Taxing Subnetworks ⋆

Martin Hoefer, Lars Olbrich, and Alexander Skopalik

Department of Computer Science, RWTH Aachen University, Germany

Abstract. We study taxes in the well-known game theoretic traffic model due to Wardrop. Given a network and a subset of edges, on which we can impose taxes, the problem is to find taxes inducing an equilibrium flow of minimal network-wide latency cost. If all edges are taxable, then marginal cost pricing is known to induce the socially optimal flow for arbitrary multi-commodity networks. In contrast, if only a strict subset of edges is taxable, we show NP-hardness of finding optimal taxes for general networks with linear latency functions and two commodities. On the positive side, for single-commodity networks with parallel links and linear latency function, we provide a polynomial time algorithm for finding optimal taxes.

1 Introduction

An important problem in traffic management is to set incentives for rational users to act in a favorable manner. An effective means to achieve this is to set appropriate taxes. In this paper, we study the problem of computing optimal taxes in the Wardrop model, a well-studied model for traffic routing with important applications in road networks and computer networks. In this model, we are given a network equipped with non-decreasing non-negative latency functions mapping flow on the edges to latency. For each of several commodities a fixed demand has to be routed between a source-sink pair. The cost of a flow assignment is the weighted sum of travel times between the source and target nodes. A flow that minimizes the total latency is called *(socially) optimal*. A common interpretation of the Wardrop model is that flow is controlled by an infinite number of selfish users each of which carries an infinitesimal amount of flow. Each user aims at minimizing its path latency. An allocation, in which no user can improve its situation by unilaterally deviating from its current path is called *Wardrop equilibrium*. In general a Wardrop equilibrium is not socially optimal, i.e, it does not minimize the total latency. The inefficiency of selfish flows has been extensively studied in previous work [2, 19, 20, 22].

Taxing can be successful in improving total latency of equilibria. In this case users are assumed to minimize the sum of their latencies and taxes. A fundamental result is that using *marginal cost pricing* to tax every edge results in equilibrium flows that are optimal with respect to total latency [1]. A serious drawback of marginal cost pricing is that it requires *every* edge of the network to be taxable. In many situations there might be technical or legal restrictions that prevent an operator from imposing a tax on all

⋆ Supported by the DFG GK/1298 "AlgoSyn" and by the German Israeli Foundation (GIF) under contract 877/05.

edges. Therefore, we adjust the model to a more realistic case in which only a subset of edges can be taxed. The problem is to find a set of taxes for the subset of taxable edges that minimizes the total latency of the resulting Wardrop equilibrium. To the best of our knowledge, this generalization has not been considered before.

Taxing subnetworks can be more difficult and non-trivial. Consider a parallel link network of two links and linear latency functions. If one can tax only one edge, the latency cost is generally not monotone in the imposed tax. Using this insight, we carefully construct networks with one taxable edge and several distinct optimal taxes. A combination of these networks establishes NP-hardness of the problem for two commodities and linear latency functions in Sect. 3. On the other hand, for parallel link networks with linear latency functions, we derive a precise structural analysis of optimally taxed equilibrium flows in Sect. 4. This allows to construct a polynomial-time algorithm to find optimal taxes. Most proof details are omitted and will be given in the full version of the paper.

Related Work There is a huge amount of work addressing the inefficiency of equilibria in the Wardrop model. Therefore, we only give a rough overview and concentrate on the classical results and recent developments. The game theoretic traffic model considered in this paper was introduced by Wardrop [25]. Beckmann et al. [1] observe that such an equilibrium flow is an optimal solution to a related convex program. They give existence and uniqueness results for traffic equilibria (see also [7] and [20]). Dafermos and Sparrow [7] show that the equilibrium state can be computed efficiently under some assumptions on the latency functions.

The inefficiency of Wardrop equilibria is a well-known phenomenon [17], which is exemplified by Braess paradox [2]. Bounding the inefficiency of equilibria, however, has only recently been considered, initiated by Koutsoupias and Papadimitriou [15], and for the Wardrop model by Roughgarden and Tardos [20]. Roughgarden [22] provides a cumulative overview of the most important results that have been obtained.

