

Selfish routing with oblivious users

George Karakostas^{1*}, Taeyon Kim^{1**}, Anastasios Viglas², and
Hao Xia^{1***}

¹ McMaster University, Dept. of Computing and Software, 1280 Main St. West,
Hamilton, Ontario L8S 4K1, Canada, {karakos,kimt22,xiah}@mcmaster.ca.

² University of Sydney, School of Information Technologies, Madsen Building F09,
The University of Sydney, NSW 2006, Australia, tastos@it.usyd.edu.au.

Abstract. We consider the problem of characterizing user equilibria and optimal solutions for selfish routing in a given network. We extend the known models by considering users *oblivious to congestion*. While in the typical selfish routing setting the users follow a strategy that minimizes their individual cost by taking into account the (dynamic) congestion due to the current routing pattern, an *oblivious* user ignores congestion altogether. Instead, he decides his routing on the basis of cheapest routes on a network without any flow whatsoever. These cheapest routes can be, for example, the shortest paths in the network without any flow. This model tries to capture the fact that routing tables for at least a fraction of the flow in large scale networks such as the Internet may be based on the physical distances or hops between routers alone. The phenomenon is similar to the case of traffic networks where a certain percentage of travelers base their route simply on the distances they observe on a map, without thinking (or knowing, or caring) about the delays experienced on this route due to their fellow travelers. In this work we study the price of anarchy of such networks, i.e., the ratio of the total latency experienced by the users in this setting over the optimal total latency if all users were centrally coordinated.

Keywords: Selfish routing, price of anarchy, oblivious users.

* Research supported by an NSERC Discovery grant and MITACS.

** Research supported by an NSERC Discovery grant and MITACS.

*** Research supported by an NSERC Discovery grant and MITACS.

1 Introduction

The general framework of a system of non-cooperative users can be used to model many different optimization problems such as network routing, traffic or transportation problems, load balancing and distributed computing, auctions and many more. Game-theoretic techniques can be used to model and analyze such systems in a natural way. The performance of a system of non-cooperative users is measured by an appropriate cost function which depends on the behaviour, or strategies of the users. For example in the case of network routing, the total, system-wide cost can be defined as the total routing cost, or the total latency experienced by all the users in the network. On the other hand, there is also a cost associated with each individual user (for example the latency experienced by the user). It is a well known fact that if each user optimizes her own cost, then they might choose a strategy that does not give the optimal total cost for the entire system, also known as *social cost*. In other words, the *selfish* behaviour of the users leads to a sub-optimal performance.

Koutsoupas and Papadimitriou [3] initiated the study of the *coordination ratio* (also referred to as the price of anarchy): How much worse is the performance of a system of selfish users where each user optimizes her own cost, compared to the best possible performance that can be achieved on the same system? In particular, this question was first studied in the setting of selfish network routing by Roughgarden and Tardos [5]. In this model, the network users experience edge latencies that depend on the congestion on each edge according to some latency function. Given a particular flow pattern, the users decide to route their flow through paths of minimum latency. 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 [6] 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.

In the past, several variations of this basic model have been considered. For example, Roughgarden [4] studied the case of combining selfish and centrally coordinated users on the same network, proposing *Stackelberg strategies* for the latter that would improve the price of anarchy. Karakostas and Viglas [2] studied the combination of selfish and malicious users: A malicious user will choose a strategy that will cause the worst possible performance for the entire network. These models try to capture a richer set of paradigms in networks such as the Internet, where traffic

does not consist by users of the same profile or behavior. In this work we introduce a new paradigm that is based on the following observation: A fundamental assumption in the basic selfish routing model is that each user is able to measure the latencies of *all* paths available to him at any moment, in order to pick the best possible path currently available for his flow. It is clear that in very large networks this assumption is probably quite unrealistic, since it may not be possible to measure these latencies or measure them as often as needed. Hence it may be easier for a fraction α of the network users to consult predefined routing tables based on *non-dynamic* parameters of the network, such as the physical distances between nodes. The price these users pay for the convenience is a degree of naivety in their decisions, since they are completely *oblivious* to congestion phenomena. We call such users *oblivious*.

