Equilibria for Networks with Malicious Users*

George Karakostas[1] and Anastasios Viglas[2]

[1] McMaster University, Dept. of Computing and Software, 1280 Main St. West, Hamilton, Ontario L8S 4K1, Canada, gk@cas.mcmaster.ca
[2] University of Toronto, Computer Science Department, 10 King's College Road, Toronto, ON M5S 3G4, Canada aviglas@cs.toronto.edu

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 malicious behaviour. While selfish users follow a strategy that minimizes their individual cost, a *malicious* user will use his flow through the network in an effort to cause the maximum possible damage to this cost. We define a generalized model, present characterizations of flows at Wardrop equilibrium and prove bounds for the ratio of the social cost of a flow at Wardrop equilibrium over the cost when centralized coordination among users is allowed.

1 Introduction

Koutsoupias and Papadimitriou [5] initiated the study of the *coordination ratio* (also referred to as the price of anarchy): How much worse is the performance of a network of selfish users where each user optimizes her own cost, compared to the best possible performance that can be achieved on the same system? This question has been studied in various different models (e.g. [11], [12]) and bounds for the coordination ratio have been shown for many interesting cases.

A basic assumption of the models considered so far is that the users are considered to be selfish and *non-malicious*: the user optimizes her own utility or payoff, and does not care about the performance of the system or the cost induced to other users by her strategy. We extend these models by considering *malicious users*. A malicious user will choose a strategy that will cause the worst possible performance for the entire network. Such malicious behaviour can be found in practice in settings such as the internet (for example in 'denial of service' attacks, or malicious flow in peer-to-peer networks). While in terms of Wardrop equilibria, the extension of the selfish model considered before is quite straightforward, the existence of malicious users forces us to a different model for the 'social cost'. We no longer have an objective function that can be minimized by the centralized coordination among the users, since in our setting some of the users still can be coordinated to *minimize* it, but at the same time there is a (malicious) user that tries to *maximize* it. This leads naturally to the formulation of the 'social cost' objective as a *minimax* problem instead of just a minimization

* Extended Abstract. Detailed proofs appear in the full version.

T. Ibaraki, N. Katoh, and H. Ono (Eds.): ISAAC 2003, LNCS 2906, pp. 696–704, 2003.

problem. As a result, we cannot refer to an ‘optimal social cost’ that is a global minimizer of the social cost objective. Instead, we have to compare the worst Wardrop equilibrium to the *saddle-points* of the minimax problem. We define the ‘optimal social cost’ as the minimum cost achieved by the set of saddle-points. The fact that this set is (usually) *non-convex* makes the exact characterization of the ‘optimal social cost’ (and therefore the coordination ratio) more difficult to characterize than the previous models. Nevertheless, in this paper we show that in the very general setting considered by Roughgarden and Tardos [11], their results can be extended to the case of systems with malicious users.

Previous Work: Many of the Game Theoretic tools used for analyzing systems of non-cooperative users derive from results in traffic models and transportation, including work of Dafermos and Sparrow [4], Beckmann, McGuire and Winsten [2] and Aashtiani and Magnanti [1]. More recently, Nash equilibria and their applications were used for routing problems and the internet. Koutsoupias and Papadimitriou [5] considered the coordination ratio for load balancing problems (routing on a network of parallel links). The model they considered allowed multiple equilibria, and the coordination ratio compared the worst case equilibrium cost to the optimal routing cost. Their bounds were improved in subsequent work on the same model by Mavronicolas and Spirakis [6], and Czumaj and Vöcking [3]. Roughgarden and Tardos [11] considered a different model for selfish routing, where there is a unique Wardrop equilibrium and proved bounds for the coordination ratio, including results for the special case of linear utility functions. Other work in this model includes results on the topology of the underlying network [8,10], and algorithms and bounds for Stackelberg scheduling strategies [9].

2 The Model

We are given a directed network $G = (V, E)$ and k source-sink pairs of nodes $(s_i, t_i), i = 1 \ldots k$. There are also two special nodes s_M, t_M connected to G with edges $(s_M, s_i), (t_i, t_M), i = 1 \ldots k$. A commodity i with demand r_i is associated with each pair $(s_i, t_i), i = 1 \ldots k$, and a commodity M of demand F is associated with pair (s_M, t_M). Let $\mathcal{P}_i$ ($\mathcal{P}_M$) be the set of acyclic paths from s_i to t_i (s_M to t_M). A latency function $l_P(\cdot)$ is associated with 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. We use the shorthand (G, r, F, l) to describe an instance of the model.

