The price of anarchy: The case of an exponential multi-server

Moshe Haviv[*]   Tim Roughgarden[†]

Abstract

We consider a single multi-server memoryless service station. Servers have heterogeneous service rates. Arrivals are routed to one of the servers. The routing decisions are not based on the queue lengths. There are two criteria for routing selection. The first is the (Nash) equilibrium, under which each customer minimizes his own mean waiting time, given the behavior of the others. The second is social optimization, where the routing is designed to minimize the average mean waiting time across all arrivals. The ratio between the social costs in both cases in called the price of anarchy (PoA). We show that the PoA is bounded from above by the number of servers utilized in the socially optimal outcome. In particular, it is bounded by the number of available servers. We also show that this bound is tight.

1 Introduction

Consider a situation in which many individuals (players) participate, and each incurs some cost. The social cost is defined as the sum (or average)

[*]Department of Statistics, The Hebrew University of Jerusalem, 91905 Jerusalem, Israel. The research leading to this paper was conducted in the School of Mathematics, Statistics and Computer Science at Victoria University of Wellington, New Zealand and at the School of Economics and Political Science at the University of Sydney, Sydney, NSW, Australia. Supported in part by The Israel Science Foundation Grant no. 237/02. E-mail: haviv@mscc.huji.ac.il.

[†]Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305. Supported in part by ONR grant N00014-04-1-0725, DARPA grant W911NF-05-1-0224, and an NSF CAREER Award. Email: tim@cs.stanford.edu.

over the players' individual costs. Each of the players selects one action from his set of possible actions. The individual's cost is a function of the actions taken by all of the players (and not only his own).

In the search for a strategy profile, two criteria are usually considered in the literature. The first criterion is a *Nash equilibrium.* A strategy profile is a Nash equilibrium if each individual chooses an action (or a distribution over the actions) which minimizes his cost, given the actions chosen by the other players. Nash equilibria are not necessarily unique. The second criterion is that of *social optimization.* Here a central planner assigns an action (or a distribution over the set of actions) to each of the players to minimize the social cost. It is possible that the socially optimal profile is not unique but of course the corresponding optimal value is.

The *price of anarchy (PoA)* [8], or the *coordination ratio* [7], is defined as the ratio between the social cost of a Nash equilibrium and the social cost of a socially optimal profile. When different equilibria can have different social costs, the price of anarchy is usually defined in terms of the equilibrium with the largest cost [7]. In this paper, however, we will only consider games in which all equilibria have equal social cost.

The PoA is problem dependent. Moreover, a universal bound on it does not exist. To see this, consider the two-person game known as the *Prisoner's Dilemma* [9], described by the following cost matrix:

$$\begin{pmatrix} (x,x) & (0,x+1) \\ (x+1,0) & (1,1) \end{pmatrix}.$$

Assume that $x > 1$, that both players want to minimize their cost, and that the social cost is the sum of the individual costs. In the unique Nash equilibrium, both players choose their first strategy and the social cost is $2x$. In the optimal solution, both players choose their second strategy and the social cost is 2. The PoA is therefore $2x/2 = x$, which of course is unbounded.

Since the PoA is in general unbounded, the literature has sought bounds on the PoA in restricted classes of games. The class of *congestion games* has been extensively studied. In a congestion game, the cost incurred by a player is a function only of the player's strategy and of the number of other players that choose the same or some "interfering" strategy. The PoA remains unbounded in such games unless further restrictions are imposed on the cost functions [12]. This fact has motivated the study of linear, polynomial, and queueing-theoretic cost functions (e.g. [2, 10]). This paper contributes tight bounds on the PoA with this last type of function.

This paper considers a single multi-server service station. Next we state the model in general terms; full details are given in the next section. We also state here our main results. We then conclude the introduction with a short literature review on bounding the PoA.

Consider a single service station with a number of (not necessarily identical) servers. Customers' costs are reflected by their expected waiting times. We assume that the time in service is considered as part of the waiting time. The common action set of the customers corresponds to server selection and no further information (such as the queue lengths upon arrival) is given. Selfish customers wish to minimize their own mean waiting time but ignore the social consequences of their actions. These consequences, or negative externalities, are in the form of added waiting costs on others, costs which vary with the line customers join. Thus, customers' selections may not be socially optimal. On the other end, a central planner wishes to direct customers to servers to minimize the average mean waiting time (which may not be the best choice for each individual customer). As these two selections need not coincide, a PoA is paid by society if individuals behave selfishly.