More specifically, we consider oblivious users that route their flow through the shortest path connecting its origin to its destination, as measured in the network without flow. We study the price of anarchy in case of linear edge latency functions, first for a (single commodity) single pair of nodes connected by a set of parallel edges, and then for general topologies with an arbitrary number of origin-destination pairs. Unlike the case of selfish routing without oblivious users where the price of anarchy is bounded by $4/3$ [5], our bounds are not independent of the network parameters. For the both the cases of parallel links and general topologies, the bounds depend on the coefficients a_e of the linear latency functions $l_e(f_e) = a_e f_e + b_e$ for edges e , where f_e is the total flow through edge e . In addition, the general topology bound depends on the minimum fraction of total demand that the optimal routing sends through any edge. Although these bounds can be very large, if, for example, there are network edges with vastly different behavior under congestion (as is the case in a traffic network with both highways and side-streets), this seems to be unavoidable in view of the fact that the myopic behavior of oblivious users may lead to great congestion of ‘wide’ edges by them, and in that way directing the selfish users to non-congested but ‘narrow’ paths. Also, the dependence of the general topology bound on the ‘spread’ of the optimal flow seems to be necessary, given that the oblivious flow concentrates the oblivious users on specific (initially fastest, but possibly very slow after the selfish users have been added) paths, which may be orthogonal to what the optimal flow does.

Organization: In Section 2 we define the model, in Section 3 we study linear latency functions in simple networks of two nodes connected by

parallel links, and in Section 4 we study linear functions in general, multicommodity networks. We conclude with a discussion in Section 5.

2 Preliminaries

We are given a directed network $G = (V, E)$ with a latency function $l_P(\cdot)$ associated to each path P . For a flow f on G , $l_P(f)$ is the latency (cost) of path P for this particular flow. Notice that in general this latency depends on the whole flow f , and not only on the flow f_e through each edge $e \in P$. In this paper we adopt the *additive model* for the path latencies, i.e., $l_P(f) = \sum_{e \in P} l_e(f_e)$, where l_e is the latency function for edge e and f_e is the amount of flow that goes through e . We also let \mathcal{P} be the set of all available paths in the network and assume that for every source-sink pair there is at least one path joining the source to the sink. In this work we assume that the latency functions are linear functions of the edge flow f_e , i.e., $l_e(f_e) = a_e f_e + b_e$, $\forall e \in E$. The total cost of a flow f is defined as $C(f) = \sum_{e \in E} f_e l_e(f_e)$.

We consider the case where for every origin-destination commodity of demand d , a fraction α of it consists of an infinite number of *oblivious* users, each carrying an infinitesimal amount of flow through the shortest path connecting the source to the destination when there is no flow routed on G . If there are more than one shortest paths, we will assume that all these users pick the smallest in a lexicographic ordering. The rest $(1 - \alpha)d$ of the demand consists of an infinite number of *selfish* users, each carrying an infinitesimal amount of flow.

3 Parallel links

Let G be a network consisting of parallel links connecting two nodes s, t . We will assume that the edge latency functions are strictly increasing, i.e., for every edge e with $l_e(f_e) = a_e f_e + b_e$ we have $a_e > 0$. Note that in this setting, both the traffic equilibrium flow and the optimal flow are unique. In what follows, we use $0 \leq \alpha \leq 1$ to denote the fraction of total flow from s to t that is oblivious. We will use the term ‘traffic equilibrium’ for flows with $\alpha = 0$ that are at traffic equilibrium, while we reserve the term ‘oblivious equilibrium’ for flows with $\alpha > 0$ and with their selfish users at traffic equilibrium in the network that results *after* routing the oblivious users. The following observation is true due to Wardrop’s principle and since the latency functions are increasing:

Proposition 1 Let f^d and $f^{d+\delta}$ be flows at traffic equilibrium, of demand d and $d + \delta$ respectively, with $\delta \geq 0$. Then we have $f^d \leq f^{d+\delta}$.

In what follows, we denote the total demand from s to t with d , and the optimal flow of demand d with f^{opt} . We denote the flow of demand d at traffic equilibrium with f^* , and the flow of demand d at oblivious equilibrium with \tilde{f} , where the flow of oblivious users with total demand αd is denoted with \tilde{f}^o and the flow of selfish users (with total demand $(1 - \alpha)d$) with \tilde{f}^* , i.e., $\tilde{f} = \tilde{f}^* + \tilde{f}^o$. Obviously, the oblivious flow will be routed through the edge e with the smallest $l_e(0) = b_e$ (or the first such edge in a lexicographic ordering, if there are more than one). Let e_s be this edge.

