Optimizing Throughput and Makespan of Queuing Systems by Information Design

Svenja M. Griesbach
Institute of Mathematics, Technische Universität Berlin, Germany

Max Klimm
Institute of Mathematics, Technische Universität Berlin, Germany

Philipp Warode
School of Business and Economics, Humboldt-Universität zu Berlin, Germany

Theresa Ziemke
Institute of Mathematics, Technische Universität Berlin, Germany

Abstract

We study the optimal provision of information for two natural performance measures of queuing systems: throughput and makespan. A set of parallel links (queues) is equipped with deterministic capacities and stochastic offsets where the latter depend on a realized state, and the number of states is assumed to be constant. A continuum of flow particles (agents) arrives at the system at a constant rate. A system operator knows the realization of the state and may (partially) reveal this information via a public signaling scheme to the flow particles. Upon arrival, the flow particles observe the signal issued by the system operator, form an updated belief about the realized state, and decide on which link they use. Inflow into a link exceeding the link's capacity builds up in a queue that increases the cost (total travel time) on the link. Dynamic inflow rates are in a Bayesian dynamic equilibrium when the expected cost along all links with positive inflow is equal at every point in time and not larger than the expected cost of any unused link. For a given time horizon T, the throughput induced by a signaling scheme is the total volume of flow that leaves the links in the interval $[0, T]$. The public signaling scheme maximizing the throughput may involve irrational numbers. We provide an additive polynomial time approximation scheme (PTAS) that approximates the optimal throughput by an arbitrary additive constant $\varepsilon > 0$. The algorithm solves a Lagrangian dual of the signaling problem with the Ellipsoid method whose separation oracle is implemented by a cell decomposition technique. We also provide a multiplicative fully polynomial time approximation scheme (FPTAS) that does not rely on strong duality and, thus, allows to compute the optimal signals. It uses a different cell decomposition technique together with a piecewise convex under-estimator of the optimal value function. Finally, we consider the makespan of a Bayesian dynamic equilibrium which is defined as the last point in time when a total given value of flow leaves the system. Using a variational inequality argument, we show that full information revelation is a public signaling scheme that minimizes the makespan.

2012 ACM Subject Classification Theory of computation → Routing and network design problems; Theory of computation → Nonconvex optimization; Applied computing → Transportation

Keywords and phrases Information Design, Dynamic Flows, Public Signals, Convex Envelope

Digital Object Identifier 10.4230/LIPIcs.ESA.2024.62

Related Version *Full Version*: https://arxiv.org/abs/2310.07656

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

Acknowledgements We thank Martin Hoefer and the other participants of Dagstuhl seminar 22192 "Dynamic Traffic Models in Transportation Science" for inspiration and discussions.

1 Introduction

Imagine you are managing an airport with several security lanes: Should you inform the passengers about the current state of the system in order to maximize its performance? On the one hand, informing passengers about less congested lanes may reduce queue times, but on the other hand it may lead to inefficiencies if lanes with delays remain underutilized.

In this paper, we study this question using the framework of Bayesian persuasion for a dynamic queuing model first studied (in a more basic and deterministic setting) by Vickrey [29].

Specifically, we consider a set $[m] \coloneqq \{1,\dots,m\}$ of parallel links. The state of each link is stochastic and depends on a realized state $\theta \in \Theta$ where $d \coloneqq |\Theta|$ is constant. First, each link i has a state-independent *capacity* $\nu_i \in \mathbb{Q}_{>0}$. Second, each link i has an *offset* (free-flow travel time) $b_{i,\theta} \in \mathbb{R}_{\geq 0}$ that depends on the realized state θ. Initially, the links have no queues. A steady flow of infinitesimally small agents arrives at the links with a constant inflow rate of $u \in \mathbb{Q}_{>0}$. Upon arrival, each of the flow particles chooses a link leading to an inflow rate of $f_i(t)$ into each link i for all times t. When at some time t the inflow rate $f_i(t)$ into a link i exceeds its capacity ν_i, a queue forms at link i at a rate of $f_i(t) - \nu_i$. When at some time t the inflow rate into a link i is less than the capacity ν_i and the link has a queue, the queue depletes at a rate of $\nu_i - f_i(t)$. This queuing dynamics lead to uniquely defined queue lengths $z_i(t)$. The *cost* of a flow particle that arrives at the system at time t and chooses link i when the realized state is θ is given by $c_{i,\theta}(t) = \frac{z_i(t)}{\nu_i} + b_{i,\theta}$.

The flow particles, however, do not know the realized state θ and, instead, form stochastic beliefs about the state of the system. A belief is a vector $\boldsymbol{\mu} = (\mu_1,\dots,\mu_d)^\top \in [0,1]^d$ such that $\sum_{\theta\in\Theta} \mu_\theta = 1$ and μ_θ is the (anticipated) probability that state θ is realized. In the following, we denote by Δ the set of all such beliefs. All flow particles have initially a true prior belief $\boldsymbol{\mu}^* = (\mu_1^*,\dots,\mu_d^*)^\top \in \Delta$, e.g., from previous observations of the system. If the flow particles do not receive any further information on the state of the system, they each choose a link that minimizes their expected delay (where the expectation is taken according to the prior belief $\boldsymbol{\mu}^*$), i.e., when arriving at time t they choose a link i that minimizes $\sum_{\theta\in\Theta} \mu_\theta^* b_{i,\theta} + z_i(t)/\nu_i$. We say that a vector of inflow functions $\mathbf{f} = (f_i(\cdot))_{i\in[m]}$ is a *dynamic equilibrium* (with respect to the belief $\boldsymbol{\mu}^*$) if this property of the particles' behavior is satisfied for almost all t.

The operator of the system, however, knows the realized state and, hence, the travel times of the links and can determine in how far this information should be shared with the flow particles. To this end, the system operator commits to a *public signaling scheme* Φ. Such a signaling scheme consists of a finite set of signals Σ, as well as probabilities $(\varphi_{\theta,\sigma})_{\theta\in\Theta,\sigma\in\Sigma}$ where $\varphi_{\theta,\sigma}$ is the combined probability that state θ is realized and signal σ is issued. Since the prior belief μ_θ^* represents the true probability that state θ is realized, we have the constraint $\sum_{\sigma\in\Sigma} \varphi_{\theta,\sigma} = \mu_\theta^*$ for all $\theta \in \Theta$. When arriving at the system, the flow particles observe the issued signal σ (but not the realized state θ) and perform a Bayesian update of their belief. In particular, after having received signal σ, their posterior belief is given by $\boldsymbol{\mu}^\sigma = (\mu_1^\sigma,\dots,\mu_d^\sigma)^\top \in [0,1]^d$ defined as $\mu_\theta^\sigma = \varphi_{\theta,\sigma}/\sum_{\theta'\in\Theta} \varphi_{\theta',\sigma}$. After this Bayesian update, the flow particles choose a link i that minimizes the updated expected delays given by $\sum_{\theta\in\Theta} \mu_\theta^\sigma b_{i,\theta} + z_i(t)/\nu_i$. It is a standard observation in the field of Bayesian persuasion that there is a one-to-one correspondence between public signaling schemes and convex decompositions of the prior $\boldsymbol{\mu}^*$ (Kamenica and Gentzkow [20]). Specifically, every public signaling scheme yields posterior beliefs $(\boldsymbol{\mu}^\sigma)_{\sigma\in\Sigma}$ as well as corresponding probabilities $\varphi_\sigma \coloneqq \sum_{\theta\in\Theta} \varphi_{\theta,\sigma}$ that signal σ is issued such that $\boldsymbol{\mu}^* = \sum_{\sigma\in\Sigma} \varphi_\sigma \boldsymbol{\mu}^\sigma$. Conversely, for every such (finite) convex decomposition of the prior, there is a corresponding public signaling scheme with a finite set of signals such that the updated beliefs (after receiving one of the signals) correspond exactly to the beliefs in the convex decomposition.

Suppose, for each belief $\boldsymbol{\mu} \in \Delta$ there is a unique way to choose a corresponding dynamic equilibrium $\mathbf{f}(\boldsymbol{\mu})$. Further assume, the system operator aims to maximize a certain functional $H(\mathbf{f})$ of the flow vector $\mathbf{f}$. Then, a natural question is to find the public signaling scheme that optimizes the expected value of the functional of the resulting dynamic equilibrium. Using the one-to-one correspondence of public signaling schemes and convex decompositions of the prior, this can be phrased as

$$\sup\left\{\sum_{\sigma\in\Sigma}\varphi_\sigma\, H\big(\mathbf{f}(\boldsymbol{\mu}^\sigma)\big) : |\Sigma|<\infty, \varphi_\sigma\in[0,1], \boldsymbol{\mu}^\sigma\in\Delta \text{ for all } \sigma\in\Sigma \text{ and } \sum_{\sigma\in\Sigma}\varphi_\sigma\boldsymbol{\mu}^\sigma = \boldsymbol{\mu}^*\right\}. \quad (1)$$

Our Results and Techniques. We study the generic optimization problem (1) for two natural objectives of a system operator. First, we study the objective of maximizing the expected *throughput* of a system. Informally, for a given time horizon $T \in \mathbb{Q}_{>0}$, the throughput is the amount of flow that has left the links up to time T. The full version of this paper contains an example exhibiting that the signaling scheme maximizing the throughput may involve irrational numbers, even though all input numbers are rational. To avoid the issue of representing irrational numbers with finite precision, we resort to approximating the maximum achievable throughput by public signaling schemes. To this end, we first provide an additive polynomial-time approximation scheme (PTAS), i.e., for any $\varepsilon > 0$, we provide an algorithm that runs in polynomial time for constant $|\Theta|$ and computes a value p such that $p \in [\textsc{Opt} - \varepsilon, \textsc{Opt}]$, where $\textsc{Opt}$ is the maximal throughput that can be achieved by public signals (Theorem 9). We stress that, unlike other PTAS for signaling (e.g., [4, 9]), we do *not* require that the functional is normalized, i.e., that $||H||_\infty = 1$.

