

Network Pricing: How to Induce Optimal Flows Under Strategic Link Operators

JOSÉ CORREA, Universidad de Chile

CRISTÓBAL GUZMÁN, Pontificia Universidad Católica de Chile

THANASIS LIANEAS, National Technical University of Athens

EVDOKIA NIKOLOVA, University of Texas at Austin

MARC SCHRÖDER, RWTH Aachen University.

Network pricing games provide a framework for modeling real-world settings with two types of strategic agents: owners (operators) of the network and users of the network. Owners of the network post a price for usage of the link they own so as to attract users and maximize profit; users of the network select routes based on price and level of use by other users. We point out that an equilibrium in these games may not exist, may not be unique and may induce an arbitrarily inefficient network performance.

Our main result is to observe that a simple regulation on the network owners market solves all three issues above. Specifically, if an authority could set appropriate caps (upper bounds) on the tolls (prices) operators can charge, then: the game among the link operators has a unique and strong Nash equilibrium and the users' game results in a Wardrop equilibrium that achieves the optimal total delay. We call any price vector with these properties a great set of tolls. As a secondary objective, we want to compute great tolls that minimize total users' payments and we provide a linear program that does this. We obtain multiplicative approximation results compared to the optimal total users' payments for arbitrary networks with polynomial latencies of bounded degree, while in the single-commodity case we obtain a bound that only depends on the topology of the network. Lastly, we show how the same mechanism of setting appropriate caps on the allowable prices extends to the model of elastic demands.

CCS Concepts: •**Theory of computation** → **Algorithmic game theory and mechanism design; Quality of equilibria; Network games; Exact and approximate computation of equilibria;**

1 INTRODUCTION

Network pricing games provide a framework for modeling real-world settings with two types of strategic agents: users of the network and owners (operators) of the network. Owners of the network post a price for usage of the links they own, while users of the network select routes based on price and level of use by other users. A landmark example of this type of interaction, that sets the ground in which we describe our work, occurs in transportation networks. Here, link owners

This work was partially supported by CONICYT through grants ISCI FB0816, PCI PII 20150140, FONDECYT 11160939, and FONDECYT 1160079; by NSF through grants CCF-1216103, CCF-1331863, CCF-1350823 and CCF-1733832; and by ICM through the Millennium Nucleus Information and Coordination in Networks RC130003, and the Millennium Institute for Research in Market Imperfections and Public Policy IS130002. Part of the work was done while a subset of the authors was visiting the Simons Institute for the Theory of Computing, Berkeley, CA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ACM EC'18, June 18–22, 2018, Ithaca, NY, USA.

© 2018 Association of Computing Machinery. ACM ISBN 978-1-4503-4529-3/18/06..\$15.00.

<https://doi.org/10.1145/3219166.3219190>

are road operators and may toll the usage of their road. Users are travelers that seek to minimize their travel time plus tolls.

The challenge in these games is that there are two levels of competition: one, among the owners to attract users to their link so as to maximize profit; and second, among users of the network to select routes that are cheap yet not too congested. In the absence of self-interested link owners, these games reduce to the well-studied network congestion games—a widely accepted and practically useful model for selfish resource allocation in transportation and communication networks (see, e.g., [4, 13, 33] and references therein). In congestion games, self-interested users travel in a network from their origin to their destination with the goal of minimizing their own travel cost. The natural solution concept of the game is the so-called Wardrop equilibrium flow, under which all users route along shortest paths, given the strategic choices of other users.

Since selfish behavior usually drives systems to socially inefficient situations, a central authority is typically interested in optimizing the social welfare for the network users—a task that, in the absence of self-interested link owners, can be implemented by appropriate *marginal tolls* [4], which simply charge each user the negative externality she imposes on the system. However, the presence of link owners, who set tolls so as to maximize their profits, is at odds with user welfare. For instance, under marginal tolls, some operators may want to increase their toll in order to make a higher profit. In this regard, an equilibrium for the link owners is a toll vector, such that, a change in a single toll does not increase the profit of the corresponding toll operator (under the implied user Wardrop equilibrium flow).

In this more complex game-theoretic environment, (i) an equilibrium may fail to exist; (ii) it might not be unique and; (iii) the network performance at equilibrium can be arbitrarily inefficient. Specifically, the only case where an equilibrium is proved to exist is in parallel link networks with affine latency functions [1, 20, 21]. Additionally, equilibria might not be unique [2, 20], implying that we might not be able to predict the outcome of the game. Further, the network's performance at the Wardrop equilibrium flow under a Nash equilibrium toll vector can be arbitrarily degraded, which implies that the network can behave arbitrarily worse than the case where tolls (prices) are completely absent.

In the face of these challenges, we set out to find ways to mitigate the effect of selfish toll operator behavior. We introduce competition regulation by allowing a regulator to set specific price caps on the toll values that each toll operator can set on her link. Different price caps for different operators is consistent with the practice in some privately operated networks of highways. For example, in Santiago de Chile there are currently 12 different operators who set tolls on different urban highways, as shown in Fig. 1. The current regulation sets a price cap that is unique to each highway and the toll operators are allowed to set tolls below the caps [19]. As it turns out, introducing such regulation can resolve all of the above issues, as there are caps for which the game has a unique Nash equilibrium, under which the Wardrop equilibrium is optimal.

Contribution

We first study the network pricing game in its original form, i.e. without regulation, providing examples regarding the nonexistence of Nash equilibria and the inefficiency of the resulting Wardrop equilibrium flow. The first example, Example 3.1, shows that Nash equilibria may not exist in the case of extension-parallel networks¹ with affine latencies. Interestingly, the source of the problem is twofold. On the one hand, the best-response mapping cannot be bounded a priori: there are instances of the game in which the best response of operators is to increase their toll and thus,

¹Extension-parallel networks are recursively defined as a single link or two extension-parallel networks connected in parallel or a single link connected in series with an extension-parallel network [16].



Fig. 1. Privately operated highways in Santiago de Chile. Each highway has its own toll cap.

iterated best responses lead to unbounded tolls. On the other hand, and quite surprisingly, the best-response mapping may turn out to be discontinuous so that iterated best responses actually cycle. From the latter it turns out that even if we impose a priori caps on the tolls, equilibria may fail to exist. The second example (Example 3.2) provides an instance where a Nash equilibrium exists yet the network performance, quantified by total users' delay of the Wardrop flow, is arbitrarily worse than optimal.

Since the unregulated network pricing game behaves so poorly, in Section 4 we set out to study mechanisms to improve it. In particular, we investigate the simple regulatory policy of setting upper bounds (caps) on the toll values that each operator is allowed to set. Note that there is a tension between the toll charged by the operator and the amount of flow she will get. It is plausible for a toll operator to gain from decreasing her toll since her link will attract more flow, which may result in an overall higher profit. For large enough caps it indeed happens that it is optimal for the operator to set a toll below the cap. Our main result, Theorem 4.4, shows that using such a vector as toll caps leads to every operator charging precisely the cap. Moreover, we show this equilibrium is robust to coalitions, a concept known as *strong Nash equilibrium*. A direct corollary of this result is that if the caps are set to be the marginal tolls, then the unique Nash equilibrium is the set of marginal tolls itself, which is known to induce the optimal flow as a Wardrop equilibrium. A surprising fact is that even though the flow on a given link is a decreasing function of the toll on that link (as we show in Lemmas 4.1 and 4.2), the profit of the toll operator as a function of the toll she charges is an increasing function up to a certain point (Lemma 4.3), which is a key tool to prove Theorem 4.4. Also, we give an example to show that our results do not extend to the setting in which a regulator owns multiple links of the network.

Inspired by Theorem 4.4 we study *great tolls* in more detail. These are toll vectors that induce the optimal flow as a Wardrop equilibrium and when set as caps are themselves the unique Nash equilibrium. As great tolls are guaranteed to exist and need not be unique, a basic question is that of finding the great tolls minimizing the users' payments. Note that the users' total delay under great tolls is fixed, since all great tolls induce the optimal flow; thus, our question is equivalent to that of finding great tolls that minimize the total (delays plus payments) users' cost. By the results in Section 4, any vector that induces the optimal flow and is upper bounded by the marginal toll vector, is a great set of tolls, yet simple examples show that these are not the only ones. While a characterization of great tolls seems to be out of reach, by restricting to those which are upper bounded by marginal tolls we formulate a linear program for optimizing the total users' cost over this restricted set of great tolls. Section 5 is devoted to finding bounds on the performance of such great tolls in terms of the induced total users' cost. To this end, we use as a benchmark the

minimum payment tolls, defined as those that while inducing optimal flows, minimize the total users' cost. Minimum payment tolls can also be computed by a linear program.

We prove that the ratio of these linear programs can be bounded in two fundamentally different ways, and these bounds are tight. First, we prove a functional bound (Theorem 5.2) for instances where latency functions satisfy a bounded relative growth condition, which for example gives a $(d + 1)$ -approximation on the total cost incurred by users, when latencies are polynomial functions of degree at most d . Notice that all the results so far hold for multicommodity instances. Second, we prove a structural bound (Theorem 5.5) for single-commodity instances that only depends on the structure of the network but not on the underlying latency functions. As a corollary, we get that for series-parallel networks the previous ratio is actually one and thus our linear programming approach is optimal.

