Game-Theoretic Analysis of Internet Switching with Selfish Users[*]

Alex Kesselman[1,**], Stefano Leonardi[2], and Vincenzo Bonifaci[2]

[1] Max Planck Institut für Informatik, Saarbrücken, Germany
akessel@mpi-sb.mpg.de
[2] Dipartimento di Informatica e Sistemistica,
Universita' di Roma "La Sapienza", Roma, Italy
leon@dis.uniroma1.it

Abstract. We consider the problem of Internet switching, where traffic is generated by selfish users. We study a packetized (TCP-like) traffic model, which is more realistic than the widely used fluid model. We assume that routers have First-In-First-Out (FIFO) buffers of bounded capacity managed by the drop-tail policy. The utility of each user depends on its transmission rate and the congestion level. Since selfish users try to maximize their own utility disregarding the system objectives, we study Nash equilibria that correspond to a steady state of the system. We quantify the degradation in the network performance called the price of anarchy resulting from such selfish behavior. We show that for a single bottleneck buffer, the price of anarchy is proportional to the number of users. Then we propose a simple modification of the Random Early Detection (RED) drop policy, which reduces the price of anarchy to a constant.

1 Introduction

If all Internet users voluntarily deploy a congestion-responsive transport protocol (e.g. TCP [19]), one can design this protocol so that the resulting network would achieve certain performance goals such as high utilization or low delay. However, with fast growth of the Internet users population, the assumption about cooperative behavior may not remain valid. Users are likely to behave "selfishly", that is each user makes decisions so as to optimize its own utility, without coordination with the other users. Buffer sharing and bandwidth allocation problems are prime candidates for such a selfish behavior.

 If a user does not reduce its sending rate upon congestion detection, it can get a better share of the network bandwidth. On the other hand, all users suffer during congestion collapse, since the network delay and the packet loss increase drastically. Therefore, it is important to understand the nature of congestion resulting from selfish behavior. A natural framework to analyze this class of problems is that of non-cooperative games, and an appropriate solution concept is that of Nash equilibrium [24]. Strategies of the

users are at a Nash equilibrium if no user can gain by unilaterally deviating from its current policy.

The subject of this paper is game-theoretic analysis of the Internet switching problem. We consider a bottleneck First-In-First-Out (FIFO) buffer shared by many users. The users compete for the buffer share in order to maximize their throughput, but they suffer from the created congestion. We assume that each user knows its buffer usage, the queue length and the buffer size. In reality, these parameters can be estimated only approximately (e.g. such a mechanism exists in TCP Vegas). However, in this work we rather study *lack of coordination* and not *lack of information*. The goal of a user is to maximize its utility. We assume that the utility function of a user increases when its goodput increases and decreases when the network congestion increases.

We assume that there are n users and the buffer capacity is B while $B \gg n$. In our model all packets have a constant size. Time is slotted. Every time step each user may or may not send a packet to the buffer.[1] These packets arrive in an arbitrary order and are processed by the buffer management policy one by one. Then the first packet in the FIFO order is transmitted on the output link.

The drop policy has to decide at each time step which of the packets to drop and which to accept. It can also preempt (drop) accepted packets from the buffer. Under the *drop-tail* policy, all arriving packets are accepted if the buffer is not full and dropped otherwise (when the buffer is full).

We denote by w_i^t the number of packets of user i in the buffer at the beginning of time step t. We call it the *buffer usage* of user i. We denote by $W^t = \sum_i w_i^t$ the *queue length* at time t. We also denote by W_{-i}^t the buffer usage of all users but user i, i.e., $W_{-i}^t = W^t - w_i^t$. We assume that the users are greedy, i.e., they always have data to send and the queue is never empty.

We denote by $r_i^t = w_i^t / W^t$ the *instant transmission rate* of user i at time t. We say that the system is in a *steady state* if the queue length remains constant. In a steady state, the *average transmission rate* of user i is $r_i = w_i / W$, where w_i is the average buffer usage of user i and W is the average queue length.