To prove the result, we consider a Lagrangian dual of the primal signaling problem and show that strong duality holds. The proof of strong duality is non-trivial since the objective of the primal is non-convex and non-concave such that standard constraint qualifications such as Slater's cannot be applied. Duality has been used before in the context of signaling by Bhaskar et al. [4], but they use standard linear programming duality for an approximate version of the primal problem such that they can only show a (small) bound on the duality gap. Our dual signaling problem has a finite number of variables but an uncountable number of constraints. Yet, we are able to show that the separation problem for the dual signaling problem is solvable in polynomial time. To this end, we show that the separation problem for the dual can be reduced to finding the global maximum of a piecewise quadratic function whose quadratic parts have a polytopal domain. Using a cell decomposition technique together with the reverse search algorithm by Avis and Fukuda [2] allows to compute the global maximum exactly, thus, solving the separation problem. Finally, we use the Ellipsoid method and the equivalence of optimization and separation to obtain the result.

While the *additive* PTAS yields a compelling approximation of the optimal throughput achievable by signaling, it does not allow to compute the corresponding signals. The underlying reason is that the approximately optimal *dual* solution obtained by the Ellipsoid method does not seem to provide any useful information on how approximately optimal *primal* solutions may look like. To close this gap, we propose a fully polynomial-time approximation scheme (FPTAS) for constant $|\Theta|$ with a *multiplicative* approximation guarantee that allows to compute the corresponding signals (Theorem 16). For the multiplicative FPTAS, again a main issue is that the objective is a non-convex and non-concave function on the space of beliefs Δ. We propose a non-uniform discretization of Δ that leads to a piecewise affine under-estimator of the objective. By controlling the approximation error of the under-estimator, we are able to compute signals such that the expected throughput ALG achieved by the signals satisfies ALG $\geq (1-\varepsilon)$OPT.

The second objective that we study is the expected *makespan*. For a given time horizon $T \in \mathbb{Q}_{>0}$, the makespan is the latest point in time a flow particle that departed in the time interval $[0, T]$ leaves the system. To analyze optimal signaling for the minimization of the makespan, we show a general property for dynamic equilibria. Suppose we are given a system with a vector of deterministic offsets $\mathbf{b} = (b_e)_{e\in E}$. Further, let $\mathbf{b}' = (b'_e)_{e\in E}$ be a vector of (potentially different) deterministic offsets and let $\mathbf{f}(\mathbf{b}')$ be a dynamic equilibrium where particles act as if the travel times were $\mathbf{b}'$. Then, we show that the makespan of the dynamic equilibrium $\mathbf{f}(\mathbf{b}')$ is minimized when $\mathbf{b}' = \mathbf{b}$ (up to constant shifts). This implies in particular that full information revelation is always an optimal signaling scheme for makespan minimization (Theorem 21).

Related Work. Optimal signaling for congestion-prone systems is primarily studied in Wardrop's static equilibrium model (e.g., [1, 13, 23, 22, 28, 30]). For the Wardrop model with affine costs, Bhaskar et al. [4] showed that it is NP-hard to compute a public signaling scheme that approximates the total travel time better than a factor of 4/3. Griesbach et al. [17] proved that optimal information revelation is always optimal if and only if the underlying network is series-parallel and provided an algorithm computing the optimal public signaling scheme for parallel links when the number of states and commodities is constant. For atomic congestion games, Castiglioni et al. [8] studied information design, but considered a different model where players commit to following the signal before they receive it. Zhou et al. [32] computed public and private signals in singleton games with a constant number of resources.

A cell decomposition related to ours has been used in the context of information design by Xu [31]. In contrast to our work, their model features receivers with a binary choice only. In addition, in the work of Xu it is assumed that the utility of a receiver does not depend on the actions of other agents. Also the meaning of the cells in the decomposition is quite different. In Xu's work in a cell the response of the receivers is constant whereas in our cells a certain ordering of the links with respect to the equilibrium flow is constant.

The dynamic flow model that we use here dates back to Vickrey [29]. It has been studied in more detail, e.g., by Koch and Skutella [21] and Cominetti et al. [10]. Koch and Skutella [21] also showed that the price of anarchy with respect to the throughput objective, i.e., the worst-case ratio of the throughput of an arbitrary dynamic flow and that of a dynamic equilibrium, is unbounded on general networks. For the makespan objective, Bhaskar et al. [5] showed that the price of anarchy is $\mathrm{e}/(\mathrm{e}-1)$ when one is allowed to reduce the capacity of the links arbitrarily (but still compares to the optimal flow for the original capacities). Correa et al. [12] showed that this bound on the price of anarchy also holds when the inflow rate at the source can be reduced. For parallel link networks (as considered in this work), they showed that the price of anarchy is 4/3. For both objectives, bounds on the price of anarchy are relevant for information design since, if they exist, they yield an approximation guarantee for the signaling scheme of full information revelation. The long-term behavior of dynamic equilibria in this model has been explored (see [11, 24]). Also, further variants of the model with multiple commodities, more complicated queuing behavior, or further side constraints have been studied (e.g., [16, 19, 26, 27]), but they do not have any implications for mechanism or information design. Graf et al. [15] study a model where users use machine learning to predict future traffic states, and compare different predictors empirically. Oosterwijk et al. [25] study a dynamic model on parallel paths where players need to meet a certain time deadline and try to minimize the costs of the used links. They obtain tight bounds on the price of anarchy both for the makespan and the throughput objective.

2 Preliminaries

For an integer $m \in \mathbb{N}$, let $[m] := \{1, \dots, m\}$ and $[m]_0 := \{0, \dots, m\}$. For $x \in \mathbb{R}$, we denote by $[x]^+ := \max\{x, 0\}$ the positive part and by $[x]^- := \min\{x, 0\}$ the negative part of x. We denote vectors and matrices with bold face and assume that all vectors are column vectors. Further, $\mathbf{e}_i$ denotes the i-th unit vector, $\mathbf{1}$ the all-ones vector, and $\mathbf{0}$ the all-zeros vector (of the appropriate dimension). We first introduce the dynamic equilibrium model with deterministic travel times and then turn to the dynamic equilibrium model with stochastic travel times.

Dynamic Equilibrium with Deterministic Travel Times. Consider a set $[m]$ of parallel links where each link $i \in [m]$ has a *capacity* $\nu_i \in \mathbb{Q}_{>0}$ and a constant *offset* (free-flow travel time) $b_i \in \mathbb{Q}_{\geq 0}$. There is a continuum of flow particles arriving at the links with a constant rate of $u \in \mathbb{Q}_{>0}$. A *dynamic flow* is a family of measurable functions $\mathbf{f} = (f_i)_{i \in [m]}$ with $f_i \colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ satisfying $\sum_{i \in [m]} f_i(t) = u$ for almost all times $t \geq 0$. The value $f_i(t)$ describes the inflow into link i at time t. Each link operates with the following queuing dynamics: if the inflow into a link i is higher than its capacity ν_i, a queue builds up. Particles in the queue are processed with rate ν_i. After passing the queue, it takes an additional amount of time b_i to traverse the link. We denote by $z_i(t)$ the length of the queue at any given time $t \geq 0$. The queue dynamics are described via the differential equation

$$z_i'(t) = \begin{cases} f_i(t) - \nu_i & \text{if } z_i(t) > 0, \\ [f_i(t) - \nu_i]^+ & \text{if } z_i(t) = 0. \end{cases} \tag{2}$$

A flow particle that enters link i at time t waits for time $\frac{z_i(t)}{\nu_i}$ in the queue and then experiences a free-flow travel time of b_i. Therefore, the cost of a flow particle entering link i at time t is $c_i(t) := \frac{z_i(t)}{\nu_i} + b_i$ and thus its *exit time* (total cost) is $C_i(t) := t + \frac{z_i(t)}{\nu_i} + b_i$. A flow is a *dynamic equilibrium* if almost all particles have no incentive to deviate to a different link, i.e., if $c_i(t) = \min_{j \in [m]} c_j(t)$ for all $i \in [m]$ with $f_i(t) > 0$ for almost all $t \geq 0$. A link i with minimal cost $c_i(t)$ at time t is called *active* and we denote the set of active links at time t by $\mathcal{A}(t) := \big\{i \in [m] : c_i(t) = \min_{j \in [m]} c_j(t)\big\}$. In general, the dynamic equilibrium may not be unique, but the exit times and the set of active links are (cf. [10, 24]). Further, we refer to the set of links with positive inflow at time t as the *support* of the flow and denote it by $S(t) := \big\{i \in [m] : f_i(t) > 0\big\}$. For a given time horizon $T > 0$, the *throughput* of a flow $\mathbf{f}$ is defined as $F_T(\mathbf{f}) := \int_0^T \sum_{i \in [m]} f_i^-(t) \,\mathrm{d}t$, where $f_i^-(t)$ is the outflow of link i at time t that can be computed as

$$f_i^-(t + b_i) = \begin{cases} \min\{f_i(t), \nu_i\} & \text{if } z_i(t) = 0, \text{ and} \\ \nu_i & \text{if } z_i(t) > 0. \end{cases}$$

For a given time horizon $T > 0$, the *makespan* of a flow $\mathbf{f}$ is defined as $M_T(\mathbf{f}) := \sup\big\{C_i(t) : t \in [0, T], i \in S(t)\big\}$.

