

The efficiency of optimal taxes

George Karakostas¹ *

Stavros G. Kolliopoulos^{2**}

¹ Department of Computing and Software, McMaster University

² Department of Informatics and Telecommunications, University of Athens and
Department of Computing and Software, McMaster University

Abstract. It is well known that the selfish behavior of users in a network can be regulated through the imposition of the so-called *optimal taxes* on the network edges. Any traffic equilibrium reached by the selfish users who are conscious of both the travel latencies and the taxes will minimize the social cost, i.e., will minimize the total latency.

Optimal taxes incur desirable behavior from the society point of view but they cause disutility to the network users since the users' total cost is in general increased [4]. Excessive disutility due to taxation may be undesirable from the societal perspective as well. In this work we examine the efficiency of taxation as a mechanism for achieving the desired goal of minimizing the social cost. We show that for large classes of latency functions the total disutility due to taxation that is caused to the users and/or the system is bounded with respect to the social optimum. In addition, we show that if the social cost takes into account both the total latency and the total taxation in the network, the coordination ratio for certain latency functions is better than the coordination ratio when taxation is not used.

1 Introduction

In the *selfish routing* setting, we are given a directed network $G = (V, E)$ and a set of k *classes of users* (commodities), each with its own origin and destination, and with a fixed total demand (traffic) rate per class $d_i > 0$, $i = 1, \dots, k$. Individual users are thought as carrying each an infinitesimal amount of a commodity. We are also given a nonnegative latency function l_P describing the delay experienced by users wishing to travel on the path P as a function of the total flow through the edges of the path. In this work we assume that the *additive model* holds, i.e., for every edge e there is a latency function $l_e(f_e)$ that describes the latency on this edge due to the flow f_e that crosses it. Then the latency for a path is defined as $l_P(f) := \sum_{e \in P} l_e(f_e)$. Each user tries to selfishly route

* (karakos@mcmaster.ca). Research supported by an NSERC Discovery grant.

** (www.cas.mcmaster.ca/~stavros). Research partially supported by NSERC Research Grant 227809-00.

his flow so that his path cost is minimized. A *traffic equilibrium* is an assignment of traffic to paths so that no user can unilaterally switch her flow to a path of smaller cost. Wardrop's principle [16] for selfish routing postulates that

at equilibrium, for each origin-destination pair, the travel costs on all the routes actually used are equal, or less than the travel costs on all nonused routes.

A widely used measure of the network performance is the *social cost* (or *total latency*), defined as $\sum_{\text{path } P} f_P l_P(f) = \sum_{e \in E} f_e l_e(f_e)$ for a flow f that routes f_P units of traffic through path P . Although it must obey Wardrop's principle at equilibrium, the unregulated choice of paths by individual users may incur a social cost which in general can be higher than the social optimum. In fact the latency of an equilibrium can be arbitrarily larger than the social optimum [14]. A classic way of dealing with this problem is the introduction of *taxation* on the edges of the network, so that the users' path cost has both a travel time and a budgetary component. Without taxation, users experience only their own traffic delay as their cost. With taxation users are also charged for the right to use a path. This technique has been studied by the traffic community for a long time (cf. [5] and the references therein), especially in the context of *marginal costs* (see, for example, [2],[8],[15]). Each selfish user of class i using path P will experience the following path cost:

$$\text{path cost}(P) := \text{latency}(P) + a(i) \cdot \text{taxation}(P).$$

The $\text{taxation}(P)$ is the sum of taxes along the edges of the path. The factor $a(i) > 0$, denotes the sensitivity of user class i to the taxes. In the *homogeneous* case all user classes have the same sensitivity to the taxation (i.e. $a(i) = 1$, for all i), while in the *heterogeneous* case $a(i)$ can take different positive values for different classes. Through edge taxation, we would like to force *all* equilibria on the network to induce flow that minimizes the social cost $\sum_{e \in E} f_e l_e(f_e)$. We refer to a set of edge taxes that achieves this as *optimal taxes*. In the homogeneous case, marginal costs have been shown (cf. [2],[8],[15]) to be optimal taxes. In the heterogeneous case, the existence and calculation of such optimal taxes were shown for the single source-destination pair case by Cole et al. [5], and were later extended to the multicommodity setting in [9, 11].

