

Original citation:

Cao, Zhigang, Chen, Bo, Chen, Xujin and Wang, Changjun (2017) A network game of dynamic traffic. ACM Transactions on Economics and Computation.

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/88188>

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions. Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

"© ACM, 2017. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. <http://doi.org/10.1145/3033274.3085101>

A note on versions:

The version presented here may differ from the published version or, version of record, if you wish to cite this item you are advised to consult the publisher's version. Please see the 'permanent WRAP url' above for details on accessing the published version and note that access may require a subscription.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

A Network Game of Dynamic Traffic *

Zhigang Cao[†] Bo Chen[‡] Xujin Chen[§] Changjun Wang[¶]

May 5, 2017

Abstract

We study a network congestion game of discrete-time dynamic traffic of atomic agents with a single origin-destination pair. Any agent freely makes a dynamic decision at each vertex (e.g., road crossing) and traffic is regulated with given priorities on edges (e.g., road segments). We first constructively prove that there always exists a subgame perfect equilibrium (SPE) in this game. We then study the relationship between this model and a simplified model, in which agents select and fix an origin-destination path simultaneously. We show that the set of Nash equilibrium (NE) flows of the simplified model is a proper subset of the set of SPE flows of our main model. We prove that each NE is also a strong NE and hence weakly Pareto optimal. We establish several other nice properties of NE flows, including global First-In-First-Out. Then for two classes of networks, including series-parallel ones, we show that the queue lengths at equilibrium are bounded at any given instance, which means the price of anarchy of any given game instance is bounded, provided that the inflow size never exceeds the network capacity.

1 Introduction

Selfish routing is one of the fundamental models in the study of network traffic systems [25, 26, 31]. Most of its literature assumes *essentially static* traffic flows, failing to capture many dynamic characteristics of traffics, despite the usual justification that static flows are good approximations of steady states of the traffic system. Recently, models of *dynamic* flows have been drawing much attention in various research areas of selfish routing [4, 15, 18, 27, 32].

*This work was supported in part by the National Natural Science Foundation of China (NSFC), under grant numbers 11601022, 11531014 and 11471326.

[†]Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Beijing, 100190, China, Email: zhigangcao@amss.ac.cn

[‡]Warwick Business School, University of Warwick, Coventry, CV4 7AL, United Kingdom, Email: Bo.Chen@wbs.ac.uk

[§]Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Beijing, 100190, China, Email: xchen@amss.ac.cn

[¶]Beijing Institute for Scientific and Engineering Computing, Beijing University of Technology, Beijing, 100124, China, Email: wcj@bjut.edu.cn

1.1 Model

Let us start with an informal and intuitive description of our model, with a formal definition postponed to Section 2. We are given an acyclic directed network, with two special vertices called the origin o and the destination d , respectively, such that each edge is on at least one o - d path. Each edge e of this network is associated with an integer capacity c_e and a positive integer free-flow transit time t_e . Time horizon is infinite and discretized as $1, 2, \dots$. At each time point, a set of selfish agents enter the network from the origin o , trying to reach the destination d as quickly as possible. When an agent uses an edge e , two costs are incurred to him: a fixed transit cost t_e and a variable waiting cost dependent on the volume of the traffic flow on e and the capacity c_e as well as the agent's position in the queue of agents waiting at e . The total cost to each agent, also referred to as his *latency*, is the sum of the two costs on all edges he uses (which form an o - d path in the network). The time an agent reaches the destination is simply his departure time plus his latency.

A critical setting of our model is how the waiting time of each agent is determined in a way called *deterministic queuing*. At each time step, for each edge e of the network, there is a (possibly empty) queue of agents waiting at the tail of e , and at most a capacity number c_e of agents who have the highest priority (according to the queue order) can start to move along e . After a fixed period of transit time t_e , each of these agents reaches the head of e , and enters the next edge (waiting at its tail for traveling along it) unless the destination d has been reached. The queue at each edge is updated according to the First-In-First-Out (FIFO) rule and pre-specified Edge Priorities – if two agents enter an edge e simultaneously, their queue priorities are determined by the priorities of the preceding edges from which they enter e . To break ties when agents enter an edge simultaneously from the same edge, we divide each edge e into a number c_e of lanes, and assign priorities among these lanes. Initially, the agents who enter the network at the same time are associated with original priorities *among them*, which are *temporarily valid* only when they enter the network. This is a crucial difference between our model and that in [15, 27] where pre-specified priorities *among all agents* are *permanent*, working globally throughout the game.

Each agent makes a routing decision at *every* nonterminal vertex he reaches. This is the second crucial and fundamental difference of our model from almost all in the literature [15, 27, 32], in which agents are assumed to make their routing decisions *only* at the origin o as to which o - d path to take. In addition, it is usually assumed that agents' decisions upon entering the system are independent of the decisions of those agents who are currently in the system. In game-theoretic terminology, while the dynamic traffic problem is usually modeled as a game of normal form (or a simultaneous-move game), we model it as a game of extensive form, which allows for agents' *online* adjustments and error corrections.

1.2 Results

Let Γ denote our *extensive-form* game of dynamic traffic, for which we apply the solution concept of (pure strategy) *Subgame Perfect Equilibrium* (SPE). The analysis of SPE of Γ is much more complicated than that of *Nash Equilibrium* (NE) of the *normal-form* game, denoted as Γ^n , which corresponds to the simplified setting where agents choose their one-off o - d paths simultaneously

at the very start of the game. We relate our study of SPEs for Γ to that of NEs for Γ^n , and establish four sets of results as specified below.

SPE existence in Γ . We demonstrate that Γ always possesses a very special SPE (which particularly implies the *NE existence* of Γ^n). To the best of our knowledge, this is the first SPE existence result in the field of dynamic atomic routing games. Since typically Γ is not finite and multiple agents may move simultaneously at each time step, the usual backward induction method does not work here. Instead, we construct iteratively an intuitively simple SPE such that, *in any subgame of Γ* , the path profile realized according to the SPE is *iteratively dominating* in the following sense: given dominating paths determined for a set D of agents (initially none), the next dominating path determined for another agent i (from his current location to the destination d) satisfies the property that as long as agents in $D \cup \{i\}$ follow their dominating paths, no agent outside $D \cup \{i\}$ can overtake i at any vertex of i 's dominating path (particularly the destination d), regardless of the choices of all the other agents.

Equilibrium relationship between Γ^n and Γ . Given any NE in Γ^n , we construct an SPE of Γ such that its realized path profile is the given NE. This result along with a counterintuitive example shows that the NE set of Γ^n is a *proper* subset of the SPE outcome set of Γ . On the one hand, the proper containment reaffirms the intuition that model Γ is more flexible than Γ^n . On the other hand, the result builds a useful bridge between the more frequently analyzed model Γ^n and the more realistic but much more complicated model Γ .

NE properties in Γ^n . We show that in Γ^n the set of NEs and that of *strong NEs* are identical, which particularly implies that all NEs of Γ^n are *weakly Pareto Optimal*. The result is a corollary of the following stronger properties for any NE routing of Γ^n . Let agents be grouped according to their arrival times at the destination d .

- *Hierarchal Independence*: The arrival times at any vertex of agents in earlier groups are independent of *any* choices of agents in later groups (even if the later agents do not care about their own latencies and move in a coordinated way).
- *Hierarchal Optimality*: The equilibrium latency of any agent is minimum for him among all routings in which all agents in earlier groups follow their NE paths.
- *Global FIFO*: vertex of the network (particularly the origin o), then the former will reach the destination d no later than the latter.

In contrast to the above global property, FIFO in the literature usually serves as an *assumption* to regulate the queue on each *edge*. We would also like to emphasize that the satisfactions of all the above nice properties are by *all NEs on all networks*, while usually in the related literature only properties of some *special NEs* are studied.

The above three sets of results remain valid even if the network has *multiple origins* and a single destination. This is one of the key reasons that we are able to prove the existence of an SPE for Γ (note that the off-equilibrium scenarios are essentially new problems with multiple origins).

Bounded queue lengths. Since there may be infinitely many agents in dynamic routing models like ours, a natural question is whether the queue length at each edge at equilibria can be bounded by a constant that depends *only* on the input *finite* network, when the inflow size never exceeds the minimum capacity of an *o - d* cut in the network. This interesting theoretical challenge has been listed as an important open problem in the literature [4, 27]. For a special class of series-parallel networks, the so-called chain-of-parallel networks, Scarsini et al. [27] proved that the queue lengths at UFR equilibrium (i.e., NE with earliest arrivals) are bounded for seasonal inflows.

We prove that equilibrium queue lengths are indeed bounded in our games on two classes of networks: Γ^n on series-parallel networks, and Γ on networks in which the in-degree does not exceed the out-degree at any nonterminal vertex. The result of bounded queue length implies that the price of anarchy (PoA) and price of stability (PoS) w.r.t. certain naturally defined social welfare measures [19, 28] are also bounded for the two classes of networks.

1.3 Related literature

Classical models of dynamic flows, a.k.a. flows over time, were pioneered by Ford and Fulkerson [9, 10], who studied in a discrete-time setting the problem of maximizing the amount of goods that can be transported from the origin to the destination within a given period of time. They showed that this dynamic problem can be reduced to the standard (static) minimum-cost flow problem and thus polynomially solvable. Philpott [24] and Fleischer and Tardos [8] showed that many results can be carried over to the continuous counterpart. More recent developments along this line of research can be found in [7, 29].

Games of dynamic flows, a.k.a. routing games over time, were initially studied by [30, 33], which focused on analyzing Nash equilibria for small-sized concrete examples. As summarized by [17, 23], most subsequent studies can be classified into four categories in terms of methodologies: mathematical programming, optimal control, variational inequalities, and simulation-based approaches. Variational inequality formulations [11] turn out to be the most successful in investigating Nash equilibria of the *simultaneous departure-time route-choice models*, where each player has to make a decision on his departure time as well. Unfortunately, little is known about the existence, uniqueness and characterizations of equilibria under the general formulation [22].

Our model is a variant of the *deterministic queuing* (DQ) model which was introduced by [30], developed by [13], and recently revived by [17]. The notational and inconsequential differences are in the locations used to store queues – the standard DQ model uses head parts of edges, while our storages, as well as those in [4], are located on tail parts. The DQ model usually assumes that the waiting queue has no physical length, and the problem seems to be harder to analyze if otherwise assumed [5].

Non-atomic dynamic traffic games. Koch and Skutella [17, 18] are the first to use a DQ model to study non-atomic dynamic flow games. They investigated the continuous-time single-origin single-destination case with uniform inflow rates, i.e., the *temporal routing game* as named by [3]. It was shown in [18] that the Nash equilibria, i.e., Nash flows over time, can be viewed as concatenations of special static flows, the so-called *thin flows with resetting*. The Nash flow was characterized by the so-called *universal FIFO condition* that no flow overtakes another, and

equivalently by an analogue to Wardrop principle that flow is only sent along currently shortest paths. Cominetti et al. [4] proved in a constructive way the existence and uniqueness of the Nash equilibrium of temporal routing games in a more general setting with piecewise constant inflow rates. For multi-origin multi-destination cases, Cominetti et al. [4] proved in a nonconstructive way the NE existence when the inflow rates belong to the space of p -integrable functions with $1 < p < \infty$.

Building on the temporal routing model and results of [18], Bhaskar et al. [3] investigated the price of anarchy in Stackelberg routing games, where the network manager acts as the leader picking a capacity for each edge within its given limit, and the agents act as followers, each of whom picks a path. The authors gave a polynomial-time computable strategy for the leader in an acyclic directed network to reduce edge capacities such that the PoA w.r.t. minimum completion time is $e/(e - 1)$, and the PoA w.r.t. minimum total delay is upper bounded by $2e/(e - 1)$, provided that the original network is saturated by its earliest arrival flow. Macko et al. [21] showed that the PoA of the temporal routing game w.r.t. to the minimum maximum delay can be as large as $n - 1$ in networks with n vertices. Anshelevich and Ukkusuri [1] considered a dynamic routing game whose monotone increasing edge-delay functions are more general than DQ models but still obey the local FIFO principle. They studied how a single splittable flow unit present at each origin at time 0 would travel across the network to the corresponding destination, assuming that each flow particle is controlled by a different agent. They showed that in the single-origin single-destination case, there is a unique Nash equilibrium, and the efficiency of this equilibrium can be arbitrarily bad; in the multi-origin multi-destination case, however, the existence of an NE is not guaranteed.

Atomic dynamic traffic games. The most related work to ours is [27], with several notable differences. First, to break ties Scarsini et al. [27] placed priorities on agents rather than on edges. Second, they only allowed each agent to make a decision at the origin rather than at each intermediary vertex. Third, they focused on seasonal inflows and how the transient phases impact the long-run steady outcomes, whereas their notion of steady outcome does not apply in our model because the inflows we consider are not restricted to be seasonal. Finally, they concentrated on a special kind of NE named *Uniformly Fastest Route* (UFR) equilibrium, which is an NE such that each agent reaches all vertices on his route as early as possible. Using an argument similar to that in [32], Scarsini et al. [27] proved that the game admits at least one UFR equilibrium. When the inflow is a constant not exceeding the capacity of the network, they characterized the flows and costs generated by optimal routings in general networks, and those generated by UFR equilibria in parallel networks and more general chain-of-parallel networks. The results on parallel networks were extended to the setting with seasonal inflows.

In [32], more variants of atomic games of dynamic traffic were considered for finitely many agents under discrete-time DQ models. Apart from the sum-type of latency functions as considered in [27] and in this paper, Werth et al. [32] also studied the bottleneck-type objective functions for agents, where the cost of each agent equals his expense on the slowest edge of his chosen path. To break ties, the global priorities placed on agents as in [27] were discussed for both the sum-objective and bottleneck-objective models, while the local priorities placed on edges as in this paper were investigated only for the bottleneck-objective model. Werth et al. [32] focused on computational issues on NE and optimization problems, while we concentrate

on SPE existence and NE properties.

As one of the earliest papers studying dynamic atomic routing games, [14, 15] was concerned with computational complexity properties of NE and best-response strategies of a finite network congestion model, where each edge of the network is viewed as a machine with a processing speed and a scheduling policy, and each agent is viewed as a task with a positive length which has to be processed by the machines one after another along the path the agent chooses. While the positive transit costs of agents are determined by machine speeds and task lengths, their waiting costs are determined by scheduling policies. Apart from the (local) FIFO policy (on each machine) that tasks are processed non-preemptively in order of arrival, the policy of (non)-preemptive global ranking and that of fair time-sharing were also investigated. For agents (tasks) with uniform lengths in a single-origin directed network, when the FIFO policy is coupled with the global agent priorities for tie-breaking, the network congestion game turns out to be a generalization of the sum-objective model of [32], and admit a strong NE which can be computed efficiently. Somewhat surprisingly, for the more restrictive case of single-origin single-destination networks, computing best responses is NP-hard. Kulkarni and Mirrokni [20] studied the robust PoA of a generalization of the model in [15] under the so-called *highest-density-first* forwarding policy.

The global priority scheme introduced in [6] was considered by Harks et al. [12] for a dynamic route-choice game under the discrete-time DQ model without the FIFO queuing rule on each edge. They analyzed the impact of agent priority ordering on the efficiency of NEs for minimizing the total latency of all k agents, as well as their computability. They proved that the asymmetric game has its PoS in $\Omega(\sqrt{k})$ and PoA upper bounded by $1 + k^3/2$, while the symmetric game has its PoS and PoA equal to 1 and $(k + 1)/2$, respectively. They also showed that an NE and each agent's best response are polynomially computable. Koch [16] analyzed the dynamic atomic routing game on a restricted class of the discrete DQ model with global agent priorities, where each edge has zero free-flow transit time and unit capacity.

The rest of this paper is organized as follows. In Section 2 we present a formal definition of our network game model along with its extensive-form setting Γ and normal-form game setting Γ^n . In Section 3 we establish the existence of an SPE in Γ . In Section 4 we study the properties of all NEs in Γ^n , and then investigate the relationship between the equilibrium flows of Γ and Γ^n . In Section 5 we bound the lengths of equilibrium queues for two special classes of networks. Finally, we conclude the paper with remarks on future research directions in Section 6. All missing proofs and details are provided in the appendix.

2 The Model

In this section we formally present our network game model for dynamic traffic followed with some necessary notations and preliminaries. By edge subdivisions and duplications, we assume w.l.o.g. in the rest of this paper that *each edge of the networks has a unit capacity and a unit length*.

2.1 Dynamic traffic in directed networks

Let G be a finite *acyclic* directed multi-graph with two distinguished vertices, the origin o and the destination d , such that every edge of E is on some o - d path in G . The dynamic network traffic is represented by *players* moving through G . The infinite time is discretized as $0, 1, 2, \dots$ and at each integer time point $r \geq 1$, a (possibly empty) set Δ_r of finitely many players enter G from their common origin o and each of them will go through some o - d path in G and leave G from d . Each player, when reaching a vertex v ($\neq d$), *immediately* selects an edge e outgoing from v and enters e *at once*. At any *integer* time point r , all players (if any) who have entered e but not exited yet *queue at the tail part of e* , and only the *unique* head of the queue will leave the queue. This queue head will spend one time unit in traversing e from its tail to its head and exit e at time $r + 1$.

