Worst-case Equilibria

Elias Koutsoupias* Christos Papadimitriou†

April 29, 2009

Abstract

In a system where noncooperative agents share a common resource, we propose the price of anarchy, which is the ratio between the worst possible Nash equilibrium and the social optimum, as a measure of the effectiveness of the system. Deriving upper and lower bounds for this ratio in a model where several agents share a very simple network leads to some interesting mathematics, results, and open problems[1].

1 Introduction

Internet users and service providers act selfishly and spontaneously, without an authority that monitors and regulates network operation in order to achieve some "social optimum" such as minimum total delay [1]. *How much performance is lost because of this?* This question appears to exemplify a novel and timely genre of algorithmic problems, in which we are investigating the cost of the lack of *coordination* —as opposed to the lack of *information* (on-line algorithms) or the lack of *unbounded computational resources* (approximation algorithms). As we show in this paper, this point of view leads to some interesting algorithmic and combinatorial questions and results.

It is nontrivial to arrive at a compelling mathematical formulation of this question. Independent, noncooperative agents obviously evoke *game theory* [15], and its main concept of rational behavior, the *Nash equilibrium:* In an environment where each agent (or player) is aware of the situation facing all other agents, a Nash equilibrium is a combination of choices (deterministic *or* randomized), one for each agent, from which no agent has an incentive to unilaterally move away. Nash equilibria are known not to always optimize overall performance, with the Prisoner's Dilemma [15, 17] being the best-known example. Conditions under which Nash equilibria can achieve or approximate the overall optimum have been studied extensively ([17]; see also [6, 10, 19] for studies on networks). However, this line of previous work compares the overall optimum with the *best* Nash equilibrium, not the *worst*, as befits our line of reasoning. To put it otherwise, this previous research aims at achieving or approximating the social optimum by implicit acts of coordination, whereas we are interested in evaluating the loss to the system due to its deliberate lack of coordination.

Game-theoretic aspects of the Internet have also been considered by researchers associated with the Internet Society [1, 20], with an eye towards designing variants of the Internet Protocols

*University of Athens, elias@di.uoa.gr. This work was done while the first author was affiliated with the University of California, Los Angeles.

†University of California, Berkeley, christos@cs.berkeley.edu

[1]The conference version of this work [9] appeared a decade ago and it has been followed by a large amount of work on the concept of the price of anarchy (as witnessed by the extensive coverage in [14]). In this journal version we tried to keep as much as possible the text of the original paper. There are, though, important changes because results that were substantially improved in the meantime are omitted. The other notable change is that here we use the term "price of anarchy" instead of the original term "coordination ratio". The use of the latter term faded away in the literature, replaced by the term "price of anarchy" which was first introduced in [16].

which are more resilient to video-like traffic. Their point of view is also that of the mechanism design aspect of game theory, in that they try to design games (strategy spaces and reward tables) that encourage behaviors close to the social optimum. Understanding the worst-case distance of a Nash equilibrium from the social optimum in simple situations, which is the focus of the present paper, is a prerequisite for making rigorous progress in that project.

The model

Let us make the general game-theoretic framework more precise. Consider a network in which each link has a law (curve) whereby traffic determines delay. Each of several agents wants to send a particular amount of traffic along a path from a fixed source to a fixed destination. This immediately defines a game-theoretic framework, in which each agent has as many pure strategies as there are paths from its origin to its destination, and the cost for each agent is the delay experienced by the agent, as determined by the traffic on the links. There is also a well-defined optimization problem, in which we wish to minimize the *social* or *overall optimum*, the sum of all delays over all agents, say. The question we want to ask is, how far from the optimum total delay can the total delay achieved by a Nash equilibrium be? Numerical experiments reported in [7] imply that there are Nash equilibria which can be more than 20% off the overall optimum.

In this paper we address a very simple special case of this problem, in which the network is just a set of m parallel links from an origin to a destination, all with the same capacity (similar special cases are studied in other works in this field, e.g. [10]; we also briefly examine the case of two parallel links with unequal capacity). We model the delay of these links in a very simple way. Since the capacity is unit, we assume that the delay suffered by each agent using a link equals the total flow through this link. We assume that n agents have each an amount of traffic w_i, $i = 1, \ldots, n$ to send from the origin to the destination. Hence the resulting problem is essentially a task allocation problem with m machines and n independent tasks with lengths w_i, $i = 1, \ldots, n$. The set of pure strategies for agent i is therefore $\{1, \ldots, m\}$, and a mixed strategy is a distribution on this set. Let $(j_1, \ldots, j_n) \in \{1, \ldots, m\}^n$ be a combination of pure strategies, one for each agent; its *cost* for agent i, denoted $c_i = c_i(j_1, \ldots, j_n)$, is simply

