

Dynamic Congestion Games: The Price of Seasonality

Marco Scarsini*

Engineering and System Design Pillar
Singapore University of Technology and Design
20 Dover Drive
Singapore 138682
`marco_scarsini@sutd.edu.sg`

Tristan Tomala†

HEC Paris and GREGHEC
1 rue de la Libération
78351 Jouy-en-Josas France
`tomala@hec.fr`

July 12, 2013

*On leave from Dipartimento di Economia e Finanza, LUISS, Roma. Partially supported by PRIN 20103S5RN3 and SRG ESD 2012 044.

†The support of the HEC foundation and of the Agence Nationale de la Recherche under grants ANR JEUDY, ANR-10-BLAN 0112 is gratefully acknowledged.

Abstract

In this paper, we propose a model of discrete time dynamic congestion games with atomic players. This approach allows to give a precise description of the dynamics induced by the individual strategies of players and to study how the steady state is reached, either when players act selfishly, or when the traffic is controlled by a planner. We model also seasonalities by assuming that departure flows fluctuate periodically with time. We focus mostly on simple networks and give closed form formulas for the long-run equilibrium and optimal latencies, as functions of the seasonality. We then derive computations and bounds on the price of anarchy. We also characterize optimal and equilibrium flows and show that, although they produce different costs, they coincide from some time onwards.

Keywords: Network games, price of anarchy, dynamic flows.

OR/MS Subject Classification: networks/graphs: multicommodity, theory; games/group decisions: noncooperative; transportation: models, network.

1 Introduction

The analysis of transportation networks naturally leads to the consideration of congestion games, where each agents selfishly behaves as to minimize her own time on the road without regard for the effects that this behavior has on the other agents' traveling time. The social outcome of the individual selfish behavior can be compared to the outcome that a benevolent planner would choose. A way of comparison is, for instance, the price of anarchy (see, e.g. [Roughgarden, 2005, 2007](#), [Roughgarden and Tardos, 2007](#)), namely the ratio of the worst social cost induced by selfish behavior to the optimal social cost.

Although the matter of this theory is traffic flows, most of the existing literature is actually static. The commonly adopted justification is that the static game represents the steady state of a dynamic model where the flow over the network is constant over time. At least two reasons make the static model unsatisfactory: first, even if the flow of travelers is constant, it may be important to determine how the steady state is reached; second, in real life traffic flows are rarely constant, although they are often periodic or nearly periodic.

In this paper, we study a dynamic model of congestion where the flow of departures fluctuates with time. An example is routing with seasonal flow of traffic, where more drivers enter the system at rush hours. One can also think of firms who use common facilities and who face seasonal demands. The aim is to characterize long-run average latencies imposed on the system by selfish players, socially optimal average latencies and to compute the price of anarchy. In particular, we are interested in analyzing the long-run effect of selfish behavior, and the impact of seasonality. The consideration of periodic rather than constant departures in a dynamic congestion game is novel.

1.1 Model

We consider a dynamic congestion game on a simple network in which finitely many parallel edges link a source to a destination. Time is discrete and at each stage, a generation of finitely many players enters at the source. The number of players entering the network per stage is a periodic function of time. Players are ordered by priorities, and successively choose edges. Each edge e is endowed with a free-flow transit time τ_e and a capacity γ_e . When a player chooses an edge, she enters it right-away and travels on that edge

at a constant speed for τ_e periods. At most γ_e players can exit the edge at the same time and priorities define who goes first. When reaching the edge, the player waits until there remains less than γ_e players with higher priority ahead of her. The latency suffered by the player is thus the sum of the transit and waiting time.

An equilibrium of the game is the result of the selfish behavior of each player who seeks to minimize her own latency. An optimal strategy is the choice of a social planner who allocates the players over the network in order to minimize the limit average total latency. As usual, the price of anarchy is the ratio of the worst equilibrium latency to the optimal latency.

1.2 Results

As a benchmark, in Section 3 we study uniform departures where the same number of players depart at each time and fill up the capacity of the network. We prove that the optimal flow is unique: at each stage, the flow on each edge matches its capacity. Then, we prove that in all equilibria, the flow coincides with the optimal flow from some stage on. Yet, in the worst equilibrium, all players pay eventually the highest transit cost. The intuition is that the first players all choose the fastest edges and induce congestion. Eventually, all edges get so congested that all latencies become equal to the highest transit time. This result shows the impact of the dynamic nature of the model on latencies. While optimal and equilibrium behavior eventually coincide, the selfish behavior of the first players have consequences that cannot be undone and are paid by all future generations.

Then, in Section 4, we study sequences of departures which are periodic over time. We assume that departures are sometimes above and sometimes below capacity, but, on average over a period, they fill the capacity. We show that, even in this case, the flows induced by optima and equilibria, averaged over a period, eventually coincide. However, since at some stages players enter above capacity, waiting queues form. We introduce a measure of discrepancy between a periodic sequence and the uniform sequence and prove that, compared with the case of uniform departures, both optimum and equilibrium latencies are incremented by exactly this amount. We provide two alternative ways of computing this measure.

From our characterizations of optimal and equilibrium latencies, we derive computation of the price of

anarchy for this model. The main findings are the following. Firstly, the price of anarchy is bounded by the ratio of the highest to the lowest free-flow transit time, irrespectively of capacities. The bound is tight and is approximately achieved for networks for which fast roads are wide and slow roads are narrow. For these networks, achieving optimality entails that some players accept to use slow roads. Secondly, we are able to compute the price of anarchy as a function of the seasonality of departures and find that the irregularity of the departure process can only reduce the price of anarchy. Intuitively, when too many players depart at the same time, the capacity constraints imply that the planner cannot avoid queues to be created. In the worst case, when the number of players departing at the same time is very large, the equilibrium is nearly optimal.

1.3 Related literature

Our model belongs to the class of congestion games with atomic players. In his fundamental paper, [Wardrop \(1952\)](#) modeled selfish behavior of a huge number of agents on an network as a nonatomic flow and introduced a concept of equilibrium that has become the standard reference in the literature. [Charnes and Cooper \(1961\)](#) showed the relation between Nash and Wardrop equilibria and [Haurie and Marcotte \(1985\)](#) proved that, under some conditions, the Wardrop equilibrium in a nonatomic model can be obtained as a limit of Nash equilibria of atomic models. The relation between atomic and nonatomic games has been recently studied by [Bhaskar, Fleischer, and Huang \(2010\)](#). A nice survey of Wardrop equilibria can be found in [Correa and Stier-Moses \(2010\)](#). General congestion games with a finite number of players were introduced by [Rosenthal \(1973\)](#), who proved that they have pure Nash equilibria; they are actually isomorphic to potential games (see [Monderer and Shapley, 1996](#)). The issue of multiplicity of equilibria in atomic congestion games was studied by [Harker \(1988\)](#), [Bhaskar, Fleischer, Hoy, and Huang \(2009\)](#). Consistently with this literature, we find multiple equilibria for our game.

Importantly, our model is dynamic. Equilibrium concepts in dynamic network models date back to [Vickrey \(1969\)](#) in the economic literature and to [Yagar \(1971\)](#) in the transportation literature. We refer the reader to [Koch and Skutella \(2011\)](#) for an extensive list of references on this topic. Some of the recent mathematical formulations of the model resort to deterministic queueing theory, as introduced originally

by [Vickrey \(1969\)](#) and later developed by [Hendrickson and Kocur \(1981\)](#). In this stream of literature [Akamatsu \(2000, 2001\)](#), [Akamatsu and Heydecker \(2003\)](#), [Mounce \(2006, 2007\)](#), [Anshelevich and Ukkusuri \(2009\)](#), [Hoefler, Mirrokni, Röglin, and Teng \(2009\)](#), and especially [Koch and Skutella \(2011\)](#) extended some results known for static congestion games to dynamic congestion games. These last authors used a deterministic queueing model to study dynamic flows and characterize Nash equilibria, showing the relation between dynamic and static models, and computed the price of anarchy for the dynamic model. Along these lines, [Cominetti, Correa, and Larré \(2011\)](#) studied equilibria for flows over time in the single-source single-sink deterministic queueing model and constructively proved existence and uniqueness of equilibria when the inflow rate is piecewise constant. [Koch, Nasrabadi, and Skutella \(2011\)](#) used measure-theoretic techniques to combine continuous and discrete time models of flow over time and, among other things, extended to this general setting the classical max-flow min-cut theorem. Compared with this literature, an important novelty of our work is to study the impact of the seasonality of flows of players into the system. Moreover most of the literature deals with nonatomic games, the exception being [Hoefler et al. \(2009\)](#). The dynamic of our game is deliberately simple to allow explicit computations of optima and equilibria, and to analyze precisely the impact of the behavior during the transient phase on the steady state. Relatedly, [Shah and Shin \(2010\)](#) considered the transient phase of a dynamic network before a steady state equilibrium is reached. Although their model is stochastic, some of the questions they consider are close in spirit to our model.

For measuring the efficiency of a game, we use the now famous price of anarchy, namely the ratio between the worst Nash equilibrium latency and the social optimum latency. This measure was introduced by [Koutsoupias and Papadimitriou \(1999\)](#), the name was coined by [Papadimitriou \(2001\)](#) and since then, inefficiency of equilibria in routing games has been studied by several authors (see among others [Roughgarden and Tardos \(2002, 2004\)](#), [Correa, Schulz, and Stier-Moses \(2004, 2008, 2007\)](#)).

1.4 Organization of the paper

The paper is organized as follows. Section 2 presents the model. Section 3 considers the case of uniform departures over time. Section 4 deals with the case of periodic departures at capacity over the period.

Section 5 contains several extensions of the main model. All proofs are relegated in the Appendix.

2 Model

2.1 Description of the game

We consider a dynamic congestion game, based on the model of [Ford and Fulkerson \(1958\)](#) of flows on networks. We are given a simple directed network where a source is linked to a destination by finitely many parallel edges (see [Figure 1](#)).




Figure 1: Parallel network.

Each edge has a transit cost and a capacity. Time is discrete and at each stage t , finitely many players depart from the source and choose a path to the destination. The number of departures per stage is throughout assumed to be periodic. The data of the model are thus as follows:

- A directed network $\mathcal{N} = (E, (\tau_e)_{e \in E}, (\gamma_e)_{e \in E})$, composed of
 - a finite \prec -ordered set of edges $E = \{e_1, \dots, e_n\}$ that connect a source to a destination,
 - for each $e \in E$, a *free-flow transit cost* $\tau_e \in \mathbb{N}$ (simply called transit cost in the sequel); without any loss of generality we assume that $e_i \prec e_j$ implies $\tau_i \leq \tau_j$,

- for each $e \in E$, a *capacity* $\gamma_e \in \mathbb{N}$; we denote $\gamma := \sum_{e \in E} \gamma_e$ the capacity of the network.
- An integer K and a vector $\boldsymbol{\delta} = (\delta_1, \dots, \delta_K) \in \mathbb{N}^K$. With an abuse of notation, we use the same symbol $\boldsymbol{\delta}$ to denote the periodic sequence obtained by infinitely concatenating $\boldsymbol{\delta}$. This will be the sequence of departures from the source.

The dynamic of the model is the following.

- At each stage $t \in \mathbb{N}_+$, a finite set G_t of players departs from the source. This is called the *generation* at time t and we assume that $\#G_t = \delta_k$ if $t = k \pmod K$. We denote the i -th player in generation G_t by $[it]$ (when there is no risk of confusion, the square brackets are removed). We thus have an infinite set of players $G := \cup_t G_t$. We order this set (anti-lexicographically) by \triangleleft as follows:

$$[js] \triangleleft [it] \quad \text{iff} \quad s < t \quad \text{or} \quad (s = t \text{ and } j < i).$$

This order represents priorities, if $[js] \triangleleft [it]$ and if these two players are on the same edge at the same time, then $[js]$ exits first.

- Upon departure, each player $[it]$ observes the choices of players $[js] \triangleleft [it]$ and chooses an edge $\varepsilon_{it} \in E$. Denote h_{it} the *history* of information available to player $[it]$

$$h_{it} = \{\varepsilon_{js} : [js] \triangleleft [it]\},$$

and \mathcal{H}_{it} the set of all possible $[it]$ -histories. A *strategy* for player $[it]$ is a mapping $\sigma_{it} : \mathcal{H}_{it} \rightarrow E$, that is $\varepsilon_{it} = \sigma_{it}(h_{it})$. We let \mathcal{S}_{it} be the set of all strategies available to player $[it]$ and $\mathcal{S} := \times_{[it] \in G} \mathcal{S}_{it}$ the set of strategy profiles.

- Player $[it]$ then progresses along edge ε_{it} by taking a step of size 1 per unit of time. At stage $t + \tau_{\varepsilon_{it}}$ she arrives at the exit of the edge where a queue might have formed.
- The rules for exiting the queue are the following:

- All players $[js] \triangleleft [it]$ that are in the queue are ahead of $[it]$.
- At most γ_e players can exit e simultaneously. When she arrives at the end of ε_{it} , if player $[it]$ finds less than $\gamma_{\varepsilon_{it}}$ players in the queue, then she exits immediately; otherwise, only the first $\gamma_{\varepsilon_{it}}$ players exit at this stage and player $[it]$ waits for one stage. This process repeats until there remains less than $\gamma_{\varepsilon_{it}}$ players in the queue ahead of player $[it]$.

