



Mathematics of Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Convergence of a Packet Routing Model to Flows over Time

Leon Sering, Laura Vargas Koch, Theresa Ziemke

To cite this article:

Leon Sering, Laura Vargas Koch, Theresa Ziemke (2023) Convergence of a Packet Routing Model to Flows over Time. Mathematics of Operations Research 48(3):1741-1766. <https://doi.org/10.1287/moor.2022.1318>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2022, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Convergence of a Packet Routing Model to Flows over Time

 Leon Sering,^{a,*} Laura Vargas Koch,^{a,*} Theresa Ziemke^{b,c,*}

^aInstitute for Operations Research, ETH Zürich, 8092 Zürich, Switzerland; ^bCombinatorial Optimization and Graph Algorithms, Technische Universität Berlin, 10623 Berlin, Germany; ^cTransport Systems Planning and Transport Telematics, Technische Universität Berlin, 10587 Berlin, Germany

*Corresponding authors

Contact: sering@math.ethz.ch,  <https://orcid.org/0000-0003-2953-1115> (LS); laura.vargas@math.ethz.ch,
 <https://orcid.org/0000-0002-7499-5958> (LVK); tziemke@vsp.tu-berlin.de,  <https://orcid.org/0000-0001-8812-9041> (TZ)

Received: September 13, 2021
 Revised: August 4, 2022
 Accepted: September 8, 2022
 Published Online in Articles in Advance:
 October 26, 2022

MSC2020 Subject Classification: Primary:
 05C21; secondary: 90B10, 91A25

<https://doi.org/10.1287/moor.2022.1318>

Copyright: © 2022 INFORMS

Abstract. Mathematical approaches for modeling dynamic traffic can be roughly divided into two categories: discrete *packet routing models* and continuous *flow over time models*. Despite very vital research activities on models in both categories, their connection was poorly understood so far. We build this connection by specifying a (competitive) packet routing model, which is discrete in terms of flow and time, and 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 cannot be guaranteed in the multicommodity setting.

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

Keywords: flow over time • packet routing • multi-commodity • dynamic equilibrium • convergence of discrete to continuous • discretization error • approximate equilibrium • dynamic traffic assignment

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 (PTV AG [23], Rosenthal [24], Wardrop [30]). These models are relatively simple and well-understood, for example, 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 exist various mathematical approaches as well as simulation tools in this category. Some consider discrete time steps (Harks et al. [12], Horni et al. [14], Scarsini et al. [25], Werth et al. [31]), whereas others model continuous time (Hoefler et al. [13], Koch and Skutella [18], Scheffler et al. [26]). Some model traffic as discrete travelers (vehicles/packets) (Harks et al. [12], Hoefler et al. [13], Horni et al. [14], Scarsini et al. [25], Scheffler et al. [26], Werth et al. [31]), whereas others model it as continuous splittable flows (Adamik and Sering [1], Frascaria and Olver [10], Graf et al. [11], Koch and Skutella [18]). In this paper, we consider two of these approaches: on the one hand, the intensively studied flow over time model with deterministic queuing (Koch and Skutella [18]), 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 (multiagent transport simulation, see Horni et al. [14]), and it is very similar to the model studied by Werth et al. [31].

The flows over time model yields exact user equilibria, called *dynamic equilibria* or *Nash flows over time* (Koch and Skutella [18]). They are guaranteed to exist (Cominetti et al. [4, 5]), and their structure is studied intensively (Cominetti et al. [6], Correa et al. [8], Israel and Sering [16], Kaiser [17], Olver et al. [20], Pham and Sering [22], Sering and Vargas Koch [29]). Even though most of the work is done for single origin–destination networks, the existence results and further structural insights are transferred to the multiterminal setting as well (Cominetti et al. [5], Sering and Skutella [28]); see also Sering [27] for an extended analysis on all facets of Nash flows over time. Even more general, Meunier and Wagner [19] 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 coevolutionary algorithm is used to compute approximate user equilibria. Despite the different perspectives of the two approaches, experiments by Ziemke et al. [32] 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 multicommodity 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. [31].

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.

On a high level, the proof works as follows. We assign particles to packets and do an induction over time to show that the travel time difference between particles and corresponding packets is bounded. For this, we consider each arc of the network individually and show how to bound the arrival time deviation between particles and packets at the head of the arc in dependency on the deviation given at the tail. As traversing an arc takes at least a minimum amount of time, the induction hypothesis can be applied.

Note that this convergence result focuses on the network loading aspect of the models and does not consider any game-theoretical aspects yet; that is, the routes for the commodities are fixed. 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 because, to our knowledge—the existence of (approximate) pure Nash equilibria in a multicommodity setting is not proven for other similar competitive packet routing models. Note that *exact* multicommodity 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 toward 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 [21] consider the Vickrey model and focus on experimental comparisons, and Cantarella and Watling [2] compare discrete and continuous time models by presenting mathematical tools and applications and unifying both in a single model.

In static routing games, convergence has been shown very recently by Cominetti et al. [7]. They prove convergence of equilibria in routing games from two perspectives: on the one hand, when player sizes decrease 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 [9] 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 Cao et al. [3], Ismaili [15], Scarsini et al. [25], and Werth et al. [31]. However, each model has a different focus. Werth et al. [31] 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. Whereas Ismaili [15] analyzes the computational complexity of computing a best response and a Nash equilibrium, Scarsini et al. [25] analyze the price of anarchy of a packet routing model with periodic inflow rates. Finally, Cao et al. [3] 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, in which 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-style merging 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. [31]. The only differences are the tie-breaking rules in the merging process and that we allow fractional capacities.

2.1. Preliminaries

Units. Because we consider convergence later on, we keep the time and packet size flexible. As help for orientation, we introduce some units for time and 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 $G = (V, E)$, in which every arc is equipped with a transit time $\tau_e > 0$ and a capacity $v_e > 0$. 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, that is, 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. Because the point in time corresponding to a time step $t \in \mathbb{N}_0$ is given by at , 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 , that is,

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

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 $\hat{\tau}_e := \lceil \frac{\tau_e}{\alpha} \rceil_\alpha \in \mathbb{N}$. The maximum number of packets that are allowed to leave an arc at each time step is given by $\hat{v}_e := \frac{v_e \alpha}{\beta}$. Note that \hat{v}_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 determining the packets that leave each arc. Afterward, 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. Because all transit times are strictly positive, a packet can only enter one arc per time step. This justifies processing all arc dynamics first and computing the node transitions afterward at 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 FIFO principle. How ties are broken for packets that enter an arc at the same time step is specified by the node transition algorithm given subsequently. 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 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 are able to leave the arc at time step t as their volume might exceed the current arc capacity. Because the discretized capacity \hat{v}_e can be noninteger, 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 $\hat{v}_e(t)$. It is initialized by $\hat{v}_e(0) := \hat{v}_e$ and then defined iteratively by

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

Note that $\lfloor \hat{v}_e(t-1) \rfloor$ describes the number of packets that is allowed to leave arc e at time step $t-1$ as we specify in the next paragraph. Hence, $\hat{v}_e(t-1) - \lfloor \hat{v}_e(t-1) \rfloor$ 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 $\hat{v}_e(t-1)$.


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 leave the arc at 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 $\hat{v}_e(t)$.

In other words, these are the first $\lfloor \hat{v}_e(t) \rfloor$ packets in line that 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 at time step t , that is, that are not contained in $L_e^-(t)$. Following the naming of the flow over time model we refer to $Z_e(t)$ as the (*waiting*) *queue*.

Figure 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.



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. At each time step, every leaving packet is moved from its current arc to its subsequent arc according to 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 in a zipper-style way similar to 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.

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 $e' \in \delta_v^-$; 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 $e \in \delta_v^+$.

$L_e^+ \leftarrow ()$ for all $e \in \delta_v^+$

for each $e \in \delta_v^+$ **do**

 //initialization of incoming arcs:

for each $e' \in \delta_v^-$ **do**

$Y_{e'} \leftarrow$ ordered sublist of packets i from $L_{e'}^-$ with (e', e) on path P_i

if $|Y_{e'}| \geq 1$ **then**

$y_{e'} \leftarrow |Y_{e'}|$ //number of packets going from e' to e

$a_{e'} \leftarrow \frac{1}{y_{e'}}$ //priority counter (small means higher priority)

end

end

$A \leftarrow \{e' \in \delta_v^- \mid |Y_{e'}| \geq 1\}$ //arcs that contain packets moving into e

 //initialization of released packets:

$Y_v \leftarrow$ ordered sublist of packets i from L_v^- where e is the first arc of path P_i

if $|Y_v| \geq 1$ **then**

$y_v \leftarrow |Y_v|$ //number of packets released into e

$a_v \leftarrow \frac{1}{y_v}$ //priority counter

$A \leftarrow A \cup \{v\}$ // v is treated as an incoming arc

end

 //transition:

while $A \neq \emptyset$ **do**

$e^* \leftarrow \arg \min_{e' \in A} a_{e'}$ //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^+

$a_{e^*} \leftarrow a_{e^*} + \frac{1}{y_{e^*}}$

if $|Y_{e^*}| = 0$ **then**

$A \leftarrow A \setminus \{e^*\}$

end

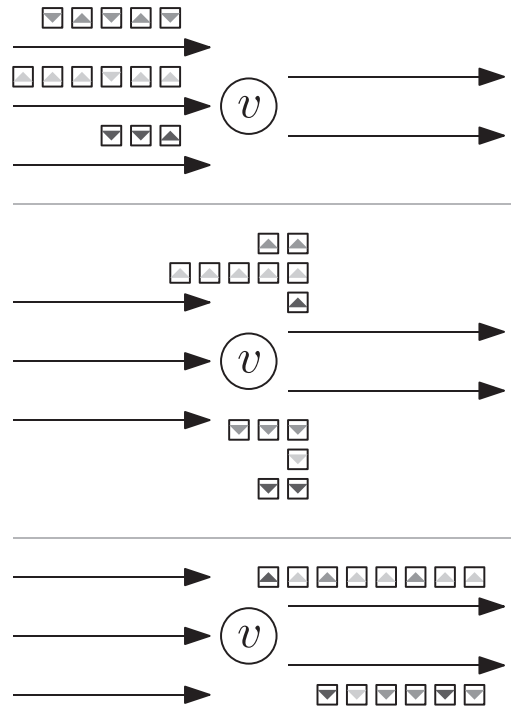
end

end

return $(L_e^+)_{e \in \delta_v^+}$.