Now we proceed to define a utility function. Since in general the rate and the delay are incomparable values, it is natural to normalize one of these parameters. We introduce the notion of *congestion level*, that is the congestion level at time t is $L^t = \frac{W^t}{B}$. Note that the congestion level is zero when the buffer is empty and one when the buffer is full.

The utility of user i at time t is $u_i(w_i^t, W^t) = r_i^t \cdot (1 - L^t)$. We assume that user i sends a packet to the buffer at time t if its utility increases, i.e., $u_i(w_i^t + 1, W^t + 1) > u_i(w_i^t, W^t)$. In a steady state, the utility of user i is $u_i(w_i, W) = r_i \cdot (1 - L)$, where $L = W/B$. Observe that when the buffer is almost empty, the utility of each user approximately equals its transmission rate. On the other hand, when the buffer is nearly full, all users have utility close to zero. The latter situation can be viewed as "congestion collapse".

The *strategy* of each user is its buffer usage while the strategies of the other users define the background buffer backlog. Now we define a Nash equilibrium.

---

[1] We note that our results can be extended to the case in which each user can send an arbitrary number of packets.

Definition 1. *The system is said to be in a* Nash equilibrium *if no user can benefit by changing its buffer usage.*

The total utility of the users in a Nash equilibrium is $\sum_{i=1}^{n} u_i(w_i, W)$. Observe that under an optimal centralized policy, all users have equal sending rates of $1/n$ and experience almost negligible delay, which results in the total utility close to 1. We define the *price of anarchy* in a Nash equilibrium to be $1/\sum_{i=1}^{n} u_i(w_i, W)$. We are also concerned with the fairness of a Nash equilibrium. A Nash equilibrium is said to be *fair* if all users have the same buffer usage.

A Nash equilibrium in a networking environment is interesting only if it can be reached efficiently (in polynomial time). We define the *convergence time* to a Nash equilibrium as the maximum number of time steps required to reach a Nash equilibrium starting from an arbitrary state of the buffer.

We demonstrate that the drop-tail buffering policy imposes a fair Nash equilibrium. However, the price of anarchy is proportional to the number of users. We also show that the system converges to a Nash equilibrium in polynomial time, namely after $O(B^2)$ time steps. Then we propose a simple modification of the Random Early Detection (RED) policy [13] called Preemptive RED (PRED) that achieves a constant price of anarchy. We note that PRED is in the spirit of CHOKe [26] (see Section 4).

Paper organization. The rest of the paper is organized as follows. In Section 2 we discuss the related work. Analysis of a single switch appears in Section 3. The PRED policy is presented in Section 4. We conclude with Section 5.

2   Related Work

Shenker [29] analyzes switch service disciplines with a $M/M/1$ model and Markovian arrival rates. The utility of each user is increasing in its rate and decreasing in the network congestion. Shenker shows that the traditional FIFO policy does not guarantee efficiency and fairness and proposes a policy called Fair Share that guarantees both of them.

Garg at al. [17] study a switching problem using a continuous fluid-flow based traffic model, which is amenable to analysis of an arbitrary network. The user's utility is an increasing function of its goodput only. Garg at al. show that selfish behavior leads to congestion collapse and propose a rate inverse scheduling service discipline under which a max-min fair rate allocation is a Nash equilibrium. Contrary to [17], we consider only a simple FIFO scheduling policy.

Unfortunately, the complexity of the policies proposed in [29, 17] is too high to implement them in the core of the Internet since they have to maintain per-flow state. Dutta et al. [9] analyze simple state-less buffer management policies under the assumption that the traffic sources are Poisson and the utility of each user is proportional to its goodput. Dutta et al. demonstrate that drop-tail and RED do not impose Nash equilibria and present a modification of RED that enforces an efficient Nash equilibrium. Differently from [9], we assume that the utility of each user depends on the network congestion as well.