Designing optimal taxes is a classic instance of mechanism design, a central topic in game theory (see, e.g., [13]). A set of outcomes is fixed (here achieving the social optimum) and users are paid or penalized, (here

they pay taxes) in order to achieve the desired outcome in equilibrium. One can actually see the taxation cost in two different contexts. One is the context already discussed, which is as monetary cost. The other is to see the tax for every edge as part of the edge latency function itself. Then, instead of taxation, we can speak about *artificial delays* introduced possibly at the entrance of each edge, in order to minimize the total amount of time users actually spend on the edges themselves. For example, this is the technique used at some highway exits, where traffic lights have been installed in order to better control traffic. Whatever the context though, taxation increases in general the user cost, as was shown in [4] for the case of marginal cost taxes. The natural question that arises is whether taxes are an efficient mechanism for achieving the desired result. Is the additional disutility caused through taxation proportionate to the desired goal, i.e., a routing that minimizes the total latency? In this paper we tackle this problem by comparing the social cost of the traffic equilibria when taxation is used against (i) the social optimum without taxation and (ii) the social optimum when taxation is taken into account.

In Section 4 we show that in the homogeneous case the ratio of social cost at equilibrium with taxation to the social optimum without taxation is not much bigger than the worst case ratio without any taxation for many important families of latency functions, like linear or low-degree polynomial ones. In particular for strictly increasing linear latency functions we show that, if b is the vector of optimal edge taxes (in this case, the marginal cost taxes),

$$\frac{\sum_e f_e^*(l_e(f_e^*) + b_e)}{\sum_e \hat{f}_e l_e(\hat{f}_e)} \leq 2$$

for any equilibrium flow f^* and flow \hat{f} that achieves the social optimum. This bound is tight, and is not far from the $4/3$ upper bound on the *coordination ratio* in the case without taxes, shown by Roughgarden and Tardos [14]. The coordination ratio ρ was defined by Koutsoupias and Papadimitriou in [12] as follows

$$\rho := \sup_{f^*} \frac{\sum_e f_e^* l_e(f_e^*)}{\sum_e \hat{f}_e l_e(\hat{f}_e)}$$

for the worst case (in terms of social cost) equilibrium f^* , and \hat{f} as before. Hence we show that with a small increase in network inefficiency, we achieve, at equilibrium, a flow pattern that minimizes the total latency of the users. Note that, in principle, the tax b_e on an edge e could be very

big compared to the latency part $l_e()$ of the edge cost function. Hence it is rather surprising that taxation does not drive the social cost further than a small constant factor away from the social (without taxation) optimum.

This approach in bounding the inefficiency of taxation as a mechanism to achieve minimum social cost is influenced by the notion of coordination mechanisms. This concept was recently introduced by Christodoulou, Koutsoupias and Nanavati [3]. Informally speaking a coordination mechanism is a cost function experienced by the users, chosen from a family of possible cost functions called a coordination model. The measure of the efficiency of a coordination mechanism is the supremum over all possible demand sets of the ratio of the social cost of the worst-case equilibrium to the social optimum achieved with the *original* cost function. See [3] for the precise mathematical definitions. Note that in our work the demands are fixed.