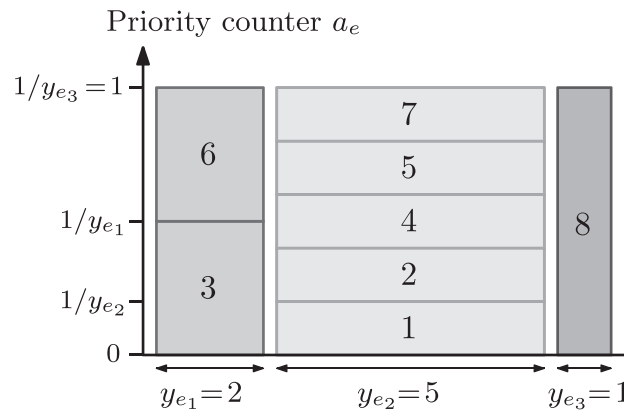
The algorithm considers a node v at a fixed time step t and describes how packets traverse this node. The basic idea is to consider each outgoing arc e individually. All incoming arcs e' that contain at least one leaving packet that wants to continue its journey on arc e are collected in the set A . 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 incoming arcs e' that operate proportionally to the number of packets $y_{e'}$ transiting from e' to e at the current time step. Initially, these counters are set to $\frac{1}{y_{e'}}$, and the incoming arc with the lowest counter

Figure 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 are 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 the top incoming arc to the bottom incoming arc.



(highest priority) can send the next packet. Afterward, 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 this merging procedure is mainly important for the ordering on the outgoing arcs.

Figure 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 a 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). Note that ties are broken in favor of the arc with the lower index.



Algorithm 2 (Network Loading)

Input: discretized network $(V, E, (\hat{v}_e)_{e \in E}, (\hat{\tau}_e)_{e \in E})$; set of packets N with $(o_i, d_i, P_i, \hat{\tau}_i)_{i \in N}$.

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  $e \in E$ 
 $\hat{v}_e(0) \leftarrow \hat{v}_e$  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  $\hat{\tau}_i = t$  and  $o_i = v$  ordered by packet index
         $(L_e^+(t))_{e \in \delta_v^+} \leftarrow \text{NodeTransition}((L_e^-(t))_{e \in \delta_v^-}, L_v^-(t))$ 
    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 & append entering packets
        update the current capacity  $\hat{v}_e(t+1)$  according to Equation (1)
    end
     $t \leftarrow t + 1$ 
end
return  $(\hat{c}_i)_{i \in N}$ 
    