An extended infinite network. For ease of description, we make a technical extension of G to accommodate all players of $\Delta := \cup_{r \geq 1} \Delta_r$ at the very beginning. Suppose that $|\Delta_r|$ is upper bounded by a fixed integer F for all $r \geq 1$. We construct an infinite network \bar{G} from G as follows: add a new vertex \bar{o} , and connect \bar{o} to o with F internally disjoint \bar{o} - o paths: $P^f := \bar{o} \cdots o_r^f o_{r-1}^f \cdots o_2^f o_1^f o$, $f = 1, 2, \dots, F$. Each path P^f has a set of an infinite number of edges $\{o_{r+1}^f o_r^f \mid r \geq 0\}$, where $o_0^f = o$, and intersects G only at o . At time 0, all players of Δ are queuing at and sets off from distinct edges in $\cup_{f=1}^F P^f \subseteq \bar{G}$ (so they are all queue heads) such that for any $r \geq 1$, players of Δ_r are on the tail parts of edges $o_r^f o_{r-1}^f$, $f = 1, \dots, |\Delta_r|$, respectively; they are all at a distance r from vertex o , and set off from these edges for their common destination d . The traffic on \bar{G} naturally corresponds to the one on G in a way that the restriction of the former to its subnetwork G is exactly the latter—for any $r \geq 1$, all players in Δ_r reach o at time r .

For every $v \in \bar{V}$, we use $\bar{E}^+(v)$ and $\bar{E}^-(v)$ to denote the set of outgoing edges from v and the set of incoming edges to v in \bar{G} , respectively. For every $e \in \bar{E}$, we use $u_e, v_e \in \bar{V}$ to denote the tail and head of e , respectively, i.e., $e = u_e v_e$.

Traffic regulation. A complete priority order \prec_v is pre-specified over all edges in $\bar{E}^-(v)$ for any $v \in \bar{V}$ ($e_1 \prec_v e_2$ means that e_1 has a higher priority than e_2). If players queue at (the tail part of) edge e , they are prioritized for entering and therefore exiting the queue according to the following *queuing rule*: For any pair of players, whoever enters e earlier has a higher priority; if they enter e at the same time, then they must do so through two different edges of $\bar{E}^-(u_e)$, in which case their priorities are determined by the priority order \prec_{u_e} on the two edges.

Configurations. We use Q_e^r to denote the queue on edge e at time r , which is both a sequence of players and the corresponding set. We call $c_r = (Q_e^r)_{e \in \bar{E}}$ a *configuration* at time r if $Q_e^r \cap Q_{e'}^r = \emptyset$ for different e and e' . Throughout this paper,

- $\Delta(c_r) := \cup_{e \in \bar{E}} Q_e^r$ denotes the set of players involved in configuration c_r .
- c_0 denotes the unique *initial configuration* given by queues at time 0.

Let \mathcal{C}_r denote the set of configurations at time r .

Action sets. Given configuration $c_r = (Q_e^r)_{e \in \bar{E}}$, the *action set* of player $i \in Q_e^r$, denoted $\bar{E}(i, c_r)$, is defined as follows:

- If the head v_e of e is the destination d and i queues first in Q_e^r , then $\bar{E}(i, c_r) := \emptyset$, i.e., i simply exits the system at time $r + 1$;
- If $v_e \neq d$ and i queues first in Q_e^r , then $\bar{E}(i, c_r) := \bar{E}^+(v_e)$, i.e., player i selects the next edge that is available at v_e ;
- Otherwise (i.e., i is not the head of Q_e^r), player i has to stay on e with $\bar{E}(i, c_r) := \{e\}$.

Consecutive configurations. Given a configuration c_r and an action profile $\mathbf{a} = (a_i)_{i \in \Delta(c_r)}$ with $a_i \in \bar{E}(i, c_r)$, the traffic rule leads to a new configuration $c_{r+1} = (Q_e^{r+1})_{e \in \bar{E}}$ at time $r + 1$, referred to as a *consecutive configuration* of c_r :

- As a set, $Q_e^{r+1} = \{i \in \Delta(c_r) \mid a_i = e\}$ consists of players choosing e in action profile \mathbf{a} .
- As a sequence, Q_e^{r+1} equals Q_e^r with its head removed followed by $Q_e^{r+1} \setminus Q_e^r$ whose positions are according to the traffic regulation priority \prec_{u_e} at the tail vertex u_e of edge e .

2.2 Extensive form game setting Γ

Given a dynamic traffic problem on network \bar{G} with initial configuration c_0 (which is equivalent to the problem on network G with incoming flows $\{\Delta_r : r \geq 0\}$), we study a natural extensive-form game $\Gamma = \Gamma(\bar{G})$ as specified below. Given any nonnegative integer k , we write $[k]$ for the set of positive integers at most k . Particularly $[0] = \emptyset$.

Histories. For each time point $r \geq 0$, a sequence of consecutive configurations $h_r = (c_0, c_1, \dots, c_r)$ starting from the initial configuration c_0 is called a *history* at time r . The set of all possible histories at time r is denoted as \mathcal{H}_r .

Strategies. Each player $i \in \Delta$ needs to make a decision at every configuration in histories $h_r = (c_0, \dots, c_r)$ with $i \in \Delta(c_r)$.¹ The *strategy* of player $i \in \Delta$ is a mapping σ_i that maps each history $h_r = (c_0, \dots, c_r)$ with $i \in \Delta(c_r)$ to an edge $\sigma_i(h_r) \in \bar{E}(i, c_r)$. The *strategy set* of player i is denoted as Σ_i . A vector $\sigma = (\sigma_i)_{i \in \Delta}$ is called a *strategy profile* of Γ .

Latencies. Each strategy profile σ of Γ gives each player $i \in \Delta$ a latency equal to his exiting time, denoted as $t_i(\sigma)$, from the system, i.e., the length of time he stays in \bar{G} . Each player tries to minimize his latency, which is equivalent to minimizing his travelling time from o to d .²

¹Since \bar{G} is acyclic, i stays in \bar{G} for a finite period of time, i.e., $i \notin \Delta(c_s)$ for all c_s with large enough s .

²Note that for any time r and any player in Δ_r , his arrival time at o will always be r under any strategy profile σ .

Game tree. The game tree of Γ is typically an infinite tree with nodes corresponding to histories. At each game tree node $h_r = (c_0, c_1, \dots, c_r)$, players in $\Delta(c_r)$ need to make decisions simultaneously, and their action profile leads to a new node h_{r+1} , which is a child of h_r . Each subtree of the game tree rooted at history h_r can be viewed as a separate game, referred to as a *subgame* of Γ . The restriction of each strategy σ_i to a subgame tree is also a strategy of the subgame. Given strategy profile σ of Γ , the time when player $i \in \Delta(c_r)$ exits \bar{G} in the subgame started from h_r is denoted as $t_i(\sigma|h_r)$.

Definition 1 (Subgame Perfect Equilibrium). A strategy profile $\sigma = (\sigma_i)_{i \in \Delta}$ is a *Subgame Perfect Equilibrium* (SPE) of Γ if for any $r \geq 0$ and any history $h_r \in \mathcal{H}_r$, $t_i(\sigma|h_r) \leq t_i(\sigma'_i, \sigma_{-i}|h_r)$ holds for all $i \in \Delta(c_r)$ and $\sigma'_i \in \Sigma_i$, where σ_{-i} is the partial strategy profile of players other than i .

2.3 Normal form game setting Γ^n

A variant game model Γ^n is more popular in the literature and will also be studied in this paper. Γ^n serves both as an independent model and as a technical approach facilitating our analysis of the main model Γ . The key assumption in Γ^n is that all players select an origin-destination path simultaneously at time 0. They should follow the chosen paths and are not allowed to deviate at any vertex. All the other settings are the same as in game Γ .

We frequently analyze the *interim game* $\Gamma^n(c_r)$ of Γ^n for each configuration $c_r = (Q_e^r)_{e \in \bar{E}}$ defined as follows. The player set is $\Delta(c_r)$. For each player $i \in \Delta(c_r)$, suppose $e_i(c_r)$ is the edge at which i queues, i.e., $i \in Q_{e_i(c_r)}^r$, and $o_i(c_r)$ is the tail of $e_i(c_r)$. The strategy set of player $i \in \Delta(c_r)$, denoted as $\mathcal{P}_i[o_i(c_r), d]$, is the set of $o_i(c_r)$ - d paths containing $e_i(c_r)$.

For any player $i \in \Delta(c_r)$ and any path profile $\mathbf{p} = (P_i)_{i \in \Delta(c_r)}$ with $P_i \in \mathcal{P}[o_i(c_r), d]$, we use $t_i^d(\mathbf{p})$ to denote the arrival time of i at destination d under the corresponding routing determined by \mathbf{p} .

Definition 2 (Nash Equilibrium). A path profile $\boldsymbol{\pi}$ of $\Delta(c_r)$ is a *Nash Equilibrium* (NE) of $\Gamma^n(c_r)$ if no player can gain by unilateral deviation, i.e., $t_i^d(\boldsymbol{\pi}) \leq t_i^d(P_i, \boldsymbol{\pi}_{-i})$ for all $i \in \Delta(c_r)$ and $P_i \in \mathcal{P}[o_i(c_r), d]$, where $\boldsymbol{\pi}_{-i}$ is the partial path profile of players in $\Delta(c_r)$ other than i .

2.4 A warmup example

Each strategy profile σ of Γ induces a realized path profile, which is a strategy profile of Γ^n . When σ is an SPE of Γ , Example 1 below shows that the induced path profile is not necessarily an NE of Γ^n , demonstrating that richer phenomena may be observed in Γ than in Γ^n . In our illustrations, we only show the restrictions of games Γ and Γ^n to G , which we denote as $\Gamma(G)$ and $\Gamma^n(G)$ respectively. Accordingly, the *travelling cost* a player spends in G from o to d is his arrival time at d minus that at o .

Example 1. The network G is as depicted in Figure 1. Edge e_1 (resp. e_2) has a higher priority than e_3 (resp. e_4). There are only two players in the game, who enter G via origin o at the same time $r = 1$. Player 1 has a higher original priority than player 2. In Γ , player 1 takes a vicious strategy in the following sense. He initially chooses edge ov and then tries to block

player 2 by choosing $v \rightarrow w_1 \rightarrow d$ if player 2 used edge ou_1 and $v \rightarrow w_2 \rightarrow d$ otherwise. Player 2 always follows $o \rightarrow u_1 \rightarrow w_1 \rightarrow d$. It is easy to check that these yield a strategy profile that is an SPE of $\Gamma(G)$ (off-equilibrium behaviors of the two players can be easily defined), incurring a travelling cost 3 for player 1 and 4 for player 2. However, the induced path profile, ovw_1d for player 1 and ou_1w_1d for player 2, is not an NE of $\Gamma^n(G)$. Indeed, $\Gamma^n(G)$ admits in total six NEs, all bringing the two players the same travelling cost of 3.

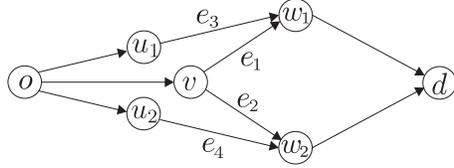


Figure 1: An SPE of Γ may not induce an NE of Γ^n .

3 Existence of SPE

In this section, we establish the existence of an SPE for game $\Gamma = \Gamma(\bar{G})$ by using an algorithmic method. Given any feasible configuration c_r , our algorithm (Algorithm 1 below) computes a path profile, which is a special NE, of $\Gamma^n(c_r)$. The special NEs computed for all configurations will assemble an SPE of Γ in a way that for each configuration c_r the path profile induced by c_r and the SPE is exactly the NE computed for $\Gamma^n(c_r)$.

The high-level idea behind our algorithm is a natural greedy best-response strategy: iteratively computing the unique best path based on all paths that have been obtained. Similar approaches with noticeable differences have been adopted to deal with various routing games, e.g., [15, 27, 32]. The special difficulties in our setting are imposed by subgames that start from any feasible configurations (which are essentially multi-origin problems) and priorities that are placed on edges rather than on players.

For any $S \subseteq \Delta(c_r)$, we also use $\mathbf{p} = (P_i)_{i \in S}$ with $P_i \in \mathcal{P}[o_i(c_r), d]$ to denote a partial path profile, corresponding to a (partial) routing. Given any vertex $v \in \bar{V}$, edge $uv \in \bar{E}$, and player $i \in S$,

- $t_i^v(\mathbf{p})$ (resp. $t_i^{uv}(\mathbf{p})$) denotes the time when i reaches v (resp. enters uv) under \mathbf{p} ;
- $t_i^v(\mathbf{p}) := \infty$ if $v \notin P_i$, and $t_i^{uv}(\mathbf{p}) := \infty$ if $uv \notin P_i$.

If $S' \subseteq S$, we often write $(P_i)_{i \in S'}$ as $\mathbf{p}_{S'}$.

Given a feasible configuration c_r , our algorithm goes roughly as follows. Initially, let player subset D of $\Delta(c_r)$ and partial routing $\boldsymbol{\pi}_D$ of players in D be empty. Then recursively, assuming players in D go along their *selected* paths specified in $\boldsymbol{\pi}_D$, we enlarge D with a new player $i \in \Delta(c_r) \setminus D$ and an associated path $\Pi_i \in \mathcal{P}[o_i(c_r), d]$ in the following way. For each $j \in \Delta(c_r) \setminus D$, define $\tau_j^v = \min\{t_j^v(\boldsymbol{\pi}_D, R_j) \mid R_j \in \mathcal{P}[o_j(c_r), d]\}$ as the “ideal latency” of player j at vertex v w.r.t. $\boldsymbol{\pi}_D$. Let $i \in \Delta(c_r) \setminus D$ and $\Pi_i \in \mathcal{P}[o_i(c_r), d]$ be such that (i) $t_i^d(\boldsymbol{\pi}_D, \Pi_i)$ equals the smallest “ideal latency” at d among all players in $\Delta(c_r) \setminus D$, i.e., $t_i^d(\boldsymbol{\pi}_D, \Pi_i) = \min_{j \in \Delta(c_r) \setminus D} \tau_j^v$, (ii) if there

are more than one of such candidates (i, Π_i) , then the choice of (i, Π_i) is made such that Π_i has the highest priority w.r.t. \prec_d , (iii) if still more than one candidates satisfy (ii), in which case the paths involved must share the same ending edge and suppose u is the tail vertex of this edge, then $t_i^u(\pi_D, \Pi_i[o_i(c_r), u])$ is the smallest. Repeat the above process until $o_i(c_r)$ is reached, in which case all the remaining players must all queue at $e_i(c_r)$ and i is the head of this queue. The process is repeated while D and π_D become larger and larger. A formal description of the process is presented in Algorithm 1.

ALGORITHM 1: (ITERATIVE DOMINATING PATH PROFILE)

Input: a feasible configuration c_r at time r .

Output: the iterative dominating path profile (routing) $\pi = (\Pi_i)_{i \in \Delta(c_r)}$ for $\Delta(c_r)$ along with the corresponding player indices $1, 2, \dots$

1. Initiate $D \leftarrow \emptyset$, $\pi_{[0]} \leftarrow \emptyset$, $i \leftarrow 0$.
 2. $i \leftarrow i + 1$ (NB: Start to search for a new dominator and his dominating path).
 3. For each player $j \in \Delta(c_r) \setminus D$ and vertex $v \in \bar{G}$,
 - let $\tau_j^v = \min\{t_j^v(\pi_{[i-1]}, R_j) \mid R_j \in \mathcal{P}[o_j(c_r), d]\}$ be the earliest time for j to reach vertex v from his *current location in c_r* , assuming that *all other* players in \bar{G} are those in D and they go along their paths specified in $\pi_{[i-1]}$;
 - let \mathcal{P}_j^v denote the set of all the *corresponding* $o_j(c_r)$ - v paths for j to reach v at time τ_j^v .
 - if there is no such path, then set $\tau_j^v \leftarrow \infty$ and $\mathcal{P}_j^v \leftarrow \emptyset$.
 4. $w \leftarrow d$, $D' \leftarrow \Delta(c_r) \setminus D$, $\Pi \leftarrow \emptyset$. (NB: Steps 3–10 are to help identify player i and path Π_i in Steps 11 and 12; D' holds the candidates for i ; Π is a subpath of Π_i that will grow edge by edge starting from d .)
 5. **While** $\tau := \min_{j \in D'} \tau_j^w \geq r + 1$ **do**
 6. $D' \leftarrow \{j \in D' \mid \tau_j^w = \tau\}$;
 7. let uw be the edge of highest priority among all last edges of paths in $\cup_{j \in D'} \mathcal{P}_j^w$;
 8. $\Pi \leftarrow \Pi \cup \{uw\}$;
 9. $w \leftarrow u$;
 10. **End-While** (NB: at the end of the while-loop all players in D' are queuing on the starting edge of Π in configuration c_r)
 11. Let $i \in D'$ be the player who stands first (among all players in D') in line on the starting edge of Π in configuration c_r .
 12. Let i *select* $\Pi_i \leftarrow \Pi$.
 13. $D \leftarrow D \cup \{i\}$, $\pi_{[i]} \leftarrow (\pi_{[i-1]}, \Pi_i)$. (NB: the algorithm outputs player i and his path Π_i .)
 14. Go to Step 2.
-

It is worth noting that, for each $i \geq 1$, Algorithm 1 can identify player i and path Π_i in finite time via ignoring players and vertices that are sufficiently far from the origin o in configuration c_r .