$$c_i = \sum_{k\,:\,j_k = j_i} w_k,$$

the cost of the link chosen by i. This calculation assumes that, if agent i's traffic ends up in link j, the agent experiences delay equal to all traffic on link j; this is realistic if traffic is broken in packets, which are then sent in a round-robin way. Finally, the cost to agent i of a combination of mixed strategies is the expected cost of the corresponding experiment in which a pure strategy is chosen independently for each agent, with the probability assigned to it by the mixed strategy. The overall optimum in this situation, against which we propose to compare the Nash equilibria of the game just described, would be the optimum solution of the m-way load balancing (partition into m sets) problem for the n lengths $w_1, \ldots, w_n$.

The costs in our model are a simplification of the delays incurred in a network link when agents inject traffic into it. The actual delays are in fact not the sums of the individual delays, but nonlinear functions, as increased traffic causes increased loss rates and delays.

In this paper we show upper and lower bounds on the price of anarchy, the ratio between the worst Nash equilibrium and the overall optimum solution. Some results in the original conference version were improved in the meantime (see for example Chapter 20 in the book [14]) and we omit them. And a few of the original results are more or less intact, and these are:

- For m parallel links we analyze the structure of Nash equilibria.

- In a network with two identical parallel links, we show that the price of anarchy is $\frac{3}{2}$ (both upper bound and lower bound), independent of the number n of agents (Theorems 1 and 3).
- If the two links have different capacities, then the worst-case ratio increases to the golden ratio $\phi \approx 1.618$ (lower bound, Theorem 4).
- In a network of m parallel links, the expected traffic experienced by every individual agent at a Nash equilibrium is at most $(2 - 1/m)$ times the optimum (Theorem 2).
- However, when we consider the expected maximum over all agents which is the cost for the system, the bound is much higher. More precisely, in a network of m parallel links, the price of anarchy is $\Omega(\frac{\log m}{\log \log m})$ (Theorem 1). In the conference version of the paper we had a nonmatching upper bound. It was later shown that this lower bound is essentially tight [3, 8].

2 All Nash equilibria

We consider the case of n agents sharing m identical links. Before describing all Nash equilibria, we need a few definitions. We usually use subscripts for agents and superscripts for links. For example, for a Nash equilibrium, we denote the probability that agent i selects link j with p_i^j. Let M^j denote the *expected traffic* on link j. It is easy to see that

$$M^j = \sum_i p_i^j w_i. \tag{1}$$

From the point of view of agent i, its cost when its own traffic w_i is assigned to link j is

$$c_i^j = w_i + \sum_{k \neq i} p_k^j w_k = M^j + (1 - p_i^j) w_i. \tag{2}$$

Th probabilities p_i^j define a Nash equilibrium if there is no incentive for agent i to change its strategy. Thus, agent i will assign nonzero probabilities only to links j that minimize c_i^j. We will denote this minimum value by c_i, i.e.,

$$c_i = \min_j c_i^j,$$

and we will call the set of links $S_i = \{j : p_i^j > 0\}$ the *support* of agent i. More generally, let S_i^j be an indicator variable that takes the value 1 when $p_i^j > 0$.

Conversely, a Nash equilibrium is completely defined by the supports $S_1, \ldots, S_n$ of all agents. More precisely, if we fix the S_i^j, the strategies in a Nash equilibrium are given by

$$p_i^j = (M^j + w_i - c_i)/w_i \tag{3}$$

subject to

$$\text{for all } j\text{: } M^j = \textstyle\sum_i S_i^j (M^j + w_i - c_i)$$
$$\text{for all } i\text{: } \textstyle\sum_j S_i^j (M^j + w_i - c_i) = w_i.$$

To see that these constraints indeed define an equilibrium, notice that the first set of equations is equivalent to (2). The constraints are equivalent to (1), and to the fact that the probabilities of agent i should sum up to exactly 1. Notice also that the set of constraints specify, in general, a unique solution for c_i and M^j (there are $n + m$ constraints and $n + m$ unknowns). If the resulting probabilities p_i^j are in the interval $(0, 1]$, then the above equations define an equilibrium with support S_i^j. Thus, an equilibrium is completely defined by the supports of the agents (although not all supports give rise to a feasible equilibrium). As a result, the number of equilibria is, in general, exponential in n and m.

