

Worst-case Equilibria

Elias Koutsoupias¹ and Christos Papadimitriou²

¹ Univ of California, Los Angeles
elias@cs.ucla.edu

² Univ of California, Berkeley
christos@cs.berkeley.edu

Abstract. In a system in which noncooperative agents share a common resource, we propose 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 in which several agents share a very simple network leads to some interesting mathematics, results, and open problems.

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, non-cooperative agents obviously evoke *game theory* [8], and its main concept of rational behavior, the *Nash equilibrium*: In an environment in which each agent 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 [8, 10] being the best-known example. Conditions under which Nash equilibria can achieve or approximate the overall optimum have been studied extensively ([10]; see also [5, 7, 11] 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, 12], with an eye towards designing variants of the Internet Protocols 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 to an agent of a combination of strategies (one for each agent) is the negative of the total delay for each 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 be the total delay achieved by a Nash equilibrium? Numerical experiments reported in [6] 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. [7]; 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 capacity of flow through this link. We assume that n agents have each an amount of traffic w_i , $i = 1, \dots, n$ to send from the origin to the destination. Hence the resulting problem is essentially a scheduling problem with m links and n independent tasks with lengths w_i , $i = 1, \dots, n$. The set of pure strategies for agent i is therefore $\{1, \dots, m\}$, and a mixed strategy is a distribution on this set. Let $(j_1, \dots, j_n) \in \{1, \dots, m\}^n$ be a combination of pure strategies, one for each agent; its *cost* for agent i , denoted $C_i(j_1, \dots, j_n)$, is simply

$$L^{j_i} + \sum_{j_k=j_i} w_k,$$

the finish time of the link chosen by i ; here we assume that link j has in the beginning an initial task of length L^j scheduled, so it will be available for scheduling the agents' tasks only after L^j time units. This calculation assumes that, if agent i 's task ends up in link j , it ends when all tasks on link j end; this is realistic if the tasks are broken in packets, which are then sent in a round-robin way. We also examine the alternative model, in which the tasks scheduled in link j are executed in a random batch order, and hence the cost to agent i is

$C_i(j_1, \dots, j_n) = L^{j_i} + \frac{1}{2} \sum_{j_k=j_i} w_k$. We call this *the batch model*. 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, \dots, 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. We discuss briefly in the last section the open problems suggested by our work that are associated with more accurate modeling of network delays.

The results of this paper

In this paper we show upper and lower bounds on the ratio between the worst Nash equilibrium and the overall optimum solution.

- In a network with two parallel links, we show that the worst-case ratio is $\frac{3}{2}$ (both upper and lower bound), independent of the number n of agents (Theorems 1 and 2).
- The above result assumes that the two link speeds are the same. If the two links have different speeds, then the worst-case ratio increases to the golden ratio $\phi = 1.618\dots$ (lower bound, Theorem 3).
- Also, in the batch model of two links, the worst-case ratio is lower bounded by $\frac{29}{18} = 1.6111\dots$ which is also an upper bound if we have two agents (Theorem 4).
- We have not been able to determine the answers for three or more links. However, the worst-case ratio (in all of the above models) is bounded from below by the ratio suggested by the load-balancing aspect of the problem, that is to say, $\Omega(\frac{\log m}{\log \log m})$ (Theorem 6). Using the Azuma-Hoeffding inequality, we establish an $O(\sqrt{m \log m})$ upper bound (Theorem 8). A similar bound holds for links of different speeds (Theorem 9).

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 agents i selects link j with p_i^j . Let M^j denote the *expected traffic* on link j . If L^j is the initial load on link j , it is easy to see that

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

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

$$c_i^j = w_i + L^j + \sum_{t \neq i} p_t^j w_t = M^j + (1 - p_i^j)w_i. \quad (2)$$

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 value 1 when $p_i^j > 0$.

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

$$p_i^j = (M^j + w_i - c_i)/w_i \quad (3)$$

subject to

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

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

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:

$$\text{cost} = \sum_{j_1=1}^m \cdots \sum_{j_n=1}^m \prod_{i=1}^n p_i^{j_i} \max_{j=1, \dots, m} \{L^j + \sum_{t: j_t=j} w_t\}. \quad (4)$$

We call it the *social cost* and we wish to compare it with the social optimum opt. More precisely, we want to estimate the *coordination ratio* which is the worst-case ratio $R = \max \text{cost}/\text{opt}$ (the maximum is over all equilibria). 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: $\text{opt} \geq \max\{w_1, \sum_j M^j/m\} = \max\{w_1, (\sum_j L^j + \sum_i w_i)/m\}$ (we shall be assuming that $w_1 \geq w_2 \geq \dots \geq w_n$).

