Demand-Independent Optimal Tolls

Riccardo Colini-Baldeschi
Core Data Science Group, Facebook Inc., 1 Rathbone Place, London, W1T 1FB, UK
rickuz@fb.com
https://orcid.org/0000-0001-5739-1178

Max Klimm
School of Business and Economics, HU Berlin, Spandauer Str. 1, 10099 Berlin, Germany
max.klimm@hu-berlin.de
https://orcid.org/0000-0002-9061-2267

Marco Scarsini
Dipartimento di Economia e Finanza, LUISS, Viale Romania 32, 00197 Roma, Italy
marco.scarsini@luiss.it
https://orcid.org/0000-0001-6473-794X

Abstract

Wardrop equilibria in nonatomic congestion games are in general inefficient as they do not induce an optimal flow that minimizes the total travel time. Network tolls are a prominent and popular way to induce an optimum flow in equilibrium. The classical approach to find such tolls is marginal cost pricing which requires the exact knowledge of the demand on the network. In this paper, we investigate under which conditions demand-independent optimum tolls exist that induce the system optimum flow for any travel demand in the network. We give several characterizations for the existence of such tolls both in terms of the cost structure and the network structure of the game. Specifically we show that demand-independent optimum tolls exist if and only if the edge cost functions are shifted monomials as used by the Bureau of Public Roads. Moreover, non-negative demand-independent optimum tolls exist when the network is a directed acyclic multi-graph. Finally, we show that any network with a single origin-destination pair admits demand-independent optimum tolls that, although not necessarily non-negative, satisfy a budget constraint.

2012 ACM Subject Classification Theory of computation → Network games

Keywords and phrases nonatomic congestion games, efficiency of equlibria, tolls

Digital Object Identifier 10.4230/LIPIcs.ICALP.2018.151

Related Version A full version of the paper is available at `https://arxiv.org/abs/1708.02737`.

Acknowledgements Riccardo Colini-Baldeschi and Marco Scarsini are members of GNAMPA-INdAM. Max Klimm gratefully acknowledges the support and hospitality of LUISS during a visit in which this research was initiated. The research of Max Klimm was carried out in the framework of Matheon supported by Einstein Foundation Berlin.

1 Introduction

The impact of selfish behavior on the efficiency of traffic networks is a longstanding topic in the algorithmic game theory and operations research literature. Already more than half a century ago, Wardrop [41] stipulated a main principle of a traffic equilibrium that – in light of the omnipresence of modern route guidance systems – is as relevant as ever: "The

© Riccardo Colini-Baldeschi, Max Klimm, and Marco Scarsini;
licensed under Creative Commons License CC-BY

journey times on all the routes actually used are equal, and less than those which would be experienced by a single vehicle on any unused route." This principle can be formalized by modeling traffic as a flow in a directed network where edges correspond to road segments and vertices correspond to crossings. Each edge is endowed with a cost function that maps the total amount of traffic on it to a congestion cost that each flow particle traversing the edge has to pay. Further, we are given a set of commodities, each specified by an origin, a destination, and a flow demand. In this setting, a Wardrop equilibrium is a multi-commodity flow such that for each commodity the total cost of any used path is not larger than the total cost of any other path linking the commodity's origin and destination. Popular cost functions, put forward for the use in traffic models by the U.S. Bureau of Public Roads [40] are of the form

$$c_e(x_e) = t_e\left(1 + \alpha\left(\frac{x_e}{k_e}\right)^{\beta}\right), \tag{1}$$

where x_e is the traffic flow along edge e, $t_e \geq 0$ is the free-flow travel time, k_e is the capacity of edge e and α and β are parameters fitted to the model.

While Wardrop equilibria are guaranteed to exist under mild assumptions on the cost functions [4], it is well known that they are inefficient in the sense that they do not minimize the overall travel time of all commodities [34]. A popular mechanism to decrease the inefficiency of selfish routing are congestion tolls. A toll is a payment that the system designer defines for each edge of the graph and that has to be paid by each flow particle traversing the edge. By carefully choosing the edges' tolls, the system designer can steer the Wardrop equilibrium in a favorable direction. A classic approach first due to Pigou [34] is *marginal cost pricing* where the toll of each edge is equal to difference between the marginal social cost and the marginal private cost of the system optimum flow on that edge. Marginal cost pricing induces the system optimal flow – the one that minimizes the overall travel time – as a Wardrop equilibrium [4]. Congestion pricing is not only an interesting theoretical issue that links system optimal flows and traffic equilibria, but also a highly relevant tool in practice, as various cities of the world, including Stockholm, Singapore, Bergen, and London, implement congestion pricing schemes to mitigate congestion [22, 39].

The problem. Marginal cost pricing is an elegant way to induce the system optimum flow as a Wardrop equilibrium, but the concept crucially relies on the exact knowledge of the travel demand. As an example consider the Pigou network in Figure 1a for an arbitrary flow demand of $\mu > 0$ going from o to d. The optimal flow only uses the lower edge with cost function $c(x) = x$ as long as $\mu \leq 1/2$. For demands $\mu > 1/2$, a flow of $1/2$ is sent along the lower edge and the remaining flow of $\mu - 1/2$ is sent along the upper edge. Using marginal cost pricing the toll of the lower edge is $\min\{\mu, 1/2\}$ and no toll is to be payed for the upper edge, see Figure 1b.

This toll scheme is problematic as it depends on the flow demand μ in the network. In particular, when the demand is estimated incorrectly, the resulting tolls will be sub-optimal. Assume the network designer expects a flow of $\mu = 1/4$ and, thus, sets a (marginal cost) toll of $\tau_2 = 1/4$ on the lower link. When the actual total flow demand is higher than expected and equal to $\mu = 1$, a fraction $1/4$ of the flow uses the upper edge and $3/4$ of the flow uses the lower edge resulting in a cost of $1/4 + 9/16$. However, it is optimal to induce an equal split between the edge which can be achieved by a toll of $1/2$ on the lower edge. Quite strikingly, the toll of $\tau_2 = 1/2$ is optimal for all possible demand values as it always induces the optimal flow. For demands less than $1/2$, a toll of $1/2$ on the lower edge does not hinder any flow particle from using the lower edge, which is optimal. On the other hand, for any demand

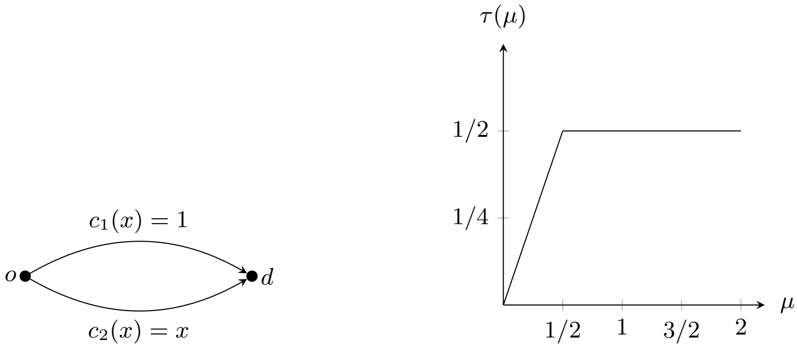

(a) Pigou network. **(b)** Marginal cost toll for lower edge.