A natural quantity associated with an equilibrium is the *expected maximum traffic* over all links:

$$\text{cost} = \sum_{j_1=1}^{m} \cdots \sum_{j_n=1}^{m} \prod_{i=1}^{n} p_i^{j_i} \max_{j=1,\ldots,m} \sum_{k:j_k=j} w_k. \tag{4}$$

We call this the *social cost* and we wish to compare it with the social optimum opt. More precisely, we want to estimate the *price of anarchy* or *coordination ratio* which is the worst-case ratio $R = \max \text{cost}/\text{opt}$ (the maximum is over all equilibria). If we view the problem as scheduling [2], then we want to estimate the worst-case ratio of the makespan at a Nash equilibrium over the optimal makespan. Computing the social optimum opt is an NP-complete problem (partition problem), but for the purpose of upper bounding R here, it suffices to use two simple approximations of it: $\text{opt} \geq \max\{w_1, \sum_j M^j/m\} = \max\{w_1, \sum_i w_i/m\}$ (we shall be assuming that $w_1 \geq w_2 \geq \cdots \geq w_n$).

3 Worst-case equilibria

Our first theorem is almost trivial:

Theorem 1. *The price of anarchy for m links is $\Omega(\log n/\log\log n)$. In particular, for $m = 2$ links, it is at least* $3/2$.

Proof. Consider the case where there are $n = m$ agents, each with unit traffic, i.e., $w_i = 1$. It is easy to see that the set of uniform strategies with $p_i^j = 1/m$ for $i, j = 1, \ldots, m$ is a Nash equilibrium. To compute the social cost of the equilibrium we see this as the problem of throwing m balls into m bins. The social cost of the equilibrium is equal to the expected maximum number of balls in a bin which is well known to be $\Theta(\log m/\log\log m)$ [11]. Given that the optimal solution has cost 1, the lower bound follows.

For $m = 2$, this gives a lower bound of $3/2$. □

Our main technical result is a matching upper bound for $m = 2$. To prove it, we find a way to upper bound the complicated expression (4) for the social cost. In fact, it is relatively easy to compute the strategies of a Nash equilibrium. There are 2 types of agent: *pure strategy agents* with support of size 1 and *stochastic agents* with support of size 2. Let d^j be the traffic of pure strategy agents assigned to link j. Also, let $k > 1$ denote the number of stochastic agents. It is not difficult to verify that the system of equations (3) gives the following probabilities of a stochastic agent i:

$$p_i^j = \frac{1}{2} - \frac{d^1 + d^2 - 2d^j}{2(k-1)w_i}. \tag{5}$$

When there is only one stochastic agent, $k = 1$, the probabilities are $1/2$.

We do not see how to use this expression to upper bound (4). Instead, we use an indirect way. But first, we give an important bound, which holds for any number of agents and links, and it is interesting in its own right:

Theorem 2. *In a network of m identical parallel links, at every Nash equilibrium the expected cost c_i for each agent i is at most $(2 - 1/m)$ times the optimum, and more precisely,*

$$c_i \leq \frac{\sum_k w_k}{m} + \frac{m-1}{m} w_i. \tag{6}$$

Proof. This follows from

$$\begin{aligned} c_i &= \min_j c_i^j \\ &\leq \frac{1}{m} \sum_j c_i^j \\ &= \frac{1}{m} \sum_j (M^j + (1 - p_i^j) w_i) \\ &= \frac{\sum_j M_j}{m} + \frac{m-1}{m} w_i \\ &= \frac{\sum_k w_k}{m} + \frac{m-1}{m} w_i. \end{aligned}$$

□

We can now show a tight bound for the price of anarchy of $m = 2$ links.

Theorem 3. *The price of anarchy for any number of agents and $m = 2$ links is at most $3/2$.*

Proof. Central to the proof of the upper bound is the notion of *contribution probability*: The contribution probability q_i of agent i is equal to the probability that its traffic goes to the link of maximum load (if there is more than one maximum load link, we consider the lexicographically first such link, say). Clearly, the social cost is given by cost $= \sum_i q_i w_i$. The key idea in the proof is to consider the pairwise contribution to social cost. In particular, let t_{ik} be the *collision probability* of agents i and k; that is, the probability that the traffic of both agents goes to the same link: $t_{ik} = \sum_j p_i^j p_k^j$. Observe then that both agents i and k can contribute to the social cost only if they collide; that is,

$$q_i + q_l \leq 1 + t_{ik}. \tag{7}$$

We also observe that the expected cost c_i experienced by agent i is

$$c_i = w_i + \sum_{k \neq i} t_{ik} w_k,$$

since the expected contribution of agent k to the link of agent i is $t_{ik} w_k$.