3 Worst-case equilibria for 2 links

We shall assume that there are no initial loads —that is, all L^j 's are zero. This is no restriction at all for the standard model, because initial loads can be considered as jobs of m additional agents, each with a pure strategy. However, this may not be true for other models. In particular, in the batch model (the one with the $\frac{1}{2}$ factor in front of $\sum w_i$) it follows from our results that initial loads result in strictly worse ratio.

Our first theorem is trivial:

Theorem 1. *The coordination ratio for 2 links is at least 3/2.*

Proof. Consider two agents with traffic $w_1 = w_2 = 1$. It is easy to check that probabilities $p_i^j = 1/2$ for $i, j = 1, 2$ give rise to a Nash equilibrium. The expected maximum load is $\text{cost} = 3/2$ and the social optimum is $\text{opt} = 1$ achieved by allocating each job to its own link.

Our main technical result of this section is a matching upper bound. 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 agents: *pure strategy agents* with support of size one and *stochastic agents* with support of size 2. Let d^j be the sum of all jobs 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 :

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

However, we don't see how to use this expression to upper bound (4).

Central to our proof of the upper bound is the notion *contribution probability*: The contribution probability q_i of agent i is equal to the probability that its job goes to the link of maximum load (if there are more than one maximum load links, we consider the lexicographically first such link, say). Clearly, the social cost is given by $\text{cost} = \sum_i q_i w_i$. The key idea in our 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. Observe then that both agents i and k can contribute to the social cost only if they collide, that is,

$$q_i + q_k \leq 1 + t_{ik}. \quad (6)$$

The following lemma provides a crucial property of collision probabilities. It holds for any number of links.

Lemma 1. *The collision probabilities of a Nash equilibrium of n agents and m links satisfy*

$$\sum_{k \neq i} t_{ik} w_k = c_i - w_i.$$

Proof. Observe first that $t_{ik} = \sum_j p_i^j p_k^j$. Therefore, we have

$$\sum_{k \neq i} t_{ik} w_k = \sum_j p_i^j \sum_{k \neq i} p_k^j w_k = \sum_j p_i^j (M^j - p_i^j w_i).$$

It follows from (3) that we can use $p_i^j w_i = M^j + w_i - c_i$. There is a minor technical point to be made here: the equality $p_i^j w_i = M^j + w_i - c_i$ holds only if link j is in the support of agent i ($p_i^j > 0$). However, observe that when $p_i^j = 0$ there is no harm in replacing $p_i^j w_i$ with any expression. We get

$$\sum_{k \neq i} t_{ik} w_k = \sum_j p_i^j (c_i - w_i) = c_i - w_i.$$

A final ingredient for the proof is the bound (which also holds for any number of agents and links):

$$c_i \leq \frac{\sum_i w_i}{m} + \frac{m-1}{m} w_i. \quad (7)$$

This follows from $c_i = \min_j c_i^j \geq \frac{1}{m} \sum_j (M^j + (1 - p_i^j) w_i) = \frac{\sum_j M_j}{m} + \frac{m-1}{m} w_i = \frac{\sum_k w_k}{m} + \frac{m-1}{m} w_i$.

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

Proof. We have seen that pairwise the contribution probabilities satisfy $q_i + q_k \leq 1 + t_{ik}$. Therefore, $\sum_{k \neq i} (q_i + q_k) w_k \leq \sum_{k \neq i} (1 + t_{ik}) w_k$. Using Lemma 1 and bound (7), we get $\sum_{k \neq i} (q_i + q_k) w_k \leq \frac{3}{2} \sum_{k \neq i} w_k$. From this we can compute

$$\text{cost} = \sum_k q_k w_k = \left(\frac{3}{2} - q_i\right) \sum_k w_k + \left(2q_i - \frac{3}{2}\right) w_i.$$

Recall that $\text{opt} \geq \max\{\frac{1}{2} \sum_k w_k, w_i\}$. If for some agent i , $q_i \geq \frac{3}{4}$, then $(2q_i - \frac{3}{2}) w_i \leq (2q_i - \frac{3}{2}) \text{opt}$ and $\text{cost} \leq (\frac{3}{2} - q_i) 2\text{opt} + (2q_i - \frac{3}{2}) \text{opt} = \frac{3}{2} \text{opt}$. Otherwise, when all contribution probabilities are at most $\frac{3}{4}$, $\text{cost} = \sum_k q_k w_k \leq \frac{3}{4} \sum_k w_k \leq \frac{3}{2} \text{opt}$.