```

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.

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 the player’s destination.

Strategies and Costs. We are given a discretized network $(V, E, (\hat{v}_e)_{e \in E}, (\hat{\tau}_e)_{e \in E})$ and a set of packets N , each packet $i \in N$ equipped with an (o_i, d_i) -pair and a release time $\hat{\tau}_i$. 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 the player aims to minimize.


Equilibria. We consider exact pure Nash equilibria in this model. These are states in which no player can arrive strictly earlier by unilaterally changing path. This equilibrium concept has been considered in related models before (Cao et al. [3], Werth et al. [31]).

In our model—similar to other packet routing models with the FIFO rule—pure Nash equilibria do not exist in general. This follows from the multicommodity example presented by Ismaili [15], which we adapt to our model here.

Proposition 1. *There are competitive packet routing games in our model for which no pure Nash equilibrium exists.*

Proof. Consider a game with six players played on the discretized network depicted in Figure 4 with $\alpha = \beta = 1$. We have two main players, a pursuer 1 going from o_P to d_P and an evader 2 going from o_E to d_E . Each of these two players has exactly two paths from which to choose, which we denote by top and bottom. Moreover, we have four additional players 3, . . . , 6 with origins $o_3 = o_4 = o_E$ and $o_5 = o_6 = o_P$ and destinations d_3, \dots, d_6 . Note that these players have a unique path and, hence, only a single strategy. Their purpose is to transfer the information between the main players.

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



All players have a release time of zero and are ordered by their index in the beginning. In a tie-break in the merging procedure, the four long arcs $(v_1v_6, v_5v_2, v_7v_4, v_3v_8)$ are prioritized. Note that all arc capacities are one, which implies that also the current capacities of all arcs are constant one (as there are no remaining capacities). Therefore, at most one packet can leave an arc per time step. Furthermore, if two packets enter into an arc at the same time, the priority counter always ends up in a tie (both counters are one), which means that the packet coming from the prioritized arc (long arc) is always transferred first.

First, we consider the network loading of the state in which both the pursuer and the evader choose their respective top path. In this case, the evader delays player 3 at the end of arc o_Ev_5 by one time step, and the pursuer delays player 5 at o_Pv_1 . Thus, the pursuer enters arc v_2d_3 before player 3 (as the pursuer arrives there at time 3 and player 3 arrives at time 4). Hence, the pursuer reaches d_P at time 5. On the other side, the evader is delayed by player 5 when leaving v_6d_5 (because both players entered the arc at the same time, but because of the tie break, player 5 is in front of the evader in the arc queue). Thus, the evader reaches d_E at time 6.

Now, consider the state in which the pursuer chooses the top path and the evader chooses the bottom path. In this case, the evader delays player 4 at the end of arc o_Ev_7 by one time step, and the pursuer still delays player 5 at o_Pv_1 . Thus, the pursuer enters arc v_2d_3 at the same time as player 3. By the tie-breaking rule of the arcs, player 3 is inserted in the arc queue of v_2d_3 before the pursuer and, thus, delays the pursuer when leaving the arc. This leads to an arrival time of the pursuer of 6. The evader reaches node v_8 at time 3 when player 6 passed already and can, thus, arrive without delay at time 5 in d_E .

Because the game is highly symmetric, this analysis gives us that the pursuer prefers to play strategy profiles top/top and bottom/bottom, whereas the evader prefers strategy profiles top/bottom and bottom/top. Because this is a classic matching pennies situation, the game has no pure Nash equilibrium. □

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

Definition 1. For $\epsilon > 0$, 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

$$c_i(\pi) \leq c_i(P_{-i}, P'_i) + \epsilon.$$

Here, (P_{-i}, P'_i) denotes the strategy profile in which all players except player i act as in P and only player i plays a different strategy, namely, P'_i .

In the example of Proposition 1, every state is a one-equilibrium, but there do not exist any ϵ -equilibria for $\epsilon < 1$. 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) and any given $\epsilon > 0$, 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 multicommodity flow over time model, which is studied in Cominetti

et al. [5], Graf et al. [11], Sering [27], and Sering and Skutella [28]. 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 $(V, E, (v_e)_{e \in E}, (\tau_e)_{e \in E})$ 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 $u_j : [0, \infty) \rightarrow [0, \infty)$. The inflow rate function u_j must have bounded support and must satisfy $\int_0^\infty u_j(\xi) \, d\xi = m_j$.

A multicommodity 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. Because flow does not enter the network before time 0, we set $f_{j,e}^+(\theta) = f_{j,e}^-(\theta) = 0$ for $\theta < 0$ and all $j \in J, e \in E$. We denote the *cumulative flow* by

$$F_{j,e}^+(\theta) := \int_0^\theta f_{j,e}^+(\xi) \, d\xi \quad \text{and} \quad F_{j,e}^-(\theta) := \int_0^\theta f_{j,e}^-(\xi) \, d\xi.$$

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, we need commodity-wise flow conservation at every node at every point in time:

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

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

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

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

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

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 :

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

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

If Conditions (2)–(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

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

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. Multicommodity 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 containing v . Formally,

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

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, that is,

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

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 2. A feasible flow over time is a *multicommodity Nash flow over time* if, for all arcs $e = uv \in E$, all commodities $j \in J$, and almost all points in time $\theta \in [0, \infty)$, it holds that

$$f_{j,e}^+(\theta) > 0 \implies \theta \in \{\ell_{j,u}(\phi) \mid \phi \in M_j \text{ such that } e \in E'_{j,\phi}\}.$$

For all networks and all supply rate functions $u_j \in L^p$ with bounded support such a Nash flow over time exists (Cominetti et al. [5]). Because 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 a 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 $u_j : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ that has bounded support. Let $m_j := \int_0^\infty u_j(\xi) d\xi$ and $\bar{u}_j := \sup_{\theta \in \mathbb{R}_{\geq 0}} u_j(\theta)$. 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 toolbox 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 multicommodity 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 $N_j := \{1, 2, \dots, \lfloor m_j / \beta \rfloor\}$ 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 , that is,

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

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

$$g_{j,e}^+(\theta) := \frac{\beta}{\alpha} \cdot |\{i \in N_j | i \in L_e^+(\theta/\alpha)\}| \quad \text{and} \quad g_{j,e}^-(\theta) := \frac{\beta}{\alpha} \cdot |\{i \in N_j | i \in L_e^-(\theta/\alpha)\}|.$$

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

$$G_{j,e}^+(\theta) := \sum_{\xi \in \Theta_\alpha, \xi \leq \theta} g_{j,e}^+(\xi) \cdot \alpha \quad \text{and} \quad G_{j,e}^-(\theta) := \sum_{\xi \in \Theta_\alpha, \xi \leq \theta} g_{j,e}^-(\xi) \cdot \alpha.$$

The set of all commodities that use an arc e is denoted by $J_e := \{j \in J | e \in P_j\}$, and the set of all commodities that pass node v is denoted by $J_v := \{j \in J | v \text{ lies on } P_j\}$. 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^- 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 $\mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ as follows.

Refined Time. In order to describe the arrival times and cumulative in- and outflows for the packet model, we use continuous piecewise 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 the left side of Figure 5. The in- and outflow functions $g_{j,e}^+, g_{j,e}^- : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ 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

$$G_{j,e}^+(\theta) := \int_0^\theta g_{j,e}^+(\xi) \, d\xi \quad \text{and} \quad G_{j,e}^-(\theta) := \int_0^\theta g_{j,e}^-(\xi) \, d\xi.$$


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

$$\ell_{j,v}^D(i) := \min\{\theta \in \mathbb{R}_{\geq 0} | G_{j,uv}^-(\theta) \geq i \cdot \beta\} = \min\{\theta \in \mathbb{R}_{\geq 0} | G_{j,vw}^+(\theta) \geq i \cdot \beta\}$$

for $uv, vw \in P_j$. Note that, for $v = o_j$, the arrival time $\ell_{j,v}^D(i)$ is defined by $G_{o_j w}^+$ with $o_j w \in P_j$, and analogously, for $v = d_j$ it is defined by $G_{u d_j}^-$ with $u d_j \in P_j$. The superscript D indicates that we consider the discrete packet model. We also call $\ell_{j,v}^D(i)$ 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 $\ell_{j,v}^D(i) < \ell_{j,v}^D(i')$ for $i < i'$. Note that packet i is processed in the node transition of v at time $\lceil \ell_{j,v}^D(i) \rceil_\alpha$. The refined arrival time also contains the information about the position $k \in \{1, 2, \dots\}$ of packet i among all packets of

Figure 5. To exemplarily 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 packets 1 and 2 enter arc $e = vw$ at time 0.5; packet 3 at time 1; packets 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 $\ell_{j,v}^D$ 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 $\ell_{j,v}^D(5) = 1.25$. It enters the arc at time step $\lceil \ell_{j,v}^D(5) \rceil_\alpha = 1.5$ on position $k = (1.25 - (1.5 - 0.5))/0.5 = 4 = 2$.



commodity j that enter arc vw at this time. More precisely, the position k is obtained by

$$k = \frac{1}{\alpha} \cdot (\ell_{j,v}^D(i) - (\lceil \ell_{j,v}^D(i) \rceil_\alpha - \alpha)) \cdot \overbrace{g_{j,vw}^+(\ell_{j,v}(i))}^{\text{#packets of } j \text{ in } L_{vw}^+} \cdot \frac{\alpha}{\beta}.$$