These data and rules define a game $\Gamma(\mathcal{N}, K, \delta)$ with perfect information. Each strategy profile $\sigma \in \mathcal{S}$ induces a choice of edge $\varepsilon_{it}(\sigma) \in E$ for each player $[it]$. We denote $\ell_{it}(\sigma)$ the *latency* suffered by player $[it]$, defined as

$$\ell_{it}(\sigma) = c_{it}(\sigma) + w_{it}(\sigma),$$

where $c_{it}(\sigma) := \tau_{\varepsilon_{it}(\sigma)}$ is the *transit cost* paid by player $[it]$ and $w_{it}(\sigma)$ is the *waiting cost* paid by player $[it]$, namely, the number of stages that $[it]$ spends queueing.

We define the *total transit cost* c_t , the *total waiting cost* w_t , and the *total latency* ℓ_t at stage t as follows:

$$\begin{aligned} c_t(\sigma) &= \sum_{[it] \in G_t} c_{it}(\sigma), \\ w_t(\sigma) &= \sum_{[it] \in G_t} w_{it}(\sigma), \\ \ell_t(\sigma) &= \sum_{[it] \in G_t} \ell_{it}(\sigma) = c_t(\sigma) + w_t(\sigma). \end{aligned}$$

For each integer p , the total latency over the period $\{pK + 1, \dots, (p + 1)K\}$ is

$$L_p(\sigma) = \sum_{t=pK+1}^{(p+1)K} \ell_t(\sigma).$$

The *average total latency* over P periods is

$$\bar{L}_P(\sigma) = \frac{1}{P} \sum_{p=1}^P L_p(\sigma).$$

If $\lim_{P \rightarrow \infty} \bar{L}_P(\sigma)$ exists, then it is called *asymptotic average total latency* for the strategy σ .

Definition 2.1. A strategy profile σ is an *equilibrium* if

$$\ell_{it}(\sigma) \leq \ell_{it}(\sigma'_{it}, \sigma_{-it}) \text{ for all } [it] \in G, \text{ for all } \sigma'_{it} \in \mathcal{S}_{it},$$

where σ_{-it} indicates the profile of strategies of all players different from $[it]$. We denote \mathcal{E} the set of equilibria of the game. The quantity

$$\text{Eq}(\mathcal{N}, K, \delta) := \inf_{\sigma \in \mathcal{E}} \lim_{P \rightarrow \infty} \bar{L}_P(\sigma)$$

is called the *worst equilibrium latency*.

Definition 2.2. A strategy profile σ is *socially optimal* (henceforth optimal) if

$$\liminf_{P \rightarrow \infty} \bar{L}_P(\sigma') \geq \limsup_{P \rightarrow \infty} \bar{L}_P(\sigma) \text{ for all } \sigma' \in \mathcal{S}. \quad (2.1)$$

If (2.1) holds for some $\sigma \in \mathcal{S}$, then

$$\text{Opt}(\mathcal{N}, K, \delta) := \liminf_{P \rightarrow \infty} \bar{L}_P(\sigma) = \limsup_{P \rightarrow \infty} \bar{L}_P(\sigma)$$

is called the *optimal latency*.

Observe that the existence of an equilibrium is guaranteed. Indeed, each player cares only about the strategies of the finitely many players ahead of him. Therefore, finding the best-response of a player amounts to solving a finite minimization problem, and finding an equilibrium amounts to solving these problems sequentially. However, with the above definition of optimal strategy, existence has to be proved.

There is obviously an asymmetry between the two definitions, which stems from the type of rationality driving the two concepts. An equilibrium is a strategy profile such that each finitely lived player optimizes given the choices of the other players. By contrast, an optimal strategy is the choice that a long-lived planner would like to take in order to optimize the long-run social welfare. Since we assume that departures are periodic, the model is stationary, and one would expect an optimal strategy to be periodic as well. This statement holds true, as we shall see, but this requires a proof.

3 Uniform departures

In this section, we assume that departures are uniform over time that is, $K = 1$ and there exists an integer $\delta \leq \gamma$ such that the infinite sequence of departures is $\boldsymbol{\delta} := (\delta, \delta, \dots)$. To simplify notation, we write $(\mathcal{N}, \boldsymbol{\delta})$ rather than $(\mathcal{N}, 1, \boldsymbol{\delta})$. We study this benchmark case in details. In the following subsections we characterize optima and equilibria at capacity, i.e., when $\delta = \gamma$, derive estimates of the price of anarchy, and then generalize the results to departures below capacity, i.e., $\delta < \gamma$.

3.1 Departures at capacity

Assume that the sequence of departures is $\boldsymbol{\gamma} := (\gamma, \gamma, \dots)$. Our first main result gives a characterization of optimal and equilibrium strategies and of the associated latencies. We introduce another piece of notation: let $x_t(\sigma) = (x_t^e(\sigma))_{e \in E}$ be the *allocation* of generation t , defined by

$$x_t^e(\sigma) = \#\{[it] \in G_t : \varepsilon_{it}(\sigma) = e\}.$$

The total latency at stage t can be re-written as,

$$\ell_t(\sigma) = \sum_e x_t^e(\sigma) \tau_e + w_t(\sigma)$$

and if we denote

$$\bar{z}_T := \frac{1}{T} \sum_{t=1}^T z_t$$

the average of a finite sequence (z_1, \dots, z_T) , then the average total latency becomes

$$\bar{\ell}_T(\sigma) = \sum_e \bar{x}_T^e(\sigma) \tau_e + \bar{w}_T(\sigma).$$

Theorem 3.1. (a) *The optimal latency $\text{Opt}(\mathcal{N}, \boldsymbol{\gamma})$ exists and*

$$\text{Opt}(\mathcal{N}, \boldsymbol{\gamma}) = \sum_e \gamma_e \tau_e.$$

A strategy σ^{Opt} is optimal if and only if for all $t \in \mathbb{N}_+$ and for all $e \in E$ we have $x_t^e(\sigma^{\text{Opt}}) = \gamma_e$.

(b) There exists a time T_0 such that, for each $t \geq T_0$, in each equilibrium σ , $x_t^e(\sigma) = \gamma_e$ for each edge e .

Furthermore, there exists an equilibrium σ^{Eq} such that for each $t \geq T_0$, $\ell_{it}(\sigma^{\text{Eq}}) = \max_e \tau_e$ for each player $[it] \in G_t$. This strategy σ^{Eq} yields the worst equilibrium latency

$$\text{Eq}(\mathcal{N}, \gamma) = \left(\sum_e \gamma_e \right) \max_e \tau_e.$$

Theorem 3.1 says that optimal strategies exist and are the ones that fill exactly the capacity of each edge at each stage. Regarding equilibria, the main findings are that the flows, or allocations, of any equilibrium coincide eventually with those generated by optimal strategies. The insight is that the selfish behavior of the players in the first generations creates permanent queues that increase the latency of edges for future generations. In the worst case, the latency incurred on each edge is eventually equal to the highest transit cost, which is paid by all players from some time on.

To see that the equilibrium latency may not be unique, consider a network with two edges e_1, e_2 , each having capacity 1, and with transit costs $\tau_{e_i} = i$. There is an equilibrium in which the first two players go on e_1 . The first pays 1 and the second pays 2. From that point on, in each generation, there is one player per edge, and each player pays 2. This is the worst equilibrium, and the associated latency is 2. But since the second player in the first generation is indifferent between the two edges, there is another equilibrium where the first player [11] chooses e_1 and the second player [21] chooses e_2 . Their payoffs are again 1 and 2 respectively. Yet, since the first generation creates no congestion, the next two players can also adopt these strategies and pay respectively 1 and 2. Repeating this at each generation, we get a strategy profile which is both optimal and an equilibrium with long-run latency 1.5.

The formal proof of Theorem 3.1 can be found in the Appendix. The intuition is as follows. First consider point (a). Suppose that a strategy profile sends on some edge an average number of players that is higher than the capacity. This creates an exploding queue on this edge and therefore a diverging average waiting cost. As a result, such a strategy cannot be optimal. Now, if the long-run average number of players does not exceed the capacity on any edge, this means that the strategy is similar to a stationary

one where the allocation exactly matches the capacities at each stage. We conclude that these are the only optimal strategies.

Regarding point (b), the idea is that the selfish players first fill cheap (or short) edges, thereby creating queues. As a result, the latency of these edges increases for future generations, and, at some point, all latencies become equal. From that point on, players are basically indifferent and allocations match capacities.

3.2 The price of anarchy

A direct consequence of Theorem 3.1 is a computation of the price of anarchy.

Corollary 3.2. *The price of anarchy of the game $\Gamma(\mathcal{N}, \gamma)$ is*

$$\text{Poa}(\mathcal{N}, \gamma) := \frac{\text{Eq}(\mathcal{N}, \gamma)}{\text{Opt}(\mathcal{N}, \gamma)} = \frac{(\sum_e \gamma_e) \max_e \tau_e}{\sum_e \gamma_e \tau_e} \leq \frac{\max_e \tau_e}{\min_e \tau_e}.$$

The last inequality is straightforward and shows that the price of anarchy admits an upper bound that does not depend on capacities but only on the relative lengths of edges. This bound is tight. To see this, consider a network with two edges, such that $\tau_1 = 1$, $\tau_2 = N$, $\gamma_{e_1} = N^2$, $\gamma_{e_2} = 1$. The number of players per stage is $N^2 + 1$. The price of anarchy for this network is $(1 + N^2)N / (N^2 + N)$ which is roughly N for N large.

This example shows that the price of anarchy is unbounded over this class of networks. Intuitively, in the worst network the first road is very fast and very large, the second road being slow and narrow. In equilibrium, all players wish to take the fast road. But because of the capacity constraint, a queue builds up on the fast road, so its latency increases until it becomes as costly as the slow road. In steady state, this queue becomes permanent, and all players pay the cost of the slow road. The social optimum is quite different. If one player per generation uses the slow road, then all other players benefit from the fast road without queue.

On the other hand, if roads have comparable length, then the equilibrium is not very different from the social optimum and the price of anarchy is close to one.

3.3 Below capacity

We now assume that the sequence of departures is $\boldsymbol{\delta} = (\delta, \delta, \dots)$ with $\delta < \gamma$ and derive characterizations of optima, equilibria, and price of anarchy, from Theorem 3.1. For convenience, assume that the transit costs are different from one edge to the other, that is, for $e, f \in E$, $e \prec f$ implies $\tau_e < \tau_f$ (this is without loss of generality, see Lemma 6.2).

For $f \in E$ denote $f^\prec := \{e : e \prec f\}$ and $f^\succ := f^\prec \cup \{f\}$, where $e_1^\prec = \emptyset$. For each $\delta < \gamma$, there is a unique $f \in E$ such that

$$\sum_{e \in f^\prec} \gamma_e < \delta \leq \sum_{e \in f^\succ} \gamma_e.$$

Define the (sub-)network \mathcal{N}_δ with set of edges f^\succ , such that each edge $e \in f^\prec$ has transit cost τ_e and capacity γ_e , and edge f has transit cost τ_f and capacity $\delta - \sum_{e \in f^\prec} \gamma_e \leq \gamma_f$. The total capacity of \mathcal{N}_δ is precisely δ . The next result states that optima and equilibria are those of \mathcal{N}_δ deduced from Theorem 3.1.

Proposition 3.3. (a) *The optimal latency $\text{Opt}(\mathcal{N}, \boldsymbol{\delta})$ exists and*

$$\text{Opt}(\mathcal{N}, \boldsymbol{\delta}) = \sum_{e \in f^\prec} \gamma_e \tau_e + \left(\delta - \sum_{e \in f^\prec} \gamma_e \right) \tau_f.$$

A strategy σ^{Opt} is optimal if and only if for all $t \in \mathbb{N}_+$

$$x_t^e(\sigma^{\text{Opt}}) = \gamma_e \quad \text{for } e \prec f \quad \text{and} \quad x_t^f(\sigma^{\text{Opt}}) = \delta - \sum_{e \in f^\prec} \gamma_e.$$

(b) *The worst equilibrium latency $\text{Eq}(\mathcal{N}, \boldsymbol{\delta})$ exists and*

$$\text{Eq}(\mathcal{N}, \boldsymbol{\delta}) = \delta \tau_f.$$

There exists a time T_0 such that in any equilibrium σ^{Eq} , for each $t \geq T_0$,

$$x_t^e(\sigma^{\text{Eq}}) = \gamma_e \quad \text{for } e \prec f \quad \text{and} \quad x_t^f(\sigma^{\text{Eq}}) = \delta - \sum_{e \in f^\prec} \gamma_e.$$

(c) *The price of anarchy*

$$\text{Poa}(\mathcal{N}, \delta) = \frac{\delta \tau_f}{\sum_{e \in f^{\prec}} \gamma_e \tau_e + (\delta - \sum_{e \in f^{\prec}} \gamma_e) \tau_f} \quad (3.1)$$

is decreasing over $\delta \in \{\sum_{e \in f^{\prec}} \gamma_e + 1, \dots, \sum_{e \in f^{\succ}} \gamma_e\}$.