We assume a never-ending stream of arrivals which follows a Poisson process, and exponentially-distributed (server-dependent) service times. We assume that steady-state conditions have been reached. In particular, we assume that the arrival rate is smaller than the total service rate. Since a unique equilibrium exists in this model, there is no ambiguity regarding the definition of the PoA. Friedman [5] gives an example that shows that a universal bound on the PoA in a multi-server exponential service station does not exist. In this example, the number of servers is unbounded and indeed the PoA increases with the number of servers. Hence, our next hope for a finite bound on the PoA in this model lies in limiting the number of servers.

In this paper, we show that the PoA in this model is bounded from above by the number of servers used in the optimal outcome (and hence by the number of servers). We also give examples that demonstrate that this upper bound is the best possible.

We next describe some existing results on the PoA in related models. The literature on the PoA in congestion games is already quite large, and we therefore confine our attention to only the most closely related works. For further discussion of and many more references to this literature, see e.g. [11].

Although different terminology is used, Koutsoupias and Papadimitriou [7] study a multi-server system. Their model is different from ours, however, as they consider a finite number of players that each control a non-negligible fraction of the arriving traffic. They also define the social cost to be the maximum individual cost, whereas we consider the average player cost. For surveys of more recent work in this model, see [3, 4].

Our model can be viewed as a special case of the traffic routing model that is often referred to as "traffic equilibria" in the transportation science literature (see e.g. [13]) and "selfish routing" by the theoretical computer science community (see e.g. [11]). This general routing model studies arbitrary flow networks, with potentially multiple origins and destinations for traffic. This model also allows an arbitrary continuous, nondecreasing cost function on each edge of the network. The PoA in the general routing model is unbounded, even in networks with only one origin, one destination, and two parallel links connecting them [12]. The PoA is only 4/3 when all edge cost functions are affine [12], and is also bounded by a small constant for several other classes of restricted cost functions [2, 10].

Our model corresponds to networks with one traffic origin, one destination, and a number of parallel links (each representing a server) connecting the origin to the destination, each with a cost function of a restricted (queueing-theoretic) form. The only previously known positive results for a multi-server service station assumed that the station was lightly loaded in some sense [5, 10]. However, previous results on the PoA did not restrict the number of allowable edges (servers) and therefore give no finite bound on the PoA in our model. While Roughgarden [10] proved a general result stating, in essence, that the price of anarchy in two-node networks with two parallel links can be just as high as in networks with an arbitrary number of parallel links, this result assumes that edges are allowed to possess constant cost functions. For this reason, this result does not directly apply to a multi-server service station. Indeed, we will see that the number of parallel links (servers) in our model completely controls the worst-possible PoA.

The rest of the paper is organized as follows. In the next section we give full details on the model and describe some existing results that are needed for our derivations. Section 3 then contains our main findings.

2 The Model and Preliminaries

Suppose there are n exponential servers who man a single service station. Let μ_i be the service rate of server i, $1 \leq i \leq n$. Assume without loss of generality that $\mu_1 \geq \mu_2 \geq \cdots \geq \mu_n > 0$. For convenience assume that $\mu_{n+1} = 0$. Denote $\sum_{j=1}^{i} \mu_j$ by $\mu_{(i)}$. To this service station there is a single Poisson arrival process with rate λ. In order to guarantee stability we assume that $\lambda < \mu_{(n)}$. Customers are routed to one of the servers and wait in line there for their turn. The service regime in each line is first-come first-served. Jockeying between the lines is not allowed.

2.1 Equilibrium

Suppose customers wish to minimize their mean waiting time and select a line to join accordingly. They do not observe the actual queue lengths at the time of deciding which line to join. Customers are allowed to randomize regarding their choice of servers. Thus, each line corresponds to a pure strategy and every distribution over the lines is a mixed (randomized) strategy.