Bayesian Dynamic Equilibrium with Stochastic Travel Times. Consider a set Θ of states with $d := |\Theta|$. We assume that the offsets $b_{i,\theta} \in \mathbb{Q}_{\geq 0}$ of each link $i \in [m]$ depend on the state $\theta \in \Theta$, and write $\mathbf{b}_i = (b_{i,\theta})_{\theta\in\Theta}$. The capacities $\nu_i \in \mathbb{Q}_{>0}$ are independent of the state. Every belief $\boldsymbol{\mu} \in \Delta$ essentially induces a system with stochastic offsets by replacing the deterministic offsets b_i with its expectation $\boldsymbol{\mu}^\top \mathbf{b}_i$. With a slight overload of notation, we use the same notation for deterministic and stochastic offsets. In particular, the *expected cost* of a flow particle entering link i at time t is $c_i(t) := \frac{z_i(t)}{\nu_i} + \boldsymbol{\mu}^\top \mathbf{b}_i$ and thus its *expected exit time* (expected total cost) is $C_i(t) := t + \frac{z_i(t)}{\nu_i} + \boldsymbol{\mu}^\top \mathbf{b}_i$. The exit time of a flow particle entering link i at time t when state θ is realized is $C_{i,\theta}(t) := t + \frac{z_i(t)}{\nu_i} + b_{i,\theta}$. We call a flow $\mathbf{f}$ a *Bayesian dynamic equilibrium* if $c_i(t) = \min_{j\in[m]} c_j(t)$ for all links $i \in [m]$ with $f_i(t) > 0$. For a fixed time horizon $T > 0$, the throughput of a flow $\mathbf{f}$ in state $\theta \in \Theta$ is defined as $F_{T,\theta}(\mathbf{f}) := \int_0^T \sum_{i\in[m]} f^-_{i,\theta}(t)\,\mathrm{d}t$ where the outflow of link i at time t in state θ is

$$f^-_{i,\theta}(t) = \begin{cases} \min\{f_i(t - b_{i,\theta}), \nu_i\} & \text{if } t - b_{i,\theta} \geq 0 \text{ and } z_i(t - b_{i,\theta}) = 0,\\ \nu_i & \text{if } t - b_{i,\theta} \geq 0 \text{ and } z_i(t) > 0\\ 0 & \text{else.} \end{cases}$$

The *expected throughput* of a flow $\mathbf{f}$ (according to belief $\boldsymbol{\mu}$) is then given by $F_T(\mathbf{f}) := \sum_{\theta\in\Theta} \mu_\theta F_{T,\theta}(\mathbf{f})$. For a time horizon $T > 0$ and a state θ, the makespan of a flow $\mathbf{f}$ in state θ is defined as $M_{T,\theta}(\mathbf{f}) := \sup\{C_{i,\theta}(t) : t \in [0, T], i \in [m] \text{ with } f_i(t) > 0\}$ and the expected makespan is then $M_T(\mathbf{f}) := \sum_{\theta\in\Theta} \mu_\theta M_{T,\theta}(\mathbf{f})$.

Information Design. We assume flow particles to have a prior belief $\boldsymbol{\mu}^*$. There is a one-to-one correspondence between public signaling schemes and convex decompositions of the prior, and – due to Caratheodory's Theorem – the convex decomposition requires at most d signals [14, 20]. We note that for a belief $\boldsymbol{\mu}$, the corresponding Bayesian dynamic equilibrium may not be unique. We here assume that in case multiple equilibria exist, we choose the one that yields in expectation the best objective for the system designer; this is a standard assumption in the information design literature and is, e.g., justified since the information designer may signal the best equilibrium play, see also the discussion by Bergemann and Morris [3]. Formally, for the throughput objective, for a fixed time horizon T, let us define $F\colon \Delta \to \mathbb{R}_{\geq 0}$ be the expected throughput of a Bayesian dynamic equilibrium for belief $\boldsymbol{\mu}$ that maximizes the throughput among all such equilibria. To compute the public signaling scheme that maximizes the expected throughput, we are interested in solving

$$\sup\left\{\sum_{\sigma\in\Sigma} \varphi_\sigma F(\boldsymbol{\mu}_\sigma) : |\Sigma| \leq d, \varphi_\sigma \in [0,1], \boldsymbol{\mu}_\sigma \in \Delta \text{ for all } \sigma \in \Sigma \text{ such that } \sum_{\sigma\in\Sigma} \varphi_\sigma \boldsymbol{\mu}_\sigma = \boldsymbol{\mu}^*\right\}.$$

Similarly, for a fixed time horizon T, let us define $M\colon \Delta \to \mathbb{R}_{\geq 0}$ as the expected makespan of a Bayesian dynamic equilibrium for belief $\boldsymbol{\mu}$ that minimizes the makespan among all such equilibria. To compute the public signaling scheme that minimizes the expected makespan, we are interested in solving

$$\inf\left\{\sum_{\sigma\in\Sigma} \varphi_\sigma M(\boldsymbol{\mu}_\sigma) : |\Sigma| \leq d, \varphi_\sigma \in [0,1], \boldsymbol{\mu}_\sigma \in \Delta \text{ for all } \sigma \in \Sigma \text{ such that } \sum_{\sigma\in\Sigma} \varphi_\sigma \boldsymbol{\mu}_\sigma = \boldsymbol{\mu}^*\right\}.$$

Before we proceed with the structural results we give an example for the throughput and the makespan objective.

▶ **Example 1** (Throughput maximization). We consider an example for the throughput objective. The instance consists of $m = 2$ edges, $|\Theta| = 2$ states, and a time horizon $T = 5$. The edges have capacities $\nu_1 = \frac{1}{3}$ and $\nu_2 = \frac{2}{3}$ and the offsets depend on the two states θ_1, θ_2 with $\mathbf{b}_1 = (1, 5)^\top$ and $\mathbf{b}_2 = (4, 3)^\top$. The arrival rate is set to $u = 1$ and the instance is depicted in Figure 1a.

The throughput function F is shown in Figure 1b. It is a continuous piecewise quadratic function over the set of beliefs $\mu \in [0, 1]$, where $\mu := \mu_{\theta_2}$ indicates the probability of state θ_2 being realized. The function F has two breakpoints at $\mu \in \left\{\frac{1}{5}, \frac{3}{5}\right\}$. For $\mu \in \left[0, \frac{1}{5}\right)$, the upper edge e_1 has a lower expected offset and, thus, the first flow particles only use edge e_1. Since the inflow rate u exceeds the capacity ν_1, a queue starts to build up, increasing the cost of edge e_1. Thus, at some point, flow particles start deviating to the lower edge e_2. With increasing μ, flow particles start to deviate earlier in time to edge e_2 as its expected offset decreases in μ while the expected offset of edge e_1 increases. For $\mu = \frac{1}{5}$, it is guaranteed, that the first particle, that deviates to edge e_2 will leave the queuing system before time horizon T no matter which state is realized. Hence, for $\mu \in \left[\frac{1}{5}, \frac{3}{5}\right]$, the throughput increases linearly in μ. For $\mu = \frac{3}{5}$, both edges have the same expected offset and, thus, both edges are used from the very first point in time. This results in the maximization of the total throughput F. For $\mu \in \left(\frac{3}{5}, 1\right]$, the first flow particles use only edge e_2 until a sufficiently large queue has built up such that the cost of edge e_2 equals the offset of edge e_1. Hence, flow particles deviate to edge e_1.

Additionally, Figure 1b depicts two convex decompositions of a prior $\boldsymbol{\mu}^*$ with the posterior beliefs μ^1, μ^2, μ^3 and $\tilde{\mu}^1, \tilde{\mu}^2$. The signaling scheme inducing the three posterior beliefs μ^i is an optimal signaling scheme. The signaling scheme inducing the two posterior beliefs $\tilde{\mu}^i$ corresponds to full information revelation and is suboptimal. The dashed lines give the expected throughput that can be achieved if the respective signaling scheme is used. ◀

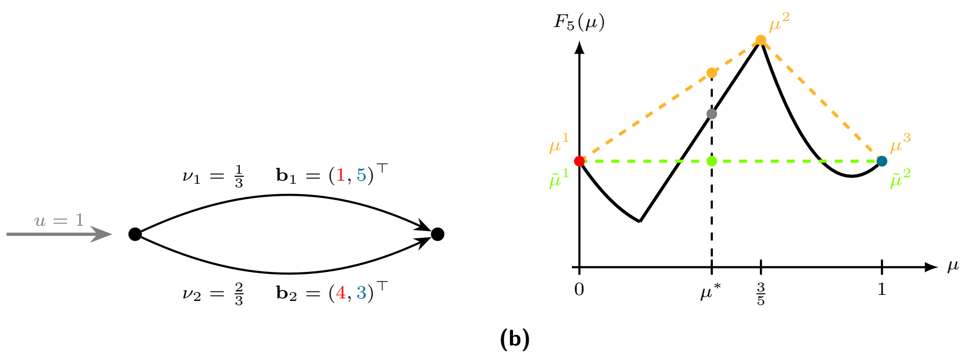

Figure 1 The illustration of Example 1 for the throughput objective. (a) The instance with $m = 2$ edges and 2 states θ_1, θ_2 in red and blue. (b) The throughput function $F_5(\mu)$ for $T = 5$ of the dynamic equilibrium as a function of the belief described by parameter $\mu := \mu_{\theta_2} \in [0, 1]$. The throughput in state θ_1, i.e., $\mu = 0$, is shown in red and the throughput in state θ_2, i.e., $\mu = 1$, is shown in blue. Full information revelation with posterior beliefs $\tilde{\mu}^1$ and $\tilde{\mu}^2$ in green is *not optimal*; a signaling scheme in orange incorporating the indifference point at $\mu^2 = \frac{3}{5}$ is optimal.

▶ **Example 2** (Makespan minimization). We consider an example for the makespan objective. The instance consists of $m = 3$ edges, $|\Theta| = 2$ states, and a time horizon $T = \frac{1}{2}$. The edges have capacities $\nu_1 = \frac{1}{2}$, $\nu_2 = \frac{1}{3}$, and $\nu_3 = \frac{1}{2}$, and the offsets depend on the two states θ_1, θ_2 with $\mathbf{b}_1 = (0, 5)^\top$, $\mathbf{b}_2 = (1, 1)^\top$, and $\mathbf{b}_3 = (4, 0)^\top$. The arrival rate is set to $u = 1$ and the instance is depicted in Figure 2a.