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 $z_{j,e}^D$, waiting times $q_{j,e}^D$ and exit times $T_{j,e}^D$ for every commodity j corresponding to the flow over time model.

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

$$z_{j,e}^D(\theta + \lceil \tau_e \rceil_\alpha) := G_{j,e}^+(\theta) - G_{j,e}^-(\theta + \lceil \tau_e \rceil_\alpha).$$

Note that, for $\theta \in \Theta_\alpha$, the queue size $z_{j,e}^D(\theta)$ describes the total volume of packets in $Z_e(\theta/\alpha) \cap N_j$. These are exactly the packets of commodity j that have spent 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 $z_e^D := \sum_{j \in J_e} z_{j,e}^D$.

The waiting time function $q_{j,e}^D$ is obtained by setting

$$q_{j,e}^D(\theta) := \min \left\{ q \geq 0 \mid \int_{\theta + \lceil \tau_e \rceil_\alpha}^{\theta + \lceil \tau_e \rceil_\alpha + q} g_{j,e}^-(\xi) d\xi \geq z_{j,e}^D(\theta + \lceil \tau_e \rceil_\alpha) \right\}$$

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

$$T_{j,e}^D(\theta) := \theta + \lceil \tau_e \rceil_\alpha + q_{j,e}^D(\theta).$$

The following lemma shows that this function indeed describes the exit time of a packet that entered e at refined time θ .

Lemma 1. *If packet $i \in N_j$ enters $e = uv$ at refined time $\ell_{j,u}^D(i) \in \mathbb{R}_{\geq 0}$, it leaves e at refined time*

$$\ell_{j,v}^D(i) = T_{j,e}^D(\ell_{j,u}^D(i)).$$

Proof. Let $\theta := \ell_{j,u}^D(i)$. By definition of $q_{j,e}^D$ together with the fact that $G_{j,e}^+(\theta) = i\beta$, we have

$$\begin{aligned} T_{j,e}^D(\theta) &= \theta + \lceil \tau_e \rceil_\alpha + \min \left\{ q \geq 0 \mid \int_{\theta + \lceil \tau_e \rceil_\alpha}^{\theta + \lceil \tau_e \rceil_\alpha + q} g_{j,e}^-(\xi) d\xi \geq z_{j,e}^D(\theta + \lceil \tau_e \rceil_\alpha) \right\} \\ &= \min \left\{ T \geq \theta + \lceil \tau_e \rceil_\alpha \mid \int_{\theta + \lceil \tau_e \rceil_\alpha}^T g_{j,e}^-(\xi) d\xi \geq G_{j,e}^+(\theta) - G_{j,e}^-(\theta + \lceil \tau_e \rceil_\alpha) \right\} \\ &= \min \{ T \geq 0 \mid G_{j,e}^-(T) \geq G_{j,e}^+(\theta) \} \\ &= \min \{ T \geq 0 \mid G_{j,e}^-(T) \geq i\beta \} \\ &= \ell_{j,v}^D(i). \quad \square \end{aligned}$$

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; that is, the number of time steps between two successive packets would increase (because the current capacity \hat{v}_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 $\alpha \rightarrow 0$ and $\frac{\beta}{\alpha} \rightarrow 0$ from now on.

As key component for convergence we prove the following theorem.

Theorem 1. *Let $(V, E, (v_e)_{e \in E}, (\tau_e)_{e \in E})$ 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*

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

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 $\tau^* C^{-\lceil H/\tau^* \rceil}$, 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 one and small enough such that there are at least two packets allowed to leave e at each time step; that is, $\hat{v}_e = v_e \cdot \alpha/\beta \geq 2$ (including when entering the network, i.e., $\bar{u}_j \cdot \alpha/\beta \geq 2$ for all commodities j).

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 road map of the proof goes as follows. We do an induction on K ; that is, 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 $v_e := \bar{u}_j$.

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

$$\kappa_e := \max \left\{ \sum_{e' \in \delta_u^-} (v_{e'} + 1), v_e + 1 \right\}.$$

This is an instance constant that only depends on the arc capacities, including the upper bounds of the supply rates. Because the number of leaving packets of each predecessor arc e' (including the auxiliary arcs for newly released packets) is upper bounded by $\hat{v}_{e'} + 1$ in every time step, we have that $g_{e'}^-(\theta) \leq v_{e'} + \beta/\alpha$. Because β/α is assumed to be smaller than one, 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. The first lemma shows that the waiting time in the packet model can be approximated by the queue size divided by the capacity.

Lemma 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

$$\left| q_{j,e}^D(\theta) - \frac{z_e^D(\theta + \lceil \tau_e \rceil \alpha)}{v_e} \right| \leq 2\alpha + \frac{\alpha \kappa_e}{v_e} + \frac{\beta}{v_e}.$$

Proof. By the definition of the waiting times, we obtain for each commodity j' separately that

$$z_{j',e}^D(\theta + \lceil \tau_e \rceil \alpha) = \int_{\theta + \lceil \tau_e \rceil \alpha}^{\theta + \lceil \tau_e \rceil \alpha + q_{j',e}^D(\theta)} g_{j',e}^-(\xi) d\xi.$$

Hence, because all $g_{j',e}^-$ are nonnegative, we obtain with $j_{\max} := \arg \max_{j' \in J} q_{j',e}^D(\theta)$ that

$$\begin{aligned} z_e^D(\theta + \lceil \tau_e \rceil \alpha) &\leq \int_{\theta + \lceil \tau_e \rceil \alpha}^{\theta + \lceil \tau_e \rceil \alpha + q_{j_{\max},e}^D(\theta)} g_e^-(\xi) d\xi \\ &= \int_{\theta + \lceil \tau_e \rceil \alpha}^{\theta + \lceil \tau_e \rceil \alpha + q_{j',e}^D(\theta)} g_e^-(\xi) d\xi + \int_{\theta + \lceil \tau_e \rceil \alpha + q_{j',e}^D(\theta)}^{\theta + \lceil \tau_e \rceil \alpha + q_{j_{\max},e}^D(\theta)} g_e^-(\xi) d\xi. \end{aligned} \quad (5)$$

For the right term, observe that, for two packets $i \in N_j$ and $i' \in N_{j'}$ entering at the same refined time (so, in particular, at the same time step), the total flow volume of packets entering between them is bounded by $\alpha \cdot \sum_{e' \in \delta_u^-} (v_{e'} + \beta/\alpha) \leq \alpha \kappa_e$ as this is the maximal volume that can transition at node u within one time step. As we have FIFO on the arc, the same is true for cumulative outflow at the exit times, that is,

$$|G_e^-(T_{j,e}^D(\ell_{j,u}^D(i))) - G_e^-(T_{j',e}^D(\ell_{j',u}^D(i')))| \leq \alpha \kappa_e.$$

In particular, we have

$$\int_{\theta + \lceil \tau_e \rceil \alpha + q_{j',e}^D(\theta)}^{\theta + \lceil \tau_e \rceil \alpha + q_{j_{\max},e}^D(\theta)} g_e^-(\xi) d\xi = G_e^-(T_{j_{\max},e}^D(\theta)) - G_e^-(T_{j',e}^D(\theta)) \leq \alpha \kappa_e.$$

For the left term of (5) observe that, for all time steps t with $\theta + \lceil \tau_e \rceil \alpha < t\alpha \leq \theta + \lceil \tau_e \rceil \alpha + q_{j',e}^D(\theta) + \alpha$, the number of leaving packets $L_e^-(t)$ is bounded by the current capacity $\hat{v}_e(t)$.

The number of time steps in this interval (of length $q_{j,e}^D(\theta) + \alpha$) is upper bounded by $q_{j,e}^D(\theta)/\alpha + 2$. By the definition of the current capacity in (1), we obtain that the number of packets that leave e within this interval is upper bounded by

$$\left(\frac{q_{j,e}^D(\theta)}{\alpha} + 2\right) \cdot \hat{v}_e + 1.$$

Here, that +1 is due to some potential remaining capacity that is transferred to the very first time step in the interval. Hence, we obtain with $\hat{v}_e = \frac{\alpha}{\beta} v_e$ that

$$\int_{\theta + \lceil \tau_e \rceil \alpha}^{\theta + \lceil \tau_e \rceil \alpha + q_{j,e}^D(\theta)} g_e^-(\xi) d\xi \leq \beta \cdot \left(\frac{q_{j,e}^D(\theta)}{\alpha} + 2\right) \cdot \hat{v}_e + \beta = q_{j,e}^D(\theta) \cdot v_e + 2 \cdot \alpha \cdot v_e + \beta.$$

In conclusion, (5) yields the following upper bound:

$$z_e^D(\theta + \lceil \tau_e \rceil \alpha) \leq q_{j,e}^D(\theta) \cdot v_e + \alpha \cdot (\kappa_e + 2 \cdot v_e) + \beta.$$

For the lower bound, we consider $j_{\min} := \arg \min_{j' \in J} q_{j',e}^D(\theta)$, which yields

$$\begin{aligned} z_e^D(\theta + \lceil \tau_e \rceil \alpha) &\geq \int_{\theta + \lceil \tau_e \rceil \alpha}^{\theta + \lceil \tau_e \rceil \alpha + q_{j_{\min},e}^D(\theta)} g_e^-(\xi) d\xi \\ &= \int_{\theta + \lceil \tau_e \rceil \alpha}^{\theta + \lceil \tau_e \rceil \alpha + q_{j,e}^D(\theta)} g_e^-(\xi) d\xi - \int_{\theta + \lceil \tau_e \rceil \alpha + q_{j_{\min},e}^D(\theta)}^{\theta + \lceil \tau_e \rceil \alpha + q_{j,e}^D(\theta)} g_e^-(\xi) d\xi. \end{aligned} \tag{6}$$

With the same argument as before, we obtain for the right term that

$$\int_{\theta + \lceil \tau_e \rceil \alpha + q_{j_{\min},e}^D(\theta)}^{\theta + \lceil \tau_e \rceil \alpha + q_{j,e}^D(\theta)} g_e^-(\xi) d\xi \leq \alpha \cdot \kappa_e.$$

For the left term of (6), we consider time steps t with $\theta + \lceil \tau_e \rceil \alpha + \alpha \leq t\alpha < \theta + \lceil \tau_e \rceil \alpha + q_{j,e}^D(\theta)$. The number of time steps in this interval (of length $q_{j,e}^D(\theta) - \alpha$) are bounded from below by $q_{j,e}^D(\theta)/\alpha - 1$. Furthermore, in all these time steps, the set of leaving packets $L_e^-(t)$ has maximal size, that is, $|L_e^-(t)| = \lfloor \hat{v}_e(t) \rfloor$. This is true because packet i of commodity j that has entered at θ is in $B_e(t)$ but not in $L_e^-(t)$ (because it leaves e only at time $\lceil \theta + \lceil \tau_e \rceil \alpha + q_{j,e}^D(\theta) \rceil \alpha$). Hence, with the definition of the current capacity (see Equation (1)), the number of packets that leave the arc within this interval is bounded from below by

$$\left(\frac{q_{j,e}^D(\theta)}{\alpha} - 1\right) \cdot \hat{v}_e - 1.$$

Here, the last -1 is due to the rounding in the last time step of the interval. We obtain that

$$\int_{\theta + \lceil \tau_e \rceil \alpha}^{\theta + \lceil \tau_e \rceil \alpha + q_{j,e}^D(\theta)} g_e^-(\xi) d\xi \geq \beta \cdot \left(\frac{q_{j,e}^D(\theta)}{\alpha} - 1\right) \cdot \hat{v}_e - \beta = q_{j,e}^D(\theta) \cdot v_e - \alpha \cdot v_e - \beta.$$

In conclusion, (6) yields that

$$z_e^D(\theta + \lceil \tau_e \rceil \alpha) \geq q_{j,e}^D(\theta) \cdot v_e - \alpha \cdot (\kappa_e + v_e) - \beta.$$

Dividing the upper and the lower bound by v_e results in

$$-\alpha \frac{\kappa_e}{v_e} - \alpha - \frac{\beta}{v_e} \leq \frac{z_e^D(\theta + \lceil \tau_e \rceil \alpha)}{v_e} - q_{j,e}^D(\theta) \leq \alpha \frac{\kappa_e}{v_e} + 2\alpha + \frac{\beta}{v_e},$$

which completes the proof. \square

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 $h_{j,e}^- : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ 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 $H_{j,e}^+(\theta) := \int_0^\theta h_{j,e}^+(\xi) d\xi$ and $H_{j,e}^-(\theta) := \int_0^\theta h_{j,e}^-(\xi) d\xi$. 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 functions for the hypothetical inflow rate are denoted by z_e^H , q_e^H , and T_e^H . We can even define hypothetical arrival times at u and v for all $\phi \in M_j \setminus \{0\}$ by

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

With this definition, the particles $(\beta(i-1), \beta i] \subseteq M_j$ correspond to packet $i \in N_j$. Particle 0 is a special case for which we explicitly include it in packet 1 by defining

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