Figure 1 Dependence of the marginal cost tolls for the Pigou network.

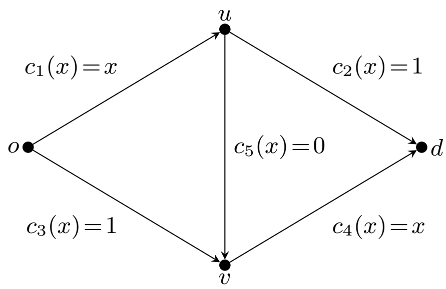

Figure 2 Braess network.

larger 1/2, the toll of 1/2 on the lower edge forces the flow on this edge to not exceed 1/2, which is optimal as well.

The situation is even more severe for the Braess network in Figure 2. When the system designer expects a traffic demand of 1 going from o to d, marginal cost pricing fixes a toll of 1 on both the upper left and the lower right edge (both with cost function $c(x) = x$). When the demand is lower than expected, say, $\mu = 1/2$, under marginal cost pricing, the flow is split equally between the lower and the upper path leading to a total cost of 5/4. The optimal flow with flow value 1/2, however, only uses the zig-zag-path $o \to u \to v \to d$ with cost 1. It is interesting to note that this flow is actually equal to the Wardrop equilibrium without any tolls. To conclude this example, marginal cost pricing may actually *increase* the total cost of the Wardrop equilibrium when the travel demand is estimated incorrectly. We note that also in the Braess graph, there is a distinct toll vector that enforces the optimum flow as a Wardop equilibrium *for any demand*. We will see that by setting a toll of 1 on the central edge from u to v, the Wardrop equilibrium for any flow demand is equal to the respective optimum flow.

We conclude that for both the Pigou network and the Braess network, marginal cost pricing is not robust with respect to changes in the demand since wrong estimates of the travel demand lead to sub-optimal tolls. Since such changes may occur frequently in road networks (e.g., due to sudden weather changes, accidents, or other unforeseen events), marginal cost pricing does not use the full potential of congestion pricing and may even be harmful for the traffic. On the other hand, for both networks, there exist tolls that enforce the optimum flow as an equilibrium for *any* flow demand. In this paper, we systematically study conditions for such demand-independent optimal tolls to exist.

Our results. In this paper we study the existence of demand-independent optimum tolls (DIOTs) that induce the optimum flow as an equilibrium for *any* flow demand. We give a precise characterization in terms of the cost structure on the edges for DIOTs to exists. Specifically, we show that DIOTs exist for any network where the cost of each edge is a BPR-type function, i.e.,

$$c_e(x) = t_e + a_e x_e^{\beta} \quad \text{for all } e \in \mathcal{E}, \tag{2}$$

where $t_e, a_e \in \mathbb{R}_+$ are arbitrary while $\beta \in \mathbb{R}_+$ is a common constant for all edges $e \in \mathcal{E}$. This existence result holds regardless of the topology of the network and on the number of origin-destination (O/D) pairs. On the other hand, for any cost function that is not of the form as in (2), there is a simple network consisting of two parallel edges with cost function c and cost function $c + b$ for some $b > 0$ that does not admit a DIOT. Our existence result for networks with BPR-type cost functions is proven in terms of a characterization that uniquely determines the sum of the tolls along each path that is used by the optimum flow for some demand.

In general the DIOTs used in the characterization may use negative tolls as well. We provide an example of a network with BPR-type cost functions where a non-negative DIOT does not exist. Besides conditions on the costs, conditions on the network are needed to guarantee the existence of non-negative DIOTs. We show that non-negative DIOTs exist for directed acyclic multi-graphs (DAMGs) with BPR-type cost functions, like the Pigou network and the Braess network discussed in Section 1. This condition on the network is sufficient, but not necessary for the existence of non-negative DIOTs.

Under a weaker condition than DAMG, we prove the existence of DIOTs that follow a budget constraint of non-negativity of the total amount of tolls. This condition is satisfied by networks with a single O/D pair.

Due to space constraint, some of the proofs are omitted and can be found in the arXiv version [15].

Related work. Marginal cost pricing as a means to reduce the inefficiency of selfish resource allocation was first proposed by Pigou [34] and subsequently discussed by Knight [31]. Wardrop [41] introduced the notion of a traffic equilibrium where each flow particle only uses shortest paths. Beckmann et al. [4] showed that marginal cost pricing always induces the system optimal flow as a Wardrop equilibrium. The set of feasible tolls that induce optimal flows was explored in [5, 26, 32]. They showed that the set of optimal tolls can be described by a set of linear equations and inequalities.

This characterization led to various developments regarding the optimization of secondary objectives of the edge tolls, such as the minimization of the tolls collected from the users [1, 16, 17], or the minimzation of the number of edges that have positive tolls [2, 3]. A problem closely related to the latter is to compute tolls for a given subset of edges with the objective to minimize the total travel time of the resulting equilibrium. Hoefer et al. [27] showed that this problem is NP-hard for general networks, and gave an efficient algorithm for parallel edges graphs with affine cost functions. Harks et al. [25] generalized their result to arbitrary cost functions satisfying a technical condition. Bonifaci et al. [8] studied generalizations of this problem with further restrictions on the set of feasible edge tolls.

For heterogenous flow particles that trade off money and time differently, marginal cost pricing cannot be applied to find tolls that induce the system optimum flow. In this setting, Ciole et al. [13] showed the existence of a set of tolls enforcing the system optimal flow, when there is a single commodity in the network. Similarly, Yang and Huang [42] studied how to design toll structure when there are users with different toll sensitivity. Fleischer [18] showed that in single source series-parallel networks the tolls have to be linear in the latency of the maximum latency path. Karakostas and Kolliopoulos [28] and Fleischer et al. [19] independently generalized this result to arbitrary networks. Karakostas and Kolliopoulos [30] showed similar results for players with elastic demands. Han et al. [24] extended the previous results to different classes of cost functions.

Most of the literature assumes that the charged tolls cause no disutilty to the network users. For the case where tolls contribute to the cost, Cole et al. [14] showed that marginal cost tolls do not improve the equilibrium flow for a large class of instances, including all instances with affine costs. They further showed that for these networks it is NP-hard to approximate the minimal total cost that can be achieved as a Wardrop equilibrium with tolls. Karakostas and Kolliopoulos [29] proved that the total disutility due to taxation is bounded with respect to the social optimum for large classes of latency functions. Moreover, they showed that, if both the tolls and the latency are part of the social cost, then for some latency functions the coordination ratio improves when taxation is used. For networks of parallel edges, Christodoulou et al. [12] studied a generalization of edge tolls where cost functions are allowed to increase in an arbitrary way. They showed that for affine cost functions, the price of anarchy is strictly better than in the original network, even when the demand is not known.