Commodities $i = 1 \ldots k$ model selfish, but otherwise ‘good’ users who want to just use the network in order to satisfy their demands with the smallest possible cost (i.e. latency for every unit of flow routed). Commodity M models a selfish 'malicious' user who wants to use his own flow F in such a way that will do the biggest possible damage to the total cost of the good players.

For our equilibrium model, we use the following general formulation by Aashtiani and Magnanti [1]:

Definition 1 *A flow* $f = \cup_{P\in\mathcal{P}} f_P$ *is at* Wardrop equilibrium *for instance* (G, r, F, l) *iff it satisfies the following constraints:*

$$\begin{aligned}
(T_P(f) - u_i)f_P &= 0 \quad \text{for all } P \in \mathcal{P}_i, i = 1 \dots k \\
(T_P(f) - u_M)f_P &= 0 \quad \text{for all } P \in \mathcal{P}_M \\
T_P(f) - u_i &\geq 0 \quad \text{for all } P \in \mathcal{P}_i, i = 1 \dots k \\
T_P(f) - u_M &\geq 0 \quad \text{for all } P \in \mathcal{P}_M \\
\sum_{P\in\mathcal{P}_i} f_P - r_i &= 0 \quad \text{for all } i = 1 \dots k \\
\sum_{P\in\mathcal{P}_M} f_P - F &= 0 \\
f \geq 0, \quad u &\geq 0
\end{aligned} \tag{1}$$

where T_P is the delay time or general disutility for path P, f_P is the flow through path P, and $u = (u_1, \dots, u_k, u_M)$ is the vector of shortest travel times (or generalized costs) for the commodities.

T_P does not need to be the same function for all paths P (it will be a different function for the good and the malicious users). Also we emphasize that T_P is *not* the path latency (the latter is given by function l_P). In what follows we define precisely the functions T_P for all users, and thus we define completely the equilibrium model of Definition 1.

The first four equations are the conditions for the existence of a Wardrop traffic equilibrium. They require that the general disutility for all paths P that carry flow $f_P > 0$ is the same and equal to u for every user, and less or equal to the disutility of any path with zero flow. Any flow that complies with this definition of a Wardrop equilibrium, also satisfies the following alternative characterization:

Lemma 1. *A flow that is feasible for instance (G, r, F, l) is a Wardrop equilibrium iff for every commodity i (i can be the malicious commodity M) and every pair of paths $P_1, P_2 \in \mathcal{P}_i$ with $f_{P_1} > 0$, $T_{P_1}(f) \leq T_{P_2}(f)$.*

2.1 Existence of Wardrop Equilibrium

The model of Definition 1 is very general. It turns out that the existence of a Wardrop equilibrium in this model can also be proved under very general assumptions. More specifically, the following theorem follows immediately from Theorem 5.4 in [1]:

Theorem 1. *Suppose that T_P is a positive continuous function for all $P \in \mathcal{P}$. Then there is a flow that satisfies the conditions of Definition 1.*

A function is positive if its values are positive. In order to make sure that a Wardrop always exists, from now on we make the following assumption:

Assumption 1 *The disutility function for every path is a* positive *function of the total flow, and that the disutility functions for the* good *users are increasing functions of the flow, i.e. as the congestion increases for a good user's path, its disutility also increases.*

3 Social Cost When Malicious Users Are Present

The existence of a malicious user forces us to redefine the notion of 'social cost' [5]. In addition to a set of users that collectively strive to *minimize* their collective cost (the 'social cost', as defined earlier [5], [11]), we have a user who strives to *maximize* this same cost. Therefore we define the 'socially best' flow in terms of a *minimax* problem. Note that in such a setting the notion of an "optimal flow" is replaced by the notion of a flow "in equilibrium". Therefore our work compares a Wardrop equilibrium to a minimax equilibrium (as opposed to the comparison of a Wardrop equilibrium to an optimal solution of a minimization problem, as in [11]).