Clearly, we have that $k_{j,u}(i\beta) = \ell_{j,u}^D(i)$ for all $i \in N_j$. Because the hypothetical flow on arc e follows the continuous flow dynamics, we have $k_{j,v}(\phi) = k_{j,u}(\phi) + \tau_e + q_e^H(k_{j,u}(\phi))$. 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 between the arrival times at u (hypothetical flow compared with continuous flow) is not too large for all particles arriving earlier than ϕ , then this is also true for the difference in the arrival times of particle ϕ at node v .

Lemma 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

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

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

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

Proof. The proof is structured as follows. We first bound the difference in the cumulative inflow rates (between the continuous and hypothetical flows) at node u at the entrance time; afterward, we bound the difference in the cumulative outflow rates at node v at the queue entrance time. Finally, we can bound the difference in the arrival time functions at node v .

We start the proof with the following claim. Let $\theta := \ell_{j,u}(\phi)$ and $\zeta := k_{j,u}(\phi)$.

Claim 1. It holds that

$$|F_e^+(\theta) - H_e^+(\zeta)| \leq 3\kappa_e \cdot \delta.$$

Proof of Claim 1. The continuous and hypothetical flows are symmetric when considering a single arc only. For that reason, we focus on the proof of $H_e^+(\theta) - F_e^+(\zeta) \leq 3\kappa_e \cdot \delta$.

First, note that F_e^+ and H_e^+ are both nondecreasing functions whose slopes are bounded by κ_e . For each commodity $j' \in J_e$ it holds that $F_{j',e}^+(\ell_{j',u}(\varphi)) = \varphi = H_{j',e}^+(k_{j',u}(\varphi))$ for all $\varphi \in M_{j'}$. We partition J_e into two sets:

$$J_1 := \{j' \in J_e \mid H_{j',e}^+(\zeta) = 0\} \quad \text{and} \quad J_2 := \{j' \in J_e \mid H_{j',e}^+(\zeta) > 0\}.$$

In other words, J_2 is the set of commodities for which flow of the first package has already arrived at u at time ζ , and J_1 are commodities for which no flow has arrived there yet. Note that, for all commodities $j' \in J_1$, it holds that $F_{j',e}^+(\theta) \geq 0 = H_{j',e}^+(\zeta)$.

For all $j' \in J_2$, we denote by $\phi_{j'}$ the last particle of $M_{j'}$ with $\phi_{j'} \in \beta\mathbb{N}$ and $k_{j',u}(\phi_{j'}) \leq \zeta$. Because of the requirements of the lemma applied to ϕ and $\phi_{j'}$, the arrival time of this particle $\phi_{j'}$ in the continuous flow satisfies

$$\theta \geq \zeta - \delta \geq k_{j',u}(\phi_{j'}) - \delta \geq \ell_{j',u}(\phi_{j'}) - 2\delta.$$

Combining these observations, yields

$$\begin{aligned} F_e^+(\theta) &= \sum_{j' \in J_1} F_{j',e}^+(\theta) + \sum_{j' \in J_2} F_{j',e}^+(\theta) \\ &\geq \sum_{j' \in J_1} H_{j',e}^+(\zeta) + \sum_{j' \in J_2} F_{j',e}^+(\ell_{j',u}(\phi_{j'})) - 2\delta\kappa_e \\ &= \sum_{j' \in J_1} H_{j',e}^+(\zeta) + \sum_{j' \in J_2} H_{j',e}^+(k_{j',u}(\phi_{j'})) - 2\delta\kappa_e \\ &\geq \sum_{j' \in J} H_{j',e}^+(\zeta) - 2\delta\kappa_e - \alpha\kappa_e \geq \sum_{j' \in J} H_{j',e}^+(\zeta) - 3\delta\kappa_e. \end{aligned}$$

The first inequality holds because the slope of the total inflow F_e^+ is bounded by κ_e and for each commodity $j' \in J_2$ the difference $\ell_{j',u}(\phi_{j'}) - \theta$ is upper bounded by 2δ . Finally, note that the last inequality holds because $k_{j',u}(\phi_{j'})$ and ζ correspond to the same time step or no flow of commodity j' enters e between these points in time.

Equivalently, by replacing ζ by θ , k by ℓ , and F^+ and by H^+ , we obtain

$$H_e^+(\zeta) \geq F_e^+(\theta) - 3\delta\kappa_e. \quad \square$$

Because the deviation of cumulative inflow is bounded, we can derive a bound on the difference in the cumulative outflow functions at the time the particle ϕ enters the respective queues.

Claim 2. *We have*

$$|F_e^-(\theta + \tau_e) - H_e^-(\zeta + \tau_e)| \leq 7\delta\kappa_e.$$

Proof of Claim 2. Assume for contradiction that we have $F_e^-(\theta + \tau_e) - H_e^-(\zeta + \tau_e) > 7\delta\kappa_e$.

Let ζ' be the latest point in time before ζ such that there was no queue at time $\zeta' + \tau_e$, that is,

$$\zeta' := \max\{\xi \leq \zeta \mid H_e^+(\xi) - H_e^-(\xi + \tau_e) = 0\}.$$

Note that this condition is satisfied for $\xi = 0$ because $H_e^+(0) = H_e^-(\tau_e) = 0$. Because either $\zeta' = \zeta$ or a queue builds up immediately after ζ' (implying H_e^+ is strictly increasing), there has to be a commodity j' and a particle $\phi' \in \beta\mathbb{N}$ with $k_{j',u}(\phi') \in [\zeta', \zeta' + \alpha]$. Note that this immediately implies that

$$H_e^+(k_{j',u}(\phi')) - H_e^+(\zeta') \leq \alpha\kappa_e \leq \delta\kappa_e.$$

Because there is no queue for ζ' , it holds that $H_e^-(\zeta' + \tau_e) = H_e^+(\zeta')$. Furthermore, within $(\zeta' + \tau_e, \zeta + \tau_e]$, there is a queue present on e , and thus, we have $h_e^-(\xi) = v_e$ within this interval. Hence,

$$H_e^-(\zeta + \tau_e) = H_e^-(\zeta' + \tau_e) + v_e \cdot (\zeta - \zeta').$$

To proceed, we distinguish two cases.

Case 1. $\theta \geq \zeta' + \delta$.

Because queues can never be negative ($F_e^+(\vartheta) - F_e^-(\vartheta + \tau_e) \geq 0$ for all ϑ) and $\zeta \in [\theta - \delta, \theta + \delta]$, we obtain

$$F_e^+(\zeta' + \delta) \geq F_e^-(\zeta' + \delta + \tau_e) \geq F_e^-(\theta + \tau_e) - v_e \cdot (\theta - (\zeta' + \delta)) \geq F_e^-(\theta + \tau_e) - v_e \cdot (\zeta - \zeta').$$

Note for the second inequality that the slope of F_e^- is bounded by v_e .

By the requirement of the lemma, it follows that $\ell_{j',u}(\phi') \in [\zeta' - \delta, \zeta' + 2\delta]$, and because F_e^+ cannot grow faster than κ_e , we obtain

$$|F_e^+(\zeta' + \delta) - F_e^+(\ell_{j',u}(\phi'))| \leq 2\delta\kappa_e.$$

Putting it all together, we get

$$\begin{aligned} F_e^+(\ell_{j',u}(\phi')) + 3\delta\kappa_e - H_e^+(k_{j',u}(\phi')) &\geq F_e^+(\zeta' + \delta) - H_e^+(\zeta') \\ &\geq F_e^-(\theta + \tau_e) - v_e \cdot (\zeta - \zeta') - H_e^-(\zeta' + \tau_e) \\ &= F_e^-(\theta + \tau_e) - v_e \cdot (\zeta - \zeta') - (H_e^-(\zeta + \tau_e) - v_e \cdot (\zeta - \zeta')) \\ &= F_e^-(\theta + \tau_e) - H_e^-(\zeta + \tau_e) \\ &> 7\delta\kappa_e. \end{aligned}$$

This is a contradiction to Claim 1 (applied to particle ϕ' of commodity j' , which satisfies the requirement of Lemma 3).

Case 2. $\theta < \zeta' + \delta$.

By the requirement of the lemma, we have $\zeta' \leq \zeta \leq \theta + \delta$, which gives us $|\theta - \zeta'| \leq \delta$. Using this and the fact that queues cannot be negative yields

$$F_e^+(\zeta') \geq F_e^-(\zeta' + \tau_e) \geq F_e^-(\theta - \delta + \tau_e) \geq F_e^-(\theta + \tau_e) - v_e\delta.$$

By the requirement of the lemma, it follows that $\ell_{j',u}(\phi') \in [\zeta' - \delta, \zeta' + 2\delta]$, and because F_e^+ cannot grow faster than κ_e , we obtain

$$|F_e^+(\zeta') - F_e^+(\ell_{j',u}(\phi'))| \leq 2\delta\kappa_e.$$

Again, putting everything together gives

$$\begin{aligned} F_e^+(\ell_{j,u}(\phi')) + 3\delta\kappa_e - H_e^+(k_{j,u}(\phi')) &\geq F_e^+(\zeta') - H_e^+(\zeta') \\ &= F_e^+(\zeta') - H_e^-(\zeta' + \tau_e) \\ &\geq F_e^-(\theta + \tau_e) - \nu_e\delta - H_e^-(\zeta + \tau_e) \\ &> 7\delta\kappa_e - \nu_e\delta \\ &\geq 6\delta\kappa_e. \end{aligned}$$

This again contradicts Claim 1. Hence, in both cases, we obtain

$$F_e^-(\theta + \tau_e) - H_e^-(\zeta + \tau_e) \leq 7\delta\kappa_e.$$

The assumption $H_e^-(k_{j,u}(\phi) + \tau_e) - F_e^-(\ell_{j,u}(\phi) + \tau_e) > 7\delta\kappa_e$ can lead to a contradiction in the same way by using the symmetry of flow F and H . \square

We can now use Claims 1 and 2 to finish the proof of the lemma. The difference in the arrival times of particle $\phi \in \beta\mathbb{N}$ at node v in the continuous and hypothetical flows can be bounded as follows:

$$\begin{aligned} |\ell_{j,v}(\phi) - k_{j,v}(\phi)| &= |\ell_{j,u}(\phi) + \tau_e + q_{j,e}(\ell_{j,u}(\phi)) - (k_{j,u}(\phi) + \tau_e + q_{e,j}^H(k_{j,u}(\phi)))| \\ &\leq \delta + |q_{j,e}(\ell_{j,u}(\phi)) - q_{e,j}^H(k_{j,u}(\phi))| \\ &= \delta + \left| \frac{1}{\nu_e} \cdot (F_e^+(\ell_{j,u}(\phi)) - F_e^-(\ell_{j,u}(\phi) + \tau_e) - H_e^+(k_{j,u}(\phi)) + H_e^-(k_{j,u}(\phi) + \tau_e)) \right| \\ &\leq \delta + \frac{1}{\nu_e}(3\delta\kappa_e + 7\delta\kappa_e) \\ &\leq 11\delta \frac{\kappa_e}{\nu_e}. \quad \square \end{aligned}$$

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

Lemma 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*

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

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

Proof. This lemma bounds the error between the arrival times in the packet model and the continuous model on a single arc. To do so, we introduce the hypothetical flow, which has exactly the same inflow functions as the packet routing. In the following, we bound the difference between the arrival time of a packet and the corresponding particle entering arc e at the same time.

As we only consider a single arc on which we have FIFO in both models, we suppose for the beginning, that all particles and packets belong to a single commodity with particles $\phi \in M := [0, m] \subset \mathbb{R}_{\geq 0}$ with $m := \sum_{j \in J} m_j$ and packets $i \in N := \{1, 2, \dots, \sum_{j \in J} N_j\}$, respectively. In other words, we focus on the total flow.

We define the arrival times at u and v by

$$\begin{aligned} k_u(\phi) &:= \min\{\theta \in \mathbb{R}_{\geq 0} \mid H_e^+(\theta) = \phi\}, & k_v(\phi) &:= \min\{\theta \in \mathbb{R}_{\geq 0} \mid H_e^-(\theta) = \phi\}, \\ \ell_u^D(i) &:= \min\{\theta \in \mathbb{R}_{\geq 0} \mid G_e^+(\theta) = i\beta\}, & \ell_v^D(i) &:= \min\{\theta \in \mathbb{R}_{\geq 0} \mid G_e^-(\theta) = i\beta\}. \end{aligned}$$

Clearly, as we have $H_e^+ = G_e^+$, it holds that $k_u(i\beta) = \ell_u^D(i)$ for all $i \in N$.

From now on, we consider a particle–packet pair ϕ and $i := \lceil \frac{\phi}{\beta} \rceil$. We show by induction over the particles that the following holds.

Claim 3. *There exist constants \tilde{C}_1 and \tilde{C}_2 that only depend on the instance properties such that, for all $\phi \in M$ with $i = \lceil \frac{\phi}{\beta} \rceil \in N$, it holds that*

$$|k_v(\phi) - \ell_v^D(i)| \leq \delta(\phi) \quad \text{with} \quad \delta(\phi) \leq \tilde{C}_1 \cdot \sqrt{\alpha} + \phi \cdot \tilde{C}_2 \cdot \sqrt{\alpha}.$$

Proof of Claim 3. The proof idea is to do an induction over the particles $M \subseteq \mathbb{R}_{\geq 0}$. In every induction step, we extend the statement for a flow volume of at least $\sqrt{\alpha} - \beta$. The major part of the proof deals with the induction step, and the induction start is shown at the very end.