Suppose Algorithm 1 indexes the players of $\Delta(c_r)$ as $1, 2, \dots$ with the associated path profile

$\boldsymbol{\pi} = (\Pi_1, \Pi_2, \dots)$. As can be seen from Lemma 1 below, each player $i \in \Delta(c_r)$ is a *dominator* in $\Delta(c_r) \setminus [i-1]$ and Π_i is a *dominating path* w.r.t. $\boldsymbol{\pi}_{[i-1]}$ in the following sense: under the assumption that players in $[i-1]$ all follow $\boldsymbol{\pi}_{[i-1]}$, as long as i takes Π_i , he will be among the first (within $\Delta(c_r) \setminus [i-1]$) to reach the destination d and no player in $\Delta(c_r) \setminus [i]$ will be able to reach any of Π_i 's intermediary vertices earlier than i does, regardless of the choices of players in $\Delta(c_r) \setminus [i]$. We may also call the above $\boldsymbol{\pi} = (\Pi_1, \Pi_2, \dots)$ an *iterative dominating path profile*.

Lemma 1. *Given any feasible configuration c_r at time r , let players $1, 2, \dots$ of $\Delta(c_r)$ be as indexed and path profile $\boldsymbol{\pi} = (\Pi_i)_{i \in \Delta(c_r)}$ be as computed in Algorithm 1. If players $i, j \in \Delta(c_r)$ and player subset S satisfy $j \in S \subseteq \Delta(c_r) \setminus [i-1]$, then for any vertex $v \in \Pi_i$ and any path profile \boldsymbol{p} of $\Gamma^n(c_r)$, it holds that*

$$t_i^v(\boldsymbol{\pi}_{[i]}, \boldsymbol{p}_{S \setminus \{i\}}) = \min\{t_i^v(\boldsymbol{\pi}_{[i-1]}, R_i) \mid R_i \in \mathcal{P}[o_i(c_r), d]\} \leq t_j^v(\boldsymbol{\pi}_{[i-1]}, \boldsymbol{p}_S).$$

The iterative dominance of $\boldsymbol{\pi}$ particularly implies that given the selected paths in $\boldsymbol{\pi}_{[i-1]}$ of players in $[i-1]$, player i has no incentive to deviate from Π_i for each i , saying that $\boldsymbol{\pi}$ is an NE of $\Gamma^n(c_r)$. Furthermore, a special SPE of Γ can be constructed, building on the iterative dominating path profiles for all feasible configurations.

Theorem 1. *Game Γ admits an SPE.*

Proof. Given any history $h_r = (c_0, \dots, c_r) \in \mathcal{H}_r$ for any time point $r \geq 0$, suppose the players in $\Delta(c_r)$ are named as $1_r, 2_r, \dots$ such that player i_r is the i th player added to D in Step 13 of Algorithm 1 (with input c_r). For each player $i_r \in \Delta(c_r)$, let $\Pi_{i_r}^{c_r}$ denote the path selected by i_r in Algorithm 1; recall that under c_r player i queues at the first edge of $\Pi_{i_r}^{c_r}$. A configuration in \mathcal{C}_{r+1} will result from c_r according to action profile \boldsymbol{a}^{c_r} defined as follows:

$$\text{The action of } i = \begin{cases} \text{the first edge of } \Pi_{i_r}^{c_r}, & \text{if } i \text{ queues after someone else;} \\ \emptyset, & \text{if } i \text{ queues first on the last edge of } \Pi_{i_r}^{c_r}; \\ \text{the second edge of } \Pi_{i_r}^{c_r}, & \text{otherwise.} \end{cases}$$

The set $\cup_{r \geq 0} \cup_{c_r \in \mathcal{C}_r} \boldsymbol{a}^{c_r}$ of action profiles defines a strategy profile $\boldsymbol{\sigma}^* = (\sigma_i^*)_{i \in \Delta}$ of Γ . We will prove that $\boldsymbol{\sigma}^*$ is an SPE of Γ .

Let $(c_r, c_{r+1}, c_{r+2}, \dots)$ be the list of configurations and $(P_i^*)_{i \in \Delta(c_r)}$ be the path profile induced by h_r and $\boldsymbol{\sigma}^*$. It can be deduced from Lemma 1 and Algorithm 1 that

- For any $s \geq r+1$, player sequence $1_s, 2_s, \dots$ is a subsequence of $1_{s-1}, 2_{s-1}, \dots$ such that $\Delta(c_{s-1}) \setminus \Delta(c_s)$ consists of the first $|\Delta(c_{s-1}) \setminus \Delta(c_s)|$ players of $1_{s-1}, 2_{s-1}, \dots$
- For any $s \geq r+1$ and $i \in \Delta(c_s)$, $\Pi_i^{c_s} \subseteq \Pi_i^{c_{s-1}}$.

Therefore, the path P_i^* formed by the actions of each player $i \in \Delta(c_r)$ is exactly $\Pi_i^{c_r}$. According to Lemma 1, we have, for any $i \geq 1$,

$$t_{i_r}(\boldsymbol{\sigma}^* | h_r) = \min\{t_{i_r}^d(\Pi_{1_r}^{c_r}, \dots, \Pi_{(i-1)_r}^{c_r}, P_{i_r}) \mid P_{i_r} \in \mathcal{P}[o_{i_r}(c_r), d]\} \text{ for each } i \geq 1.$$

Moreover, for any $j \geq 1$ and any strategy profile $\boldsymbol{\sigma}'$ of Γ with $\sigma'_{i_r} = \sigma_{i_r}^*$ for all $i \in [j]$, the path profile $(P'_i)_{i \in \Delta(c_r)}$ induced by h_r and $\boldsymbol{\sigma}'$ satisfies $P'_{i_r} = \Pi_{i_r}^{c_r}$ for all $i \in [j]$.

Now given any $k \geq 1$ and any $\sigma'_{k_r} \in \Sigma_{k_r}$, we consider strategy profile $\sigma' = (\sigma'_{k_r}, \sigma^*_{-k_r})$ and the path profile $\mathbf{p}' = (P'_i)_{i \in \Delta(c_r)}$ induced by h_r and σ' . We have $P'_{i_r} = \Pi_{i_r}^{c_r}$ for all $i \in [k-1]$, and

$$t_{k_r}(\sigma'_{k_r}, \sigma^*_{-k_r} | h_r) = t_{k_r}^d(\mathbf{p}') = t_{k_r}^d(\Pi_{1_r}^{c_r}, \dots, \Pi_{(k-1)_r}^{c_r}, \mathbf{p}'_{\Delta(c_r) \setminus \{1_r, \dots, (k-1)_r\}}).$$

It follows from Lemma 1 that

$$t_{k_r}(\sigma'_{k_r}, \sigma^*_{-k_r} | h_r) \geq \min\{t_{k_r}^d(\Pi_{1_r}, \dots, \Pi_{(k-1)_r}, P_{k_r}) \mid P_{k_r} \in \mathcal{P}[o_{k_r}(c_r), d]\} = t_{k_r}(\sigma^* | h_r).$$

The arbitrary choices of k and σ'_{k_r} imply that σ^* is an SPE of Γ . \square

We would like to remark that placing priorities on edges is crucial to the existence of an SPE. As the following example shows, the guarantee of SPE existence would be impossible if priorities were placed on players.

Example 2. Consider a subgame starting with a configuration illustrated in Figure 2. There are 9 players in total represented by small rectangles on the edges, with players i, j, k being our focus (the remaining 6 players do not need to make substantial decisions). The global priorities placed on the players are such that i ranks higher than j and j higher than k . However, k reaches v_1 (or v_2) one time unit earlier than i (if they choose to pass the same vertex) and adds 1 to the waiting time of i . Therefore, i, j, k have a Rock-Paper-Scissors like relationship. It can be checked that this subgame does not have any NE. Since each of the three players i, j, k only needs to make one decision, it follows that an SPE does not exist in this subgame.

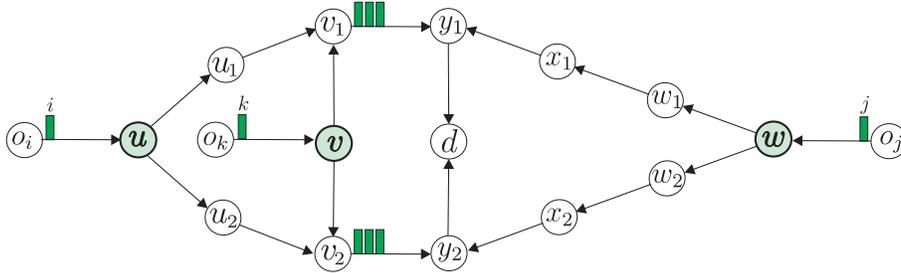


Figure 2: SPE might not exist if priorities were placed on players. In the subgame, player i ranks higher than player j , and j ranks higher than player k .

In contrast, the subgame in Figure 2 will have an NE if priorities are placed on edges as in our model. Suppose, e.g., $v_1 y_1 \prec_{y_1} x_1 y_1$ and $x_2 y_2 \prec_{y_2} v_2 y_2$, then it can be seen that the following implies an NE: k chooses path $o_k v v_1 y_1 d$, j chooses path $o_j w w_2 x_2 y_2 d$, and i makes an arbitrary choice. If in addition $y_2 d \prec_d y_1 d$, then

$$(v_2 y_2 d, v_1 y_1 d, v_2 y_2 d, v_1 y_1 d, v_2 y_2 d, v_1 y_1 d, o_j w w_2 x_2 y_2 d, o_k v v_1 y_1 d, o_i u u_2 v_2 y_2 d)$$

is the iteratively dominating path profile for the configuration.

4 Equilibria of interim games

Being natural and reasonable, models that are quite similar to Γ^n have been the focus in the studies of dynamic traffic games [14, 27, 32]. While only special NEs were studied in the literature, we establish some general properties of all NEs of Γ^n . Using these properties, we prove that the NE outcomes of game Γ^n form a proper subset of the SPE outcomes of game Γ . Therefore, Γ^n can serve not only as an independent model but also as a technical approach that facilitates our analysis of the main game model Γ .

Given any configuration c_r at time r , when studying some special SPE of game Γ in Section 3, we obtain an iterative dominating path profile, which is a *special* NE of the interim game $\Gamma^n(c_r)$, where players in $\Delta(c_r)$ are *completely ordered* such that the ones with smaller indices have advantages over those with larger indices. For studying *every* NE of Γ^n , it turns out that batching players according to their arrival times at the destination d is useful. For any NE π of $\Gamma^n(c_r)$, let $\tau(\pi, 1) < \tau(\pi, 2) < \tau(\pi, 3) < \dots$ be the arrival times of all players in $\Delta(c_r)$ at d under π . For each integer $k \geq 1$, let

$$\Delta(\pi, k) := \{i \in \Delta(c_r) \mid t_i^d(\pi) = \tau(\pi, k)\}$$

be the set of players in $\Gamma^n(c_r)$ who reach d at the k th earliest time, i.e., $\tau(\pi, k)$ under π ; we often refer to $\Delta(\pi, k)$ as the *kth batch*. We use

$$\Delta(\pi, [k]) := \cup_{j \in [k]} \Delta(\pi, j)$$

to denote the set of players reaching d no later than $\tau(\pi, k)$, i.e., those in the first k batches. For notational convenience, we set $\Delta(\pi, [0]) := \emptyset$ to be the 0th batch, and let $\Delta(\pi, [\infty]) := \Delta(c_r)$ denote the disjoint union of all batches.

4.1 A motivating example

In studying NE of Γ^n , we need to frequently analyze what happens if one player unilaterally deviates by choosing a different path. This is a quite tricky issue in general. Players may affect one another in quite unexpected ways. As the following extreme example shows, adding a player to the system may weakly improve the performances of other players (nobody is worse off and at least one is better off), or equivalently, removing a player may weakly harm the performances of other players.

Example 3. The configuration c_r at time r is as depicted in Figure 3. There are 10 players (shown as small rectangles on edges) in total, with i, j, k being our focus (note that Figure 3 shows only a part of the whole network). Edge e_1 has a higher priority than e_2 . Let player k choose the top path $u_1u_2u_3u_4u_5d$, player i choose the middle path $u_1u_2v_3v_4d$, and player j follow the bottom path $v_1v_2v_3v_4v_5d$. It can be checked that all the three players reach d at time $r + 5$ (and the three chosen paths along with the trivial paths of the other 7 players constitute an NE for $\Gamma^n(c_r)$). Suppose now player k is removed from the system, and i and j keep their chosen paths. This removing makes i arrive at vertices u_2, v_3 and v_4 one unit of time earlier. Since i has to pay one extra unit of waiting cost on edge v_4d , he reaches d still at time $r + 5$. However, the earlier arrival of i at v_3 delays j , making him reach d at time $r + 6$. Note that both path profiles are NEs.

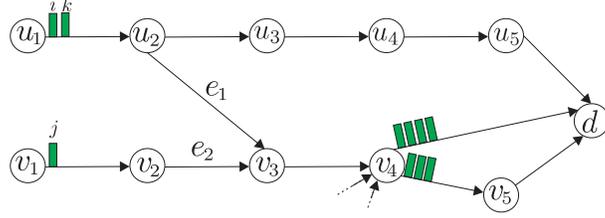


Figure 3: Removing a player (weakly) harms the performances of other players.

4.2 A critical lemma

Throughout this subsection, we fix a strategy profile $\alpha = (A_i)_{i \in \Delta(c_r)}$ of $\Gamma^n(c_r)$ and consider a fixed player $\zeta \in \Delta(c_r)$ and the scenario where only player ζ changes his strategy. Let $\alpha_{-\zeta} = (A_i)_{i \in \Delta(c_r) \setminus \{\zeta\}}$ denote the partial strategy profile of players in $\Delta(c_r)$ other than ζ . For each vertex v , define

$$\tau^v := \min_{A'_\zeta \in \mathcal{P}[o_\zeta(c_r), d]} \{t_\zeta^v(A'_\zeta, \alpha_{-\zeta})\}$$

as the earliest time at which player ζ can reach vertex v by unilaterally changing his path (if there is no path from $o_\zeta(c_r)$ to v , then set $\tau^v := +\infty$). Analogously, for each player $j \in \Delta(c_r) \setminus \{\zeta\}$ and vertex $v \in A_j$, define

$$\tau_j^v := \min_{A'_\zeta \in \mathcal{P}[o_\zeta(c_r), d]} \{t_j^v(A'_\zeta, \alpha_{-\zeta})\}$$

as the earliest time at which player j can reach vertex v when player ζ unilaterally changes his path. We emphasize that j keeps his path at A_j in the definition of τ_j^v .

In the following, for any non-degenerated path P in \bar{G} and vertex $v \in P$, we use $e_v(P)$ to denote the edge with head v in P .

Definition 3. For every player $j \in \Delta(c_r) \setminus \{\zeta\}$ and vertex $v \in A_j$, we say player ζ *dominates* player j at vertex v under α if

- (i) $\tau^v < \tau_j^v$; or
- (ii) $\tau^v = \tau_j^v$, and there exists an $o_\zeta(c_r)$ - v - d path A_ζ^* such that $t_\zeta^v(A_\zeta^*, \alpha_{-\zeta}) = \tau^v$ and $e_v(A_\zeta^*) \preceq_v e_v(A_j)$.

In words, player ζ dominates j at v if and only if either the earliest possible time that ζ reaches v is earlier than the earliest possible time that j reaches v , or the two earliest arrival times are equal and ζ has a corresponding path that allows him to enter v from an edge with a higher or equal priority over the edge taken by j , assuming that only ζ is allowed to deviate his path from the given path profile α .

Lemma 2. For any player $j \in \Delta(c_r) \setminus \{\zeta\}$ and vertex $v \in A_j$, if there exists a path $A'_\zeta \in \mathcal{P}[o_\zeta(c_r), d]$ such that $t_j^v(A'_\zeta, \alpha_{-\zeta}) \neq t_j^v(\alpha)$, i.e., player ζ can influence player j 's arrival time at v by unilaterally changing his path, then ζ dominates j at v .

The above lemma implies in particular that if player ζ can influence the arrival time of player j at vertex v , then ζ is able to reach v earlier than j does. Note that this result would be impossible if priorities were placed on players, which can be seen from Figure 2. We focus on player k and suppose the path profile α is such that i and j both choose their upper paths and k chooses his lower path. Then k is able to influence player j 's arrival time at node y_1 by deviating to the upper path. However, k is unable to reach y_1 earlier than j does.

From the proof of Lemma 2, we also have the following stronger result.

Corollary 1. *If player ζ dominates some player $j \in \Delta(c_r) \setminus \{\zeta\}$ at vertex $v \in A_j$, then ζ dominates j at all the other vertices on the subpath $A_j[v, d]$.*

As an immediate result of Corollary 1, we have the following lemma.

