

Atomic Splittable Flow Over Time Games

Antonia Adamik 

Technische Universität Berlin, Germany

Leon Sering  

ETH Zürich, Switzerland

Abstract

In an atomic splittable flow over time game, finitely many players route flow dynamically through a network, in which edges are equipped with transit times, specifying the traversing time, and with capacities, restricting flow rates. Infinitesimally small flow particles controlled by the same player arrive at a constant rate at the player's origin and the player's goal is to maximize the flow volume that arrives at the player's destination within a given time horizon. Here, the flow dynamics are described by the deterministic queuing model, i.e., flow of different players merges perfectly, but excessive flow has to wait in a queue in front of the bottle-neck. In order to determine Nash equilibria in such games, the main challenge is to consider suitable definitions for the players' strategies, which depend on the level of information the players receive throughout the game. For the most restricted version, in which the players receive no information on the network state at all, we can show that there is no Nash equilibrium in general, not even for networks with only two edges. However, if the current edge congestions are provided over time, the players can adapt their route choices dynamically. We show that a profile of those strategies always lead to a unique feasible flow over time. Hence, those atomic splittable flow over time games are well-defined. For parallel-edge networks Nash equilibria exists and the total flow arriving in time equals the value of a maximum flow over time leading to a price of anarchy of 1.

2012 ACM Subject Classification Theory of computation → Network flows; Theory of computation → Network games; Mathematics of computing → Network flows; Theory of computation → Quality of equilibria

Keywords and phrases Flows Over Time, Deterministic Queuing, Atomic Splittable Games, Equilibria, Traffic, Cooperation

Digital Object Identifier 10.4230/LIPIcs.SAND.2022.4

Related Version *Full Version:* <https://arxiv.org/abs/2010.02148>

Funding *Leon Sering:* Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project ID: 390685689).

Acknowledgements We have considered atomic splittable flow over time games in different settings and under various assumptions in collaboration with several people. Unfortunately, most of these research directions were more challenging than expected and not as successful as the work at hand. Nonetheless, we want to thank Laura Vargas Koch, Veerle Timmermans, Björn Tauer, Tim Oosterwijk and Dario Frascaria for the excellent collaboration and inspiring discussions.

1 Introduction

In *static routing problems*, traffic is to be routed through a network at minimum total cost. The cost or traveling time on each edge depends on its congestion. However, the assumption that an optimal routing might be implemented by some superordinate authority is not realistic in many settings. More likely, each network participant selfishly chooses a path in order to minimize their own traveling time. In general, the lack of coordination causes a higher total traveling time. To quantify this decrease in performance, the total traveling time



© Antonia Adamik and Leon Sering;

licensed under Creative Commons License CC-BY 4.0

1st Symposium on Algorithmic Foundations of Dynamic Networks (SAND 2022).

Editors: James Aspnes and Othon Michail; Article No. 4; pp. 4:1–4:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

of a *Wardrop equilibrium* [45] is compared to the total traveling time of the *system optimum*. The ratio between a worst equilibrium and the system optimum is the *price of anarchy* [38]. This basic model can be extended in several ways. In this research work we want to focus on two aspects.

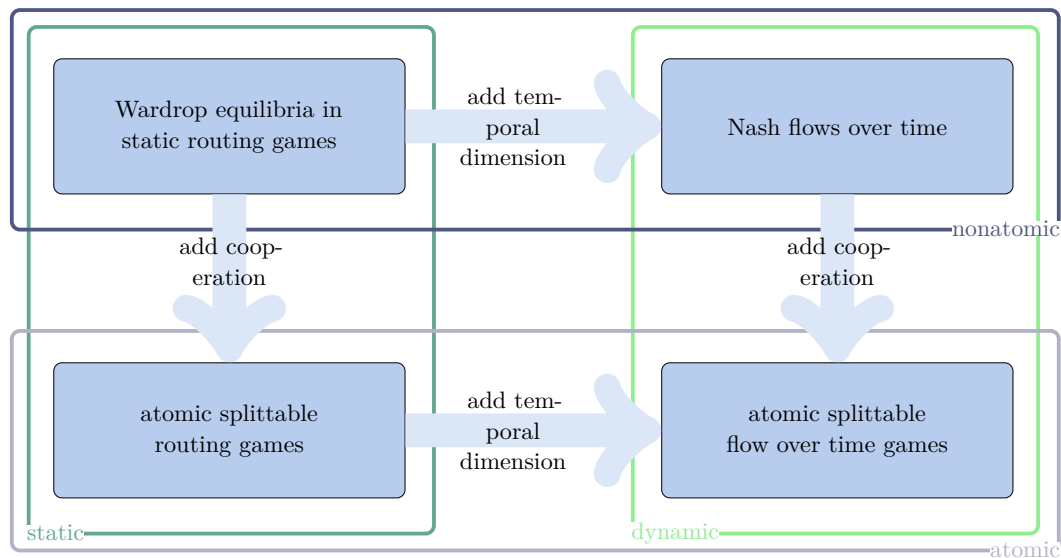
The first aspect is the temporal dimension. Vehicles in real traffic need time to move from the origin to the destination and the traveling time increases with the degree of congestion, which varies over time. In other words, the traffic flow does not traverse the network instantaneously, but progresses at a certain pace. In addition, the effects of a routing decision in one part of the network take some time to spread across the network as a whole. In order to mathematically model this, we add a time component, transforming static flows into *flows over time*. Here, every infinitesimally small flow particle needs time to traverse the network and the flow rates of the edges are restricted by capacities. By assuming that each particle acts selfishly, we can consider dynamic equilibria, which are called *Nash flows over time* [29].

For the second aspect, note that in real-world traffic the activity of a single road user has in most cases a negligible impact on the performance of the system as a whole. Furthermore, the assumption of independent and selfish particles is not justifiable in all applications: Networks where participants control a flow of positive measure are not covered. For instance, in transportation networks freight units might not act selfishly; they are controlled by freight companies that each control a significant amount of traffic. This leads to the second aspect, cooperative behavior among groups of network participants. To integrate this into the mathematical model, flow particles are allowed to form coalitions. In so-called *atomic splittable routing games* we consider a finite number of *atomic* players (the coalitions), each controlling a positive amount of flow volume that has to be *routed* through the network but can be *split up* and divided over different routes.

In this paper we want to combine both aspects, as depicted in Figure 1. That means, in contrast to Nash flows over time, sets of particles form coalitions which will be represented by superordinate players. In contrast to atomic splittable routing games, flow is modeled by a flow over time and players' decisions might be adapted to new situations. This extension covers a greater variety of scenarios. For example, road traffic models in which most drivers are guided by navigation systems (e.g. Google Maps, TomTom, Here, Garmin) can be modeled by covering the strategic behavior of the firms: The decisions of a single driver do not have a big impact on the city's traffic, but Google Maps decisions do; and TomTom might actually want to react. The situation will even intensify in the future with the rise of autonomous driving, as the decision making process is shifted to the navigation systems. We would like to point out that – in contrast to previous models – the interests of navigation companies and the general public are assumably in line: The cooperation of the users of a navigation system could reduce the average driving time per company on the one hand and the driving time in general on the other hand. This would lead to lower energy consumption, and therefore, lower emissions of polluting substances.

It turned out to be surprisingly challenging to consider equilibria in atomic splittable flow over time games. For this reason the overall goal is to define a solid model on these dynamic games and to present some preliminary observations, as well as some non-trivial first results, which serve as a basis for further research.