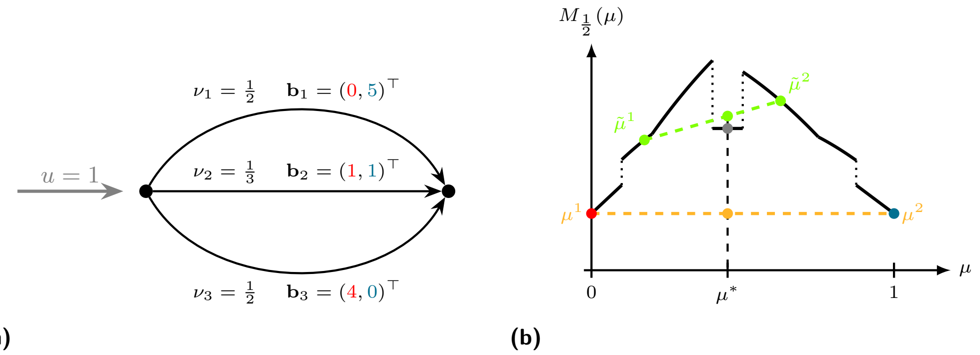

Figure 2 The illustration of Example 2 for the makespan objective. (a) The instance with $m = 3$ edges and 2 states θ_1, θ_2 in red and blue. (b) The makespan function $M_{\frac{1}{2}}(\mu)$ for $T = \frac{1}{2}$ of the dynamic equilibrium as a function of the belief described by parameter $\mu := \mu_{\theta_2} \in [0, 1]$. The makespan in state θ_1, i.e., $\mu = 0$, is shown in red and the throughput in state θ_2, i.e., $\mu = 1$, is shown in blue. Full information revelation with posterior beliefs μ^1 and μ^2 in orange is an optimal signaling scheme; the other signaling scheme with posterior beliefs $\tilde{\mu}^1$ and $\tilde{\mu}^2$ in green is suboptimal.

The makespan function M is shown in Figure 2b. It is a piecewise quadratic function over the set of beliefs $\mu \in [0, 1]$, where $\mu := \mu_{\theta_2}$ indicates the probability of state θ_2 being realized. The function M has six breakpoints at $\mu \in \left\{\frac{1}{10}, \frac{1}{5}, \frac{2}{5}, \frac{1}{2}, \frac{3}{4}, \frac{7}{8}\right\}$. Roughly speaking, M is continuous at a breakpoint, if the order in which edges are chosen by the flow particles change. In contrast to this, M is not continuous at a breakpoint, if the last flow particle entering the queuing system is indifferent between two edges and thus the makespan cannot be uniquely defined. In fact, the makespan depends on which edge the last particle chooses. In more detail, for $\mu \in \left[0, \frac{1}{10}\right)$, all flow particles use the upper edge e_1. Since the inflow rate $u = 1$ into edge e_1 exceeds its capacity, a queue builds up and, hence, M grows linearly in μ. When $\mu = \frac{1}{10}$, the last particle entering the queuing system is indifferent between choosing edge e_1 and the middle edge e_2. However, choosing edge e_2 increases the makespan M_{T,θ_1} for state θ_1 from 1 to $\frac{3}{2}$, resulting in an immediate increase in the makespan $M_T(\mathbf{f})$ and, thus, in a discontinuity point. With increasing μ, the flow particles start to deviate from edge e_1 to edge e_2 earlier in time. When $\mu = \frac{1}{5}$, already the very first particle entering the queuing system is indifferent between edges e_1 and e_2. Thus, for $\mu \in \left(\frac{1}{5}, \frac{2}{5}\right)$, the first flow particles use edge e_2 and only later in time, particles start to deviate to edge e_1. For $\mu = \frac{2}{5}$, the expected offset of edge e_1 is sufficiently large, such that the last particle entering the queuing system is the first particle indifferent between edges e_1 and e_2. In contrast to the case when $\mu = \frac{1}{10}$, this now results in an immediate decrease of the makespan. For $\mu \in \left(\frac{2}{5}, \frac{1}{2}\right)$, all flow particles use edge e_2. Since the offset of edge e_2 is deterministic, this gives a constant makespan for $\mu \in \left(\frac{2}{5}, \frac{1}{2}\right)$. We can make similar observations for the remaining values of μ. In short, for $\mu \in \left(\frac{1}{2}, \frac{3}{4}\right)$, flow particles start by using edge e_2 and then deviate to the lower edge e_3 later in time. For $\mu \in \left(\frac{3}{4}, \frac{7}{8}\right)$, flow particles start by using edge e_3 and then deviate to edge e_2 later in time. Finally, for $\mu \in \left(\frac{7}{8}, 1\right]$, all flow particles only use edge e_3.

Additionally, Figure 2b depicts two convex decompositions of a prior μ^* with posterior beliefs μ^1, μ^2 and $\tilde{\mu}^1, \tilde{\mu}^2$. The signaling scheme inducing the posterior beliefs μ^1, μ^2 corresponds to full information revelation and is optimal. The signaling scheme inducing the posterior beliefs $\tilde{\mu}^1, \tilde{\mu}^2$ corresponds to a suboptimal signaling scheme. The dashed lines give the expected makespan if the respective signaling scheme is used. ◀

3 Structural Results

In this section, we develop explicit formulas for the Bayesian dynamic equilibrium flows maximizing that maximizes the expected throughput or minimizes the expected makespan. Before we obtain these results for stochastic travel times, we first develop explicit formulas for the dynamic equilibria for deterministic travel times.

Deterministic Travel Times. In this section, we consider deterministic offsets and construct a flow $\mathbf{f}$ that is a dynamic equilibrium. In general, the dynamic equilibrium may not be unique. However, the exit times C_i and, thus, also the set of active links $\mathcal{A}$ are the same for any dynamic equilibrium (cf. [10, 24]). We assume for now that the links are ordered by their offsets, i.e., $b_1 \leq \dots \leq b_m$. We denote by $\bar{\nu}(i) := \sum_{j\in[i]} \nu_j$ the sum of the first i capacities, where $\bar{\nu}(0) := 0$. We define points in time t_i^* for $i \in [m]$ recursively as $t_1^* = 0$ and

$$t_{i+1}^* = \begin{cases} t_i^* + \frac{\bar{\nu}(i)}{u-\bar{\nu}(i)}(b_{i+1} - b_i) & \text{if } \bar{\nu}(i) < u, \\ \infty & \text{if } \bar{\nu}(i) \geq u, \end{cases} \qquad \text{for } i = 1, \dots, m-1. \tag{3}$$

As we will show, the time t_i^* is the point in time when link i becomes active. More precisely, link i is in the support for all times $t > t_i^*$. Let

$$k := \max\{j \in [m] : \bar{\nu}(j) < u\} \tag{4}$$

be the maximum index of a link such that the total capacities up to that link are strictly less than the inflow rate u of the instance. Hence, links 1 to $k+1$ suffice to handle the total inflow of u, and thus, we have $t_i^* = \infty$ for all $i > k+1$. For every $i \in [m]$, we define an inflow function $f_i \colon \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ by

$$f_i(t) := \begin{cases} \frac{u}{\bar{\nu}(j)} \cdot \nu_i & \text{if } i \leq k \text{ and } t \in (t_j^*, t_{j+1}^*], j = i, \dots, k \\ \nu_i & \text{if } i \leq k \text{ and } t > t_{k+1}^*, \\ u - \bar{\nu}_k & \text{if } i = k+1 \text{ and } t > t_{k+1}^*, \\ 0 & \text{otherwise}, \end{cases} \tag{5}$$

where $t_{m+1}^* = \infty$. Note that the inflow defined in (5) depends on the values t_i^*, which, in turn, depend on the offsets b_i. We proceed by showing that this flow is indeed a dynamic equilibrium. For the proof, we simply check the dynamic equilibrium conditions. This proof and all other missing proofs are deferred to the full version.

▶ **Lemma 3.** *Let $\mathbf{f} = (f_i)_{i\in[m]}$ be the flow defined in (5). Then,*

1. *the flow $\mathbf{f}$ is a feasible flow with queue lengths*

$$z_i(t) = \begin{cases} \nu_i \left(\frac{u - \bar{\nu}(|S(t)|)}{\bar{\nu}(|S(t)|)} \left(t + \sum_{j=1}^{|S(t)|} \frac{u\,\nu_j\,b_j}{(u - \bar{\nu}(j))(u - \bar{\nu}(j-1))} \right) - b_i \right) & \text{if } t_i^* < t \leq t_{k+1}^*, i \leq k, \\ \nu_i (b_{k+1} - b_i) & \text{if } t_{k+1}^* < t, i \leq k, \\ 0 & \text{otherwise;} \end{cases}$$

2. *the flow $\mathbf{f}$ is a dynamic equilibrium;*
3. *the queue length $z_i(t)$ on link i is continuous in $\mathbf{b}$;*
4. *the queue lengths $z_i(t)$ are non-increasing in b_i and non-decreasing in b_j for $j \neq i$. If $i \leq k$ and $t_i^* < t$, then $z_i(t)$ is strictly decreasing in b_i. If, additionally, $j \leq k$ and $t_j^* < t$, then $z_i(t)$ is strictly increasing in b_j.*

Stochastic Travel Times. In this subsection, we obtain structural results for the dynamic equilibrium $\mathbf{f}$ as a function of the belief $\boldsymbol{\mu} \in \Delta$. For a belief $\boldsymbol{\mu} \in \Delta$, let $\tau_b(\cdot\,;\boldsymbol{\mu}) : [m] \to [m]$ be an arbitrary permutation of the links that orders them non-decreasingly by their expected offsets (according to the belief $\boldsymbol{\mu}$), i.e., $\boldsymbol{\mu}^\top \mathbf{b}_{\tau_b(1;\boldsymbol{\mu})} \leq \cdots \leq \boldsymbol{\mu}^\top \mathbf{b}_{\tau_b(m;\boldsymbol{\mu})}$. Analogously to the deterministic case (3), the times $t^*_{\tau_b(i;\boldsymbol{\mu})}$ when the edges become active can be computed in dependence of $\boldsymbol{\mu}$ as $t^*_{\tau_b(1;\boldsymbol{\mu})}(\boldsymbol{\mu}) = 0$ and

$$t^*_{\tau_b(i+1;\boldsymbol{\mu})}(\boldsymbol{\mu}) = \begin{cases} t^*_{\tau_b(i;\boldsymbol{\mu})} + \dfrac{\bar{\nu}(i;\boldsymbol{\mu})}{u - \bar{\nu}(i;\boldsymbol{\mu})}\, \boldsymbol{\mu}^\top \left(\mathbf{b}_{\tau_b(i+1;\boldsymbol{\mu})} - \mathbf{b}_{\tau_b(i;\boldsymbol{\mu})}\right) & \text{if } \bar{\nu}(i;\boldsymbol{\mu}) < u, \\ \infty & \text{if } \bar{\nu}(i;\boldsymbol{\mu}) \geq u, \end{cases} \tag{6}$$

for all $i = 1, \ldots, m-1$, where $\bar{\nu}_b(i;\boldsymbol{\mu}) \coloneqq \sum_{j\in[i]} \nu_{\tau_b(j;\boldsymbol{\mu})}$ is the total capacity of the i links used first, depending on the belief $\boldsymbol{\mu}$.