Proposition 2 $f_e^* \geq \tilde{f}_e^*$, $\forall e \in E$.

Proof: If $f_{e_s}^* \geq \alpha d$ then $f^* = \tilde{f}$. Otherwise we have $\alpha d > f_{e_s}^*$. In this case, no selfish users will flow through e_s because of Wardrop's principle, i.e., $\tilde{f}_{e_s}^* = 0$. By removing e_s from G together with the portion of flow on it, we get two new Nash flows $f^{*'}$ and $\tilde{f}^{*'}$, of demand $d - f_{e_s}^*$ and $d - \alpha d$. As $\alpha d > f_{e_s}^*$, from Proposition 1, we have $f^{*'} \geq \tilde{f}^{*'}$. Then since $\tilde{f}_{e_s}^* = 0$, we get $f^* \geq \tilde{f}^*$. \square

Lemma 1. $C(f^*) \leq \frac{4}{3}(1 - \alpha)C(f^{opt})$.

Proof: Since latency functions are increasing and $f^* \geq \tilde{f}^*$ from Proposition 2, we know that $\forall e$, $l_e(f_e^*) \geq l_e(\tilde{f}_e^*)$. Also Wardrop's principle for f^* implies that $l_e(f_e^*) = L(f^*)$, $\forall e : f_e^* > 0$ and $l_e(f_e^*) \geq L(f^*)$, $\forall e$, where $L(f^*)$ is the common latency of the paths used by the traffic equilibrium flow f^* . Then

$$\begin{aligned} C(f^*) - C(\tilde{f}^*) &= \sum_e \left(l_e(f_e^*)f_e^* - l_e(\tilde{f}_e^*)\tilde{f}_e^* \right) \\ &\geq \sum_e (f_e^* - \tilde{f}_e^*)l_e(f_e^*) \\ &\geq \sum_e (f_e^* - \tilde{f}_e^*)L(f^*) = \alpha d L(f^*) = \alpha C(f^*). \end{aligned}$$

Then Theorem 4.5 of [5] implies the lemma. \square

The last proposition together with Wardrop's principle implies that $l_e^*(f_e^{opt}) = l_{e_s}^*(f_{e_s}^{opt})$, $\forall e : f_e^{opt} > 0$ and $l_e^*(f_e^{opt}) \geq l_{e_s}^*(f_{e_s}^{opt})$, $\forall e$.

Proposition 3 $f_{e_s}^{opt} > 0$.

Proof: Let e be an edge with $f_e^{opt} > 0$. The KKT conditions imply that f^{opt} is a traffic equilibrium for latency functions $l_e^*(f_e) = 2a_e f_e + b_e$. By applying Wardrop's principle for this equilibrium, we get

$$\begin{aligned} l_{e_s}^*(f_{e_s}^{opt}) &\geq l_e^*(f_e^{opt}) = 2a_e f_e^{opt} + b_e \\ &> b_e \\ &\geq b_{e_s} = l_{e_s}^*(0). \end{aligned}$$

Therefore $f_{e_s}^{opt} > 0$. □

In what follows, let $\mathcal{E}^{opt} = \{e : f_e^{opt} > 0\}$.

Proposition 4 *For any edge e , we have $l_{e_s}(f_{e_s}^{opt}) \leq l_e(f_e^{opt})$. For any edge $e \in \mathcal{E}^{opt}$, we have $a_{e_s} f_{e_s}^{opt} \geq a_e f_e^{opt}$.*

Proof: For any edge e , we have

$$\begin{aligned} l_e(f_e^{opt}) &= \frac{1}{2} (l_e^*(f_e^{opt}) + b_e) \\ &\geq \frac{1}{2} (l_{e_s}^*(f_{e_s}^{opt}) + b_e) \\ &\geq \frac{1}{2} (l_{e_s}^*(f_{e_s}^{opt}) + b_{e_s}) \\ &= l_{e_s}(f_{e_s}^{opt}). \end{aligned}$$

Similarly, we get the second part of the proposition for any edge $e \in \mathcal{E}^{opt}$. □

Lemma 2. $C(\tilde{f}^o) \leq \max\{\alpha, \alpha^2 r\} C(f^{opt})$, where $r = \sum_{e \in \mathcal{E}^{opt}} (a_{e_s}/a_e)$.