Proving that taxation incurs a small increase to the cost of an equilibrium compared to the social optimum without taxation is a satisfying result. However it is possible that once taxes are imposed by some central authority they are considered to be part of the social cost. In other words taxation may incur disutility to society as a whole. To address this issue, we compare the worst-case cost with taxation of an equilibrium against the social optimum with taxation. In other words we consider the coordination ration in the standard sense [12] of the game with taxes. In Section 3 we show that, for certain families of strictly increasing and continuous latency functions (like linear or polynomial ones), the coordination ratio of the network actually *decreases* when optimal taxes are introduced. In particular for the linear latency functions case, we show that

$$\frac{\sum_e f_e^*(l_e(f_e^*) + b_e)}{\sum_e \bar{f}_e(l_e(\bar{f}_e) + b_e)} \leq \frac{5}{4}$$

for any equilibrium flow f^* and social optimum \bar{f} when taxes b are used in both cases. This is significantly better than the $4/3$ bound of [14] for general linear functions. The gap between $4/3$ and $5/4$ quantifies the beneficial effect of taxation on the behavior of the selfish users, specifically the reduction in their resistance to coordination. Our bound holds for heterogeneous users as well, and its proof is based on the definition of the parameter $\beta(\mathcal{L})$ for a family of functions \mathcal{L} by Correa, Schulz and Stier Moses [6].

The two results in combination show that imposing the optimal taxes on a selfish routing game not only yields a game with better coordination ratio, but also the added disutility to the users is bounded with respect

to the original system optimum. In addition we emphasize that our two approaches together provide a stronger guarantee on the worst-case cost of an equilibrium with taxation than each one of them taken separately. For a given tax vector b , let \bar{f} be the social optimum with taxation and \hat{f} be the social optimum without taxation. There does not seem to be any a priori information about which of the two quantities

$$\frac{5}{4} \sum_e \bar{f}_e (l_e(\bar{f}_e) + b_e), \quad 2 \sum_e \hat{f}_e l_e(\hat{f}_e)$$

is smaller.

2 The model

Let $G = (V, E)$ be a directed network (possibly with parallel edges but with no self-loops), and a set of *users*, each with an infinitesimal amount of traffic (flow) to be routed from an origin node to a destination node of G . Moreover, each user α has a positive *tax-sensitivity* factor $a(\alpha) > 0$. We will assume that the tax-sensitivity factors for all users come from a finite set of possible positive values. We can bunch together into a single *user class* all the users with the same origin-destination pair and with the same tax-sensitivity factor; let k be the number of different such classes. We denote by $d_i, \mathcal{P}_i, a(i)$ the total flow of class i , the flow paths that can be used by class i , and the tax-sensitivity of class i , for all $i = 1, \dots, k$ respectively. We will also use the term ‘commodity i ’ for class i . Set $\mathcal{P} \doteq \cup_{i=1, \dots, k} \mathcal{P}_i$. Each edge $e \in E$ is assigned a *latency function* $l_e(f_e)$ which gives the latency experienced by any user that uses e due to congestion caused by the total flow f_e that passes through e . In other words, as in [5], we assume the additive model in which for any path $P \in \mathcal{P}$ the latency is $l_P(f) = \sum_{e \in P} l_e(f_e)$, where $f_e = \sum_{P \ni e} f_P$ and f_P is the flow through path P . If every edge is assigned a per-unit-of-flow tax $\beta_e \geq 0$, a selfish user in class i that uses a path $P \in \mathcal{P}_i$ experiences total cost of

$$\sum_{e \in P} l_e(f_e) + a(i) \sum_{e \in P} \beta_e$$

hence the name ‘tax-sensitivity’ for the $a(i)$ ’s: they quantify the importance each user assigns to the taxation of a path.

Let \hat{f} be a flow that satisfies the users’ demands and minimizes the social cost $\sum_{e \in E} f_e l_e(f_e) = \sum_i \sum_{P \in \mathcal{P}_i} f_P l_P(f)$, i.e., \hat{f} is a solution of the following mathematical program:

$$\min \sum_{e \in E} f_e l_e(f_e) \quad \text{s.t.} \quad \text{(MP)}$$

$$\begin{aligned}
\sum_{P \in \mathcal{P}_i} f_P &= d_i && \forall i \\
f_e &= \sum_{P \in \mathcal{P}: e \in P} f_P && \forall e \in E \\
f_P &\geq 0 && \forall P
\end{aligned}$$

Note that, although in certain cases (e.g. when the latency functions l_e are convex) the flow \hat{f} can be computed efficiently, for more general latency functions it may be extremely difficult to compute \hat{f} (see Section 4 in [5]). We will assume that such an \hat{f} exists and that it induces finite latency on every edge.