The idea of the proof is simple and we only give informal arguments. It is sub-optimal, both from the individual and from the social perspective, to use edges with transit cost above τ_f . Thus, optimal and equilibrium latencies are computed by applying Theorem 3.1 to the sub-network \mathcal{N}_δ whose capacity is precisely the number of departures. That is, we can find an optimal strategy where players are confined to the sub-network. Regarding equilibria, we remark that each player can choose a best-response in the sub-network, and this for every strategy profile of other players.

Notice that the monotonicity of the price of anarchy holds only within each interval $\sum_{e \in f^{\prec}} \gamma_e < \delta \leq \sum_{e \in f^{\succ}} \gamma_e$. The sequence of prices at the cut-points,

$$\text{Poa}(\mathcal{N}, \sum_{e \in f^{\succ}} \gamma_e) = \frac{(\sum_{e \in f^{\succ}} \gamma_e) \tau_f}{\sum_{e \in f^{\succ}} \gamma_e \tau_e}$$

need not be monotonic with f . To see it, consider the following example.

Example 3.4. Consider a network with $E = \{e_1, e_2, e_3\}$, where $\gamma_{e_i} = 2$ and $\tau_{e_1} = 1, \tau_{e_2} = 5, \tau_{e_3} = 6$. A simple application of (3.1) shows that

$$\text{Poa}(\mathcal{N}, \delta) = \begin{cases} 1 & \text{for } \delta = 2, \\ 10/6 & \text{for } \delta = 4, \\ 9/6 & \text{for } \delta = 6. \end{cases}$$

4 Periodic departures

In this section, we consider periodic sequences of departures $(\delta_1, \dots, \delta_K)$, characterize optimal and equilibrium latencies and derive a computation of the price of anarchy. For most of this section, we assume that the average number of departures equals the capacity of the network: $\sum_{k=1}^K \delta_k = \gamma K$.

4.1 Optimal and equilibrium latencies

To get a feeling of how seasonality affects optimal and equilibrium latencies, consider the following example.

Example 4.1. Consider a network having two edges e_1, e_2 with $\gamma_{e_1} = \gamma_{e_2} = 1$ and $\tau_{e_i} = i$, and assume that the 3-periodic sequence of departures is $\delta = (6, 0, 0)$ where all players depart at the same time. Then, the following strategy which allocates three player per period to each edge is optimal.

$$\sigma^{\text{Opt}}(it) = \begin{cases} e_1 & \text{for } i \text{ odd,} \\ e_2 & \text{for } i \text{ even.} \end{cases}$$

To see it, consider the first two players and send the first one to e_1 , the second one to e_2 . Then, the next two players have to queue at least for one period, so it is as if they had departed one period later and it is optimal to send one of them to e_1 and the other to e_2 . Now, the remaining two players have to queue at least two periods, so it is as if they had departed two periods later, and it is again optimal to send one of them to e_1 and the other to e_2 .

The total latency over a period of time $\{1, 2, 3\}$ (modulo 3) is 15, that is, 3 times the single-period optimal total latency that we would have if departures were uniform $(2, 2, 2)$ plus the added cost of 6 induced by the waiting times: two players pay an extra cost of 1 and two players pay an extra cost of 2.

Consider now the following equilibrium strategy σ^{Eq} . For $t = 1$ we let

$$\sigma^{\text{Eq}}(it) = \begin{cases} e_1 & \text{for } i = 1 \text{ or } i \text{ even,} \\ e_2 & \text{for } i > 2, \text{ odd,} \end{cases}$$

therefore the latencies for the first six players are

$$\begin{aligned} \ell_{11}(\sigma^{\text{Eq}}) &= 1, \quad \ell_{21}(\sigma^{\text{Eq}}) = 2, \quad \ell_{31}(\sigma^{\text{Eq}}) = 2, \\ \ell_{41}(\sigma^{\text{Eq}}) &= 3, \quad \ell_{51}(\sigma^{\text{Eq}}) = 3, \quad \ell_{61}(\sigma^{\text{Eq}}) = 4. \end{aligned}$$

For $t \geq 4$ (again when $t = 1 \pmod{3}$),

$$\sigma^{\text{Eq}}(it) = \begin{cases} e_1 & \text{for } i \text{ odd,} \\ e_2 & \text{for } i \text{ even,} \end{cases}$$

and

$$\begin{aligned} \ell_{1t}(\sigma^{\text{Eq}}) &= 2, \ell_{2t}(\sigma^{\text{Eq}}) = 2, \ell_{3t}(\sigma^{\text{Eq}}) = 3, \\ \ell_{4t}(\sigma^{\text{Eq}}) &= 3, \ell_{5t}(\sigma^{\text{Eq}}) = 4, \ell_{6t}(\sigma^{\text{Eq}}) = 4. \end{aligned}$$

It is easy to check that this is an equilibrium. It is constructed in such a way that each player chooses e_1 when he is indifferent between the two edges. This choice makes it the worst equilibrium. In the steady state, the total equilibrium payoff over a 3-period is 18, that is 3 times the single-period equilibrium total latency when departures are uniform plus the added cost of 6 induced by the waiting times.

In the same network, when the 3-periodic sequence of departures is $\delta = (3, 2, 1)$, then the following strategy σ^{Opt} is still optimal

$$\sigma^{\text{Opt}}(it) = \begin{cases} e_1 & \text{for } i \text{ odd,} \\ e_2 & \text{for } i \text{ even.} \end{cases}$$

The last player at stage 1 has to wait for one period. The situation is then as if she had arrived at stage 2, thus the last player at stage 2 also has to wait for one period. The total latency over any period of time $\{1, 2, 3\}$ (modulo 3) is thus 11, that is, 3 times the single-period optimal total latency that we would have if departures were uniform $(2, 2, 2)$ plus the added cost of 2 induced by the waiting times.

Consider now the (worst) equilibrium strategy σ^{Eq} . We have

$$\begin{aligned} \sigma^{\text{Eq}}(11) &= \sigma^{\text{Eq}}(21) = \sigma^{\text{Eq}}(12) = \sigma^{\text{Eq}}(13) = e_1, \\ \sigma^{\text{Eq}}(31) &= \sigma^{\text{Eq}}(22) = e_2, \end{aligned}$$

therefore the latencies for the first six players are

$$\begin{aligned} \ell_{11}(\sigma^{\text{Eq}}) &= 1, \ell_{21}(\sigma^{\text{Eq}}) = 2, \ell_{31}(\sigma^{\text{Eq}}) = 2, \\ \ell_{12}(\sigma^{\text{Eq}}) &= 2, \ell_{22}(\sigma^{\text{Eq}}) = 2, \ell_{13}(\sigma^{\text{Eq}}) = 2, \end{aligned}$$

and then, for $t \geq 4$,

$$\sigma^{\text{Eq}}(it) = \begin{cases} e_1 & \text{for } i \text{ odd,} \\ e_2 & \text{for } i \text{ even,} \end{cases}$$

and

$$\begin{aligned} \ell_{1t}(\sigma^{\text{Eq}}) &= 2, \ell_{2t}(\sigma^{\text{Eq}}) = 2, \ell_{3t}(\sigma^{\text{Eq}}) = 3, & \text{for } t = 1 \pmod{3}, \\ \ell_{1t}(\sigma^{\text{Eq}}) &= 2, \ell_{2t}(\sigma^{\text{Eq}}) = 3, & \text{for } t = 2 \pmod{3}, \\ \ell_{1t}(\sigma^{\text{Eq}}) &= 2, & \text{for } t = 3 \pmod{3}. \end{aligned}$$

So in steady state the total equilibrium payoff over a 3-period is 14, that is 3 times the single-period equilibrium total latency when departures are uniform plus the added cost of 2 induced by the waiting times.

We see from this example that, since departures are not uniform, some queues have to be created and some players must pay a positive waiting cost. We define now a quantity that measures the total waiting cost.

Let $\mathbb{N}_K(\gamma)$ be the set of K -periodic sequences $(\delta_1, \dots, \delta_K)$ such that $\sum_{k=1}^K \delta_k = \gamma K$. Define the following binary relation on $\mathbb{N}_K(\gamma)$:

Definition 4.2. For any two elements $\delta, \delta' \in \mathbb{N}_K(\gamma)$, we say that δ' is obtained from δ by an elementary

operation (denote it $\delta \rightarrow \delta'$), if there exist times i, j , with $i < j$, such that

$$\begin{aligned} \delta_i &> \gamma, \\ \delta_l &= \gamma, \text{ for all } l \text{ such that } i < l < j, \\ \delta_j &< \gamma, \\ \delta'_i &= \delta_i - 1, \\ \delta'_{i+1} &= \delta_{i+1} + 1, \\ \delta'_l &= \delta_j \text{ for } l \notin \{i, i+1\}, \end{aligned}$$

where indices i, j are considered modulo K .

Consider the directed graph representing this binary relation \rightarrow and denote $D(\delta)$ the distance in this graph from δ to the uniform sequence $\gamma_K = (\gamma, \dots, \gamma)$. An elementary operation $\delta \rightarrow \delta'$ consists in moving one unit from a slot where the capacity is over-filled, to the next slot. For any $\delta \neq \gamma_K$, there exist i, j such that

$$\delta_i > \gamma, \quad \delta_{i+1} = \dots = \delta_{j-1} = \gamma, \quad \delta_j < \gamma,$$

so that an elementary operation can be performed (note that indices are considered modulo K , so this definition is invariant under circular permutation). In other words, any $\delta \neq \gamma_K$ has at least one successor in the graph and γ_K is the only element with no successor. Then, $D(\delta)$ is the minimal number of elementary operations one has to perform to transform δ into γ_K (in the next subsection, we prove that $D(\delta) < +\infty$ for any δ). See Figures 2,3, and 4.

Our next main result, Theorem 4.3, states that the quantity $D(\delta)$ measures the total waiting time incurred by the players due to non-uniform departures. We need an additional piece of notation. For each integer p , the total allocation over the period $\{pK + 1, \dots, (p+1)K\}$, is $X_p(\sigma) = (X_p^e(\sigma))_e$, with

$$X_p^e(\sigma) = \sum_{t=pK+1}^{(p+1)K} x_t^e(\sigma).$$




Figure 2: Operations needed to transform $(6, 0, 0)$ into $(2, 2, 2)$.

Theorem 4.3. *Let $\delta \in \mathbb{N}_K(\gamma)$. Then*

(a) *The optimal latency $\text{Opt}(\mathcal{N}, K, \delta)$ exists and*

$$\text{Opt}(\mathcal{N}, K, \delta) = K \sum_e \gamma_e \tau_e + D(\delta).$$

A strategy profile $\sigma^{\text{Opt}} \in \mathcal{S}$ is optimal if and only if for each edge $e \in E$, and for each integer $p \in \mathbb{N}_+$, we have $X_p^e(\sigma^{\text{Opt}}) = \gamma_e K$.

(b) *The worst equilibrium latency $\text{Eq}(\mathcal{N}, K, \delta)$ exists and*

$$\text{Eq}(\mathcal{N}, K, \delta) = K \left(\sum_e \gamma_e \right) \max_e \tau_e + D(\delta).$$




Figure 3: Operations needed to transform $(3, 2, 1)$ into $(2, 2, 2)$.




Figure 4: Operations needed to transform $(3, 1, 2)$ into $(2, 2, 2)$.

There exists an integer P_0 such that in any equilibrium σ^{Eq} , for each edge $e \in E$, and each integer $p \geq P_0$, we have $X_p^e(\sigma^{\text{Eq}}) = \gamma_e K$.

This result gives a computation of optimal and equilibrium latencies. There are two main insights. First, as in the uniform case, optimal and equilibrium strategies are eventually identical. Namely in both cases, the average number of players that choose an edge over a period is equal to its capacity. Second, non-uniform departures create periodic queues. The length of these queues depends on the measure $D(\delta)$ of discrepancy between the distribution of departures and the uniform distribution. Remarkably, the effect of seasonality on optimal and equilibrium latencies is the same.

The formal proof is in the appendix. The idea is to prove the result by induction on the number $D(\delta)$. The main argument is that if δ' is obtained from δ by an elementary operation, then the departure sequence δ creates one more unit of average waiting time than δ' does. The reason is that under δ , there is a departure over capacity at some stage, and δ' is obtained from δ by letting a player postpone her departure by one unit of time, thus saving her one unit of waiting time.

4.2 On periodic queues

The main finding of Theorem 4.3 is that, to account for periodic departures, we have to compute the additional waiting time $D(\boldsymbol{\delta})$ induced by the sequence $\boldsymbol{\delta}$. This is defined as the minimal number of elementary operations one has to perform to transform the sequence $\boldsymbol{\delta}$ into the uniform sequence. In this section, we give a “dual” formula for $D(\boldsymbol{\delta})$.