Proof: We have

$$\begin{aligned} C(f^{opt}) &= \sum_e l_e(f_e^{opt}) f_e^{opt} \\ &\geq l_{e_s}(f_{e_s}^{opt}) \sum_e f_e^{opt} \quad (\text{Proposition 4}) \\ &= (a_{e_s} f_{e_s}^{opt} + b_{e_s}) d. \end{aligned}$$

From the second part of Proposition 4 we have that $(a_{e_s}/a_e) f_{e_s}^{opt} \geq f_e^{opt}$, $\forall e \in \mathcal{E}^{opt}$. By summing over all $e \in \mathcal{E}^{opt}$, we get $a_{e_s} f_{e_s}^{opt} \geq a_e d/r$. Thus

$$C(f^{opt}) \geq \left(\frac{a_{e_s}}{r} d + b_{e_s} \right) d.$$

Therefore

$$\frac{C(\tilde{f}^o)}{C(f^{opt})} \leq \frac{(a_{e_s} \alpha^2 d + \alpha b_{e_s}) d}{\left(\frac{a_{e_s}}{r} d + b_{e_s}\right) d} \leq \max\{\alpha, \alpha^2 r\}.$$

□

Theorem 1. *If $\tilde{f}_{e_s}^o = \alpha d \geq f_{e_s}^*$, then $\frac{C(\tilde{f})}{C(f^{opt})} \leq \frac{4}{3}(1 - \alpha) + \max\{\alpha, \alpha^2 r\}$, otherwise $\frac{C(\tilde{f})}{C(f^{opt})} \leq \frac{4}{3}$.*

Proof: If $\tilde{f}_{e_s}^o = \alpha d < f_{e_s}^*$ then $\tilde{f} = f^*$ and the second part of the theorem follows. In the case $\alpha d \geq f_{e_s}^*$, edge e_s which is used by the oblivious users is no longer attractive to selfish users, i.e., $\tilde{f}_{e_s}^* = 0$. Thus \tilde{f}^* and \tilde{f}^o are actually orthogonal, i.e., $\tilde{f}^{*T} \tilde{f}^o = 0$. Then, if $A > 0$ is the $|E| \times |E|$ diagonal matrix whose diagonal elements are the a_e 's, we have

$$\begin{aligned} C(\tilde{f}) &= (A(\tilde{f}^* + \tilde{f}^o) + b)^T (\tilde{f}^* + \tilde{f}^o) \\ &= \tilde{f}^{*T} A \tilde{f}^* + \tilde{f}^{oT} A \tilde{f}^o + 2\tilde{f}^{*T} A \tilde{f}^o + b^T (\tilde{f}^* + \tilde{f}^o) \\ &= (A \tilde{f}^* + b)^T \tilde{f}^* + (A \tilde{f}^o + b)^T \tilde{f}^o \end{aligned}$$

and the first part of the theorem now comes from Lemmata 1 and 2. □

3.1 An example for parallel links

We provide an example to show that in networks with parallel links, the price of anarchy can be as bad as our bound in Theorem 1 in case $\alpha = 1$, i.e., all users are oblivious. The network has only two links, namely e_1 and e_2 , with latency functions $l_1(x) = 10x$ and $l_2(x) = x + \epsilon$, where $0 < \epsilon < 1$. The total demand is $d = 1$.

The optimal cost in this setting is $C_{opt} = 10/11 + (40\epsilon - \epsilon^2)/44$. When $\alpha \geq (1 + \epsilon)/11$, the cost of the oblivious equilibrium is $C_{eq} = 11\alpha^2 - (2 + \epsilon)\alpha + (1 + \epsilon)$. One can see that when α is one (all users are oblivious), and ϵ tends to zero, the price of anarchy is

$$\lim_{\epsilon \rightarrow 0} \frac{C_{eq}}{C_{opt}} = 11,$$

which is exactly the bound we get in Theorem 1.

However, this example is not tight when $\alpha < 1$. The loss of tightness comes from the $4/3$ which is the upper bound for the selfish routing price of anarchy [5]. We used this result directly in the last step of Lemma 1

and in the first case of Theorem 1. While this example is a tight example for Lemma 2, it is not tight for the $4/3$. The price of anarchy here is very close to 1. Thus a real tight example for our bound would be one that is tight for both $4/3$ and Lemma 2. Unfortunately such an ideal example does not exist since the tightness of Lemma 2 requires very small b_e/a_e for all links, but in order to make $4/3$ tight we need a relatively large b_e/a_e to make a distinction between the selfish flow and the optimal flow. This implies that the bound in Theorem 1 is not tight, and a tighter bound remains as an open problem.