Lemma 3. *Let $\pi = (\Pi_i)_{i \in \Delta(c_r)}$ be an NE of $\Gamma^n(c_r)$. For every $k \geq 1$ and every player $j \in \Delta(c_r) \setminus \Delta(\pi, [k])$, j cannot dominate any player $i \in \Delta(\pi, [k])$ at any vertex of path Π_i .*

Proof. Suppose on the contrary that there does exist a player $j \in \Delta(c_r) \setminus \Delta(\pi, [k])$ and player $i \in \Delta(\pi, [k])$ such that j dominates i at some vertex $v \in \Pi_i$. Then by Corollary 1, j also dominates i at vertex d . This means that there exists a path Π_j^* such that

$$t_j^d(\Pi_j^*, \pi_{-j}) \leq \min_{\Pi'_j \in \mathcal{P}[o_j(c_r), d]} \{t_i^d(\Pi'_j, \pi_{-j})\} \leq t_i^d(\pi) = \tau(\pi, k).$$

However, $t_j^d(\pi) > \tau(\pi, k)$ due to $j \in \Delta(c_r) \setminus \Delta(\pi, [k])$, indicating that j has an incentive to deviate to Π_j^* and hence violating the fact that π is an NE. \square

4.3 Computing the earliest-arrival best response

For any configuration c_r and the corresponding normal-form game $\Gamma^n(c_r)$, suppose now we are given a path profile α . We show in this subsection how to compute a special best-response defined below.

Definition 4. Given a partial path profile $\alpha_{-\zeta}$, an $o_\zeta(c_r)$ - d path P^* of player $\zeta \in \Delta(c_r)$ is called the *earliest-arrival best response* if (i) ζ arrives at each node of P^* the earliest among all $o_\zeta(c_r)$ - d paths, and (ii) when there are more than one earliest ways to arrive at a node of P^* , P^* always selects an entering edge with the highest priority.

We shall see that the earliest-arrival best response exists and is unique. Observe that each dominating path computed in Algorithm 1 is an earliest-arrival best response. It can also be seen that, when all players take the earliest-arrival best response to each other in α , then α is exactly the NE computed in Algorithm 1.

We find that the earliest-arrival best responses can be computed via a dynamic programming that resembles the classical Dijkstra's algorithm for the shortest-path problem. This algorithm is polynomial when $\Delta(c_r)$ is finite. Note that there may well be best-responses that are not earliest-arrival. However, the earliest-arrival best response possesses several nice properties that are not owned by general best responses.

Under the path profile α , for any edge $e = uv$, recall that Q_e^t is the set of players queuing on edge e at time $t \geq r$. For an edge $e' = wu$ whose ending vertex u is the starting one of e , we use

$$Q_e^t(t, e') \subseteq Q_e^t$$

to denote the subset of players in Q_e^t who entered $e = uv$ at time t and from edges with priorities no higher than $e' = wu$. For any vertex $v \in \bar{V}$ with $\tau^v \neq 0$, denote $e^*(v)$ as the edge with the highest priority w.r.t vertex v that ζ can use to reach v at time τ^v . Initially, we set $Q_e^r(r, e^*(o_\zeta(c_r)))$ as the set of players in Q_e^r that queue after ζ when $e = e_\zeta(c_r)$, where $e_\zeta(c_r)$ is the initial edge that ζ rests on in c_r .

Lemma 4. *For any vertex $v \in \bar{V}$, if $\tau^v \neq \infty$, then*

$$\tau^v = \min_{uv \in \bar{E}^-(v)} \left\{ \tau^u + 1 + |Q_{uv}^{\tau^u}| - |Q_{uv}^{\tau^u}(\tau^u, e^*(u))| \right\}$$

and

$$e^*(v) = \arg \min_{uv \in \bar{E}^-(v)}^* \left\{ \tau^u + 1 + |Q_{uv}^{\tau^u}| - |Q_{uv}^{\tau^u}(\tau^u, e^*(u))| \right\},$$

where $\arg \min^*$ selects the edge with the highest priority among all those achieving the minimum.

Suppose now $\Delta(c_r)$ is finite. Then, the values of Q_e^t and $Q_e^t(t, e')$ are polynomially computable through simulating the traffic process. Therefore, the earliest-arrival best response of each player is polynomially computable in this case.

4.4 NE properties

The following lemma is very powerful. It helps us to show that the interactions between players of different batches at an NE are relatively clear and hierarchal, bearing many similarities to those in the special SPE studied in Section 3 and Appendix A.3.

Lemma 5. *Let π be an NE of $\Gamma^n(c_r)$. For any batch index $k \geq 1$, player $i \in B := \Delta(\pi, [k])$, vertex $v \in \Pi_i$, player $j \in \Delta(c_r) \setminus B$, and partial path profile \mathbf{r} for players in $\Delta(c_r) \setminus B$, it holds that*

$$t_i^v(\pi) = t_i^v(\pi_B, \mathbf{r}) \leq t_j^v(\pi_B, \mathbf{r}) \text{ and } t_j^d(\pi_B, \mathbf{r}) \geq \tau(\pi, k+1) > t_i^d(\pi_B, \mathbf{r}). \quad (4.1)$$

One of our main results in this section is the following theorem. Recall that a strong NE is an NE such that no subset of players are able to strictly better off via group deviation.

Theorem 2. *Let π be an NE of $\Gamma^n(c_r)$. The following properties are satisfied.*

- (i) *Hierarchal Independence. If players in a batch and those in earlier batches all follow their equilibrium strategies as in π , then their arrival times at any vertex are independent of other players' strategies.*
- (ii) *Hierarchal Optimality. The latency of each player in the first batch $\Delta(\pi, 1)$ is the smallest among the latencies of all players under any routing of $\Gamma^n(c_r)$. In general, for all $k \geq 2$, the latency of each player in the k th batch $\Delta(\pi, k)$ is the smallest among the latencies of all players in $\Delta(c_r) \setminus \Delta(\pi, [k-1])$ under any routing of $\Gamma^n(c_r)$ in which players in the first $k-1$ batches $\Delta(\pi, [k-1])$ follow their routes specified by π .*

- (iii) Global FIFO. Under π , if there exists a vertex $v \in \bar{V}$ (particularly the origin o) such that player i reaches v earlier than j does or they reach v at the same time but i comes from an edge with a higher priority than j does, then i reaches the destination d no later than j does.
- (iv) Strong NE. In $\Gamma^n(c_r)$, every NE is also a strong NE. That is, in any NE of $\Gamma^n(c_r)$, no set of players can be strictly better off by deviating together to other paths. In particular, this means that each NE of $\Gamma^n(c_r)$ is weakly Pareto Optimal.

4.5 Equilibrium relationship between Γ and Γ^n

In this subsection, we establish in a constructive way that each NE outcome of Γ^n is an SPE outcome of Γ . Given any NE π of Γ^n , we construct for every history $h_r = (c_0, \dots, c_r)$, $r \geq 0$, an NE $\pi(h_r)$ of interim game $\Gamma^n(c_r)$ with the particular setting $\pi(h_0) := \pi$. Then, we construct an SPE of Γ by assembling these NEs such that starting from any history h_r ($r \geq 0$) the outcome of the SPE is exactly the NE $\pi(h_r)$. (Note that the reference of each NE constructed is a history rather than a configuration. Since different histories may have the same ending configuration, we may construct multiple NEs for the same interim game.)

Such an NE-based assembling is more difficult than the one in Section 3, which constructs a special SPE restricted to nothing. What is more complicated here is that we are unable to design a Markovian SPE (cf. Appendix A.3). In particular, the natural idea of constructing the NEs $\pi(h_r)$, $r \geq 1$ rather than $\pi(h_0) = \pi$ directly using Algorithm 1 does not work, because, e.g., players not in the first batch under π may have incentives to deviate at the game tree root.

Our construction of the NEs is done iteratively on the game tree of Γ from the root $h_0 = (c_0)$. Initially, the constructed NE $\pi(h_0)$ for h_0 is simply π . For each $r \geq 1$, suppose inductively that for a history $h_{r-1} = (c_0, \dots, c_{r-1}) \in \mathcal{H}_{r-1}$, the NE $\pi(h_{r-1})$ of game $\Gamma^n(c_{r-1})$, written for convenience as $\rho = (P_i)_{i \in \Delta(c_{r-1})}$, has been constructed. We construct below the NE $\pi(h_r)$, denoted $\rho' = (P'_i)_{i \in \Delta(c_r)}$, for each child history $h_r = (c_0, \dots, c_{r-1}, c_r)$ of h_{r-1} . Suppose that

- $(e_i)_{i \in \Delta(c_{r-1})}$ is the action profile at game tree node h_{r-1} determined by ρ , i.e., no action in the profile deviates from ρ , and
- $(e'_i)_{i \in \Delta(c_{r-1})}$ is the action profile that leads history h_{r-1} to its child history h_r (or equivalently leads c_{r-1} to c_r).

We construct ρ' in two steps as follows.

CONSTRUCTION I: Let $\mathbb{k} \geq 0$ be the maximum nonnegative integer such that the action of each player of $\Delta(\rho, [\mathbb{k}])$ under $(e'_i)_{i \in \Delta(c_{r-1})}$ is the same as under $(e_i)_{i \in \Delta(c_{r-1})}$, i.e.,

$$\mathbb{k} := \sup\{k \mid e_i = e'_i \text{ for all } i \in \Delta(\rho, [k])\}. \quad (4.2)$$

(Note that it is possible $\mathbb{k} = 0$ with $\Delta(\rho, [0]) = \emptyset$ or $\mathbb{k} = \infty$ with $\Delta(\rho, [\infty]) = \Delta(c_{r-1})$.) We let players in the first \mathbb{k} batches who are still in the system at time r under ρ keep their paths under ρ' . To be more specific, we set

$$P'_i := P_i[o_i(c_r), d] \text{ for all } i \in \Delta(\rho, [\mathbb{k}]) \cap \Delta(c_r). \quad (4.3)$$

CONSTRUCTION II: Based on the equilibrium strategies kept for players in $\Delta(\boldsymbol{\rho}, [\mathbb{k}]) \cap \Delta(c_r)$ as specified in (4.3), which particularly guarantees invariant arrival times at any vertex for these players regardless of other players' choices (see the *hierarchical independence* in Theorem 2(i)), we find an *iterative dominating path profile* $(P'_i)_{i \in \Delta(c_r) \setminus \Delta(\boldsymbol{\rho}, [\mathbb{k}])}$ for the remaining players, using a process that is more general than Algorithm 1 (see Appendices A.2 and B.4 for more details).

Intuitively, $\boldsymbol{\rho}'$ is a combination of a part of $\boldsymbol{\rho}$ and an iterative dominating partial path profile. It can be shown that the constructed $\boldsymbol{\rho}'$ is indeed an NE of $\Gamma^n(c_r)$ (see Lemma 13 in the Appendix), which completes our *inductive* construction. Furthermore, the partial hierarchical independence and iterative domination guaranteed by Constructions I and II enable us to accomplish our task of assembling all the NEs constructed into an SPE of Γ .

Theorem 3. *If $\boldsymbol{\pi}$ is an NE of game Γ^n , then there exists an SPE σ of game Γ , such that the path profile induced by the initial history h_0 and σ is exactly $\boldsymbol{\pi}$.*

Combining with Example 1, the above theorem shows that the NE outcome set of Γ^n is typically a proper subset of the SPE outcome set of Γ , reaffirming an intuition that model Γ is more flexible than Γ^n . Since model Γ^n is relatively easier to study, also natural and more frequently analyzed in the literature, Theorem 3 can serve as a useful bridge between Γ and Γ^n .

5 Bounding equilibrium queue lengths

In this section, we consider two classes of networks, and prove that the queue lengths at any NE of the game $\Gamma^n(G)$ are bounded above by a finite number that *only* depends on the number $m = |E|$ of edges in G , provided that the inflow size never exceeds the network capacity.

Definition 5. A network G with origin o and destination d is *o - d series-parallel* or simply *series-parallel* if

- (i) G consists of a single edge od ; or
- (ii) G is obtained by connecting two smaller o_i - d_i series-parallel networks G_i , $i = 1, 2$, *in series* — identifying d_1 and o_2 , and naming o_1 as o , and d_2 as d ; or
- (iii) G is obtained by connecting two smaller o_i - d_i series-parallel networks G_i , $i = 1, 2$, *in parallel* — identifying o_1 and o_2 to form o and identifying d_1 and d_2 to form d .

We reserve L for the length of a longest o - d path in G , and Λ for the maximum in-degree of vertices in G . Then clearly, $\max\{L, \Lambda\} \leq m$. The following observation is trivial but important.

Observation 1. Under any routing, at most Λ players can reach the same vertex (in particular, the destination d) at the same time.

In view of Definition 5, the series-parallel network G is obtained from m edges by performing a sequence of $m - 1$ series or parallel connection operations. Each of these operations connects two series-parallel subnetworks G_1 and G_2 of G into a bigger series-parallel subnetwork G_3 of G . Fix any such sequence of $m - 1$ connection operations that leads to G and let \mathfrak{S} be the set of all the subnetworks G_1 , G_2 and G_3 that appear during the whole process of the sequence of $m - 1$ connection operations. Then clearly,

$$G \in \mathfrak{S} \text{ and } |\mathfrak{S}| = 2m - 1.$$

Henceforth, we study an arbitrary NE, denoted as $\pi = (\Pi_i)_{i \in \Delta}$, of game Γ^n . For each player $i \in \Delta$ and each vertex $v \in \Pi_i$, we use $\Delta_i^v := \{j \in \Delta \mid t_j^v(\pi) = t_i^v(\pi)\}$ to denote the set of players in Δ who reach v under π at the same time as i does. Note from Observation 1 that

$$|\Delta_i^v| \leq \Lambda \text{ for every } i \in \Delta \text{ and each vertex } v \in \Pi_i.$$

Using the observation and global FIFO property in Theorem 2(iii), we can upper bound the ratio between the population sizes in any pair of subnetworks connecting in parallel, as the following lemma states.

Lemma 6. *Suppose that $G_1, G_2 \in \mathfrak{S}$ are connected in parallel to form a series-parallel network in \mathfrak{S} . Under π and at any time point t , if G_1 and G_2 accommodate n_1 and n_2 players respectively, then $n_i \leq 2\Lambda L(2\Lambda + n_j) \leq 2m^2(2m + n_j)$ for $\{i, j\} = \{1, 2\}$.*

By an o - d cut or simply a cut of G we mean a set of edges whose removal from G leaves the graph unconnected from o to d . We say that a minimal cut of G is *full* at time t if each edge of the cut has some player on it at time t (i.e., there is a nonempty queue at each edge). Let $C = \{e_1, \dots, e_k\}$ and $C' = \{e'_1, \dots, e'_k\}$ be two minimum cuts of G . We say that C is *on the left* of C' if for each $i \in [k]$ there exists an o - d path in G which contains e_i, e'_i and visits e_i before e'_i .

For any o' - d' series-parallel network G' , let $\Xi(G')$ denote the leftmost minimum o' - d' cut of G' . In particular, $|\Xi(G')|$ is the minimum cut size of G' . The cut $\Xi(G')$ partitions G' into two parts: the left part G'^l that contains o' , and the right G'^r that contains d' .

Definition 6. Let \mathfrak{F} denote the set of series-parallel subnetworks $G_i \in \mathfrak{S}$ such that there exist finite integers F_i^l and F_i^r satisfying the following two conditions (for any inflow):

- (i) $\Xi(G_i)$ is full as long as G_i^l accommodates more than F_i^l players,³ and
- (ii) G_i^r can accommodate at most F_i^r players at any time.

First, we can see the set \mathfrak{F} is nonempty because every single edge apparently belongs to \mathfrak{F} . In the two lemmas below, we are given $G_1, G_2 \in \mathfrak{F}$, such that G_i is an o_i - d_i series-parallel network for $i = 1, 2$.

Lemma 7. *If $G_s \in \mathfrak{S}$ is the combination of G_1 and G_2 in series, then $G_s \in \mathfrak{F}$.*

Lemma 8. *If $G_p \in \mathfrak{S}$ is the combination of G_1 and G_2 in parallel, then $G_p \in \mathfrak{F}$.*

Proof. It is clear that $\Xi(G_p) = \Xi(G_1) \cup \Xi(G_2)$, $G_p^l = G_1^l \cup G_2^l$, and $G_p^r = G_1^r \cup G_2^r$. Therefore, G_p^r accommodates at most $F_p^r \equiv F_1^r + F_2^r$ players at any time. To prove $G_p \in \mathfrak{F}$, we only need to show that $\Xi(G_1) \cup \Xi(G_2)$ is full as long as $G_1^l \cup G_2^l$ accommodates more than a certain finite number of players. As $G_1, G_2 \in \mathfrak{F}$, it suffices to consider the time t when one of G_1^l and G_2^l , say G_1^l , accommodates at most F_1^l players. Suppose that G_2^l accommodates F players at time t . It follows from Lemma 6 that G_1 accommodates at least $F/(2m^2) - 2m$ players at time t . As $G_1 \in \mathfrak{F}$, there are at least $F/(2m^2) - 2m - F_1^r$ players inside G_1^l at time t . It follows from $F/(2m^2) - 2m - F_1^r \leq F_1^l$ that $F \leq 2m^2(2m + F_1^l + F_1^r)$. \square

³It means that the full cut and the accommodation are observed at the same time.