There are several approaches that have been proposed to address the inefficiency of equilibria. The effectiveness of taxes has been observed by Pigou [17] and generalized by Beckmann et al. [1]. They show that *marginal cost pricing* completely eliminates the inefficiency of selfish routing. Cole et al. [6] show existence of taxes inducing the optimal flow for single-commodity networks and heterogeneous users that value tax versus latency in an individual way. Fleischer [8] reduces the required taxes to linear functions. In the more general setting of multi-commodities, Fleischer et al. [9] and Karakostas and Kolliopoulos [11] independently prove the existence of optimal taxes.

Other approaches for coping with selfishness are, for example, proposed by Korilis et al. [14], who give methods for improving system performance by adding additional capacity to system resources. Cocchi et al. [4] study the role of various pricing policies in networks with selfish users. Roughgarden [21] studies designing networks that exhibit good performance when used selfishly and proves tight inapproximability results. Cole et al. [5] show hardness of computing taxes minimizing the total user disutility (latency plus tax) at equilibrium.

Korilis et al. [13] consider the problem of a Stackelberg leader, who in a first phase can fix the routes for a certain fraction of the demand. In a second phase, selfish users enter the system and route their own flow on top of the leader demand. The objective of the leader is to minimize the resulting total cost of the total (both leader and selfish) flow. Roughgarden [18] shows that it is weakly NP-hard to compute the optimal leader strategy even for parallel links with linear latency functions. Kumar and Marathe [16] give a FPAS for this problem. Kaporis and Spirakis [10] show that for single-commodity networks the minimal fraction of flow needed by the leader to induce optimal cost can be computed in polynomial time. Subsequent papers [24, 23, 12] consider Stackelberg routing in different variants for more general networks.

2 Preliminaries

We consider Wardrop's traffic model originally introduced in [25]. We are given a directed graph $G = (V, E)$ with vertex set V, edge set E, a set of commodities $[k] = \{1, \ldots, k\}$ specified by source-sink pairs $(s_i, t_i) \in V \times V$, and flow demands $d_i > 0$. For single-commodity networks we normalize the demand to one. Considering only parallel edges, we speak of parallel link networks and denote the set of links by $[n] = \{1, \ldots, n\}$. The edges are equipped with non-decreasing, continuous latency functions $\ell_e : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. We allow a set of non-negative taxes $\{\tau_e\}_{e \in T}$ to be imposed on a subset of edges $T \subset E$. We call edges in T *taxable* and edges in $N = E \setminus T$ *non-taxable*.

Let $\mathcal{P}_i$ denote the admissible paths of commodity i, i. e., all paths connecting s_i and t_i, and let $\mathcal{P} = \bigcup_{i \in [k]} \mathcal{P}_i$. A non-negative path flow vector $(f_P)_{P \in \mathcal{P}}$ is *feasible* if it satisfies the flow demands $\sum_{P \in \mathcal{P}_i} f_P = d_i$ for all $i \in [k]$. Throughout this paper, we will consider only feasible path flow vectors. A path flow vector $(f_P)_{P \in \mathcal{P}}$ induces an edge flow vector $f = (f_e)_{e \in E}$ with $f_e = \sum_{i \in [k]} \sum_{P \in \mathcal{P}_i : e \in P} f_P$. For single-commodity networks, we drop the index i. The latency of an edge $e \in E$ is given by $\ell_e(f_e)$ and the latency of a path P is given by the sum of the edge latencies $\ell_P(f) = \sum_{e \in P} \ell_e(f_e)$. The latency cost of a flow is defined as $C(f) = \sum_{P \in \mathcal{P}} \ell_P(f) f_P = \sum_{e \in E} \ell_e(f_e) f_e$. A flow f of minimal latency cost is called *(socially) optimal*. The cost of a path is defined as latency plus tax, i.e., $\ell_P(f) + \sum_{e \in P} \tau_e$. Finally, we call the quadruple $(V, T, N, (d_i))$ an *instance*.

A flow vector is considered stable when no fraction of the flow can improve its sustained cost by moving unilaterally to another path. Such a stable state is generally known as *Nash equilibrium*. In our model a flow is stable if and only if all used paths within a commodity have the same minimal cost, whereas unused paths may have larger cost. We call such a flow *Wardrop equilibrium*.