We introduce some notation. A sequence $\boldsymbol{\delta} \in \mathbb{N}_K(c)$ is identified with a vector $\boldsymbol{\delta} = (\delta_1, \dots, \delta_K)$ such that $\sum_{m=1}^K \delta_m = \gamma K$. This can be thought of as a “distribution” over the integers $\{1, \dots, K\}$. The “cumulative distribution function” of this distribution is the following vector,

$$\text{Cdf}(\boldsymbol{\delta}) = (\delta_1, \delta_1 + \delta_2, \dots, \delta_1 + \dots + \delta_K), \quad (4.1)$$

whose m -th component is $\text{Cdf}_m(\boldsymbol{\delta}) = \sum_{k=1}^m \delta_k$. We denote $\mathbb{I}(\boldsymbol{\delta})$ the “integral” of the cumulative distribution function,

$$\mathbb{I}(\boldsymbol{\delta}) = \sum_{m=1}^K \sum_{k=1}^m \delta_k. \quad (4.2)$$

As a special case, for the uniform distribution $\boldsymbol{\gamma} = (\gamma, \dots, \gamma)$, we have $\text{Cdf}_k(\boldsymbol{\gamma}) = k\gamma$ and

$$\mathbb{I}(\boldsymbol{\gamma}) = \frac{K(K+1)}{2}\gamma.$$

Lastly, for each $m \in \{1, \dots, K\}$, we define the circular permutation π_m :

$$\boldsymbol{\delta} = (\delta_1, \dots, \delta_K) \mapsto \pi_m \boldsymbol{\delta} = (\delta_{m+1}, \dots, \delta_K, \delta_1, \dots, \delta_m).$$

The dual formula for $D(\boldsymbol{\delta})$ is given in the following proposition.

Proposition 4.4. *For each $\boldsymbol{\delta} \in \mathbb{N}_K(c)$,*

$$D(\boldsymbol{\delta}) = \max_{m=1, \dots, K} \mathbb{I}(\pi_m \boldsymbol{\delta}) - \mathbb{I}(\boldsymbol{\gamma}). \quad (4.3)$$

To prove the result we show that $D(\boldsymbol{\delta})$ coincides with the dual expression for $\boldsymbol{\delta} = \boldsymbol{\gamma}$ and that the dual

formula decreases by one along each elementary operation. To get an intuition, consider the following example.

Example 4.5. Consider $\gamma = (2, 2, 2)$, then, by (4.1), $\text{Cdf}(\gamma) = (2, 4, 6)$ and, by (4.2), $\mathbb{I}(\gamma) = 12$.

If $\delta = (6, 0, 0)$, then $\text{Cdf}(\delta) = (6, 6, 6)$ and $\mathbb{I}(\delta) = 18$.

If $\delta' = (3, 2, 1)$, then $\text{Cdf}(\delta') = (3, 5, 6)$ and $\mathbb{I}(\delta') = 14$.

Moreover $D(\delta) = 6$, since to go from $\delta = (6, 0, 0)$ to $\gamma = (2, 2, 2)$ two players have to be moved from position 1 to position 2 and two players have to be moved from 1 to 3. A similar argument shows that $D(\delta') = 2$, since one player has to be moved from 1 to 3. These results are in accordance with (4.3).

4.3 Price of anarchy

We derive a computation of the price of anarchy from the results obtained in the previous sections.

Corollary 4.6. *Let $\delta \in \mathbb{N}_K(\gamma)$. Then*

$$\text{Poa}(\mathcal{N}, K, \delta) = \frac{K \sum_e \gamma_e \tau_e + D(\delta)}{K (\sum_e \gamma_e) \max_e \tau_e + D(\delta)} \leq \frac{\sum_e \gamma_e \tau_e}{(\sum_e \gamma_e) \max_e \tau_e} = \text{Poa}(\mathcal{N}, K, \gamma).$$

It is clear that $\text{Poa}(\mathcal{N}, K, \delta)$ is a decreasing function of $D(\delta)$, that is, the price of anarchy decreases when departures are periodic. The intuition is that queues have to be created at stages where departures are above capacity. So at the optimum, the planner may want to create these queues on the edges with smallest transit cost, thereby choosing a strategy profile that resembles the equilibrium. The extreme case is $\delta = (K\gamma, 0, \dots, 0)$. Using equation 4.3,

$$D(\delta) = K^2\gamma - \frac{K(K+1)}{2}\gamma = \frac{K(K-1)}{2}\gamma$$

and $\text{Poa}(\mathcal{N}, K, \delta)$ is close to 1 when K is large. In that case, the equilibrium is almost optimal.

4.4 Below capacity

We assume now that the departures over a period are below the capacity of the network: $\sum_{k=1}^K \delta_k < \gamma K$. As in Subsection 3.3, consider the sub-network \mathcal{N}_δ with capacity $\delta = (1/K) \sum_{k=1}^K \delta_k$. When departures

are uniform over time, Proposition 3.3 states that it is enough to apply the results for the network \mathcal{N}_δ at capacity. This is no longer the case when departures are not uniform.

To see this, consider a network with two edges e_1, e_2 with $\tau_{e_1} = 1, \gamma_{e_1} = N, \tau_{e_2} = M, \gamma_{e_2} = 1$, that is, e_1 is short and wide, e_2 is long and narrow. Consider the K -periodic departure sequence $\delta = (KN, 0, \dots, 0)$. The average number of departures over a period is N which is less than the capacity of the network $N + 1$. The sub-network \mathcal{N}_δ is simply $\{e_1\}$ whose capacity is filled by these departures. On this sub-network, optimum and equilibrium latencies at capacity obviously coincide, and from Theorem 4.3, they are equal to

$$\text{Opt}(\mathcal{N}_\delta, K, \delta) = \text{Eq}(\mathcal{N}_\delta, K, \delta) = 1 + K(K - 1)/2$$

Now, if $M < K$, it is not optimal, nor an equilibrium, to use e_1 only. The first N players pay 1, the second batch of N players pay 2, \dots , the M -th batch of N -players pay M . Then, the next player pays $M + 1$ if he chooses e_1 and M if he chooses e_2 . We thus see that in any equilibrium, at least one player chooses e_2 and that the strategy profile where all players choose e_1 is not socially optimal: letting the last player switch to e_2 improves latencies for all players.

However, if $M \geq K$, then the optimal and equilibrium latencies are those of \mathcal{N}_δ at capacity.

This example shows that, for departures below capacity, the extra cost due to seasonality is more delicate to characterize, as it depends on relations between the departure process and the parameters of the network. We give a sufficient condition, under which it is enough to study the subnetwork \mathcal{N}_δ at capacity. Assume that all capacities are equal to one, let f be the edge such that the set of edges of \mathcal{N}_δ is $f^\preceq := \{e : e \preceq f\}$ and let f^+ be the next edge in \mathcal{N} .

Proposition 4.7. *If $\tau_{f^+} \geq \tau_f + K$, then equilibrium, optimal latencies and price of anarchy of $\Gamma(\mathcal{N}, K, \delta)$ are those of $\Gamma(\mathcal{N}_\delta, K, \delta)$.*

5 Extensions

We have considered dynamic congestion games with simple networks, simple latency functions, discrete time and atomic players. We have introduced a new dimension to congestion models: the non-stationarity

of the departure process, and characterized optimal and equilibrium latencies. An obvious direction of generalization is to consider more general networks. The formal model is easily adapted. The network is now

$$\mathcal{N} = (V, E, (\tau_e)_{e \in E}, (\gamma_e)_{e \in E})$$

where (V, E) is a directed graph with set of vertices V and set of edges E . The source and the destination are two special vertices. Each edge is endowed with a transit time and a capacity. The choice of a player is now a path in the graph from the source to the destination. The total latency suffered by a player is the sum of latencies over edges.

We choose as a convention, used in Koch and Skutella (2011), that queues are formed at the end of an edge. That is, all the players in a given generation who choose a given edge enter it at the same time, and progress along the edge at the same speed. The queue is formed at the exit, where priorities are used to define who exits first. With simple networks, forming queues at the entry would not change the results. However, with general networks, this convention is important.

Given a network $\mathcal{N} = (V, E, (\tau_e)_{e \in E}, (\gamma_e)_{e \in E})$, a *cut* is a subset of edges such that each path from the source to the destination goes through one of its edges. A cut is *minimal* if the sum of capacities of its edges is minimal among all cuts of the graph. We define the capacity of a network as the capacity of a minimal cut.

A general question is how to compute the price of anarchy for general networks and general departure process. Our main results answered this question for simple networks. In what follows, we provide a simple generalization and explain the difficulties of handling the two dimensions together.

5.1 Series of parallel networks

In this section we consider series of parallel networks and uniform departures.

For $h \in \{1, \dots, H\}$ let $\mathcal{N}^{(h)} = (E^{(h)}, (\tau_e)_{e \in E^{(h)}}, (\gamma_e)_{e \in E^{(h)}})$ be a parallel network and consider the network $\mathcal{N}_{\text{ser}}(H)$ obtained by concatenating $\mathcal{N}^{(1)}, \dots, \mathcal{N}^{(H)}$ in series (see Figure 5). We call this a series-of-parallel network.

Clearly, any subnetwork $\mathcal{N}^{(h)}$ is a cut of $\mathcal{N}_{\text{ser}}(H)$. Let $\mathcal{N}^{(*)}$ be a minimal cut of $\mathcal{N}_{\text{ser}}(H)$ and let $\gamma^{(*)}$




Figure 5: Series of parallel networks.

be the capacity of $\mathcal{N}^{(*)}$. Optimal and equilibrium values are characterized as follows.

Theorem 5.1. (a) *The optimal latency $\text{Opt}(\mathcal{N}_{\text{ser}}(H), \gamma^{(*)})$ exists and*

$$\text{Opt}(\mathcal{N}_{\text{ser}}(H), \gamma^{(*)}) = \sum_{h=1}^H \text{Opt}(\mathcal{N}^{(h)}, \gamma^{(*)}).$$

A strategy is optimal if and only if its restriction to $(\mathcal{N}^{(h)}, \gamma^{()})$ is optimal for $h \in \{1, \dots, H\}$.*

(b) *The worst equilibrium latency $\text{Eq}(\mathcal{N}_{\text{ser}}(H), \gamma^{(*)})$ exists and*

$$\text{Eq}(\mathcal{N}_{\text{ser}}(H), \gamma^{(*)}) = \sum_{h=1}^H \text{Eq}(\mathcal{N}^{(h)}, \gamma^{(*)}).$$

Any equilibrium strategy σ^{Eq} is such that its restriction to $(\mathcal{N}^{(h)}, \gamma^{()})$ is an equilibrium for $h \in \{1, \dots, H\}$.*

The insights are as follows. First, the capacity of the network is naturally defined as the capacity of its minimal cut. Second, due to the modular structure of the graph, each sub-network can be analyzed separately. As a result, optimal and equilibrium latencies are additive over the parallel modules.

It follows directly that the price of anarchy of this network is dictated by the sub-network with the highest price of anarchy.

Corollary 5.2. *Given a series-of-parallel network $\mathcal{N}_{\text{ser}}(H)$, we have*

$$\text{Poa}(\mathcal{N}_{\text{ser}}(H), \gamma^{(*)}) \leq \max_{h \in \{1, \dots, H\}} \text{Poa}(\mathcal{N}^{(h)}, \gamma^{(*)}).$$

5.2 General networks and Braess paradox

Consider a Wheatstone network as in Figure 6.




Figure 6: Wheatstone network.

The set of edges is $E = \{(sb), (sc), (bc), (bd), (cd)\}$ with the following transit costs and capacities

edge	τ	γ
(sb)	1	1
(sc)	3	2
(bc)	1	10
(bd)	3	2
(cd)	1	1

Notice that $\{(sb), (cd)\}$ is a minimal cut of the network and the capacity of this cut is 2. Consider the case of constant departure at capacity, i.e., $\delta = \gamma = 2$. The social optimum is achieved by a strategy σ^{Opt} such that

$$\sigma_{it}^{\text{Opt}}(h_{it}) = \begin{cases} (sb)(bd) & \text{for } i \text{ odd,} \\ (sc)(cd) & \text{for } i \text{ even.} \end{cases}$$

The optimal latency is therefore 8.

Consider now the following strategy¹ σ^{Eq} .

$$\sigma_{it}^{\text{Eq}}(h_{it}) = \begin{cases} (sb)(bc)(cd) & \text{for } [it] = [11], \\ (sb)(bc)(cd) & \text{for } [it] = [21], \\ (sb)(bc)(cd) & \text{for } [it] = [12], \\ (sc)(cd) & \text{for } [it] = [22], \\ (sb)(bd) & \text{for } [it] = [1t], \quad t \geq 3 \\ (sc)(cd) & \text{for } [it] = [2t], \quad t \geq 3. \end{cases}$$

The induced latencies are

$$\ell_{it}(\sigma^{\text{Eq}}) = \begin{cases} 3 & \text{for } [it] = [11], \\ 4 & \text{for } [it] = [21], \\ 4 & \text{for } [it] = [12], \\ 5 & \text{for } [it] = [22], \\ 5 & \text{for } [it] = [1t], \quad t \geq 3 \\ 5 & \text{for } [it] = [2t], \quad t \geq 3. \end{cases}$$

That is, the first players choose the zig-zag path $(sb)(bc)(cd)$ which is the fastest. This creates congestion to the point that the fourth player [22] prefers to take the right path $(sc)(cd)$. Thereafter, the zig-zag path

¹We define only the choices made by players at equilibrium. A complete description of the strategies would specify the choices for all histories.

is too congested and in each subsequent generation, one player takes the left path $(sb)(bd)$ and the other player takes the right path.