In a recent paper, Gao et al. [16] propose efficient drop policy that in case of congestion drops packets of the highest-rate sender. Gao et al. show that if all sources are

Poisson, this policy results in a Nash equilibrium which is a max-min fair rates allocation. In addition, it is demonstrated that the throughput of a TCP source is a constant factor of its max-min-fairness value when competing with Poisson sources. In contrast to [16], we do not make any assumptions regarding the traffic pattern.

Unlike the works mentioned above, in this paper we study *packetized* traffic model in which sources do not control the sending rate explicitly, but rather make it on per-packet basis. This model is implicit in the TCP protocol.

Nash equilibria for network games have been extensively studied in the literature. Douligeris and Mazumdar [8] determine conditions for a Nash equilibrium for a $M/M/1$ system when the utility of a user is a function the throughput and the delay. Orda et al. [25] investigate uniqueness of Nash equilibria in communication networks with selfish users. Korilis and Lazar [21] study Nash equilibria of a non-cooperative flow control game. Congestion control schemes utilizing pricing based on explicit feedback are proposed by Kelly et al. [20]. Gibbens and Kelly [18] explore the implementation of network pricing in which users are charged for marks that routers place on packets in order to achieve congestion control. Akella et al. [1] present a game-theoretic analysis of TCP congestion control. Qiu et al. [27] consider selfish routing in intra-domain network environments.

Traditionally, in Computer Science research has been focused on finding a global optimum. With the emerging interest in computational issues in game theory, the *price of anarchy* introduced by Koutsoupias and Papadimitriou [22] has received considerable attention [6, 7, 14, 28, 11]. The price of anarchy is the ratio between the cost of the worst possible Nash equilibrium (the one with the maximum social cost) and the cost of the social optimum (an optimal solution with the minimum social cost). In some cases the price of anarchy is small, and thus good performance can be achieved even without a centralized control.

If a Nash equilibrium imposed by selfish users is not satisfiable, one can deploy resource allocation policies to improve the situation. Recently, Christodoulou at al. [5] introduced the notion of coordination mechanism, which is a resource allocation policy whose goal is to enforce a better Nash equilibrium. They show that even simple local coordination mechanisms can significantly reduce the price of anarchy.

Efficient convergence to a Nash equilibrium is especially important in the network environment, which is highly variable. The question of convergence to a Nash equilibrium has received significant attention in the game theory literature [15]. Altman at al. [2] and Boulogne at al. [4] analyze the convergence to a Nash equilibrium in the limit for a routing and a scheduling game, respectively. Even-Dar et al. [10] consider deterministic convergence time to a pure Nash equilibrium for a load balancing game (routing on parallel links).

3   Single Switch

In this section we consider a single switch. We first characterize Nash equilibria and derive the price of anarchy. Then we establish an upper bound on the convergence time. Finally, we study the Nash traffic model.

3.1   Price of Anarchy

The intuition explaining why congestion collapse does not happen is that at some point users start to suffer from the delay due to their own packets in the bottleneck buffer and would not further increase their sending rates. The next theorem shows the existence of a unique Nash equilibrium and derives the price of anarchy.

Theorem 1. *There exists a unique Nash equilibrium that has the price of anarchy of n and is fair.*

Proof. By the definition of a Nash equilibrium, no user can increase its utility by changing the buffer usage. Therefore, a Nash equilibrium is a local maximum point for the utility function of each user i. We have that the first derivative $(u_i(w_i, W))'$ must be equal to zero:

$$\frac{-w_i^2 - 2W_{-i}w_i - W_{-i}^2 + BW_{-i}}{B(w_i + W_{-i})^2} = 0,$$

and thus

$$w_i = \sqrt{BW_{-i}} - W_{-i}. \tag{1}$$

It is easy to verify that the second derivative $(u_i(w_i, W))''$ is negative: $\frac{-2W_{-i}}{B(w_i+W_{-i})^3} < 0$.