In what follows, we denote the flow of the good users by f^G, and the flow of the malicious user by f^M (recall that we denote by f the total flow). We consider the following minimax formulation:

$$\begin{aligned}
&\max_{f^M} \min_{f^G} \sum_{e \in E} c_e(f_e^M, f_e^G) \quad \text{subject to:} && \text{(MINMAX)}\\
&\sum_{P \in \mathcal{P}_i} f_P^G = r_i && \forall i \in \{1, \dots, k\}\\
&\sum_{P \in \mathcal{P}_M} f_P^M = F &&\\
&f_e^G = \sum_{P \in \mathcal{P}: e \in P} f_P^G && \forall e \in E\\
&f_e^M = \sum_{P \in \mathcal{P}: e \in P} f_P^M && \forall e \in E\\
&f_P^G \geq 0 && \forall P \in \mathcal{P}\\
&f_P^M \geq 0 && \forall P \in \mathcal{P}
\end{aligned}$$

where $c_e(f_e^M, f_e^G)$ is the cost of flow (f_e^M, f_e^G) passing through edge e. In our case we have

$$c_e(f_e^M, f_e^G) = f_e^G \cdot l_e(f_e^G, f_e^M)$$

We call this minimax formulation (MINMAX), and its objective function $C(f^M, f^G) = \sum_{e \in E} c_e(f_e^M, f_e^G)$. The solution(s) to (MINMAX) are called *saddle-points*, defined as follows:

Definition 2 *A flow* $(\bar{f}^G, \bar{f}^M)$ *is said to be a* saddle-point *of* C *(with respect to maximizing in* f^M *and minimizing in* f^G*) if*

$$C(\bar{f}^G, f^M) \le C(\bar{f}^G, \bar{f}^M) \le C(f^G, \bar{f}^M), \quad \forall f^M,\ \forall f^G. \tag{2}$$

We also refer to (MINMAX) saddle-points as *(MINMAX) equilibria.*

3.1 Existence of Saddle-Points

A saddle-point is not always guaranteed to exist. But under certain assumptions, we can show that (at least one) saddle-point exists. We assume the following for the cost function $C(f^M, f^G)$:

Assumption 2 *The functions* $c_e(f_e^M, f_e^G)$ *are* continuous, differentiable, convex *with respect to* f^G*, and* concave *with respect to* f^M *for all* $e \in E$.

Following the methods of Dafermos and Sparrow [4], and under Assumption 2, we can prove the following theorem for the existence and properties of saddle-points for (MINMAX).

Theorem 2. *Under Assumption 2, a feasible flow* $\bar{f} = (\bar{f}^M, \bar{f}^G)$ *is a solution (saddle-point) to the minimax problem (MINMAX) if and only if it has the following properties:*

$$\sum_{e\in P} \frac{\partial c_e}{\partial f_e^G}(\bar{f}) = \sum_{e\in P'} \frac{\partial c_e}{\partial f_e^G}(\bar{f}) = A_i, \quad \forall P, P' \in \mathcal{P}_i,\ \ \bar{f}_P^G, \bar{f}_{P'}^G > 0 \tag{3}$$

$$\sum_{e\in P} \frac{\partial c_e}{\partial f_e^M}(\bar{f}) = \sum_{e\in P'} \frac{\partial c_e}{\partial f_e^M}(\bar{f}) = B, \quad \forall P, P' \in \mathcal{P}_M,\ \ \bar{f}_P^M, \bar{f}_{P'}^M > 0 \tag{4}$$

The conditions of Theorem 2 are simply the Kuhn-Tucker conditions for problem (MINMAX) [7].

4 Wardrop vs. Minimax Equilibria

We define natural *selfish* behaviors for both the good and malicious users, in accordance with the general model of Definition 1. Our aim will be to estimate how far can selfishness push the total cost from the optimal coordinated one (i.e. the best saddle-point of (MINMAX)). In order to do this, we modify the definition of the *price of anarchy* or *coordination ratio*, defined by Koutsoupias and Papadimitriou [5] and used by Roughgarden and Tardos [11].