4 General topologies

In this section we study the price of anarchy of oblivious equilibria for general topologies, arbitrary number of origin-destination pairs (commodities) and linear latency functions. We will use the concept of β -function defined in [1]. Let \mathcal{L} be a family of continuous and nondecreasing latency functions. For every function $l \in \mathcal{L}$ and every value $v \geq 0$, let us define:

$$\beta(v, l) := \frac{1}{vl(v)} \max_{x \geq 0} \{x(l(v) - l(x))\}.$$

In addition, let us define

$$\beta(l) := \sup_{v \geq 0} \beta(v, l),$$

and

$$\beta(\mathcal{L}) := \sup_{l \in \mathcal{L}} \beta(l).$$

We will denote the inner product of two vectors x, y by $\langle x, y \rangle$.

We will also use an alternative characterization of a traffic equilibrium f^* of demand d , as a solution to the following variational inequality:

$$\langle l(f^*), f - f^* \rangle \geq 0, \quad \forall f \in \{f : f \text{ is a flow satisfying demand } d\}.$$

By applying this formulation to the selfish part of an oblivious equilibrium $\tilde{f} = f^* + \tilde{f}^o$ in the network obtained after the oblivious users have been routed³, we get

$$\langle l(\tilde{f}), f - \tilde{f}^* \rangle \geq 0, \quad \forall f \in \{f : f \text{ is a flow satisfying demand } (1 - \alpha)d\}.$$

By setting $f := (1 - \alpha)f^{opt}$ we have

$$\langle l(\tilde{f}), \tilde{f}^* \rangle \leq (1 - \alpha) \langle l(\tilde{f}), f^{opt} \rangle. \quad (1)$$

³ Note that the new latency functions in this network are $l'_e(f_e) = a_e(\tilde{f}_e^o + f_e) + b_e = l_e(f_e + \tilde{f}_e^o)$.

Lemma 3. $\langle l(\tilde{f}), \tilde{f}^* \rangle \leq (1 - \alpha)\beta(\mathcal{L})C(\tilde{f}) + (1 - \alpha)C(f^{opt})$.

Proof:

$$\begin{aligned} \langle l(\tilde{f}), \tilde{f}^* \rangle &\stackrel{(1)}{\leq} (1 - \alpha)\langle l(\tilde{f}), f^{opt} \rangle \\ &\leq (1 - \alpha) \left(\sum_e \beta(\tilde{f}_e, l_e) l_e(\tilde{f}_e) \tilde{f}_e + \sum_e l_e(f_e^{opt}) f_e^{opt} \right) \\ &\leq (1 - \alpha) \left(\beta(\mathcal{L})C(\tilde{f}) + C(f^{opt}) \right). \end{aligned}$$

□

Lemma 4. $\langle l(\tilde{f}), \tilde{f}^o \rangle \leq \frac{n\alpha d_i \gamma_a}{f_{min}^{opt}} C(f^{opt})$, where $n = |V|$, $\gamma_a = \frac{\max_e a_e}{\min_e a_e}$, and $f_{min}^{opt} = \min_e f_e^{opt}$.

Proof: Let P_{s_i} be the path used by the oblivious users corresponding to the i -th origin-destination pair (commodity). Note that this is the shortest path amongst all possible paths \mathcal{P}_i connecting this pair when we define the edge distances as $l_e(0) = b_e$. Also let d_i be the demand for this pair, therefore αd_i is the amount of oblivious flow routed through P_{s_i} . Let also $a_{min} = \min_e a_e$, $a_{max} = \max_e a_e$.