We obtain that a Nash equilibrium is fair since for each user i, $W = \sqrt{BW_{-i}}$, where W is constant. Summing Equation (1) for all users, we get: $W = B(n-1)/n$. Therefore, the unique Nash equilibrium is $w_i = B(n-1)/n^2$ for each user i. The price of anarchy in the Nash equilibrium is $1/\sum_{i=1}^{n}\left(\frac{1}{n} \cdot \frac{1}{n}\right) = n$. ∎

For simplicity, we assume that $B(n-1)/n^2$ is an integer. Otherwise, there would be an approximate Nash equilibrium in which the buffer usage of each user varies between $\lfloor B(n-1)/n^2 \rfloor$ and $\lceil B(n-1)/n^2 \rceil$ as our simulations indeed show.

3.2   Convergence

Now we analyze convergence to a Nash equilibrium. The next theorem demonstrates that the convergence time to a Nash equilibrium is proportional to the square of the buffer size.

Theorem 2. *The convergence time to a Nash equilibrium is $O(B^2)$.*

Proof. Let $w^* = B(n-1)/n^2$ be the buffer usage of each user in a Nash equilibrium. Observe that until a Nash equilibrium is reached, each time step either a user with the maximum buffer usage (greater than w^*) can benefit from decreasing its buffer usage or a user with the minimum buffer usage (smaller than w^*) can benefit from increasing its buffer usage. Note that a user can increase its buffer usage at any time if the buffer is not full, but can decrease its buffer usage only when one of its packets is transmitted out of the buffer.

Initially, we let the system run for B time steps so that the users would be able to fill the buffer. Note that at this point no user with the maximum buffer usage wishes to

increase its buffer usage. In the sequel, the maximum buffer usage will only decrease and the minimum buffer usage will only increase until they become equal to w^*.

We divide time into *phases* of B time steps. Consider a phase that starts at time step t. We argue that if a Nash equilibrium is not reached by time $t + B$, then either the maximum buffer usage decreases or the minimum buffer usage increases. Otherwise, there would be a Nash equilibrium contradicting Theorem 1, which states that in a Nash equilibrium all users have the same buffer usage. Therefore, the system reaches a Nash equilibrium after B phases or in B^2 time steps. ∎

We also perform a convergence simulation. We consider different settings, where initially the buffer is filled randomly with packets of different users. The dynamics of the minimum and the maximum buffer usage for 8 and 16 users appears in Figure 1. The simulation results show that the actual convergence time is roughly proportional to the buffer size.
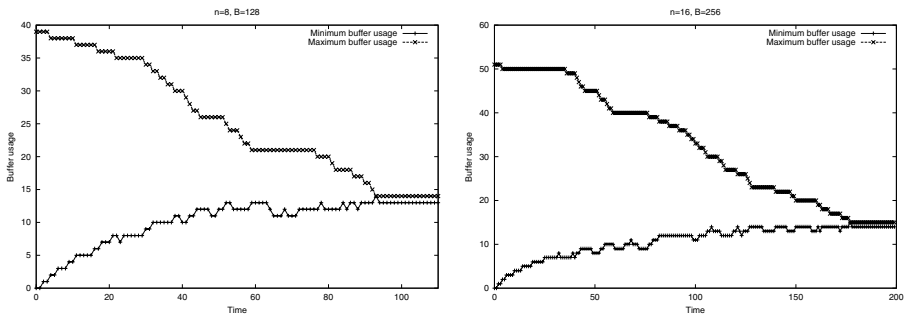

Fig. 1. Convergence to a Nash equilibrium

3.3   Nash Traffic Model

The traffic model of any QoS architecture is the regulation of the rate at which a user is allowed to inject packets into the network. There are two important policing criteria: *average rate* – limit of the long-term average rate at which packets can be injected into the network and *burst size* – the limit of the maximum number of packets that can be sent into the network over an extremely short time interval.

Definition 2. *The (σ, ρ) leaky bucket mechanism is an abstraction modeling a source with an average rate of ρ and burst size of at most σ. Therefore, at any time interval of length l there are at most $\rho \cdot l + \sigma$ packets entering the system.*

We show that the traffic of a user in a Nash equilibrium conforms to the leaky bucket model.