Definition 1 *A feasible flow vector f is at* Wardrop equilibrium *if for every commodity $i \in [k]$ and paths $P_1, P_2 \in \mathcal{P}_i$ with $f_{P_1} > 0$ it holds that $\ell_{P_1}(f) + \sum_{e \in P_1} \tau_e \leq \ell_{P_2}(f) + \sum_{e \in P_2} \tau_e$.*

In particular, without taxes, if f is at Wardrop equilibrium then all used paths in commodity i have equal latency $L_i(f)$ and the latency cost can be expressed as $C(f) = \sum_{i \in [k]} L_i(f) \cdot d_i$ (see [20, 25]). A classical result on taxing selfish flow, called *marginal cost pricing*, is that with taxes $\tau_e = x_e \cdot \ell'_e(x_e)$ for all $e \in E$ the resulting equilibrium flow minimizes the latency cost. With $\ell_e^*(x) = (x \cdot \ell_e(x))' = \ell_e(x) + x \cdot \ell_e'(x)$ denoting the marginal cost of increasing flow in edge e we have the following lemma.

Lemma 1 ([1, 7, 20]) *Let $(V, T, \emptyset, (d_i))$ denote an instance in which $x \cdot \ell_e(x)$ is a convex function for each edge e. Then a flow f minimizes the latency cost w.r.t. $(\ell_e)_{e\in T}$ if and only if it is at Wardrop equilibrium w.r.t. $(\ell_e^*)_{e\in T}$.*

In the restricted case with only a subset of edges being taxable such a result is obviously out of reach. This directly leads us to the following definition.

Definition 2 *Given an instance $(V, T, N, (d_i))$, a set of taxes $\{\tau_e\}_{e\in T}$ is called* optimal*, if there is an equilibrium flow f_τ w.r.t. $\ell + \tau$ with $C(f_\tau) \leq C(f_{\tau'})$ for all equilibrium flows $f_{\tau'}$ w.r.t. $\ell + \tau'$ for any $\{\tau_e'\}_{e\in T}$.*

3 NP-Hardness for Multi-Commodity Networks

In this section we study the optimization problem of computing an optimal set of taxes. We show that this turns out to be NP-hard even for the two-commodity case with linear latency functions. We start with an observation which allows us to discretise the problem and enables us to prove the main result of this section.

Lemma 2 *There is a family of instances $(V, T, N_A, d_A)_{A\in\mathbb{N}}$ with parallel link networks allowing for two separated optimal tax values.*

Proof. Consider a parallel link network, in which two nodes s and t are connected via three links with $\ell_1(x) = x + A$ and $\ell_2(x) = \ell_3(x) = x$. Suppose we can only tax the third link. Set $d_A = A(1 + \frac{\sqrt{3}}{2})$. For tax $0 \leq \tau \leq A(1 - \frac{\sqrt{3}}{2})$, the total demand is split among links two and three at equilibrium. Since both used links are identical, $\tau = 0$ is optimal with an induced cost of $(\frac{7}{8} + \frac{\sqrt{3}}{2})A^2$. For $A(1 - \frac{\sqrt{3}}{2}) < \tau < A(1 + \frac{\sqrt{3}}{4})$ all links are used and the corresponding cost function $\frac{2}{3}\tau^2 - \frac{1}{3}A\tau + (\frac{11}{12} + \frac{\sqrt{3}}{2})A^2$ yields an optimal tax of $A/4$ with cost $(\frac{7}{8} + \frac{\sqrt{3}}{2})A^2$ as well. For $\tau \geq A(1 + \frac{\sqrt{3}}{4})$ the latency cost at equilibrium is $(\frac{11}{8} + \frac{3\sqrt{3}}{4})A^2$. Thus, both $\tau = 0$ and $\tau = A/4$ are optimal. □

Theorem 3 *Given an instance $(V, T, N, (d_i))$, the problem of computing optimal taxes is NP-hard, even for only two commodities and linear latency functions.*