The key observation is that the oblivious flow \tilde{f}^o is a traffic equilibrium for the original network, if we define its latency functions as $l_e^o(f) = b_e$. From the discussion above, this implies that

$$\langle l(0), f - \tilde{f}^o \rangle \geq 0, \quad \forall f \in \{f : f \text{ is a flow satisfying demand } \alpha d\},$$

or, if we set $f := \alpha f^{opt}$, we get

$$\sum_{e \in E} b_e \tilde{f}_e^o \leq \alpha \sum_{e \in E} b_e f_e^{opt}. \quad (2)$$

Then

$$\begin{aligned} \langle l(\tilde{f}), \tilde{f}^o \rangle &= \sum_{e \in E} \left(a_e(\tilde{f}_e^* + \tilde{f}_e^o) \tilde{f}_e^o + b_e \tilde{f}_e^o \right) \\ &\stackrel{(2)}{\leq} \alpha \sum_i d_i \sum_{e \in P_{s_i}} a_e(\tilde{f}_e^* + \tilde{f}_e^o) + \alpha \sum_{e \in E} b_e f_e^{opt}. \end{aligned} \quad (3)$$

To get an upper bound of the first term:

$$\alpha \sum_i d_i \sum_{e \in P_{s_i}} a_e(\tilde{f}_e^* + \tilde{f}_e^o) = \alpha \sum_i d_i \sum_{e \in P_{s_i}} a_e(\tilde{f}_e^* + \alpha d_i)$$

$$\begin{aligned}
&\leq n\alpha^2 a_{max} \sum_i d_i^2 + \alpha \sum_i d_i \sum_{e \in P_{s_i}} a_e \tilde{f}_e^* \\
&\leq n\alpha^2 a_{max} \sum_i d_i^2 + n\alpha a_{max} (1 - \alpha) \sum_i d_i^2 \\
&\leq n\alpha d \gamma_a (a_{min} d) \\
&\leq n\alpha d \gamma_a \sum_{e \in E} a_e f_e^{opt} \\
&\leq \frac{n\alpha d \gamma_a}{f_{min}^{opt}} \sum_{e \in E} a_e f_e^{opt^2}. \tag{4}
\end{aligned}$$

Since $\frac{n\alpha d \gamma_a}{f_{min}^{opt}} \geq \alpha$ the combination of (3),(4) proves the lemma. □

Theorem 2.

$$\frac{C(\tilde{f})}{C(f^{opt})} \leq \frac{4 \left(1 - \alpha + n\alpha d \gamma_a / f_{min}^{opt} \right)}{3 + \alpha}$$

Proof: By combining Lemma 3 with Lemma 4, we have

$$\begin{aligned}
C(\tilde{f}) &= \langle l(\tilde{f}), \tilde{f}^o \rangle + \langle l(\tilde{f}), \tilde{f}^* \rangle \\
&\leq (1 - \alpha) C(f^{opt}) + (1 - \alpha) \beta(\mathcal{L}) C(\tilde{f}) + \frac{n\alpha d \gamma_a}{f_{min}^{opt}} C(f^{opt}),
\end{aligned}$$

hence

$$\frac{C(\tilde{f})}{C(f^{opt})} \leq \frac{1 - \alpha + n\alpha d \gamma_a / f_{min}^{opt}}{1 - (1 - \alpha) \beta(\mathcal{L})}.$$

For \mathcal{L} being the set of nondecreasing linear functions $\beta(\mathcal{L}) = \frac{1}{4}$ [1], and the theorem follows. □

5 Discussion and open problems

The obvious open problem is the tightening of the bounds of Theorems 1,2. One method of doing so seems to be the avoidance of relating the cost of oblivious equilibria to the optimal cost via traffic equilibria. It is precisely this intermediate step that doesn't allow us yet to have a tight analysis for Theorem 1.

Especially Theorem 2 for general topologies may be possible to be improved by removing its dependence on the minimum optimum path flow f_{min}^{opt} . Although the optimum flow is a parameter of the network, it may be very difficult to be determined by the network designer, while the other parameters of the network (G, n, d, a_e, b_e) can be set directly. Finally, it would be interesting to get non-trivial bounds (if they exist) for general latency functions.

References

1. J. R. Correa, A. S. Schulz, and N. E. Stier Moses. Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29:961–976, 2004.
2. G. Karakostas and A. Viglas. Equilibria for networks with malicious users. *Mathematical Programming Series A*, published online July 29, 2006, DOI: 10.1007/s10107-006-0015-2.
3. 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.
4. T. Roughgarden. Stackelberg scheduling strategies. *SIAM Journal on Computing*, 33:332–350, 2004.
5. T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49:236–259, 2002.
6. J. G. Wardrop. Some theoretical aspects of road traffic research. *Proc. Inst. Civil Engineers, Part II*, 1:325–378, 1952.