Theorem 4. *Let G be a series-parallel network. If $|\Delta_t| \leq |\Xi(G)|$ for all $t \geq 1$, then there exists a finite number such that, for any NE of $\Gamma^n(G)$, the number of players in G is upper bounded by this number, implying that the latency of any player at any NE is upper bounded.*

Proof. Since $|\mathfrak{S}| = 2m - 1$, combining Lemmas 7 and 8, an inductive argument shows that $\mathfrak{S} = \mathfrak{F}$. In particular, $G \in \mathfrak{F}$ says that $\Xi(G)$ is full as long as G^l accommodates more than F^l players, and G^r can accommodate at most F^r players at any time, where F^l and F^r are finite numbers. Similar to the argument used in the proof of Lemma 7, we consider any time point t such that G^l accommodates at most F^l players at time t , and more than F^l players at time $t + 1$. Then G^l accommodates at most $F^l + |\Delta_{t+1}| \leq F^l + |\Xi(G)|$ players and $\Xi(G)$ is full at time $t + 1$. Thus $|\Xi(G)|$ players leave G^l at time $t + 2$ while $|\Delta_{t+2}|$ ($\leq |\Xi(G)|$) players enter G^l . It follows that the number of players inside G^l is nonincreasing unless the number decreases below $F^l + 1$. Therefore, at any time G^l can accommodate at most $F^l + |\Xi(G)|$ players, and G can accommodate at most $F^l + |\Xi(G)| + F^r$ players. This also means the total latency for any player travelling from o to d under any NE is bounded by a finite number since any queue length is upper bounded by the finite number $F^l + m + F^r$. \square

If we take the average traveling time from o to d for all players as a measure of the social welfare as in [27], then the boundedness of the queues implies the PoA of game $\Gamma^n(G)$ is also bounded. By Theorem 3, this means that the PoS of game $\Gamma(G)$ is also bounded.

In closing this section, we establish boundedness of any *SPE queue lengths* for another type of networks.

Theorem 5. *Given a dynamic routing game Γ on an acyclic network $G = (V, E)$ with origin o and destination d , suppose $|E^+(v)| \geq |E^-(v)|$ for any vertex $v \in V \setminus \{o, d\}$. Let σ be any SPE of game $\Gamma(G)$. If the inflow size $|\Delta_r|$ never exceeds the size of a minimum o - d cut of G , then there exists a finite number U depending only on $|E|$ such that $t_i(\sigma) \leq U$ all $i \in \Delta$.*

6 Concluding remarks

In this paper, we have studied an atomic network congestion game of discrete-time dynamic traffic. This is a relatively unexplored area in the study of congestion games. The most prominent feature of our model is the great flexibility agents enjoy so as to make online decisions at all intermediary vertices and, accordingly, SPE serves as the default solution concept. We have shown that this more flexible model has close connections with the corresponding game of normal form, which is more often studied in the literature. We have identified many surprisingly nice properties of the NE flows of the latter model, the equivalence between NEs and strong NEs and a global FIFO, to name a few.

This paper is our first attempt in understanding the consequences of the introduction of agents' flexibility of online decision making in dynamic traffic games. Many interesting problems are widely open. For example, is there an upper bound on the SPE (or NE) queue lengths for general networks? What are more accurate bounds on PoA and PoS w.r.t. either NEs or SPEs? Does a long-run steady state exist when the inflow is constant or seasonal? How efficient is this steady state if it does exist? What if agents have multi-origins and multi-destinations? What if

agents are allowed to choose their departure times? Exploring these problems will undoubtedly help us better understand atomic games of dynamic traffic.

References

- [1] Elliot Anshelevich and Satish Ukkusuri. Equilibria in dynamic selfish routing. In *International Symposium on Algorithmic Game Theory*, pages 171–182. Springer, 2009.
- [2] Robert J Aumann. Acceptable points in general cooperative n-person games. *Contributions to the Theory of Games (AM-40)*, 4:287, 1959.
- [3] Umang Bhaskar, Lisa Fleischer, and Elliot Anshelevich. A stackelberg strategy for routing flow over time. *Games and Economic Behavior*, 92:232–247, 2015.
- [4] Roberto Cominetti, José Correa, and Omar Larré. Dynamic equilibria in fluid queueing networks. *Operations Research*, 63(1):21–34, 2015.
- [5] Carlos F Daganzo. Queue spillovers in transportation networks with a route choice. *Transportation Science*, 32(1):3–11, 1998.
- [6] Babak Farzad, Neil Olver, and Adrian Vetta. A priority-based model of routing. *Chicago Journal of Theoretical Computer Science*, 1, 2008.
- [7] Lisa Fleischer and Martin Skutella. Quickest flows over time. *SIAM Journal on Computing*, 36(6):1600–1630, 2007.
- [8] Lisa Fleischer and Éva Tardos. Efficient continuous-time dynamic network flow algorithms. *Operations Research Letters*, 23(3):71–80, 1998.
- [9] D. R. Ford and D. R. Fulkerson. *Flows in Networks*. Princeton University Press, Princeton, NJ, USA, 1962.
- [10] Lester R Ford and Delbert Ray Fulkerson. Constructing maximal dynamic flows from static flows. *Operations Research*, 6(3):419–433, 1958.
- [11] Terry L. Friesz, David Bernstein, Tony E. Smith, Roger L. Tobin, and B. W. Wie. A variational inequality formulation of the dynamic network user equilibrium problem. *Operations Research*, 41(1):179–191, 1993.
- [12] Tobias Harks, Britta Peis, Daniel Schmand, and Laura Vargas Koch. Competitive Packet Routing with Priority Lists. In Piotr Faliszewski, Anca Muscholl, and Rolf Niedermeier, editors, *41st International Symposium on Mathematical Foundations of Computer Science (MFCS 2016)*, volume 58 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 49:1–49:14, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [13] Chris Hendrickson and George Kocur. Schedule delay and departure time decisions in a deterministic model. *Transportation Science*, 15(1):62–77, 1981.

- [14] 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.
- [15] Martin Hoefer, Vahab S Mirrokni, Heiko Röglin, and Shang-Hua Teng. Competitive routing over time. *Theoretical Computer Science*, 412(39):5420–5432, 2011.
- [16] Ronald Koch. *Routing games over time*. Ph.d. thesis, Technische Universität Berline, 2012.
- [17] Ronald Koch and Martin Skutella. Nash equilibria and the price of anarchy for flows over time. In *International Symposium on Algorithmic Game Theory*, pages 323–334. Springer, 2009.
- [18] 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.
- [19] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Annual Symposium on Theoretical Aspects of Computer Science*, pages 404–413. Springer, 1999.
- [20] Janardhan Kulkarni and Vahab Mirrokni. Robust price of anarchy bounds via lp and fenchel duality. In *Proceedings of the Twenty-sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '15, pages 1030–1049, Philadelphia, PA, USA, 2015. Society for Industrial and Applied Mathematics.
- [21] Martin Macko, Kate Larson, and Lubovs Steskal. *Braess's Paradox for Flows over Time*, pages 262–275. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.
- [22] Frédéric Meunier and Nicolas Wagner. Equilibrium results for dynamic congestion games. *Transportation Science*, 44(4):524–536, 2010.
- [23] Srinivas Peeta and Athanasios K Ziliaskopoulos. Foundations of dynamic traffic assignment: The past, the present and the future. *Networks and Spatial Economics*, 1(3):233–265, 2001.
- [24] A. B. Philpott. Continuous-time flows in networks. *Math. Oper. Res.*, 15(4):640–661, October 1990.
- [25] Tim Roughgarden. Routing games. *Algorithmic Game Theory*, 18:459–484, 2007.
- [26] Tim Roughgarden and Éva Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.
- [27] Marco Scarsini, Marc Schröder, and Tristan Tomala. Dynamic atomic congestion games with seasonal flows. *HEC Paris Research Paper No. ECO/SCD-2013-1016*. Available at SSRN: <https://ssrn.com/abstract=2278203> or <http://dx.doi.org/10.2139/ssrn.2278203>, 2016.
- [28] Andreas S Schulz and NS Moses. On the performance of user equilibria in traffic networks. In *Proc. 14th ACM-SIAM Symposium on Discrete Algorithms*, pages 12–14, 2003.

- [29] Martin Skutella. An introduction to network flows over time. In *Research Trends in Combinatorial Optimization*, pages 451–482. Springer, 2009.
- [30] William S Vickrey. Congestion theory and transport investment. *The American Economic Review*, 59(2):251–260, 1969.
- [31] John Glen Wardrop. Road paper. some theoretical aspects of road traffic research. In *ICE Proceedings: engineering divisions*, volume 1, pages 325–362. Thomas Telford, 1952.
- [32] TL Werth, M Holzhauser, and SO Krumke. Atomic routing in a deterministic queuing model. *Operations Research Perspectives*, 1(1):18–41, 2014.
- [33] Samuel Yagar. Dynamic traffic assignment by individual path minimization and queuing. *Transportation Research*, 5(3):179–196, 1971.

Appendix

A Details in Section 3

A.1 Iterative dominations

To facilitate our discussions, the proof of Lemma 1 in particular, we introduce several notations. Given any (directed) path P in \bar{G} , and vertices a, b in P such that a is passed no later than b by path P , we use $P[a, b]$ to denote the sub-path of P from a to b . We write $P(a, b) = P[a, b] \setminus \{a\}$, $P[a, b) = P[a, b] \setminus \{b\}$ and $P(a, b) = P[a, b] \setminus \{a, b\}$.

Given any feasible configuration c_r at time r , let players $1, 2, \dots$ of $\Delta(c_r)$ be as indexed and path profile $\boldsymbol{\pi} = (\Pi_i)_{i \in \Delta(c_r)}$ be as computed in Algorithm 1. For any player indices i, j with $i < j$ and any vertex $v \in \bar{G}$, let

$$\tau_j^v \llbracket i \rrbracket := \min\{t_j^v(\boldsymbol{\pi}_{[i]}, R_j) \mid R_j \in \mathcal{P}[o_j(c_r), d]\} \quad (\text{A.1})$$

denote the value τ_j^v computed for player j in Step 3 at the $(i + 1)$ th iteration of Algorithm 1, i.e., the earliest time for player j to reach vertex v starting from edge $e_j(c_r)$, based only on the partial routing $\boldsymbol{\pi}_{[i]}$ of players in $[i]$.

Lemma 9 (Restatement of Lemma 1). *Let player indices i, j and player subset S satisfy $j \in S \subseteq \Delta(c_r) \setminus [i - 1]$. Then*

$$t_i^v(\boldsymbol{\pi}_{[i]}, \boldsymbol{p}_{S \setminus \{i\}}) = \tau_i^v \llbracket i - 1 \rrbracket \leq t_j^v(\boldsymbol{\pi}_{[i-1]}, \boldsymbol{p}_S) \quad (\text{A.2})$$

holds for every vertex $v \in \Pi_i$ and every path profile \boldsymbol{p} of $\Gamma^n(c_r)$.

Proof. We prove by induction on i . Consider first the base case $i = 1$, whose proof is quite similar to the general case. First, $t_1^v(\boldsymbol{\pi}_{[1]}, \boldsymbol{p}_{S \setminus \{1\}}) \geq \tau_1^v \llbracket 0 \rrbracket$ is apparent, because $\tau_1^v \llbracket 0 \rrbracket$ is the shortest path length without waiting cost from $o_1(c_r)$ to v . To see that $t_1^v(\boldsymbol{\pi}_{[1]}, \boldsymbol{p}_{S \setminus \{1\}}) = \tau_1^v \llbracket 0 \rrbracket$, it suffices to prove that player 1 never queues on any edge under routing $(\boldsymbol{\pi}_{[1]}, \boldsymbol{p}_{S \setminus \{1\}})$. Suppose

on the contrary that player 1 queues after some player $k \in S$ on edge xy , and let xy be the first of such edge on Π_1 . Then, under $(\boldsymbol{\pi}_{[1]}, \boldsymbol{p}_{S \setminus \{1\}})$, either player k enters xy earlier than player 1 or they enter xy at the same time but player k comes from an edge with a higher priority. Define a path of k as $X_k := P_k[o_k(c_r), y] \cup \Pi_1[y, d] \in \mathcal{P}[o_k(c_r), d]$. Then, $t_k^w(X_k) \leq t_1^w(\Pi_1)$ for all vertices $w \in \Pi_1[y, d]$. Note that $\tau_1^w[[0]] = t_1^w(\boldsymbol{\pi}_{[1]}, \boldsymbol{p}_{S \setminus \{1\}})$ and $\tau_k^w[[0]] \leq t_k^w(\boldsymbol{\pi}_{[1]}, \boldsymbol{p}_{S \setminus \{1\}})$ for all $w \in \Pi_1[y, d]$. Then, either (i) $\tau_1^x[[0]] > \tau_k^x[[0]]$, or (ii) $\tau_1^x[[0]] = \tau_k^x[[0]]$ and via X_k player k is able to enter xy from an edge with a higher priority and follow the remaining path of Π_1 , or (iii) $e_1(c_r) = e_k(c_r) = xy$ and player 1 queues after k on this edge, all contradicting the choice of player 1. This finishes the first part of (A.2) for $i = 1$. The second part of (A.2) for $i = 1$ is obvious, whether $j = i$ or not.

Suppose now $i \geq 2$ and (A.2) is valid for smaller values. This means that, as long as players in $[i - 1]$ follow $\boldsymbol{\pi}_{[i-1]}$, they will never be affected by other players; in fact, no player $j \geq i$ can reach any vertex earlier than any one in $[i - 1]$. This also means that, as long as players in $[i - 1]$ follow $\boldsymbol{\pi}_{[i-1]}$, they will exert *invariant influences* on the movements of other players: the set $Q_e^s \cap [i - 1]$ depends only on time $s \geq r$ and edge e but not on the choices of players in $S \setminus [i - 1]$; in addition, players in $Q_e^s \cap [i - 1]$ (if nonempty) always queue before other players. It is this invariant influence property that makes the proof of the general case almost the same as the base case, as demonstrated below.

First, due to the above invariant influences, it's not hard to see that $t_i^v(\boldsymbol{\pi}_{[i]}, \boldsymbol{p}_{S \setminus \{i\}}) \geq \tau_i^v[[i-1]]$: players in $S \setminus [i]$ cannot speed up i , because the latency caused by players in $[i - 1]$ alone is always $\tau_i^v[[i - 1]]$, regardless of the choices of $S \setminus [i]$. Combining $\tau_i^v[[i - 1]] \leq \tau_j^v[[i - 1]]$ (because otherwise, using $X_j := P_j[o_j(c_r), v] \cup \Pi_i[v, d] \in \mathcal{P}[o_j(c_r), d]$, player j would be able to reach v earlier than i does and take the remaining part of Π_i , contradicting the choice of i and Π_i) and $\tau_j^v[[i - 1]] \leq t_j^v(\boldsymbol{\pi}_{[i-1]}, \boldsymbol{p}_S)$ (due to the invariant influences from $[i - 1]$) also gives the inequality part of (A.2).

So it remains to show that $t_i^v(\boldsymbol{\pi}_{[i]}, \boldsymbol{p}_{S \setminus \{i\}}) \leq \tau_i^v[[i-1]]$, i.e., players in $S \setminus \{i\}$ will not slow down i . Suppose on the contrary that $t_i^v(\boldsymbol{\pi}_{[i]}, \boldsymbol{p}_{S \setminus \{i\}}) > \tau_i^v[[i - 1]] = t_i^v(\boldsymbol{\pi}_{[i]})$ for some vertex $v \in \Pi_i$. Let v be the first such vertex along Π_i , indicating that

$$(1) \quad t_i^w(\boldsymbol{\pi}_{[i]}, \boldsymbol{p}_{S \setminus \{i\}}) = \tau_i^w[[i - 1]] \text{ for every vertex } w \in \Pi_i[o_i(c_r), v).$$

In view of the invariant influences from players in $[i - 1]$, there must exist some player $k \in S \setminus [i]$ and an edge $xy \in \Pi_i[o_i(c_r), v]$ such that k enters xy earlier than i does under $(\boldsymbol{\pi}_{[i]}, \boldsymbol{p}_{S \setminus \{i\}})$. Let xy be the *first* such edge along $\Pi_i[o_i(c_r), v]$. By (1), we have

$$(2) \quad \text{under routing } (\boldsymbol{\pi}_{[i]}, \boldsymbol{p}_{S \setminus \{i\}}), \text{ player } k \text{ reaches vertex } x \text{ and enters edge } xy \text{ at time } t_k^x(\boldsymbol{\pi}_{[i]}, \boldsymbol{p}_{S \setminus \{i\}}) \leq t_i^x(\boldsymbol{\pi}_{[i]}, \boldsymbol{p}_{S \setminus \{i\}}) = \tau_i^x[[i - 1]].$$

Construct a path $X_k := P_k[o_k(c_r), x] \cup \Pi_i[x, d] \in \mathcal{P}[o_k(c_r), d]$ for k . Since $t_k^x(\boldsymbol{\pi}_{[i-1]}, X_k) = t_k^x(\boldsymbol{\pi}_{[i-1]}, P_k) \leq t_k^x(\boldsymbol{\pi}_{[i]}, \boldsymbol{p}_{S \setminus \{i\}})$ (where the equality is simply from definition of X_k and the inequality is due to the invariant influences of $[i - 1]$ that k cannot be sped up by players in $S \cup \{i\}$), which together with (2) implies $t_k^x(\boldsymbol{\pi}_{[i-1]}, X_k) \leq \tau_i^x[[i - 1]] = t_i^x(\boldsymbol{\pi}_{[i-1]}, \Pi_i)$. Consequently,

$$(3) \quad t_k^w(\boldsymbol{\pi}_{[i-1]}, X_k) \leq \tau_i^w[[i - 1]] \text{ for each vertex } w \in X_k[x, d] = \Pi_i[x, d].$$

By definition of player i from Algorithm 1 and (3), we have $t_k^w(\boldsymbol{\pi}_{[i-1]}, X_k) = \tau_i^w[[i - 1]]$ for each vertex $w \in X_k[x, d] = \Pi_i[x, d]$. Therefore, from (3), we know either (if X_k and Π_i have

different incoming edges into x) X_k has a higher priority incoming edge into x than Π_i does, or (by the choice of edge xy) $e_k(c_r) = xy = e_i(c_r)$, $X_k = \Pi_i$, and player k queues before player i at $e_k(c_r) = e_i(c_r)$. However, the choice made at the i th iteration of Algorithm 1 excludes the possibilities of both cases. This finishes the proof. \square

A.2 Generalized iterative dominations

As can be seen from the proof of Lemma 1, our induction hypothesis only involves the equation (the first part) of (A.2), which guarantees the critical invariant influences property. This leads us to the following generalization of Algorithm 1, which computes an iterative dominating *partial* path profile *based* on the fixed routing of some players.