Proof. We reduce from the PARTITION problem: given n positive integers $a_1, \ldots, a_n$, is there a subset $S \subseteq \{1, 2, \ldots, n\}$ satisfying $\sum_{i\in S} a_i = \frac{1}{2}\sum_{i=1}^n a_i$? We will show that deciding the PARTITION problem reduces to deciding if a given 2-commodity instance $(V, T, N, (d_i))$ with latency functions admits taxes inducing a Wardrop equilibrium with a given cost. Given an arbitrary instance of PARTITION specified by positive integers $a_1, \ldots, a_n$, we define an instance $(V_{\{a_i\}}, T_{\{a_i\}}, N_{\{a_i\}}, (d_{\{a_i\}}))$ as depicted in Fig. 1. Let the set of taxable edges T consist of the bold edges. Commodity one has a demand of $A = \prod_{i=1}^n a_i$ to route between $s_1 = v_1$ and $t_1 = v_{n+1}$, the second commodity has to route a demand of $\sum_i a_i$ between s_2 and t_2. For $i \in [n]$ define the following

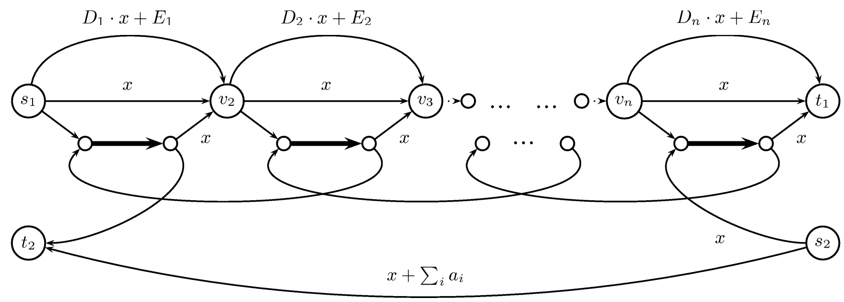

Fig. 1. The network of an instance $(V_{\{a_i\}}, T_{\{a_i\}}, N_{\{a_i\}}, (d_{\{a_i\}}))$. The edges are labeled with the latency functions. Unlabeled edges have latency 0. Taxes can be imposed on the set of bold edges only.

constants: $A_{-i} = \prod_{j \neq i}^{n} a_j$, $D_i = \frac{2-4A_{-i}+A_{-i}^2}{4A_{-i}-2}$ and $E_i = 2a_i(D_i + 1) = \frac{AA_{-i}}{2A_{-i}-1}$. We show that $\{a_1, \ldots, a_n\}$ is a YES instance if and only if there are taxes for instance $(V_{\{a_i\}}, T_{\{a_i\}}, N_{\{a_i\}}, (d_{\{a_i\}}))$ inducing a Wardrop equilibrium with cost of at most $C = \frac{n}{2}A^2 + \frac{7}{8}(\sum_i a_i)^2$. The idea is that the minimal latency cost is reached if and only if the tax between v_i and v_{i+1} is 0 or a_i (inducing a latency cost of $A^2/2$ for this set of edges) and the sum of all taxes is exactly $\sum_i a_i/2$. □

4 Parallel Links with Linear Latency Functions

We have seen that the latency cost is generally not monotone in the imposed tax even in case of linear latency functions and one taxable link. Further, such instances do not necessarily admit a unique optimal tax. These observations indicate that studying optimal taxes in parallel link networks might be intriguing. Our main goal in this section is to provide an algorithm for finding optimal taxes in single-commodity parallel link networks $(V, T, N, 1)$ in which every link $i \in [n]$ has a linear latency function $\ell_i(x) = a_i x + b_i$. This setting has been of special interest in the related problem of computing a Stackelberg leader strategy [18] described in the introduction. While this problem is already NP-hard in this setting, it may be surprising that we will be able to formulate a polynomial time algorithm for computing optimal taxes. Suppose the links are numbered by $N = \{1, \ldots, k\}$ and $T = \{k+1, \ldots, n\}$, such that $b_1 \leq \ldots \leq b_k$ and $b_{k+1} \leq \ldots \leq b_n$. We use this labelling for convenience, but note that the ordering conditions apply only *within* N and T. We do not require $b_i \leq b_j$ for all $i \in N$ and $j \in T$ or any other restriction or relation between the links of N and T. W.l.o.g. we assume at most one constant latency link in $N \cup T$.

4.1 Candidate Supports Sets

A flow f is at Wardrop equilibrium if and only if there is a constant $L > 0$, s.t. all used links $i \in [n]$ have the same latency $L = \ell_i(f_i)$, whereas $L \leq \ell_{i'}(0) = b_{i'}$ for unused links $i' \in [n]$. Lemma 1 shows that a flow f is socially optimal if and only if there is a constant $C > 0$, s.t. $C = \ell_j^*(f_j) = 2a_j f_j + b_j$ for all used links $j \in [n]$, whereas $C \leq \ell_{j'}^*(0) = b_{j'}$ for unused links $j' \in [n]$.