Definition 3 (Coordination ratio) *Let* (G, r, F, l) *be an instance of the routing problem on network* G *with latency function* $l_e(\cdot)$ *for every edge* e*, with* k *good users with demands* r_i*,* $i = 1, \ldots, k$ *and a malicious user with flow* F*. Then the* coordination ratio $\rho(G, r, F, l)$ *for this instance is defined as follows:*

$$\rho(G, r, F, l) = \frac{\textit{worst Wardrop equilibrium}}{\textit{best saddle-point of (MINMAX)}}. \tag{5}$$

According to the model of Definition 1, the selfish users will base their decisions for picking flow paths on their individual notion of general disutility T_P, for every path P. This disutility is very easy to be defined for the 'good' users: it is simply the latency of the path, i.e.

$$T_P(f^G, f^M) := l_P(f^G, f^M) \ (= \sum_{e \in P} l_e(f_e)), \quad \forall i = 1, \dots, k, \ \forall P \in \mathcal{P}_i \tag{6}$$

For the malicious user though, the form of his general disutility in fact determines how powerful or weak this user can be. In this paper we study malicious players that base their decisions exclusively on the costs of *individual paths*. The malicious player exhibits a rather greedy behavior, and does not (or cannot[1]) take into account the impact of his decisions on the whole network (e.g. by solving (MINMAX) so that his allocation of flow will have the worst impact on the 'social cost' he might be able to achieve more damage than looking greedily at the costs of individual paths). Let $M(f^G) = \sum_{e \in E} f_e^G \cdot \frac{\partial l_e}{\partial f_e^M}(f_e^G, 0)$. Then the general disutility for the malicious user paths is defined as follows:

$$T_P(f^G, f^M) := M(f^G) - \sum_{e \in P} f_e^G \cdot \frac{\partial l_e}{\partial f_e^M}(f_e^G, f_e^M), \quad \forall P \in \mathcal{P}_M \tag{7}$$

In other words, the malicious player always tries to send his flow through a path with the biggest possible congestion increase for every unit of flow he allocates to this path, i.e. the malicious player follows a "best value for your money" policy.

The quantity $M(\cdot)$ is introduced so that Assumption 1 holds and therefore Wardrop equilibria exist.

4.1 Bicriteria Bound

As in the case of [11] we can prove a "bicriteria" result that gives an upper bound for the ratio between the cost at Wardrop equilibrium and the cost of the saddle-point solution.

Theorem 3. *If $f = (f^G, f^M)$ is a flow at Wardrop Equilibrium for (G, r, F, l) and $\hat{f} = (\hat{f}^G, \hat{f}^M)$ is a saddle-point of (MINMAX) for $(G, 2r, F, l)$ then $C(f) \leq C(\hat{f})$.*

Proof. The (social) cost of flow f is defined as

$$C(f) = \sum_e f_e^G \cdot l_e(f_e^G + f_e^M).$$

If f is at Wardrop equilibrium, then the total latency along any flow path P for good user i from s_i to t_i, $i = 1 \dots, k$ is the same, denoted by $L_i(f)$, and the total cost can be expressed as $C(f) = \sum_i L_i(f) r_i$. Define a new latency function $\bar{l}_e(x, y)$ as follows:

[1] maybe because of lack of resources, e.g. time in an on-line scenario

$$\bar{l}_e(x,y) = \begin{cases} l_e(x,y) & \text{if } x > f_e^G \text{ and } y > f_e^M \\ l_e(x,f_e^M) & x > f_e^G \text{ and } y \le f_e^M \\ l_e(f_e^G,y) & x \le f_e^G \text{ and } y > f_e^M \\ l_e(f_e^G,f_e^M) & x \le f_e^G \text{ and } y \le f_e^M \end{cases} \tag{8}$$

Note that the difference $\bar{l}_e(x,f_e^M) - l_e(x,f_e^M)$ is zero for $x \ge f_e^G$. Therefore the following is true for all $x \ge 0$:

$$x(\bar{l}_e(x,f_e^M) - l_e(x,f_e^M)) \le l_e(f_e^G,f_e^M)f_e^G. \tag{9}$$

The new latency functions give a new cost (cost with respect to $\bar{l}$) that is not too far from the real cost:

$$\begin{aligned} &\sum_e \bar{l}_e(\hat{f}_e^G,f_e^M)\hat{f}_e^G - C(\hat{f}^G,\hat{f}^M) \le \sum_e \bar{l}_e(\hat{f}_e^G,f_e^M)\hat{f}_e^G - C(\hat{f}^G,f^M) = \\ &\sum_e \hat{f}_e^G(\bar{l}_e(\hat{f}_e^G,f_e^M) - l_e(\hat{f}_e^G,f_e^M)) \le \sum_e f_e^G l_e(f_e^G,f_e^M) = C(f) \end{aligned} \tag{10}$$