We find therefore that the equilibrium latency is 10. This example shows a significant difference between this case and all the results in the previous sections. Although optimal and equilibrium flows eventually coincide, it is no longer the case that all players eventually pay the maximal free-flow transit cost. As a matter of fact, all players end up paying even more than that. This is due to the fact that in the transient phase, they use the path $(sb)(bc)(cd)$ that they abandon later on. But their initial use of this path creates congestion on two different edges (sb) and (cd) that will be part of the steady state equilibrium flow.

We see that the topology of the network strongly affects the price of anarchy. We leave as an open problem to compute optimal and equilibrium values for general networks.

Remark 5.3. In static congestion games, a paradox due to Braess (1968, 2005) can arise: adding a new edge to a network may increase the worst equilibrium latency. A Braess-type paradox holds also in dynamic games. If we remove the edge (bc) in the network of Figure 6, then the equilibrium and the social optimum coincide. This appears even more paradoxical since the edge (bc) is not used in equilibrium in steady state. A version of Braess’s paradox for dynamic congestion games can be found in Macko, Larson, and Steskal (2010). Relatedly, an analysis of a dynamic Braess-type paradox in communication networks is provided by Xia and Hill (2013).

6 Appendix

Proof of Theorem 3.1

We start with a series of preparatory definitions and lemmata.

Definition 6.1. Two networks $\mathcal{N}, \mathcal{N}'$ are said to be *equivalent* if for each $K \in \mathbb{N}_+$ and for each $\boldsymbol{\delta} \in \mathbb{N}^K$, given the games $\Gamma(\mathcal{N}, K, \boldsymbol{\delta})$ and $\Gamma(\mathcal{N}', K, \boldsymbol{\delta})$, for each strategy profile σ of one game, there exists a strategy profile σ' of the other game such that each player in the first game pays the same cost as the corresponding player in the other game.

Lemma 6.2. *A simple network such that $\tau_{e_i} = \tau_{e_{i+1}}$ is equivalent to a network where the two edges e_i, e_{i+1} are merged into a single edge with capacity $\gamma_{e_i} + \gamma_{e_{i+1}}$.*

Proof. Given a network \mathcal{N} , let $\Gamma(\mathcal{N}, K, \delta)$ be the associated game. Let us assume that e_i, e_{i+1} have the same transit cost and consider the network \mathcal{N}' where the two edges are replaced by a single edge f with capacity $\gamma_{e_i} + \gamma_{e_{i+1}}$. By associating e_i and e_{i+1} with f , each strategy σ in $\Gamma(\mathcal{N}, K, \delta)$ is naturally mapped to a strategy σ' in $\Gamma(\mathcal{N}', K, \delta)$ with the same payoff profile. Conversely, each strategy σ' in $\Gamma(\mathcal{N}', K, \delta)$ can also be mapped to a strategy σ in $\Gamma(\mathcal{N}, K, \delta)$ with the same payoff profile, provided that players who choose f in $\Gamma(\mathcal{N}', K, \delta)$ at some stage, choose e_i or e_{i+1} in suitable proportions at the same stage. Namely, let N_f be the number of those players. Write $N_f = k(\gamma_{e_i} + \gamma_{e_{i+1}}) + r$ with $r < \gamma_{e_i} + \gamma_{e_{i+1}}$, and $r = r_{e_i} + r_{e_{i+1}}$ with $r_{e_i} \leq \gamma_{e_i}, r_{e_{i+1}} \leq \gamma_{e_{i+1}}$ (with one strict inequality). In σ , $k\gamma_{e_i} + r_{e_i}$ players choose e_i , and $k\gamma_{e_{i+1}} + r_{e_{i+1}}$ choose e_{i+1} . \square

As a consequence, whenever convenient, we may either assume that all transit costs are different, or that all capacities are 1.

Definition 6.3. For each edge e , denote $q_t^e(\sigma)$ the *length of the queue* on edge e at the end of stage t . This is the waiting cost paid by player $[1(t+1)]$ to enter edge e . We say that queues are *vanishing* if for each e

$$\limsup_T \frac{q_T^e(\sigma)}{T} = 0.$$

Lemma 6.4. *Let $\sigma \in \mathcal{S}$ be a strategy profile. If queues are not vanishing then*

$$\liminf_T \bar{\ell}_T(\sigma) = +\infty.$$

Proof. Assume that there exists e such that

$$\limsup_T \frac{q_T^e(\sigma)}{T} > 0.$$

Then, there exists $\alpha > 0$ and a subsequence $\{T_k\}$ such that along this subsequence we have $q_{T_k}^e(\sigma) \geq \alpha T_k$. This implies that there exists a player who has a waiting time of $w = \lfloor \alpha T_k \rfloor$. This in turn implies that for

each integer $s \leq w$, there exists a player who has waiting time s . Thus, the total waiting time adds up to at least

$$1 + \dots + w = \frac{w(w+1)}{2}$$

and the average waiting time at stage T_k is such that

$$\bar{w}_{T_k}(\sigma) \geq \frac{\lfloor \alpha T_k \rfloor (\lfloor \alpha T_k \rfloor + 1)}{2T_k}.$$

The r.h.s. diverges as $k \rightarrow \infty$, which concludes the proof. \square

Lemma 6.5. *Let $\sigma \in \mathcal{S}$ be a strategy profile. If queues are vanishing then*

$$\liminf_T \bar{\ell}_T(\sigma) \geq \sum_e \gamma_e \tau_e.$$

In the proof of Lemma 6.5 and of the following Claim 6.6, we use Lemma 6.2 and assume for simplicity that each edge e has capacity $\gamma_e = 1$, so that the number of edges equals the size of any generation and $\sum_e \gamma_e \tau_e = \sum_e \tau_e$.

Claim 6.6. *If $\limsup_T q_T^e(\sigma)/T = 0$, then $\limsup_T \bar{x}_T^e(\sigma) \leq 1$.*

Proof. Otherwise, there exists $\alpha > 0$ and a subsequence $\{T_k\}$ such that along this subsequence, $\bar{x}_{T_k}^e(\sigma) > 1 + \alpha$, that is,

$$\#\{[it] : \sigma_{it} = e\} \geq T_k + \alpha T_k.$$

This means that over the time interval $\{1, \dots, T_k\}$, $T_k + \alpha T_k$ players entered edge e . Since the capacity is one, at most T_k players could go out of e over this time interval, and therefore at least αT_k players are still in e . Now, these players could be at different distances from the exit, but at least $\alpha T_k / (\max_e \tau_e)$ of them are exactly at the same distance from the exit and will reach it at the same time. Thus, at some stage t with $T_k \leq t \leq T_k + \max_e \tau_e$, the queue is at least $\alpha T_k / (\max_e \tau_e)$. This implies

$$\frac{\limsup_T q_T^e(\sigma)}{T} \geq \frac{\alpha}{(\max_e \tau_e)} > 0,$$

contradicting the assumption. □

Proof of Lemma 6.5. Thanks to Claim 6.6, for any $\eta < 1$, there exists a time $T(\eta)$ such that for all $T \geq T(\eta)$ and for each edge e

$$\bar{x}_T^e(\sigma) \leq 1 + \eta. \tag{6.1}$$

Since $\sum_e \bar{x}_T^e(\sigma) = \gamma$, for each edge e , we have $\bar{x}_T^e(\sigma) = \gamma - \sum_{f \neq e} \bar{x}_T^f(\sigma)$. Inequality (6.1) then implies

$$\bar{x}_T^e(\sigma) \geq \gamma - (\gamma - 1)(1 + \eta) = 1 - (\gamma - 1)\eta.$$

As a consequence,

$$\bar{\ell}_T(\sigma) = \sum_e \bar{x}_T^e(\sigma)\tau_e + \bar{w}_T(\sigma) \geq \sum_e (1 - (\gamma - 1)\eta)\tau_e = \sum_e \tau_e - \eta(\gamma - 1) \sum_e \tau_e.$$

Setting $\eta' = \eta(\gamma - 1) \sum_e \tau_e$, we find that for all $\eta' > 0$, there exists T' such that for all $T \geq T'$

$$\bar{\ell}_T(\sigma) \geq \sum_e \tau_e - \eta'.$$

That is,

$$\liminf_T \bar{\ell}_T(\sigma) \geq \sum_e \tau_e,$$

as desired. □

Lemma 6.7. *Let $\sigma \in \mathcal{S}$ be a strategy profile. If at some stage t , there exists an edge e such that $x_t^e(\sigma) > \gamma_e$, then for all $t' \geq t$, there exists e' with $q_{t'}^{e'}(\sigma) > 0$.*

Proof. Assume that at stage t we have $x_t^e(\sigma) > \gamma_e$. This creates a queue $q_t^e(\sigma) > 0$. At the next stage, either $x_{t+1}^e(\sigma) \geq \gamma_e$, in which case

$$q_{t+1}^e(\sigma) = q_t^e(\sigma) - \gamma_e + x_{t+1}^e(\sigma) \geq q_t^e(\sigma) > 0,$$

or $x_{t+1}^e(\sigma) < \gamma_e$, but then there exists e' such that $x_{t+1}^{e'}(\sigma) > \gamma_{e'}$ and thus $q_{t+1}^{e'}(\sigma) > 0$. □

Proof of Theorem 3.1(a). Since the size of a generation equals the capacity of the network, a strategy σ that sends γ_e players on each edge e at each stage, satisfies $\limsup_T \bar{\ell}_T(\sigma) = \sum_e \gamma_e \tau_e < +\infty$. We prove now that such a strategy is optimal.

From Lemma 6.4, an optimal strategy must induce queues with negligible lengths and Lemma 6.5 shows that the optimal latency is larger or equal than $\sum_e \gamma_e \tau_e$. Thus, a strategy σ is optimal if $x_t^e(\sigma) = \gamma_e$ for all $t \in \mathbb{N}_+$ and for all $e \in E$. It remains to prove that these are the only optimal strategies.

Consider t_1 the first stage where there exists e such that $x_{t_1}^e(\sigma) > \gamma_e$. At this moment, a queue $q_1 \geq 1$ is created and, from Lemma 6.7, this queue will not disappear. It follows that $\liminf_T \bar{w}_T(\sigma) \geq q_1$. Now, at $t_1 + 1$, we are facing the same optimization problem except that the objective function is translated upwards by some constant $w_1 \geq 1$. From Lemma 6.5, the optimal latency is $\sum_e \gamma_e \tau_e + w_1$. Therefore, for any optimal strategy, there is no such stage t_1 . This ends the proof of Theorem 3.1(a). \square

Proof of Theorem 3.1(b). In this proof, it is convenient to assume $\tau_1 < \dots < \tau_n$, which we can assume without loss of generality from Lemma 6.2. We proceed by induction on the number n of edges. The result being obvious for $n = 1$, consider $n = 2$.

Lemma 6.8. *In any equilibrium, the allocation matches the capacities from some stage on.*

Proof. Take a given equilibrium.

Claim 6.9. *There exists a first player who is indifferent between the two edges.*

Proof. The number of departures per generation is $\gamma > \gamma_{e_1}$. In any equilibrium, the first players choose e_1 . The latency of a player is given by the expression,

$$f(x) = \tau_1 + \left\lfloor \frac{(x - \gamma_{e_1})^+}{\gamma_{e_1}} \right\rfloor,$$

with x the number of people in the queue at the same time as the player, and ahead of her. This latency function is weakly increasing, and increases by 1 on each multiple of γ_{e_1} . Since $\gamma > \gamma_{e_1}$, there must exist a first stage T_0 where some player $[iT_0]$ would pay τ_2 if she chose e_1 . Since the number of departures per generation is $\gamma = \gamma_{e_1} + \gamma_{e_2}$, this stage T_0 is the smallest t such that $t(\gamma - \gamma_{e_1}) + \tau_1 \geq \tau_2$, and this for any

equilibrium. Indeed, before that stage, all players choose e_1 . Since only γ_{e_1} players per stage can exit, $t(\gamma - \gamma_{e_1})$ players remain in the queue, yielding a cost of $\tau_1 + t(\gamma - \gamma_{e_1})$ which eventually reaches τ_2 .

As a consequence, player $[iT_0]$ is the first player to be indifferent between choosing e_1 or e_2 . \square

Now, we consider two cases.

Case 1. Assume first that $[iT_0]$ or one her followers $[jT_0]$ in her generation chooses e_1 , that is, there exists $j \geq i$ such that $[jT_0]$ chooses e_1 . Then, at the next generation, exactly γ_{e_1} players choose e_1 and γ_{e_2} players choose e_2 . Indeed, the first players to choose e_1 pays $\tau_2 + 1 - 1$. So if more than γ_{e_1} players choose e_1 , the last one pays $\tau_2 + 1$ and could have paid τ_2 by choosing e_2 . This scenario repeats at the next generation since queues are stationary from that point on.