Now consider an instance and increase the demand. The characterization yields that in equilibrium and in optimum links j will be filled with flow in order of increasing b_j. Regarding cost the set of taxes will induce an equilibrium assigning flow to some link set $S \subset N \cup T$. All used non-taxable links have the same latency L. Since we allow for non-negative taxes only, the used taxable links will not have higher latency. This property allows us to parametrize the problem by the set of taxable and non-taxable links filled with flow. These sets turn out to be *candidate support sets* defined as follows.

Definition 3 *Every set of the form* $S = \{1, \ldots, l_1\} \cup \{k+1, \ldots, l_2\}$ *with* $1 \leq l_1 \leq k$ *and* $k+1 \leq l_2 \leq n$ *is called a* candidate support set.

Note that there are at most $n^2/4$ candidate support sets for any instance.

Lemma 4 *Let f denote a socially optimal flow for a parallel link network in which every edge is taxable. Then* $\ell_1(f_1) \leq \ell_2(f_2) \leq \ldots \leq \ell_n(f_n)$.

Proof. The set of used links is of the form $\{1, \ldots, l\}$ for some $l \leq n$. Since f is a minimal latency flow, all links $j \in \{1, \ldots, l\}$ have equal marginal cost, and there is a constant $C > 0$ with $2a_j f_j + b_j = C$. Thus, $\ell_j(f_j) = a_j f_j + b_j = C/2 + b_j/2$. □

Let us first argue that the consideration of candidate support sets is indeed sufficient to find optimal taxes. Imagine two separate commodities, routing demands d_N and $1 - d_N$ exclusively over N and T, resp. In such an instance, it would be optimal to set marginal cost taxes on T, and the set of used links form a candidate support set.

The difference to our setting is that demand can change between N and T, and thus we also need to ensure that latency and taxes create an equilibrium. If the optimal flow in T yields latencies smaller than L, then we can satisfy the latency constraint by setting appropriate non-negative taxes. Otherwise, the latency restriction reduces the flow on some used links. However, if the flow on a link is smaller than in the optimum due to the latency constraint, the marginal cost on this link is also smaller. Therefore, it is still optimal to fill the link with flow to the maximal possible extent (see Lemma 5). For all links not affected by the latency restriction, however, it is optimal to equalize the marginal costs, and the allocation of flow follows the ordering of offsets. In conclusion, the set of links allocated with flow remains a candidate support set.

4.2 Problem Parametrization

Fixing numbers n_S and t_S yields a candidate support set $S = N_S \cup T_S$ with $N_S = \{1, \ldots, n_S\}$ and $T_S = \{k+1, \ldots, t_S\}$. For S denote by d_{N_S} and $1 - d_{N_S}$ the demand routed over N_S and T_S, respectively. $C_{N_S}(d_{N_S})$ is the latency cost for an equilibrium flow $(f_i)_{i \in N_S}$ of demand d_{N_S}. Denote by $C_{T_S}(1 - d_{N_S})$ the latency cost for an optimal flow $(f_j)_{j \in T_S}$ of demand $1-d_{N_S}$ additionally fulfilling the latency restriction $\ell_j(f_j) \leq L(d_{N_S})$, where $L(d_{N_S})$ denotes the unique latency of all used links in N_S for a demand of d_{N_S}. Let $C(d_{N_S}) = C_{N_S}(d_{N_S}) + C_{T_S}(1-d_{N_S})$ denote the latency cost of the flow.

The problem of finding a set of optimal taxes for a fixed set S can be formulated as follows: Minimize the cost function C, s.t. the flow for N is at equilibrium and the remaining flow on T is optimal subject to the additional constraint $\ell_j(f_j) \leq L(d_{N_S})$.

We will show that, if this minimization problem has a solution, the cost function $C(d_{N_S})$ is piecewise quadratic with at most n breakpoints and the optimal demand distribution $(d^*_{N_S}, 1-d^*_{N_S})$ for N_S and T_S is efficiently computable. Iterating this for all possible sets S enables us to find optimal taxes.