Related work. Static network flows have been studied for quite a while. A lot of pioneer work is due to Ford and Fulkerson, who also were the first to introduce *flows over time* [15, 16]. They provided an efficient algorithm for a *maximum flow over time*, which sends the maximal flow volume from a source to a sink given a finite time horizon. Closely related, a *quickest flow*



■ **Figure 1** Relationship between equilibrium situations in static routing games and atomic splittable flow over time games.

minimizes the arrival time of the latest particle for a given flow volume. This can be achieved by combining the algorithm of Ford and Fulkerson with a binary search framework [6, 14]. For single-source and single-sink networks it is furthermore possible to construct a flow over time that is maximal for all time horizons (and quickest for all flow volumes) simultaneously. The existence of these so-called *earliest arrival flows* was shown by Gale in 1959 [17]. They can be computed algorithmically by using the successive shortest path algorithm in the residual networks [31]. For more details and further references to literature on optimization problems in the flow over time setting, we refer to the survey of Skutella [43].

Koch and Skutella [29] approach flows over time from a game theoretic perspective by introducing *Nash flows over time*. In their model, every infinitesimally small flow particle is considered to be a player aiming to reach the common destination as early as possible. As the flow rate entering an edge could exceed its capacity, they considered the *deterministic queuing model* [44], which causes the excess flow to wait in a queue in front of the bottle-neck. Existence of these *dynamic equilibria* were shown by Cominetti et al. [8]. Several other aspects, including uniqueness, continuity, long term behavior, multi-terminals, spillback and price of anarchy, were studied in recent years [3, 9, 10, 12, 27, 32, 34, 40, 41]; see [39] for an overview. A slightly different approach for user equilibria was presented by Graf et al. [18, 19, 20]. They use the same flow over time model, except that particles do not anticipate the future evolution of the flow, but instead choose quickest routes according to current waiting times. As these delays may be subject to change, each particle can adapt its route choice along the way.

Atomic splittable congestion games for static network flows can be described as Wardrop equilibria [45] with coalitions [25, 30]; see also the survey of Correa and Stier-Moses [13]. For these games, Nash equilibria always exist, which can be shown by standard fixed point techniques [35]. Altman et al. [1] showed that equilibria are unique if the delay functions are polynomials of degree less than 3. Regarding more general delay functions, Bhaskar et al. [4] showed that for two players a unique equilibrium exists if, and only if, the network is a *generalized series-parallel graph*. Harks and Timmermans [24] showed uniqueness of

equilibria when the players' strategy space has a bidirectional flow polymatroid structure. Roughgarden [36] showed that the inefficiency of a system decreases with an increasing degree of cooperation. He showed that the price of anarchy for classes of traveling time functions in the atomic case is bounded by the price of anarchy for the same class of functions in the nonatomic case. Further research on the price of anarchy in static atomic splittable games is due to Cominetti et al. [11], Harks [21], and Roughgarden and Schoppmann [37]. Computational-wise Cominetti et al. [11] showed that equilibria can be computed efficiently when the cost functions are affine and player-independent. Regarding player-specific affine costs, Harks and Timmermans [23] described a polynomial algorithm for parallel-edge networks, and Bhaskar and Lolakapuri [5] presented an exponential algorithm for general convex functions. Very recently, Klimm and Warode showed that the computation with player-specific affine costs is PPAD-complete for general networks [28].

We should also mention that the combination of cooperation and temporal dimension has been considered for discrete packet routing games; see Peis et al. [33]. Here, each player controls a finite amount of packets, which has to be routed through a network in discrete time steps. For more results on competitive packet routing models we refer to Hoefer et al. [26] (continuous-time packets model) and Harks et al. [22] (discrete-time packet model).

Contribution and overview. In Section 2, we introduce all notations and formally describe atomic splittable flow over time games for general networks. The players' strategies determine how much flow they assign to each edge for every point in time during the game depending on available information on the current state. We consider two very natural sets of information. The first consists solely of the current time. In Section 3 we show that this setting does not allow for a Nash equilibrium in general, not even in a network with only two parallel edges. This motivates to consider more complex information models. Hence, Section 4 is dedicated to the second set of information which additionally comprises the current congestion of the edges in form of the exit times. As the first main result we show that every strategy profile results in a unique feasible flow over time by formulating the conditions as initial value problem and applying the Picard-Lindelöf theorem. For parallel-edge networks we show that Nash equilibria always exist by explicitly stating a strategy profile. Furthermore, we prove for that setting that all Nash equilibria have the same objective equal to the system optimum (i.e., the value of a maximum flow over time) implying that the price of anarchy for those networks is 1. Finally, we suggest further areas of research in Section 5.


2 Atomic Splittable Flow Over Time Games

In this section, we are going to properly define atomic splittable flow over time games. The two main aspects are the multi-commodity flow dynamics (see [40]) and the players' strategies, which depend on the information received over time.

Game setting. A *network* consists of a directed graph $G = (V, E)$, where every edge $e \in E$ is equipped with a *transit time* $\tau_e > 0$ and a *capacity* $\nu_e > 0$. For a node $v \in V$ we denote the set of all incoming edges by δ_v^- and the set of outgoing edges by δ_v^+ .

For an *atomic splittable flow over time game* we consider a finite set of players P , each with an origin-destination pair s_j - t_j and a *supply rate* $d_j > 0$, as well as a time horizon $H > 0$. We assume that s_j can reach t_j within the network.

The flow of player j enters the network via node s_j at a rate of d_j from time 0 onwards. The goal is to maximize the cumulative flow volume reaching node t_j before the end of the game at time H .



■ **Figure 2** Representation of an edge in the deterministic queuing model: If more flow particles enter edge $e = uv$ within the total inflow rate $\sum_{j \in P} f_{e,j}^+(\theta)$ than its capacity ν_e allows to process, they build up a queue, whose current length is given by $z_e(\theta)$. Whenever the queue is non-empty at time θ the total outflow rate $\sum_{j \in P} f_{e,j}^-(\theta + \tau_e)$ at time $\theta + \tau_e$ equals the capacity ν_e .

Flow dynamics. In the deterministic queuing model the total inflow rate into an edge is bounded by the capacity. If the capacity is exceeded, a queue builds up in which all entering particles have to wait in line. Afterwards each particle needs τ_e time to traverse the edge before it can enter the next edge along the path; see Figure 2. The dynamics of this process are formally defined as follows:

A *flow over time* is described by a family of Lebesgue-integrable functions $f = (f^+, f^-) = (f_{e,j}^+, f_{e,j}^-)_{e \in E, j \in P}$, where $f_{e,j}^+, f_{e,j}^- : [0, H) \rightarrow \mathbb{R}_{\geq 0}$ denote the rate at which flow controlled by player j enters and leaves edge e . The flow rates summed over all players $\sum_{j \in P} f_{e,j}^+(\theta)$ and $\sum_{j \in P} f_{e,j}^-(\theta)$ are called *total in- and outflow rates*. The *cumulative in- and outflows*, i.e., the amount of player j 's flow that has entered or left an edge e up to time θ , is denoted by $F_{e,j}^+(\theta) = \int_0^\theta f_{e,j}^+(\vartheta) d\vartheta$ and $F_{e,j}^-(\theta) = \int_0^\theta f_{e,j}^-(\vartheta) d\vartheta$. Finally, f^+, f^-, F^+ and F^- denote the vectors of (cumulative) in- and outflows with one entry per edge-player-pair.