Case 2. The second possible case is that all players $[jT_0]$, with $j \geq i$, choose e_2 . This can happen within an equilibrium only if the number of those players is no more than τ_2 (otherwise, we are in the previous case). Then, the situation faced by the next generation G_{T_0+1} is the same, that is, players $[j(T_0 + 1)]$ choose e_1 for $j < i$, and player $[i(T_0 + 1)]$ is indifferent. Now, either $i - 1 = \gamma_{e_1}$ and there is an equilibrium where in each generation, players $[jT]$, $j < i$, choose e_1 , and the other players choose e_2 ; or, in some generation, there is a player $[jT]$ with $j \geq i$ who chooses e_1 and we are in Case 1 from that stage on.

This concludes the proof of Lemma 6.8 for $n = 2$. \square

Now we argue that $\gamma \max_e \tau_e$ is the worst equilibrium latency. Remark that the multiplicity of equilibria stems from the indifference conditions, as seen in the previous proof. Consider the equilibrium where each player, when indifferent between several edges, chooses the one with least transit cost, i.e., the first edge in the order \prec . It follows from the previous proof that all player pay $\max_e \tau_e$ from some stage on. Remark also that in all equilibria considered above, no player pays more than $\max_e \tau_e$.

Assume now the result true for $n - 1$ and consider a network with n edges. Until the latencies have not reached τ_n , players use only the $n - 1$ first edges. By the induction hypothesis, if the number of players is $\sum_{e \prec e_n} \tau_e$, then, from some point on, the allocation of players on edges matches the capacities. Since the

number of departures is $\gamma > \sum_{e \prec e_n} \tau_e$, then the latencies of all edges $e \prec e_n$ eventually become equal to τ_{n-1} at some stage. Subsequent players may regard the edges e_1, \dots, e_{n-1} as a single edge with latency τ_{n-1} . Thus, we are back to the two-edge case and may apply the previous argument. \square

Proof of Theorem 4.3

Claim 6.10. *If $\liminf_P \bar{L}_P(\sigma) < +\infty$, then for each $e \in E$, $\limsup_P \bar{X}_P^e(\sigma) \leq \gamma_e$.*

Proof. The proof is as in Lemma 6.4 and Claim 6.6, treating each period as one stage. As before, assume for simplicity that all capacities are one. If $\limsup_P \bar{X}_P^e(\sigma) > 1$, then on a subsequence of stages, say T_k , the average number of people on e exceeds the capacity by some $\alpha > 0$. This implies that at least αT_k people are queuing which results in an unbounded average waiting time. \square

Lemma 6.11. *For any strategy profile $\sigma \in \mathcal{S}$,*

$$\liminf_P \bar{L}_P(\sigma) \geq \sum_e \gamma_e \tau_e + D(\delta),$$

and there exists a strategy $\sigma \in \mathcal{S}$ such that

$$\limsup_P \bar{L}_P(\sigma) \leq \sum_e \gamma_e \tau_e + D(\delta),$$

which is therefore optimal.

Proof. We consider only strategies $\sigma \in \mathcal{S}$ such that $\liminf_P \bar{L}_P(\sigma) < +\infty$. From Claim 6.10, and as in the proof of Lemma 6.5, for any $\eta > 0$, there exists $P(\eta)$, such that for all $P \geq P(\eta)$, and each edge $e \in E$,

$$|\bar{X}_P^e(\sigma) - K\gamma_e| \leq \eta.$$

As a consequence, for $P \geq P(\eta)$, the average latency is such that

$$\bar{L}_P(\sigma) \geq \sum_e \bar{X}_P^e(\sigma) \tau_e \geq K \sum_e \gamma_e \tau_e + \eta \left(\sum_e \tau_e \right).$$

Thus, the optimal latency is at least $K \sum_e \gamma_e \tau_e$. Consider the relaxed optimization problem where the players are allowed to choose their departure time in the system. There, we could force players to depart uniformly over time, and therefore, the optimal latency of the relaxed problem is $K \sum_e \gamma_e \tau_e$.

Now, we proceed by induction on the integer $D(\boldsymbol{\delta})$. If $D(\boldsymbol{\delta}) = 0$ then $\boldsymbol{\delta} = \boldsymbol{\gamma}$, and the result follows from Theorem 3.1. Assume $D(\boldsymbol{\delta}) = 1$, then the sequence of departures is either,

$$(\gamma, \dots, \gamma, \gamma + 1, \gamma - 1, \gamma, \dots, \gamma)$$

or

$$(\gamma - 1, \gamma, \dots, \gamma, \gamma + 1).$$

These two sequences are similar, up to circular permutations. Yet, since queues might be created at the first stages, the starting point of the sequence is relevant. However, whenever departures are no more than the capacity, it is possible to pay a total cost of at most $\sum_e \gamma_e \tau_e$ without creating queues. The impact of these first stages disappears in the long run. It is thus without loss of generality to assume that the sequence of departures is

$$(\gamma + 1, \gamma - 1, \gamma, \dots, \gamma).$$

Consider $[\delta_1 1]$, the last player in the first generation. In the relaxed problem, we may postpone by one stage the departure time of player $[\delta_1 1]$, and be back to the uniform departure case where the optimal latency is $K \sum_e \gamma_e \tau_e$. Now, at the first stage, since $\gamma + 1$ players depart, there exists at least one player that must wait for at least one period. This situation repeats at each period and therefore the long run average latency for any strategy is at least $K \sum_e \gamma_e \tau_e + 1$. Now, it is possible to achieve exactly this amount. Consider the optimal solution of the relaxed problem where player $[\delta_1 1]$ postpones her departure by one stage and becomes the first player in the second generation. An optimal strategy consists simply in playing this strategy. Namely, the first γ players before $[\delta_1 1]$ fill the capacities of the network. Then, player $[\delta_1 1]$ chooses, say, the first edge and waits for one period. The second generation fills the capacities of the network, considering that one slot has been already taken by $[\delta_1 1]$. Subsequent generations with size γ fill capacities. It is then clear that the only waiting cost is 1 and is paid by $[\delta_1 1]$. Observe that the

choice of player $[\delta_1 1]$ is immaterial. In any optimal strategy, the players in the second generation must fill the capacities of the network, taking into account the choice of $[\delta_1 1]$.

Now we consider the induction step. Assume the result to be true for any sequence δ' with $D(\delta') < n$ and consider δ, δ' such that $\delta \rightarrow \delta'$, $D(\delta) = n$ and $D(\delta') = n - 1$. We want to prove that the long run average latency is at least $K \sum_e \gamma_e \tau_e + D(\delta') + 1$ and that this amount can actually be achieved. Since $\delta \rightarrow \delta'$, either there exist $i < j - 1$, $\delta'_i \geq \gamma$, $\delta'_j < \gamma$ such that δ and δ' can be written as follows:

$$\begin{aligned}\delta' &= (\dots, \delta'_i, \gamma + 1, \gamma, \dots, \gamma, \delta'_j, \dots) \\ \delta &= (\dots, \delta'_i + 1, \gamma, \gamma, \dots, \gamma, \delta'_j, \dots)\end{aligned}$$

or there exist i , $\delta'_i \geq \gamma$, $\delta'_{i+1} < \gamma + 1$ such that δ and δ' can be written as follows:

$$\begin{aligned}\delta' &= (\dots, \delta'_i, \delta'_{i+1}, \dots) \\ \delta &= (\dots, \delta'_i + 1, \delta'_{i+1} - 1, \dots)\end{aligned}$$

Consider the first case. If $t^* = i \pmod K$, under δ , players depart above capacity. There is thus at least one player $[jt^*]$ who has to wait for at least one stage. Consider the relaxed problem were we are allowed to postpone the departure of this player by one stage, that is, to transform δ into δ' . At stage $t^* + 1$, players still depart above capacity so the postponed player $[jt^*]$ still induces a waiting time of at least one. But without delaying the departure of $[jt^*]$, this waiting time is paid two times, at stages t^* and $t^* + 1$. This shows that the long run average latency is at least $K \sum_e \gamma_e \tau_e + D(\delta') + 1$. This lower bound is achievable by letting player $[jt^*]$ wait for one stage and then choose the strategy she would have in the relaxed problem. In the second case, postponing the departure of $[jt^*]$ also saves one unit of waiting cost since at stage $t^* + 1$, players depart below capacity. As before, it can be remarked that the choice of $[jt^*]$ is not important, as long as subsequent players adjust for his choice so as to respect the capacity constraint over the period. This concludes the induction proof. \square

Now, we characterize optimal strategies.

Lemma 6.12. *If a strategy is optimal, then for each edge e , the number of players that choose that edge over a period of K consecutive stages is $K\gamma_e$.*

We start by the following observation.

Claim 6.13. For each $\boldsymbol{\delta} \in \mathbb{N}_K(c)$, for each $m \in \{1, \dots, K\}$,

$$\mathbb{I}(\boldsymbol{\delta}) - \mathbb{I}(\pi_m \boldsymbol{\delta}) = K(\text{Cdf}_m(\boldsymbol{\delta}) - \text{Cdf}_m(\boldsymbol{\gamma})).$$

Proof. We start from $\mathbb{I}(\boldsymbol{\delta}) = \sum_{k=1}^K \sum_{j=1}^k \delta_j$ and rewrite it as,

$$\mathbb{I}(\boldsymbol{\delta}) = K\delta_1 + (K-1)\delta_2 + \dots + (K-m+1)\delta_m + (K-(m+1)+1)\delta_{m+1} + \dots + \delta_K.$$

Similarly,

$$\begin{aligned} \mathbb{I}(\pi_m \boldsymbol{\delta}) &= K\delta_{m+1} + (K-1)\delta_{m+2} + \dots \\ &\quad + (K - (K-m) + 1)\delta_K + (K - (K-m+1) + 1)\delta_1 + \dots + \delta_m. \end{aligned} \quad (6.2)$$

It follows directly that,

$$\mathbb{I}(\boldsymbol{\delta}) - \mathbb{I}(\pi_m \boldsymbol{\delta}) = \sum_{k=1}^m (K-m)\delta_k - \sum_{k=m+1}^K m\delta_k.$$

Replacing $\sum_{k=m+1}^K \delta_k$ by $\gamma K - \sum_{k=1}^m \delta_k$, we get,

$$\mathbb{I}(\boldsymbol{\delta}) - \mathbb{I}(\pi_m \boldsymbol{\delta}) = (K-m)\text{Cdf}_m(\boldsymbol{\delta}) - m(K\gamma - \text{Cdf}_m(\boldsymbol{\delta})) = K(\text{Cdf}_m(\boldsymbol{\delta}) - \text{Cdf}_m(\boldsymbol{\gamma})),$$

as desired. □

We get the direct corollary:

Corollary 6.14.

$$\mathbb{I}(\boldsymbol{\delta}) = \max_{m=1, \dots, K} \mathbb{I}(\pi_m \boldsymbol{\delta}) \text{ iff, for all } m \in \{1, \dots, K\}, \text{Cdf}_m(\boldsymbol{\delta}) \geq \text{Cdf}_m(\boldsymbol{\gamma}).$$

The consequence is that we can choose the starting point of the cycle in such a way that, at the optimum, no queues are left at the end of each cycle.

Claim 6.15. Let $\delta = (\delta_1, \dots, \delta_K)$ be the sequence of departures. Let k be such that $\mathbb{I}(\pi_k \delta) = \max_m \mathbb{I}(\pi_m \delta)$. Then the strategy in the initial phase $\{1, \dots, k\}$ can be chosen such that no queues remain at the beginning of stage $k + 1$.

Proof. The sequence of departures is as follows:

$$(\delta_1, \dots, \delta_k, \delta_{k+1}, \dots, \delta_K, \delta_1, \dots, \delta_k, \dots)$$

which we decompose as an initial phase $(\delta_1, \dots, \delta_k)$ and a period

$$(\delta_{k+1}, \dots, \delta_K, \delta_1, \dots, \delta_k).$$

From the choice of k , for each $j = 1, \dots, k$,

$$\sum_{m=k+1}^{K+j-1} \delta_m \geq (K - k + j + 1)\gamma,$$

and thus $\sum_{m=j}^k \delta_m \leq (k - j + 1)\gamma$. That is,

$$\delta_k \leq \gamma, \delta_k + \delta_{k-1} \leq 2\gamma, \dots, \delta_1 + \dots + \delta_k \leq k\gamma.$$

If for all $j = 1, \dots, k$, $\delta_j \leq \gamma$, then we can send players at capacity on each edge and no queues are formed. Otherwise, there exists $j < k$ such that $\delta_j > \gamma$. But then, $\delta_j + \dots + \delta_k \leq (k - j + 1)\gamma$ and $\delta_{j+1} + \dots + \delta_k \leq (k - j)\gamma$. It follows that

$$0 < \delta_j - \gamma \leq (k - j)\gamma - (\delta_{j+1} + \dots + \delta_k),$$

which means that the excess of departures $\delta_j - \gamma > 0$ at stage j is compensated by the deficit of departures at stages $\{j + 1, \dots, k\}$. It follows that all queues formed at some stage $j < k$, can be undone before stage k . □

Claim 6.16. *For any optimal strategy, for each edge $e \in E$ and period $p \in \mathbb{N}_+$, if $X_p^e(\sigma) = K\gamma_e$, then the total latency over each period is*

$$L_p(\sigma) = K \sum_e \gamma_e \tau_e + D(\boldsymbol{\delta})$$

and no queues are left at the end of each period.