ALGORITHM 2: (ITERATIVE DOMINATING PARTIAL PATH PROFILE WITH A BASE)

Input: a feasible configuration c_r at time r ; a partial path profile $\mathbf{b} = (B_\ell)_{\ell \in U}$ for players in a (possibly empty) finite subset $U \subseteq \Delta(c_r)$ such that, for all $\ell \in U$ and all $v \in B_\ell$, $t_\ell^v(\mathbf{b}, \mathbf{p}_S)$ is the same over all path profile \mathbf{p} of $\Gamma^n(c_r)$ and $S \subseteq \Delta(c_r) \setminus U$.

Output: the iterative dominating partial path profile (routing) $\boldsymbol{\pi} = (\Pi_i)_{i \in \Delta(c_r) \setminus U}$ for $\Delta(c_r) \setminus U$ along with the corresponding player indices $1, 2, \dots$.

1. Initiate $D \leftarrow U$, $\boldsymbol{\pi}_{[0]} \leftarrow \emptyset$, $i \leftarrow 0$.
 2. $i \leftarrow i + 1$ (NB: Start to search for the new dominator and his dominating path).
 3. For each player $j \in \Delta(c_r) \setminus D$ and vertex $v \in \bar{G}$,
 - let $\tau_j^v = \min\{t_j^v(\mathbf{b}, \boldsymbol{\pi}_{[i-1]}, R_j) \mid R_j \in \mathcal{P}[o_j(c_r), d]\}$ be the earliest time for j to reach vertex v from his *current location in c_r* , assuming that *all other* players in \bar{G} are those in D and they go along their paths specified in $(\mathbf{b}, \boldsymbol{\pi}_{[i-1]})$;
 - let \mathcal{P}_j^v denote the set of all the *corresponding* $o_j(c_r)$ - v paths for j to reach j at time τ_j^v .
 - if there is no path in \bar{G} to v from the current location of player j , then set $\tau_j^v \leftarrow \infty$ and $\mathcal{P}_j^v \leftarrow \emptyset$.
 4. Run Steps 4– 12 of Algorithm 1 to identify dominator i and his dominating path Π_i .
 5. $D \leftarrow D \cup \{i\}$, $\boldsymbol{\pi}_{[i]} \leftarrow (\boldsymbol{\pi}_{[i-1]}, \Pi_i)$. (NB: the algorithm outputs player i and his path Π_i .)
 6. Go to Step 2.
-

The verbatim adaption of the proof of Lemma 1 gives the following generalization for iterative domination, which will play a critical role in the discussion of Section 4.

Lemma 10. *Given the input and output of Algorithm 2, if $j \in S \subseteq \Delta(c_r) \setminus (U \cup [i - 1])$, then for every vertex $v \in \Pi_i$ and path profile \mathbf{p} of $\Gamma^n(c_r)$, it holds that*

$$t_i^v(\mathbf{b}, \boldsymbol{\pi}_{[i]}, \mathbf{p}_{S \setminus [i]}) = \min_{R_i \in \mathcal{P}[o_i(c_r), d]} t_i^v(\mathbf{b}, \boldsymbol{\pi}_{[i-1]}, R_i) \leq t_j^v(\mathbf{b}, \boldsymbol{\pi}_{[i-1]}, \mathbf{p}_S).$$

A.3 Nice properties of the special SPE

Definition 7. [Preemption] Given a strategy profile \mathbf{p} of interim game $\Gamma^n(c_r)$, we say that i *preempts* j at vertex v (under \mathbf{p}) if they both pass v and i reaches v earlier than j does (under

\mathbf{p}); we say i *weakly preempts* j at vertex v if either i preempts j at vertex v , or i and j reach v at the same time but i comes from an edge (with head v) with a higher priority than j does.

Definition 8. A player $i \in \Delta$ is said to have a *higher original priority* than player $j \in \Delta$ if i preempts or weakly preempts j at the origin o .

It can be summarized from the above discussions that the special SPE σ^* which we construct in the proof of Theorem 1 has the following related but distinct nice properties.

Sequential Independence. If a player i and all players with higher original priorities than i fix their strategies as in the SPE, then their realized paths as well as the arrival times at all vertices are independent of other players' strategies.

Sequential Optimality. The latency of any player realized under the SPE is minimum among all the feasible flows in which players with higher original priorities follow their strategies in the SPE.

Pareto Optimality. No group of players can be strictly better off by deviating together from σ^* . That is, the SPE outcome from any configuration is a strong NE ([2]) of the corresponding subgame. Thus the SPE outcome of σ^* is weakly Pareto optimal.

Global FIFO. In any routing realized under the SPE, if player i weakly preempts player j at some vertex of G (particularly its origin o), then i leaves the system no later than j .

No Overtaking. If player i has a higher original priority than player j and both pass through some vertex v in the realization of the SPE, then i weakly preempts j at v .

Earliest Arrival. Given the other players' strategies in the SPE, each player using his realized path under the SPE is guaranteed to reach any vertex on the path (not only the destination d) at an earliest time among all of his possible choices of paths.

Markov and Anonymity. According to the SPE, the action each player takes at each node of the game tree of Γ depends *only* on the immediately previous configuration but not on earlier configurations in the history, and the identities of other players do not matter.⁴

B Details in Section 4

B.1 Proof of Lemma 2

Proof. For each vertex $v \in \bar{V}$, we denote $I_v \subset \Delta(c_r)$ as the set of players apart from ζ whose arrival times at v can be influenced by player ζ 's unilateral path changing. Apparently, if $j \in I_v$, then it must be the case that $v \in A_j$. Thus, we need to prove, for every player $j \in I_v$, that player ζ dominates player j at vertex v . Suppose edge $(o_\zeta(c_r), v_0)$ is the starting edge of path A_ζ .

Claim 1. If $I_v \neq \emptyset$, then there exists a directed path from v_0 to v , i.e., $\tau^v \neq \infty$.

Proof. Suppose $j \in I_v$. Let $e_v(A_j) = (u, v)$. Since there exists $A'_\zeta \in \mathcal{P}[o_\zeta(c_r), d]$ such that $t_j^v(A'_\zeta, \alpha_{-\zeta}) \neq t_j^v(\alpha)$, one of the following cases must happen:

- (i) $(u, v) \in A'_\zeta$ or $(u, v) \in A_\zeta$;

⁴A strategy σ_i of player i is called *Markovian* if $\sigma_i(h_r) = \sigma_i(h_{r'})$ holds for all histories $h_r = (c_0, c_1, \dots, c_r)$ and $h_{r'} = (c'_0, c'_1, \dots, c'_{r'})$ with $c'_r = c_{r'}$ and $i \in \Delta(c_r)$.

(ii) $I_u \neq \emptyset$.

This is true because, if $I_u = \emptyset$ and $(u, v) \notin A'_\zeta \cup A_\zeta$, then j 's arrival time at u will be a constant and j 's waiting time on edge (u, v) will also be a constant, contradicting the fact that $j \in I_v$. If case (i) happens, then obviously there is a path from v_0 to v . Otherwise, $I_u \neq \emptyset$, we can find a directed path from v_0 to v by backward induction. \square

Define $V_\zeta := \{v \in \bar{V} \mid \tau^v \neq \infty\}$. It follows from Claim 1 that $\{v \mid I_v \neq \emptyset\} \subseteq V_\zeta$. Since the graph \bar{G} is acyclic, there exists a full order among the vertices in V_ζ , such that each directed edge's tail vertex's order is smaller than the head vertex's order. Apparently, v_0 's order is the smallest. So it suffices to prove that, for any vertex $v \in V_\zeta$, player ζ dominates every player $j \in I_v$ at vertex v . We will prove this by induction on the order of the vertices in V_ζ . The base case is obvious because $I_{v_0} = \emptyset$. Now for some vertex $w \in V_\zeta$, assume the above statement is true for all vertices with orders smaller than w in V_ζ , we prove that it's also the case for vertex w .

Since the case $I_w = \emptyset$ is trivial, we suppose now $I_w \neq \emptyset$. For any player $j \in I_w$, suppose edge $(u, w) \in A_j$. In the following, we prove first that player ζ dominates player j at vertex u , then show the dominance at vertex w . If $j \in I_u$, since u has a smaller order than w , then by the assumption, player ζ dominates player j at vertex u . If $j \notin I_u$, then no matter how ζ changes his path, player j 's arrival time at u is a constant. However, since $j \in I_w$, there exists a path $A'_\zeta \in \mathcal{P}[o_\zeta(c_r), d]$ such that $t_j^w(A', \alpha_{-\zeta}) \neq t_j^w(\alpha)$. Suppose w.l.o.g. that $t_j^w(A', \alpha_{-\zeta}) < t_j^w(\alpha)$. Then, combining the above two facts of $j \notin I_u$ and $j \in I_w$, we know one of the two following cases must happen:

- (i) $\exists i \in I_u$ with $uw \in A_i \cap A_j$, s.t., $t_i^u(\alpha) < t_j^u(\alpha)$, or $t_i^u(\alpha) = t_j^u(\alpha)$ and $e_u(A_i) \prec_u e_u(A_j)$.
- (ii) $uw \in A_\zeta \cap A_j$, and $t_\zeta^u(\alpha) < t_j^u(\alpha)$, or $t_\zeta^u(\alpha) = t_j^u(\alpha)$ and $e_u(A_\zeta) \prec_u e_u(A_j)$.

If it is the case (i), then combining the assumption that ζ dominates all players in I_u at vertex u and the fact that $\tau_j^u = t_j^u(\alpha)$, we can deduce that ζ dominates player j at vertex u ; If it is the case (ii), then apparently ζ still dominates player j at vertex u . Next we prove ζ dominates j at vertex w .

It can be observed from the above analysis that no matter whether $j \in I_u$ or not, player ζ always dominates player j at vertex u . Now suppose path A_ζ^* satisfies $\tau^u = t_\zeta^u(A_\zeta^*, \alpha_{-\zeta})$ (if there are more than one such paths, then let A_ζ^* be the one with the highest edge priority at vertex u), and $A'_\zeta \in \mathcal{P}[o_\zeta(c_r), d]$ satisfies $\tau_j^w = t_j^w(A'_\zeta, \alpha_{-\zeta})$. Define $\bar{A}_\zeta = A_\zeta^*[o_\zeta(c_r), u] \cup A_j[u, d]$, then apparently $\bar{A}_\zeta \in \mathcal{P}[o_\zeta(c_r), d]$. Under the strategy profile $(\bar{A}_\zeta, \alpha_{-\zeta})$, consider first the case that ζ travels along the edge uw and there is no queue. In this case, $t_\zeta^w(\bar{A}_\zeta, \alpha_{-\zeta}) = \tau^u + 1 \leq \tau_j^u + 1 \leq t_j^u(A'_\zeta, \alpha_{-\zeta}) + 1 \leq t_j^w(A'_\zeta, \alpha_{-\zeta}) = \tau_j^w$. Combining with the facts that $\tau^w \leq t_\zeta^w(\bar{A}_\zeta, \alpha_{-\zeta})$ and $e_w(\bar{A}_\zeta) = e_w(A_j)$, we can deduce that player ζ dominates player j at vertex w . Now we are left with the case that ζ travels along the edge uw under the strategy profile $(\bar{A}_\zeta, \alpha_{-\zeta})$ and there is a queue. Let I' be the set of players queuing before ζ and those who pass through uw earlier than that queue. Let $i \in I'$ be the player that queues immediately before ζ , i.e., $t_\zeta^w(\bar{A}_\zeta, \alpha_{-\zeta}) = t_i^w(\bar{A}_\zeta, \alpha_{-\zeta}) + 1$. Since ζ dominates all players in I_u and $t_\zeta^u(\bar{A}_\zeta, \alpha_{-\zeta}) = \tau^u$, we can see that $I' \cap I_u = \emptyset$ and ζ cannot dominate any player in I' at vertex u . So no matter how ζ changes his strategy, the time that players in I' pass through edge uw will never be influenced

by ζ or players in I_u , which also means $I' \cap I_w = \emptyset$. Thus, $t_i^w(A'_\zeta, \alpha_{-\zeta}) = \tau_i^w = t_i^w(\bar{A}_\zeta, \alpha_{-\zeta})$. Recall that ζ dominates j at vertex u and $uw \in A_j \cap A_i$. Therefore, no matter how ζ chooses his path strategy, player j will always arrive at vertex w at least one unit of time later than i does. So, by the definition of path A'_ζ , we have $\tau_j^w = t_j^w(A'_\zeta, \alpha_{-\zeta}) \geq t_i^w(A'_\zeta, \alpha_{-\zeta}) + 1 = t_i^w(\bar{A}_\zeta, \alpha_{-\zeta})$. This along with the facts that $\tau^w \leq t_\zeta^w(\bar{A}_\zeta, \alpha_{-\zeta})$ and $e_w(\bar{A}_\zeta) = e_w(A_j)$ still implies that player ζ dominates player j at vertex w . By the arbitrariness of player j in I_w , we can see that player ζ dominates all players in I_w at vertex w . \square

B.2 Proof of Lemma 4

Proof. We only need to show the first equality because the second one is true by definition and the correctness of the first. Since the network is acyclic, we have a natural full order among all vertices that player ζ can reach, i.e., the vertices v with $\tau^v \neq \infty$. We prove by induction on the order of these vertices. The base case that v is the head of the initial edge $e_\zeta(c_r)$ is obvious due to the initial setting. Let us now consider the case that v is not the head of $e_\zeta(c_r)$, and suppose the lemma is true for all vertices with orders smaller than v .

We claim that, for the players in $Q_{uv}^{\tau^u} - Q_{uv}^{\tau^u}(\tau^u, e^*(u))$, no matter how ζ chooses his path, their arrival times at u will never be influenced. Suppose the contrary. Then, by Lemma 2, ζ dominates at least one player $j \in Q_{uv}^{\tau^u} - Q_{uv}^{\tau^u}(\tau^u, e^*(u))$. Note first from the definition of $Q_{uv}^{\tau^u}$ that $\tau_j^u \leq \tau^u$, where τ_j^u is the earliest time that j can reach u when ζ changes his path. By the definition of domination, it can only be the case that $\tau_j^u = \tau^u$ and ζ is able to arrive at u at time τ^u via an edge e' that has a priority no lower than the one taken by j . By definition of $e^*(u)$, the priority of $e^*(u)$ is at least that of e' , and therefore at least that of the edge taken by j . However, this is impossible because $j \notin Q_{uv}^{\tau^u}(\tau^u, e^*(u))$. Hence the claim is valid. It follows from the claim and the definition of τ^u that ζ cannot influence the arrival times of players in $Q_{uv}^{\tau^u} - Q_{uv}^{\tau^u}(\tau^u, e^*(u))$ at v either.

Consequently, if ζ uses edge $uv \in E^-(v)$ to reach v , his arrival time at v is at least $\tau_u + 1 + |Q_{uv}^{\tau^u}| - |Q_{uv}^{\tau^u}(\tau^u, e^*(u))|$. On the other hand, the above value is obtainable by reaching u at τ^u via $e^*(u)$. It follows that the earliest time that ζ can reach v via edge uv is exactly $\tau_u + 1 + |Q_{uv}^{\tau^u}| - |Q_{uv}^{\tau^u}(\tau^u, e^*(u))|$. Since ζ must use one edge in $E^-(v)$ to reach v , the minimization is correct for vertex v . This finishes the proof. \square

B.3 NE properties of interim games

Recall the notions of preemption and weak preemption in Definition 7. If player i weakly preempts player j at a vertex u and both i and j choose to enter the same edge uv , then i preempts j at vertex v . It is possible that player i preempts player j at vertex x and j preempts i at another vertex y (even under equilibrium routings). Note that while the notion of domination compares the arrival times of two players at the same node under possibly different strategy settings, weak preemption compares two arrival times under the same strategy setting.

Letting $\Gamma^n(c_r)$ be an interim game with $c_r = (Q_e^r)_{e \in \bar{E}}$ and S a subset of $\Delta(c_r)$, we use $\Gamma^n(c_r, S)$ to denote the game where only the ones in S play the game (the other players are assumed to disappear and the orders of players of S in any queue Q_e^t , $e \in \bar{E}$ are modified in accordance). In contrast to Example 3, we have the following useful lemma.

