

A Game-Theoretic Analysis of TCP Vegas

Tuan Anh Trinh, Sándor Molnár *

High Speed Networks Laboratory
Department of Telecommunications and Media Informatics
Budapest University of Technology and Economics
E-mail: {trinh,molnar}@tmit.bme.hu

Abstract. We use the tools from game theory to understand the impacts of the *inherent congestion pricing schemes* in TCP Vegas as well as the problems of parameter setting of TCP Vegas on its performance. It is shown how these inherent pricing schemes result in a rate control equilibrium state that is a Nash equilibrium which is also a global optimum of the all-Vegas networks. On the other hand, if the TCP Vegas' users are assumed to be selfish in terms of setting their desired number of backlogged packets in the buffers along their paths, then the network as a whole, in certain circumstances, would operate *very inefficiently*. This poses a serious threat to the possible deployment of Vegas-based TCP (such as FAST TCP) in the future Internet.

1 Introduction

The Internet has been a huge success since its creation in the early 70's. It has a big impact on the way we interact and communicate. As the Internet evolves, it is shared, used by millions of end-points and many kinds of applications. They compete with each other for the shared resources and their demand for resources (such as bandwidth) is growing rapidly. As a result, congestion at certain points of the network is inevitable. The TCP protocol suite was originally designed to control congestion in the Internet and to protect it from congestion collapse. Basically, TCP is a closed loop control scheme. Congestion in the network is fed back to the source in the form of losses (Reno-like versions) or delay (such as TCP Vegas) The source then reacts to the congestion signal from the network by reducing its transmitting rate. In other words, we can consider packet loss and high queueing delay as the *cost* of (aggressively) sending packets into the network. The higher the rate, the higher the cost (certainly, the relationship is not necessarily linear in nature), given a fix network. Furthermore, as the Internet has been gradually transforming from a government sponsored project to a private enterprise (or even a commodity), the economics of the Internet becomes more and more important issue. Consequently, Internet connectivity

* This work is supported by the Inter-University Center for Telecommunications and Informatics (ETIK)

and services will have to confront issues of pricing and cost recovery. In this perspective, the cost of congestion can be in monetary form. Introducing cost of congestion into the network creates balance, stability and high utilization of resource usage.

From what has been discussed so far, the cost of congestion, in our case, can be either price or delay (the application is delay-sensitive). It is suggested in [13] that congestion pricing could be implemented by using "smart market" where price for sending a packet varies on a very short time scale. Specifically, for the TCP implementation, the congestion price is updated every round-trip time.

A natural question arises then: Why TCP Vegas? There is a number of reasons that motivate us to re-examine TCP Vegas. Firstly, it is because TCP Vegas has *inherent pricing schemes* in its design that resemble the congestion pricing schemes proposed in the literature. We believe that by better understanding TCP Vegas' inherent pricing schemes, we will have a better insight into understanding and designing pricing schemes for TCP traffic in general. Secondly, the emergence of very large bandwidth-delay product networks such as the transatlantic link with a capacity in the range of 1 Gbps - 10 Gbps, new transport protocols have been proposed to better utilize the network in these circumstances. One promising proposal is the FAST TCP [18]. Since the design of FAST TCP is heavily based on the design of TCP Vegas, there is a need to reconsider the benefits as well as the drawbacks of TCP Vegas in order to have an insight into the performance and possible deployment of FAST TCP in the future Internet.

Given the pricing schemes (parameter setting), another natural question arises then: Are these schemes efficient? Is there any equilibrium state from where no one has the incentive to deviate? Game theory (see [4] for a comprehensive introduction) provides us the tools to answer these questions as we will illustrate later in the paper. We use these game-theoretic tools to investigate the impact of the pricing schemes and parameter setting in TCP on the performance of each user as well as the network as a whole.

Regarding the related work, we would like to divide it into two classes. The first class mainly deals with game-theoretic analysis of flow control mechanisms in the Internet. Scott Schenker in his pioneering paper [2] used game-theoretic approach to analyze the flow control mechanisms (for Poisson arrivals) with different queueing disciplines at the routers. Korilis *et al* in [3] also used game-theoretic approach to study the existence of equilibria in noncooperative optimal flow control, especially those with QoS constraints. Recently, Akella *et al* in [1] also used the tools from game-theory to examine the behavior of TCP Reno-like (loss-based) flow controls under selfish parameter setting. Our work is different from their work in the sense that we study delay-based versions of TCP. We provide an extensive analysis of the parameter setting problem of the traditional TCP Vegas, the modified version of TCP Vegas (TCP Vegas under REM) as well as FAST TCP. The second class deals with the mechanisms and issues of congestion pricing in the Internet. We would mention here the work of MacKie-Mason *et al* [13], [14]. These papers apply economic theory ("club theory") to

study basic issues of congestion pricing in the Internet. Incentive-compatible pricing strategies in noncooperative networks are introduced and analyzed in [8]. A survey on Internet pricing and charging in general can be found in [6]. These papers deal with the general pricing problems. Our paper, on the other hand, deals specifically, from game-theoretic point of view, with TCP Vegas and its variants.