Proof. First argue that such a strategy exists. The proof of this fact is similar to that of Lemma 6.11 and proceeds by induction on $D(\boldsymbol{\delta})$. By Theorem 3.1, the result holds for $D(\boldsymbol{\delta}) = 0$. If $D(\boldsymbol{\delta}) = 1$, then $\boldsymbol{\delta} = (\gamma, \dots, \gamma, \gamma+1, \gamma-1, \gamma, \dots, \gamma)$ and, without loss of generality, we can choose $\boldsymbol{\delta} = (\gamma+1, \gamma-1, \gamma, \dots, \gamma)$. The strategy is deduced from an optimal strategy in the uniform case. The last player in G_1 waits one stage and plays the strategy she would have played, if her departure had been postponed by one stage. For the induction step, assume that $\boldsymbol{\delta} \rightarrow \boldsymbol{\delta}'$ and that the result is true for $\boldsymbol{\delta}'$. We write,

$$\begin{aligned} \boldsymbol{\delta}' &= (\dots, \delta'_i, \gamma+1, \gamma, \dots, \gamma, \delta'_j, \dots) \\ \boldsymbol{\delta} &= (\dots, \delta'_i + 1, \gamma, \gamma, \dots, \gamma, \delta'_j, \dots) \end{aligned}$$

with $\delta'_i \geq \gamma$, $\delta'_j < \gamma$. We deduce an optimal strategy in $\Gamma(\mathcal{N}, K, \boldsymbol{\delta})$ from an optimal strategy in $\Gamma(\mathcal{N}, K, \boldsymbol{\delta}')$ by letting the last player in the i -th generation wait for one stage and then play the strategy she would have played, if her departure had been postponed by one stage. As in the proof of Lemma 6.11, the total latency is $K \sum_e \gamma_e \tau_e + D(\boldsymbol{\delta}') + 1$. If no queues were left at the end of the period under $\boldsymbol{\delta}'$, then, with this strategy, no queues are left under $\boldsymbol{\delta}$, either. We have thus proved that there exists a strategy as in the statement of the claim. \square

Proof of Lemma 6.12. Thanks to Claim 6.15, we assume from now on that $k+1 = 1$ modulo K , or equivalently, $\text{Cdf}(\boldsymbol{\delta}) \geq \text{Cdf}(\boldsymbol{\gamma})$ and we use Claim 6.16 to prove that any strategy that respects the capacity of edges is optimal.

Finally, suppose that a strategy leaves a queue at the end of a period. This means that next period, the network has to accommodate at least $K\gamma + 1$ players which is not possible, the total capacity over K stages being $K\gamma$. Thus, there is an extra permanent queue which results in a long run average latency of at least $K \sum_e \gamma_e \tau_e + D(\boldsymbol{\delta}) + 1$, which precludes the strategy from being optimal. \square

Proof of Theorem 4.3(a). The first part of the proof is similar to that of Theorem 3.1.

Lemma 6.11 generalizes Lemma 6.5, taking into account the queues created by non-uniform departures. It shows that the optimal latency exists and is $K \sum_e \gamma_e \tau_e + D(\boldsymbol{\delta})$. Finally Lemma 6.12 shows that for any optimal strategy, for each edge e , the number of players that choose that edge over a period of K consecutive stages is $K\gamma_e$. \square

Proof of Theorem 4.3(b). We also proceed by induction on $D(\boldsymbol{\delta})$. If $D(\boldsymbol{\delta}) = 0$ then $\boldsymbol{\delta} = \boldsymbol{\gamma}$, and the result follows from Theorem 3.1. Recall that in that case, there exists a stage T_0 such that, for each $t \geq T_0$, in each equilibrium σ , $x_t^e(\sigma) = \gamma_e$ for each edge e . This stage T_0 falls into one period of the sequence of departures and let P_0 be the next period, that is, the integer such that $(P_0 - 1)K + 1 \leq T_0 \leq P_0K$. Clearly, for each integer $p \geq P_0$ and each edge $e \in E$, we have $X_p^e(\sigma^{\text{Eq}}) = \gamma_e K$.

Let us do the induction step. We take two sequences $\boldsymbol{\delta}, \boldsymbol{\delta}'$ such that $\boldsymbol{\delta} \rightarrow \boldsymbol{\delta}'$, assume the result to be true for $\boldsymbol{\delta}'$, and prove it for $\boldsymbol{\delta}$. We want to prove that the long run average latency is at least $K \sum_e \gamma_e \tau_e + D(\boldsymbol{\delta}') + 1$ and that this amount can actually be achieved. Since $\boldsymbol{\delta} \rightarrow \boldsymbol{\delta}'$, either there exist $i < j - 1$, $\delta'_i \geq \gamma$, $\delta'_j < \gamma$ such that $\boldsymbol{\delta}$ and $\boldsymbol{\delta}'$ can be written as follows:

$$\begin{aligned} \boldsymbol{\delta}' &= (\dots, \delta'_i, \gamma + 1, \gamma, \dots, \gamma, \delta'_j, \dots) \\ \boldsymbol{\delta} &= (\dots, \delta'_i + 1, \gamma, \gamma, \dots, \gamma, \delta'_j, \dots) \end{aligned}$$

or there exist i , $\delta'_i \geq \gamma$, $\delta'_{i+1} < \gamma + 1$ such that $\boldsymbol{\delta}$ and $\boldsymbol{\delta}'$ can be written as follows:

$$\begin{aligned} \boldsymbol{\delta}' &= (\dots, \delta'_i, \delta'_{i+1}, \dots) \\ \boldsymbol{\delta} &= (\dots, \delta'_i + 1, \delta'_{i+1} - 1, \dots) \end{aligned}$$

We only treat the first case, the second case being similar. Denote $t^* \in \{1, \dots, K\}$ the first stage where $\delta'_i + 1$ players depart and let $t_p^* = t^* + pK$ for each integer p . Then $[(\gamma + 1)t_p^*]$ is the last player in the generation $G_{t_p^*}$. Consider the game Γ^* where the departure of each player $[(\gamma + 1)t_p^*]$ is postponed by one stage so that $[(\gamma + 1)t_p^*]$ is the first player in generation $G_{t_{p+1}^*}$. This game is identical to $\Gamma(\mathcal{N}, K, \boldsymbol{\delta}')$. We claim that this auxiliary game Γ^* and the actual game $\Gamma(\mathcal{N}, K, \boldsymbol{\delta})$ have the same equilibria. Clearly, for all players that are ahead of $[(\gamma + 1)t_1^*]$, the two games are identical, so their equilibrium strategies are

the same. Then in $\Gamma(\mathcal{N}, K, \boldsymbol{\delta})$, player $[(\gamma + 1)t_1^*]$ can at least play her equilibrium strategy, say e , in Γ^* . Therefore, her best-response in the actual game $\Gamma(\mathcal{N}, K, \boldsymbol{\delta})$ is either e , or $e' \prec e$. If it is $e' \prec e$, it would also be a best-reply in Γ^* since $\tau_{e'} < \tau_e$ and the queues that $[(\gamma + 1)t_1^*]$ faces in Γ^* cannot be larger than those she faces in $\Gamma(\mathcal{N}, K, \boldsymbol{\delta})$. Thus, the best-reply of $[(\gamma + 1)t_1^*]$ is the same in both games. Finally, for all players that come between $[(\gamma + 1)t_1^*]$ and $[(\gamma + 1)t_2^*]$, the two games are identical and thus they have the same best-replies. By a simple induction, this holds for any $[(\gamma + 1)t_p^*]$.

As a consequence, from the induction hypothesis, the number of player who choose edge e over a period is $K\gamma_e$, and the worst equilibrium latency is $\text{Eq}(\mathcal{N}, K, \boldsymbol{\delta}') + 1$. We also get by induction that for each integer $p \geq P_0$ and each edge $e \in E$, we have $X_p^e(\sigma^{\text{Eq}}) = \gamma_e K$. \square

Proof of Proposition 4.4

The proof builds on Corollary 6.14 from which we can derive the following findings. For any distribution $\boldsymbol{\delta}$, there exists a circular permutation π_m such that $\text{Cdf}_m(\pi_m \boldsymbol{\delta}) \geq \text{Cdf}_m(\boldsymbol{\gamma})$. This happens if and only if $\pi_m \boldsymbol{\delta}$ has the largest integral among all circular permutations. Since, from its definition, $D(\boldsymbol{\delta})$ is invariant with respect to circular permutations, we assume from now on that $\mathbb{I}(\boldsymbol{\delta}) = \max_{m=1, \dots, K} \mathbb{I}(\pi_m \boldsymbol{\delta})$. The next lemma establishes the link between $D(\boldsymbol{\delta})$ and $\mathbb{I}(\boldsymbol{\delta})$.

Lemma 6.17. *Let $\boldsymbol{\delta} \in \mathbb{N}_K(\gamma)$ such that $\text{Cdf}_m(\boldsymbol{\delta}) \geq \text{Cdf}_m(\boldsymbol{\gamma})$ for all $m \in \{1, \dots, K\}$ and let $\boldsymbol{\delta} \rightarrow \boldsymbol{\delta}'$. Then $\text{Cdf}_m(\boldsymbol{\delta}') \geq \text{Cdf}_m(\boldsymbol{\gamma})$ for all $m \in \{1, \dots, K\}$ and $\mathbb{I}(\boldsymbol{\delta}') = \mathbb{I}(\boldsymbol{\delta}) - 1$.*

Proof. Let $\boldsymbol{\delta} \in \mathbb{N}_K(\gamma)$ be such that $\text{Cdf}_m(\boldsymbol{\delta}) \geq m\gamma$ for all $m \in \{1, \dots, K\}$. Take $i < j$ such that $\delta_i > \gamma$, $\delta_l = \gamma$ for all $l \in \{i, \dots, j\}$, and $\delta_j < \gamma$. Necessarily, $1 \leq i < j \leq K$, that is, i and j are within the same period. In other words, it is not possible to have the following situation:

$$(\gamma, \dots, \gamma, \delta_j, \dots, \delta_i, \gamma, \dots, \gamma)$$

for if it were the case, denoting $M + 1$ the rank of δ_i in this sequence, we would have $\sum_{m=M+1}^K \delta_m > (K - M)\gamma$. Thus, $\sum_{m=1}^M \delta_m < M\gamma$, a contradiction.

Define then $\boldsymbol{\delta}'$ by letting $\delta'_i = \delta_i - 1$, $\delta'_{i+1} = \delta_{i+1} + 1$, all other components remaining unchanged. Since

$\sum_{m=1}^{i-1} \delta_m \geq (i-1)\gamma$ and $\delta_i > \gamma$, we have $\sum_{m=1}^i \delta'_m = \sum_{m=1}^i \delta_m - 1 \geq i\gamma$, and thus $\text{Cdf}_m(\boldsymbol{\delta}') \geq m\gamma$ for all $m \in \{1, \dots, K\}$. Now,

$$\begin{aligned} \mathbb{I}(\boldsymbol{\delta}') &= \sum_{m=1}^K \sum_{k=1}^m \delta'_k \\ &= \sum_{k=1}^K (K - k + 1) \delta'_k \\ &= \mathbb{I}(\boldsymbol{\delta}) - (K - i + 1) + (K - (i + 1) + 1) \\ &= \mathbb{I}(\boldsymbol{\delta}) - 1, \end{aligned}$$

as claimed. □

Proof of Proposition 4.4. We may now conclude the proof of Proposition 4.4. Recall that $D(\boldsymbol{\delta})$ is the minimal number of elementary operations needed for transforming $\boldsymbol{\delta}$ into $\boldsymbol{\gamma}$. For $\boldsymbol{\delta}$ such that $\text{Cdf}(\boldsymbol{\delta}) \geq \text{Cdf}(\boldsymbol{\gamma})$, we have $\boldsymbol{\delta} \neq \boldsymbol{\gamma}$ if and only if $\mathbb{I}(\boldsymbol{\delta}) > \mathbb{I}(\boldsymbol{\gamma})$. So, as long as $\text{Cdf}(\boldsymbol{\delta}) \geq \text{Cdf}(\boldsymbol{\gamma})$ and $\boldsymbol{\delta} \neq \boldsymbol{\gamma}$, an elementary operation is possible, and the result $\boldsymbol{\delta}'$ is such that $\text{Cdf}(\boldsymbol{\delta}') \geq \text{Cdf}(\boldsymbol{\gamma})$ and $\mathbb{I}(\boldsymbol{\delta}') = \mathbb{I}(\boldsymbol{\delta}) - 1$. Therefore along each chain of elementary operations

$$\boldsymbol{\delta} \rightarrow \boldsymbol{\delta}' \rightarrow \dots \rightarrow \boldsymbol{\delta}^{(n)} \rightarrow \dots$$

the integral decreases by one until $\boldsymbol{\gamma}$ is reached. As a consequence, $\boldsymbol{\gamma}$ is reached eventually and along each chain $\boldsymbol{\delta} \rightarrow \dots \rightarrow \boldsymbol{\gamma}$, the number of operations is $\mathbb{I}(\boldsymbol{\delta}) - \mathbb{I}(\boldsymbol{\gamma})$. □

Proof of Proposition 4.7

Proof of Proposition 4.7. Consider first optimal strategies. We argue that an optimal strategy in $\Gamma(\mathcal{N}_\delta, K, \boldsymbol{\delta})$ is also optimal in $\Gamma(\mathcal{N}, K, \boldsymbol{\delta})$. Such a strategy sends K players per period on each edge (we assume that all capacities are one). It follows that the length of the queue on each edge is at most K at each point in time and that each player pays at most $\tau_{f+} + K$. Thus, the total latency cannot be improved by transferring players from edges in \mathcal{N}_δ to edges outside of the sub-network. To justify it, consider two cases. First,

consider a player who is paying $\tau_{f^+} + K$ in \mathcal{N}_δ , in which case, he is last in the queue. Transferring this player to f^+ does not improve her own cost, nor the latencies of other players. Second, consider a player who is paying less than $\tau_{f^+} + K$ in \mathcal{N}_δ . Transferring this player to f^+ increases her own cost by at least one, whereas it only decreases queues by at most one, so this does not improve the total latency.