To test the robustness of our results, we finally consider the setting with elastic traffic demands. Here users have a valuation for traveling through the network and may opt out from traveling if the travel cost exceeds their valuation. We show that, similar to the fixed demand model, any toll vector that induces the optimal flow and is at most equal to the marginal toll vector is a great set of tolls (Theorem 6.1). Moreover, even though the consumer surplus under great tolls can be arbitrary worse than the optimal surplus, we prove that it is at least as high as the optimal surplus if each user had half the original valuation (Theorem 6.5).

Related Work

Acemoglu and Ozdaglar [1] introduced a model of price competition between link operators where each user has some fixed reservation value for travel. They show that increasing competition among operators from a monopoly to an oligopoly may cause a reduction in efficiency, measured as the difference between users' willingness to pay and delay costs, and provide a (tight) bound on efficiency in pure strategy equilibria. In a follow up work, Acemoglu and Ozdaglar [2] generalized the above study to slightly more general topologies in which parallel paths with multiple links may replace the parallel links. They showed that even this slight generalization can make the game arbitrarily inefficient, where the efficiency is measured as mentioned above.

Hayrapetyan et al. [21] considered instances on parallel links where the demand to be routed is elastic and decreases in a concave way as the cost for using the network increases. The social cost in that work is the sum of the players' profits plus a term that represents the utility gathered by the traffic that gets routed. For that game they showed that in a network with parallel links and linear latencies, there is always a pure Nash equilibrium with the price of anarchy, i.e. the measure for inefficiency, being bounded by a constant factor even when the latency functions are relaxed to be convex. For the case where latencies have zero value under zero flow, they improve the constant in the above bound. Following that work, Ozdaglar [29] studied the same model and managed to prove tight bounds on the efficiency of that game. Musacchio [26] and Musacchio and Wu [27] rederived and generalized those (upper) bounds for the case of series-parallel networks via a connection to electrical circuits; see the survey by Ozdaglar and Srikant [30] for further discussion.

With respect to the congestion model itself, the most closely related work to ours is that of Harks et al. [20]. There, a regulator is able to set a unique price cap for all toll operators. As it turns out this restricts the regulator enough so that the induced network performance is not always optimal. For two-link parallel networks this reduction in performance is characterized for different classes of latency function.

Our price competition model corresponds to Bertrand competition in a network setting [15, Chapter 6]. Under this setting Chawla et al. [11] addressed questions regarding the price of anarchy and price of stability with respect to two objectives: the social welfare of all the players (users

and sellers), and the total profit obtained by all the sellers. The main differences with our work are that only capacity-based congestion is considered,² and no regulation is imposed on the game. Their results show a dependence of the price of anarchy/stability on the number of monopolistic links, namely the links whose removal disconnects an origin-destination pair. In contrast, our model does not suffer from monopolies: setting caps on prices prevents monopolistic links from charging arbitrarily large prices. Following the same model and focusing on the social welfare of the consumers as the objective, Chawla and Niu [10] extended the results of Chawla et al. [11].

An interesting and related model is raised by Anshelevich and Sekar [3]. In that work, the edges of a network correspond to goods, and each edge is owned by a different profit maximizing seller. In the first stage of the game, sellers set prices for the use of their edges and have production costs depending on the level of use. In the second stage, the users of the network, i.e., the buyers, choose origin to destination paths so as each of them maximizes her utility minus the payments to the sellers. The main difference with the model presented here (in particular, the extension to the elastic demand case in Section 6) is that users impose an externality on the sellers, via the production costs, and not on the other users, as is the case in our model. This crucially affects the equilibrium pricing and the social welfare (e.g., for single commodity networks, if no monopolies are present, there is always an optimal equilibrium, which is not the case in our model).

In other related work, Papadimitriou and Valiant [31] consider the case where the routing is no longer selfish, but is controlled by the edges of the network, and each edge either minimizes its average latency, or announces a suitable price to its neighbors in order to maximize its profit. Caragiannis et al. [8] consider a model of buyers and sellers of a similar product, that under some reformulation can be seen as a variant of the parallel links model of Acemoglu and Ozdaglar [1] with heterogeneous buyers but constant latencies. Instead of minimizing traffic costs, maximizing the profit from tolls, has been considered in the past [7, 9, 24]. There, a central authority/unique owner has control of all the tollable edges, yet, more importantly, the edge costs are constants rather than flow dependent.

The study of network congestion games where a central operator is allowed to charge tolls in order to improve efficiency has a long history, starting with Beckmann et al. [4]. Cole et al. [12] and Fleischer [17] provided upper bounds on tolls that induce the optimal flow as an equilibrium, and Dial [14] considered the objective of minimizing users' payments among optimal flow-inducing tolls. The study of network users games where for each link a (potentially adversarially chosen) upper bound on the toll is present was first considered by Bonifaci et al. [6] and later by Jelinek et al. [22] and Fotakis et al. [18]. Results in these papers show that when upper bounds are present, optimality cannot in general be achieved, yet, on the positive side, algorithms are proposed under rather restrictive settings with provable guarantees regarding the efficiency of the network.

2 PRELIMINARIES

We study a network pricing game, where nonatomic players, which we call *users*, selfishly minimize their cost (delays plus tolls) across a network; on top of this, each network link is operated by a different selfish agent which maximizes profit by charging tolls on users traversing her link.

The Network Users' Game: Selfish Routing

Let $G = (V, E)$ be a network, with V the set of nodes, and E the set of directed edges/links of the network. We consider a multi-commodity flow instance, described by origin-destination node pairs $\{(o^k, d^k)\}_{k \in K}$, for a finite set of commodities K . For each commodity k , $r^k > 0$ units of demand

²Capacity-based congestion corresponds to latency functions which are identically zero until capacity is reached, and then jump to infinity.

need to be routed from o^k to d^k . For each link $e \in E$, there is a latency function $\ell_e : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, that represents the delay experienced by users traversing this link, as a function of the total flow on the link; we will assume this function to be strictly increasing, convex and smooth.³

Paths and Flows. For each commodity $k \in K$, let \mathcal{P}^k denote the set of $o^k - d^k$ paths and let $\mathcal{P} = \cup_k \mathcal{P}^k$ be the union of all these paths. A *flow* for commodity k is a nonnegative vector $\mathbf{x}^k = (x_p^k)_{P \in \mathcal{P}^k}$ such that $\sum_{P \in \mathcal{P}^k} x_p^k = r^k$. A *flow* \mathbf{x} is a vector $(\mathbf{x}^k)_{k \in K}$, where each \mathbf{x}^k is a flow for commodity k . For a flow \mathbf{x} and $e \in E$, let $x_e^k = \sum_{P \in \mathcal{P}^k: e \in P} x_p^k$ be the amount of flow that \mathbf{x}^k routes on each link e and let $x_e = \sum_{k \in K} x_e^k$ be the amount of flow that \mathbf{x} routes on e . With a slight abuse of notation, we will also denote $\mathbf{x} = (x_e)_{e \in E}$, the link-wise description of a flow.

Wardrop Equilibria and Optimal Flows. Given flow \mathbf{x} , the delay experienced on e is $\ell_e(x_e)$. In the case a toll $t_e \geq 0$ is charged for link usage, the combined cost of traversing e is $[\ell_e(x_e) + \alpha \cdot t_e]$, where α represents the trade-off factor between delay and tolls and is assumed to be identical for all users. Without loss of generality, we can assume α to be equal to 1 because we can always divide all tolls by α . A flow is a *Wardrop equilibrium* if it is supported on paths of minimum cost. Formally, \mathbf{x} is a Wardrop equilibrium if, for every k , for every path $P \in \mathcal{P}^k$ with $x_P^k > 0$, and every path $P' \in \mathcal{P}^k$, $\sum_{e \in P} [\ell_e(x_e) + t_e] \leq \sum_{e \in P'} [\ell_e(x_e) + t_e]$. Given such a flow, any $o^k - d^k$ path achieving the minimum end-to-end cost will be called an *active path*, and any link that belongs to an active path will be called an *active link*.

For any toll vector $\mathbf{t} \geq 0$, a Wardrop equilibrium exists, moreover it minimizes the convex Wardrop potential $\Phi_t(\mathbf{x})$ [4], and thus is given by:

$$\mathbf{x}(\mathbf{t}) \triangleq \arg \min_{\mathbf{x} \geq 0 \text{ flow}} \left\{ \Phi_t(\mathbf{x}) \triangleq \sum_{e \in E} \int_0^{x_e} (\ell_e(\mathbf{x}) + t_e) dx \right\}. \quad (1)$$

Notice that by our choice of latency functions, Φ_t is strictly convex, implying that the link-wise description of the Wardrop equilibrium is unique, and therefore $(x_e(\mathbf{t}))_{e \in E}$ is a well-defined function, which is moreover continuous by Berge's theorem [5].

Given flow \mathbf{x} , the total delay experienced by users is $\sum_{e \in E} x_e \ell_e(x_e)$, and it is the standard measure of network performance. Notice this function is strictly convex in \mathbf{x} . The optimal flow \mathbf{x}^* is the flow that minimizes the total delay. In general, it is known [32] that the untolled Wardrop flow $\mathbf{x}(0)$ is not necessarily optimal. However, a classical result [4] shows that the vector of marginal tolls $\hat{\mathbf{t}}$, defined as $\hat{t}_e = x_e^* \ell'_e(x_e^*)$, induces the optimal flow, that is $\mathbf{x}(\hat{\mathbf{t}}) = \mathbf{x}^*$. Any toll vector with this property will be called *optimal*.