Theorem 3. *In a Nash equilibrium the traffic of each user conforms to the leaky bucket model with $\sigma = B(n - 1)/n^2$ and $\rho = 1/n$.*

The theorem follows due to the fact that in order to maintain the buffer usage of $B(n-1)/n^2$, each user must send a burst of $B(n-1)/n^2$ packets and keep sending packets at the average rate of $1/n$.

4   PRED Policy

In this section we consider the problem of designing a coordination mechanism (drop policy) that will improve the price of anarchy. The main goal of a drop policy is to control the average queue length. This can be done by dropping arriving packets probabilistically when the average queue length exceeds some pre-defined threshold. Such a policy called Random Early Detection (RED) has been introduced by Floyd and Jacobson [13]. We will show that a simple modification of RED reduces the price of anarchy to a constant. Again, we assume that the users are aware of the drop policy.

The RED policy calculates the average queue length, using a low-pass filter with an exponential weighted moving average. The average queue length is compared to two thresholds, a minimum threshold and a maximum threshold. When the average queue length is less than the minimum threshold T_l, no packets are dropped. When the average queue length is greater than the maximum threshold T_h, every arriving packet is dropped. When the average queue length is between the minimum and the maximum threshold, each arriving packet is dropped with probability p, which is a function of the average queue length.

Pan et al. [26] devise a simple packet dropping scheme, called CHOKe, that discriminates against unresponsive or misbehaving flows aiming to approximate the fair queuing policy. When a packet arrives at a congested router, CHOKe picks up a packet at random from the buffer and compares it with the arriving packet. If they both belong to the same flow, then they are both dropped. Otherwise, the randomly chosen packet is left intact and the arriving packet is admitted into the buffer with the same probability as in RED.

Our goal is to ensure a small queue length. Unfortunately, CHOKe is not aggressive enough in penalizing misbehaving flows since the probability that the buffer usage of a non-responsive user decreases is inversely proportional to the queue length. We propose a modification of RED in the spirit of CHOKe, called Preemptive RED (PRED). The PRED policy is presented in Figure 2. The main feature of PRED is extra drop mechanism that drops an additional packet of the same user from the buffer when its packet is dropped by RED. Intuitively, we try to penalize users that do not respond to congestion signals. Note that if there is no penalty associated with dropped packets, users will greedily send new packets as long as they can benefit from increasing their buffer usage.

---

1. When a new packet arrives, apply the RED policy (**regular** drop).
2. If the packet is dropped by RED, preempt (drop) the earliest packet of the same user from the buffer, if any (**extra** drop).

Fig. 2. The Preemptive RED (PRED) drop policy

The next theorem analyzes Nash equilibria imposed by PRED.

Theorem 4. *Under the PRED policy, the price of anarchy in a Nash equilibrium is at most $\frac{B}{B-T_l}$ and there exists a fair Nash equilibrium.*

Proof. First we show how to select the drop probability function so that in a Nash equilibrium the queue length is bounded by T_l. Consider the expected utility of user i after sending a new packet when the queue length is $X \geq T_l$. The buffer usage of user i will increase and decrease by one with probability $1 - p$ and p, respectively. Therefore, the expected utility of user i is

$$(1 - p(X)) \cdot u_i^t(w_i + 1, X + 1) + p(X) \cdot u_i(w_i - 1, X - 1).$$

User i would refrain from sending additional packets if the expected utility is less than its current utility $u_i(w_i, X)$, and we get that $p(X)$ must be greater than

$$\frac{u_i(w_i + 1, X + 1) - u_i(w_i, X)}{u_i(w_i + 1, X + 1) - u_i(w_i - 1, X - 1)}. \tag{2}$$

Suppose that the drop probability function satisfies inequality (2) for each user i and for each $w_i \geq T_l/n$ (e.g. $p \approx 1/2$ satisfies the above requirements).

If $T_l \geq B(n - 1)/n$ then clearly there is a unique Nash equilibrium identical to that imposed by the drop-tail policy. If $T_l < B(n-1)/n$, then we argue that $w^* = T_l/n$ is a fair Nash equilibrium. We have that no user can benefit from increasing its buffer usage above T_l/n by the selection of the drop probability function and no user can benefit from decreasing its buffer usage below T_l/n since in the Nash equilibrium imposed by the drop-tail policy the buffer usage of each user is greater than T_l/n.