Brown and Marden [9, 10] studied how marginal tolls can create perverse incentives when users have different sensitivity to the tolls and how it possible to circumvent this problem. Caragiannis et al. [11] studied the optimal toll problem for atomic congestion games. They proved that in the atomic case the optimal system performance cannot be achieved even in very simple networks. On the positive side they shown that there is a way to assign tolls to edges such that the induced social cost is within a factor of 2 to the optimal social cost. Singh [38] observed that marginal tolls weakly enforce optimal flows. Fotakis and Spirakis [21] showed that in series-parallel networks with increasing cost functions the optimal social cost can be induced with tolls. Fotakis et al. [20] were the first to consider the problem of defining tolls for heterogeneous users in atomic congestion games with unsplittable flow. In [7] the problem of defining tolls for atomic congestion games with polynomial cost functions was considered for the first time. Meir and Parkes [33] discussed how in atomic congestion games with marginal tolls multiple equilibria are near-optimal when there is a large number of players.

Sandholm [36, 37] studied tolls from a mechanism design perspective where the social planner has no information over the preferences of the users and has limited ability to observer the users' behavior. Bhaskar et al. [6] studied the problem of achieving target flows as equilibria of a nonatomic routing game without knowing the underlying latency functions and showed that a given target flow can be achieved using a polynomial number of queries to an oracle that takes tolls as input and outputs the resulting equilibrium flow. Roth et al. [35] considered the problem of finding optimal tolls in a polynomial number of rounds when latency functions are unknown. They solved the problem by embedding it in a broad class of Stackelberg game.

2 Model and preliminaries

In this section, we present some notation and basic definitions that are used in the sequel. We start with the underlying network model and will then introduce equilibria and tolls.

Network model. We consider a finite directed multi-graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with vertex set $\mathcal{V}$ and edge set $\mathcal{E}$. We call $(v \to v')$ the set of all edges e whose tail is v and whose head is v'. We assume that there is a finite set of *origin-destination* (O/D) pairs $i \in \mathcal{I}$, each with an individual *traffic demand* $\mu^i \geq 0$ that has to be routed from an origin $o^i \in \mathcal{V}$ to a destination $d^i \in \mathcal{V}$ via $\mathcal{G}$. Denote the demand vector by $\boldsymbol{\mu} = (\mu^i)_{i\in\mathcal{I}}$. We call $\mathcal{P}^i$ the set of (simple) paths joining o^i to d^i, where each path $p \in \mathcal{P}^i$ is a finite sequence of edges such that the head of each edge meets the tail of the subsequent edge. For as long as all pairs (o^i, d^i) are different, the sets $\mathcal{P}^i$ are disjoint. Call $\mathcal{P} := \bigcup_{i\in\mathcal{I}} \mathcal{P}^i$ the union of all such paths.

Each path p is traversed by a *flow* $f_p \in \mathbb{R}_+$. Call $\boldsymbol{f} = (f_p)_{p\in\mathcal{P}}$ the vector of flows in the network. The set of feasible flows for $\boldsymbol{\mu}$ is defined as

$$\mathcal{F}(\boldsymbol{\mu}) = \Big\{ \boldsymbol{f} \in \mathbb{R}_+^{\mathcal{P}} : \textstyle\sum_{p\in\mathcal{P}^i} f_p = \mu^i \text{ for all } i \in \mathcal{I} \Big\}. \tag{3}$$

In turn, a routing flow $\boldsymbol{f} \in \mathcal{F}(\boldsymbol{\mu})$ induces a *load* on each edge $e \in \mathcal{E}$ as

$$x_e = \sum_{p\ni e} f_p. \tag{4}$$

We call $\boldsymbol{x} = (x_e)_{e\in\mathcal{E}}$ the corresponding *load profile* on the network. For each $e \in \mathcal{E}$ consider a nondecreasing, continuous *cost function* $c_e \colon \mathbb{R}_+ \to \mathbb{R}_+$. Denote $\boldsymbol{c} = (c_e)_{e\in\mathcal{E}}$. If $\boldsymbol{x}$ is the load profile induced by a feasible routing flow $\boldsymbol{f}$, then the incurred delay on edge $e \in \mathcal{E}$ is given by $c_e(x_e)$; hence, with a slight abuse of notation, the associated cost of path $p \in \mathcal{P}$ is given by the expression $c_p(\boldsymbol{f}) \equiv \sum_{e\in p} c_e(x_e)$. We call the tuple $\Gamma = (\mathcal{G}, \mathcal{I}, \boldsymbol{c})$ a (*nonatomic*) *routing game*.

Equilibrium Flows and Optimal Flows. A routing flow $\boldsymbol{f}^*$ is a *Wardrop equilibrium* (WE) of Γ if, for all $i \in \mathcal{I}$, we have:

$$c_p(\boldsymbol{f}^*) \leq c_{p'}(\boldsymbol{f}^*) \quad \text{for all } p, p' \in \mathcal{P}^i \text{ such that } f_p^* > 0.$$

This concept was introduced by Wardrop [41]. Beckmann et al. [4] showed that Wardrop equilibria are the optimal solutions to the convex optimization problem

$$\begin{aligned} \min \; & \sum_{e\in\mathcal{E}} \int_0^{x_e} c_e(s)\,\mathrm{d}s \\ \text{s.t.:} \; & x_e = \sum_{p\ni e} f_p \\ & \boldsymbol{f} \in \mathcal{F}, \end{aligned}$$

and, thus, are guaranteed to exist. A *social optimum* (SO) is a flow that minimizes the overall travel time, i.e., it solves the following total cost minimization problem:

$$\begin{aligned} \min \; & L(\boldsymbol{f}) = \sum_{p\in\mathcal{P}} f_p c_p(\boldsymbol{f}), \\ \text{s.t.:} \; & \boldsymbol{f} \in \mathcal{F}. \end{aligned} \tag{5}$$

As shown in [4], all Wardrop equilibria have the same social cost. We write $\mathrm{Eq}(\Gamma) = L(\boldsymbol{f}^*)$ and $\mathrm{Opt}(\Gamma) = \min_{\boldsymbol{f}\in\mathcal{F}} L(\boldsymbol{f})$, where $\boldsymbol{f}^*$ is a Wardrop equilibrium of Γ. The game's *price of anarchy* (PoA) is then defined as $\mathrm{PoA}(\Gamma) = \mathrm{Eq}(\Gamma)/\,\mathrm{Opt}(\Gamma)$. It is known that Wardrop equilibria need not minimize the social cost, in that case $\mathrm{PoA}(\Gamma) > 1$. For a pair $i \in \mathcal{I}$, we denote the set of paths that are eventually used in a optimum flow for some demand vector $\boldsymbol{\mu}$ by

$$\tilde{\mathcal{P}}^i = \{p \in \mathcal{P}^i : \tilde{\boldsymbol{f}}_p(\boldsymbol{\mu}) > 0 \text{ for some demand } \boldsymbol{\mu} \text{ and corresponding social optimum } \tilde{\boldsymbol{f}}(\boldsymbol{\mu})\}.$$

Here and in the following we write $\boldsymbol{f}(\boldsymbol{\mu})$ instead of $\boldsymbol{f}$ when we want to indicate the corresponding demand vector $\boldsymbol{\mu}$.

Tolls. We want to explore the possibility of imposing tolls on the edges of the network in such a way that the equilibrium flow of the game with tolls produces a flow that is a solution of the original miminization problem (5). In other words, we want to see whether it is possible to achieve an optimum flow as an equilibrium of a modified game.

