröder, and N. Stier-Moses. 2021. Approximation and Convergence of Large Atomic Congestion Games. arXiv:cs.GT/2001.02797

[7] J. Correa, A. Cristi, and T. Oosterwijk. 2019. On the price of anarchy for flows over time. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC*. ACM, Phoenix, Arizona, USA, 559–577.

[8] L. Fleischer and É. Tardos. 1998. Efficient continuous-time dynamic network flow algorithms. *Operations Research Letters* 23, 3-5 (1998), 71–80.

[9] D. Frascaria and N. Olver. 2020. Algorithms for flows over time with scheduling costs. In *International Conference on Integer Programming and Combinatorial Optimization, IPCO*. Springer, London, UK, 130–143.

[10] L. Graf, T. Harks, and L. Sering. 2020. Dynamic flows with adaptive route choice. *Mathematical Programming* 183, 1 (2020), 309–335.

[11] T. Harks, B. Peis, D. Schmand, B. Tauer, and L. Vargas Koch. 2018. Competitive packet routing with priority lists. *ACM Transactions on Economics and Computation (TEAC)* 6, 1 (2018), 1–26.

[12] M. Hoefer, V.S. Mirrokni, H. Röglin, and S.-H. Teng. 2011. Competitive routing over time. *Theoretical Computer Science* 412, 39 (2011), 5420–5432.

[13] A. Horni, K. Nagel, and K. Axhausen (Eds.). 2016. *Multi-agent transport simulation MATSim*. Ubiquity Press, London, UK.

[14] A. Ismaili. 2017. Routing Games over Time with FIFO Policy. In *International Conference on Web and Internet Economics, WINE*. Springer International Publishing, Bangalore, India, 266–280.

[15] J. Israel and L. Sering. 2020. The impact of spillback on the price of anarchy for flows over time. In *International Symposium on Algorithmic Game Theory, SAGT*. Springer, Springer, Augsburg, Germany, 114–129.

[16] M. Kaiser. 2020. Computation of Dynamic Equilibria in Series-Parallel Networks. arXiv:cs.GT/2002.11428

[17] R. Koch and M. Skutella. 2011. Nash equilibria and the price of anarchy for flows over time. *Theory of Computing Systems* 49, 1 (2011), 71–97.

[18] Frédéric Meunier and Nicolas Wagner. 2010. Equilibrium results for dynamic congestion games. *Transportation science* 44, 4 (2010), 524–536.

[19] H. Otsubo and A. Rapoport. 2008. Vickrey's model of traffic congestion discretized. *Transportation Research Part B: Methodological* 42, 10 (2008), 873–889.

[20] H.M. Pham and L. Sering. 2020. Dynamic equilibria in time-varying networks. In *International Symposium on Algorithmic Game Theory, SAGT*. Springer, Springer, Augsburg, Germany, 130–145.

[21] PTV AG 2016. *VISUM 16 – User Manual*. PTV AG, Karlsruhe, Germany.

[22] R.W. Rosenthal. 1973. The network equilibrium problem in integers. *Networks* 3, 1 (1973), 53–59.

[23] M. Scarsini, M. Schröder, and T. Tomala. 2018. Dynamic atomic congestion games with seasonal flows. *Operations Research* 66, 2 (2018), 327–339.

[24] R. Scheffler, M. Strehler, and L. Vargas Koch. 2018. Nash equilibria in routing games with edge priorities. arXiv:cs.GT/1803.00865

[25] L. Sering. 2020. *Nash flows over time*. Ph.D. Dissertation. Technische Universität Berlin.

[26] L. Sering and M. Skutella. 2018. Multi-source multi-sink Nash flows over time. In *18th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems, ATMOS*, Vol. 65. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Helsinki, Finland, 12:1–12:20.

[27] L. Sering and L. Vargas Koch. 2019. Nash flows over time with spillback. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*. SIAM, San Diego, California, USA, 935–945.

[28] J.G. Wardrop. 1952. Some theoretical aspects of road traffic research. *Proceedings of the Institute of Civil Engineers* 1, 3 (1952), 325–378.

[29] T.L. Werth, M. Holzhauser, and S.O. Krumke. 2014. Atomic routing in a deterministic queuing model. *Operations Research Perspectives* 1, 1 (2014), 18–41.

[30] T. Ziemke, L. Sering, L. Vargas Koch, M. Zimmer, K. Nagel, and M. Skutella. 2021. Flows over time as continuous limit of packet-based network simulations. *Transportation Research Procedia* 52 (2021), 123–130.