The Network Operators' Game: Price Competition on Tolls

In our model, additionally, every link $e \in E$ is operated by a different operator: these are the players of the price competition game. Each player e is allowed to charge a nonnegative toll t_e for its usage. Under the resulting toll vector \mathbf{t} , each link gets flow $x_e(\mathbf{t})$ according to the induced Wardrop equilibrium, and thus the profit of player e is given by $\pi_e(\mathbf{t}) \triangleq t_e x_e(\mathbf{t})$. We are interested in the equilibrium outcomes of this game.

Profit functions. For each player $e \in E$, her strategy is given by toll t_e , and her profit is given by $\pi_e(t_e, \mathbf{t}_{-e}) = t_e x_e(t_e, \mathbf{t}_{-e})$, where we have used the standard game-theoretic notation $\mathbf{t} = (t_e, \mathbf{t}_{-e})$. Whenever \mathbf{t}_{-e} is clear from context, we will simply write $x_e(t_e)$ and $\pi_e(t_e)$.

Regulated Network Pricing Game and Nash Equilibria. The regulated network pricing game we consider is the following. A central planner may choose a cap vector $\bar{\mathbf{t}} \geq 0$ for tolls, and each

³More general assumptions have been studied in the literature, but we will restrict to this setting in order to avoid some minor technicalities.

player wants to maximize her own profit under this constraint. We study two solution concepts, (pure) Nash equilibrium and strong (pure) Nash equilibrium. Tolls \mathbf{t} are a Nash equilibrium for the network pricing game if for every e , t_e is the best response of player e to \mathbf{t}_{-e} , i.e. we have $t_e \in \mathcal{BR}_e(\mathbf{t}_{-e})$, where the best response mapping $\mathcal{BR}_e(\mathbf{t}_{-e})$ is defined as

$$\mathcal{BR}_e(\mathbf{t}_{-e}) \triangleq \arg \max \{\pi_e(s_e, \mathbf{t}_{-e}) : 0 \leq s_e \leq \bar{t}_e\}.$$

Tolls \mathbf{t} are a strong Nash equilibrium if there is no possible coalition that jointly deviates, resulting in an improvement of their individual profits. Formally, tolls \mathbf{t} are a strong Nash equilibrium if for any set $E^\Delta \subset E$ (the coalition) and \mathbf{t}' (the deviation) such that $E^\Delta = \{e \in E : t'_e \neq t_e\}$, if there is a player $e \in E^\Delta$ who gains by deviating with $\pi_e(\mathbf{t}') > \pi_e(\mathbf{t})$ then there exists another player $f \in E^\Delta$ who is worse off with $\pi_f(\mathbf{t}') < \pi_f(\mathbf{t})$. Notice that any strong Nash equilibrium is a Nash equilibrium.

Definition 2.1 (Cap equilibrium and great tolls). Given an instance of the profit maximization game, we say that a (nonnegative) vector $\bar{\mathbf{t}} = (\bar{t}_e)_{e \in E}$ is

- (a) a cap equilibrium if when restricting the strategy space for every player e to tolls $s_e \in [0, \bar{t}_e]$, then $(s_e)_{e \in E} = (\bar{t}_e)_{e \in E}$ is a strong Nash equilibrium;
- (b) a great set of tolls if it is optimal (i.e., induces the optimal flow) and a cap equilibrium.

The case where $\bar{t}_e = \infty$ for every $e \in E$ corresponds to the unregulated game, which will be the main focus of the next section. Observe that there are situations where the unregulated game may not have equilibrium at all, for instance in the presence of a monopoly, i.e. a link that cuts a commodity. In order to avoid such cases, we will assume that each commodity has at least two disjoint $o^k - d^k$ paths. Notice this will only be necessary for the unregulated game, as when imposing caps we immediately prevent monopolies from setting arbitrarily large prices.

3 UNREGULATED NETWORK PRICING GAME

We address the question of existence of equilibria in the network pricing game without caps. In this setting, the only case where equilibria are known to exist is for parallel link networks and affine latencies, and, additionally, if one drops the affine latencies requirement while keeping the parallel links topology, then equilibria may fail to exist (see e.g. [1, 20]). To highlight and add on the brittleness of existence of equilibria for the unregulated game, we show that if one drops the network topology requirement, in the simplest nontrivial way, while keeping the affine latencies requirement, then equilibria may again fail to exist. We conclude the section by showing that even when equilibria exist, the resulting performance may be arbitrarily bad with respect to the social optimum.

Nonexistence of Equilibria. To begin with, we show that even for affine latencies and rather simple topologies, an equilibrium is not guaranteed to exist.

Example 3.1. Consider the extension-parallel network on the left side of Figure 2, with vertex set $\{o, u, d\}$, link set $E = \{e_1 = (o, u), e_2 = (u, d), e_3 = (u, d), e_4 = (o, d)\}$, latencies $\ell_{e_1}(x) = \ell_{e_3}(x) = \ell_{e_4}(x) = x$ and $\ell_{e_2}(x) = 2$, and a single o, d commodity with demand $r = 2$. This instance does not have a Nash equilibrium.⁴

⁴ In this example, as well as in other examples, constant latencies are present. This violates the assumption that latencies are strictly increasing, yet this assumption was adopted for technical simplicity in order to have uniqueness of Wardrop flow, so that profits are well defined. All examples present in our work indeed satisfy such uniqueness, and moreover they can be modified to satisfy the strictly increasing assumption, e.g. by adding an ϵx factor to each of the constant latencies for $\epsilon > 0$ sufficiently small.

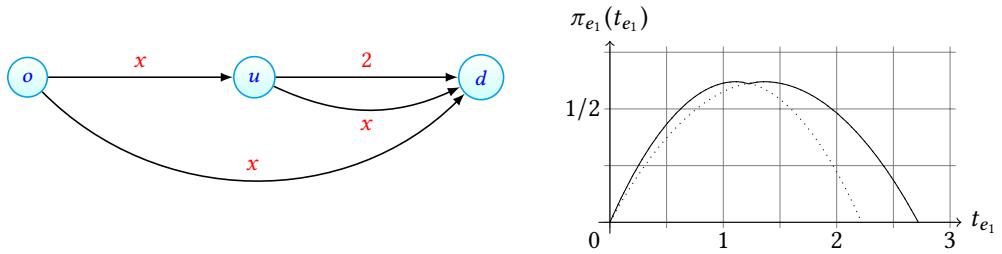


Fig. 2. On the left, the network of Example 3.1. On the right, the profit function of player 1 when $t_{e_2} = 0$, $t_{e_3} = 3/2$ and $t_{e_4} = (2 + \sqrt{6})/2$.

To get an intuition why no Nash equilibrium exists, suppose that $t_{e_2} = 0$, $t_{e_3} = 3/2$ and $t_{e_4} = (2 + \sqrt{6})/2$. Figure 2 on the right, shows the profit function of player 1. Notice that the profit function is not concave and has two maxima, one at $(2 + \sqrt{6})/4$ and one at $(3 + \sqrt{6})/4$. This implies that the best-response correspondence of player 1 is not convex. Yet, convexity of the best-response correspondences seems to be a crucial property as, on the one hand, it is necessary for the existence proof for the parallel links with affine latencies case (which is using Kakutani's fixed-point theorem), and, on the other hand, it is also absent from other nonexistence examples with parallel links [1, 20]. \square

Unboundedness of Total Delay. Next, we show that even when Nash equilibria exist, the resulting Wardrop flows may be arbitrarily inefficient, in terms of total delay. Our example is a parallel link network with two links and affine latencies.

Example 3.2. Consider a single-commodity instance where $r = 1$ units of flow is routed through a two-link parallel network with $\ell_{e_1}(x) = Ax$ and $\ell_{e_2}(x) = a$ for $0 < a < A$. By direct computation, we get the optimal flow $\mathbf{x}^* = (\frac{a}{2A}, \frac{2A-a}{2A})$, achieving total delay $[4aA - a^2]/[4A]$.

On the other hand, for any toll vector \mathbf{t} , the Wardrop flow is $\mathbf{x}(\mathbf{t}) = ([a + t_{e_2} - t_{e_1}]/A, [A - a - t_{e_2} + t_{e_1}]/A)$ (provided these values lie in $[0, 1]$). By the first-order conditions on the profit maximization problem for players 1 and 2, we get a Nash equilibrium $(t_{e_1}, t_{e_2}) = ([a + A]/3, [2A - a]/3)$, and the respective Wardrop flow $\mathbf{x}(\mathbf{t}) = ([A + a]/[3A], [2A - a]/[3A])$, resulting in total delay $[A^2 + 8Aa - 2a^2]/[9A]$.

As a consequence, the ratio of the total delay at the Nash equilibrium over the optimal delay is $\frac{4[A^2 + 8Aa - 2a^2]}{9[4aA - a^2]}$. Keeping $A > 0$ fixed and taking $a \rightarrow 0$, we have that the ratio diverges. \square

4 REGULATED NETWORK PRICING GAME

In this section, we study the profit maximization game and some of its structural properties. We prove that the marginal tolls, when used as caps, are always a Nash equilibrium for the profit maximization game, thus resolving the issues of equilibrium existence, uniqueness and inefficiency raised in the previous section. Our main result, presented in Theorem 4.4, strengthens the above by showing that a potentially larger set of optimal tolls, when used as caps, leads to a unique Nash equilibrium, which is furthermore robust to coalitions.