Such a family of functions f is a *flow over time* if the following two conditions hold for all $e \in E$ and $j \in P$:

Flow conservation is fulfilled for all $\theta \in [0, H)$:

$$\sum_{e \in \delta_v^+} f_{e,j}^+(\theta) - \sum_{e \in \delta_v^-} f_{e,j}^-(\theta) = \begin{cases} d_j & \text{for } v = s_j, \\ 0 & \text{for } v \in V \setminus \{s_j, t_j\}. \end{cases} \quad (1)$$

Non-deficit constraints are satisfied for all $\theta \in [0, H - \tau_e)$:

$$F_{e,j}^+(\theta) - F_{e,j}^-(\theta + \tau_e) \geq 0. \quad (2)$$

To track the net flow that is not yet processed and remains in the *queue*, we introduce $z_e(\theta)$ to denote the *queue length* at time θ . Formally, it is defined as $z_e(\theta) = \sum_{j \in P} F_{e,j}^+(\theta) - \sum_{j \in P} F_{e,j}^-(\theta + \tau_e)$ for $e \in E$. For a *feasible* flow over time we require that, whenever flow is waiting in the queue, the edge operates at capacity rate. In other words, for all $\theta \in [0, H - \tau_e)$ and $e \in E$, we require

$$\sum_{j \in P} f_{e,j}^-(\theta + \tau_e) = \begin{cases} \nu_e & \text{if } z_e(\theta) > 0, \\ \min \{ \sum_{j \in P} f_{e,j}^+(\theta), \nu_e \} & \text{else.} \end{cases} \quad (3)$$

The *waiting time* $q_e(\theta)$ experienced by a particle entering edge e at time θ is generally defined as the time needed to process the flow present in the queue when the particle enters it. In other words, it is the time span between its entrance to and its exit from the queue just before traversing the edge. That is

$$q_e(\theta) := \min \left\{ q \geq 0 \mid \int_\theta^{\theta+q} \sum_{j \in P} f_{e,j}^-(\vartheta + \tau_e) d\vartheta = z_e(\theta) \right\} = \frac{z_e(\theta)}{\nu_e}.$$

4:6 Atomic Splittable Flow Over Time Games

The *exit time* $T_e(\theta)$ of a particle entering an edge $e \in E$ at time θ is given by the sum of the entrance time θ , its waiting time in the queue $q_e(\theta)$ and the transit time τ_e . Hence, we have

$$T_e(\theta) := \theta + \frac{z_e(\theta)}{\nu_e} + \tau_e.$$

Since z_e can at most decrease at rate ν_e , it follows that $q'_e(\theta) \geq -1$ and $T'_e(\theta) \geq 0$, which induces that T_e is non-decreasing. Note that these derivatives exist for almost all θ due to Lebesgue's differentiation theorem.

Finally, in a feasible flow over time, the flow of different players should merge seamlessly. This means that, at any point in time, a player's share of the total outflow rate is equal to her share of the total inflow rate back at the time when the flow entered the edge. More precisely, we require

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

for all $\theta \in [0, H)$ and $\phi := \max T_e^{-1}(\theta)$. Here, we set $f_{e,j}^+(\phi) := 0$ for $\phi < 0$. Note that ϕ denotes the edge-entrance times of all particles leaving the edge at time θ . Taking the maximum of $T_e^{-1}(\theta)$ simply ensures well-definedness which is required since T_e might not be strictly increasing.

To conclude we denote the set of all *feasible flows over time* by \mathcal{F} , i.e., all $f = (f^+, f^-)$ that satisfy (1), (2), (3) and (4). As the outflow rates are uniquely defined by the inflow rates, we refer to a feasible flow over time only by the corresponding inflow f^+ , and write $f^+ \in \mathcal{F}$.

Atomic splittable flow over time games. Let $\rho_j : \mathcal{F} \rightarrow \mathbb{R}$ be the function indicating player j 's *payoff* for a given $f \in \mathcal{F}$, which is to be maximized. In general, ρ can be set to various objective functions (e.g. arrival time of the player's latest particle or the average arrival time), but in this paper we will focus on the *maximum flow over time problem*. Each player wants to maximize her amount of flow routed from s_j to t_j before the end of the game at time H :

$$\rho_j(f) = \sum_{e \in \delta_{t_j}^-} F_{e,j}^-(H) - \sum_{e \in \delta_{t_j}^+} F_{e,j}^+(H).$$

We choose this maximum flow objective as it seems to be the most straight-forward payoff-function. It is conceivable though, that most results might transfer to quickest flow payoff-functions via a binary-search framework (cf. in non-competitive settings quickest flows are constructed from maximum flows over time via binary-search). But we leave this for future research.

The *strategy space* is a player's set of viable options in order to maximize her payoff. A single *strategy* is a complete instruction determining the player's inflow rates of all times and for all situations possibly occurring. Formally, the strategy space of player j is a set of functions

$$\Sigma_j = \left\{ g_j : \mathcal{I} \rightarrow [0, 1]^E \mid \begin{array}{l} g_{e,j} \text{ is Lebesgue-integrable for all } e \in E \text{ and} \\ \sum_{e \in \delta_v^+} g_{e,j}(I) = 1 \text{ for all } I \in \mathcal{I}, v \in V \setminus \{t_j\} \end{array} \right\},$$

where \mathcal{I} is the set of information available to the players. This set is not well-defined yet, but we will discuss this extensively, and in the end, we will consider two separate definitions for \mathcal{I} , one in Section 3 and one in Section 4. Informally, the set of information is used to delineate what defines a situation and how it is perceived by the players. The interpretation is as follows. For every information $I \in \mathcal{I}$, the value $g_{e,j}(I)$ determines which proportion of player j 's flow arriving at v is distributed onto the outgoing edges $e \in \delta_v^+$. We use proportions that sum up to 1 instead of the inflow rates, as this easily ensures that flow conservation is fulfilled at all times. Note that the received information I can depend on the current time θ and on the flow over time f itself. As we normally do not want the players to see the future, it should only depend on the flow over time up to time θ . In general it might be player dependent, hence, we write $I_j(\theta, f)$.

In order to turn a *strategy profile* $g = (g_j)_{j \in P} \in \times_{j \in P} \Sigma_j$ into a feasible flow over time f , we consider the following system of equations that has to be satisfied for all $\theta \in [0, H)$:

$$\begin{aligned} f_{e,j}^+(\theta) &= g_{e,j}(I_j(\theta, f)) \cdot \left(\sum_{e' \in \delta_{s_j}^-} f_{e',j}^-(\theta) + d_j \right) \quad \text{for } j \in P, e \in \delta_{s_j}^+, \\ f_{e,j}^+(\theta) &= g_{e,j}(I_j(\theta, f)) \cdot \sum_{e' \in \delta_u^-} f_{e',j}^-(\theta) \quad \text{for } j \in P, e = uv \in E \setminus (\delta_{s_j}^+ \cup \delta_{t_j}^+), \\ f_{e,j}^+(\theta) &= 0 \quad \text{for } j \in P, e \in \delta_{t_j}^+. \end{aligned} \tag{5}$$

Note that in order to keep it as simple as possible, we assume that flow of player j reaching t_j leaves the network immediately, hence, $f_{e,j}^+(\theta) = 0$ for all $e \in \delta_{t_j}^+$. This leads to a simpler payoff for all players j of $\rho_j(f) = \sum_{e \in \delta_{t_j}^-} F_{e,j}^-(H)$.