Throughout the proof, we need some particular particles and packets, which we denote as follows. Let ϕ' be the particle that leaves e at the moment when particle ϕ enters the queue. Formally, ϕ' is a unique particle with $k_v(\phi') = k_u(\phi) + \tau_e$. Note that, whenever $z_e^H(k_u(\phi) + \tau_e) > 0$, there has to be a particle leaving the queue of e at that very moment. Otherwise, $\phi' = \phi$. Let i' be the particle corresponding to ϕ' , formally $i' := \lceil \frac{\phi'}{\beta} \rceil$.

Let, furthermore, j be the last packet that left the queue before packet i lines up in the packet routing; in other words, $j := \max\{j' \leq i \mid \ell_v^D(j') \leq \ell_u^D(i) + \lceil \tau_e \rceil_\alpha\}$. This set is not empty as it contains packet 1. We denote the corresponding particle by $\psi := j \cdot \beta$.

Overall, we have the following correspondences between packets and particles:

$$i \leftrightarrow \phi \quad i' \leftrightarrow \phi' \quad j \leftrightarrow \psi.$$

In order to bound the difference between $k_v(\phi)$ and $\ell_v^D(i)$, we carefully analyze the arrival times of these particles and packets.

First of all, we assume that $\frac{\beta}{\alpha}$ is small enough such that there is at least one packet allowed to leave e at each time step, that is, $\hat{v}_e \geq 1$. This way, we ensure that

$$|\ell_u^D(i) + \lceil \tau_e \rceil_\alpha - \ell_v^D(j)| \leq \alpha. \tag{7}$$

Furthermore, because $G_e^+ = H_e^+$, we have that $|k_u(\phi) - \ell_u^D(i)| \leq \alpha$, and hence,

$$|k_v(\phi') - \ell_v^D(j)| \stackrel{(7)}{\leq} |k_u(\phi) + \tau_e - (\ell_u^D(i) + \lceil \tau_e \rceil_\alpha)| + \alpha \leq 3\alpha. \tag{8}$$

In the following, we distinguish three cases in dependency of the queue lengths at the time particle ϕ or packet i , respectively, queues up:

1. The continuous and discrete queues are small: $z_e^H(k_u(\phi) + \tau_e) \leq \sqrt{\alpha}$ and $z_e^D(\ell_u^D(i) + \lceil \tau_e \rceil_\alpha) \leq \sqrt{\alpha}$.
2. The continuous queue is larger than $\sqrt{\alpha}$ and larger than the discrete queue: $z_e^H(k_u(\phi) + \tau_e) > \sqrt{\alpha}$ and $z_e^D(\ell_u^D(i) + \lceil \tau_e \rceil_\alpha) \leq z_e^H(k_u(\phi) + \tau_e)$.
3. The discrete queue is larger than $\sqrt{\alpha}$ and larger than the continuous queue: $z_e^D(\ell_u^D(i) + \lceil \tau_e \rceil_\alpha) > \sqrt{\alpha}$ and $z_e^D(\ell_u^D(i) + \lceil \tau_e \rceil_\alpha) > z_e^H(k_u(\phi) + \tau_e)$.

Case 1. $z_e^H(k_u(\phi) + \tau_e) \leq \sqrt{\alpha}$ and $z_e^D(\ell_u^D(i) + \lceil \tau_e \rceil_\alpha) \leq \sqrt{\alpha}$.

Because both queues are upper bounded by $\sqrt{\alpha}$, we obtain

$$k_v(\phi) \in \left[k_v(\phi'), k_v(\phi') + \frac{\sqrt{\alpha}}{v_e} \right].$$

With Lemma 2 and (7), we obtain

$$\ell_v^D(i) \in \left[\ell_v^D(j), \ell_v^D(j) + \frac{\sqrt{\alpha}}{v_e} + 3\alpha + \frac{\kappa_e \alpha}{v_e} + \frac{\beta}{v_e} \right].$$

If we put this together with (8), it follows that

$$|k_v(\phi) - \ell_v^D(i)| \leq \frac{\sqrt{\alpha}}{v_e} + 6\alpha + \frac{\kappa_e \alpha}{v_e} + \frac{\beta}{v_e}.$$

Case 2. $z_e^H(k_u(\phi) + \tau_e) > \sqrt{\alpha}$ and $z_e^D(\ell_u^D(i) + \lceil \tau_e \rceil_\alpha) \leq z_e^H(k_u(\phi) + \tau_e)$.

For the sake of simplicity, we use the intuitive notation and write $x \pm y$ for the interval $[x - y, x + y]$.

As a first step, notice that

$$k_v(\phi) = k_v(\phi') + \frac{(\phi - \phi')}{v_e} \in k_v(\phi') + \frac{\beta(i - i')}{v_e} \pm \frac{\beta}{v_e},$$

And similar to Lemma 2,

$$\ell_v^D(i) \in \ell_v^D(j) + q_e^D(\ell_u^D(i)) \pm \alpha \subseteq \ell_v^D(j) + \frac{\beta(i - j)}{v_e} \pm \left(3\alpha + \frac{\kappa_e \alpha}{v_e} + \frac{\beta}{v_e} \right).$$

This gives us that

$$\begin{aligned} |k_v(\phi) - \ell_v^D(i)| &\leq \left| k_v(\phi') - \ell_v^D(j) + \frac{\beta((i-i') - (i-j))}{v_e} \right| + 3\alpha + \frac{\kappa_e \alpha}{v_e} + \frac{2\beta}{v_e} \\ &= \left| k_v(\phi') - \ell_v^D(j) + \frac{\beta(j-i')}{v_e} \right| + 3\alpha + \frac{\kappa_e \alpha}{v_e} + \frac{2\beta}{v_e}. \end{aligned} \tag{9}$$

Furthermore, in the case of $j \geq i'$, we observe that the difference between the arrival times at node v from packet j and i' can be bounded from below by the time the packets between them necessarily need to leave the arc e . Hence, we apply Lemma 2 on a fictional packet routing in which packet i' leaves the queue exactly at the time packet j enters the queue (consisting of all particles between i' and j). As the time between i' and j arriving at v in the actual packet routing is at least that large, we obtain

$$\ell_v^D(j) - \ell_v^D(i') \geq \left| \frac{\beta(j-i')}{v_e} \right| - \left(2\alpha + \frac{\kappa_e \alpha}{v_e} + \frac{\beta}{v_e} \right). \tag{10}$$

In the case that $i' > j$, we have $i' \leq j + 2$ as we explain in the following. We have

$$(i' - 1)\beta \leq \phi' = \phi - z_e^H(k_u(\phi) + \tau_e) \leq i\beta - z_e^H(k_u(\phi) + \tau_e)$$

and also

$$z_e^D(\ell_u^D(i) + \lceil \tau_e \rceil_\alpha) = G_e^+(\ell_u^D(i)) - G_e^-(\ell_u^D(i) + \lceil \tau_e \rceil_\alpha) \geq i\beta - (j+1)\beta.$$

Hence, we obtain with the assumption of Case 4 that

$$\begin{aligned} (i' - j)\beta &\leq i\beta - z_e^H(k_u(\phi) + \tau_e) + \beta - (i\beta - z_e^D(\ell_u^D(i) + \lceil \tau_e \rceil_\alpha) - \beta) \\ &= z_e^D(\ell_u^D(i) + \lceil \tau_e \rceil_\alpha) - z_e^H(k_u(\phi) + \tau_e) + 2\beta \\ &\leq 2\beta. \end{aligned}$$

In other words, there is at most one additional packet between i' and j . Furthermore, it holds that $i' \in [j+1, i]$, and thus, i' is in a queue at time $\ell_u^D(i) + \lceil \tau_e \rceil_\alpha$. Assuming that $\hat{v}_e \geq 2$ (true for sufficiently small $\frac{\beta}{\alpha}$) yields $\ell_v^D(j) - \ell_v^D(i') \geq -\alpha$. With

$$\left| \frac{(j-i')\beta}{v_e} \right| \leq \frac{2\beta}{v_e} \leq \frac{\alpha \kappa_e}{v_e} + \frac{\beta}{v_e},$$

Equation (10) also holds in this case of $i' > j$. Finally, combining all observations, we obtain

$$\begin{aligned} &|k_v(\phi) - \ell_v^D(i)| \\ &\stackrel{(9)}{\leq} \max \left\{ \ell_v^D(j) - k_v(\phi') + \frac{2\beta}{v_e}, k_v(\phi') - \ell_v^D(j) + \left| \frac{\beta(j-i')}{v_e} \right| \right\} + 3\alpha + \frac{\kappa_e \alpha}{v_e} + \frac{2\beta}{v_e} \\ &\stackrel{(10)}{\leq} \max \left\{ \ell_v^D(j) - k_v(\phi') + \frac{2\beta}{v_e}, k_v(\phi') - \ell_v^D(i') + 2\alpha + \frac{\alpha \kappa_e}{v_e} + \frac{\beta}{v_e} \right\} + 3\alpha + \frac{\kappa_e \alpha}{v_e} + \frac{2\beta}{v_e} \\ &\stackrel{(8)}{\leq} \max \left\{ 3\alpha, k_v(\phi') - \ell_v^D(i') + 2\alpha + \frac{\alpha \kappa_e}{v_e} \right\} + 3\alpha + \frac{\kappa_e \alpha}{v_e} + \frac{4\beta}{v_e} \\ &\leq \delta(\phi - \sqrt{\alpha}) + 6\alpha + \frac{2\kappa_e \alpha}{v_e} + \frac{4\beta}{v_e}. \end{aligned}$$

Here, the last inequality follows by induction because $\phi - \sqrt{\alpha} \geq \phi'$ (case assumption) and by monotonicity of δ .

Case 3. $z_e^D(\ell_u^D(i) + \lceil \tau_e \rceil_\alpha) > \sqrt{\alpha}$ and $z_e^D(\ell_u^D(i) + \lceil \tau_e \rceil_\alpha) > z_e^H(k_u(\phi) + \tau_e)$.

We begin with

$$k_v(\phi) = k_v(\phi') + \frac{(\phi - \phi')}{v_e}.$$

Lemma 2 and (7) provide

$$\begin{aligned} \ell_v^D(i) &= \ell_u^D(i) + \lceil \tau_e \rceil_\alpha + q_e^D(\ell_u^D(i)) \\ &\in \ell_v^D(j) + \frac{z_e^D(\ell_u^D(i) + \lceil \tau_e \rceil_\alpha)}{v_e} \pm \left(3\alpha + \frac{\kappa_e \alpha}{v_e} + \frac{\beta}{v_e} \right) \\ &\subseteq \ell_v^D(j) + \beta \cdot \frac{i-j}{v_e} \pm \left(3\alpha + \frac{\kappa_e \alpha}{v_e} + \frac{2\beta}{v_e} \right) \\ &\subseteq \ell_v^D(j) + \frac{\phi - \psi}{v_e} \pm \left(3\alpha + \frac{\kappa_e \alpha}{v_e} + \frac{3\beta}{v_e} \right). \end{aligned}$$

This gives us that

$$|\ell_v^D(i) - k_v(\phi)| \leq \left| \ell_v^D(j) - k_v(\phi') + \frac{(\phi - \psi) - (\phi - \phi')}{v_e} \right| + \left(3\alpha + \frac{\kappa_e \alpha}{v_e} + \frac{3\beta}{v_e} \right). \quad (11)$$

The case assumption states that the discrete queue is longer than the hypothetical queue. Together with $\phi' = \phi - z_e^H(k_u(\phi) + \tau_e)$ and $\psi = j\beta \leq i\beta - z_e^D(\ell_u^D(i) + \lceil \tau_e \rceil_\alpha)$, this implies

$$\phi' - \psi \geq z_e^D(\ell_u^D(i) + \lceil \tau_e \rceil_\alpha) - z_e^H(k_u(\phi) + \tau_e) + \phi - i\beta > -\beta. \quad (12)$$

Furthermore, because queues in the continuous flow model cannot operate faster than v_e , we have for $\phi' \geq \psi$ that

$$k_v(\phi') - k_v(\psi) \geq \frac{\phi' - \psi}{v_e}. \quad (13)$$

For $\phi' < \psi$, there is a queue for all particles in $[\phi, \phi'] \ni \psi$. Thus, (13) holds with equality. Combining these inequalities, we obtain

$$\begin{aligned} & |\ell_v^D(i) - k_v(\phi)| \\ & \stackrel{(11),(12)}{\leq} \max \left\{ k_v(\phi') - \ell_v^D(j) + \frac{\beta}{v_e}, \ell_v^D(j) - k_v(\phi') + \frac{(\phi' - \psi)}{v_e} \right\} + 3\alpha + \frac{\kappa_e \alpha}{v_e} + \frac{3\beta}{v_e} \\ & \stackrel{(13)}{\leq} \max \left\{ k_v(\phi') - \ell_v^D(j) + \frac{\beta}{v_e}, \ell_v^D(j) - k_v(\phi') + k_v(\phi') - k_v(\psi) \right\} + 3\alpha + \frac{\kappa_e \alpha}{v_e} + \frac{3\beta}{v_e} \\ & \stackrel{(8)}{\leq} \max \left\{ 3\alpha + \frac{\beta}{v_e}, \ell_v^D(j) - k_v(\psi) \right\} + 3\alpha + \frac{\kappa_e \alpha}{v_e} + \frac{3\beta}{v_e} \\ & \leq \delta(\phi - \sqrt{\alpha} + \beta) + 6\alpha + \frac{\kappa_e \alpha}{v_e} + \frac{4\beta}{v_e}. \end{aligned}$$

The last inequality holds by monotonicity of $\delta(\cdot)$ and because we have $\psi = j\beta \leq i\beta - \sqrt{\alpha} \leq \phi + \beta - \sqrt{\alpha}$, which holds by the case assumption as the discrete queue is longer than $\sqrt{\alpha}$.

Case Conclusion. To conclude the case distinction and prove the induction step, we need a monotonically increasing function $\delta : M \rightarrow \mathbb{R}_{>0}$ that bounds $|\ell_v^D(i) - k_v(\phi)|$, that is, satisfies

- $\delta(\phi) \geq 6\alpha + \frac{\sqrt{\alpha}}{v_e} + \frac{\kappa_e \alpha}{v_e} + \frac{\beta}{v_e}$.
- $\delta(\phi) \geq \delta(\phi - \sqrt{\alpha}) + 6\alpha + \frac{2\kappa_e \alpha}{v_e} + \frac{4\beta}{v_e}$.
- $\delta(\phi) \geq \delta(\phi - \sqrt{\alpha} + \beta) + 6\alpha + \frac{\kappa_e \alpha}{v_e} + \frac{4\beta}{v_e}$.
- $\delta(\phi) \rightarrow 0$ for $\alpha \rightarrow 0$ and $\frac{\beta}{\alpha} \rightarrow 0$ for all $\phi \in M$.

The function

$$\delta(\phi) = 6\alpha + \frac{\sqrt{\alpha}}{v_e} + \frac{\kappa_e \alpha}{v_e} + \frac{2\beta}{v_e} + \frac{2\phi}{\sqrt{\alpha}} \left(6\alpha + \frac{2\kappa_e \alpha}{v_e} + \frac{4\beta}{v_e} \right) \leq \tilde{C}_1 \cdot \sqrt{\alpha} + \phi \cdot \tilde{C}_2 \cdot \sqrt{\alpha},$$

fulfills all requirements, in which \tilde{C}_1 and \tilde{C}_2 are constants that only depend on instance properties, namely, v_e and κ_e .