We adopt the solution concept of a symmetric Nash equilibrium under steady-state conditions. Specifically, suppose all customers decide to join line i with probability p_i, $1 \leq i \leq n$. Under the resulting steady-state conditions, the arrival rate to a server i is Poisson with rate λp_i and the mean waiting time is $1/(\mu_i - \lambda p_i)$ (assuming of course that $\mu_i > \lambda p_i$). The mixed strategy $\{p_i\}_{i=1}^{n}$ is a symmetric Nash equilibrium if and only if when all customers follows this strategy, no customer can decrease its mean waiting time by unilaterally deviating to a different strategy. An alternative interpretation of a symmetric equilibrium $\{p_i\}_{i=1}^{n}$, known as a *Wardrop equilibrium* [14] in the transportation literature, is that all customers pick pure strategies, with a p_i fraction of the arriving customers deterministically selecting the ith server.

In a symmetric Nash equilibrium, the mean waiting times in all utilized servers are identical and their common value is smaller than or equal to the mean service time in each of the idle servers (if such servers exist). Clearly, if $i < j$ and server j is utilized under equilibrium, then so is server i. It is shown in [1] (see also [6, P.63]) that at equilibrium, the only servers that are utilized are the first i_e servers, where

$$i_e = \min\left\{ i \geq 1 \,:\, \mu_{i+1} \leq \frac{\mu_{(i)} - \lambda}{i} \right\} . \tag{1}$$

Moreover, the Poisson arrival rate to server i, denoted here by λ_i^e, is

$$\lambda_i^e = \mu_i - \frac{\mu_{(i_e)} - \lambda}{i_e}, \quad 1 \leq i \leq i_e .$$

In particular, the equilibrium strategy is to join queue i with probability λ_i^e/λ, $1 \leq i \leq i_e$, and not to join queue i, $i_e + 1 \leq i \leq n$. Finally, the common mean waiting time in queues 1 through i_e is

$$W^e = \frac{i_e}{\mu_{(i_e)} - \lambda}$$

and the total social cost per unit of time, which is the mean number of customers in the system, under this behavior equals

$$L^e = \lambda \frac{i_e}{\mu_{(i_e)} - \lambda} \quad . \tag{2}$$

The social cost under equilibrium behavior is therefore a function only of the number of and the sum of the service rates of the utilized servers.

Equation (1) implies that i_e is non-decreasing in λ. Hence, we can ask the inverse question: For a given integer $m \leq n$ what is the infimum value for λ such that the first m servers are utilized at equilibrium. Denote this value by $\lambda_e^{\min}(m)$. Equation (1) implies that

$$\lambda_e^{\min}(m) = \mu_{(m-1)} - \mu_m(m-1), \quad 1 \leq m \leq n.^1 \tag{3}$$

As expected, $\lambda_e^{\min}(m)$ is nondecreasing in m. Note that $\lambda_e^{\min}(m+1)$ is also the supremum of the arrival rates for which precisely the first $m \leq n$ servers are utilized at equilibrium.

Remark 2.1 With one server with service rate $\mu_{(i_e)}$, the common mean waiting time would be $1/(\mu_{(i_e)} - \lambda)$. This is only a $1/i_e$ fraction of the mean waiting time in an equilibrium that utilizes i_e servers with total service rate $\mu_{(i_e)}$. Note that this increase in the mean waiting time is a function only of the number i_e of utilized servers, and not a function of how the $\mu_{(i_e)}$ total service rate is partitioned among these servers. Nonetheless, as we will see below, the PoA depends on the partition in the total capacity.

[1] Equation (1) gives

$$\lambda_e^{\min}(m) = \min\{\lambda \geq 0 \,:\, \mu_{i+1} \geq \frac{\mu_{(i)} - \lambda}{i}, 1 \leq i \leq m-1\}.$$

Since μ_i is nondecreasing in i, the above inequality holds for all $i \in \{1, 2, \ldots, m-1\}$ if and only if it holds for $i = m-1$. This implies (3).

Remark 2.2 Remark 2.1 is reminiscent of the following fact about queue splitting. Suppose instead of one server with an arrival rate of λ and a service rate of μ, there are n servers such that the arrival rate and service rate at each server $i \in \{1, 2, \ldots, n\}$ are, respectively, $p_i\lambda$ and $p_i\mu$ for some positive probabilities $\{p_i\}_{i=1}^n$ with $\sum_{i=1}^n p_i = 1$. Let $\rho = \lambda/\mu$. Then, for any choice of p_i's, all servers are utilized at the level of ρ and the mean queue length in *each* is $\rho/(1-\rho)$, as in the single queue. This implies that the total mean number in the system after the split is n times larger then the original mean and likewise, by Little, the mean waiting time is n times larger then the corresponding mean before the split. Note also that the social cost after such a split coincides with the social cost under equilibrium (2) if all n servers are utilized in the equilibrium outcome.