We also claim that in a Nash equilibrium the queue length is bounded by T_l. If it is not the case, at least one user has buffer usage greater than T_l/n. However, the selection of the drop probability function implies that its buffer usage will eventually drop to or below T_l/n, which contradicts to the stability of a Nash equilibrium. Therefore, the price of anarchy is at most $\frac{B}{B-T_l}$. ∎

Note that if the queue length exceeds T_l, new users with zero buffer usage would still keep sending new packets since they can only increase their utility. That allows to avoid Denial of Service (DoS) when the buffer is completely monopolized by old users.

We also present a simulation that illustrates the effect of PRED on the queue length. The setting is similar to that of Section 3. The dynamics of the queue length for 8 and 16 users is presented in Figure 3. It turns out that the queue length in a steady state is very close to T_l and has a small variance, as our analysis indeed shows.
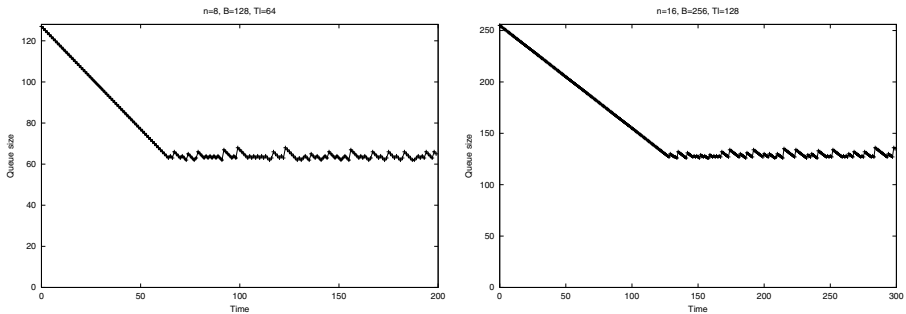
Fig. 3. The effect of PRED on the queue length

5   Concluding Remarks

The Internet is stable today mostly due to the fact that the majority of the users voluntarily use a congestion-responsive TCP protocol. However, some users can benefit from not reducing their transmission rate during congestion. Thus, if users behave selfishly, the assumption about cooperative behavior may not remain valid. Therefore, it is important to understand the nature of congestion resulting from selfish behavior.

We analyze the users' behavior by means of game theory. We consider a single bottleneck buffer under a packetized traffic model, which makes our approach more practical and applicable to real networks such as the Internet. We show that there exist efficient and fair Nash equilibria imposed by a simple FIFO buffering policy. However, the congestion created by the users is rather high. Then we propose a simple modification of RED policy, which decreases the congestion at a Nash equilibrium. Finally, we consider some natural extensions of our model.

We believe that our results can help to shed more light on the stability of the existing Internet infrastructure in presence of selfish users. Some interesting open problems include analysis of routing in general networks, alternative utility functions and interaction of greedy and non-greedy users (e.g. VBR and CBR applications).