We now have all the bounds needed to prove the theorem. Fix now some agent i. We estimate

$$\begin{aligned} \sum_{k \neq i} (q_i + q_k) w_k &\leq \sum_{k \neq i} (1 + t_{ik}) w_k \\ &= \sum_{k \neq i} w_k + \sum_{k \neq i} t_{ik} w_k \\ &= \sum_{k \neq i} w_k + c_i - w_i \\ &\leq \sum_{k \neq i} w_k + \frac{\sum_k w_k}{2} + \frac{w_i}{2} - w_i \\ &\leq \frac{3}{2} \sum_{k \neq i} w_k, \end{aligned}$$

where we used the bound (6) for $m = 2$, in the second to last inequality. By rearranging the above terms we get

$$\sum_{k \neq i} q_k w_k \leq \left(\frac{3}{2} - q_i \right) \sum_{k \neq i} w_k.$$

We can now bound the cost which is equal to $\sum_k q_k w_k$.

$$\text{cost} \le \left(\frac{3}{2} - q_i\right) \sum_{k \neq i} w_k + q_i w_i$$
$$= \left(\frac{3}{2} - q_i\right) \sum_{k} w_k + \left(2q_i - \frac{3}{2}\right) w_i.$$

Recall that opt $\ge \max\{\frac{1}{2}\sum_k w_k, w_i\}$. If for some agent i, $q_i \ge \frac{3}{4}$, then both coefficients in the last expression are nonnegative and we can bound both terms

$$\text{cost} \le \left(\frac{3}{2} - q_i\right) 2\text{opt} + \left(2q_i - \frac{3}{2}\right) \text{opt} = \frac{3}{2}\text{opt}.$$

If there is no $q_i \ge \frac{3}{4}$, the second coefficient is negative and this approach does not work. But in this case it is trivial to bound the cost: when all contribution probabilities are at most $\frac{3}{4}$,

$$\text{cost} = \sum_k q_k w_k \le \frac{3}{4} \sum_k w_k \le \frac{3}{2}\text{opt}.$$

□

Links with different capacities

So far, we have assumed that all links have the same capacity. We now consider the general problem where links may have different capacities or speeds. Let s_j be the speed of link j. Without loss of generality, we shall assume $s_1 \le \cdots \le s_m$. We can estimate all Nash equilibria again. Equation (2) now becomes

$$c_i^j = (M^j + (1 - p_i^j) w_i)/s_j. \tag{8}$$

and the equilibria are given by

$$p_i^j = (M^j + w_i - s_j c_i)/w_i \tag{9}$$

subject to

for all j: $M^j = \sum_i S_i^j (M^j + w_i - s_j c_i)$
for all i: $\sum_j S_i^j (M^j + w_i - s_j c_i) = w_i$.

We can extend the lower bound Theorem 1 to this case for $m = 2$:

Theorem 4. *The price of anarchy for two links with speeds $s_1 \le s_2$ is at least $R = 1 + s_2/(s_1+s_2)$ when $s_2 \le \phi\, s_1$, where $\phi = (1+\sqrt{5})/2$. The expression R achieves its maximum value ϕ when $s_2/s_1 = \phi$.*

Proof. We first describe the equilibria for $m = 2$ and any number of agents, generalizing (5). Again let d^j be the sum of all traffic assigned to link j by pure agents and let $k > 1$ be the number of stochastic agents. We give the probabilities p_i^1 of the stochastic agents ($p_i^2 = 1 - p_i^1$).

$$p_i^1 = \frac{s_2}{s_1 + s_2} - \frac{(s_2 - s_1)\sum_k w_k + (s_2 d^1 - s_1 d^2)}{(k-1)(s_1 + s_2) w_i}.$$

When there is only one stochastic agent, $k = 1$, the probability is $p_i^1 = \frac{s_2}{s_1+s_2}$. It is not hard to verify that these probabilities indeed satisfy (9). To prove the theorem, we consider the case of no initial loads and two agents with jobs $w_1 = s_2$ and $w_2 = s_1$. The probabilities are $p_1^1 = \frac{s_1^2}{s_2(s_1+s_2)}$ and $p_2^1 = 1 - \frac{s_2^2}{s_1(s_1+s_2)}$. We can then compute cost $= (p_1^1 p_2^1/s_1 + p_1^2 p_2^2/s_2)(w_1 + w_2) + (p_1^1 p_2^2/s_1 + p_1^2 p_2^1/s_2) w_1 = (s_1 + 2s_2)/(s_1 + s_2)$ and opt $= 1$. The lower bound follows.