Since the inflow rates $f^+(\theta)$ might depend on the flow over time itself up to time θ , it is not guaranteed that this system of equations yields a feasible flow over time as solution.

To illustrate this issue assume for example a game with a single player in a network with two parallel edges e_1 and e_2 where e_1 has a tiny capacity. If the player's strategy is to send everything into e_1 as long as there is no waiting time on e_1 , and otherwise send everything into e_2 , this would not result in a feasible flow over time. To see this, suppose that the inflow into e_1 would be positive on a measurable set, which would immediately cause a positive waiting time on e_1 . Hence, the strategy says that no flow is sent into e_1 . On the other hand, if no flow is sent into e_1 at all there would not be any waiting time, leading again to a contradiction.


For that reason, the key challenge is to find a reasonable set of information \mathcal{I} and strategy spaces Σ_j , such that, on the one hand, there exists a unique (up to a null set) feasible flow over time satisfying (5) for every given strategy profile g , and on the other hand, a Nash equilibrium exists. Note that we only consider *pure* Nash equilibria.

In the following two sections we discuss two very natural sets of information.

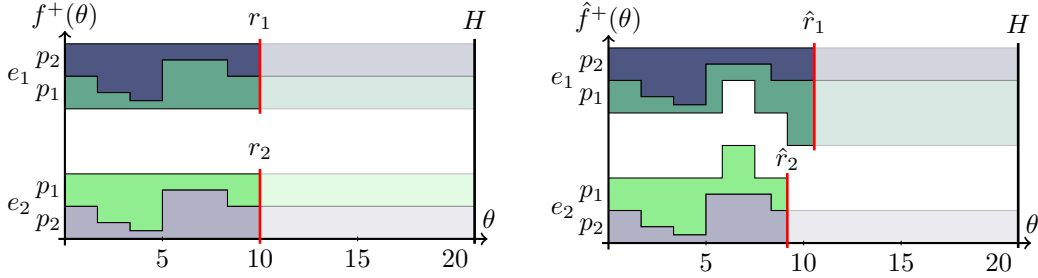
3 Temporal Information Only

First, we want to examine the simplest set of information, namely the current point in time only: We set $I_j(\theta, f) := \theta$ for all players $j \in P$. That means the players do not receive any information about the current state of the flow, but instead have to decide at the beginning of the game along which routes their flow particles are routed. In this case it is guaranteed that there exists a unique feasible flow over time that satisfies (5), as we can simply set $f_{e,j}^+(\theta)$ to the right sides of (5) (formally this also follows from Theorem 2). However, Nash equilibria do not exist in general, which is already true for very simple networks.

4:8 Atomic Splittable Flow Over Time Games



(a) A parallel network with two identical edges and two identical players.



(b) Step 1: p_1 mirrors p_2 's strategy. We have $r_1 = r_2 = 10$ and both queues build up from $\theta = 0$ onward. W.l.o.g., we have $f_{1,2}^+(\theta) \leq d_2 = 4$ for $\theta \in [10 - \varepsilon, 10)$.

(c) Step 2: p_1 shifts δ flow units from e_1 to e_2 causing $\hat{r}_1 > 10$ and $\hat{r}_2 < 10$. By pumping at maximum rate into e_1 after \hat{r}_2 her payoff is $\hat{\rho}_1 > 40 = \text{OPT} / 2$.

■ **Figure 3** Construction of a response strategy that always yields strictly more than half of the total optimum in a network of two parallel edges. Here, r_i and \hat{r}_i denote the time when the last particles arriving at t just in time, enter e_i .

► **Theorem 1.** *With temporal information only, there exists no Nash equilibrium in an atomic splittable flow over time game of two players p_1 and p_2 with identical supply rates $d_1, d_2 > 0$, on a network with two identical parallel edges e_1, e_2 from s ($= s_{p_1} = s_{p_2}$) to t ($= t_{p_1} = t_{p_2}$) with $\nu_{e_1} = \nu_{e_2} < d_1 = d_2$ and $\tau_{e_1} = \tau_{e_2} < H$.*

The key proof idea is the following. To every strategy of the competitor, a player can choose a response strategy that yields a payoff of strictly more than half the total optimum, i.e., the flow value of a maximum flow over time with inflow rate $d_1 + d_2$. This can be achieved by first *mirroring* the competitor's strategy (copying the strategy but interchanging the roles of e_1 and e_2) and then shifting some flow in the beginning of the game. This shift changes the points in time r_i when the last particles arriving at t just in time, enter edge e_i for $i = 1, 2$. As the new values \hat{r}_1 and \hat{r}_2 are not equal anymore the responding player can modify the inflow rates between \hat{r}_1 and \hat{r}_2 in order to squeeze in a little more flow than the competitor, which will then be a little more than $\text{OPT} / 2$.

As in turn the competitor can again choose a strategy with a payoff of more than half the optimum, this immediately implies that a Nash equilibrium cannot exist. This idea for the network given in Figure 3a is visualized in Figures 3b and 3c.

Proof of Theorem 1. For the sake of simplicity we replace indexes e_i by i and p_j by j . E.g. $f_{1,2}$ denotes the inflow rate function of p_2 into e_1 . Furthermore, r_1 and r_2 denote the points in time when the last particles arriving at t before H enter e_1 and e_2 , respectively. We will use a hat to denote parameters that have altered with the change in strategy. For example, ρ will denote the initial and $\hat{\rho}$ the new payoff. The optimal value of a maximum flow is $\text{OPT} = \nu_1 \cdot (H - \tau_1) + \nu_2 \cdot (H - \tau_2)$. We show that to every strategy of player p_2 , player p_1 can choose a response strategy that yields a payoff of strictly more than half the optimum.

As in turn p_2 can again choose a strategy with a payoff of more than half the optimum, this immediately implies that a Nash equilibrium cannot exist. To achieve this, player p_1 mirrors p_2 's strategy. By this we mean that p_1 has the same strategy as p_2 , but the roles of the edges are interchanged. Figure 3b illustrates this behavior.

$$f_{1,1}^+(\theta) = f_{2,2}^+(\theta) \quad \text{and} \quad f_{2,1}^+(\theta) = f_{1,2}^+(\theta) \quad \text{for } \theta \in [0, H).$$