A function $g(x)$ is *positive* if $g(x) > 0$ when $x > 0$. We assume that the functions l_e are strictly increasing, i.e., $x > y \geq 0$ implies $l_e(x) > l_e(y)$, and that $l_e(0) \geq 0$. This implies that $l_e(f_e) > 0$ when $f_e > 0$, i.e., the function l_e is positive. Similar assumptions on monotonicity are made in [5].

Let

$$K := \{f : 0 \leq f_P, \forall P \wedge \sum_{P \in \mathcal{P}_i} f_P = d_i, \forall i\}$$

be the set of all flows that satisfy the users' demands.

Definition 1. A flow f is called feasible iff $f \in K$.

A traffic (or Wardrop) equilibrium is a feasible flow $f^* \in K$ such that

$$\langle T(f^*), f - f^* \rangle \geq 0, \quad \forall f \in K. \quad (1)$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product, and $T_P(f)$ is the function that gives the generalized cost of a user that uses path P when the network flow is f .

3 Improving the coordination ratio

By extending the results of [5] from the single source-sink to the multi-commodity setting in [11], we have shown that there is a set of per-unit taxes b which forces the users to induce on the edges the same edge flow as \hat{f} . This result was also independently obtained by Fleischer, Jain and Mahdian [9].

Theorem 1. (Theorem 1 in [11]) Consider the selfish routing game with the latency function seen by the users in class i being

$$T_P(f) := \sum_{e \in E} l_e(f_e) + a(i) \sum_{e \in P} b_e, \quad \forall i, \quad \forall P \in \mathcal{P}_i.$$

If for every edge $e \in E$ l_e is a strictly increasing function with $l_e(0) \geq 0$, then there is a vector of per-unit taxes $b \in \mathbb{R}_+^{|E|}$ such that, if f^* is a traffic equilibrium for this game, $f_e^* = \hat{f}_e$, $\forall e \in E$. Therefore f^* minimizes the social cost $\sum_{e \in E} f_e l_e(f_e)$.

We also can compute this b in polynomial time when \hat{f} is given to us.

Correa et al. [6, 7] define the following quantity β for a continuous nondecreasing latency function $l : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ and every value $u \geq 0$:

$$\beta(u, l) := \frac{1}{ul(u)} \max_{x \geq 0} \{x(l(u) - l(x))\}, \quad (2)$$

where by convention $0/0 = 0$. In addition, they define $\beta(u) := \sup_{l \in \mathcal{L}} \beta(u, l)$ and $\beta(\mathcal{L}) := \sup_{l \in \mathcal{L}} \beta(l)$, where \mathcal{L} is a family of latency functions. Note that $\beta(l) \geq \beta(u, l)$, $\forall u, l$, and $\beta(\mathcal{L}) \geq \beta(l)$, $\forall l \in \mathcal{L}$. Also $0 \leq \beta(\mathcal{L}) < 1$.

Define the cost $C(f)$ of flow f to be $C(f) := \sum_P f_P T_P(f)$, and assume the additive model for l , i.e., $l_P(f) := \sum_{e \in P} l_e(f_e)$, where l_e is a function that gives the latency for edge e when flow f_e passes through it.

Theorem 2. *Let f^* be any traffic equilibrium for the game of Theorem 1, and \bar{f} be a flow that minimizes $C(f)$. Then, for latency functions l_e that belong to a family \mathcal{L}*

$$\frac{C(f^*)}{C(\bar{f})} \leq 1 + \beta(\mathcal{L}).$$

Proof. First we note that $C(f^*) = \langle T(f^*), f^* \rangle$. Also, because of Theorem 1 we have that $f_e^* = \hat{f}_e$. Then, for any feasible flow $\bar{f} \in K$ we have

$$\begin{aligned} \langle T(f^*), \bar{f} \rangle &= \sum_i \sum_{P \in \mathcal{P}_i} \bar{f}_P \left(l_P(f^*) + a(i) \sum_{e \in P} b_e \right) \\ &= \sum_{e \in E} \bar{f}_e l_e(f_e^*) + \sum_i \sum_{P \in \mathcal{P}_i} \bar{f}_P a(i) \sum_{e \in P} b_e \\ &= \sum_{e \in E} \bar{f}_e l_e(\hat{f}_e) + \sum_i \sum_{P \in \mathcal{P}_i} \bar{f}_P a(i) \sum_{e \in P} b_e \end{aligned} \quad (3)$$

where the second equality is due to the additive model.