We start by studying the behavior of the Wardrop equilibrium flow as a function of tolls. The following result states some monotonicity properties that are standard. Its proof is omitted since we prove a more general result right after.

LEMMA 4.1 (MONOTONICITY). *Let $e \in E$ and \mathbf{t}_{-e} be a fixed toll vector for the rest of the players. If $t'_e > t_e$ then*

- (i) $x_e(t'_e) \leq x_e(t_e)$, and
- (ii) $\ell_e(x_e(t'_e)) + t'_e \geq \ell_e(x_e(t_e)) + t_e$.

In this work we need a generalization of the previous result, establishing how flows may change as several tolls decrease. This generalization is stated and proved below.

LEMMA 4.2. *Let $\mathbf{t}, \mathbf{t}' \geq 0$ be two toll vectors such that $\mathbf{t} \leq \mathbf{t}'$ and $E^< = \{e \in E : t_e < t'_e\}$ is nonempty. Then, there exist $e_1, e_2 \in E^<$ such that:*

- (i) $x_{e_1}(\mathbf{t}') \leq x_{e_1}(\mathbf{t})$,
- (ii) $[x_{e_2}(\mathbf{t}') - x_{e_2}(\mathbf{t})][\ell_{e_2}(x_{e_2}(\mathbf{t}')) + t'_{e_2} - \ell_{e_2}(x_{e_2}(\mathbf{t})) - t_{e_2}] \leq 0$.

PROOF. To prove (i) we compare flows $\mathbf{x}(\mathbf{t})$ and $\mathbf{x}(\mathbf{t}')$ with respect to the Wardrop potentials, $\Phi_t(\mathbf{x})$ and $\Phi_{t'}(\mathbf{x}) = \Phi_t(\mathbf{x}) + \sum_{e \in E^<} [t'_e - t_e]x_e$. By the optimality of the Wardrop flow on its respective potential, we get the following inequalities

$$\begin{aligned} \Phi_t(\mathbf{x}(\mathbf{t})) &\leq \Phi_t(\mathbf{x}(\mathbf{t}')) \\ \Phi_t(\mathbf{x}(\mathbf{t}')) + \sum_{e \in E^<} [t'_e - t_e]x_e(\mathbf{t}') &\leq \Phi_t(\mathbf{x}(\mathbf{t})) + \sum_{e \in E^<} [t'_e - t_e]x_e(\mathbf{t}). \end{aligned}$$

Combining these inequalities we get

$$\sum_{e \in E^<} [t'_e - t_e][x_e(\mathbf{t}') - x_e(\mathbf{t})] \leq 0,$$

and thus there must exist $e_1 \in E^<$ such that $x_{e_1}(\mathbf{t}') \leq x_{e_1}(\mathbf{t})$, proving (i).

Let us now prove (ii). Notice that since $\mathbf{x}(\mathbf{t}) + (\mathbf{x}(\mathbf{t}') - \mathbf{x}(\mathbf{t})) = \mathbf{x}(\mathbf{t}')$ is a feasible flow, then $\mathbf{x}(\mathbf{t}') - \mathbf{x}(\mathbf{t})$ is a feasible direction for $\mathbf{x}(\mathbf{t})$. By the first-order optimality conditions,

$$\sum_e [\ell_e(x_e(\mathbf{t})) + t_e][x_e(\mathbf{t}) - x_e(\mathbf{t}')] \leq 0.$$

Analogously, $\mathbf{x}(\mathbf{t}) - \mathbf{x}(\mathbf{t}')$ is a feasible direction for $\mathbf{x}(\mathbf{t}')$, and thus

$$\sum_e [\ell_e(x_e(\mathbf{t}')) + t'_e][x_e(\mathbf{t}') - x_e(\mathbf{t})] \leq 0.$$

Adding up these inequalities, we obtain

$$\sum_{e \in E^=} [x_e(\mathbf{t}') - x_e(\mathbf{t})][\ell_e(x_e(\mathbf{t}')) - \ell_e(x_e(\mathbf{t}))] + \sum_{e \in E^<} [x_e(\mathbf{t}') - x_e(\mathbf{t})][\ell_e(x_e(\mathbf{t}')) + t'_e - \ell_e(x_e(\mathbf{t})) - t_e] \leq 0,$$

where $E^= = \{e \in E : t_e = t'_e\}$. Observe now that the first summation term is nonnegative, as ℓ_e is increasing for all e . Thus, $\sum_{e \in E^<} [x_e(\mathbf{t}') - x_e(\mathbf{t})][\ell_e(x_e(\mathbf{t}')) + t'_e - \ell_e(x_e(\mathbf{t})) - t_e] \leq 0$, implying that there exists $e_2 \in E^<$ such that $[x_{e_2}(\mathbf{t}') - x_{e_2}(\mathbf{t})][\ell_{e_2}(x_{e_2}(\mathbf{t}')) + t'_{e_2} - \ell_{e_2}(x_{e_2}(\mathbf{t})) - t_{e_2}] \leq 0$, proving (ii). \square

The following result gives an intriguing inequality satisfied by any profit maximizing toll. This property is a consequence of the first-order optimality conditions, in combination with the monotonicity properties stated above. We will later see this lemma is crucial for our main result.

LEMMA 4.3. *Let $\mathbf{t} \geq 0$ be a toll vector and $e \in E$ with $t_e > 0$. If t_e is a local optimum for the profit maximization problem (that is, for objective $\pi_e(\cdot, \mathbf{t}_{-e})$) then $x_e(t_e) \cdot \ell'_e(x_e(t_e)) \leq t_e$.*

PROOF. Notice first that in the case $x_e(t_e) = 0$ the result obviously holds, thus we may restrict ourselves to the case $x_e(t_e) > 0$. In order to use the first-order optimality conditions, we need to ensure that the flow on e neither suddenly drops to zero nor it remains constant. By continuity, there exists an interval $[t_e, t_e + \delta']$ where $x_e(\cdot) > 0$, with $\delta' > 0$; furthermore, by local optimality,

we may choose $0 < \delta < \delta'$ so that t_e is a profit maximizing toll on the interval. This in particular implies that $0 < x_e(s_e) < x_e(t_e)$ for any $s_e \in [t_e, t_e + \delta]$.

Now, since t_e is a local maximizer, we can use the first-order optimality conditions, $D^+[\pi_e(t_e)] \leq 0$, where $D^+[\cdot]$ is the upper-right Dini derivative.⁵ Using the linearity of Dini derivatives, we get

$$x_e(t_e) + t_e D^+[x_e(t_e)] = D^+[\pi_e(t_e)] \leq 0; \quad (2)$$

notice in particular $D^+[x_e(t_e)] < 0$.

On the other hand, since $x_e(s_e) < x_e(t_e)$ for any $s_e \in [t_e, t_e + \delta]$, by Lemma 4.1 (ii), we have that $[\ell_e(x_e(s_e)) + s_e - \ell_e(x_e(t_e)) - t_e]/[x_e(s_e) - x_e(t_e)] \leq 0$. After re-arranging terms, we get

$$\frac{\ell_e(x_e(s_e)) - \ell_e(x_e(t_e))}{x_e(s_e) - x_e(t_e)} \leq -\frac{s_e - t_e}{x_e(s_e) - x_e(t_e)}.$$

Taking $\limsup_{s_e \rightarrow t_e^+}$ in the expression above, we get that the left hand side converges to $\ell'_e(x_e(t_e))$, whereas the right hand side converges to $-(D^+[x_e(t_e)])^{-1}$ (recall the Dini derivative is nonzero); we conclude that $\ell'_e(x_e(t_e)) \leq -\frac{1}{D^+[x_e(t_e)]}$. This, in combination with (2), gives

$$x_e(t_e) \ell'_e(x_e(t_e)) \leq -\frac{x_e(t_e)}{D^+[x_e(t_e)]} \leq t_e,$$

which proves the result. \square

Our main result below shows a strong consequence of the lemma above. Consider tolls that induce the optimal flow, and are at most the marginal tolls (in particular, marginal tolls satisfy these two conditions). Then these tolls, when used as caps, lead to a unique and strong Nash equilibrium.