We call $\boldsymbol{\tau} = (\tau_e)_{e\in\mathcal{E}} \in \mathbb{R}^{\mathcal{E}}$ a *toll vector*. Note that we allow both for positive and negative tolls. We call c_e^τ the cost of edge e under the toll τ, i.e., $c_e^\tau(x_e) := c_e(x_e) + \tau_e$. Similarly $c_p^\tau(\boldsymbol{f}) := \sum_{e\in p} c_e^\tau(x_e)$. Define $\Gamma^{\boldsymbol{\tau}} := (\mathcal{G}, \mathcal{I}, \boldsymbol{c}^{\boldsymbol{\tau}})$. A toll vector $\boldsymbol{\tau}$ that for each demand vector $\boldsymbol{\mu} \in \mathbb{R}_+^{\mathcal{I}}$ enforces the corresponding system optimum as the equilibrium in $\Gamma^{\boldsymbol{\tau}}$ is called demand-independent optimal toll.

▶ **Definition 1** (Demand-independent optimal toll (DIOT))**.** Let $\Gamma = (\mathcal{G}, \mathcal{I}, \boldsymbol{c})$. A toll vector $\boldsymbol{\tau} \in \mathbb{R}^{\mathcal{E}}$ is called demand-independent optimal toll (DIOT) for Γ if for every demand vector $\boldsymbol{\mu} \in \mathbb{R}_+^{\mathcal{I}}$ every corresponding equilibrium with tolls $\boldsymbol{f}^{\boldsymbol{\tau}}(\boldsymbol{\mu}) \in \mathrm{Eq}(\Gamma^{\boldsymbol{\tau}})$ is optimal for Γ, i.e., $L(\boldsymbol{f}^{\boldsymbol{\tau}}(\boldsymbol{\mu})) = \sum_{p\in\mathcal{P}} f_p^{\tau}(\boldsymbol{\mu}) c_p(\boldsymbol{f}^{\boldsymbol{\tau}}(\boldsymbol{\mu})) \leq L(\boldsymbol{f}(\boldsymbol{\mu})) = \sum_{p\in\mathcal{P}} f_p c_p(\boldsymbol{f}(\boldsymbol{\mu}))$ for all $\boldsymbol{f}(\boldsymbol{\mu}) \in \mathcal{F}(\boldsymbol{\mu})$.

In Section 1 we visited two games, the Pigou networks and Braess' paradox, that admit a DIOT. The aim of this paper is to characterize the networks $\Gamma = (\mathcal{G}, \mathcal{I}, \boldsymbol{c})$ for which DIOTs exist.

3 BPR-type cost functions

In this section, we give a complete characterization of the sets of cost functions that admit a DIOT. On the positive side, we will show that any network with BPR-type cost functions admits a DIOT, independently of the number of commodities and the network topology. On the other hand, we show a strong lower bound proving that for any non-BPR cost function there is a single-commodity game on two parallel edges with costs functions c and $c + t$ for some $t \in \mathbb{R}_+$ that does not admit a DIOT. Formally, for $\beta > 0$, let

$$\mathcal{C}_{\mathrm{BPR}}(\beta) = \{c : \mathbb{R}_+ \to \mathbb{R}_+ : c(x) = t_c + a_c x^\beta \text{ for all } x \geq 0,\, a, t \in \mathbb{R}_+\}$$

be the set of BPR-type cost functions with degree β.

The following theorem gives a sufficient condition for games with BPR-type cost functions to admit a DIOT.

▶ **Theorem 2.** *Consider a game $\Gamma = (\mathcal{G}, \mathcal{I}, \boldsymbol{c})$ such that there is $\beta \in \mathbb{R}_+$ with $c_e \in \mathcal{C}_{\mathrm{BPR}}(\beta)$ for all $e \in \mathcal{E}$. Let $\boldsymbol{\tau}$ be a toll vector such that*

$$\sum_{e\in p}\Big(\tau_e + \tfrac{\beta}{\beta+1} t_e\Big) \leq \sum_{e\in p'}\Big(\tau_e + \tfrac{\beta}{\beta+1} t_e\Big). \tag{6}$$

for all $i \in \mathcal{I}$ and all $p \in \tilde{\mathcal{P}}^i$ and all $p' \in \mathcal{P}^i$. Then, $\boldsymbol{\tau}$ is a DIOT.

Proof. Fix a demand vector $\boldsymbol{\mu} \in \mathbb{R}_+^{\mathcal{I}}$ and a corresponding optimum flow $\tilde{\boldsymbol{f}}(\boldsymbol{\mu})$ in Γ arbitrarily. We denote by $\tilde{x}_e = \sum_{p \in \mathcal{P}: e \in p} \tilde{f}_p(\boldsymbol{\mu})$ the load imposed on edge e by $\tilde{\boldsymbol{f}}(\boldsymbol{\mu})$. The local optimality conditions of $\tilde{\boldsymbol{f}}(\boldsymbol{\mu})$ imply that for all $i \in \mathcal{I}$ and all $p, p' \in \mathcal{P}^i$ with $\tilde{f}_p(\boldsymbol{\mu}) > 0$ we have

$$\sum_{e \in p} c_e(\tilde{x}_e) + c_e'(\tilde{x}_e)\tilde{x}_e \leq \sum_{e \in p'} c_e(\tilde{x}_e) + c_e'(\tilde{x}_e)\tilde{x}_e.$$

This implies that for all $i \in \mathcal{I}$ and all $p, p' \in \mathcal{P}^i$ with $\tilde{f}_p(\boldsymbol{\mu}) > 0$, there exists a non-negative constant $\lambda(p, p') \geq 0$ such that

$$\lambda(p, p') + \sum_{e \in p} \Big((\beta+1) a_e \tilde{x}_e^{\beta} + t_e \Big) = \sum_{e \in p'} \Big((\beta+1) a_e \tilde{x}_e^{\beta} + t_e \Big). \tag{7}$$

We proceed to show that when the toll vector $\boldsymbol{\tau}$ satisfies (6), then $\tilde{\boldsymbol{f}}(\boldsymbol{\mu})$ is a Wardop equilibrium of Γ^{τ}. To this end, consider arbitrary $p, p' \in \mathcal{P}^i$ with $\tilde{f}_p(\boldsymbol{\mu}) > 0$. We have

$$\begin{aligned}
\sum_{e \in p} \Big(c_e(\tilde{x}_e) + \tau_e \Big) &= \sum_{e \in p} \Big(a_e \tilde{x}_e^{\beta} + t_e + \tau_e \Big) \\
&= \sum_{e \in p} \Big((\beta+1) a_e \tilde{x}_e^{\beta} + t_e \Big) - \sum_{e \in p} \Big(\beta a_e \tilde{x}_e^{\beta} - \tau_e \Big) \\
&= \sum_{e \in p'} \Big((\beta+1) a_e \tilde{x}_e^{\beta} + t_e \Big) - \sum_{e \in p} \Big(\beta a_e \tilde{x}_e^{\beta} - \tau_e \Big) - \lambda(p, p'), \\
&= \sum_{e \in p'} \Big(c_e(\tilde{x}_e) + \tau_e \Big) + \sum_{e \in p'} \Big(\beta a_e \tilde{x}_e^{\beta} - \tau_e \Big) - \sum_{e \in p} \Big(\beta a_e \tilde{x}_e^{\beta} - \tau_e \Big) - \lambda(p, p'),
\end{aligned}$$