We also have by definition that

$$\beta(\hat{f}_e, l_e) := \frac{1}{\hat{f}_e l_e(\hat{f}_e)} \max_{x \geq 0} \{x(l_e(\hat{f}_e) - l_e(x))\},$$

from which we get for $x := \bar{f}_e$ that

$$\beta(\hat{f}_e, l_e) \hat{f}_e l_e(\hat{f}_e) \geq \bar{f}_e l_e(\hat{f}_e) - \bar{f}_e l_e(\bar{f}_e)$$

or, by the definition of $\beta(\mathcal{L})$,

$$\beta(\mathcal{L}) \hat{f}_e l_e(\hat{f}_e) \geq \bar{f}_e l_e(\hat{f}_e) - \bar{f}_e l_e(\bar{f}_e). \quad (4)$$

By the definition of \hat{f} we know that $\sum_{e \in E} \hat{f}_e l_e(\hat{f}_e) \leq \sum_{e \in E} \bar{f}_e l_e(\bar{f}_e)$, therefore (4) implies that

$$(1 + \beta(\mathcal{L})) \cdot \sum_{e \in E} \bar{f}_e l_e(\bar{f}_e) \geq \sum_{e \in E} \bar{f}_e l_e(\hat{f}_e),$$

and this, in turn, together with (3) implies that

$$\begin{aligned} \langle T(f^*), \bar{f} \rangle &\leq (1 + \beta(\mathcal{L})) \cdot \sum_{e \in E} \bar{f}_e l_e(\bar{f}_e) + \sum_i \sum_{P \in \mathcal{P}_i} \bar{f}_P a(i) \sum_{e \in P} b_e \\ &\leq (1 + \beta(\mathcal{L})) \cdot \sum_i \sum_{P \in \mathcal{P}_i} \bar{f}_P \left(l_P(\bar{f}) + a(i) \sum_{e \in P} b_e \right) \\ &= (1 + \beta(\mathcal{L})) \cdot C(\bar{f}). \end{aligned} \quad (5)$$

Since f^* is an equilibrium, (1) implies that

$$\langle T(f^*), f \rangle \geq \langle T(f^*), f^* \rangle, \quad \forall f \in K,$$

which, together with (5) and for $f := \bar{f}$, implies that

$$(1 + \beta(\mathcal{L})) \cdot C(\bar{f}) \geq C(f^*).$$

Correa et al. [6] give upper bounds of $\beta(\mathcal{L})$ for some function families \mathcal{L} . We repeat here these bounds (Corollaries 4.3, 4.4 in [6]):

Corollary 1. [6] *If the set \mathcal{L} of continuous and nondecreasing latency functions is contained in the set $\{l(\cdot) : l(cx) \geq cl(x) \text{ for } c \in [0, 1]\}$, then $\beta(\mathcal{L}) \leq 1/4$.*

Note that Corollary 1 holds for the family of nondecreasing linear functions, hence the worst case bound of Theorem 2 for linear functions is $5/4$ which is better than the tight $4/3$ bound [14] achieved without the use of the b 's from Theorem 1.

Corollary 2. [6] *If the set \mathcal{L} of continuous and nondecreasing latency functions is contained in the set $\{l(\cdot) : l(cx) \geq c^n l(x) \text{ for } c \in [0, 1]\}$ for some positive number n , then*

$$\beta(\mathcal{L}) \leq \frac{n}{(n+1)^{1+1/n}}.$$

Therefore for $n = 2, 3$ and 4 the upper bound becomes $1.385, 1.472$ and 1.535 respectively, as opposed to $1.626, 1.896$ and 2.151 respectively when b is not used [6]. Also note that as n increases, the bound goes to 2 .