THEOREM 4.4. *Let $\bar{t} \leq \hat{t}$ be an optimal toll vector. For the profit maximization problem with caps \bar{t} , there exists a unique Nash equilibrium, given by \bar{t} , which, moreover, is a strong Nash equilibrium.*

PROOF. First we prove that \bar{t} is a strong Nash equilibrium for the game. By way of contradiction, let $E^<$ be a set of links for which the corresponding players, without losing profit, deviate to some smaller toll value and let $t \leq \bar{t}$ be the toll vector after these deviations. In particular, for all $e \in E^<$ we have $t_e x_e(t) \geq \bar{t}_e x_e(\bar{t}) > t_e x_e(\bar{t})$, which implies $t_e > 0$ along with $x_e(t) > x_e(\bar{t})$ and

$$t_e \geq \frac{\bar{t}_e x_e(\bar{t})}{x_e(t)}. \quad (3)$$

By Lemma 4.2 (ii), there exists $e \in E^<$ such that $[x_e(\bar{t}) - x_e(t)][\ell_e(x_e(\bar{t})) + \bar{t}_e - \ell_e(x_e(t)) - t_e] \leq 0$, which in combination with $x_e(t) > x_e(\bar{t})$ gives

$$\ell_e(x_e(\bar{t})) + \bar{t}_e \geq \ell_e(x_e(t)) + t_e. \quad (4)$$

Let $e \in E^<$ be a link satisfying (4). We have the following inequalities

$$\ell'_e(x_e(\bar{t}))(x_e(t) - x_e(\bar{t})) \leq \ell_e(x_e(t)) - \ell_e(x_e(\bar{t})) \leq \bar{t}_e - t_e \leq \bar{t}_e \left(1 - \frac{x_e(\bar{t})}{x_e(t)}\right), \quad (5)$$

where the first inequality follows from convexity of ℓ_e , the second from (4), and the third from (3). Since $\bar{t}_e \leq \hat{t}_e = \ell'_e(x_e(\bar{t}))x_e(\bar{t})$, we obtain from (5) that

$$\ell'_e(x_e(\bar{t}))(x_e(t) - x_e(\bar{t})) \leq \ell'_e(x_e(\bar{t}))x_e(\bar{t}) \left(\frac{x_e(t) - x_e(\bar{t})}{x_e(t)}\right).$$

Since $\ell'_e(x_e(\bar{t})) > 0$ and $x_e(t) > x_e(\bar{t})$, we conclude that $x_e(t) \leq x_e(\bar{t})$, a contradiction.

⁵Recall the definition of the upper-right Dini derivative, $D^+[f(x)] \triangleq \limsup_{h \rightarrow 0^+} \frac{f(x+h) - f(x)}{h}$.

Now that we have proved existence, we show there is a unique Nash equilibrium. By way of contradiction, suppose there exists another Nash equilibrium $\mathbf{t} \neq \bar{\mathbf{t}}$ for the profit maximization game with caps $\bar{\mathbf{t}}$. Since all players are playing their best response, we may assume $t_e > 0$ for all players e such that $\bar{t}_e > 0$, and thus we can use Lemma 4.3, concluding that $x_e(\mathbf{t})\ell'_e(x_e(\mathbf{t})) \leq t_e$ for all $e \in E^< = \{e \in E : t_e < \bar{t}_e\}$ (nonempty by assumption) which in turn gives

$$x_e(\mathbf{t})\ell'_e(x_e(\mathbf{t})) \leq t_e < \bar{t}_e \leq \hat{t}_e = x_e^*\ell'_e(x_e^*),$$

concluding that $x_e(\mathbf{t}) < x_e^* = x_e(\bar{\mathbf{t}})$ for all $e \in E^<$. But from Lemma 4.2 (i) there exists $e_1 \in E^<$ such that $x_{e_1}(\mathbf{t}) \geq x_{e_1}(\bar{\mathbf{t}})$, a contradiction. \square

REMARK 4.5. *Theorem 4.4 cannot be extended to the setting where an operator owns more than one link. For an example, consider the network of Figure 3 with unit demand and suppose that all links are owned by a single operator. The marginal tolls, and thus the caps, are given by $(\hat{t}_{ou}, \hat{t}_{uv}, \hat{t}_{vd}) = (2/3, 1/3, 2/3)$, yielding a total profit of 1 for the operator (links (o, v) and (u, d) have zero caps). If, however, the operator charges a toll vector of $(t_{ou}, t_{uv}, t_{vd}) = (2/3, 0, 2/3)$, then it induces a flow of $2/9$ on each of the upper and lower paths, and a flow of $5/9$ on the zig-zag path, gaining a profit of $28/27$. Hence, the marginal tolls are not a cap equilibrium.*

5 MINIMIZING USERS' COSTS

Theorem 4.4 implies that any opt-inducing toll vector that is not above the marginals is itself a cap equilibrium, and thus is a great set of tolls. The fact that great tolls need not be unique, motivates the question of whether it is possible to compute the great tolls vector that minimizes the total users' costs. If all great tolls were not above the marginals, then computing the great tolls vector that minimizes the total users' costs would reduce to solving a linear program (see program (BMT) below). Yet, this is not the case, since, as the following example suggests, there are great tolls that are above the marginals for some edges and additionally incur smaller costs for the users.

Example 5.1. Consider the Braess network of Figure 3 with unit demand. The optimal flow is obtained by sending one third of the flow on each of the upper and lower paths, and one third of the flow on the zig-zag path; leading to a total delay of $21/9$. The marginal tolls are given by $(\hat{t}_{ou}, \hat{t}_{uv}, \hat{t}_{vd}) = (2/3, 1/3, 2/3)$, yielding a total profit of 1. Hence the users' costs under the marginal tolls are $21/9 + 1$. Notice that the only great tolls that are not above the marginal tolls are the marginal tolls themselves.

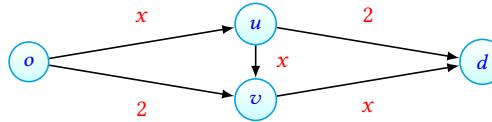


Fig. 3. The network of Remark 4.5 and Example 5.1. Under unit demand, the optimal flow routes one third of the flow to each of the upper, zig-zag and lower paths.

Taking $S = \{o, v\}$ as a cut, we can decrease the tolls by some $\varepsilon > 0$ for outgoing links (namely, (o, u) and (v, d)) and increase them by ε on incoming links (i.e., (u, v)). This procedure maintains optimality of tolls, however it is unclear whether the cap equilibrium property is maintained. Solving the profit maximization problem explicitly for link (u, v) we have that these modified tolls are a cap equilibrium if and only if $\varepsilon \leq 1/6$, and therefore the minimum users' cost among great tolls is $21/9 + 5/6$. \square

The mathematical program for obtaining great tolls that minimize users' costs while keeping \mathbf{x}^* as a Wardrop equilibrium flow is (BGT stands for *Best Great Tolls*):

$$(BGT) \left\{ \begin{array}{llll} \min & \sum_{e \in E} [\ell_e(x_e^*) + t_e] x_e^* & & \\ & v_u^k - v_v^k + t_e & = & -\ell_e(x_e^*) & \forall k, e = (u, v) : x_e^{*k} > 0 \\ & v_u^k - v_v^k + t_e & \geq & -\ell_e(x_e^*) & \forall k, e = (u, v) : x_e^{*k} = 0 \\ & t & & \text{is a cap equilibrium} & \\ & t & \geq & 0 & \end{array} \right.$$

The program, in variables (t, v) , defines a potential v^k for each commodity k , in such a way that any flow-carrying path is indeed a shortest path. A careful reader may also notice that we included in the objective the constant term $\sum_e \ell_e(x_e^*) x_e^*$, corresponding to the total delay experienced by users. This is of course part of the cost experienced by users, but it also turns out to be important in order to obtain some of the approximation bounds we are aiming to prove.

In the formulation above, the problem (BGT) does not have any amenable structure for efficient methods, particularly due to the cap equilibrium constraint. However, by Theorem 4.4, optimal tolls upper bounded by the marginal tolls are always a cap equilibrium, thus within this restricted set of tolls we can write the following LP, that encodes all possible great set of tolls upper bounded by marginals (BMT stands for *Below Marginal Tolls*):

$$(BMT) \left\{ \begin{array}{llll} \min & \sum_{e \in E} [\ell_e(x_e^*) + t_e] x_e^* & & \\ & v_u^k - v_v^k + t_e & = & -\ell_e(x_e^*) & \forall k, e = (u, v) : x_e^{*k} > 0 \\ & v_u^k - v_v^k + t_e & \geq & -\ell_e(x_e^*) & \forall k, e = (u, v) : x_e^{*k} = 0 \\ & t_e & \leq & \hat{t}_e & \forall e \in E \\ & t_e & \geq & 0 & \forall e \in E. \end{array} \right.$$

Since by Example 5.1 the value of (BMT) does not necessarily coincide with the value of (BGT), the next natural question to ask is how efficient program (BMT) can be. In order to answer this question, we can use as benchmark the value of *Minimum Payment Tolls*: These are the optimal tolls that minimize the total payments. These tolls can be computed by a similar linear program, where we only drop the cap equilibrium constraints (MPT stands for *Minimum Payment Tolls*):

$$(MPT) \left\{ \begin{array}{llll} \min & \sum_{e \in E} [\ell_e(x_e^*) + t_e] x_e^* & & \\ & v_u^k - v_v^k + t_e & = & -\ell_e(x_e^*) & \forall k, e = (u, v) : x_e^{*k} > 0 \\ & v_u^k - v_v^k + t_e & \geq & -\ell_e(x_e^*) & \forall k, e = (u, v) : x_e^{*k} = 0 \\ & t_e & \geq & 0 & \forall e \in E. \end{array} \right.$$

In what follows, \bar{t} and t^* are optimal solutions of (BMT) and (MPT), respectively.

5.1 Functional Approximation Bound on Cost

We consider an approximation bound based on properties of the latency functions.