We note that the ordering $\tau_b(\cdot;\boldsymbol{\mu})$ need not be unique since there is a degree of freedom in how to order links that have the same expected offsets. Every ordering will result in a corresponding Bayesian dynamic equilibrium, but the ordering may influence the expected throughput or makespan of the system. For the throughput objective and a given belief $\boldsymbol{\mu} \in \Delta$, we first compute an arbitrary ordering $\tau_b(\cdot;\boldsymbol{\mu})$ by non-decreasing expected offsets and obtain values $t^*_{\tau_b(\cdot;\boldsymbol{\mu})}$. Assume that two links $i, i+1 \in [m]$ have the same expected offset, then using (6) we have $t^*_{\tau_b(i;\boldsymbol{\mu})} = t^*_{\tau_b(i+1;\boldsymbol{\mu})}$. We reorder them non-increasingly by their contribution to the throughput objective, i.e., we assume that

$$\boldsymbol{\mu}^\top [\mathbf{1}(T - t^*_{\tau_b(i;\boldsymbol{\mu})}) - \mathbf{b}_{\tau_b(i;\boldsymbol{\mu})}]^+ \geq \boldsymbol{\mu}^\top [\mathbf{1}(T - t^*_{\tau_b(i+1;\boldsymbol{\mu})}) - \mathbf{b}_{\tau_b(i+1;\boldsymbol{\mu})}]^+$$

whenever $\boldsymbol{\mu}^\top \mathbf{b}_{\tau_b(i;\boldsymbol{\mu})} = \boldsymbol{\mu}^\top \mathbf{b}_{\tau_b(i+1;\boldsymbol{\mu})}$ for all $i \in [m-1]$. Intuitively, the contribution of a link i in state θ is the total time span in which flow using link i exits the system in state θ, i.e., $[T - t^*_i - b_{i,\theta}]^+$. We then take the sum over all states of the state-dependent timespan multiplied with the probability μ_θ that this state is realized. Like this, we obtain an ordering such that the corresponding Bayesian dynamic equilibrium maximizes the expected throughput among all such equilibria. In particular, we consider the equilibrium flow $\mathbf{f}$ with respect to the given belief $\boldsymbol{\mu}$, i.e., $\mathbf{f}$ is the flow defined in (5) with respect to the offsets $\boldsymbol{\mu}^\top \mathbf{b}_i$. For the capacity $\bar{\nu}_b(i;\boldsymbol{\mu})$ of the first i links used by the dynamic equilibrium flow $\mathbf{f}$ for a given belief $\boldsymbol{\mu}$, we simply write $\bar{\nu}_b(i)$, if $\boldsymbol{\mu}$ is clear from context. Further, we set $k_b \coloneqq \max\{i \in [m]_0 : \bar{\nu}_b(i) < u\}$.

For the makespan objective, there does not seem to be an explicit formula for the ordering that minimizes the expected makespan. However, all our results for the makespan objective hold regardless of which ordering in the case of identical expected offsets is selected and, thus, which equilibrium emerges.

We are interested in partitioning the set of beliefs Δ into subsets such that the ordering of the links by expected offsets is fixed within each subset. A naive bound on the number of these sets is $m!$, but this bound is not sufficient for us since our algorithms will iterate over these sets. To obtain a better bound, we will resort to the theory of *hyperplane arrangements*. To this end, for every pair of links $i, j \in [m]$ with $i < j$, we define by $H_{i,j} \coloneqq \{\boldsymbol{\mu} \in \mathbb{R}^d : \boldsymbol{\mu}^\top \mathbf{b}_i = \boldsymbol{\mu}^\top \mathbf{b}_j\}$ the (possibly empty) hyperplane containing all $\boldsymbol{\mu}$ such that the expected offsets on links i and j are the same. Then, $\mathcal{H} \coloneqq \{H_{i,j} : i, j \in [m] \text{ with } i < j\}$ is an arrangement of $|\mathcal{H}| = \frac{m(m-1)}{2}$ linear hyperplanes in $\mathbb{R}^d$ (where we allow that one or more of the hyperplanes are empty). The hyperplanes of the arrangement $\mathcal{H}$ partition Δ into a number of open regions whose closures are called the $(d-1)$-*cells* of the arrangement, i.e., the $(d-1)$-cells are the closures of the maximal connected subsets of $\Delta \setminus \bigcup_{i,j\in[m],i<j} H_{i,j}$. Every $(d-1)$-cell is a polyhedron in Δ. For $k \in \{0, \ldots, d-1\}$, a k-cell of the arrangement is a k-dimensional face of one of its $(d-1)$-cells. The following theorem of Buck [7] bounds the number of k-cells of a hyperplane arrangement.

▶ **Theorem 4** (Buck [7]). *For any hyperplane arrangement of n hyperplanes in $\mathbb{R}^{d-1}$ and any $k \in \{0, \dots, d-1\}$, the number of k-cells is at most*

$$\binom{n}{d-1-k} \sum_{i=0}^{k} \binom{n-(d-1)+k}{i}.$$

For a k-cell P of $\mathcal{H}$ with $k \in \{1, \dots, d-1\}$, we denote by P° its interior. We slightly overload notation by writing $P^\circ = P$ when P is a 0-cell, and we then also call P the interior of the 0-cell. Like this, the set of beliefs Δ is partitioned into interiors of k-cells with $k \in \{0, \dots, d-1\}$. In the interior P° of each k-cell P of $\mathcal{H}$ the ordering of the links $\tau_b(\cdot; \boldsymbol{\mu})$ remains the same and, therefore, every function $\boldsymbol{\mu} \mapsto t_i^*(\boldsymbol{\mu})$ is affine within P° by formula (6). From this, we obtain the following result.

▶ **Lemma 5.** *For every link $i \in [m]$, the function $t_i^* \colon \Delta \to \mathbb{R}_{\geq 0}$ is piecewise affine. In particular, the function is affine on the interior of every k-cell of $\mathcal{H}$ with $k \in \{0, \dots, d-1\}$.*

For a link i, $t_i^*(\boldsymbol{\mu})$ is the last point in time at which no flow enters link i yet, i.e., $f_i(t) > 0$ if and only if $t_i^* < t$. Given belief $\boldsymbol{\mu}$ and state $\theta \in \Theta$, we define for every link i its *first exit time* as $\omega_{i,\theta}(\boldsymbol{\mu}) \coloneqq t_i^*(\boldsymbol{\mu}) + b_{i,\theta}$. We obtain as an immediate corollary of Lemma 5 the following result.

▶ **Corollary 6.** *For every link $i \in [m]$ and every state $\theta \in \Theta$, the function $\omega_{i,\theta} \colon \Delta \to \mathbb{R}_{\geq 0}$ is piecewise affine. In particular, the function is affine on the interior of every k-cell of $\mathcal{H}$ with $k \in \{0, \dots, d-1\}$.*

We proceed to introduce another permutation of the links. For a given state $\theta \in \Theta$, let $\tau_\theta(\,\cdot\,; \boldsymbol{\mu}) : [m] \to [m]$ be a permutation of the links that orders them non-decreasingly with respect to their exit times in state θ, i.e., $\omega_{\tau_\theta(1;\boldsymbol{\mu}),\theta}(\boldsymbol{\mu}) \leq \dots \leq \omega_{\tau_\theta(m;\boldsymbol{\mu}),\theta}(\boldsymbol{\mu})$. For ease of notation, we write $\omega_{\tau_\theta(i;\boldsymbol{\mu})} \coloneqq \omega_{\tau_\theta(i;\boldsymbol{\mu}),\theta}(\boldsymbol{\mu})$. Since the inflow functions f_i of the dynamic equilibrium $\mathbf{f}$ are piecewise constant (by (5)), so are the outflows $f_{i,\theta}^-$ in state θ. We define $\bar{\nu}_\theta(i; \boldsymbol{\mu}) \coloneqq \sum_{j \in [i]} \nu_{\tau_\theta(j;\boldsymbol{\mu})}$ for $i \in [m]$ and set $\bar{\nu}_\theta(0; \boldsymbol{\mu}) \coloneqq 0$, $\tau_\theta(0; \boldsymbol{\mu}) \coloneqq 0$, $\tau_\theta(m+1; \boldsymbol{\mu}) \coloneqq m+1$, $\omega_{0,\theta}(\boldsymbol{\mu}) \coloneqq 0$, and $\omega_{m+1,\theta}(\boldsymbol{\mu}) \coloneqq \infty$. For $\theta \in \Theta$ and $\boldsymbol{\mu} \in \Delta$, let

$$\eta_\theta(\boldsymbol{\mu}) \coloneqq \min\bigl\{\max\bigl\{i \in [m]_0 : \omega_{\tau_\theta(i;\boldsymbol{\mu})} < T\bigr\}, \min\bigl\{i \in [m] : \bar{\nu}_\theta(i; \boldsymbol{\mu}) \geq u\bigr\}\bigr\}$$

be the number of links that contribute to the throughput in state θ. We proceed to compute explicit formulas for the outflows and the throughput.