where the third equality comes from (7). By assumption, the toll vector $\boldsymbol{\tau}$ satisfies equation (6) which implies $\sum_{e \in p} \tau_e - \sum_{e \in p'} \tau_e \leq -\frac{\beta}{\beta+1} \Big(\sum_{e \in p} t_e - \sum_{e \in p'} t_e \Big)$. Therefore,

$$\begin{aligned}
\sum_{e \in p} \Big(c_e(\tilde{x}_e) + \tau_e \Big) &\leq \sum_{e \in p'} \Big(c_e(\tilde{x}_e) + \tau_e + \beta a_e \tilde{x}_e^{\beta} + \tfrac{\beta}{\beta+1} t_e \Big) - \sum_{e \in p} \Big(\beta a_e \tilde{x}_e^{\beta} + \tfrac{\beta}{\beta+1} t_e \Big) - \lambda(p, p') \\
&= \sum_{e \in p'} \Big(c_e(\tilde{x}_e) + \tau_e \Big) + \tfrac{\beta}{\beta+1} \lambda(p, p') - \lambda(p, p') \\
&= \sum_{e \in p'} \Big(c_e(\tilde{x}_e) + \tau_e \Big) - \tfrac{1}{\beta+1} \lambda(p, p'),
\end{aligned}$$

where the first equality stems from (7). The statement of the theorem follows from the fact that $\lambda(p, p') \geq 0$. ◀

As an immediate corollary of this theorem, we obtain the existence of DIOTs in arbitrary multi-commodity networks with BPR-type cost functions as long as negative tolls are allowed.

▶ **Corollary 3.** *Consider a game $\Gamma = (\mathcal{G}, \mathcal{I}, \boldsymbol{c})$ with BPR-type cost functions. Then there exists a DIOT for Γ.*

Proof. Setting the toll vector $\hat{\boldsymbol{\tau}} = (\hat{\tau}_e)_{e \in \mathcal{E}}$ as

$$\hat{\tau}_e = -\frac{\beta}{\beta+1} t_e, \tag{8}$$

obviously satisfies (6), so that the claim follows from Theorem 2. ◀

In the following, we call $\hat{\boldsymbol{\tau}}$ the *trivial DIOT*. The necessary condition of Theorem 2 implies in particular that $\sum_{e \in p} \tau_e + \frac{\beta}{\beta+1} t_e$ must be equal for all paths $p \in \tilde{\mathcal{P}}^i$ and all pairs $i \in \mathcal{I}$. We proceed to show that this is actually a necessary condition for a DIOT.

▶ **Theorem 4.** *Consider a game $\Gamma = (\mathcal{G}, \mathcal{I}, \boldsymbol{c})$ with BPR-type cost functions. If $\boldsymbol{\tau}$ is a DIOT for Γ, then*

$$\sum_{e \in p} \Big(\tau_e + \tfrac{\beta}{\beta+1} t_e \Big) = \sum_{e \in p'} \Big(\tau_e + \tfrac{\beta}{\beta+1} t_e \Big).$$

for all $i \in \mathcal{I}$ and all $p, p' \in \tilde{\mathcal{P}}^i$.

Proof. Let $i \in \mathcal{I}$ and $p^*, p^{**} \in \tilde{\mathcal{P}}^i$ be arbitrary. By definition of $\tilde{\mathcal{P}}^i$ there are demands vectors $\boldsymbol{\mu}, \boldsymbol{\mu}' \in \mathbb{R}_+^{\mathcal{I}}$ such that $\tilde{f}_{p^*}(\boldsymbol{\mu}) > 0$ and $\tilde{f}_{p^{**}}(\boldsymbol{\mu}') > 0$. Before we prove that $\sum_{e \in p^*} (\frac{\beta}{\beta+1} t_e + \tau_e) = \sum_{e \in p^{**}} (\frac{\beta}{\beta+1} t_e + \tau_e)$, we need some observation regarding the continuity of the optimal path flow functions.

Hall [23] showed that the path flow functions of a Wardrop equilibrium are continuous functions in the travel demand. For $\lambda \in [0,1]$, let $\boldsymbol{\mu}(\lambda) = (1-\lambda)\boldsymbol{\mu} + \lambda \boldsymbol{\mu}'$ parametrize the travel demands on the convex combination of $\boldsymbol{\mu}$ and $\boldsymbol{\mu}'$. Then, by Hall's result, there are continuous path flow functions $f_p(\boldsymbol{\mu}(\cdot)) : [0,1] \to \mathbb{R}_+^{\mathcal{E}}$ for all $p \in \mathcal{P}$ such that $\boldsymbol{f}(\boldsymbol{\mu}(\lambda))$ is a Wardrop equilibrium for the travel demand vector $\boldsymbol{\mu}(\lambda)$ for all $\lambda \in [0,1]$. As the system optimal flow $\tilde{\boldsymbol{f}}$ is a Wardrop equilibrium with respect to the marginal cost function (cf. [4]), the same holds for the system optimum flow vector $\tilde{\boldsymbol{f}}$, i.e., there are continuous path flow functions $\tilde{f}_p(\boldsymbol{\mu}(\cdot)) : [0,1] \to \mathbb{R}_+^{\mathcal{E}}$ for all $p \in \mathcal{P}$ such that $\tilde{\boldsymbol{f}}(\boldsymbol{\mu}(\lambda))$ is an optimal flow for the travel demand vector $\boldsymbol{\mu}(\lambda)$ for all $\lambda \in [0,1]$.

For a flow vector $\boldsymbol{f}$, let $\mathcal{S}^i(\boldsymbol{f}) = \{p \in \mathcal{P}^i : f_p > 0\}$ denote the support of $\boldsymbol{f}$ for $i \in \mathcal{I}$. For a possible support set $S \in 2^{\mathcal{P}^i}$, let

$$L^i(S) = \big\{ \lambda \in [0,1] : S = \mathcal{S}^i(\tilde{\boldsymbol{f}}(\boldsymbol{\mu}(\lambda))\big\}$$

denote the (possibly empty) set of values of λ for which the optimal flow $\tilde{\boldsymbol{f}}(\boldsymbol{\mu}(\lambda))$ has support S for i.