The first inequality is due to the fact that $\hat{f} = (\hat{f}^G,\hat{f}^M)$ is a saddle-point for $(G,2r,F,l)$, i.e. $C(\hat{f}^G,f^M) \le C(\hat{f}^G,\hat{f}^M)$ since $(\hat{f}^G,f^M)$ is a feasible solution for (MINMAX). The second inequality comes from (9) for $x := \hat{f}_e^G$.

Consider any path $P \in \mathcal{P}_i$. From the definition of $\bar{l}_e$ we have that

$$\sum_{e\in P} \bar{l}_e(0,f_e^M) \ge \sum_{e\in P} l_e(f_e^G,f_e^M) = L_i(f).$$

and from the fact that $\bar{l}_e(x,f_e^M)$ is an increasing function of x we get

$$\sum_{e\in P} \bar{l}_e(\hat{f}_e^G,f_e^M) \ge \sum_{e\in P} \bar{l}_e(0,f_e^M).$$

Therefore:

$$\begin{aligned} &\sum_{e\in E}\bar{l}_e(\hat{f}_e^G,f_e^M)\cdot \hat{f}_e^G \ge \sum_i \sum_{P\in\mathcal{P}_i} \hat{f}_P^G \sum_{e\in P} \bar{l}_e(\hat{f}_e^G,f_e^M) \ge \\ &\sum_i \sum_{P\in\mathcal{P}_i} L_i(f)\hat{f}_P^G = \sum_i 2L_i(f)r_i = 2C(f) \end{aligned} \tag{11}$$

By combining (10) with (11) we get $C(f) \le C(\hat{f})$.

The same proof also gives the following result:

Theorem 4. *If $f = (f^G,f^M)$ is a flow at Wardrop Equilibrium for (G,r,F,l) and $\hat{f} = (\hat{f}^G,\hat{f}^M)$ is a saddle-point of (MINMAX) for $(G,(1+\gamma)r,F,l)$, $\gamma > 0$ then $C(f) \le \frac{1}{\gamma}C(\hat{f})$.*

At a first glance, it seems rather surprising that the bicriteria bounds of [11] are quite robust against the existence of a malicious user. But if we look closer to the quantities compared in the theorems above, we see that while the demands of the good users are increased, the flow quantity at the disposal of the malicious user remained the same. Intuitively, the malicious user has the same power to disrupt the good users in both cases, and therefore if he settles with some strategy to do so for the initial good demands, this strategy should work about as well when the latter demands increase. The same goes for the good users' strategies as well.

4.2 Special Case: Linear Latency Functions

In this section we deal with the special case of linear edge latency functions, i.e. for every edge $e \in E$, $l_e(f_e^G, f_e^M) = a_e(f_e^G + f_e^M) + b_e$ for some $a_e \geq 0, b_e > 0$. Note that we assume that the latency for an edge is positive even if no flow passes through it. This is a quite natural assumption (in all physical systems there is always some delay in moving from point A to point B, even if there is no congestion at all), and allows Theorem 1 to apply in this case. We modify our shorthand notation to (G, r, F, a, b) to include the linear coefficient vectors. In this special case we have

- $T_P(f^G, f^M) := \sum_{e \in P}(a_e f_e^G + a_e f_e^M + b_e),\ \forall i = 1, \ldots, k,\ \forall P \in \mathcal{P}_i$
- $T_P(f^G, f^M) := \sum_{e \in E} a_e f_e^G - \sum_{e \in P} a_e f_e^G,\ \forall P \in \mathcal{P}_M$

Lemma 1 and Theorem 2 take a more specific form for the linear case:

Lemma 2. *Let $l_e(f_e^G, f_e^M) = a_e(f_e^G + f_e^M) + b_e$ with $a_e \geq 0, b_e > 0$ be the latency function for every edge $e \in E$ of G.*