▶ **Lemma 7.** *For a fixed state $\theta \in \Theta$ and a given belief $\boldsymbol{\mu} \in \Delta$, the total outflow at time t is*

$$\sum_{j \in [m]} f_{j,\theta}^-(t) = \min\bigl\{\bar{\nu}_\theta(i; \boldsymbol{\mu}), u\bigr\}$$

whenever $\omega_{\tau_\theta(i;\boldsymbol{\mu})} < t \leq \omega_{\tau_\theta(i+1;\boldsymbol{\mu})}$ for some $i \in [m]_0$. The throughput for state $\theta \in \Theta$ is given by the equation

$$F_\theta(\boldsymbol{\mu}) = uT + T\bigl[\bar{\nu}_\theta(\eta_\theta(\boldsymbol{\mu}); \boldsymbol{\mu}) - u\bigr]^- + \omega_{\tau_\theta(\eta_\theta(\boldsymbol{\mu}); \boldsymbol{\mu})}\bigl[\bar{\nu}_\theta(\eta_\theta(\boldsymbol{\mu}); \boldsymbol{\mu}) - u\bigr]^+ - \sum_{i \in [\eta_\theta(\boldsymbol{\mu})]} \nu_{\tau_\theta(i;\boldsymbol{\mu})} \omega_{\tau_\theta(i;\boldsymbol{\mu})}.$$

We continue toward dividing Δ into cells, such that $F_\theta(\boldsymbol{\mu})$ is affine on the interior of each cell. To this end, note that as long as the ordering $\tau_\theta(\,\cdot\,; \boldsymbol{\mu})$ remains unchanged, $\omega_{\tau_\theta(i;\boldsymbol{\mu})}$ is affine in $\boldsymbol{\mu}$ and the capacities $\nu_{\tau_\theta(i;\boldsymbol{\mu})}$ are constant in $\boldsymbol{\mu}$. As long as the same number of links has an exit time smaller or equal than T, the number $\eta_\theta(\boldsymbol{\mu})$ is constant as well. Therefore, in this case, $F_\theta(\boldsymbol{\mu})$ is affine in $\boldsymbol{\mu}$. We define new hyperplanes

$$H_{i,j,\theta} := \{\boldsymbol{\mu} \in \mathbb{R}^{|\Theta|} : \omega_{i,\theta}(\boldsymbol{\mu}) = \omega_{j,\theta}(\boldsymbol{\mu})\} \qquad \text{for } \theta \in \Theta \text{ and } i, j \in [m], i < j, \text{and}$$
$$H_{i,\theta,T} := \{\boldsymbol{\mu} \in \mathbb{R}^{|\Theta|} : \omega_{i,\theta}(\boldsymbol{\mu}) = T\} \qquad \text{for } \theta \in \Theta \text{ and } i \in [m].$$

For every interior P° of a k-cell of the hyperplane arrangement $\mathcal{H}$, consider the hyperplane arrangement

$$\mathcal{H}^* := \{H_{i,j,\theta} : \theta \in \Theta, i, j \in [m] \text{ with } i < j\} \cup \{H_{i,\theta,T} : \theta \in \Theta, i \in [m]\}.$$

This hyperplane arrangement further subdivides every interior P° of a k-cell of $\mathcal{H}$. The following lemma gives the main structural insights for the behavior of the functions F_θ on the cells of the hyperplane arrangements $\mathcal{H}$ and $\mathcal{H}^*$.

▶ **Lemma 8.** *For every state $\theta \in \Theta$, the function $F_\theta \colon \Delta \to \mathbb{R}_{\geq 0}$ is piecewise affine. In particular, let P° be the interior of a k-cell of $\mathcal{H}$ with $k \in \{0, \dots, d-1\}$ and let P^* be a $(d-1)$-cell of $\mathcal{H}^*$; then F_θ is affine on $P^\circ \cap P^*$.*

4 Additive PTAS for Throughput Maximization

In this subsection, we give an additive PTAS for computing the optimal throughput achievable by a public signaling scheme. For ease of notation, let us write $\mathbf{M} \in [0,1]^{d \times d}$ for the matrix that has the vectors $(\boldsymbol{\mu}_\sigma)_{\sigma \in \Sigma}$ as column vectors. Here, we assume that $|\Sigma| = d$. This is possible, as $|\Sigma| \leq d$ follows from Caratheodory's theorem and if $|\Sigma| < d$ holds, we can add artificial signals, that are never issued. We write $\mathcal{M}$ for the set of left-stochastic matrices whose rows sum to 1. The primal signaling problem is then rephrased as

$$\textsc{Opt} := \sup\left\{\sum_{\sigma \in \Sigma} \varphi_\sigma F(\boldsymbol{\mu}_\sigma) : \mathbf{M} \in \mathcal{M}, \boldsymbol{\varphi} \in [0,1]^d \text{ such that } \mathbf{M}\boldsymbol{\varphi} = \boldsymbol{\mu}^*\right\}, \qquad (P)$$

where F is the expected throughput as a function of the belief $\boldsymbol{\mu} \in \Delta$ and the supremum is taken both over $\mathbf{M} \in \mathcal{M}$ and $\boldsymbol{\varphi} \in [0,1]^d$. The main result of this section is the following.

▶ **Theorem 9.** *For every constant d and every $\varepsilon^* > 0$, there is a polynomial-time algorithm computing $p \in [\textsc{Opt} - \varepsilon^*, \textsc{Opt}]$.*

Instead of (approximately) solving the primal signaling problem (P) directly, our algorithm relies on the following Lagrangian dual.

▶ **Lemma 10.** *The dual signaling problem is*

$$d^* = \inf\left\{\mathbf{w}^\top \boldsymbol{\mu}^* : \mathbf{w} \in \mathbb{R}^d \textit{ with } \mathbf{w}^\top \boldsymbol{\mu} \geq F(\boldsymbol{\mu}) \textit{ for all } \boldsymbol{\mu} \in \Delta\right\}. \qquad (D)$$

In particular, weak duality holds.

We show that no duality gap exists, i.e., the optimal values for the primal and dual signaling problem are attained and coincide. The proof uses the definition and properties of the concave envelope of F.

▶ **Lemma 11.** *The optimal values of the primal signaling problem (P) and the dual signaling problem (D) are attained at finite values and coincide.*

The general idea for the additive PTAS is to apply the Ellipsoid method on the dual signaling problem and to use the equivalence of optimization and separation. Note that the feasible region of the Lagrange dual is always convex (see, e.g., [6], § 5.2). Formally, for a set $K \subseteq \mathbb{R}^n$ and any $\varepsilon > 0$, let $B(K, \varepsilon) \coloneqq \{\mathbf{x} \in \mathbb{R}^n : ||\mathbf{x} - \mathbf{y}||_2 \leq \varepsilon \text{ for some } \mathbf{y} \in K\}$ be the set of points that are within an ε-distance to a point in K. Furthermore, let $B(K, -\varepsilon) \coloneqq \{\mathbf{x} \in K : B(\{x\}, \varepsilon) \subseteq K\}$ be the set of points that are in the ε-interior of K. The following definitions are taken from Grötschel et al. ([18], Def. 2.1.10 and 2.1.13), and adapted to minimization instead of maximization.

▶ **Definition 12** (Weak optimization problem). *Given a vector $\mathbf{c} \in \mathbb{Q}^n$ and a rational number $\varepsilon > 0$, the* weak optimization problem *is to*

1. *compute a vector $\mathbf{y} \in \mathbb{Q}^n$ with $\mathbf{y} \in B(K, \varepsilon)$ and $\mathbf{c}^\top \mathbf{y} \leq \mathbf{c}^\top \mathbf{x} + \varepsilon$ for all $\mathbf{x} \in B(K, -\varepsilon)$, or*
2. *assert that $B(K, -\varepsilon)$ is empty.*

▶ **Definition 13** (Weak separation problem). *Given a vector $\mathbf{y} \in \mathbb{Q}^n$ and a rational number $\delta > 0$, the weak separation problem is to either*

1. *assert that $\mathbf{y} \in B(K, \delta)$, or*
2. *compute a vector $\mathbf{c} \in \mathbb{Q}^n$ with $||\mathbf{c}||_\infty = 1$ such that $\mathbf{c}^\top \mathbf{x} \leq \mathbf{c}^\top \mathbf{y} + \delta$ for all $\mathbf{x} \in B(K, -\delta)$.*

Roughly speaking, the equivalence of optimization and separation implies that a polynomial algorithm for the weak separation problem yields a polynomial algorithm for the weak optimization problem (see [18, Corollary 4.2.7] for a formal statement). Thus, in order to solve the weak optimization problem, it is enough to solve the weak separation problem.

We proceed to show that for the dual signaling problem, we can even solve the exact separation problem (a stronger version of the weak separation problem where $\delta = 0$) which is defined as follows. Given a vector $\mathbf{w} \in \mathbb{R}^d$, find $\boldsymbol{\mu} \in \Delta$ such that $\mathbf{w}^\top \boldsymbol{\mu} < F(\boldsymbol{\mu})$, or decide that no such $\boldsymbol{\mu} \in \Delta$ exists. For the proof, we use Buck's formula (Theorem 4) to show that the subdivision into sets $Q \cap Q^*$ where Q is a k-cell of $\mathcal{H}$ and Q^* is a k^*-cell of $\mathcal{H}^*$ for some $k, k^* \in \{0, \ldots, d-1\}$ produces only a polynomial number of sets. This allows us to iterate over these sets in polynomial time while checking whether there is a candidate for an extreme point of the function $F(\boldsymbol{\mu}) - \mathbf{w}^\top \boldsymbol{\mu}$ in the relative interior of $Q \cap Q^*$. The function is piecewise quadratic by Lemma 8. Thus, showing that it is differentiable in $Q \cap Q^*$ allows first-order conditions that reduce to a linear system.

▶ **Lemma 14.** *Given $\mathbf{w} \in \mathbb{R}^n$, we can compute $\boldsymbol{\mu} \in \Delta$ in polynomial time such that $\mathbf{w}^\top \boldsymbol{\mu} < F(\boldsymbol{\mu})$, or decide that no such $\boldsymbol{\mu} \in \Delta$ exists.*

Immediately following, the weak separation problem for the dual signaling problem is also solvable in polynomial time.