Consider an arbitrary support set S with $L^i(S) \neq \emptyset$ and an arbitrary $\lambda \in L^i(S)$. Since $\boldsymbol{\tau}$ is a DIOT for Γ, the optimal flow $\tilde{\boldsymbol{f}}(\boldsymbol{\mu}(\lambda))$ is a Wardrop equilibrium with respect to $\boldsymbol{\tau}$, i.e., $\sum_{e \in p} (c_e(\tilde{x}_e(\boldsymbol{\mu}(\lambda))) + \tau_e) = \sum_{e \in p'} (c_e(\tilde{x}_e(\boldsymbol{\mu}(\lambda))) + \tau_e)$ for all $p, p' \in S$ which implies

$$\sum_{e \in p} \big(a_e \tilde{x}_e(\boldsymbol{\mu}(\lambda))^{\beta} + t_e + \tau_e \big) = \sum_{e \in p'} \big(a_e \tilde{x}_e(\boldsymbol{\mu}(\lambda))^{\beta} + t_e + \tau_e \big)\,. \tag{9}$$

for all $p, p' \in S$. By the local optimality conditions of $\tilde{\boldsymbol{f}}(\boldsymbol{\mu}(\lambda)))$, we further have

$$\sum_{e \in p} c_e(\tilde{x}_e(\boldsymbol{\mu}(\lambda))) + c_e'(\tilde{x}_e(\boldsymbol{\mu}(\lambda)))\tilde{x}_e(\boldsymbol{\mu}(\lambda)) = \sum_{e \in p'} c_e(\tilde{x}_e(\boldsymbol{\mu}(\lambda))) + c_e'(\tilde{x}_e(\boldsymbol{\mu}(\lambda)))\tilde{x}_e(\boldsymbol{\mu}(\lambda)), \tag{10}$$

which is equivalent to

$$\sum_{e \in p} \Big((\beta+1) a_e \tilde{x}_e(\boldsymbol{\mu}(\lambda))^{\beta} + t_e \Big) = \sum_{e \in p'} \Big((\beta+1) a_e \tilde{x}_e(\boldsymbol{\mu}(\lambda))^{\beta} + t_e \Big) \tag{11}$$

for all $p, p' \in S$. Subtracting (11) from $(\beta+1)$ times (9) and dividing by $\beta + 1$ we obtain

$$\sum_{e \in p} \Big(\tfrac{\beta}{\beta+1} t_e + \tau_e \Big) = \sum_{e \in p'} \Big(\tfrac{\beta}{\beta+1} t_e + \tau_e \Big) \tag{12}$$

for all $p, p' \in S$.

By the continuity of the path flow functions $\tilde{f}_p(\boldsymbol{\mu}(\cdot))$, for each path $p \in \mathcal{P}$ the set $L^i(p) = \bigcup_{S \in 2^{\mathcal{P}^i} : p \in S} L^i(S)$ is open in $[0,1]$, i.e., it is an open set in the relativ topology of $[0,1]$. In addition, we have that neither $\mu_i(0)$ nor $\mu_i(1)$ are zero, so that for all $\lambda \in [0,1]$ there is a path $p \in \mathcal{P}^i$ such that $\tilde{f}_p(\boldsymbol{\mu}(\lambda)) > 0$. We conclude that $\bigcup_{p \in \mathcal{P}} L^i(p) = [0,1]$. Since the sets $L^i(p)$, $p \in \mathcal{P}^i$ are open and cover the compact $[0,1]$, there is a finite set $\{p_1, \dots, p_t\} \subseteq \mathcal{P}^i$ such that $0 \in L^i(p_1)$, $1 \in L^i(p_t)$, $L^i(p_1) \cup \cdots \cup L^i(p_t) = [0,1]$ and $L^i(p_j) \cap L^i(p_{j+1}) \neq \emptyset$ for all $j \in \{1, \dots, t-1\}$. The latter condition implies that for $\lambda \in L^i(p_j) \cap L^i(p_{j+1})$ we have $p_j, p_{j+1} \in \mathcal{S}(\tilde{\boldsymbol{f}}(\boldsymbol{\mu}(\lambda))$. Equation (12) then implies $\sum_{e \in p_j} (\frac{\beta}{\beta+1} t_e + \tau_e) = \sum_{e \in p_{j+1}} (\frac{\beta}{\beta+1} t_e + \tau_e)$. Iterating this argument shows $\sum_{e \in p_1} (\frac{\beta}{\beta+1} t_e + \tau_e) = \sum_{e \in p_t} (\frac{\beta}{\beta+1} t_e + \tau_e)$. Finally using that $p_1, p^* \in \mathcal{S}^i(\tilde{\boldsymbol{f}}(\boldsymbol{\mu}(0))$ and $p_t, p^{**} \in \mathcal{S}^i(\tilde{\boldsymbol{f}}(\boldsymbol{\mu}(1))$ gives the claimed result. ◀

We proceed to show that the set of BPR-type cost functions is the largest set of cost functions that guarantee the existence of a DIOT, even for single-commodity networks consisting of two parallel edges. The proof can be found in the arXiv version.

▶ **Theorem 5.** *Let c be twice continuously differentiable, strictly semi-convex and strictly increasing, but not of BPR-type. Then there is a congestion game $\Gamma = (\mathcal{G}, \mathcal{I}, \boldsymbol{c})$ with two parallel edges and cost functions $c(x)$ and $c(x) + t$ for some $t \in \mathbb{R}_+$ that does not have a DIOT.*

A similar construction as in the proof of Theorem 5 shows also that a network with two parallel edges with cost functions $c_1 \in \mathcal{C}_{\mathrm{BPR}}(\beta_1)$ and $c_2 \in \mathcal{C}_{\mathrm{BPR}}(\beta_2)$ does not admit a DIOT if $\beta_1 \neq \beta_2$.

4 Nonnegative tolls

The trivial DIOT toll $\hat{\boldsymbol{\tau}}$ is the trivial solution for both the sufficient condition for a DIOT imposed by Theorem 2 and the necessary condition for a DIOT shown in Theorem 4. However, the trivial DIOT is always negative so that the system designer needs to subsidize the traffic in order to enforce the optimum flow. One may wonder whether the conditions imposed by Theorems 2 and 4 admit also a non-negative solution. Our next result shows that, for games played on a *directed acyclic multi-graph* (DAMG), a non-negative DIOT can always be found. The proof can be found in the arXiv version.

▶ **Theorem 6.** *Consider a game $\Gamma = (\mathcal{G}, \mathcal{I}, \boldsymbol{c})$ with BPR-type cost functions where $\mathcal{G}$ is a DAMG. Then there exists a non-negative DIOT for Γ.*

Proof. Given a DAMG there exists a topological sort, namely a linear ordering $\prec$ of its vertices such that, if $v \prec v'$, then there is no path from v' to v in the DAMG. Notice that, in general the topological sort of a DAMG is not unique. Let $|\mathcal{V}| = n$ and call $\boldsymbol{v}^{\prec} = (v_{(1)}, \dots, v_{(n)})$ the vector of ordered vertices. For each edge $e \in (v_{(i)} \to v_{(j)})$, define

$$\delta_e := j - i. \tag{13}$$

Let $\hat{\boldsymbol{\tau}}$ be the trivial DIOT of the game Γ and let

$$\xi := \min_{e \in \mathcal{E}} \frac{\hat{\tau}_e}{\delta_e} \quad \text{and} \quad \chi := \xi_-, \tag{14}$$

where $\xi_- = \max\{-\xi, 0\}$ is the negative part of ξ. Define now

$$\tau_e = \hat{\tau}_e + \delta_e \chi. \tag{15}$$

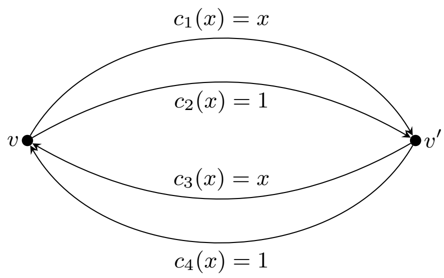

Figure 3 A cyclic network with positive tolls.