Similarly, we argue that an equilibrium in $\Gamma(\mathcal{N}_\delta, K, \delta)$ is an equilibrium in $\Gamma(\mathcal{N}, K, \delta)$. Since the number of departures per period is $K\delta$, at each point in time, there exists an edge such that the average number number of players who chose that edge in the past is at most K , and thus the length of the queue is at most K . It follows that each player can choose an edge in \mathcal{N}_δ that yields a cost of at most $\tau_{f^+} + K$. Therefore, there exists always a best-reply in the sub-network. \square

Proof of Theorem 5.1

Proof of Theorem 5.1. (a) It is clear that any subnetwork $\mathcal{N}^{(h)}$ is a cut of $\mathcal{N}_{\text{ser}}(H)$. Call $\mathcal{N}^{(*)}$ a minimal cut of $\mathcal{N}_{\text{ser}}(H)$ and call $\gamma^{(*)}$ the capacity of $\mathcal{N}^{(*)}$.

Given that the capacity of the whole network $\mathcal{N}_{\text{ser}}(H)$ is the capacity $\gamma^{(*)}$ of its minimal cut $\mathcal{N}^{(*)}$, there exists a flow such that that at any time t the input of each module $\mathcal{N}^{(h)}$ is $\gamma^{(*)}$, which is at most the capacity of $\mathcal{N}^{(h)}$. The modularity of the network assures that no flow on $\mathcal{N}_{\text{ser}}(H)$ when restricted to $\mathcal{N}^{(h)}$ can do better than the optimal flow on $\mathcal{N}^{(h)}$. Since there exists a flow on $\mathcal{N}_{\text{ser}}(H)$ that coincides with the optimal flows on each $\mathcal{N}^{(h)}$, we have the result.

(b) The same modularity argument applies to the equilibrium. \square

Proof of Corollary 5.2

Proof of Corollary 5.2. By definition

$$\begin{aligned} \text{Poa}(\mathcal{N}_{\text{ser}}(H), \gamma^{(*)}) &= \frac{\text{Opt}(\mathcal{N}_{\text{ser}}(H), \gamma^{(*)})}{\text{Eq}(\mathcal{N}_{\text{ser}}(H), \gamma^{(*)})} \\ &= \frac{\sum_{h=1}^H \text{Opt}(\mathcal{N}^{(h)}, \gamma^{(*)})}{\sum_{h=1}^H \text{Eq}(\mathcal{N}^{(h)}, \gamma^{(*)})} \\ &= \sum_{h=1}^H \frac{\text{Eq}(\mathcal{N}^{(h)}, \gamma^{(*)})}{\sum_{h=1}^H \text{Eq}(\mathcal{N}^{(h)}, \gamma^{(*)})} \frac{\text{Opt}(\mathcal{N}^{(h)}, \gamma^{(*)})}{\text{Eq}(\mathcal{N}^{(h)}, \gamma^{(*)})} \\ &\leq \max_{h \in \{1, \dots, H\}} \frac{\text{Opt}(\mathcal{N}^{(h)}, \gamma^{(*)})}{\text{Eq}(\mathcal{N}^{(h)}, \gamma^{(*)})} \\ &= \max_{h \in \{1, \dots, H\}} \text{Poa}(\mathcal{N}^{(h)}, \gamma^{(*)}), \end{aligned}$$

which gives the desired result. □

Acknowledgments

We thank Ludovic Renou for helpful suggestions and comments.

References

- AKAMATSU, T. (2000) A dynamic traffic equilibrium assignment paradox. *Transportation Res. Part B* **34**, 515–531. URL <http://www.sciencedirect.com/science/article/pii/S0191261599000363>.
- AKAMATSU, T. (2001) An efficient algorithm for dynamic traffic equilibrium assignment with queues. *Transportation Sci.* **35**, 389–404. URL <http://transci.journal.informs.org/content/35/4/389.abstract>.
- AKAMATSU, T. and HEYDECKER, B. (2003) Detecting dynamic traffic assignment capacity paradoxes in saturated networks. *Transportation Sci.* **37**, 123–138. URL <http://transci.journal.informs.org/content/37/2/123.abstract>.

- ANSHELEVICH, E. and UKKUSURI, S. (2009) Equilibria in dynamic selfish routing. In MAVRONICOLAS, M. and PAPADOPOULOU, V. G. (eds.), *Algorithmic Game Theory*, volume 5814, 171–182. Springer Berlin Heidelberg. URL http://dx.doi.org/10.1007/978-3-642-04645-2_16.
- BHASKAR, U., FLEISCHER, L., HOY, D., and HUANG, C.-C. (2009) Equilibria of atomic flow games are not unique. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, 748–757. SIAM, Philadelphia, PA.
- BHASKAR, U., FLEISCHER, L., and HUANG, C.-C. (2010) The price of collusion in series-parallel networks. In *Integer programming and combinatorial optimization*, volume 6080 of *Lecture Notes in Comput. Sci.*, 313–326. Springer, Berlin. URL http://dx.doi.org/10.1007/978-3-642-13036-6_24.
- BRAESS, D. (1968) Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung* **12**, 258–268.
- BRAESS, D. (2005) On a paradox of traffic planning. *Transportation Science* **39**, 446–450. Translation of the original German (1968) article by D. Braess, A. Nagurney and T. Wakolbinger.
- CHARNES, A. and COOPER, W. W. (1961) Multicopy traffic network models. In HERMAN, R. (ed.), *Theory of Traffic Flow*. Elsevier, Amsterdam. Proceedings of the Symposium on the Theory of Traffic Flow held at the General Motors Research Laboratories, Warren, MI.
- COMINETTI, R., CORREA, J. R., and LARRÉ, O. (2011) Existence and uniqueness of equilibria for flows over time. In *Automata, Languages and Programming. Part II*, volume 6756 of *Lecture Notes in Comput. Sci.*, 552–563. Springer, Heidelberg. URL http://dx.doi.org/10.1007/978-3-642-22012-8_44.
- CORREA, J. R., SCHULZ, A. S., and STIER-MOSES, N. E. (2004) Selfish routing in capacitated networks. *Math. Oper. Res.* **29**, 961–976. URL <http://dx.doi.org/10.1287/moor.1040.0098>.
- CORREA, J. R., SCHULZ, A. S., and STIER-MOSES, N. E. (2007) Fast, fair, and efficient flows in networks. *Oper. Res.* **55**, 215–225. URL <http://dx.doi.org/10.1287/opre.1070.0383>.
- CORREA, J. R., SCHULZ, A. S., and STIER-MOSES, N. E. (2008) A geometric approach to the price of anarchy in nonatomic congestion games. *Games Econom. Behav.* **64**, 457–469. URL <http://dx.doi.org/10.1016/j.geb.2008.01.001>.

- CORREA, J. R. and STIER-MOSES, N. E. (2010) Wardrop equilibria. In COCHRAN, J. J., COX, L. A., KESKINOCAK, P., KHAROUFEH, J. P., and SMITH, J. C. (eds.), *Wiley Encyclopedia of Operations Research and Management Science*. John Wiley & Sons, Inc. URL <http://dx.doi.org/10.1002/9780470400531.eorms0962>.
- FORD, JR., L. R. and FULKERSON, D. R. (1958) Constructing maximal dynamic flows from static flows. *Oper. Res.* **6**, 419–433.
- HARKER, P. T. (1988) Multiple equilibrium behaviors on networks. *Transportation Sci.* **22**, 39–46. URL <http://transci.journal.informs.org/content/22/1/39.abstract>.
- HAURIE, A. and MARCOTTE, P. (1985) On the relationship between Nash-Cournot and Wardrop equilibria. *Networks* **15**, 295–308. URL <http://dx.doi.org/10.1002/net.3230150303>.
- HENDRICKSON, C. and KOCUR, G. (1981) Schedule delay and departure time decisions in a deterministic model. *Transportation Sci.* **15**, 62–77. URL <http://www.jstor.org/stable/25768001>.
- HOEFER, M., MIRROKNI, V. S., RÖGLIN, H., and TENG, S.-H. (2009) Competitive routing over time. In *Proceedings of the 5th International Workshop on Internet and Network Economics*, WINE '09, 18–29. Springer-Verlag, Berlin, Heidelberg. URL http://dx.doi.org/10.1007/978-3-642-10841-9_4.
- KOCH, R., NASRABADI, E., and SKUTELLA, M. (2011) Continuous and discrete flows over time: a general model based on measure theory. *Math. Methods Oper. Res.* **73**, 301–337. URL <http://dx.doi.org/10.1007/s00186-011-0357-2>.
- KOCH, R. and SKUTELLA, M. (2011) Nash equilibria and the price of anarchy for flows over time. *Theory Comput. Syst.* **49**, 71–97. URL <http://dx.doi.org/10.1007/s00224-010-9299-y>.
- KOUTSOUPIAS, E. and PAPADIMITRIOU, C. (1999) Worst-case equilibria. In *STACS 99 (Trier)*, volume 1563 of *Lecture Notes in Comput. Sci.*, 404–413. Springer, Berlin. URL http://dx.doi.org/10.1007/3-540-49116-3_38.
- MACKO, M., LARSON, K., and STESKAL, L. (2010) Braess’s paradox for flows over time. *Computing Research Repository* [abs/1007.4864](https://arxiv.org/abs/1007.4864), 262–275.

- MONDERER, D. and SHAPLEY, L. S. (1996) Potential games. *Games Econom. Behav.* **14**, 124–143. URL <http://dx.doi.org/10.1006/game.1996.0044>.
- MOUNCE, R. (2006) Convergence in a continuous dynamic queueing model for traffic networks. *Transportation Res. Part B* **40**, 779–791. URL <http://www.sciencedirect.com/science/article/pii/S0191261505001153>.
- MOUNCE, R. (2007) Convergence to equilibrium in dynamic traffic networks when route cost is decay monotone. *Transportation Sci.* **41**, 409–414. URL <http://transci.journal.informs.org/content/41/3/409.abstract>.
- PAPADIMITRIOU, C. (2001) Algorithms, games, and the internet. In *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing*, 749–753. ACM, New York. URL <http://dx.doi.org/10.1145/380752.380883>.
- ROSENTHAL, R. W. (1973) A class of games possessing pure-strategy Nash equilibria. *Internat. J. Game Theory* **2**, 65–67.
- ROUGHGARDEN, T. (2005) *Selfish Routing and the Price of Anarchy*. MIT Press, Cambridge, Mass.
- ROUGHGARDEN, T. (2007) Routing games. In *Algorithmic Game Theory*, 461–486. Cambridge Univ. Press, Cambridge.
- ROUGHGARDEN, T. and TARDOS, É. (2002) How bad is selfish routing? *J. ACM* **49**, 236–259 (electronic). URL <http://dx.doi.org/10.1145/506147.506153>.
- ROUGHGARDEN, T. and TARDOS, É. (2004) Bounding the inefficiency of equilibria in nonatomic congestion games. *Games Econom. Behav.* **47**, 389–403. URL <http://dx.doi.org/10.1016/j.geb.2003.06.004>.
- ROUGHGARDEN, T. and TARDOS, É. (2007) Introduction to the inefficiency of equilibria. In *Algorithmic Game Theory*, 443–459. Cambridge Univ. Press, Cambridge.
- SHAH, D. and SHIN, J. (2010) Dynamics in congestion games. In *Proceedings of the ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems - SIGMETRICS '10*.

- Association for Computing Machinery / ACM-Sigmetrics, New York, New York. URL <http://dx.doi.org/10.1145/1811099.1811052>.
- VICKREY, W. S. (1969) Congestion theory and transport investment. *Amer. Econ. Rev.* **59**, 251–260. URL <http://ideas.repec.org/a/aea/aecrev/v59y1969i2p251-60.html>.
- WARDROP, J. G. (1952) Some theoretical aspects of road traffic research. In *Proceedings of the Institute of Civil Engineers, Pt. II*, volume 1, 325–378. URL <http://www.icevirtuallibrary.com/content/article/10.1680/ipeds.1952.11362>.
- XIA, Y. and HILL, D. (2013) Dynamic Braess’s paradox in complex communication networks. *IEEE Transactions on Circuits and Systems II: Express Briefs* **60**, 172–176.
- YAGAR, S. (1971) Dynamic traffic assignment by individual path minimization and queuing. *Transp. Res.* **5**, 179–196.