References

1. A. Akella, R. Karp, C. Papadimitrou, S. Seshan, and S. Shenker, "Selfish behavior and stability of the Internet: A game-theoretic analysis of TCP," *Proceedings of ACM SIGCOMM'02*.
2. E. Altman, T. Basar, T. Jimenez and N. Shimkin, "Routing into two parallel links: Game-Theoretic Distributed Algorithms," *Special Issue of Journal of Parallel and Distributed Computing on "Routing in Computer and Communication Networks"*, pp. 1367-1381, vol. 61, no. 9, September 2001.
3. D. Bertsekas and R. Gallager, "Data Networks," *Prentice-Hall*, 1987.
4. T. Boulogne, E. Altman and O. Pourtallier, "On the convergence to Nash equilibrium in problems of distributed computing," *Annals of Operation research*, 2002.
5. G. Christodoulou, E. Koutsoupias and A. Nanavati, "Coordination Mechanisms," *Proceedings of ICALP'04*.
6. A. Czumaj, P. Krysta and B. Vocking, "Selfish traffic allocation for server farms," *Proceedings of STOC'02*.
7. A. Czumaj and B. Vocking, "Tight bounds on worst case equilibria," *Proceedings of SODA'02*.
8. C. Douligeris and R. Mazumdar, "On Pareto optimal flow control in an integrated environment," *Proceedings of the 25th Allerton Conference on Communication, Control and Computing*, 1987.
9. D. Dutta, A. Goel and J. Heidemann, "Oblivious AQM and Nash Equilibria," *Proceedings of INFOCOM'03*.
10. E. Even-Dar, A. Kesselman and Y. Mansour, "Convergence Time to Nash Equilibria," *Proceedings of ICALP'03*.
11. R. Feldmann, M. Gairing, T. Lcking, B. Monien and M. Rode, "Nashification and the Coordination Ratio for a Selfish Routing Game," *Proceedings of ICALP'03*.
12. M. Florian and D. Hearn, "Network Equilibrium Models and Algorithms", *Network Routing. Handbooks in RO and MS*, M.O. Ball et al. Editors, Elsevier, pp. 485-550. 1995.

13. S. Floyd and V. Jacobson, "Random Early Detection for Congestion Avoidance," *IEEE/ACM Transactions on Networking*, August 1993.

14. D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis, "The Structure and Complexity of Nash Equilibria for a Selfish Routing Game," *Proceedings of ICALP'02*.

15. D. Fudenberg and D. Levine, "The theory of learning in games," *MIT Press*, 1998.

16. X. Gao, K. Jain, and L. J. Schulman, "Fair and efficient router congestion control," *Proceedings of SODA '04*, pp. 1050-1059.

17. R. Garg, A. Kamra and Varun Khurana, "A Game-Theoretic Approach Towards Congestion Control in Communication Networks," *ACM SIGCOMM Computer Communications Review*, vol. 32, no. 3, pp. 47-61, July 2002.

18. R. Gibbens and F. Kelly, "Resource pricing and the evolution of congestion control," *Automatica*, vol. 35, pp. 1969-1985, 1999.

19. V. Jacobson, "Congestion Avoidance and Control," *Proceedings of ACM SIGCOMM'88*.

20. F. Kelly, A. Maulloo and D. Tan, "Rate control in communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 49, pp. 237-252, 1998.

21. Y. A. Korilis and A. A. Lazar, "On the Existence of Equilibria in Noncooperative Optimal Flow Control," *Journal of the ACM*, vol. 42, pp. 584-613, 1995.

22. E. Koutsoupias, C. H. Papadimitriou, "Worst-case equilibria," *Proceedings of STACS'99*.

23. R. J. La and V. Anantharam, "Optimal Routing Control: Game Theoretic Approach," *Proceedings of the 36rd IEEE Conference on Decision and Control*, pp. 2910-2915, 1997.

24. J. F. Nash, "Non-cooperative games," *Annals of Mathematics*, vol. 54, pp. 286-295, 1951.

25. A. Orda, N. Rom and N. Shimkin, "Competitive routing in multi-user communication networks," *IEEE/ACM Transaction on Networking*, vol 1, pp. 614-627, 1993.

26. R. Pan, B. Prabhakar and K. Psounis, "CHOKe – a stateless active queue management scheme for approximating fair bandwidth allocation," *Proceedings of INFOCOM00*, pp. 942-951.

27. L. Qiu, Y. R. Yang, Y. Zhang, and S. Shenker, "On Selfish Routing in Internet-Like Environments," *Proceedings of ACM SIGCOMM'03*.

28. T. Roughgarden and E. Tardos, "How Bad is Selfish Routing?," *Journal of the ACM*, vol. 49, no. 2, pp. 236-259, 2002.

29. S. Shenker, "Making greed work in networks a game-theoretic analysis of switch service disciplines," *IEEE/ACM Transactions on Networking*, vol. 3, pp. 819-831, 1995.