We first prove that the toll vector $\boldsymbol{\tau}$ is non-negative. Notice that χ, defined as in (14) is non-negative and $\chi = 0$ only if $\hat{\tau}_e \geq 0$ for all $e \in \mathcal{E}$. Assume that there exists a $\hat{\tau}_e < 0$ and let $e^* \in \arg\min_{e \in \mathcal{E}} \hat{\tau}_e / \delta_e$. Then

$$\tau_{e^*} = \hat{\tau}_{e^*} + \delta_{e^*}\chi = \hat{\tau}_{e^*} - \delta_{e^*}\frac{\hat{\tau_{e^*}}}{\delta_{e^*}} = 0.$$

In general, whenever $\tau_e < 0$, we have

$$\tau_e = \hat{\tau}_e + \delta_e\chi = \hat{\tau}_e - \delta_e\frac{\hat{\tau_{e^*}}}{\delta_{e^*}} \geq \hat{\tau}_e - \delta_e\frac{\hat{\tau}_e}{\delta_e} = 0.$$

Now we prove that the toll vector $\boldsymbol{\tau}$ is a DIOT. By Theorem 2, this means that it satisfies equation (6). First, notice that, by construction of the δ_e, for any $i \in \mathcal{I}$, we have

$$\sum_{e \in p} \delta_e = \sum_{e \in p'} \delta_e \quad \text{for all } p, p' \in \mathcal{P}^i. \tag{16}$$

By (6) we have

$$\sum_{e \in p}\Big((\beta+1)\hat{\tau}_e + \beta t_e\Big) = \sum_{e \in p'}\Big((\beta+1)\hat{\tau}_e + \beta t_e\Big),$$

hence, by (16),

$$\sum_{e \in p}\Big((\beta+1)(\hat{\tau}_e + \delta_e\chi) + \beta t_e\Big) = \sum_{e \in p'}\Big((\beta+1)(\hat{\tau}_e + \delta_e\chi) + \beta t_e\Big),$$

that is

$$\sum_{e \in p}\left((\beta+1)\tau_e + \beta t_e\right) = \sum_{e \in p'}\left((\beta+1)\tau_e + \beta t_e\right),$$

which finishes the proof. ◀

The condition that the graph $\mathcal{G}$ is a DAMG is sufficient for the existence of a non-negative DIOT. It is not necessary, as the following counterexample shows.

▶ **Example 7.** Let $\Gamma = (\mathcal{G}, \mathcal{I}, \boldsymbol{c})$ with $\mathcal{I} = \{1, 2\}$, $\mathcal{V} = \{v, v'\}$, $o_1 = d_2 = v$, $o_2 = d_1 = v'$, $e_1, e_2 \in (o_1 \to d_1)$, $e_3, e_4 \in (o_2 \to d_2)$, and the costs are as in Figure 3. The graph $\mathcal{G}$ is not a DAMG, but the following non-negative toll is a DIOT:

$$\tau_1 = \frac{1}{2} \quad \tau_2 = 0, \quad \tau_3 = \frac{1}{2}, \quad \tau_4 = 0.$$

We proceed to show that for graphs that contain a directed cycle, non-negative DIOTs need not exist, even in networks with affine costs.

▶ **Proposition 8.** *There are networks with affine costs that do not admit a non-negative DIOT.*

5 Aggregatively non-negative Tolls

When nonnegative DIOTs do not exist, it is conceivable that a social planner may sometime want to use negative tolls in order to achieve her goal. Nevertheless, the planner may be subject to budget constraints and not be able to afford a toll system that implies a global loss. Therefore it is interesting to study the existence of conditions for a DIOT $\boldsymbol{\tau}$ such that the following budget constraint is satisfied:

$$\sum_{e\in\mathcal{E}} \tau_e x_e \geq 0, \quad \text{for any feasible flow } \boldsymbol{f}. \tag{17}$$

Intuitively, (17) requires that the social planner does not loose money for any feasible flow. In this section, we show that when the origin-destination pairs (o^i, d^i) satisfy a order condition, then a DIOT satisfying the budget constraint exists.

▶ **Theorem 9.** *Consider a game $\Gamma = (\mathcal{G}, \mathcal{I}, \boldsymbol{c})$ with BPR-type cost functions. If there exists an order $\prec$ on $\mathcal{V}$ such that for all $i \in \mathcal{I}$ we have $o^i \prec d^i$, then there exists a DIOT $\boldsymbol{\tau}$ that satisfies* (17)*.*

We obtain the existence the budget feasible DIOTs for single commodity networks as a direct corollary of Theorem 9.

▶ **Corollary 10.** *Consider a game $\Gamma = (\mathcal{G}, \mathcal{I}, \boldsymbol{c})$ with BPR-type cost and a single O/D pair. Then, there exists a DIOT $\boldsymbol{\tau}$ that satisfies* (17)*.*

Example 7 shows that the condition of Theorem 9 is only sufficient for the existence of a DIOT that satisfies (17).

References

1 Lihud Bai, Donald W. Hearn, and Siriphong Lawphongpanich. Decomposition techniques for the minimum toll revenue problem. *Networks*, 44:142–150, 2004. doi:10.1002/net.20024.

2 Lihud Bai, Donald W. Hearn, and Siriphong Lawphongpanich. A heuristic method for the minimum toll booth problem. *Journal of Global Optimization*, 48:533–548, 2010. doi:10.1007/s10898-010-9527-7.

3 Lihud Bai and Paul A. Rubin. Combinatorial Benders cuts for the minimum tollbooth problem. *Operations Research*, 57:1510–1522, 2009. doi:10.1287/opre.1090.0694.

4 Martin J. Beckmann, C. B. McGuire, and Christopher B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, New Haven, CT, 1956.

5 Pia Bergendorff, Donald W. Hearn, and Motakuri V. Ramana. Congestion toll pricing of traffic networks. In P. Pardalos, D. Hearn, and W. Hager, editors, *Network Optimization*, volume 450 of *LNE*, pages 51–71. Springer, 1997. doi:10.1007/978-3-642-59179-2_4.

6 Umang Bhaskar, Katrina Ligett, Leonard J. Schulman, and Chaitanya Swamy. Achieving target equilibria in network routing games without knowing the latency functions. In *Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 31–40. IEEE Computer Society, 2014. doi:10.1109/FOCS.2014.12.

7 Vittorio Bilò and Cosimo Vinci. Dynamic taxes for polynomial congestion games. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pages 839–856, 2016. doi:10.1145/2940716.2940750.

8 Vincenzo Bonifaci, Mahyar Salek, and Guido Schäfer. Efficiency of restricted tolls in non-atomic network routing games. In G. Persiano, editor, *Proceedings of the 4th Symposium on Algorithmic Game Theory (SAGT)*, volume 6982 of *LNCS*, pages 302–313, 2011. doi:10.1007/978-3-642-24829-0_27.

9 Philip N. Brown and Jason R. Marden. Avoiding perverse incentives in affine congestion games. In *Proceedings of the 55th IEEE Conference on Decision and Control (CDC)*, pages 7010–7015, 2016. doi:10.1109/CDC.2016.7799349.

10 Philip N. Brown and Jason R. Marden. The robustness of marginal-cost taxes in affine congestion games. *IEEE Transactions on Automatic Control*, 62:3999–4004, 2017. doi:10.1109/TAC.2016.2619674.