4 Comparison to the original latencies

In general, the values of b in Theorem 1 can be very big. It may even be the case that the part of the cost $C(f)$ due to the initial latencies $\sum_e f_e l_e(f_e)$ is negligible compared to the part due to b , which is $\sum_i \sum_P a(i) f_P \sum_{e \in P} b_e$. Therefore the improvement of the coordination ratio may come at a prohibitive increase to the overall cost. One would like to bound b so that the new overall cost is bounded by a function of the original optimal total latency $\sum_e \hat{f}_e l_e(\hat{f}_e)$.

Unfortunately, we do not know how to bound b for the general $l(\cdot)$ of Theorem 1. But we can use already known results in the case of *homogeneous* users, i.e., $a(i) = 1, \forall i$, to bound the ratio of the worst equilibrium cost when b is used to the original optimal total cost.

It is well known ([2],[8],[15]; see also [4], especially Proposition 3.1) that, for homogeneous networks with differentiable latency functions l_e , one can use the *marginal costs*³ $\hat{f}_e l'_e(\hat{f}_e)$ as b_e in Theorem 1 to achieve the following classical result:

Theorem 3. *If functions l_e are differentiable, then \hat{f} is an equilibrium for the selfish routing game with*

$$T_P(f) := \sum_{e \in P} (l_e(f_e) + \hat{f}_e l'_e(\hat{f}_e)).$$

Moreover, if we assume that l_e are strictly increasing (as in Theorem 1), then any equilibrium f^* incurs the same edge flow as \hat{f} , i.e., $f_e^* = \hat{f}_e, \forall e$ (Theorem 6.2 in [1]). Let $C_{OPT} := \sum_{e \in E} \hat{f}_e l_e(\hat{f}_e)$.

³ $l'_e(\cdot)$ is the derivative of $l_e(\cdot)$.

Theorem 4. *If $l_e(f_e) = a_e f_e + \beta_e$ with $a_e > 0, b_e \geq 0$ for all $e \in E$, and f^* is an equilibrium for the selfish routing game with taxes $b_e := a_e \hat{f}_e$, then*

$$\frac{C(f^*)}{C_{OPT}} \leq 2.$$

More generally, if the l_e 's are polynomials of degree d with positive coefficients, then

$$\frac{C(f^*)}{C_{OPT}} \leq d + 1.$$

Proof. For the case of linear latency functions, note that these functions are differentiable, therefore Theorem 3 implies that \hat{f} is an equilibrium. Since they are strictly increasing as well, we know that $f_e^* = \hat{f}_e, \forall e$, for any equilibrium f^* . Hence

$$\begin{aligned} C(f^*) &:= \sum_P f_P^* T_P(f^*) = \sum_e f_e^* (a_e f_e^* + \beta_e + a_e \hat{f}_e) \\ &= \sum_e \hat{f}_e (2a_e \hat{f}_e + \beta_e) \\ &\leq 2 \cdot C_{OPT} \end{aligned}$$

The same argument shows the upper bound for the case of degree d polynomials.

Note that the bound above is tight in the case of polynomials with all the coefficients, except for the one of the highest degree factor, being 0.

Theorem 5. *If the l_e functions are strictly increasing and continuously differentiable with a convex derivative, and f^* is an equilibrium for the selfish routing game with taxes $b_e := a_e \hat{f}_e$, then*