Links with different speeds

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

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

and the equilibria are given by:

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

subject to

$$\begin{aligned} \text{for all } j: M^j &= L^j + \sum_i S_i^j (M^j + w_i - s_j c_i) \\ \text{for all } i: \sum_j S_i^j (M^j + w_i - s_j c_i) &= w_i \end{aligned}$$

We can extend the lower bound Theorem 1 to this case:

Theorem 3. *The coordination ratio for two links with speeds $s_1 \leq s_2$ is at least $R = 1 + s_2 / (s_1 + s_2)$ when $s_2 \leq \phi s_1$, where $\phi = (1 + \sqrt{5}) / 2$. The coordination ratio R achieves its maximum value ϕ when $s_2 / s_1 = \phi$.*

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

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

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 $\text{cost} = (p_1^1 p_2^1 / s_1 + p_1^2 p_2^2 / s_2)(w_1 + w_2) + (p_1^1 p_2^2 / s_1 + p_1^2 p_2^1 / s_2) w_1 = (s_1 + 2s_2) / (s_1 + s_2)$ and $\text{opt} = 1$. 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 coordination ratio is 1.

We believe that the proof of Theorem 2 can be appropriately generalized to the case of links of different speeds.

The batch model

For the batch model with two links we can prove the following bounds (proof omitted):

Theorem 4. *In the batch model with two identical links, the coordination ratio is between $\frac{29}{18} = 1.61\dots$ and 2. The lower bound $\frac{29}{18}$ is also an upper bound in the case of $n = 2$ agents.*

When the links have no initial load, the batch model and the standard model have the same equilibria and the same coordination ratio. However, in the general case, as the above theorem demonstrates, the batch model has higher coordination ratio. But it cannot be much higher:

Theorem 5. *For m links and any number of agents, the coordination ratios of the batch model and the standard model differ by at most a factor of 2.*

We omit the details of the proof, but we point out the main idea: We can consider the initial loads L^j of the batch model as pure strategy agents of weight $2L_j$. This preserves the equilibria and changes the social optimum by at most a factor of 2.

4 Worst-case equilibria for m links

We now consider lower bounds for the coordination ratio for m links.

Theorem 6. *The coordination ratio for m identical links is $\Omega(\log m / \log \log m)$.*

Proof. Consider the case where there are m agents, each with a unit job, i.e., $w_i = 1$. If the links have no initial load, it is easy to see that the uniform strategies with $p_i^j = 1/m$ for $i, j = 1, \dots, m$ is an equilibrium. This is identical to the problem of throwing m balls into m bins and asking for the expected maximum number of balls in a bin. The answer is well-known to be $\Theta(\log m / \log \log m)$.

We believe that this lower bound is tight: That is, if T_m denotes the expected maximum number of balls in a bin, we conjecture that the coordination ratio for any number of agents and m identical links is T_m (in the standard model). Theorem 2 shows that the conjecture holds for $m = 2$.

We believe that a proof of the conjecture can be obtained by appropriately generalizing the proof technique of Theorem 2; it seems however that a substantially deeper structural theorem about the Nash equilibria, similar to Lemma 1, is needed. Here, we give a weaker upper bound. But first we need the following theorem, which is interesting on its own.

Theorem 7. *For m identical links, the expected load M^j of any link j is at most $(2 - 1/m)\text{opt}$. For links with different speeds, M_j is at most $s_j(1 + \sqrt{m - 1})\text{opt}$.*

Proof. For identical links the theorem follows directly from (7) by observing that $M^j \leq c_i \leq (\sum_i w_i)/m + (m - 1)w_i/m \leq s_j(2 - 1/m)\text{opt}$.

The proof for links with different speeds has the same flavor with (7). This time we take a weighted average over the links (the weight for machine j is $s_j / \sum_r s_r$). Thus,

$$c_i^j \leq \frac{\sum_r M^r + (m - 1)w_i}{\sum_r s_r}.$$

Also, $c_i^m \leq (M^m + w_i)/s_m \leq (\sum_r M^r + w_i)/s_m$. In summary,

$$c_i^j \leq \min\left\{\frac{\sum_r M^r + (m - 1)w_i}{\sum_r s_r}, \frac{\sum_r M^r + w_i}{s_m}\right\}.$$

However, we can lower bound the social optimum by $\max\{w_i/s_m, \sum_r M^r / \sum_r s_r\}$.

Thus, we get $c_i^j \leq \text{opt} + \min\left\{\frac{(m-1)w_i}{\sum_r s_r}, \frac{\sum_r M^r}{s_m}\right\}$. Using the obvious inequality $\min\{xa/b, c/d\} \leq \sqrt{x} \max\{a/d, c/b\}$, we get $c_i^j \leq (1 + \sqrt{m - 1})\text{opt}$. We can then conclude that $M^j \leq s_j c_i^j \leq s_j(1 + \sqrt{m - 1})\text{opt}$.

We can now prove an upper bound for the case of m identical links.