THEOREM 5.2. *Suppose all latency functions ℓ in the profit maximization game satisfy $\sup_{x \geq 0} \frac{x\ell''(x)}{\ell(x)} \leq \gamma$, then $\text{val}(\text{BMT}) \leq (\gamma + 1) \text{val}(\text{MPT})$.*

PROOF. By the assumption on latency functions,

$$\begin{aligned} \text{val}(\text{BMT}) &= \sum_e [\ell_e(x_e^*) + \bar{t}_e] x_e^* \leq \sum_e [\ell_e(x_e^*) + x_e^* \ell'_e(x_e^*)] x_e^* \\ &\leq \sum_e (1 + \gamma) \ell_e(x_e^*) x_e^* \leq (\gamma + 1) \text{val}(\text{MPT}), \end{aligned}$$

proving the result. \square

For polynomial functions with nonnegative coefficients and degree at most d it is known that one can choose $\gamma = d$ above, which gives the following corollary:

COROLLARY 5.3. *For polynomial latency functions of degree at most d and nonnegative coefficients, $\text{val}(\text{BMT}) \leq (d + 1) \cdot \text{val}(\text{MPT})$.*

The proposed approximation bound looks rather weak. However, this bound is tight even for single-commodity networks.

5.2 Topological Approximation Bound on Cost

As shown above, having a Braess type subnetwork can lead to a gap between the linear programs (BMT) and (MPT), and moreover a *generalized Braess subnetwork* will further worsen this gap. Notice however that these examples in order to achieve large gaps need both a complex network structure and large degree latencies. Is it possible that, e.g., simple networks with high-degree latencies still achieve a small gap? Fortunately, the answer is yes.

In this subsection we will restrict ourselves to single-commodity instances and will assume that, without loss of generality, they have unit demand. With this, the (BMT) linear program simplifies substantially,

$$(\text{BMT}) \left\{ \begin{array}{llll} \min & \sum_{e \in E} [\ell_e(x_e^*) + t_e] x_e^* & & \\ & v_u - v_v + t_e & = & -\ell_e(x_e^*) \quad \forall e = (u, v) : x_e^* > 0 \\ & v_u - v_v + t_e & \geq & -\ell_e(x_e^*) \quad \forall e = (u, v) : x_e^* = 0 \\ & t_e & \leq & \hat{t}_e \quad \forall e \in E \\ & t_e & \geq & 0 \quad \forall e \in E. \end{array} \right.$$

We start with a structural lemma, which allows us to upper bound the value of (BMT). This is naturally important in order to derive an approximation bound. We will first need some definitions. Consider a directed network G and the undirected network G^u that comes from G if we drop the directions of its links. Any path in G^u is called an undirected path in G . For an undirected path P , the links that are traversed in their actual direction are called *forward* links, denoted by P^+ , and the ones traversed in their reversed direction are called *backward* links, denoted by P^- . Finally, an undirected path has J alternations if, when traversing it, there are exactly J times where a forward link is followed by a backward link.⁶

LEMMA 5.4. *There exists an undirected $o - d$ path P such that the first and the last link of P belong in P^+ , all $e \in P^+$ are flow carrying, and the value of (BMT) is equal to*

$$\sum_{e \in P^+} \ell_e(x_e^*) - \sum_{e \in P^-} [\ell_e(x_e^*) + \hat{t}_e].$$

We make the following assumption on the instance: there exists a $J \geq 0$ such that any simple $o - d$ undirected path has at most J alternations. The smallest constant J satisfying this condition will be called the *alternation number*, and our approximation bound will only depend on this number.

THEOREM 5.5. *Consider a single-commodity and unit demand instance of the network pricing game whose underlying network has alternation number J . We have, $\text{val}(\text{BMT}) \leq (J + 1) \cdot \text{val}(\text{MPT})$.*

PROOF. By Lemma 5.4, we have $\text{val}(\text{BMT}) = \sum_{e \in P_+} \ell_e(x_e^*) - \sum_{e \in P_-} [\ell_e(x_e^*) + \hat{t}_e]$. Since the alternation number of G is J , we can decompose P in at most J segments of consecutive forward and backward links, $P = A_1 - B_1 - A_2 - B_2 - \dots - B_J - A_{J+1}$, from which $\text{val}(\text{BMT}) \leq \sum_{j=1}^{J+1} [\sum_{e \in A_j} \ell_e(x_e^*)]$. Let now $(\mathbf{t}^*, \mathbf{v}^*)$ be an optimal solution for (MPT); since all links $e \in A_j$, $j = 1, \dots, J + 1$ are flow

⁶For such alternating paths, see also, e.g., [23, 25, 28]

carrying, for each $j = 1, \dots, J + 1$, we have $\sum_{e \in A_j} \ell_e(x_e^*) \leq v_d^* - v_o^*$. Combining the inequalities, we obtain $\text{val}(\text{BMT}) \leq \sum_{j=1}^{J+1} \left[\sum_{e \in A_j} \ell_e(x_e^*) \right] \leq (J + 1)[v_d^* - v_o^*] = (J + 1) \cdot \text{val}(\text{MPT})$. \square

Since $J = 0$ for series-parallel networks, we get the following corollary, which implies that our approach is optimal in this case.

COROLLARY 5.6. *For instances on series-parallel networks, \bar{t} minimizes the users' cost and payments.*

The bound obtained in Theorem 5.5 is tight for single-commodity instances and it does not extend to multi-commodity instances, or even single-source multi-commodity instances.

6 ELASTIC DEMAND

In this section, we assume that the total amount of traffic for each $o^k - d^k$ pair is dependent on the costs the traffic experiences; the higher the costs, the lower the traffic. We analyze the elastic traffic demand model as introduced by Beckmann et al. [4] and later used for network pricing games by Hayrapetyan et al. [21] and Ozdaglar [29]. We model elastic demand with a *utility function* $u^k : [0, r^k] \rightarrow \mathbb{R}_+$ for each $k \in K$, where $u^k(x)$ captures the reservation value for travel of the particle of the demand (i.e., the infinitesimally small user) at level x . We assume that $u^k(\cdot)$ is nonincreasing and continuous for each $k \in K$, so that, in a sense, the users are ordered decreasingly according to their utility for traveling. Let $u = (u^k)_{k \in K}$ be the vector of all utility functions. Define the aggregate utility function $U^k : [0, r^k] \rightarrow \mathbb{R}_+$ by $U^k(x) = \int_0^x u^k(y) dy$. By definition, this function is nondecreasing, concave and continuously differentiable.

A *flow* for commodity k is a nonnegative vector $\mathbf{x}^k = (x_P^k)_{P \in \mathcal{P}^k}$ such that $\sum_{P \in \mathcal{P}^k} x_P^k \leq r^k$. For each commodity $k \in K$, let $r_x^k = \sum_{P \in \mathcal{P}^k} x_P^k$ be the amount of flow that is routed on the network by \mathbf{x}^k . A *flow* \mathbf{x} is a vector $(\mathbf{x}^k)_{k \in K}$, where each \mathbf{x}^k is a flow for commodity k .

A flow \mathbf{x} is a *Wardrop equilibrium* if all the routed traffic goes through shortest paths for the respective commodity, the utility for each traveling user is at least equal to the common shortest path cost of her commodity, and the utility of each user not traveling is at most equal to the common shortest path cost of her commodity. Mathematically, for every k , for every path $P \in \mathcal{P}^k$ with $x_P^k > 0$, and every path $P' \in \mathcal{P}^k$, $\sum_{e \in P} [\ell_e(x_e) + t_e] \leq \sum_{e \in P'} [\ell_e(x_e) + t_e]$ and $\sum_{e \in P} [\ell_e(x_e) + t_e] \leq u^k(r_x^k)$ with $r_x^k = r^k$ if $\sum_{e \in P} [\ell_e(x_e) + t_e] < u^k(r_x^k)$.

For any toll vector $\mathbf{t} \geq 0$, a Wardrop equilibrium exists, and moreover can be found by solving the following maximization problem [4]:

$$\mathbf{x}(\mathbf{t}) \triangleq \arg \max_{\mathbf{x} \geq 0 \text{ flow}} \left\{ \sum_{k \in K} U^k(r_x^k) - \sum_{e \in E} \int_0^{x_e} (\ell_e(\mathbf{x}) + t_e) dx \right\}. \quad (6)$$

Assuming that each $\ell_e(x_e)$ is strictly increasing for each $e \in E$, the link-wise description of the Wardrop equilibrium is unique, and therefore $(x_e(\mathbf{t}))_{e \in E}$ is a well-defined function, which is moreover continuous by Berge's theorem [5].

Given flow \mathbf{x} and toll vector \mathbf{t} , define the *total users' cost* by $C(\mathbf{x}, \mathbf{t}) = \sum_{e \in E} [\ell_e(x_e) + t_e] x_e$ and the *consumer surplus* by

$$CS(\mathbf{x}, \mathbf{t}, u) = \sum_{k \in K} U^k(r_x^k) - C(\mathbf{x}, \mathbf{t}).$$

We define the *social welfare* by

$$SW(\mathbf{x}, u) = \sum_{k \in K} U^k(r_x^k) - \sum_{e \in E} \ell_e(x_e) \cdot x_e.$$

Notice that tolls do not appear in the social welfare, as they are transfers from users to toll operators.

An *optimal flow* $\mathbf{x}^*(u)$ is a flow that maximizes the social welfare w.r.t. to u . Similarly to the fixed demand model, we have the classical result by Beckmann et al. [4] showing that a flow $\mathbf{x}^*(u)$ is an optimal flow if and only if $\mathbf{x}^*(u) = \mathbf{x}(\hat{\mathbf{t}}(u))$, where $\hat{\ell}_e(u) = x_e^*(u)\ell'_e(x_e^*(u))$. Also, for every network and demand vector there exist utility functions such that all users travel under the optimal solution, i.e. $r_{x^*(u)}^k = r^k$ for each $k \in K$. Observe that, in this case, this optimal flow is the same optimal flow as defined in Section 2. For a given instance, we denote this particular optimal flow by \mathbf{x}^* .