2.2 Social Optimization

Every symmetric mixed strategy leads to a set of arrival rates $\{\lambda_i\}_{i=1}^n$ to the servers. The search for a socially optimal symmetric strategy is therefore equivalent to solving the following mathematical program:

$$\min_{\lambda_1,\ldots,\lambda_n} \sum_{j=1}^n \frac{\lambda_j}{\mu_j - \lambda_j}$$

subject to

$$\sum_{j=1}^n \lambda_j = \lambda,$$

and

$$0 \leq \lambda_j < \mu_j, \quad j = 1, 2, \ldots, n.$$

It is shown in [1] that the solution to this problem is as follows. The first i_s servers are utilized, where

$$i_s = \min \left\{ i \geq 1 \,:\, \mu_{i+1} \leq \frac{(\mu_{(i)} - \lambda)^2}{(\sum_{j=1}^i \sqrt{\mu_j})^2} \right\} . \tag{4}$$

Moreover, the socially optimal arrival rates to these servers, denoted by λ_i^s, are

$$\lambda_i^s = \mu_i - \frac{\sqrt{\mu_i}}{\beta}, \quad 1 \leq i \leq i_s \tag{5}$$

where

$$\beta = \frac{\sum_{j=1}^{i_s} \sqrt{\mu_j}}{\mu_{(i_s)} - \lambda}.$$

The corresponding socially optimal policy is for each customer to join server i, $1 \le i \le i_s$, with probability λ_i^s/λ, and to join the other servers with probability zero. Finally, the mean waiting time in the utilized server i equals

$$W_i^s = \beta/\sqrt{\mu_i}, \quad 1 \le i \le i_s. \tag{6}$$

In particular, the mean waiting time is smaller in the fast servers who nevertheless receive more traffic than the slow servers. This should be compared with the equilibrium criterion in which mean waiting times in all utilized servers are identical. Note also that the utilization levels $\lambda_i^s/\mu_i = 1 - 1/\beta\sqrt{\mu_i}$, $1 \le i \le i_s$, are higher at the faster servers.

The social cost (mean number in the system) under the socially optimal behavior, denoted by L^s, is $L^s = \sum_{i=1}^{i_s} \lambda_i^s W_i^s$ where λ_i^s and W_i^s are defined as in (5) and (6), respectively. This, after minimal algebra, equals

$$L^s = \frac{(\sum_{i=1}^{i_s} \sqrt{\mu_i})^2}{\mu_{(i_s)} - \lambda} - i_s \ . \tag{7}$$

Note that the socially optimal cost is a function only of the number of utilized servers, and of the sum and the sum of the square roots of the service rates of these servers.

Equation (4) implies that i_s is non-decreasing in λ. Let $\lambda_s^{\min}(m)$ denote the infimum value of λ such that the first m servers are utilized in the socially optimal outcome. Reasoning as for equation (3), we obtain

$$\lambda_s^{\min}(m) = \mu_{(m-1)} - \sqrt{\mu_m} \sum_{j=1}^{m-1} \sqrt{\mu_j}, \quad 1 \le m \le n. \tag{8}$$

Note that $\lambda_s^{\min}(m)$ is non-decreasing in m and that $\lambda_s^{\min}(m+1)$ is the supremum of the arrival rates for which the socially optimal outcome utilizes precisely the first m servers. Comparing (3) and (8), we see that $\lambda_e^{\min}(m) \ge \lambda_s^{\min}(m)$ for all $m \in \{1, 2, \ldots, n\}$. It follows that $i_e \le i_s$.

Remark 2.3 Starting from socially optimal behavior as a point of departure, customers will migrate from queues with high mean waiting times to queues with smaller mean waiting times—from servers with high indices to those with lower indices. In particular, this may cause some servers that are utilized in the socially optimal outcome to become empty in the equilibrium. Thus, the PoA is caused by two sides of the same coin: (1) over-congestion at faster servers; and (2) servers which are slower but socially useful may stay idle with selfish behavior.

3 Main Results

Our main results are that the price of anarchy in every multi-server exponential service station is bounded above by the number of servers, and that no better upper bound is possible. The following lemma, which considers the special case of service stations in which the socially optimal and equilibrium outcomes utilize the same set of servers, is crucial for our upper bound proof.