Theorem 8. *The coordination ratio of any number of agents and m identical links is at most $T = 3 + \sqrt{4m \ln m}$.*

Proof. Using a martingale concentration bound known as the Azuma-Hoeffding inequality [4], we will show that the load of a given link j exceeds $(T-1)\text{opt}$ with probability at most $1/m^2$. Then, the probability that the maximum load on all links does not exceed $(T-1)\text{opt}$ is at least $1-1/m$. It follows that the expected maximum load is bounded by $(1-1/m)(T-1)\text{opt} + 1/m(m\text{opt}) \leq T\text{opt}$.

It remains to show that indeed the probability that the load of a given link j exceeds $(T-1)\text{opt}$ is small (at most $1/m^2$). Let X_i be a random variable denoting the contribution of agent i to the load of link j . In particular, $\Pr[X_i = w_i] = p_i^j$ and $\Pr[X_i = 0] = 1 - p_i^j$. Clearly, the random variables X_1, \dots, X_n are independent. We are interested in estimating the probability $\Pr[\sum_i X_i > (T-1)\text{opt}]$. Since the weights w_i and the probabilities p_i^j may vary a lot, we don't expect the sum $\sum_i X_i$ to exhibit the good concentration bounds of sums of binomial variables. However, we can get a weaker bound using the Azuma-Hoeffding inequality. The inequality gives very good results for probabilities around $1/2$. Unfortunately, in our case the probabilities may be very close to 0 or 1.

Let $\mu_i = E[X_i]$ and consider the martingale $Y_t = X_1 + \dots + X_t + \mu_{t+1} + \dots + \mu_n$ (it is straightforward to verify $E[Y_{t+1}|Y_t] = Y_t$). Observe that $|Y_{t+1} - Y_t| = |X_{t+1} - \mu_{t+1}| \leq w_{t+1}$. We can then apply the Azuma-Hoeffding's inequality:

$$\Pr[Y_n - Y_0 \geq x] \leq e^{-\frac{1}{2}x^2 / \sum_i w_i^2}.$$

Let $x = (T-3)\text{opt}$. Since $Y_0 = \sum_i \mu_i = M^j \leq 2\text{opt}$ (Theorem 7), we get that the load of link j exceeds $(T-1)\text{opt}$ with probability at most $e^{-\frac{1}{2}x^2 / \sum_i w_i^2}$. However, it is not hard to establish that

$$\sum_i w_i^2 \leq \max\{mw_1^2, m(\sum_i w_i/m)^2\} \leq m\text{opt}^2.$$

Thus the probability that the load of link j exceeds $(T-1)\text{opt}$ is at most $e^{-\frac{1}{2}(T-3)^2/m}$. For $T = 3 + \sqrt{4m \ln m}$, this probability becomes $1/m^2$ and the proof is complete.

It is worth noticing that the only structural property of Nash equilibria we needed in the proof of the above theorem is that the expected load of a link j is at most 2opt (and, of course, the independence of the agent strategies). We can use a similar proof to extend the theorem to the case of m links with different speeds:

Theorem 9. *The coordination ratio of any number of agents and m different links is $O(\sqrt{\frac{sm}{s_1} \sum_j \frac{s_j}{s_1} \sqrt{\log m}})$.*

5 Discussion and open problems

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.

Besides bridging the gaps left open in our theorems, there are several extensions of this work that seem interesting, namely, investigating with the same point of view more complex and realistic cost models, for example, when the cost is given by $\frac{1}{C - \min\{C, \sum w_i\}}$, where C is the capacity of a link and $\sum w_i$ its load [7]. More important is the study of realistic Internet metrics, that result from the employed protocols such as the one related to TCP and the square root of the drop frequency [3]. Finally, it would be extremely interesting, once the relative quality of the Nash equilibria in such situations is better understood, to employ such understanding in the design of improved protocols [1].

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. S. Floyd and K. Fall. Router Mechanisms to Support End-to-End Congestion Control. Technical report, Lawrence Berkeley National Laboratory, February 1997.
4. G. R. Grimmet and D. R. Stirzaker. *Probability and Random Processes*, 2nd ed.. Oxford University Press, 1992.
5. Y. Korilis and A. Lazar. On the existence of equilibria in noncooperative optimal flow control. *Journal of the ACM* 42(3):584–613, 1995.
6. Y. Korilis, A. Lazar, A. Orda. Architecting noncooperative networks. *IEEE J. Selected Areas of Comm.*, 13, 7, 1995.
7. R. La, V. Anantharam. Optimal routing control: Game theoretic approach. *Proc. 1997 CDC Conf.*
8. G. Owen. *Game Theory*, 3rd ed.. Academic Press, 1995.
9. K. Park, M. Sitharam, S. Chen. Quality of service provision in noncooperative network environments. Manuscript, Purdue Univ., 1998.
10. 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.
11. 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.
12. S. Shenker, D. Clark, D. Estrin, and S. Herzog. Pricing in Computer Network: Reshaping the Research Agenda. *Communications Policy*, 20(1), 1996.