Both total inflow rates into the edges are equal and exceed the capacities, i.e., $f_{1,1}^+(\theta) + f_{1,2}^+(\theta) = f_{2,1}^+(\theta) + f_{2,2}^+(\theta) > \nu_1 = \nu_2$ for $\theta \in [0, H)$. It is easy to see that both payoffs are $\text{OPT}/2$. As capacities are exceeded, queues build up from the very beginning. At time $r_1 = r_2 = \text{OPT}/(d_1 + d_2)$ the last particles that will reach t in time enter the network. The response strategy of p_1 now consists of reallocating a little flow to benefit from deferred \hat{r}_1 and \hat{r}_2 : W.l.o.g. let p_2 send not more than half of her supply rate d_2 at any time in $[r_1 - \varepsilon, r_1)$ for $\varepsilon > 0$, into e_1 . That means player p_1 sends not less than half of d_1 into e_1 . (If no such ε exists, due to some crazy function behavior, one can consider averages instead.) Player p_1 reallocates some $\delta > 0$ flow units from e_1 to e_2 . By doing so, \hat{r}_2 is shifted to $r_2 - \frac{\delta}{\nu_2}$. Meanwhile we want that $\frac{\delta}{\nu_2} \leq \varepsilon$, that the δ flow units are taken before time $r_2 - \varepsilon$ and that the queue of e_1 is never empty for $\theta > 0$. Since there has to exist p_1 -flow of positive measure up to time r_1 on e_1 (in case the only flow p_1 sends into e_1 is within $[r_1 - \varepsilon, r_1)$ we can choose a smaller ε), we can clearly set δ small enough to fulfill these conditions. As a result, \hat{r}_1 is postponed, i.e., $\hat{r}_1 > \hat{r}_2$. Conclusively, every flow sent until time \hat{r}_2 reaches t before H . Hence, the payoffs of both players belonging to $[0, \hat{r}_2)$ are equal. After time \hat{r}_2 , player p_1 can pump all her flow into e_1 , since it is still eligible to reach t before H . See Figure 3c for an illustration. Therefore,

$$\hat{f}_{1,1}^+(\theta) := d_1 \begin{cases} > d_2/2 \geq f_{1,2}^+(\theta) & \theta \in [\hat{r}_2, r_1) \\ \geq f_{1,2}^+(\theta) & \theta \geq r_1. \end{cases}$$

Hence, p_1 's inflow rate into e_1 is strictly greater than p_2 's during a time period of positive measure, which shows that the payoff of p_1 is strictly larger than that of p_2 . As the sum of the payoffs equals OPT we have that $\hat{r}_1 > \text{OPT}/2$. ◀

4 Information on Exit Times

The absence of a Nash equilibrium for temporal information only was mainly due to the theoretical information advantage of the deviating player. Player p_1 can respond to p_2 's strategy, while p_2 does not see p_1 's moves and is unable to react. As it is very natural to require inter-player reactions over time, we extend the information by the current congestion of the edges in form of the exit times $T(\theta) := (T_e(\theta))_{e \in E}$. Formally, we define

$$I_j(\theta, f) := (\theta, T(\theta)) \in \mathcal{I} := [0, H) \times \mathbb{R}_{\geq 0}^E \quad \text{for all } j \in P.$$

The reason for choosing exit times over waiting times or queue sizes, which both contain the same information about the congestion in the networks as the exit times, is that the exit times are non-decreasing. As the exit time $T_e(\theta) = \theta + \tau_e + \frac{1}{\nu_e} \sum_{j \in P} (F_{e,j}^+(\theta) - F_{e,j}^-(\theta + \tau_e))$ depends directly on the cumulative flows the equations system (5) becomes a system of differential equations. In order to show existence and uniqueness we use the Picard-Lindelöf theorem. For this reason, we require the strategies to be locally Lipschitz-continuous from the right in order to ensure uniqueness. More formally, we say a strategy $g_{e,j}$ is *right-Lipschitz* if for every $I \in \mathcal{I}$ there exists an $L > 0$ such that there exists an $\varepsilon > 0$ with $|g_{e,j}(I) - g_{e,j}(I + x)| \leq L \cdot \|x\|$ for all $x \in [0, \varepsilon] \times [0, \varepsilon]^E$.

Existence and uniqueness. First we show, that in this setting every right-Lipschitz strategy profile yields a unique feasible flow over time.

► **Theorem 2.** *For all right-Lipschitz strategy profiles $g = (g_j)_{j \in P}$ of an atomic splittable flow over time game with information on exit times, there exists a unique feasible flow over time f^+ satisfying (5).*

Proof. We will construct the feasible flow over time satisfying (5) step by step starting with the empty flow over time $f^+ \equiv 0$ up to time 0. We define a *restricted flow over time* on the interval $[0, a)$ to be a vector of Lebesgue-integrable functions $(f_{e,j}^+)_{e \in E, j \in P}$, such that all flow conditions hold for all times in $[0, a)$.

Suppose we are given a unique restricted feasible flow over time satisfying (5) on the interval $[0, a)$ for some $a \in [0, H]$. If $a = H$, we are done. It is possible to determine the queue lengths $z(a) = (z_e(a))_{e \in E}$, the waiting times $q(a) = (q_e(a))_{e \in E}$ and the exit times $T(a) = (T_e(a))_{e \in E}$ based on $f^+|_{[0,a)}$. Hence, we can also evaluate $g(a, T(a))$. In order to further extend the flow over time, we specify an interval $[a, b) \subseteq [0, H]$. The right endpoint $b > a$ has to be small enough to ensure that $g_{e,j}(\theta, T(\theta))$ is Lipschitz-continuous for $\theta \in [a, b)$ and for all $e \in E$ and $j \in P$. We can find such b as all $g_{e,j}$ are right-Lipschitz, the exit time functions T_e are non-decreasing and continuous in θ .

We obtain the following initial value problem:

$$\begin{aligned} f_{e,j}^+(a) &= g_{e,j}(a, T(a)) \cdot C_u(a), \\ f_{e,j}^+(\theta) &= g_{e,j}(\theta, (\theta + \tau_{e'} + \frac{1}{\nu_{e'}} \sum_{j' \in P} F_{e',j'}^+(\theta) - F_{e',j'}^-(\theta + \tau_{e'}))_{e' \in E}) \cdot C_u(\theta), \end{aligned}$$

for all $j \in P$ and $e = uv \in E$. Here, $C_u(\theta) \geq 0$ is the total inflow rate into node u except for t where it is 0 (as stated in (5)). Since the transit times are strictly positive $C_u: [a, b) \rightarrow \mathbb{R}$ is completely determined by the restricted feasible flow over time on $[0, a)$ as long as $b - a < \tau_e$. In addition, C_u is right-Lipschitz, therefore, we can choose b small enough, such that C_u is Lipschitz-continuous on $[a, b)$. Furthermore, if $q_e(\theta) > 0$ we have that $F_{e',j'}^-(\theta + \tau_{e'})$ is also determined from the past as long as $b - a < q_e(\theta)$ or in the case of $q_e(\theta) = 0$, we have that $F_{e',j'}^-(\theta + \tau_{e'}) = \sum_{j' \in P} F_{e',j'}^+(\theta)$ as required by (2).

Since the exit times T depend Lipschitz-continuous on F^+ , also the right-side depends Lipschitz-continuous on F^+ . Hence, the Picard-Lindelöf theorem guarantees the existence of a unique solution. It is easy to see, that extending f by this solution yields a restricted feasible flow over time on $[0, b)$. Flow conservation is fulfilled as $\sum_{e \in \delta_u^+} g_{e,j}(\theta, T(\theta)) = 1$ and (2), (3) and (4) are satisfied as we implicitly define the outflow rates accordingly.