Lemma 11. *Let $\Gamma^n(c_r)$ be an interim game and S a player subset of $\Delta(c_r)$. Fix \mathbf{f} as a strategy profile for players in S . If, for any player $j \in \Delta(c_r) \setminus S$, he can never weakly preempt any player in S with any $o_j(c_r)$ - d path in game $\Gamma^n(c_r, S \cup \{j\})$, then in game $\Gamma^n(c_r)$, for any strategy profile \mathbf{r} of players in $\Delta(c_r) \setminus S$, no player outside S can weakly preempt any player of S under (\mathbf{f}, \mathbf{r}) .*

Proof. Suppose on the contrary that $j \notin S$ weakly preempts $i \in S$ at some vertex v under (\mathbf{f}, \mathbf{r}) , and further that $t_j^v(\mathbf{f}, \mathbf{r})$ is the minimum. The minimality implies that players in $\Delta(c_r) \setminus (S \cup \{j\})$ do not weakly preempt those in S before $t_j^v(\mathbf{f}, \mathbf{r})$. So deleting them can only possibly reduce j 's queuing time before time $t_j^v(\mathbf{f}, \mathbf{r})$ and accelerate his arrival time at v . Therefore, j will still weakly preempt i at v under $(\mathbf{f}, \mathbf{r}_{\{j\}})$ in game $\Gamma^n(c_r, S \cup \{j\})$, contradicting the hypothesis. \square

Proof of Lemma 5. For each player $j \in \Delta(c_r) \setminus B$, define $\tau(B, j)$ as the earliest time when j can weakly preempt some player of B in game $\Gamma^n(c_r, B \cup \{j\})$ under strategy profile $(\boldsymbol{\pi}_B, R_j)$ among all possible $o_j(c_r)$ - d paths $R_j \in \mathcal{P}[o_j(c_r), d]$. If player j can never weakly preempt any player in B in this sense, we set $\tau(B, j) := \infty$. We show that $\min\{\tau(B, j) \mid j \in \Delta(c_r) \setminus B\} = \infty$, which, together with Lemma 11, will imply that the inequality of (4.1) is true. The equality of (4.1) will also be valid because as long as players in B follow $\boldsymbol{\pi}_B$, they are not affected by the remaining ones.

Assume the contrary and let B be the smallest nonempty set $\Delta(\boldsymbol{\pi}, [k])$ such that $\tau^* := \min\{\tau(B, j) \mid j \in \Delta(c_r) \setminus B\} < \infty$. This means that, before time τ^* , no player in $\Delta(c_r) \setminus B$ can weakly preempt any one in B in the sense described above. By definition, there exists a player $j \in \Delta(c_r) \setminus B$ who weakly preempts some player $i^* \in B$ under $(\boldsymbol{\pi}_B, R_j)$ for some path $R_j \in \mathcal{P}[o_j(c_r), d]$ at some vertex $u \in R_j \cap \Pi_{i^*}$. Therefore, under $(\boldsymbol{\pi}_B, R_j)$, player i^* just reaches a vertex $u' \in \Pi_{i^*}[o_{i^*}(c_r), u]$ at time τ^* . Since no weak preemption happens before time τ^* , the movement of every player in B before time τ^* will not be affected. By the minimality of τ^* , we have $\tau^* = t_{i^*}^{u'}(\boldsymbol{\pi}_B, \mathbf{r})$ for all partial path profile \mathbf{r} , which, along with the trivial relation $t_{i^*}^{u'}(\boldsymbol{\pi}_B, \mathbf{r}) \leq t_{i^*}^u(\boldsymbol{\pi}_B, \mathbf{r})$, gives

$$\tau^* \leq t_{i^*}^u(\boldsymbol{\pi}_B, \mathbf{r}), \forall \mathbf{r}.$$

We define j^* , using an adaptation of Algorithm 2 with vertex u (resp. B , $\boldsymbol{\pi}_B$) in place of destination d (resp. U , \mathbf{b}) over there, as the player who owns a dominating $o_{j^*}(c_r)$ - u path P_{j^*} with $e_{j^*}(c_r) \in P_{j^*}$ given the choices of $\boldsymbol{\pi}_B$ by players in B . That is, given $\boldsymbol{\pi}_B$, player j^* is not weakly preempted by any player in $\Delta(c_r) \setminus (B \cup \{j^*\})$ when he travels along P_{j^*} , regardless of the choices of players in $\Delta(c_r) \setminus (B \cup \{j^*\})$. Then combining with the minimalities of B and τ^* , we have $t_{j^*}^u(P_{j^*}, \boldsymbol{\pi}_{-j^*}) = \tau^*$ and j^* weakly preempts i^* at vertex u under path profile $(P_{j^*}, \boldsymbol{\pi}_{-j^*})$. Define an $o_{j^*}(c_r)$ - d path $R_{j^*} := P_{j^*} \cup \Pi_{i^*}[u, d] \in \mathcal{P}[o_{j^*}(c_r), d]$. Then $t_{j^*}^d(R_{j^*}, \boldsymbol{\pi}_{-j^*}) \leq t_{i^*}^d(R_{j^*}, \boldsymbol{\pi}_{-j^*})$, where the equality can only happen when $u = d$. If $\exists v \in \Pi_{i^*}$ such that $t_{j^*}^v(R_{j^*}, \boldsymbol{\pi}_{-j^*}) \neq t_{i^*}^v(\boldsymbol{\pi})$, then by Lemma 2 we know, j^* dominates i^* at vertex v under $\boldsymbol{\pi}$, which is a contradiction to Lemma 3 since $\boldsymbol{\pi}$ is an NE. Otherwise $\forall v \in \Pi_{i^*}$, we have $t_{j^*}^v(R_{j^*}, \boldsymbol{\pi}_{-j^*}) = t_{i^*}^v(\boldsymbol{\pi})$. It follows that $t_{j^*}^d(R_{j^*}, \boldsymbol{\pi}_{-j^*}) \leq t_{i^*}^d(R_{j^*}, \boldsymbol{\pi}_{-j^*}) = t_{i^*}^d(\boldsymbol{\pi}) = \tau(\boldsymbol{\pi}, k) < t_{j^*}^d(\boldsymbol{\pi})$, where the last inequality follows from $j^* \notin B$, contradicting the fact that $\boldsymbol{\pi}$ is an NE. This proves the correctness of (4.1).

Once the players in B have chosen their strategies as specified by $\boldsymbol{\pi}_B$, then by the correctness of (4.1), we can apply Algorithm 2 with $U := B$ and $\mathbf{b} := \boldsymbol{\pi}_B$, which provides us a player in $\Delta(c_r) \setminus B$, denoted as j' , who owns a dominating $o_{j'}(c_r)$ - d path in $\mathcal{P}[o_{j'}(c_r), d]$, denoted $R_{j'}$, among players in $\Delta(c_r) \setminus B$. Therefore, by Lemma 10 and $j' \notin B$, for any $j \in \Delta(c_r) \setminus B$ and partial

path profile \mathbf{r} of $\Delta(c_r) \setminus B$, we have

$$t_j^d(\boldsymbol{\pi}_B, \mathbf{r}) \geq t_{j'}^d(\boldsymbol{\pi}_B, R_{j'}) = t_{j'}^d(\boldsymbol{\pi}) = \tau(\boldsymbol{\pi}, k + 1).$$

This implies that the third inequality of (4.1) is valid. The fourth inequality of (4.1) is true due to $i \in B$ and the first equality of (4.1). \square

Proof of Theorem 2. (i) This is simply an interpretation of $t_i^v(\boldsymbol{\pi}) = t_i^v(\boldsymbol{\pi}_B, \mathbf{r})$ with $B = \Delta(\boldsymbol{\pi}, [k])$ for each $k \geq 1$ in Lemma 5.

(ii) For each $k \geq 1$, let $B := \Delta(\boldsymbol{\pi}, [k - 1])$. The equality stated in (4.1) enables us to apply Algorithm 2 and Lemma 10, which provides us a player $i \in \Delta(c_r) \setminus B$ who possesses a dominating path R_i provided the players in B follow their routes as in $\boldsymbol{\pi}$. Let τ^* denote the earliest time a player in $\Delta(c_r) \setminus B$ reaches d among all routings of $\Gamma^n(c_r)$ in which players in B take their routes as in $\boldsymbol{\pi}$. It follows from Lemma 10 that for any partial strategy profile \mathbf{r} of players in $\Delta(c_r) \setminus B$ and any player $j \in \Delta(c_r) \setminus B$ it holds that

$$t_j^d(\boldsymbol{\pi}_{[k-1]}, \mathbf{r}) \geq \tau^* = t_i^d(\boldsymbol{\pi}_{[k-1]}, R_i, \mathbf{r}_{\Delta(c_r) \setminus (B \cup \{i\})}).$$

In particular, the arbitrary choice of \mathbf{r} gives

$$t_j^d(\boldsymbol{\pi}) \geq \tau^* = t_i^d(R_i, \boldsymbol{\pi}_{\Delta(c_r) \setminus \{i\}}).$$

Since i cannot be better off via deviation to R_i , the minimality of τ^* enforces $t_i^d(\boldsymbol{\pi}) = \tau^*$ and $i \in \Delta(\boldsymbol{\pi}, k)$.

(iii) This can be shown easily by taking $\mathbf{r} = \boldsymbol{\pi}_{\Delta(c_r) \setminus \Delta(\boldsymbol{\pi}, [k])}$ in the first inequality of (4.1) in Lemma 5.

(iv) Suppose on the contrary that there exists a set $S \subseteq \Delta(c_r)$ of players who are able to strictly better off through deviating together from an NE $\boldsymbol{\pi}$ of $\Gamma^n(c_r)$. Let k be the smallest number such that $S \cap \Delta(\boldsymbol{\pi}, k) \neq \emptyset$. Due to Hierarchal Optimality in (ii), all players in $\Delta(\boldsymbol{\pi}, k)$ obtain their optimal latencies provided that no player in $\Delta(\boldsymbol{\pi}, [k - 1])$ deviates from $\boldsymbol{\pi}$, a contradiction. \square

Recalling original priorities defined in Definition 8, let players in Δ be indexed as $1, 2, \dots$ according to their original priorities (smaller indices correspond to higher priorities). We have the following straightforward corollary of Lemma 4 and Global FIFO stated in Theorem 2.

Corollary 2. *If $\boldsymbol{\pi}$ is an NE of Γ^n , then it satisfies the following properties:*

- (i) *Weak Earliest Arrival. For each player i and a partial path profile of players other than i , player i possesses a best response that is Earliest Arrival. This implies that any NE for a new game with the restriction that all players take Earliest Arrival paths is still an NE without this restriction.*
- (ii) *Consecutive Exiting. The indices of players within the same batch under $\boldsymbol{\pi}$ are consecutive. That is, if $i, j \in \Delta(\boldsymbol{\pi}, k)$ with $i < j$, then $h \in \Delta(\boldsymbol{\pi}, k)$ for all $i \leq h \leq j$.*
- (iii) *Temporal Overtaking. If under $\boldsymbol{\pi}$ player i overtakes player j at some vertex $v \in V \setminus \{o\}$ (i.e., i enters G with a lower original priority but reaches v earlier than j does), then under $\boldsymbol{\pi}$ they reach the destination d at the same time. Namely, the overtaking is temporal.*

B.4 Constructing an SPE from a given NE

Recall that $(e_i)_{i \in \Delta(c_{r-1})}$ is the action profile for players in $\Delta(c_{r-1})$ at configuration c_{r-1} such that no action in the profile deviates from ρ . More specifically, for each $i \in \Delta(c_{r-1})$, letting f_i denote the first edge of P_i , we have

- $e_i = f_i$ if i queues on e_i after someone else under c_{r-1} ;
- $e_i (\neq f_i)$ is the second edge of P_i if i queues first on f_i under c_{r-1} ;
- e_i is a null action if i queues first on f_i under c_{r-1} and f_i 's head is d .

For every $i \in \Delta(c_r)$, we observe that either $e'_i = f_i$ and $o_i(c_r) = o_i(c_{r-1})$, or $e'_i \neq f_i$ and $o_i(c_r)$ is the head of f_i . Therefore, given any $o_i(c_r)$ - d path $X_i \in \mathcal{P}[o_i(c_r), d]$, $\{f_i\} \cup X_i$ is an $o(c_{r-1})$ - d path in $\mathcal{P}[o_i(c_{r-1}), d]$.

Recall Construction I with the definition of \mathbb{k} in (4.2), Construction II and all the other settings in Subsection 4.5. Observe that either $\Delta(c_{r-1}) = \Delta(c_r)$, or $\Delta(c_{r-1}) \setminus \Delta(c_r) \neq \emptyset$ and each player in $\Delta(c_{r-1}) \setminus \Delta(c_r)$ exits \bar{G} at time r , giving $\Delta(c_{r-1}) \setminus \Delta(c_r) = \Delta(\rho, 1) \subseteq \Delta(\rho, [\mathbb{k}])$. In any case we have

$$\Delta(c_{r-1}) \setminus \Delta(c_r) \subseteq \mathbb{B} := \Delta(\rho, [\mathbb{k}]). \quad (\text{B.1})$$

Moreover, we recall that ρ is an NE of $\Gamma^n(c_{r-1})$ and the Hierarchal Independence of ρ stated in Theorem 2: as long as the chosen paths of players in \mathbb{B} remain as in $\rho_{\mathbb{B}}$, no matter what paths the players in $\Delta(c_{r-1}) \setminus \mathbb{B}$ choose, they have no impact on the arrival time of any player in \mathbb{B} at any vertex. This implies the following important property, which is the base of Construction II.

Lemma 12. *For any partial path profile $\chi = (X_j)_{j \in \Delta(c_r) \setminus \mathbb{B}}$, where X_j can be any $o_j(c_r)$ - d path in $\mathcal{P}[o_j(c_r), d]$, any player $i \in \mathbb{B} \cap \Delta(c_r)$ and any vertex $v \in P'_i$, it holds that $t_i^v(\rho) = t_i^v(\rho_{\mathbb{B}}, (\{f_j\} \cup X_j)_{j \in \Delta(c_r) \setminus \mathbb{B}}) = t_i^v((P'_j)_{j \in \mathbb{B} \cap \Delta(c_r)}, \chi)$.*

Proof. Notice from (B.1) that $(\rho_{\mathbb{B}}, (\{f_j\} \cup X_j)_{j \in \Delta(c_r) \setminus \mathbb{B}})$ is a strategy profile of game $\Gamma^n(c_{r-1})$. So the first equality of the conclusion follows from the Hierarchal Independence in Theorem 2(i). The second equality is straightforward from the fact that each player in $\Delta(c_{r-1}) \setminus \Delta(c_r)$ (if any) belongs to $\Delta(\rho, 1) \subseteq \mathbb{B}$, and he has only a null action under c_{r-1} in any case, which has no effect on other players. \square

For easy expression of null actions e'_j of players in $j \in \Delta(c_{r-1}) \setminus \Delta(c_r)$, we reserve symbol ϕ for the profile $(e'_j)_{j \in \Delta(c_{r-1}) \setminus \Delta(c_r)}$ of null actions.

Lemma 13. *ρ' is an NE of game $\Gamma^n(c_r)$.*

Proof. We need to prove that $t_i^d(\rho') \leq t_i^d(X'_i, \rho'_{\Delta(c_r) \setminus \{i\}})$ holds for all player $i \in \Delta(c_r)$ and all $o_i(c_r)$ - d path $X'_i \in \mathcal{P}[o_i(c_r), d]$.

Case 1. $i \in \mathbb{B}$. Suppose $i \in \Delta(\rho, k)$ for some $k \leq \mathbb{k}$. Therefore, for any path profile $\chi = (X_j)_{j \in \Delta(c_r)}$ of $\Gamma^n(c_r)$, with $B := \Delta(\rho, [k-1])$ we have

$$\begin{aligned} t_i^d(\rho') &= t_i^d(\rho) \\ &\leq t_i^d(\rho_B, (\{f_j\} \cup X_j)_{j \in \Delta(c_r) \setminus B}, \phi_{\Delta(c_{r-1}) \setminus \Delta(c_r) \setminus B}) \\ &= t_i^d(\rho'_{B \cap \Delta(c_r)}, \chi_{\Delta(c_r) \setminus B}), \end{aligned}$$

where the first equality is by Lemma 12, the inequality is from Hierarchal Optimality in Theorem 2(ii) and the last equality is due to again the definition of ρ' . In particular, when taking $X_i = X'_i$ and $X_j = P'_j$ for every $j \in \Delta(c_r) \setminus \{i\}$, we obtain $t_i^d(\rho') \leq t_i^d(X'_i, \rho'_{\Delta(c_r) \setminus \{i\}})$.

Case 2. $i \in \Delta(c_r) \setminus \mathbb{B}$. It can be seen from Lemma 10 that the path P_j of each player $j \in \Delta(c_r) \setminus \mathbb{B}$ is his best response to other players' choices, giving $t_i^d(\rho') \leq t_i^d(X'_i, \rho'_{\Delta(c_r) \setminus \{i\}})$. \square

Proof of Theorem 3. Let $\sigma = (\sigma_i)_{i \in \Delta}$ be a strategy profile of Γ defined as follows: at each history $h_r = (c_0, \dots, c_r)$, players in $\Delta(c_r)$ take actions as specified by the NE $\pi(h_r)$ constructed in Section 4.5 for h_r , where $\pi(c_0) = \pi$. Similar to the proof of Theorem 1, it can be deduced from Construction I (i.e., (4.3)) and Construction II (and Lemma 10) that for each history h_r the path profile induced by h_r and σ is exactly $\pi(h_r)$.