1. *If $l_e(\cdot) \in \{l(\cdot) : l(cx) \geq cl(x) \text{ for } c \in [0, 1]\}$, then*

$$\frac{C(f^*)}{C_{OPT}} \leq 3,$$

2. *If $l_e(\cdot) \in \{l(\cdot) : l(cx) \geq c^n l(x) \text{ for } c \in [0, 1]\}$, then*

$$\frac{C(f^*)}{C_{OPT}} \leq 2^n + 1.$$

Proof. Theorem 3 implies that \hat{f} is an equilibrium. Since the l_e 's are strictly increasing as well, we know that $f_e^* = \hat{f}_e, \forall e$, for any equilibrium f^* . Observe that the edges e for which $\hat{f}_e = 0$ do not contribute to $C(f^*)$

therefore we ignore them in the ensuing calculations. It is known that if the l'_e functions are convex and continuous in (γ, δ) , then

$$l'_e(x) \leq \frac{1}{2h} \int_{x-h}^{x+h} l'_e(t) dt \quad (6)$$

for $\gamma \leq x-h < x < x+h \leq \delta$ (fact 125, p. 98 of [10]). For $x := f_e/2, h := f_e/2$, inequality (6) becomes

$$f_e l'_e\left(\frac{f_e}{2}\right) \leq \int_0^{f_e} l'_e(t) dt. \quad (7)$$

Under the assumptions of Part 1 of the theorem, (7) gives

$$f_e l'_e(f_e) \leq 2l_e(f_e)$$

which, for $f_e := \hat{f}_e$, implies together with an argument similar to the proof of Theorem 4 that

$$\frac{C(f^*)}{\sum_e \hat{f}_e l_e(\hat{f}_e)} \leq 3$$

for any equilibrium flow f^* .

Under the assumptions of Part 2 of the theorem, (7) implies that

$$f_e l'_e(f_e) \leq 2^n l_e(f_e)$$

which, in turn, implies that in this case

$$\frac{C(f^*)}{\sum_e \hat{f}_e l_e(\hat{f}_e)} \leq 2^n + 1$$

for any equilibrium flow f^* .

Acknowledgement. Thanks to Elias Koutsoupias for useful discussions.

References

1. H. Z. Aashtiani and T. L. Magnanti. Equilibria on a congested transportation network. *SIAM Journal of Algebraic and Discrete Methods*, 2:213–226, 1981.
2. M. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.
3. G. Christodoulou, E. Koutsoupias, and A. Nanavati. Coordination mechanisms. To appear in Proc. 31st ICALP, 2004.

4. R. Cole, Y. Dodis, and T. Roughgarden. How much can taxes help selfish routing? In *Proceedings of the 4th ACM Conference on Electronic Commerce*, pages 98–107, 2003.
5. R. Cole, Y. Dodis, and T. Roughgarden. Pricing network edges for heterogeneous selfish users. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, pages 521–530, 2003.
6. J. R. Correa, A. S. Schulz, and N. E. Stier Moses. Selfish routing in capacitated networks. Technical Report Working Paper 4319-03, MIT Sloan School of Management, Cambridge, MA, June 2003.
7. J. R. Correa, A. S. Schulz, and N. E. Stier Moses. Selfish routing in capacitated networks. To appear in *Mathematics of Operations Research*, February 2004.
8. S. Dafermos and F. T. Sparrow. The traffic assignment problem for a general network. *Journal of Research of the National Bureau of Standards, Series B*, 73B:91–118, 1969.
9. L. Fleischer, K. Jain, and M. Mahdian. Taxes for heterogeneous selfish users in a multicommodity network. To appear in *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, 2004.
10. G. Hardy, J. E. Littlewood, and G. Pólya. *Inequalities - Second Edition*. Cambridge, 1934.
11. G. Karakostas and S. G. Kolliopoulos. Edge pricing of multicommodity networks for heterogeneous selfish users. To appear in *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, 2004.
12. E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, pages 404–413, 1999.
13. M. J. Osborne and A. Rubinstein. *A course in Game Theory*. The MIT Press, Cambridge, MA, 1994.
14. T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49:236–259, 2002.
15. M. J. Smith. The marginal cost taxation of a transportation network. *Transportation Research*, 13B:237–242, 1979.
16. J. G. Wardrop. Some theoretical aspects of road traffic research. *Proc. Inst. Civil Engineers, Part II*, 1:325–378, 1952.