It remains to show that it is possible to repeat the process until we cover the whole interval $[0, H)$. Let $b_i, i = 1, 2, 3, \dots$, be the right endpoints during this extension process. The sequence might converge to $\lim_{i \rightarrow \infty} b_i = b_\infty < H$. Existence and uniqueness are thus provided on $[0, b_\infty)$. But this means, we can apply the extension process for $a = b_\infty$. As we can always continue this process from a limit point, we can apply this transfinite induction to obtain a unique feasible flow over time on $[0, H)$. ◀

This theorem shows that we obtain a well-defined atomic splittable flow over time game as long as we only consider strategies that are not too wild. It is worth noting that right-Lipschitz functions can have infinitely many jumps and that we cannot hope for much more general strategy-functions as argued in the following. Suppose we allow for strategies that are not continuous from the right. We end up with the following problem: Consider a game with only one player p_1 and a network consisting of two edges e_1, e_2 both from s_{p_1} to t_{p_1}

and with $\nu_{e_1} < d_{p_1}$. The strategy with $g_{e_1, p_1}(\theta, T(\theta)) = 1$ if $T_{e_1}(\theta) \leq \theta + \tau_{e_1}$ and 0 otherwise, means, that if there is no queue on e_1 the players sends all its flow into e_1 (which causes a queue to build up) but if there is a positive queue she sends nothing (which causes any positive queue to decrease). This strategy is not right-continuous in T as $g_{e_1, 1}$ switches from 1 to 0 as soon as $T > \tau_{e_1} + \theta$, and clearly, it cannot lead to a feasible flow over time.

Existence of Nash equilibria in parallel-edge networks. Unfortunately, the task to show the existence of a Nash equilibrium in this setting turns out to be quite challenging. For this reason we only show the existence of Nash equilibria for simple networks: For the remaining of this section we consider parallel-edge networks with two nodes $s = s_j$ and $t = t_j$, $j \in P$. We obtain

$$\Sigma_j = \left\{ g_j : \mathcal{I} \rightarrow [0, d_j]^E \mid \begin{array}{l} g_{e,j} \text{ is right-Lipschitz for all } e \in E \\ \text{and } \sum_{e \in E} g_{e,j}(I) = 1 \text{ for all } I \in \mathcal{I} \end{array} \right\}.$$

This leads to an ‘‘over time’’ version of *atomic splittable singleton games* as they were studied in by Harks and Timmermans [23]. As an additional motivation, it is worth noting, that these restricted networks, become more meaningful if, instead of a routing game, we consider a throughput-scheduling problem [42]. Suppose the edges represent machines on which competing players want to maximize their throughput. The jobs can run in parallel on multiple machines up to some maximum rate of d_j and each machine has a maximal service rate of ν_e and individual time horizons $H - \tau_e$. A very similar model without the game-theoretical aspect is for example studied by Armony and Bambos [2].

In order to obtain a Nash equilibrium, it is worth noting that players do not care on which edge their flow is sent, as long as it arrives at the destination before H . To model a suitable strategy, we introduce the set of *active edges* $E'(T) := \{ e \in E \mid T_e < H \}$ on which flow still arrives at t before H , depending on the exit times. The resulting strategy for a player $j \in P$ on edge $e \in E$ could look as follows:

$$g_{e,j}(\theta, T) := \begin{cases} \frac{\nu_e}{\sum_{e' \in E'(T)} \nu_{e'}} & \text{if } e \in E'(T), \\ 0 & \text{if } E'(T) \neq \emptyset \text{ and } e \notin E'(T), \\ \frac{\nu_e}{\sum_{e' \in E} \nu_{e'}} & \text{if } E'(T) = \emptyset. \end{cases} \quad (6)$$

The third case is not of importance for the player, as none of the flow sent into the network will arrive in time anymore. As $E'(T)$ stays constant for small increases of T , the same is true for $g_{e,j}$. Since $\sum_{e \in E} g_{e,j}(\theta, T) = 1$, we have $g_j \in \Sigma_j$.

We will show that this strategy profile leads to a Nash equilibrium. For this we first show in Lemma 3 that the given strategy profile leads to a total payoff equal to the system optimum OPT (the value of a maximum flow over time with inflow rate $\sum_{j \in P} d_j$). Afterwards, we determine in Lemma 4 that the payoff for each player given the strategy profile in (6) is a fixed share of the total payoff. Finally, we argue that none of the players has an incentive to deviate from the given strategy profile, since shares cannot be increased.

For the remaining of this section, let r_e be the point in time when the (first) particle that arrives at t at time H enters edge e . Formally, $r_e := \min T_e^{-1}(H)$. Hence, $T_e(\theta) < H$ for any $\theta < r_e$ and $T_e(\theta) \geq H$ for any $\theta \geq r_e$.

► **Lemma 3.** *For the strategy profile given in (6) the sum of the payoffs of all players equals the system optimum OPT.*

4:12 Atomic Splittable Flow Over Time Games

Proof. To prove this, we allow the network to have transit times that equal 0. We split the set of instances into three cases

Case 1: $\sum_{j \in P} d_j \geq \sum_{e \in E} \nu_e$.

W.l.o.g. we assume $\tau_e < H$; otherwise the edge would be superfluous and could be deleted.

We have $\text{OPT} = \sum_{e \in E} \nu_e \cdot (H - \tau_e)$ as shown in [43]. Since all players route their flow proportionally to the capacities of the active edges, we have

$$\sum_{j \in P} f_{e,j}^+(\theta) = \frac{\nu_e}{\sum_{e' \in E'(T(\theta))} \nu_{e'}} \cdot \sum_{j \in P} d_j \geq \nu_e \quad \text{for all } \theta \text{ with } T_e(\theta) < H.$$

With this at hand, we can state the total outflow during the whole game:

$$\begin{aligned} \sum_{j \in P} \rho_j &= \sum_{e \in E} \sum_{j \in P} F_{e,j}^-(H) = \sum_{e \in E} \int_0^H \sum_{j \in P} f_{e,j}^-(\theta) d\theta \\ &= \sum_{e \in E} \left(\int_0^{\tau_e} 0 d\theta + \int_{\tau_e}^H \nu_e d\theta \right) = \sum_{e \in E} (H - \tau_e) \cdot \nu_e = \text{OPT}. \end{aligned}$$

Case 2: $\sum_{j \in P} d_j < \sum_{e \in E} \nu_e$ and $\tau_e = 0$ for all $e \in E$.

We can easily see that $\text{OPT} = H \cdot \sum_{j \in P} d_j$ and

$$\sum_{j \in J} f_{e,j}^+(\theta) = \frac{\nu_e}{\sum_{e' \in E} \nu_{e'}} \cdot \sum_{j \in P} d_j < \nu_e \quad \text{for all } \theta < H.$$

Since no queues build up, all edges stay active for all $\theta \in [0, H)$. Therefore, it holds that

$$\sum_{j \in P} \rho_j = \sum_{e \in E} \int_0^H \sum_{j \in P} f_{e,j}^-(\theta) d\theta = \sum_{e \in E} \sum_{j \in P} H \cdot d_j \cdot \frac{\nu_e}{\sum_{e' \in E} \nu_{e'}} = H \cdot \sum_{j \in P} d_j = \text{OPT}.$$

Case 3: $\sum_{j \in P} d_j < \sum_{e \in E} \nu_e$ and there exists an $e \in E$ with $\tau_e > 0$.