11 Ioannis Caragiannis, Christos Kaklamanis, and Panagiotis Kanellopoulos. Taxes for linear atomic congestion games. In Y. Azar and T. Erlebach, editors, *Proceedings of the 14th Annual European Symposium on Algorithms (ESA)*, pages 184–195. Springer, 2006. doi:10.1007/11841036_19.

12 Giorgos Christodoulou, Kurt Mehlhorn, and Evangelia Pyrga. Improving the price of anarchy for selfish routing via coordination mechanisms. *Algorithmica*, 69:619–640, 2014. doi:10.1007/s00453-013-9753-8.

13 Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. Pricing network edges for heterogeneous selfish users. In *Proceedings of the 35th Annual Symposium on Theory of Computing (STOC)*, pages 521–530, 2003. doi:10.1145/780542.780618.

14 Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. How much can taxes help selfish routing? *Journal of Computer and System Sciences*, 72:444–467, 2006. doi:10.1016/j.jcss.2005.09.010.

15 Riccardo Colini-Baldeschi, Max Klimm, and Marco Scarsini. Demand-independent optimal tolls. *arXiv*, 2017. arXiv:1708.02737.

16 Robert B. Dial. Minimal-revenue congestion pricing part I: A fast algorithm for the single-origin case. *Transportation Research, Part B*, 33:189–202, 1999. doi:10.1016/S0191-2615(98)00026-5.

17 Robert B. Dial. Minimal-revenue congestion pricing part II: An efficient algorithm for the general case. *Transportation Research, Part B*, 34:645–665, 2000. doi:10.1016/S0191-2615(99)00046-6.

18 Lisa Fleischer. Linear tolls suffice: New bounds and algorithms for tolls in single source networks. *Theoretical Computer Science*, 348:217–225, 2005. doi:10.1016/j.tcs.2005.09.014.

19 Lisa Fleischer, Kamal Jain, and Mohammad Mahdian. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 277–285, 2004. doi:10.1109/FOCS.2004.69.

20 Dimitris Fotakis, George Karakostas, and Stavros G. Kolliopoulos. On the existence of optimal taxes for network congestion games with heterogeneous users. In S. Kontogiannis, E. Koutsoupias, and P. G. Spirakis, editors, *Proceedings of the 3rd Symposium on Algorithmic Game Theory (SAGT)*, volume 6386 of *LNCS*, pages 162–173, 2010. doi:10.1007/978-3-642-16170-4_15.

21 Dimitris Fotakis and Paul G. Spirakis. Cost-balancing tolls for atomic network congestion games. *Internet Mathematics*, 5(4):343–364, 2008. doi:10.1080/15427951.2008.10129175.

22 Jose A. Gómez-Ibáñez and Kenneth A. Small. *Road Pricing for Congestion Management: A Survey of International Practice*, volume 210 of *NCHRP Synthesis*. National Academy Press, Washington, D.C., 1994.

23 Michael A. Hall. Properties of the equilibrium state in transportation networks. *Transportation Science*, 12:208–216, 1978. doi:10.1287/trsc.12.3.208.

24 Deren Han, Hong K. Lo, Jie Sun, and Hai Yang. The toll effect on price of anarchy when costs are nonlinear and asymmetric. *European Journal of Operational Research*, 186:300–316, 2008. doi:10.1016/j.ejor.2007.01.027.

25 Tobias Harks, Ingo Kleinert, Max Klimm, and Rolf H. Möhring. Computing network tolls with support constraints. *Networks*, 65:62–285, 2015. doi:10.1002/net.21604.

26 Donald W. Hearn and Motakuri V. Ramana. Solving congestion toll pricing models. In P. Marcotte and S. Nguyen, editors, *Equilibrium and advanced transportation modeling*, pages 109–124. Springer, New York, 1998. doi:10.1007/978-1-4615-5757-9_6.

27 Martin Hoefer, Lars Olbrich, and Alexander Skopalik. Taxing subnetworks. In C. Papadimitriou and S. Zhang, editors, *Proceedings of the 4th Workshop on Internet and Network Economics (WINE)*, volume 5385 of *LNCS*, pages 286–294, 2008.

28 George Karakostas and Stavros G. Kolliopoulos. Edge pricing of multicommodity networks for heterogeneous selfish users. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 268–276, 2004. doi:10.1109/FOCS.2004.26.

29 George Karakostas and Stavros G. Kolliopoulos. The efficiency of optimal taxes. In A. López-Ortiz and A. Hamel, editors, *Proceedings of the 1st Workshop on Combinatorial and Algorithmic Aspects of Networking (CAAN)*, volume 3405 of *LNCS*, pages 3–12. Springer, 2005. doi:10.1007/11527954_2.

30 George Karakostas and Stavros G. Kolliopoulos. Edge pricing of multicommodity networks for selfish users with elastic demands. *Algorithmica*, 53:225–249, 2009. doi:10.1007/s00453-008-9181-3.

31 Frank H. Knight. Some fallacies in the interpretation of social cost. *Quarterly Journal of Economics*, 38(4):582–606, 1924. doi:10.2307/1884592.

32 Torbjörn Larsson and Michael Patriksson. Side constrained traffic equilibrium models—analysis, computation and applications. *Transportation Research, Part B*, 33(4):233–264, 1999. doi:10.1016/S0191-2615(98)00024-1.

33 Reshef Meir and David Parkes. When are marginal congestion tolls optimal? In *Proceedings of the 9th Workshop on Agents in Traffic and Transportation (ATT)*, volume 1678 of *CEUR Workshop Proceedings*, 2016.

34 Arthur C. Pigou. *The Economics of Welfare.* Macmillan and Co., London, 1st edition, 1920.

35 Aaron Roth, Jonathan Ullman, and Zhiwei Steven Wu. Watch and learn: optimizing from revealed preferences feedback. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, pages 949–962. ACM, 2016. doi:10.1145/2897518.2897579.

36 William H. Sandholm. Evolutionary implementation and congestion pricing. *Review of Economic Studies*, 69(3):667–689, 2002. doi:10.1111/1467-937X.t01-1-00026.

37 William H. Sandholm. Pigouvian pricing and stochastic evolutionary implementation. *Journal of Economic Theory*, 132(1):367–382, 2007. doi:10.1016/j.jet.2005.09.005.

38 Chandramani Singh. Marginal cost pricing for atomic network congestion games. Technical report, Department of Electrical Communication Engineering, Indian Institute of Science, 2008.

39 Kenneth A. Small and Jose A. Gómez-Ibáñez. Road pricing for congestion management: the transition from theory to policy. In Tae Hoon Oum, editor, *Transport Economics*, chapter 16, pages 373–403. Routledge, London, 1997.

40 U.S. Bureau of Public Roads. Traffic assignment manual. U.S. Department of Commerce, Urban Planning Division, 1964.

41 John Glen Wardrop. Some theoretical aspects of road traffic research. In *Proceedings of the Institute of Civil Engineers, Part II*, volume 1, pages 325–378, 1952.

42 Hai Yang and Hai-Jun Huang. The multi-class, multi-criteria traffic network equilibrium and systems optimum problem. *Transportation Research, Part B*, 38:1–15, 2004. doi:10.1016/S0191-2615(02)00074-7.