6.1 Regulated Network Pricing Game

The elastic traffic demand model generalizes the fixed demand one. Fortunately, the main result of Section 4, i.e., Theorem 4.4, can be generalized as well.

THEOREM 6.1. *Let $\bar{\mathbf{t}} \leq \hat{\mathbf{t}}(u)$ be an optimal toll vector. For the profit maximization problem with caps $\bar{\mathbf{t}}$, there exists a unique Nash equilibrium, given by $\bar{\mathbf{t}}$, which, moreover, is a strong Nash equilibrium.*

For a proof, one has first to prove an analogue of Lemma 4.2. That provided, an analogue of Lemma 4.3 follows directly and the same proof as the one for Theorem 4.4 can be applied.

6.2 Maximizing Consumer Surplus

We now consider the question of which great tolls vector maximizes the consumer surplus. The main result in this subsection, Theorem 6.5, only applies to single-source, single-sink networks, and thus we will restrict ourselves to these instances.

By Theorem 6.1, optimal tolls upper bounded by the marginal tolls are always a cap equilibrium, thus within this restricted set of tolls we can write the following LP, that encodes all possible great set of tolls upper bounded by marginals:

$$(BMT2) \quad \left\{ \begin{array}{lll} \max & CS(\mathbf{x}^*(u), \mathbf{t}, u) & \\ & v_v - v_w + t_e & = -\ell_e(x_e^*(u)) \quad \forall e = (v, w) : x_e^*(u) > 0 \\ & v_v - v_w + t_e & \geq -\ell_e(x_e^*(u)) \quad \forall e = (v, w) : x_e^*(u) = 0 \\ & v_t - v_s & \leq u(r_{x^*(u)}) \quad \text{where the inequality is tight if } r_{x^*(u)} < r \\ & t_e & \leq \hat{\ell}_e(u) \quad \forall e \in E \\ & t_e & \geq 0 \quad \forall e \in E. \end{array} \right.$$

A natural question to ask is how efficient our program (BMT2) can be. In order to answer this question, we can use as benchmark the value of Minimum Payment Tolls: these are the optimal tolls that minimize the total payments. These tolls can be computed by a similar linear program, where we only drop the cap equilibrium constraints:

$$(MPT2) \quad \left\{ \begin{array}{lll} \max & CS(\mathbf{x}^*(u), \mathbf{t}, u) & \\ & v_v - v_w + t_e & = -\ell_e(x_e^*(u)) \quad \forall e = (v, w) : x_e^*(u) > 0 \\ & v_v - v_w + t_e & \geq -\ell_e(x_e^*(u)) \quad \forall e = (v, w) : x_e^*(u) = 0 \\ & v_t - v_s & \leq u(r_{x^*(u)}) \quad \text{where the inequality is tight if } r_{x^*(u)} < r \\ & t_e & \geq 0 \quad \forall e \in E. \end{array} \right.$$

In what follows, \mathbf{t}^B and \mathbf{t}^M are toll vectors that belong to some optimal solutions of (BMT2) and (MPT2), respectively. The following example shows that when all users travel, the objective of the maximization problem (BMT2) can be arbitrary higher than the objective of (MPT2).

Example 6.2. Consider the Braess network of Figure 4 with unit demand and $u(x) = 2$ for $x \in [0, 1]$.

An optimal flow in this case is to split the demand of one by half among the upper and lower paths; in particular, no flow traverses the middle uv link. This way, marginal tolls are given by

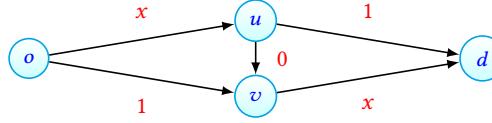


Fig. 4. The network of Example 6.2.

$(\hat{t}_{ou}, \hat{t}_{uv}, \hat{t}_{vd}) = (1/2, 0, 1/2)$, and notice this is the only feasible solution for (BMT2), achieving a value of $2 - 2 = 0$. On the other hand, it is easy to see that the LP solution for (MPT2) just needs to assign a sufficiently large value for t_{uv} (more precisely, $t_{uv} \geq 1/2$) and the rest of the tolls can be zero, therefore the value of (MPT2) is $2 - 3/2 = 1/2$. \square

The main result of this subsection is a comparison in the spirit of the bicriteria bound of Roughgarden and Tardos [34]: by how much should we decrease the utility function in order to induce the same level of consumer surplus when comparing (BMT2) to (MPT2). A corollary of the result is that multiplying the utility function by 1/2 is sufficient. In other words, the consumers' surplus under \mathbf{t}^B is at least as much as the optimal surplus if each of them had half the valuation.

First we need two lemmas:

LEMMA 6.3. *If $r_{x^*(u)} < r$, then $CS(\mathbf{x}^*(u), \mathbf{t}, u) = U(r_{x^*(u)}) - r_{x^*(u)} \cdot u(r_{x^*(u)})$ for all optimal \mathbf{t} .*

LEMMA 6.4. *Let \tilde{u} and u be utility functions so that $\tilde{u}(x) = \alpha \cdot u(x)$, where $0 < \alpha \leq 1$. Then*

- (i) $r_{x^*(u)} \geq r_{x^*(\tilde{u})}$,
- (ii) $CS(\mathbf{x}^*(u), \mathbf{t}^B, u) \geq CS(\mathbf{x}^*(\tilde{u}), \mathbf{t}^B, \tilde{u})$,
- (iii) $CS(\mathbf{x}^*(u), \mathbf{t}^M, u) \geq CS(\mathbf{x}^*(\tilde{u}), \mathbf{t}^M, \tilde{u})$.

We are now ready for our main result.

THEOREM 6.5. *Let $C(\mathbf{x}^*, \mathbf{t}^B) \leq \beta \cdot C(\mathbf{x}^*, \mathbf{t}^M)$ for some $\beta \in [1, \infty)$.⁷ Then*

$$CS(\mathbf{x}^*(u), \mathbf{t}^B, u) \geq CS(\mathbf{x}^*(\tilde{u}), \mathbf{t}^M, \tilde{u}),$$

where $\tilde{u}(x) = \frac{\beta}{2\beta-1} \cdot u(x)$.

PROOF. Recall, \mathbf{x}^* is the optimal flow when all r units are being routed. By the optimal flow characterization of Beckmann et al. [4], we have that $r_{x^*(u)} < r$ for all u with $C(\mathbf{x}^*, \hat{\mathbf{t}}) > r \cdot u(r)$, and $r_{x^*(u)} = r$ for all u with $C(\mathbf{x}^*, \hat{\mathbf{t}}) \leq r \cdot u(r)$.

We consider the following three cases: (1) $u(r) < C(\mathbf{x}^*, \hat{\mathbf{t}})/r$, (2) $C(\mathbf{x}^*, \hat{\mathbf{t}})/r \leq u(r) < \frac{2\beta-1}{\beta} \cdot C(\mathbf{x}^*, \hat{\mathbf{t}})/r$ and (3) $u(r) \geq \frac{2\beta-1}{\beta} \cdot C(\mathbf{x}^*, \hat{\mathbf{t}})/r$.

Case (1). Assume that $u(r) < C(\mathbf{x}^*, \hat{\mathbf{t}})/r$. Then $r_{x^*(u)} < r$ and thus

$$CS(\mathbf{x}^*(u), \mathbf{t}^B, u) = CS(\mathbf{x}^*(u), \mathbf{t}^M, u) \geq CS(\mathbf{x}^*(\tilde{u}), \mathbf{t}^M, \tilde{u}),$$

where the equality follows by Lemma 6.3 and the inequality by Lemma 6.4.

Case (2). Assume that $C(\mathbf{x}^*, \hat{\mathbf{t}})/r \leq u(r) < \frac{2\beta-1}{\beta} \cdot C(\mathbf{x}^*, \hat{\mathbf{t}})/r$. Then $r_{x^*(u)} = r$ and since $\tilde{u}(r) = \frac{\beta}{2\beta-1} \cdot u(r) < C(\mathbf{x}^*, \hat{\mathbf{t}})/r$, we have $r_{x^*(\tilde{u})} < r$. Thus

$$CS(\mathbf{x}^*(u), \mathbf{t}^B, u) \geq CS(\mathbf{x}^*(\tilde{u}), \mathbf{t}^B, \tilde{u}) = CS(\mathbf{x}^*(\tilde{u}), \mathbf{t}^M, \tilde{u}),$$

⁷Notice that, first, there is always such a β , since if all the demand is routed, the feasible space of MPT2 is at least as big as that of BMT2, while if the demand is partially routed we assume $\mathbf{t}^B = \mathbf{t}^M$, and, second, if $r_{x^*(u)} = r$, then $\mathbf{t}^B = \mathbf{t}^{BMT}$ and $\mathbf{t}^M = \mathbf{t}^{MPT}$, and so, for particular values of β , one can use Theorem 5.2 and Theorem 5.5.

where the inequality follows by Lemma 6.4 and the equality by Lemma 6.3.