We assume that $\tau_e < H$ for all $e \in E$. Let e^* be the first edge to drop out of $E'(T(\theta))$, i.e., r_{e^*} is minimal among all r_e . As in the first phase no queues build up (see Case 2), we have $r_{e^*} = H - \tau_{e^*}$, which means that τ_{e^*} is maximal among all τ_e . Emphasis should be put on the fact that the whole flow sent into the network up to time r_{e^*} arrives at t in time; a volume of $\sum_{j \in P} d_j \cdot (H - \tau_{e^*})$ in total. Hence, the system optimum cannot perform any better until r_{e^*} . To show that after time r_{e^*} , the summed payoffs correspond to the system optimum as well, we reduce the remaining instance. First, we remove e^* from the set of edges, i.e., $\hat{E} = E \setminus \{e^*\}$. Second, we shift the time axis r_{e^*} time units back. By that we mean that the new time 0 corresponds to r_{e^*} in the former instance and $\hat{H} = H - r_{e^*} = \tau_{e^*}$. Everything else remains untouched, in particular all queues are empty at r_{e^*} . This instance is strictly smaller, which means that eventually the reduction process must end because either we obtain $\sum_{j \in P} d_j \geq \sum_{e \in \hat{E}} \nu_e$ (Case 1), $\tau_e = 0$ for all $e \in \hat{E}$ (Case 2) or we reach $|\hat{E}| = 1$ in which case the total payoff trivially equals OPT. ◀

► **Lemma 4.** For the strategy profile given in (6), the payoff of player j is given by

$$\rho_j = \text{OPT} \cdot \frac{d_j}{\sum_{j' \in P} d_{j'}}.$$

Proof. For $\theta \in [0, r_e)$ we have $\sum_{j \in P} f_{e,j}^+(\theta) = \sum_{j \in P} d_j \cdot g_{e,j}(\theta, T(\theta)) > 0$. We want to examine the outflow rates for $\phi \in [\tau_e, H)$. With (4) and $\theta = \max T_e^{-1}(\phi) < r_e$ we obtain for all $e \in E$ and $j \in P$ that

$$\begin{aligned} f_{e,j}^-(\phi) &= \frac{f_{e,j}^+(\theta)}{\sum_{j' \in P} f_{e,j'}^+(\theta)} \cdot \sum_{j' \in P} f_{e,j'}^-(\phi) \\ &= \frac{d_j \cdot g_{e,j}(\theta, T(\theta))}{\sum_{j' \in P} d_{j'} \cdot g_{e,j'}(\theta, T(\theta))} \cdot \sum_{j' \in P} f_{e,j'}^-(\phi) = \frac{d_j}{\sum_{j' \in P} d_{j'}} \cdot \sum_{j' \in P} f_{e,j'}^-(\phi). \end{aligned}$$

Taking the integral over $[0, H]$, summing over all edges $e \in E$ and using Lemma 3, yields

$$\rho_j = \sum_{e \in E} F_{e,j}^-(H) = \frac{d_j}{\sum_{j' \in P} d_{j'}} \cdot \sum_{e \in E} \sum_{j' \in P} F_{e,j'}^-(\phi) = \frac{d_j}{\sum_{j' \in P} d_{j'}} \cdot \text{OPT}. \quad \blacktriangleleft$$

With these two lemmas we can finally prove the following theorem.

► **Theorem 5.** *The strategy profile $(g_j)_{j \in P}$ given by (6) is a Nash equilibrium.*

Proof. We want to observe what happens if one player j^* deviates from her strategy in an ex ante manner. For this let $r := \max_{e \in E} r_e$ be the point in time when the last edge becomes inactive for the modified strategy-profile and let $\rho := \sum_{j \in P} \rho_j$ be the total amount of flow arriving in time. Clearly, $\rho \leq \text{OPT}$.

During $[0, r]$ every non-deviating player $j \in P \setminus \{j^*\}$ sends all her flow into edges that are still active, and therefore, all this flow arrives in time. Hence, her payoff is given by $\rho_j = d_j \cdot r$.

Since no flow that enters after r can possibly arrive in time, player j^* can achieve at most a payoff of $d_{j^*} \cdot r$. It follows that her share of the total amount of flow arriving in time is upper bounded by $\frac{d_{j^*}}{\sum_{j \in P} d_j}$, i.e.,

$$\rho_{j^*} \leq \rho \cdot \frac{d_{j^*}}{\sum_{j \in P} d_j} \leq \text{OPT} \cdot \frac{d_{j^*}}{\sum_{j \in P} d_j}.$$


By Lemma 4 the right-side equals the payoff of player j^* when choosing g_{j^*} as strategy. Hence, she cannot improve her payoff by deviating, which shows that $(g_j)_{j \in P}$ is indeed an equilibrium. \blacktriangleleft

We observe the following: As long as the players constantly choose from the set of active edges and the summed payoff equals OPT, i.e., no queue runs dry in case of a restricted network, the players' payoffs stay the same. Hence, there is a whole class of Nash equilibria.

Next we want to show that this characterizes all possible Nash equilibria, which implies that the price of anarchy is 1.

► **Theorem 6.** *For atomic splittable flow over time games on parallel-edge networks where exit times are provided as information with right-Lipschitz strategies, the price of anarchy is 1.*

Proof. The key observation is that every player j 's share of the total payoff $\sum_{j' \in P} \rho_{j'}$ is at least $d_j / \sum_{j' \in P} d_{j'}$ as long as the player only sends flow into active edges $E'(T(\theta))$. Also, the total payoff is never decreased when a player shifts inflow from an inactive edge to an active edge, since flow sent into inactive edges does not arrive in time. Note here that due to



■ **Figure 4** A network with two non-symmetric players. We have $t = t_1 = t_2$. Player 1 must commit to a split of her inflow rate at time θ before knowing the relevant information on the congestion on e_1 and e_2 at time $\theta + 10$. Player 2 might use this to her advantage, which could lead to the non-existence of Nash equilibria.

the flow dynamics, the cumulative outflow function of an edge depends non-decreasingly on the cumulative inflow function of the edge (more precisely, $F_{j,e}^-(\theta) \geq \hat{F}_{j,e}^-(\theta)$ for all $\theta \in [0, H]$ if $F_{j,e}^+(\theta) \geq \hat{F}_{j,e}^+(\theta)$ for all $\theta \in [0, H]$). Suppose we have a strategy profile g such that the total payoff $\sum_{j \in P} \rho_j$ is strictly smaller than OPT. Either one of the players' shares of the total payoff is strictly less than $d_j / \sum_{j' \in P} d_{j'}$, so this player can improve, or all players have a share of $d_j / \sum_{j' \in P} d_{j'}$. But then there has to be an edge e where flow is wasted, which means that the queue on e is empty and the total inflow rate is strictly smaller than the capacity for a time span of positive measure when e is still active; instead, flow is sent into inactive edges or edges with a queue, which hinders later flow to get to t in time. Hence, j can improve by shifting flow from an inactive edge or edge with a queue, as this increases the total payoff but does not decrease her share. Hence, in both cases g was not a Nash equilibrium, which shows that the price of anarchy is 1. ◀

5 Further Research

The topic opens up a multitude of further research directions. First of all, either proving or disproving the existence of Nash equilibria for more general networks for games with information on the exit times. As the exit times do not cover all information that might be necessary for the players to react, a responding player might have an advantage. This is especially critical in games with non-symmetric players. To illustrate this difficulty consider the network in Figure 4.

For this reason an interesting research direction would be to identify more general network classes for which a Nash equilibrium still exists. Symmetric games where all players share the same origin and the same destination might be a necessary restriction.

Additionally, the dependencies among different objective functions have not yet been understood very well: Shedding light on whether the results for the maximum flow over time objective translate to a quickest flow objective (maybe via binary-search) would be very interesting. It might even be possible to consider the average arrived flow or the average arrival time as payoff functions, which could then be compared to an earliest arrival flow.

Finally, cooperative and non-cooperative models might be mixed in order to assess how coalitions and selfish particles behave together (as e.g. in [7] for the static case).