(a) a flow $f = (f^G, f^M)$ is at Wardrop equilibrium iff
- *for all users $i = 1, \ldots, k$ and paths $P, P' \in \mathcal{P}_i$ with $f_P > 0$*

$$\sum_{e \in P} \left(a_e f_e^G + a_e f_e^M + b_e\right) \leq \sum_{e \in P'} \left(a_e f_e^G + a_e f_e^M + b_e\right)$$

- *for all paths $P, P' \in \mathcal{P}_M$ with $f_P > 0$:* $\displaystyle\sum_{e \in P} a_e f_e^G \geq \sum_{e \in P'} a_e f_e^G$

(b) a flow $\bar{f} = (\bar{f}^G, \bar{f}^M)$ is an equilibrium for (MINMAX) *iff*
- *for all commodities $i = 1, \ldots, k$ and paths $P, P' \in \mathcal{P}_i$ with $\bar{f}_P > 0$*

$$\sum_{e \in P} \left(2a_e \bar{f}_e^G + a_e \bar{f}_e^M + b_e\right) \leq \sum_{e \in P'} \left(2a_e \bar{f}_e^G + a_e \bar{f}_e^M + b_e\right)$$

- *for all paths $P, P' \in \mathcal{P}_M$ with $\bar{f}_P > 0$:* $\displaystyle\sum_{e \in P} a_e \bar{f}_e^G \geq \sum_{e \in P'} a_e \bar{f}_e^G$

For this special form of the edge latency functions, we can prove that the saddle-point cost for (MINMAX) is unique (proof omitted). In a way similar to [11] we can prove our main theorem for the coordination ratio in the linear case:

Theorem 5. *For instance* (G, r, F, a, b), $1 \leq \rho(G, r, F, a, b) \leq \frac{4}{3}$.

Note that the lower bound for the coordination ratio is tight, since $\rho(G, r, F, a, b) = 1$ if G is just a path with the sources for all users in one end, and all the sinks in the other.

5 Open Problems

The model presented in our work gives rise to many open problems. It would be very interesting to present a result connecting the social cost of an equilibrium point in a network with malicious users and the cost in an equivalent instance without malicious users. This would give a clear characterization of the negative impact of the presence of malicious flow. For the general latency functions, it seems that it is possible to prove more tight results and extend the bicriteria result by proving a lower bound. The model defined in our work gives rise to unique saddle-points and Wardrop equilibria. It would be interesting to consider a more general model that allows multiple equilibria (for example, by adding capacities for the edges in the network [12]) and analyze the performance of the system in the presence of malicious users.

Acknowledgments. We would like to thank Tamás Terlaky and Nicola Galesi for many helpful discussions.

References

1. H. Z. Aashtiani and T. L. Magnanti. Equilibria on a congested transportation network. *SIAM J. on Algebraic and Discrete Methods*, 2(3):213–226, Sept. 1981.
2. M. Beckmann, C. B. McGuire, and C. B. Winsten. Studies in the economics of transportation. *Yale University Press*, 1956.
3. A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria. In *13th ACM-SIAM Symposium On Discrete Mathematics*, pages 413–420, 2002.
4. S. Dafermos and F. Sparrow. The traffic assignment problem for a general network. *J. of Research of the National Bureau of Standards*, 73B:91–118, 1969.
5. E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *16th Symposium on Theoretical Aspects of Computer Science, LNCS 1563*, pages 404–413, 1999.
6. M. Mavronicolas and P. Spirakis. The price of selfish routing. In *33rd Annual ACM Symposium on Theory of Computing*, pages 510–519, 2001.
7. R. T. Rockafellar. *Convex Analysis.* Princeton University Press, 1970.
8. T. Roughgarden. Designing networks for selfish users is hard. In *42nd IEEE Symposium on Foundations of Computer Science*, pages 472–481, 2001.
9. T. Roughgarden. Stackelberg scheduling strategies. In *33rd Annual ACM Symposium on Theory of Computing*, pages 104–113, 2001.
10. T. Roughgarden. The price of anarchy is independent of the network topology. In *34th Annual ACM Symposium on Theory of Computing*, pages 428–437, 2002.
11. T. Roughgarden and É. Tardos. How bad is selfish routing? *J. of the ACM*, 49(2):236–259, Mar. 2002.
12. A. S. Schulz and N. S. Moses. On the performance of user equilibria in traffic networks. In *14th Symposium on Discrete Algorithms*, pages 86–87, 2003.