Lemma 3.1 *In a multi-server exponential service station in which both the socially optimal and the equilibrium outcomes utilize the same number m of servers, the price of anarchy is at most m.*

Proof: By definition, the PoA is L^e/L^s, where L^e and L^s are defined as in (2) and (7), respectively. Assume that $i_e = i_s = m$—i.e., that the socially optimal and equilibrium outcomes both utilize precisely the first m servers. We can assume that $m > 1$, as otherwise the PoA is clearly 1. The PoA is then

$$\begin{aligned} \frac{L^e}{L^s} &= \lambda i_e \frac{\mu_{(i_s)} - \lambda}{\mu_{(i_e)} - \lambda} \left[\left(\sum_{i=1}^{i_s} \sqrt{\mu_i} \right)^2 - i_s(\mu_{(i_s)} - \lambda) \right]^{-1} \\ &= \frac{\lambda m}{(\sum_{i=1}^m \sqrt{\mu_i})^2 - m(\mu_{(m)} - \lambda)}. \end{aligned} \tag{9}$$

First, note that the socially optimal and equilibrium outcomes, along with the ratio of their social costs (2) and (7), remain unchanged if all service and arrival rates are scaled by a common positive value. We can therefore assume, without loss of generality, that $\mu_{(m)} = 1$. Second, the expression (9) is strictly decreasing in λ, and $\lambda \geq \lambda_e^{\min}(m)$ by definition. Combining our expression (3) for $\lambda_e^{\min}(m)$ with (9) then gives

$$\frac{L^e}{L^s} \leq \frac{m(1 - m\mu_m)}{(\sum_{i=1}^m \sqrt{\mu_i})^2 - m^2\mu_m}. \tag{10}$$

Our next goal is to maximize (10) where the decision variables μ_i, $1 \le i \le m$, are constrained so that $\mu_1 \ge \mu_2 \ge \cdots \ge \mu_m \ge 0$ and $\sum_{i=1}^m \mu_i = 1$. The optimal solution is of the form $\mu_i = \mu_m$ for $2 \le i \le m$ and $\mu_1 = 1-(m-1)\mu_m$.[2] Maximizing the right-hand side of (10) thus reduces to maximizing the following single variable objective:

$$\frac{m(1-m\mu_m)}{1-(3m-2)\mu_m + 2(m-1)\sqrt{\mu_m(1-(m-1)\mu_m)}}.$$

Finally, trivial algebra shows that

$$1 - mx \le 1-(3m-2)x + 2(m-1)\sqrt{x(1-(m-1)x)}$$

for every $x \in [0, 1/m]$, which proves that the PoA (9) is at most m. ■

We next show that the bound m derived in Lemma 3.1 is applicable also when more servers are utilized in the socially optimal outcome than in the equilibrium one. Thus, in a multi-server service station with n servers, the price of anarchy is at most n.

Theorem 3.1 *In every multi-server exponential service station, the price of anarchy is at most the number of servers utilized in the socially optimal outcome.*

Proof: Recall from Subsection 2.2 that the number of servers i_s utilized in the socially optimal outcome can only be greater than the number i_e utilized by the equilibrium outcome. Suppose that $i_e < i_s = m$. Define augmented service rates μ'_j as follows: $\mu'_j = \mu_j$ for $j \notin \{i_e+1, \ldots, m\}$, and

$$\mu'_j = \frac{\mu_{(i_e)} - \lambda}{i_e}, \quad i_e + 1 \le j \le m.$$

The definition of these augmented service rates ensures that the equilibrium outcome is unaffected: the rates of the unutilized servers are increased to the largest value at which customers still do not switch to a server with an

[2]To see this, fix $\mu_m \le 1/m$ and maximize (10) with respect to the other variables $\mu_1, \mu_2, \ldots, \mu_{m-1}$. This optimization problem is equivalent to minimizing $\sum_{i=1}^{m-1} \sqrt{\mu_i}$ subject to $\mu_1 \ge \mu_2 \ge \cdots \ge \mu_m$ and $\sum_{i=1}^{m-1} \mu_i = 1 - \mu_m$. Since this is minimizing a strictly concave function over a convex set, the optimal solution is on the boundary, which in this case implies that $\mu_2 = \mu_3 = \cdots = \mu_{m-1} = \mu_m$ and $\mu_1 = 1-(m-1)\mu_m$.

index $i_e + 1$ or higher. Thus, L_e (the social cost under equilibrium) is not changed but L_s (the social cost under optimal behavior) has been reduced due to the improvement in the service capacities. In particular, these capacity augmentations increase the PoA. We can therefore prove the theorem by showing that the PoA is at most m in the augmented service station.