▶ **Corollary 15.** *For any $\delta > 0$, the weak separation problem for the dual signaling problem can be solved in polynomial time.*

We are now ready to prove the main theorem of this section (Theorem 9). The general idea for the proof is to use Corollary 15 and the Ellipsoid method. However, to do so, we have to show that we can fit the dual feasible region into a ball with a polynomially bounded diameter. To this end, we show an upper bound on $||\mathbf{w}||_\infty$ of the dual optimal vector $\mathbf{w}$ based on the supergradient of F. Further calculations involving the approximation error of solutions $\mathbf{v} \in B(K, -\varepsilon)$ then yield the result.

5 Multiplicative FPTAS for Throughput Maximization

While the additive PTAS devised in the last subsection approximates the optimal throughput up to an arbitrary constant, it does not yield the corresponding approximate optimal signals. In this section, we devise a multiplicative FPTAS that also yields the corresponding signals, i.e., we show the following theorem.

▶ **Theorem 16.** *For every constant d and every $\varepsilon^* > 0$, there exists a fully polynomial-time algorithm for computing a signaling scheme, such that for its induced throughput* Alg*, it holds that* $\text{Alg} \geq (1-\varepsilon^*)\text{Opt}$.

As this result trivially holds, whenever $\text{Opt} = 0$, we assume for the remaining part of this section that $\text{Opt} > 0$. The following lemma shows that this assumption is in fact equivalent to the smallest offset $b_{i,\theta}$ over all $i \in [m]$ and $\theta \in \Theta$ being strictly smaller than the time horizon T.

▶ **Lemma 17.** *We have* $\text{Opt} > 0$ *if and only if* $T - \min\{b_{i,\theta} : i \in [m], \theta \in \Theta\} > 0$.

The next lemma will be used to define the algorithm that achieves the approximation guarantee of $(1-\varepsilon^*)$ and to bound its running time. For the proof, we bound Opt in terms of the throughput achieved by full information revelation.

▶ **Lemma 18.** *Let* $\text{Opt} > 0$. *For any* $0 < \varepsilon < 1$ *and any* $0 < \delta \leq d$, *there exists* $\kappa \in \mathbb{N}_{>0}$ *such that* $(1-\varepsilon)^{\kappa-1}|\Theta|Tu \leq \delta\text{Opt}$ *and* κ *is polynomially bounded in the input size.*

Let $0 < \varepsilon < 1$ and $0 < \delta \leq d$ be two arbitrary but fixed values. Further, let κ be defined as in Lemma 18. We proceed to define an algorithm towards proving Theorem 16. The main building block of the algorithm is to find a piecewise convex underestimator function $F_{\varepsilon,\kappa}\colon \Delta) \to \mathbb{R}_{\geq 0}$ of the total expected throughput function F. To this end, we define the following function $h_{\varepsilon,\kappa} : [0,1] \to [0,1]$ that rounds numbers to the next power of $(1-\varepsilon)$, or to 0 if the number is too small:

$$h_{\varepsilon,\kappa}(x) \coloneqq \begin{cases} 0 & \text{if } x < (1-\varepsilon)^{\kappa-1}, \\ \max\{(1-\varepsilon)^{k-1} : (1-\varepsilon)^{k-1} \leq x, k \in [\kappa]\} & \text{else.} \end{cases}$$

We define an under-estimator function $F_{\varepsilon,\kappa} : \Delta \to \mathbb{R}_{\geq 0}$ of the total expected throughput function by

$$F_{\varepsilon,\kappa}(\boldsymbol{\mu}) \coloneqq \sum_{\theta\in\Theta} h_{\varepsilon,\kappa}(\mu_\theta)F_\theta(\boldsymbol{\mu}). \tag{7}$$

To define the regions where the under-estimator is convex, we proceed to discretize Δ by a non-uniform ε-net. For this, we define for every $\theta \in \Theta$ and every $j \in [\kappa]$ the hyperplanes

$$L_{\theta,j} \coloneqq \{\boldsymbol{\mu} : \mu_\theta = (1-\varepsilon)^{j-1}\} \qquad \text{and} \quad L_{\theta,0} \coloneqq \{\boldsymbol{\mu} : \mu_\theta = 0\}.$$

We denote by $\mathcal{L}$ the union of the arrangement of hyperplanes $\mathcal{H}$, $\mathcal{H}^*$, and the ε-net, i.e.,

$$\mathcal{L} \coloneqq \mathcal{H} \cup \mathcal{H}^* \cup \big\{L_{\theta,j} : \theta \in \Theta, j \in [\kappa]_0\big\}.$$

The set $\mathcal{L}$ again defines an arrangement of hyperplanes in Δ with

$$\begin{aligned} |\mathcal{L}| &\leq |\mathcal{H}| + |\mathcal{H}^*| + \big|\big\{L_{\theta,j} : \theta \in \Theta, j \in [\kappa]_0\big\}\big| \\ &\leq \frac{m(m-1)}{2} + \frac{dm(m-1)}{2} + dm + d(\kappa+1) \end{aligned}$$

many hyperplanes. We write Δ_ε for the set of 0-cells that are determined by $\mathcal{L}$, i.e., the set of points in Δ in which $d-1$ many pairwise distinct hyperplanes of $\mathcal{L}$ intersect. More formally,

$$\Delta_\varepsilon \coloneqq \Big\{\boldsymbol{\mu} \in \Delta(\Theta) : \{\boldsymbol{\mu}\} = \bigcap\nolimits_{i \in [|\Theta|-1]} H_i \text{ with } H_i \in \mathcal{L} \text{ and } H_i \neq H_j \text{ for } i \neq j\Big\}.$$

Note, that the number of 0-cells $\ell \coloneqq |\Delta_\varepsilon|$ in $\mathcal{L}$ is polynomially bounded by the input size as

$$\ell \coloneqq |\Delta_\varepsilon| \leq \binom{|\mathcal{L}|}{d-1} = \binom{\frac{m(m-1)}{2} + \frac{|\Theta|m(m-1)}{2} + |\Theta|m + d(\kappa+1)}{|\Theta|-1}.$$

Intuitively, we restrict our algorithm to induce only posterior beliefs in $\Delta_\varepsilon(\Theta) \coloneqq \{\tilde{\boldsymbol{\mu}}_{\sigma_1}, \dots, \tilde{\boldsymbol{\mu}}_{\sigma_\ell}\}$ by a set of signals $\Sigma_\varepsilon \coloneqq \{\sigma_1, \dots, \sigma_\ell\}$. In particular, the algorithm solves the linear program

$$\textsc{Alg} \coloneqq \max\left\{\sum_{j\in[\ell]} \varphi_{\sigma_j} F(\tilde{\boldsymbol{\mu}}_{\sigma_j}) : \tilde{\boldsymbol{\mu}}_{\sigma_j} \in \Delta_\varepsilon(\Theta), \varphi_{\sigma_j} \in [0,1] \text{ for all } j \in [\ell], \sum_{j\in[\ell]} \varphi_{\sigma_j} \tilde{\boldsymbol{\mu}}_{\sigma_j} = \boldsymbol{\mu}^*\right\} \quad (8)$$

and returns a signaling scheme that induces the posterior beliefs $\tilde{\boldsymbol{\mu}}_{\sigma_j}$ that appear with positive probability $\varphi_{\sigma_j} > 0$ in (8) for all $j \in [\ell]$. Note that if the posterior beliefs $(\tilde{\boldsymbol{\mu}}_{\sigma_j})_{j\in[\ell]}$ and the appropriate coefficients $(\varphi_{\sigma_j})_{j\in[\ell]}$ are given, the signaling scheme can be recovered in polynomial time (cf. [20]).

We work towards proving that the thus defined algorithm yields the claimed approximation guarantee of $(1-\varepsilon^*)$. To this end, we first show that the underestimator function $F_{\varepsilon,\kappa}$ is convex on every k-cell of $\mathcal{L}$ for all $k \in \{0, 1, \dots, d-1\}$. The proof uses that on the interior of each k-cell of $\mathcal{L}$, the function is a linear combination of affine functions and the values on the border of the cell cannot be smaller than the continuous expansion of the affine function.

▶ **Lemma 19.** *For all $0 < \varepsilon < 1$ and $\kappa \in \mathbb{N}$, the function $F_{\varepsilon,\kappa}$ is convex on every k-cell of $\mathcal{L}$ for all $k \in \{0, 1, \dots, d-1\}$. In particular, $F_{\varepsilon,\kappa}$ is affine on the interior of every k-cell of $\mathcal{L}$.*

Next, we bound the under-estimator function $F_{\varepsilon,\kappa}$ both from below and above.

▶ **Lemma 20.** *For all $0 < \varepsilon < 1, \kappa \in \mathbb{N}$, and $\boldsymbol{\mu} \in \Delta$, we have*

$$(1-\varepsilon)F(\boldsymbol{\mu}) - d(1-\varepsilon)^\kappa F_{\max} \leq F_{\varepsilon,\kappa}(\boldsymbol{\mu}) \leq F(\boldsymbol{\mu}),$$

where $F_{\max} \coloneqq \max_{\theta\in\Theta} \sup_{\boldsymbol{\mu}\in\Delta} F_\theta(\mu)$.

With these lemmas at hand, we are ready to prove the main result of this section (Theorem 16). For the proof, we use that optimizing over the piecewise convex under-estimator $F_{\varepsilon,\kappa}$ instead of the original function F causes only a multiplicative error of $(1-\varepsilon)$ in the $(1-\varepsilon)^{\kappa-1}$-interior of Δ and an additional additive error close to the boundary of Δ. Since the under-estimator is convex on the interior of every $(d-1)$-cell of $\mathcal{L}$, every optimal convex decomposition of the prior for the convex under-estimator only uses the 0-cells of $\mathcal{L}$. Bounding their number by a polynomial of the encoding length of the input, we then obtain the result.

6 Full Information Revelation for Makespan Minimization

In this section, we consider the makespan objective. As the main result of this section, we show that full information revelation is optimal, i.e., it is optimal to choose $\Sigma = \Theta$ and have $\varphi_{\theta,\sigma} = \mu^*_\theta$ if $\theta = \sigma$ and $\varphi_{\theta,\sigma} = 0$ otherwise.