Case (3). Assume that $u(r) \geq \frac{2\beta-1}{\beta} \cdot C(\mathbf{x}^*, \hat{t})/r$. Then $r_{x^*(u)} = r$ and since $\tilde{u}(r) = \frac{\beta}{2\beta-1} \cdot u(r) \geq C(\mathbf{x}^*, \hat{t})/r$, we have $r_{x^*(\tilde{u})} = r$. Thus,

$$\begin{aligned} C(\mathbf{x}^*, \mathbf{t}^B) - C(\mathbf{x}^*, \mathbf{t}^M) &\leq \left(1 - \frac{1}{\beta}\right) \cdot C(\mathbf{x}^*, \mathbf{t}^B) \\ &\leq \left(1 - \frac{1}{\beta}\right) \cdot C(\mathbf{x}^*, \hat{t}) \\ &\leq r \cdot \frac{\beta-1}{2\beta-1} \cdot u(r) - \int_0^r u(x)dx + \int_0^r u(x)dx \\ &\leq - \int_0^r \tilde{u}(x)dx + \int_0^r u(x)dx = U(r) - \tilde{U}(r), \end{aligned}$$

where the first inequality follows from $C(\mathbf{x}^*, \mathbf{t}^B) \leq \beta \cdot C(\mathbf{x}^*, \mathbf{t}^M)$, which is true by assumption, the second from $C(\mathbf{x}^*, \mathbf{t}^B) \leq C(\mathbf{x}^*, \hat{t})$, the third from $\frac{\beta}{2\beta-1} \cdot u(r) \geq C(\mathbf{x}^*, \hat{t})/r$, and the fourth from $r \cdot u(r) \leq \int_0^r u(x)dx$ and the definition of $\tilde{u}(x)$. Rearranging terms yields

$$CS(\mathbf{x}^*(u), \mathbf{t}^B, u) = U(r) - C(\mathbf{x}^*, \mathbf{t}^B) \geq \tilde{U}(r) - C(\mathbf{x}^*, \mathbf{t}^M) = CS(\mathbf{x}^*(\tilde{u}), \mathbf{t}^M, \tilde{u}),$$

as needed. \square

REMARK 6.6. *Theorem 6.5 cannot be extended to multicommodity networks. For an example, consider the network of Figure 5. Commodities $o - d_1$, $o - d_2$ have a maximum demand of 1. If $u^1(x) = 8$ for $x \in [0, 1]$ and $u^2(x) = 4$ for $x \in [0, 1]$, then the optimal flow is 3/4 for commodity 1 and 1 for commodity 2. The unique optimal toll vector is $(3/4, 7/4)$; inducing a consumer surplus of 1/2.*

If $u^1(x) = 4$ for $x \in [0, 1]$ and $u^2(x) = 2$ for $x \in [0, 1]$, then the optimal flow is 0 for commodity 1 and 1 for commodity 2. The unique optimal toll vector is $(0, 0)$; inducing a consumer surplus of 1.

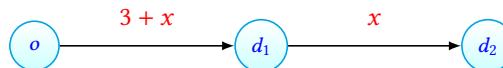


Fig. 5. The network of Remark 6.6.

REFERENCES

- [1] D. Acemoglu and A. Ozdaglar. 2007. Competition and Efficiency in Congested Markets. *Mathematics of Operations Research* 32, 1 (2007), 1–31.
- [2] D. Acemoglu and A. Ozdaglar. 2007. Competition in parallel-serial networks. *IEEE Journal on selected areas in communications* 25, 6 (2007), 1180–1192.
- [3] E. Anshelevich and S. Sekar. 2015. Price Competition in Networked Markets: How Do Monopolies Impact Social Welfare. In *WINE 2015, Amsterdam, The Netherlands, December 9–12, 2015, Proceedings*. 16–30.
- [4] M. Beckmann, C. McGuire, and C. Winsten. 1956. *Studies in the economics of transportation*. Yale Univ. Press.
- [5] C. Berge. 1963. *Topological Spaces: Including a Treatment of Multi-valued Functions, Vector Spaces, and Convexity*. Dover Publications.
- [6] V. Bonifaci, M. Salek, and G. Schäfer. 2011. Efficiency of Restricted Tolls in Non-atomic Network Routing Games. In *SAGT 2011, Amalfi, Italy, October 17–19, 2011, Proceedings*. 302–313.
- [7] P. Briest, M. Hoefer, and P. Krysta. 2012. Stackelberg Network Pricing Games. *Algorithmica* 62, 3–4 (2012).
- [8] I. Caragiannis, X. Chatzigeorgiou, P. Kanellopoulos, G.A. Krimpas, N. Protopapas, and A. Voudouris. 2017. Efficiency and Complexity of Price Competition Among Single-Product Vendors. *Artificial Intelligence* 248 (2017), 9–25.
- [9] L. Castelli, M. Labb, and A. Violin. 2017. Network pricing problem with unit toll. *Networks* 69, 1 (2017), 83–93.

- [10] S. Chawla and F. Niu. 2009. The Price of Anarchy in Bertrand Games. In *EC '09, Stanford, CA, USA, July 6-10, 2009, Proceedings*. ACM, New York, NY, USA, 305–314.
- [11] S. Chawla, F. Niu, and T. Roughgarden. 2008. Bertrand competition in networks. In *SAGT 2008, Paderborn, Germany, April/May, 2008, Proceedings*. 70–82.
- [12] R. Cole, Y. Dodis, and T. Roughgarden. 2003. Pricing network edges for heterogeneous selfish users. In *STOC '03, San Diego, CA, USA, June 9-11, 2003, Proceedings*, Vol. 3. 444–467.
- [13] J.R. Correa and N.E. Stier-Moses. 2011. Wardrop equilibria. *Wiley encyclopedia of op. research and man. science* (2011).
- [14] R. Dial. 1999. Minimal-revenue congestion pricing part I: A fast algorithm for the single-origin case. *Transportation Research Part B: Methodological* 33, 3 (1999), 189 – 202.
- [15] H. Dixon. 2001. *Surfing Economics : Essays for the Inquiring Economist*. Palgrave Macmillan.
- [16] A. Epstein, M. Feldman, and Y. Mansour. 2009. Efficient graph topologies in network routing games. *Games and Economic Behavior* 66, 1 (2009), 115 – 125.
- [17] L. Fleischer. 2005. Linear Tolls Suffice: New Bounds and Algorithms for Tolls in Single Source Networks. *Theoretical Computer Science* 348 (2005), 217–225.
- [18] D. Fotakis, D. Kalimeris, and T. Lianeas. 2015. Improving Selfish Routing for Risk-Averse Players. In *WINE 2015, Amsterdam, The Netherlands, December 9-12, 2015, Proceedings*. 328–342.
- [19] A. Gonzalez. 2016. Controversias por reducciones en la demanda en las concesiones de carreteras en Chile. *Revista de Derecho Económico* 76 (2016).
- [20] T. Harks, M. Schröder, and D. Vermeulen. 2018. Toll Caps in Privatized Road Networks. (2018). arXiv preprint arXiv:1802.10514.
- [21] A. Hayrapetyan, E. Tardos, and T. Wexler. 2007. A network pricing game for selfish traffic. *Distributed Computing* 19, 4 (2007). 255–266.
- [22] T. Jelinek, M. Klaas, and G. Schäfer. 2014. Computing Optimal Tolls with Arc Restrictions and Heterogeneous Players. In *STACS 2014, Lyon, France, March 5-8, 2014, Proceedings*.
- [23] P. Kleer and G. Schäfer. 2016. The Impact of Worst-Case Deviations in Non-Atomic Network Routing Games. In *SAGT 2016, Liverpool, UK, September 19-21, 2016, Proceedings*. 129–140.
- [24] A. Kuiteing, P. Marcotte, and G. Savard. 2016. Network Pricing of Congestion-Free Networks: The Elastic and Linear Demand Case. *Transportation Science* 51, 3 (2016), 791–806.
- [25] H. Lin, T. Roughgarden, E. Tardos, and A. Walkover. 2011. Stronger Bounds on Braess's Paradox and the Maximum Latency of Selfish Routing. *SIAM Journal on Discrete Mathematics* 25, 4 (2011), 1667–1686.
- [26] J. Musacchio. 2009. The Price of Anarchy in Parallel-Serial Competition with Elastic Demand. (2009). Technical Report No. UCSC-SOE-09-20.
- [27] J. Musacchio and S. Wu. 2007. The Price of Anarchy in a Network Pricing Game. In *Allerton Conference, Monticello, IL, USA, September 26-28, 2007, Proceedings*.
- [28] E. Nikolova and N. Stier-Moses. 2015. The Burden of Risk Aversion in Mean-Risk Selfish Routing. In *EC '15, Portland, OR, USA, June 15-19, 2015, Proceedings*. 489–506.
- [29] A. Ozdaglar. 2008. Price Competition with Elastic Traffic. *Networks* 52, 3 (2008), 141–155.
- [30] A. Ozdaglar and R. Srikant. 2007. Incentives and pricing in communication networks. In *Algorithmic Game Theory*. Cambridge Press, 571–591.
- [31] C. Papadimitriou and G. Valiant. 2010. A New Look at Selfish Routing. In *ICS 2010, Beijing, China, January 5-7, 2010, Proceedings*. 178–187.
- [32] A. Pigou. 1920. *The economics of welfare*. Macmillan.
- [33] T. Roughgarden. 2005. *Selfish Routing and the Price of Anarchy*. MIT press.
- [34] T. Roughgarden and E. Tardos. 2002. How bad is selfish routing? *J. ACM* 49, 2 (2002), 236–259.