Now consider slightly increasing the arrival rate λ to the augmented service station. At equilibrium, the added flow of jobs is evenly spread among servers $1, \ldots, i_s$, making the number of utilized servers equal to the number m of such servers in the optimal outcome. Of course, this infinitesimal change in λ has only infinitesimal effects on both L_e and L_s—formally, L_e and L_s are both continuous functions of λ. We can therefore conclude that the PoA of the augmented multi-server service station is bounded above by the PoA of a service station in which the socially optimal and equilibrium outcomes both utilize the same m servers. Lemma 3.1 then implies that the PoA of the augmented service station is at most m, which completes the proof of the theorem. ■

We conclude with an example that demonstrates that our upper bound of n in Theorem 3.1 is the best possible. In other words, we show that for every $n \geq 1$, there are multi-server exponential service stations with n servers and PoA arbitrarily close to n.

Example 3.1 Fix $n \geq 1$, an arrival rate $\lambda > n$, and a constant $\epsilon > 0$, let $\mu_1 = \lambda + 1 + \epsilon$ and $\mu_i = 1 + \epsilon$ for $i \in \{2, \ldots, n\}$. In the equilibrium outcome, all customers choose the first server, and the mean service time of all servers equals $1/(1+\epsilon)$. In particular,

$$L^e = \frac{\lambda}{1+\epsilon} \ . \tag{11}$$

Now consider assigning customers to servers so that $\lambda_1 = \lambda - n + 1$ and $\lambda_i = 1$ for every $i \in \{2, \ldots, n\}$. The mean number in the n queues in this outcome upper bounds L^s from above and equals

$$\frac{\lambda - n + 1}{n + \epsilon} + (n-1)\frac{1}{\epsilon} \ . \tag{12}$$

Now divide (11) by (12) and take λ to infinity. For each fixed $n \geq 1$ and $\epsilon > 0$, the limit of this ratio is $(n+\epsilon)/(1+\epsilon)$. This expression can be made arbitrarily close to n by choosing the constant $\epsilon > 0$ to be sufficiently small.

Hence, there are n-server service stations with PoA arbitrarily close to n, and the upper bound in Theorem 3.1 is optimal.

Acknowledgment

Thanks are due to Yoav Kerner for some helpful comments.

References

[1] C. H. Bell and S. Stidham. Individual versus optimization in allocation of customers to alternative servers. *Management Science*, 29:831–839, 1983.

[2] J. R. Correa, A. S. Schulz, and N. E. Stier Moses. Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29(4):961–976, 2004.

[3] A. Czumaj. Selfish routing on the Internet. In J. Leung, editor, *Handbook of Scheduling: Algorithms, Models, and Performance Analysis*, chapter 42. CRC Press, 2004.

[4] R. Feldmann, M. Gairing, T. Lücking, B. Monien, and M. Rode. Selfish routing in non-cooperative networks: A survey. In *Proceedings of the 28th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 2747 of *Lecture Notes in Computer Science*, pages 21–45, 2003.

[5] E. J. Friedman. Genericity and congestion control in selfish routing. In *Proceedings of the 43rd Annual IEEE Conference on Decision and Control (CDC)*, pages 4667–4672, 2004.

[6] R. Hassin and M. Haviv. *To Queue or not to Queue: Equilibrium Behavior in Queueing Systems*. Kluwer's International Series, Boston, MA, 2003.

[7] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 1563 of *Lecture Notes in Computer Science*, pages 404–413, 1999.

[8] C. H. Papadimitriou. Algorithms, games, and the Internet. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 749–753, 2001.

[9] A. Rapoport and A. M. Chammah. *Prisoner's Dilemma*. University of Michigan Press, 1965.

[10] T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.

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

[12] T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.

[13] Y. Sheffi. *Urban Transportation Networks: Equilibrium Analysis with Mathematical Programming Methods*. Prentice-Hall, 1985.

[14] J. G. Wardrop. Some theoretical aspects of road traffic research. In *Proceedings of the Institute of Civil Engineers, Pt. II*, volume 1, pages 325–378, 1952.