▶ **Theorem 21.** *For the makespan objective, full information revelation is an optimal signaling scheme.*

To prove Theorem 21, we consider dynamic equilibria with different *deterministic* offsets. We first note that among all dynamic equilibria, the cost of every edge is unique (e.g., [10, 24]), and, hence, each edge's queue length is unique. Thus, for every dynamic equilibrium, the queue lengths are determined by Lemma 3 and are non-decreasing in t. Therefore, a particle entering the system at time $t = T$ has the latest expected exit time C. This implies the following easier formula for the makespan.

▶ **Lemma 22.** *Let* $\mathbf{f}$ *be the dynamic equilibrium from* (5)*. Then it holds,*

$$M_T(\mathbf{f}) = \max\big\{C_i(T) : i \in S(T)\big\}.$$

Moreover, if the perceived and the true deterministic offsets $\mathbf{b}$ coincide, we have $M_T(\mathbf{f}) = T + b_{k+1}$. We proceed to investigate how changing the perceived offsets influences the makespan. Formally, let $\mathbf{b} = (b_i)_{i\in[m]}$ be the offsets of a given instance. Then, we denote by $\mathbf{f}(\mathbf{b})$ the dynamic equilibrium as defined by (5). For flow $\mathbf{f}$ we denote its support at time t by $S(t;\mathbf{f})$. Additionally, we consider another vector of (arbitrary) offsets $\mathbf{b}' = (b'_i)_{i\in[m]}$ with $b'_i \geq 0$ for all $i \in [m]$. (For example, the offsets $\mathbf{b}'$ could be the offsets expected by the particles given a certain belief $\boldsymbol{\mu}$.) Assume that particles behave according to $\mathbf{b}'$ rather than $\mathbf{b}$. Then, the dynamic equilibrium and, thus, the makespan changes if $\mathbf{b}'$ is changed. We denote by

$$\mathcal{F}(\mathbf{b}') := \big\{\mathbf{f} : \mathbf{f} \text{ is dynamic equilibrium with respect to the offsets } \mathbf{b}'\big\}$$

all dynamic equilibria with respect to the offsets $\mathbf{b}'$ and emphasize that these equilibria are not unique. If the flow $\mathbf{f}' \in \mathcal{F}(\mathbf{b}')$ emerges in the original instance (i.e., the instance with offsets $\mathbf{b}$) the flow particles experience the exit times

$$C_i(t;\mathbf{f}') := t + \frac{z_i(t;\mathbf{f}')}{\nu_i} + b_i,$$

where $z_i(t;\mathbf{f}')$ are the queue lengths determined by the dynamic equilibrium $\mathbf{f}'$. Then, we denote by

$$\begin{aligned}\overline{M}(\mathbf{b}') &:= \sup\big\{M_T(\mathbf{f}') : \mathbf{f}' \in \mathcal{F}(\mathbf{b}')\big\}\\ &= \sup_{\mathbf{f}'\in\mathcal{F}(\mathbf{b}')} \sup\big\{C_i(t;\mathbf{f}') : t \in [0,T], i \in S(T;\mathbf{f}')\big\}\end{aligned}$$

the (worst-case) makespan for given offsets $\mathbf{b}'$. In this setting, we are interested in finding the best possible $\mathbf{b}' \geq \mathbf{0}$ that minimizes the makespan, i.e., we want to compute

$$\inf\big\{\overline{M}(\mathbf{b}') : \mathbf{b}' \geq \mathbf{0}\big\}. \tag{9}$$

The function $\overline{M}(\mathbf{b}')$ is in general not continuous, as illustrated in Example 2. Thus, it is not clear if a minimum is attained.

Theorem 21 is then proven by showing that the infimum in (9) is actually attained for the original travel times $\mathbf{b}$. Due to space constraints, we defer the whole derivation of this result to the full version.

References

1. Daron Acemoglu, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar. Informational Braess' paradox: The effect of information on traffic congestion. *Oper. Res.*, 66(4):893–917, 2018.
2. David Avis and Komei Fukuda. Reverse search for enumeration. *Discret. Appl. Math.*, 65:21–46, 1996.
3. Dirk Bergemann and Stephen Morris. Information design: A unified perspective. *J. Econ. Lit.*, 57:44–59, 2019.
4. Umang Bhaskar, Yu Cheng, Young Kun Ko, and Chaitanya Swamy. Hardness results for signaling in Bayesian zero-sum and network routing games. In *Proc. 17th Conf. Econ. Comput. (EC)*, pages 479–496, 2016.
5. Umang Bhaskar, Lisa Fleischer, and Elliot Anshelevich. A Stackelberg strategy for routing flow over time. *Games Econ. Behav.*, 92:232–247, 2015.
6. Stephen Boyd and Lieven Vandenberghe. *Convex Optimization.* Cambridge University Press, 2004.
7. Robert C. Buck. Partition of space. *Amer. Math. Monthly*, 50(9):541–544, 1943.
8. Matteo Castiglioni, Andrea Celli, Alberto Marchesi, and Nicola Gatti. Signaling in Bayesian network congestion games: the subtle power of symmetry. In *Proc. 35th Conf. Artif. Intell. (AAAI)*, pages 5252–5259, 2021.
9. Yu Cheng, Ho Yee Cheung, Shaddin Dughmi, Ehsan Emamjomeh-Zadeh, Li Han, and Shang-Hua Teng. Mixture selection, mechanism design, and signaling. In Venkatesan Guruswami, editor, *Proc. 56th Symp. Found. Comput. Sci. (FOCS)*, pages 1426–1445, 2015.
10. Roberto Cominetti, José R. Correa, and Omar Larré. Dynamic equilibria in fluid queueing networks. *Oper. Res.*, 63(1):21–34, 2015.
11. Roberto Cominetti, José R. Correa, and Neil Olver. Long-term behavior of dynamic equilibria in fluid queuing networks. *Oper. Res.*, 70(1):516–526, 2022.
12. José R. Correa, Andrés Cristi, and Tim Oosterwijk. On the price of anarchy for flows over time. *Math. Oper. Res.*, 47(2):1394–1411, 2022.
13. Sanmay Das, Emir Kamenica, and Renee Mirka. Reducing congestion through information design. In *Proc. 55th Annual Allerton Conference on Communication, Control, and Computing*, 2017.
14. Shaddin Dughmi. On the hardness of signaling. In *Proc. 55th Symp. Found. Comput. Sci. (FOCS)*, pages 354–363, 2014.
15. Lukas Graf, Tobias Harks, Kostas Kollias, and Michael Markl. Machine-learned prediction equilibrium for dynamic traffic assignment. In *Proc. 36th Conf. Artif. Intell. (AAAI)*, pages 5059–5067, 2022.
16. Lukas Graf, Tobias Harks, and Leon Sering. Dynamic flows with adaptive route choice. *Math. Program.*, 183(1):309–335, 2020.
17. Svenja M. Griesbach, Martin Hoefer, Max Klimm, and Tim Koglin. Public signals in network congestion games. In David M. Pennock, Ilya Segal, and Sven Seuken, editors, *Proc. 23rd Conf. Econ. Comput. (EC)*, page 736, 2022.
18. Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization.* Springer, Berlin, Heidelberg, Germany, 1988.
19. Jonas Israel and Leon Sering. The impact of spillback on the price of anarchy for flows over time. In Tobias Harks and Max Klimm, editors, *Proc. 13th Symp. Algorithmic Game Theory (SAGT)*, pages 114–129. Springer, 2020.
20. Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *Amer. Econ. Rev.*, 101(6):2590–2615, 2011.
21. Ronald Koch and Martin Skutella. Nash equilibria and the price of anarchy for flows over time. *Theory Comput. Syst.*, 49(1):71–97, 2011.
22. Olivier Massicot and Cedric Langbort. Public signals and persuasion for road network congestion games under vagaries. *IFAC-PapersOnLine*, 51(34):124–130, 2019.

23 James Nachbar and Haifeng Xu. The power of signaling and its intrinsic connection to the price of anarchy. In *Proc. 3rd Intl. Conf. Distrib. Artif. Intell. (DAI)*, pages 1–20, 2021.

24 Neil Olver, Leon Sering, and Laura Vargas Koch. Continuity, uniqueness and long-term behavior of Nash flows over time. In *Proc. 62nd Symp. Found. Comput. Sci. (FOCS)*, pages 851–860, 2021.

25 Tim Oosterwijk, Daniel Schmand, and Marc Schröder. Bicriteria Nash flows over time. In Kristoffer Arnsfelt Hansen, Tracy Xiao Liu, and Azarakhsh Malekian, editors, *Proc. 18th Conf. Web and Internet Econ. (WINE)*, page 368, 2022.

26 Leon Sering and Martin Skutella. Multi-source multi-sink Nash flows over time. In Ralf Borndörfer and Sabine Storandt, editors, *Proc. 18th Workshop on Algorithmic Approaches for Transportation Modelling (ATMOS)*, volume 65, pages 12:1–12:20, 2018.

27 Leon Sering and Laura Vargas Koch. Nash flows over time with spillback. In Timothy M. Chan, editor, *Proc. 30th Symp. Discret. Algorithms (SODA)*, pages 935–945, 2019.

28 Shoshana Vasserman, Michal Feldman, and Avinatan Hassidim. Implementing the wisdom of Waze. In *Proc. 24th Int. Joint Conf. Artif. Intell. (IJCAI)*, pages 660–660, 2015.

29 William S. Vickrey. Congestion theory and transport investment. *Am. Econ. Rev.*, 59(2):161–179, 1969.

30 Manxi Wu, Saurabh Amin, and Asuman Ozdaglar. Value of information in Bayesian routing games. *Oper. Res.*, 69(1):148–163, 2021.

31 Haifeng Xu. On the tractability of public persuasion with no externalities. In *Proc. 31st Symp. Discret. Algorithms (SODA)*, pages 2708–2727, 2020.

32 Chenghan Zhou, Thanh H. Nguyen, and Haifeng Xu. Algorithmic information design in multi-player games: Possibility and limits in singleton congestion. In *Proc. 23rd Conf. Econ. Comput. (EC)*, page 869, 2022.