It remains to show the base case, which consists of all particles in $\phi \in [0, \sqrt{\alpha}]$. In this case, the total flow volume of packets on the arc and, hence, in the queue, is upper bounded by $\sqrt{\alpha} + \beta$ (note that for $\phi = \sqrt{\alpha}$, the packet i might be rounded up to at most $\sqrt{\alpha}/\beta + 1$). Applying the argumentation used in Case 1 to the situation in which queues are bounded by $\sqrt{\alpha} + \beta$ in both models yields

$$|\ell_v^D(i) - k_v(\phi)| \leq 6\alpha + \frac{\sqrt{\alpha}}{v_e} + \frac{\kappa_e \alpha}{v_e} + \frac{2\beta}{v_e} \leq \delta(\phi).$$

Our choice of δ satisfies this.

This finishes the proof of the claim. Thus, we have shown that, in a single-commodity setting, the difference in the arrival times at node v between a particle and the corresponding packet is upper bounded by δ , which goes to zero when the discretization parameters go to zero. \square

To prove the lemma, two things remain to show. First, we reassign the commodities to the packets, and second, we switch from a particle- to a time-dependent bound.

Given a particle $\phi_j \in M_j$ and corresponding packet $i_j := \lceil \frac{\phi_j}{\beta} \rceil \in N_j$, let $\vartheta := k_{j,m}(\phi_j)$ be the point in time that the particle enters e . We consider the same particle in the single commodity flow setting and denote it by ϕ . In other

words, ϕ is the particle with $k_u(\phi) = \vartheta$. The corresponding packet is $i := \lceil \frac{\phi}{\beta} \rceil$. Because the flow dynamics and the inflow rates of the single-commodity hypothetical flow equal the multicommodity hypothetical flow, we have

$$k_v(\phi) = T_e^H(k_u(\phi)) = T_e^H(\vartheta) = T_e^H(k_{j,u}(\phi_j)) = k_{j,v}(\phi_j).$$

In the following, we argue that the arrival times at node v of packet i and i_j are not too far from each other.

First, observe that packet i_j and i enter e at the same time step, which is given by $\lceil \vartheta \rceil_\alpha$. Though, the refined entrance times (which are commodity-dependent) might differ up to α , which might cause a small error for the exit times. The packet volume that enters between the refined entrance times of these two packets is bounded by $\kappa_e \alpha$. Because packets i and i_j enter the arc at the same time step, they are also assigned to Q_e at the same time step. This implies that, in the time steps between the exit times of the two packets, the capacity of the arc is fully used, that is, at least $\lfloor v_e \rfloor_\beta$ volume leaves per time step. Hence, the time steps when i and i_j leave the arc cannot be further away than $\frac{\kappa_e \alpha}{\lfloor v_e \rfloor_\beta} + \alpha$. Note that the additive α is due to the rounding to time steps. For the refined times, there might be an additional deviation of α , which gives us

$$|\ell_{j,v}^D(i_j) - \ell_v^D(i)| \leq \frac{\kappa_e \alpha}{\lfloor v_e \rfloor_\beta} + 2\alpha.$$

These considerations together with Claim 3 lead to

$$\begin{aligned} |\ell_{j,v}^D(i_j) - k_{j,v}(\phi_j)| &\leq |\ell_{j,v}^D(i_j) - \ell_v^D(i)| + |\ell_v^D(i) - k_v(\phi)| + |k_v(\phi) - k_{j,v}(\phi_j)| \\ &\leq \frac{\kappa_e \alpha}{\lfloor v_e \rfloor_\beta} + 2\alpha + \tilde{C}_1 \cdot \sqrt{\alpha} + \phi \cdot \tilde{C}_2 \cdot \sqrt{\alpha} \\ &\leq C_1 \cdot \sqrt{\alpha} + \phi \cdot \tilde{C}_2 \cdot \sqrt{\alpha} \end{aligned}$$

with $C_1 := \frac{2\kappa_e}{v_e} + 2 + \tilde{C}_1$. Note that the last inequality only holds for sufficiently small β , which fulfills $\lfloor v_e \rfloor_\beta \geq v_e/2$.

In order to switch from particle ϕ to entrance time ϑ , we use the fact that, until time ϑ , at most a total flow of $\kappa_e \cdot \vartheta$ entered arc e .

By choosing $C_2 := \kappa_e \cdot \tilde{C}_2$, we obtain

$$\begin{aligned} |\ell_{j,v}^D(i_j) - k_{j,v}(\phi_j)| &\leq C_1 \cdot \sqrt{\alpha} + \phi \cdot \tilde{C}_2 \cdot \sqrt{\alpha} \\ &\leq C_1 \cdot \sqrt{\alpha} + \vartheta \cdot \kappa_e \cdot \frac{C_2}{\kappa_e} \cdot \sqrt{\alpha} \\ &= C_1 \cdot \sqrt{\alpha} + \vartheta \cdot C_2 \cdot \sqrt{\alpha}. \end{aligned}$$

Because $\vartheta = k_{j,u}(\phi_j) \leq \theta$, this completes the proof of the lemma. \square

With the help of the previous lemmas, we can finally prove Theorem 1.

Proof of Theorem 1. We prove this theorem by an induction on K . Because of the finitely lasting network inflow rates, there is a finite time horizon H , for which the last particle has left the network. It is sufficient to prove the statement only for $K \in \{1, 2, \dots, \hat{H}\}$ with $\hat{H} := \lceil H/\tau^* \rceil$ as they cover all arrival times at all nodes for all particles. For larger values of K , the right-hand side of the bound only grows because $C > 1$, so the statement remains valid.

We choose α and β to be small enough such that $\sqrt{\alpha} < \tau^* C^{-\hat{H}}$ and $\beta/\alpha < 1$ as well as $v_e \cdot \alpha/\beta \geq 2$ for all $e \in E$. (The last two conditions are needed for Lemma 3.)

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 because, 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 $\ell_{j,v}^D(i) \leq \theta$.

For the release times, we get

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

For the induction step to $K \leq \hat{H}$, 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 $\min\{\ell_{j,v}(\phi), \ell_{j,v}^D(i)\} < K\tau^*$. 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 $\min\{\ell_{j,u}(\phi), \ell_{j,u}^D(i)\} \leq \min\{\ell_{j,v}(\phi),$

$\ell_{j,v}^D(i) - \tau_e \leq (K-1)\tau^*$, and therefore, we can apply the induction hypothesis to node u to obtain

$$\max\{\ell_{j,u}(\phi), \ell_{j,u}^D(i)\} \leq \min\{\ell_{j,u}(\phi), \ell_{j,u}^D(i)\} + \frac{\tau_e}{2} \leq \min\{\ell_{j,v}(\phi), \ell_{j,v}^D(i)\} - \frac{\tau_e}{2} \leq (K-1)\tau^*. \quad (14)$$

For the first inequality, note that we use $\sqrt{\alpha} < \tau^* C^{-\hat{H}} < \tau_e/2 \cdot C^{-(K-1)}$.

By introducing the hypothetical flow on arc e , we can apply Lemma 3 with $\delta := C^{K-1}\sqrt{\alpha}$. The preconditions of the lemma are satisfied because we can apply the induction hypothesis to every $\varphi \in \beta\mathbb{N}$ (and $i' := \varphi/\beta$) with $\min\{\ell_{j,u}(\varphi), \ell_{j,u}^D(i')\} \leq \max\{\ell_{j,u}(\phi), \ell_{j,u}^D(i)\}$ thanks to (14). Note that $k_{j,u}(\phi') = \ell_{j,u}^D(i')$ for all $i' \in N_j$ and $\phi' = i'\beta \in M_j$. Lemma 3 gives us

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

We have chosen α small enough such that $C^{K-1}\sqrt{\alpha} \leq \tau^*$. Hence, again by the induction hypothesis, we obtain

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

Applying Lemma 4 with $\theta := K\tau^*$ yields

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

Finally, by choosing C large enough, we obtain

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

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

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

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

Because 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, that is, 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 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 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 2. *Let $(V, E, (v_e)_{e \in E}, (\tau_e)_{e \in E})$ 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:*

- i. *For every $v \in V$, $j \in J_v$ and $\phi \in M_j$, the arrival time $\ell_{j,v}^D(\lceil \frac{\phi}{\beta} \rceil)$ of the packet model converges to the arrival time $\ell_{j,v}(\phi)$ of the flow model.*
- ii. *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 (i), first observe that, by the continuity of $\ell_{j,v}$, we have that $\ell_{j,v}(\lceil \frac{\phi}{\beta} \rceil) \rightarrow \ell_{j,v}(\phi)$ for $\beta \rightarrow 0$. The statement then follows immediately by Theorem 1 with K large enough such that $\ell_{j,v}(\phi) \leq K\tau^*$.

For (ii), let $\varepsilon > 0$. By Theorem 1, it is possible to choose (α, β) small enough such that

$$|\ell_{j,u}^D(i) - \ell_{j,u}(i\beta)| \leq \frac{\varepsilon - \beta}{\kappa_e}$$

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 the following for the case in which 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,

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

Because the slope of $G_{j,e}^+$ is bounded by κ_e and $F_{j,e}^+(\theta) \leq i^- \beta + \beta$, we obtain

$$G_{j,e}^+(\theta) \geq G_{j,e}^+(\ell_{j,u}^D(i^-)) - \frac{\varepsilon - \beta}{\kappa_e} \cdot \kappa_e = i^- \beta - (\varepsilon - \beta) \geq F_{j,e}^+(\theta) - \varepsilon.$$

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

$$G_{j,e}^+(\theta) \leq G_{j,e}^+(\ell_{j,u}^D(i^+)) + \frac{\varepsilon - \beta}{\kappa_e} \cdot \kappa_e = i^+ \beta + \varepsilon - \beta \leq F_{j,e}^+(\theta) + \varepsilon.$$

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 $|G_{j,e}^-(\theta) - F_{j,e}^-(\theta)| \leq \varepsilon$. \square

5.3. Example

To complement this section, we visualize the convergence result in a simple example. For this, we consider the network depicted in Figure 6. The first commodity (green) sends one flow unit within $[0, 1]$ with a rate of one along the zigzag path. The second commodity (blue) sends flow from time 0 to 1 with a rate of one along the bottom path.

We compare the arrival time of particles and packets for different discretizations in Figure 7. It becomes apparent that, for smaller discretization, the arrival times approach the arrival time of the continuous flow embedding.

In detail, the continuous flow of the green commodity reaches node w half a time unit earlier than the flow of the blue commodity. Thus, the first half of the green flow (particles $[0, \frac{1}{2}]$) can use the full capacity of arc wd , and the second half splits capacity evenly with the blue commodity. Afterward, all green flow particles have left wd , and the blue commodity can use the full capacity.

In the first packet routing instance with discretization $(1, 1)$, we only consider the network at discrete time steps. This means, at time 2, both the single blue and single green packets are located at w and enter arc wd . Because the arc of the blue packet is preferred in a tie break, the blue packet enters arc wd first, followed by the green packet (note that both packets enter the arc at the same time step; still the ordering of the packets is essential). Now, to leave the arc and to arrive at node d , a capacity of one is needed. For that reason, capacity has to be accumulated at wd for one time step before the blue packet can leave the arc at time 4. Again, capacity needs to be accumulated for one time step until the green packet leaves the arc at time 6.