We call a link $j \in T$ *full* w.r.t. some $L > 0$ if $f_j > 0$ and its latency equals the constraint value, i.e., if $\ell_j(f_j) = L$ or if $f_j = 0$ and $\ell_j(0) = b_j \geq L$. We call a link *relaxed* if $f_j > 0$ and $\ell_j(f_j) < L$. When shifting demand from N to T, the common latency L of used links in N decreases, while the demand on T increases. In the corresponding optimal flow on T respecting the constraint value, however, a full link never becomes relaxed. More formally, consider an instance $(V, T, \emptyset, d)$ and let f denote the optimal flow respecting $\ell_i(f_i) \leq L$ for all i. With Lemma 4 we can assume the full links to form a set $\{p, \ldots, n\}$ for some $p \geq 1$. Furthermore, assume there are $L' \leq L$ and $d' \geq d$ such that there is a flow of demand d' to T such that all used links have latency at most L'. For all non-constant links, we define $\ell_i^{-1}(L)$ to be the flow f_i such that $a_i f_i + b_i = L$ if $b_i \leq L$, and 0 otherwise.

Lemma 5 *The optimal flow f' respecting $\ell_i(f'_i) \leq L'$ for all i assigns $\ell_i^{-1}(L')$ flow to all non-constant links $i \in \{p_1, \ldots, n\}$ for some uniquely defined $p_1 \leq p$.*

4.3 A Polynomial-Time Algorithm for Computing Optimal Taxes

Considering an optimal flow for an increasing demand, the links become used in order of their offsets. Lemmata 4 and 5 show that the links become full w.r.t. some bound in reverse order. Thus, we can determine the lower and the upper bound $d^{\min}_{N_S}$ and $d^{\max}_{N_S}$ for d_{N_S} such that the following holds: There is an equilibrium flow of demand d_{N_S} on N using exactly the links N_S and there is an optimal flow of demand $1 - d_{N_S}$ on T respecting the bound $L(d_{N_S})$ using exactly the links T_S.
Given a candidate support set S, we compute the optimal demand distribution $(d_{N_S}, 1-d_{N_S})$. If such a distribution exists, we call S *feasible*. The corresponding demand interval $[d^{\min}_{N_S}, d^{\max}_{N_S}]$ can be computed in polynomial time by solving systems of linear equations.

Lemma 6 *The cost function $C(d_{N_S})$ is piecewise quadratic for $d_{N_S} \in [d^{\min}_{N_S}, d^{\max}_{N_S}]$ with at most n breakpoints for every feasible candidate support set S. The breakpoints can be computed in polynomial time.*

Proof. We show that while C_{N_S} is a quadratic function, C_{T_S} and therefore C is piecewise quadratic with at most n breakpoints.
Suppose f is an equilibrium flow for N_S of demand d_{N_S}. There is some $L(d_{N_S}) > 0$ with $L(d_{N_S}) = a_i f_i + b_i$ for every $i \in N_S$. With $\sum_{N_S} f_i = d_{N_S}$, we infer that $L(d_{N_S})$

Algorithm 1 OPTTAX $(V, T, N, 1)$

for every candidate support set S **do**
  if S feasible **then**
    compute the breakpoints $d_{N_S}^{\min} = d_{N_{S\,k+1}}, \dots, d_{N_{S\,1}}, d_{N_{S\,0}} = d_{N_S}^{\max}$
    $d^*_{N_S} \leftarrow \text{argmin}_{0 \le j \le k} \min_{d_{N_S} \in [d_{N_{S\,j}}, d_{N_{S\,j+1}}]} C(d_{N_S})$
  end if
end for
$S^* \leftarrow \text{argmin}_S C(d^*_{N_S})$
compute optimal flow on T_{S^*} respecting $L(d^*_{N_{S^*}})$ with $\sum_{T_{S^*}} f^*_j = 1 - d^*_{N_{S^*}}$ and set $f^*_j := 0$ for $j \in T \setminus T_{S^*}$.
set taxes $\tau_j \leftarrow L(d^*_{N_{S^*}}) - \ell_j(f^*_j)$ for $j \in T$

is linear and $C_{N_S}(d_{N_S}) = L(d_{N_S}) \cdot d_{N_S}$ is quadratic. Considering C_{T_S}, we need to respect the latency constraint for increasing $1 - d_{N_S}$. The cost function C_{T_S} turns out to be quadratic with at most n breakpoints. These breakpoints, i.e., the demand values for which the number of full links increases, can be calculated by solving systems of linear equations. □

Given that restricting to candidate support sets is sufficient for finding optimal taxes, the following result holds.