References

- 1 Eitan Altman, Tamer Bacsar, Tania Jimenez, and Nahum Shimkin. Competitive routing in networks with polynomial costs. *IEEE Transactions on Automatic Control*, 47(1):92–96, 2002.
- 2 Mor Armony and Nicholas Bambos. Queueing dynamics and maximal throughput scheduling in switched processing systems. *Queueing systems*, 44(3):209–252, 2003.

- 3 U. Bhaskar, L. Fleischer, and E. Anshelevich. A stackelberg strategy for routing flow over time. *Games and Economic Behavior*, 92:232–247, 2015.
- 4 Umang Bhaskar, Lisa Fleischer, Darrell Hoy, and Chien-Chung Huang. On the uniqueness of equilibrium in atomic splittable routing games. *Mathematics of Operations Research*, 40(3):634–654, 2015.
- 5 Umang Bhaskar and Phani Raj Lolakapuri. Equilibrium computation in atomic splittable routing games. In *26th Annual European Symposium on Algorithms (ESA 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- 6 R. E. Burkard, K. Dlaska, and B. Klinz. The quickest flow problem. *Zeitschrift für Operations Research*, 37(1):31–58, 1993.
- 7 Stefano Catoni and Stefano Pallottino. Traffic equilibrium paradoxes. *Transportation Science*, 25(3):240–244, 1991.
- 8 R. Cominetti, J. Correa, and O. Larré. Existence and uniqueness of equilibria for flows over time. In *International Colloquium on Automata, Languages, and Programming*, pages 552–563. Springer, 2011.
- 9 R. Cominetti, J. Correa, and N. Olver. Long-term behavior of dynamic equilibria in fluid queuing networks. *Operations Research*, 2021.
- 10 Roberto Cominetti, José R. Correa, and Omar Larré. Dynamic equilibria in fluid queueing networks. *Operations Research*, 63(1):21–34, 2015.
- 11 Roberto Cominetti, José R. Correa, and Nicolás E. Stier-Moses. The impact of oligopolistic competition in networks. *Operations Research*, 57(6):1421–1437, 2009.
- 12 José Correa, Andrés Cristi, and Tim Oosterwijk. On the price of anarchy for flows over time. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC '19*, pages 559–577, New York, 2019. Association for Computing Machinery.
- 13 José Correa and Nicolás E. Stier-Moses. Wardrop equilibria. *Wiley encyclopedia of operations research and management science*, 2010.
- 14 Lisa Fleischer and Éva Tardos. Efficient continuous-time dynamic network flow algorithms. *Operations Research Letters*, 23(3):71–80, 1998.
- 15 L. R. Ford and D. R. Fulkerson. Constructing maximal dynamic flows from static flows. *Operations research*, 6:419–433, 1958.
- 16 L. R. Ford and D. R. Fulkerson. *Flows in Networks*. Princeton University Press, 1962.
- 17 David Gale. Transient flows in networks. *Michigan Mathematical Journal*, 6(1):59–63, 1959.
- 18 Lukas Graf and Tobias Harks. The price of anarchy for instantaneous dynamic equilibria. In *International Conference on Web and Internet Economics*, pages 237–251. Springer, 2020.
- 19 Lukas Graf and Tobias Harks. A finite time combinatorial algorithm for instantaneous dynamic equilibrium flows. *Mathematical Programming*, pages 1–32, 2022.
- 20 Lukas Graf, Tobias Harks, and Leon Sering. Dynamic flows with adaptive route choice. *Mathematical Programming*, 2020.
- 21 Tobias Harks. Stackelberg strategies and collusion in network games with splittable flow. *Theory of Computing Systems*, 48(4):781–802, 2011.
- 22 Tobias Harks, Britta Peis, Daniel Schmand, Bjoern Tauer, and Laura Vargas Koch. Competitive packet routing with priority lists. *ACM Transactions on Economics and Computation (TEAC)*, 6(1):4, 2018.
- 23 Tobias Harks and Veerle Timmermans. Equilibrium computation in atomic splittable singleton congestion games. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 442–454. Springer, 2017.
- 24 Tobias Harks and Veerle Timmermans. Uniqueness of equilibria in atomic splittable polymatroid congestion games. *Journal of Combinatorial Optimization*, 36(3):812–830, 2018.
- 25 Alain Haurie and Patrice Marcotte. On the relationship between nash-cournot and wardrop equilibria. *Networks*, 15(3):295–308, 1985.

- 26 Martin Hoefer, Vahab S Mirrokni, Heiko Röglin, and Shang-Hua Teng. Competitive routing over time. In *International Workshop on Internet and Network Economics*, pages 18–29. Springer, 2009.
- 27 J. Israel and L. Sering. The impact of spillback on the price of anarchy for flows over time. In *International Symposium on Algorithmic Game Theory*, pages 114–129. Springer, 2020.
- 28 Max Klimm and Philipp Warode. Complexity and parametric computation of equilibria in atomic splittable congestion games via weighted block laplacians. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2728–2747. SIAM, 2020.
- 29 Ronald Koch and Martin Skutella. Nash equilibria and the price of anarchy for flows over time. *Theory of Computing Systems*, 49(1):71–97, 2011.
- 30 Patrice Marcotte. Algorithms for the network oligopoly problem. *Journal of the Operational Research Society*, 38(11):1051–1065, 1987.
- 31 E. Minieka. Maximal, lexicographic, and dynamic network flows. *Operations Research*, 21(2):517–527, 1973.
- 32 N. Olver, L. Sering, and L. Vargas Koch. Continuity, uniqueness and long-term behavior of Nash flows over time. In *2021 IEEE 62st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2021.
- 33 Britta Peis, Bjoern Tauer, Veerle Timmermans, and Laura Vargas Koch. Oligopolistic competitive packet routing. In *18th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- 34 H. Pham and L. Sering. Dynamic equilibria in time-varying networks. In *International Symposium on Algorithmic Game Theory*, pages 130–145. Springer, 2020.
- 35 J Ben Rosen. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica: Journal of the Econometric Society*, pages 520–534, 1965.
- 36 Tim Roughgarden. Selfish routing with atomic players. *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '05)*, pages 1184–1185, 2005.
- 37 Tim Roughgarden and Florian Schoppmann. Local smoothness and the price of anarchy in splittable congestion games. *Journal of Economic Theory*, 156:317–342, 2015.
- 38 Tim Roughgarden and Éva Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.
- 39 L. Sering. *Nash flows over time*. Technische Universitaet Berlin (Germany), 2020.
- 40 L. Sering and M. Skutella. Multi-source multi-sink Nash flows over time. In *18th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems*, volume 65, pages 12:1–12:20, 2018.
- 41 L. Sering and L. Vargas Koch. Nash flows over time with spillback. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 935–945. SIAM, 2019.
- 42 Jiří Sgall. Open problems in throughput scheduling. In *European Symposium on Algorithms*, pages 2–11. Springer, 2012.
- 43 Martin Skutella. An introduction to network flows over time. In *Research trends in combinatorial optimization*, pages 451–482. Springer, 2009.
- 44 W. S. Vickrey. Congestion theory and transport investment. *The American Economic Review*, 59(2):251–260, 1969. URL: <http://www.jstor.org/stable/1823678>.
- 45 John Glen Wardrop. Some theoretical aspects of road traffic research. *Proceedings of the Institution of Civil Engineers*, 1(3):325–362, 1952.