Already in the packet-routing instance with discretization $(\frac{1}{2}, \frac{1}{4})$, the packet arrival times become much more similar to the arrival times of the continuous model. First of all, because time steps have a length of $\frac{1}{2}$, the difference in the travel times from o to w along the two paths becomes noticeable, and the first two green packets enter arc wd before the first blue packet arrives at w . Second, a capacity of $\frac{1}{2}$ allows exactly one packet of size $\frac{1}{4}$ to leave every time step of length $\frac{1}{2}$, which makes the outflow of arc wd smoother.

Figure 6. In the network, the arc properties are given in the form (τ, ν) . At w , arc ow is preferred over arc vw in a tie break. We consider two commodities, one sending flow along the zigzag path (o, v, w, d) and one along the bottom path (o, w, d) . Both commodities send flow with a rate of one in $[0, 1]$.





Figure 7. The arrival times at the destination d of the particles/packets of the different commodities (green for the first commodity along the zigzag path and blue for the second on the bottom path). The solid lines represent the arrival time functions of the continuous model. On the left, the packet size and the time step size are equal to one. In the middle the discretization is given by $\alpha = \frac{1}{2}, \beta = \frac{1}{4}$ and on the right by $\alpha = \frac{1}{4}, \beta = \frac{1}{16}$. The horizontal axis shows the indices of the packets. Overall, the packets always sum up to a flow volume of one. The vertical axis corresponds to the arrival times.



In the packet-routing instance with discretization $(\frac{1}{4}, \frac{1}{16})$, we have a very similar situation as in the continuous flow. The first half of the green packets can use arc wd first. Afterward, capacity is split evenly until the last green packet has arrived at d , and the blue packets can use the full capacity for themselves.

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 multicommodity setting, in which exact pure Nash equilibria do not exist in general as is shown in Theorem 1.

Theorem 3. Consider a network $(V, E, (v_e)_{e \in E}, (\tau_e)_{e \in E})$, 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 $u_j : [0, T] \rightarrow \mathbb{R}_{\geq 0}$. 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 multicommodity Nash flow over time f , which exists because of Cominetti et al. [5, 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. Cominetti et al. [5, section 2.7]). Additionally, Cominetti et al. [5, lemma 7] show 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

$$|\ell_{P,d_P}(\phi) - \tilde{\ell}_{P,d_P}(\phi)| \leq \frac{\varepsilon}{4} \quad \text{for all particles } \phi \text{ and all paths } P \in \mathcal{P}. \tag{15}$$

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 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

$$|\ell_{P,d_P}(i\beta) - \ell_{P,d_P}^D(i)| \leq \frac{\varepsilon}{4} \quad \text{for all } P \in \mathcal{P}.$$

Additionally, we require β to be small enough such that (15) 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 arrives at least ε earlier by switching from its path P to some other $o_P\text{-}d_P$ 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 $\tilde{\ell}^D$.

Downloaded from informs.org by [88.65.1.17] on 31 August 2023, at 01:43. For personal use only, all rights reserved.

Again, with Theorem 1, we obtain that

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

Note that the constant C in Theorem 1 only depends on the instance. By choosing C large enough, it is possible to shift the flow volume without losing the bound on the arrival times.

Combining all this, we obtain

$$\begin{aligned} \ell_{P,d_P}(i\beta) - \ell_{P',d_{P'}}(i\beta) &\stackrel{(15)}{\geq} \ell_{P,d_P}(i\beta) - \tilde{\ell}_{P',d_{P'}}(i\beta) - \frac{\varepsilon}{4} \\ &\geq \ell_{P,d_P}^D(i) - \tilde{\ell}_{P',d_{P'}}^D(i) - \frac{3\varepsilon}{4} \geq \varepsilon - \frac{3\varepsilon}{4} > 0. \end{aligned}$$

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

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 (Ziemke et al. [32]) 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 probably 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 .

Our convergence result in Theorem 2 focuses on the embedding for fixed path choices of the packets/particles. An interesting but probably challenging follow-up question is the convergence of equilibria in the packet-routing model to Nash flows over time (the equilibria in the continuous flow model). More precisely, given a sequence of exact packet routing equilibria whose discretization parameters converge to zero, does the limit exist? If yes, does it constitute a Nash flow over time? As equilibria do not exist in general for competitive packet-routing games with multiple origin–destination pairs, it seems reasonable to focus on the single origin–destination case for these kinds of questions. As a first step, a deeper understanding of Nash flows over time is necessary. One essential aspect is the robustness. If slight distortions of the flow over time or small perturbations of the network lead to drastically different dynamic equilibria, there would be no hope for such a convergence result as packets (no matter how small they are) would always cause some inaccuracy. Fortunately, this essential condition is studied by Olver et al. [20] recently. The authors show continuity of Nash flows over time with respect to several parameters, such as transit times, capacities, and initial starting conditions. Still, given continuity, convergence does not follow easily. It would be necessary to overcome the challenge that packets cannot split into smaller pieces and traverse multiple paths at once, which is one key aspect of Nash flows over time. Also, the robustness results in Olver et al. [20] are not sufficient as they focus on distortion at a single point in time. However, a packet flow deviates from a Nash flow over time at almost all points in time. 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

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

References

- [1] Adamik A, Sering L (2022) Atomic splittable flow over time games. Aspnes J, Michail O, eds. *1st Sympos. Algorithmic Foundations Dynamic Networks (SAND 2022)*. *Leibniz Internat. Proc. Informatics* (Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Wadern, Germany), 4:1–4:16.
- [2] Cantarella G, Watling D (2016) Modelling road traffic assignment as a day-to-day dynamic, deterministic process: A unified approach to discrete- and continuous-time models. *EURO J. Transportation Logist.* 5(1):69–98.

- [3] Cao Z, Chen B, Chen X, Wang C (2017) A network game of dynamic traffic. *Proc. 2017 ACM Conf. Econom. Comput.* (Association for Computing Machinery, New York), 695–696.
- [4] Cominetti R, Correa J, Larré O (2011) Existence and uniqueness of equilibria for flows over time. *Internat. Colloquium Automata, Languages, Programming* (Springer-Verlag, Zurich), 552–563.
- [5] Cominetti R, Correa J, Larré O (2015) Dynamic equilibria in fluid queueing networks. *Oper. Res.* 63(1):21–34.
- [6] Cominetti R, Correa J, Olver N (2021) Long-term behavior of dynamic equilibria in fluid queueing networks. *Oper. Res.* 70(1):516–526.
- [7] Cominetti R, Scarsini M, Schröder M, Stier-Moses N (2022) Approximation and convergence of large atomic congestion games. *Math. Oper. Res.* Forthcoming.
- [8] Correa J, Cristi A, Oosterwijk T (2021) On the price of anarchy for flows over time. *Math. Oper. Res.*, 1–18.
- [9] Fleischer L, Tardos É (1998) Efficient continuous-time dynamic network flow algorithms. *Oper. Res. Lett.* 23(3–5):71–80.
- [10] Frascaria D, Olver N (2022) Algorithms for flows over time with scheduling costs. *Math. Programming* 192:177–206.
- [11] Graf L, Harks T, Sering L (2020) Dynamic flows with adaptive route choice. *Math. Programming* 183(1):309–335.
- [12] Harks T, Peis B, Schmand D, Tauer B, Vargas Koch L (2018) Competitive packet routing with priority lists. *ACM Trans. Econom. Comput.* 6(1):1–26.
- [13] Hofer M, Mirrokni V, Röglin H, Teng SH (2011) Competitive routing over time. *Theoretical Comput. Sci.* 412(39):5420–5432.
- [14] Horni A, Nagel K, Axhausen K, eds. (2016) *Multi-Agent Transport Simulation MATSim* (Ubiquity Press, London).
- [15] Ismaili A (2017) Routing games over time with FIFO policy. *Internat. Conf. Web Internet Econom.* (Springer International Publishing), 266–280.
- [16] Israel J, Sering L (2020) The impact of spillback on the price of anarchy for flows over time. *Internat. Sympos. Algorithmic Game Theory* (Springer), 114–129.
- [17] Kaiser M (2022) Computation of dynamic equilibria in series-parallel networks. *Math. Oper. Res.* 47(1):50–71.
- [18] Koch R, Skutella M (2011) Nash equilibria and the price of anarchy for flows over time. *Theory Comput. Systems* 49(1):71–97.
- [19] Meunier F, Wagner N (2010) Equilibrium results for dynamic congestion games. *Transportation Sci.* 44(4):524–536.
- [20] Olver N, Sering L, Vargas Koch L (2022) Continuity, uniqueness and long-term behavior of Nash flows over time. *2021 IEEE 62nd Annual Sympos. Foundations Comput. Sci.*, 851–860.
- [21] Otsubo H, Rapoport A (2008) Vickrey’s model of traffic congestion discretized. *Transportation Res. Part B Methodological* 42(10):873–889.
- [22] Pham H, Sering L (2020) Dynamic equilibria in time-varying networks. *Internat. Sympos. Algorithmic Game Theory* (Springer), 130–145.
- [23] PTV AG (2016) *VISUM 16—User Manual* (PTV AG, Karlsruhe, Germany).
- [24] Rosenthal R (1973) The network equilibrium problem in integers. *Networks* 3(1):53–59.
- [25] Scarsini M, Schröder M, Tomala T (2018) Dynamic atomic congestion games with seasonal flows. *Oper. Res.* 66(2):327–339.
- [26] Scheffler R, Strehler M, Vargas Koch L (2022) Routing games with edge priorities. *ACM Trans. Econom. Comput.* 10(1):1–27.
- [27] Sering L (2020) Nash flows over time. Unpublished PhD thesis, Technische Universität Berlin, Berlin.
- [28] Sering L, Skutella M (2018) Multi-source multi-sink Nash flows over time. *18th Workshop Algorithmic Approaches Transportation Model. Optim. Systems*, vol. 65 (Schloss Dagstuhl - Leibniz-Zentrum für Informatik), 1–20.
- [29] Sering L, Vargas Koch L (2019) Nash flows over time with spillback. *Proc. 30th Annual ACM-SIAM Sympos. Discrete Algorithms* (SIAM), 935–945.
- [30] Wardrop J (1952) Some theoretical aspects of road traffic research. *Proc. Inst. Civil Engrg.* 1(3):325–378.
- [31] Werth T, Holzhauser M, Krumke S (2014) Atomic routing in a deterministic queueing model. *Oper. Res. Perspect.* 1(1):18–41.
- [32] Ziemke T, Sering L, Vargas Koch L, Zimmer M, Nagel K, Skutella M (2021) Flows over time as continuous limit of packet-based network simulations. *Transportation Res. Procedia* 52:123–130.