To see that σ is an SPE of Γ , we fix an arbitrary $r \geq 0$ and an arbitrary history $h_r = (c_0, \dots, c_r) \in \mathcal{H}_r$. Let $\rho' = (P'_i)_{i \in \Delta(c_r)}$ denote the NE $\pi(h_r)$ of $\Gamma^n(c_r)$ we have constructed for h_r . In case of $r = 0$, we set $\rho' := \pi$. Moreover, we consider any $i \in \Delta(c_r)$, any $\sigma'_i \in \Sigma_i$, and the path profile $\chi = (X_j)_{j \in \Delta(c_r)}$ induced by h_r and $\sigma' := (\sigma'_i, \sigma_{-i})$. We need to verify that $t_i(\sigma|h_r) \leq t_i(\sigma'|h_r)$.

If $r = 0$, then we suppose that $i \in \Delta(\pi(c_r), k)$ and write $B = \Delta(\pi, [k-1])$. Construction I, i.e., (4.3), implies that $X_j = P'_j$ for all $j \in B$, and in turn Hierarchal Optimality in Theorem 2(ii) says that $t_i(\sigma|h_0) = t_i^d(\pi) \leq t_i^d(\pi_B, \chi_{\Delta \setminus B}) = t_i^d(\chi) = t_i(\sigma'|h_0)$.

So we assume now $r \geq 1$. Therefore h_r is a child history of some (unique) history $h_{r-1} = (c_0, \dots, c_{r-1}) \in \mathcal{H}_{r-1}$. Let $\rho = (P_j)_{j \in \Delta(c_{r-1})}$ denote the NE $\pi(h_{r-1})$ of $\Gamma^n(c_{r-1})$, \mathbb{k} be defined as in (4.2), and $\mathbb{B} := \Delta(\rho, \mathbb{k})$.

If $i \in \Delta(\rho, k)$ for some $k \leq \mathbb{k}$, then (4.3) implies that $X_j = P_i[o_i(c_r), d] = P'_j$ for all $j \in B \cap \Delta(c_r)$, where $B = \Delta(\rho, [k-1])$. As in Case 1 of the proof of Lemma 13 we deduce that $t_i(\sigma|h_r) = t_i^d(\rho') \leq t_i^d(\rho'_{B \cap \Delta(c_r)}, \chi_{\Delta(c_r) \setminus B}) = t_i^d(\chi) = t_i(\sigma'|h_r)$.

It remains to consider the case of $i \in \Delta(c_r) \setminus \mathbb{B}$. Assume w.l.o.g. that i is exactly the i th player in the ordering $1, 2, \dots$ of players in $\Delta(c_r) \setminus \mathbb{B}$ associated with the iterative dominating path profile constructed in Construction II. Again (4.3) guarantees $\chi_{\mathbb{B} \cap \Delta(c_r)} = \rho'_{\mathbb{B} \cap \Delta(c_r)}$. It follows from Lemma 10 that $\chi_{[i-1]} = \rho'_{[i-1]}$, and $t_i(\sigma|h_r) = t_i^d(\rho') \leq t_i^d(\rho'_{\mathbb{B} \cap \Delta(c_r)}, \rho'_{[i-1]}, \chi_{\Delta(c_r) \setminus \mathbb{B} \setminus [i-1]}) = t_i^d(\chi) = t_i(\sigma'|h_r)$, which completes the proof. \square

C Details in Section 5

Proof of Lemma 6. By symmetry, it suffices to prove $n_1 \leq 2\Lambda L(2\Lambda + n_2)$. For $i = 1, 2$, let S_i denote the set of n_i players who are inside $G_i = (V_i, E_i)$ at time t . Denote by o' and d' the origin and destination nodes of the resulting network from connecting G_1 and G_2 in parallel.

Take ζ to be a fixed player in S_1 who enters G_1 the latest. Then all the players in $S_1 \setminus \Delta_\zeta^{o'}$ reach o' before time $t_\zeta^{o'}$. Among them, we choose a subset H of Λ players who reach d' under π as late as possible (note that the players in H cannot reach d' at the same time). It follows that each $h \in H$ reaches d' later than those in $S_1 \setminus (\Delta^{o'} \cup \{\zeta\} \cup \Delta_h^{d'} \cup H)$. Since each time at most $|E_1^-(d')|$ players in G_1 can reach d' at the same time, and $|E_1^-(d')|$, the in-degree of d' in G_1 , is

at most $\Lambda - 1$, we have

$$t_h^{d'}(\boldsymbol{\pi}) \geq t + \left\lceil \frac{|S_1 \setminus (\Delta_\zeta^{o'} \cup \Delta_h^{d'} \cup H)|}{|E_1^-(d')|} \right\rceil \geq t + \frac{n_1 - 3\Lambda}{\Lambda - 1}, \quad h \in H. \quad (\text{C.1})$$

Moreover, from the choice of ζ and $H \subseteq S_1 \setminus \Delta_\zeta^{o'}$ we derive

$$t_h^{o'}(\boldsymbol{\pi}) < t_\zeta^{o'}(\boldsymbol{\pi}) \leq t, \quad h \in H. \quad (\text{C.2})$$

Since h reaches o' earlier than ζ , it is immediate from the FIFO property in Theorem 2(iii) that $t_h^d(\boldsymbol{\pi}) \leq t_\zeta^d(\boldsymbol{\pi})$ for all $h \in H$. Since $|\{\zeta\} \cup H| = |\Lambda| + 1 > |\Delta_\zeta^d|$, there exists $\eta \in H$ such that

$$t_\eta^d(\boldsymbol{\pi}) < t_\zeta^d(\boldsymbol{\pi}). \quad (\text{C.3})$$

Next, we distinguish between two cases depending on whether there exists a player who enters G_2 within time period $(t_\zeta^{o'}(\boldsymbol{\pi}), t + L(2\Lambda + n_2)]$.

Case 1. Under $\boldsymbol{\pi}$ no player enters G_2 within time period $(t_\zeta^{o'}(\boldsymbol{\pi}), t + L(2\Lambda + n_2)]$. Consider the routing profile $\mathbf{r} = (R_\zeta, \boldsymbol{\pi}_{\Delta \setminus \{\zeta\}})$, in which ζ changes the part of his route from $\Pi_\zeta[o', d'] \subseteq G_1$ to an o' - d' path $R_\zeta[o', d'] \subseteq G_2$, while the other parts of his route, $R_\zeta[o, o'] = \Pi_\zeta[o, o']$ and $R_\zeta[d', d] = \Pi_\zeta[d', d]$, are kept unchanged. By (C.3) and Lemma 5, we see that ζ cannot preempt η (i.e., reach some vertex earlier than η) under \mathbf{r} , which along with (C.1) gives $t + \frac{n_1 - 3\Lambda}{\Lambda - 1} \leq t_\eta^{d'}(\boldsymbol{\pi}) \leq t_\zeta^{d'}(\mathbf{r}) \leq t + L(2\Lambda + n_2)$, then $n_1 \leq 3\Lambda + (\Lambda - 1)L(2\Lambda + n_2)$ and we are done. So we assume

$$t_\zeta^{d'}(\mathbf{r}) > t + L(2\Lambda + n_2). \quad (\text{C.4})$$

Let S consist of players who enter G_2 no later than $t_\zeta^{o'}(\boldsymbol{\pi})$ under $\boldsymbol{\pi}$. According to the condition of Case 1, inequality (C.4) and the series-parallel structure of G , it is not hard to see that from time $t_\zeta^{o'}(\mathbf{r}) = t_\zeta^{o'}(\boldsymbol{\pi})$ to time $t + L(2\Lambda + n_2)$, no player outside S can queue before ζ at the same edge at the same time of ζ under \mathbf{r} . Since all players in $S \setminus \Delta_\zeta^{o'}$ preempt ζ under $\boldsymbol{\pi}$, it follows from the FIFO property in Theorem 2(iii) that the players of $S \setminus (\Delta_\zeta^{o'} \cup \Delta_\zeta^d)$ all reach d earlier than time $t_\zeta^d(\boldsymbol{\pi})$. Then by Lemma 5, players in Δ_ζ^d cannot preempt any player of $S \setminus (\Delta_\zeta^{o'} \cup \Delta_\zeta^d)$ under \mathbf{r} . Since at time t the players in $S \setminus (\Delta_\zeta^{o'} \cup \Delta_\zeta^d)$ who are inside G_2 is a subset of S_2 , we see that under \mathbf{r} player ζ reaches d' no later than time $t + |E(R_\zeta[o', d'])| \cdot |S_2 \cup \Delta_\zeta^{o'} \cup \Delta_\zeta^d| \leq t + L(n_2 + 2\Lambda)$, because under \mathbf{r} by the time $t + L(n_2 + 2\Lambda)$ the transit time of ζ through each edge of $E(R_\zeta[o', d'])$ is at most $|S_2 \cup \Delta_\zeta^{o'} \cup \Delta_\zeta^d|$, a contradiction to (C.4).

Case 2. There exists $\psi \in \Delta$ who enters G_2 at an earliest time $t_\psi^{o'}(\boldsymbol{\pi}) \in (t_\zeta^{o'}(\boldsymbol{\pi}), t + L(2\Lambda + n_2)]$. If a player $j \in \Delta$ preempts ψ at some time point in $[t_\psi^{o'}(\boldsymbol{\pi}), t_\psi^{d'}(\boldsymbol{\pi})]$ under $\boldsymbol{\pi}$, then, by Theorem 2(iii), either $j \in \Delta_\psi^d$ or j enters G_2 no later than ψ does. The definition of ψ (i.e., the minimality of $t_\psi^{o'}(\boldsymbol{\pi})$) implies that either $j \in \Delta_\psi^d \cup \Delta_\psi^{o'}$ or j enters G_2 at time $t_j^{o'}(\boldsymbol{\pi}) \leq t_\zeta^{o'}(\boldsymbol{\pi}) \leq t$. So either $j \in \Delta_\psi^d \cup \Delta_\psi^{o'}$, or $j \in S_2$, or $t_j^{d'} \leq t$ (j leaves G_2 at time t). It follows that

$$t_\psi^{d'}(\boldsymbol{\pi}) \leq \max\{t, t_\psi^{o'}(\boldsymbol{\pi})\} + |E(\Pi_\psi[o', d'])| \cdot |\Delta_\psi^d \cup \Delta_\psi^{o'} \cup S_2| \leq t + 2L(2\Lambda + n_2). \quad (\text{C.5})$$

Recalling the definition of ψ and (C.2), we have $t_{\psi}^{o'}(\boldsymbol{\pi}) > t_{\zeta}^{o'}(\boldsymbol{\pi}) > t_h^{o'}(\boldsymbol{\pi})$ for all $h \in H$. If $t_{\psi}^d(\boldsymbol{\pi}) < t_h^{o'}(\boldsymbol{\pi})$ for all $h \in H$, then we deduce from Consecutive Exiting in Corollary 2(ii) that all players in $\{\psi\} \cup H$ reach d at the same time under $\boldsymbol{\pi}$, i.e., $t_{\psi}^d(\boldsymbol{\pi})$. However, $\{\psi\} \cup H \subseteq \Delta_{\psi}^d$ and $|\{\psi\} \cup H| = \Lambda + 1$ contradict $|\Delta_{\psi}^d| \leq \Lambda$. Hence, we have $t_{\psi}^{d'}(\boldsymbol{\pi}) \geq t_h^{d'}(\boldsymbol{\pi})$ for some $h \in H$. In turn, (C.5) and (C.1) give $t + 2L(2\Lambda + n_2) \geq t_{\psi}^{d'}(\boldsymbol{\pi}) \geq t + \frac{n_1 - 3\Lambda}{\Lambda - 1}$, yielding $n_1 \leq 3\Lambda + 2L(\Lambda - 1)(2\Lambda + n_2) \leq 2\Lambda L(2\Lambda + n_2)$, as desired. \square

Proof of Lemma 7. Suppose G_s is obtained from G_1 and G_2 by identifying d_1 and o_2 . We distinguish between two cases.

In the case of $|\Xi(G_1)| \leq |\Xi(G_2)|$, clearly $\Xi(G_s)$ is just $\Xi(G_1)$. To prove $G_s \in \mathfrak{F}$, as $G_1, G_2 \in \mathfrak{F}$, it suffices to show that the number of players G_2^l can accommodate is upper bounded at any time. Since G_1^r can accommodate at most F_1^r players, and $\Xi(G_2)$ is full as long as G_2^l accommodates more than F_2^l players, we deduce that $\Xi(G_2)$ is full as long as $G_1^r \cup G_2^l$ accommodates more than $F_1^r + F_2^l$ players. Consider any time point t such that $G_1^r \cup G_2^l$ accommodates at most $F_1^r + F_2^l$ players at time t , and more than $F_1^r + F_2^l$ players at time $t + 1$. Then, at time $t + 1$, we know $G_1^r \cup G_2^l$ accommodates at most $F_1^r + F_2^l + |\Xi(G_1)|$ players and $\Xi(G_2)$ is full. So $|\Xi(G_2)|$ players leave $G_1^r \cup G_2^l$ at time $t + 2$ while at most $|\Xi(G_1)|$ players enter $G_1^r \cup G_2^l$. It follows from $|\Xi(G_1)| \leq |\Xi(G_2)|$ that the number of players inside $G_1^r \cup G_2^l$ is nonincreasing unless the number decreases below $F_1^r + F_2^l + 1$. Therefore, at any time, $G_1^r \cup G_2^l$ and consequently G_2^l can accommodate at most $F_1^r + F_2^l + |\Xi(G_1)|$ players.

In the case of $|\Xi(G_1)| > |\Xi(G_2)|$, we see that $\Xi(G_s)$ is just $\Xi(G_2)$. To prove $G_s \in \mathfrak{F}$, we only need to show that $\Xi(G_2)$ is full as long as $G_s^l (= G_1 \cup G_2^l)$ accommodates more than a certain finite number of players. Let t be a time when $\Xi(G_2)$ is not full and $G_1 \cup G_2^l$ accommodates F ($> F_1^l + F_1^r + F_2^l$) players, then it must be the case that $F \leq F_1^l + (1 + \Lambda)(F_1^r + F_2^l)$. Since $G_2 \in \mathfrak{F}$, we know G_2^l accommodates fewer than F_2^l players at time t . It follows that G_1^l accommodates at least $F - (F_1^r + F_2^l) > F_1^l$ players. Let t' be the latest time before t when G_1^l accommodates at most F_1^l players. Since from $t' + 1$ to t , at most Λ players can enter G_s from o_1 at each time, then

$$F - (F_1^r + F_2^l) - F_1^l \leq (t - t')\Lambda.$$

Since, by the choice of t' , at time $t' + 1, t' + 2, \dots, t$, there are more than F_1^l players in G_1^l , cut $\Xi(G_1)$ is full at all these consecutive $t - t'$ time points. Therefore, at time t , the number of players inside $G_1^r \cup G_2^l$ is at least

$$(t - t')(|\Xi(G_1)| - |\Xi(G_2)|) \geq t - t' \geq (F - F_1^l - F_1^r - F_2^l)/\Lambda.$$

On the other hand, by $G_1 \in \mathfrak{F}$ and the choice of t , there are at most $F_1^r + F_2^l$ players inside $G_1^r \cup G_2^l$ at time t . It follows from $(F - F_1^l - F_1^r - F_2^l)/\Lambda \leq F_1^r + F_2^l$ that $F \leq F_1^l + (1 + \Lambda)(F_1^r + F_2^l)$, proving the lemma. \square

Proof of Theorem 5. Given the conditions in the theorem, it is easy to see that $|\bar{E}^+(v)| \geq |\bar{E}^-(v)|$ for any vertex $v \in V \setminus \{o, d\}$. Since G is acyclic, there is a natural full order \prec on its vertices. For all $u, v \in V$, if there is a u - v path in G , we say that v is larger than u . We may also say that an edge $e \in E$ is larger than u if there is a directed path in G starting from u and going through e .

Note that the theorem is equivalent to stating that there is an upper bound on the lengths of the queues on all edges of G . For easy reference, we call an edge that has no such upper bound an *infinity edge*. Using the partial order \prec , we prove by induction to show that each edge has an upper bound. The base cases are easy. In fact, for each maximal vertex $u \in V \setminus \{d\}$, by the apparent boundedness on parallel networks, no outgoing edge from u is an infinity edge.

For any $u \in V \setminus \{d\}$, suppose now there is no vertex of G that is larger than u and has an infinity outgoing edge. Thus, for each edge $e \in E \setminus E^+(u)$ that is larger than u , there exists an upper bound, denoted U_e , on the queue lengths on e . For each out-neighbor v of u , we write $\mathcal{P}[v, d]$ for the set of v - d paths in G , and observe that $U_v := \max_{P \in \mathcal{P}[v, d]} \sum_{e \in P} (U_e + 1)$ is an upper bound on the costs of all the v - d paths.

Let $U_u := 1 + \max_v U_v$, where the maximum is taken over the set of all out-neighbors v of u . Since the traffic is induced by an SPE, at any time, the difference between any two queue lengths of the outgoing edges of u is at most U_u . Therefore, whenever the total number of players on the outgoing edges of u exceeds $(|E^+(u)| - 1)U_u$, each outgoing edge of u has a nonempty queue on it. Since the out-degree of u is no less than its in-degree in \bar{G} , it can be seen that the total number of players on the outgoing edges of u does not increase in the next time step, and hence it never exceeds $(|E^+(u)| - 1)U_u + |E^-(u)|$. This gives an upper bound for the queue lengths of the outgoing edges of u and finishes the proof. \square