Theorem 7 *Given an instance $(V, T, N, 1)$ with parallel links and linear latency functions, Algorithm* OptTax(V,T,N,1) *computes a set of optimal taxes $(\tau_j)_{j \in T}$ in polynomial time.*

References

[1] M. Beckmann, C. B. McGuire, and C. B. Winston. Studies in the Economics of Transportation. *Yale University Press*, 1956.

[2] D. Braess. Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung*, 12:258-268, 1968.

[4] R. Cocchi, S. Shenker, D. Estrin, L. Zhang. pricing in computer networks: motivation, formulation, and example. *IEEE/ACM Transactions on Networking*, 1(6):614–627, 1993.

[5] R. Cole, Y. Dodis, T. Roughgarden. How much can taxes help selfish routing? In *Proc. 4th EC*, 98–107, 2003.

[6] R. Cole, Y. Dodis, T. Roughgarden. Pricing edges for heterogeneous selfish users. In *Proc. 35th STOC*, 521–530, 2003.

[7] S. Dafermos, F. T. Sparrow. The Traffic Assignment Problem for a General Network. *Journal of Research of the National Bureau of Standards*, Series B, 73(2):91–118, 1969.

[8] L. Fleischer. Linear tolls suffice: New bounds and algorithms for tolls in single source networks. *Theoretical Computer Science*, 348:217–225, 2005.

[9] L. Fleischer, K. Jain, M. Mahdian. Taxes for heterogeneous selfish users in a multicommodity network. In *Proc. 45th FOCS*, 277–285, 2004.

[10] A.C. Kaporis, P.G. Spirakis. The price of optimum in stackelberg games on arbitrary single commodity networks and latency function. In *Proc. 18th SPAA*, 2006.

[11] G. Karakostas, S.G. Kolliopoulos. Edge pricing of multicommodity networks for heterogeneous selfish users. In *Proc. 45th FOCS*, 268–276, 2004.

[12] G. Karakostas, S.G. Kolliopoulos. Stackelberg strategies for selfish routing in general multicommodity networks. *Algorithmica*, to appear, 2008.

[13] Y.A. Korilis, A.A. Lazar, A. Orda. Achieving network optima using Stackelberg routing strategies. *IEEE/ACM Transactions on Networking*, 5(1):161–173, 1997.

[14] Y.A. Korilis, A.A. Lazar, A. Orda. Capacity allocation under noncooperative routing. *IEEE Transactions on Automatic Control*, 42:309–325, 1997.

[15] E. Koutsoupias, C. Papadimitriou. Worst-Case Equilibria. In *Proc. 16th STACS*, 404–413, 1999.

[16] V.S.A. Kumar, A. Marathe. Improved results for Stackelberg scheduling strategies. In *Proc. 29th ICALP*, 2002.

[17] A. C. Pigou. *The Economics of Welfare*. Macmillan, 1920.

[18] T. Roughgarden. Stackelberg scheduling strategies. In *Proc. 33rd STOC*, 104–113, 2001.

[19] T. Roughgarden. The Price of Anarchy is Independent of the Network Topology. In *Proc. 34th STOC*, 428–437, 2002.

[20] T. Roughgarden, É. Tardos. How Bad is Selfish Routing? *Journal of the ACM*, 49(2):236–259, 2002.

[21] T. Roughgarden. On the Severity of Braess's Paradox: Designing Networks for Selfish Users is Hard. *Journal of Computer and System Sciences*, 72(5):922–953, 2004.

[22] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.

[23] Y. Sharma, D.P. Williamson. Stackelberg thresholds in network routing games or the value of altruism. In *Proc. 8th EC*, 93–102, 2007.

[24] C. Swamy. The effectiveness of Stackelberg strategies and tolls for network congestion games. In *Proc. 18th SODA*, 1133–1142, 2007.

[25] J.G. Wardrop. Some Theoretical Aspects of Road Traffic Research. In *Proc. of the Institute of Civil Engineers*, Pt. II:325–378, 1952.