It is worth mentioning that when $s_2/s_1 > \phi$ the probabilities given above are outside the interval $[0, 1]$. Therefore, both agents have pure strategies and the price of anarchy is 1. □

4 Discussion

We believe that the approach introduced in this paper, namely evaluating the worst-case ratio of Nash equilibria to the social optimum, may prove a useful calculation in many contexts. Although the Nash equilibrium is not trivial to reach without coordination, it does serve as an important indicator of the kinds of behaviors exhibited by noncooperative agents.

The questions left open by the conference version of this work have been answered conclusively. The immediate important problem that was left open was to bound the price of anarchy for $m \geq 3$ links. The answer was given in [3] and [8]: the lower bound of Theorem 1 is tight. Czumaj and Vöcking [3] also gave tight bounds for links with different capacities: the price of anarchy in this case is $O(\log m / \log \log \log m)$.

A large body of work has been produced around the concept of the price of anarchy. It is infeasible to list it here, but we refer the reader to the book [14] which devotes a significant fraction to these issues. The popularity of the concept of the price of anarchy owes much to the follow-up work of Roughgarden and Tardos [18] which opened the way to studying the price of anarchy in atomic and nonatomic congestion games.

The conference version of this work together with the work of Nisan and Ronen [13, 12] on algorithmic mechanism design, which appeared at around the same time, perhaps provided the fuse for the explosive growth of algorithmic game theory in the last decade.

References

[1] B. Braden, D. Clark, J. Crowcroft, B. Davie, S. Deering, D. Estrin, S. Floyd, V. Jacobson, G. Minshall, C. Partridge, L. Peterson, K. Ramakrishnan, S. Shenker, J. Wroclawski, and L. Zhang. Recommendations on Queue Management and Congestion Avoidance in the Internet, April 1998.
`http://info.internet.isi.edu:80/in-notes/rfc/files/rfc2309.txt`

[2] Y. Cho and S. Sahni. Bounds for list schedules on uniform processors. *SIAM Journal on Computing*, 9(1):91–103, 1980.

[3] A. Czumaj and B. Vöcking. Tight Bounds for Worst-case Equilibria. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms,* pp. 413–420, January 2002.

[4] S. Floyd and K. Fall. Router Mechanisms to Support End-to-End Congestion Control. Technical report, Lawrence Berkeley National Laboratory, February 1997.

[5] G. R. Grimmet and D. R. Stirzaker. *Probability and Random Processes, 2nd ed.*. Oxford University Press, 1992.

[6] Y. Korilis and A. Lazar. On the existence of equilibria in noncooperative optimal flow control. *Journal of the ACM* 42(3):584-613, 1995.

[7] Y. Korilis, A. Lazar, A. Orda. Architecting noncooperative networks. *IEEE J. Selected Areas of Comm., 13*, 7, 1995.

[8] E. Koutsoupias, M. Mavronicolas and P. Spirakis. Approximate Equilibria and Ball Fusion. In *Proceedings of the 9th International Colloquium on Structural Information and Communication Complexity (SIROCCO)*, 2002

[9] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, pages 404-413, 1999.

[10] R. La, V. Anantharam. Optimal routing control: Game theoretic approach. *Proc. 1997 CDC Conf.*

[11] R. Motwani and P. Raghavan, *Randomized Algorithms,* Cambridge University Press, 1995.

[12] N. Nisan. Algorithms for selfish agents: Mechanism design for distributed computation. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, pages 1–15, 1999.

[13] N. Nisan and A. Ronen. Algorithmic mechanism design. In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pages 129–140, May 01-04, 1999, Atlanta, Georgia, United States.

[14] N. Nisan, T. Roughgarden, E. Tardos, and V.V. Vazirani. *Algorithmic Game Theory.* Cambridge University Press, 2007.

[15] G. Owen. *Game Theory, 3rd ed..* Academic Press, 1995.

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

[17] C. H. Papadimitriou, M. Yannakakis. On complexity as bounded rationality. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on the Theory of Computing*, pages 726-733, Montreal, Quebec, Canada, 23-25 May 1994.

[18] T. Roughgarden and E. Tardos. How bad is selfish routing? In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, pages 93–102, Redondo Beach, CA, November 12-14, 2000.

[19] S. J. Shenker. Making greed work in networks: a game-theoretic analysis of switch service disciplines. *IEEE/ACM Transactions on Networking*, 3(6):819-831, Dec. 1995.

[20] S. Shenker, D. Clark, D. Estrin, and S. Herzog. Pricing in Computer Network: Reshaping the Research Agenda. *Communications Policy*, 20(1), 1996.