The main contributions of the paper are the followings. First, we prove that, under the inherent congestion pricing schemes in TCP Vegas, there exists a unique Nash equilibrium of the rates of the TCP Vegas flows sharing a common network and this equilibrium rate vector is also system-wide optimal. Secondly, we provide an extensive game-theoretic analysis of the parameter setting problem in TCP Vegas. We conclude that, in this case, the Nash equilibria (if any), can be very inefficient. This implies that all-Vegas networks are vulnerable to selfish action of end-users posing a serious threat to the possible deployment of all-Vegas based (such as FAST TCP) in the future Internet.

The rest of the paper is organized as follows. The background on TCP is provided in Section 2. The TCP Vegas games are described and analyzed in detail in Section 3. Finally, Section 4 concludes the paper.

2 Background

2.1 TCP Vegas

TCP Vegas was first introduced by Brakmo *et al* in [16]. Basically, it is a delay-based congestion control scheme that uses both queueing delay and packet loss as congestion signal. TCP Vegas tries to control the number of packets buffered along the path with the targeted number to be between α and β ($\alpha \leq \beta$). Let $w(t)$ denote the congestion window at time t , RTT denote the round-trip time and $baseRTT$ is the smallest value of the round-trip time so far (actually, this is an *estimate* of the propagation delay). Denote $diff = \frac{RTT - baseRTT}{RTT}w$, then the dynamics of the congestion window of TCP Vegas can be expressed as follows:

$$w(t+1) = \begin{cases} w(t) + 1 & \text{if } diff < \alpha, \\ w(t) - 1 & \text{if } diff > \beta, \\ w(t) & \text{otherwise.} \end{cases} \quad (1)$$

In a TCP Vegas/REM network [15], a slight modification is introduced into the updating mechanism of the congestion window. Each link l (with capacity c_l) update the *link price* $p_l(t)$ in period t based on the aggregate input rate $x^l(t)$ and the buffer occupancy $b_l(t)$ as follows:

$$p_l(t+1) = [p_l(t) + \gamma(\mu_l b_l(t) + x^l(t) - c_l)]^+ \quad (2)$$

where $0 < \gamma$ and $0 < \mu_l < 1$ are scaling factors of REM. Each source will estimate the *total price* along its path and update its sending rate accordingly. To feed back the prices to sources, link l marks each arriving packet in period

t , that is not already marked at an upcoming stream, with probability $m_i(t)$ defined as:

$$m_i(t) = 1 - \varphi^{-p_i(t)}$$

where $\varphi > 0$ is a constant. Once a packet is marked, its mark is carried to the destination and then conveyed back to the source via acknowledgement, like ECN scheme. The source i estimates the end-to-end marking probability by the fraction $\hat{m}_i(t)$ of its packets marked in period t , and estimates the path price $p_i(t)$ by:

$$\hat{p}_i(t) = -\log_{\varphi}(1 - \hat{m}_i(t))$$

The dynamics of the congestion window of TCP Vegas/REM can be expressed as follows:

$$w_i(t+1) = \begin{cases} w_i(t) + 1 & \text{if } -\frac{w_i(t)}{RTT_i(t)} < \frac{\alpha}{\hat{p}_i(t)}, \\ w_i(t) - 1 & \text{if } -\frac{w_i(t)}{RTT_i(t)} > \frac{\alpha}{\hat{p}_i(t)}, \\ w_i(t) & \text{otherwise.} \end{cases} \quad (3)$$

2.2 Throughput models of TCP Vegas

Throughout the paper our game-theoretic analysis uses the models that are previously derived. These are:

Model 1: (Thomas Bonald's model)

In [17], the throughput of multiple flows sharing a bottleneck link is analyzed (by using fluid approximation) both for TCP Reno and TCP Vegas case. Assume N TCP flows sharing a bottleneck link with capacity μ , propagation delay τ and buffer size B . The parameters of TCP are: α and β . The main results in their paper that we use in our analysis are the following:

- If $N\alpha < B$, there exists a finite time from which no loss occurs. In addition, the window size stabilizes in finite time. If $\alpha \neq \beta$, the congestion windows converge not to a single point (but a region). This implies unfairness among flows even in equilibrium. If $\alpha = \beta$, then $w_1 = w_2 = \dots = w_N = \frac{\mu\tau}{N} + \alpha$ and the average rate $\lambda_1 = \lambda_2 = \dots = \lambda_N = \frac{\mu}{N}$. Note that in this case the link is fully utilized.
- If $N\alpha \geq B$, then TCP Vegas behaves exactly like TCP Reno. Let $\omega = \frac{\mu\tau}{B}$ and γ is the multiplicative decrease of TCP Reno (typically $\frac{1}{2}$). If $\omega \geq \frac{\gamma}{1-\gamma}$ then $\lambda_{total} = \frac{(1-\gamma^2)(\omega+1)^2}{2(1-\gamma)(\omega^2+\omega)+1}\mu < \mu$. This implies that in this case the link is not fully utilized.

Model 2: (Steven Low's model)

Steven Low *et al* in [10], [11], [12], [15] described an optimization framework to study the performance of the TCP Vegas in a general network topology and under different queue management schemes at the routers. We would mention the result regarding the throughput of TCP Vegas under REM queue management scheme (TCP Vegas/REM) that we will use in our analysis later in this paper. It is proved in [11] that the equilibrium rate of TCP Vegas can be calculated as: $\lambda_i = \frac{\alpha_i}{p^*}$, where p^* denotes the equilibrium price. Note that this result is true for a general network topology (not restricted to a single bottleneck link).

3 The TCP Vegas games

In this Section, the games regarding the inherent pricing schemes (for rate allocation) and parameter setting of TCP Vegas are described and analyzed in detail. We also investigate the impact of the results on the performance of TCP Vegas and on the network as a whole.

3.1 Game 1: Rate allocation of TCP Vegas

We consider a network that consists of a set $\mathcal{L} = \{1, 2, \dots, L\}$ of links with capacity $c_l, l \in \mathcal{L}$. Assume that the network is shared by a set of flows (sources). The set of flows is denoted by $\mathcal{N} = \{1, 2, \dots, N\}$. The rate of flow i is denoted by $x_i, i \in \mathcal{N}$. Flow i uses a subset (\mathcal{L}_i) of \mathcal{L} in its path ($\mathcal{L}_i \subseteq \mathcal{L}$). Let us define the routing matrix as follows:

$$\mathbf{R}_{li} = \begin{cases} 1 & \text{if } l \in \mathcal{L}_i, \\ 0 & \text{otherwise.} \end{cases}$$

The physical capacity constraints of the flows therefore can be defined as :

$$\mathbf{R}\mathbf{x} \leq \mathbf{c} \tag{4}$$

where $\mathbf{x} = (x_1, x_2, \dots, x_N)$ is the flow rate vector and $\mathbf{c} = (c_1, c_2, \dots, c_L)$ is the link capacity vector. In addition, flow rates cannot be negative:

$$x_i \geq 0, \quad i = 1, 2, \dots, N \tag{5}$$

The set of flow rate vectors Λ that satisfy both conditions 4 and 5 is called a feasible set.

It should be mentioned that our network, as TCP network in general, assumes feedback-based flow control. The feedback can be implicit (e.g. queueing delay) or explicit (by pricing and/or using Explicit Congestion Notification (ECN)). The sources (end-points) use Vegas-style flow control, as defined in [16], [11]. We consider the flows as the players of the game. The strategy space for a player is the range of it sending rate.

Let us define the following generic payoff function for each player

$$B_i(x_i) = \alpha_i \log(x_i) - \sum_{l \in \mathcal{L}_i} \int_0^{x_i} \pi_l(y) dy \tag{6}$$

and $\pi_l = p_l(\sum_{k \in \mathcal{L}_k} x_k)$ is defined as the function of the total flow rates on link l . This function is actually the *price* that is fed back to the player i sending at rate x_i , which is an increasing function. The higher the rate, the higher the price. Hence, the second term in Equation 6 can be interpreted as the bandwidth *cost* fed back to player i when it attempts to transmit at rate x_i . The first term in Equation 6 reflects the *gain* of player i when transmitting at rate x_i , [12] (note that this is a concave function of x_i). As a result, the payoff $B_i(x_i)$

represents the *net* benefit of player i when transmitting at rate x_i . The price (cost) can be communicated to the end-user (the player) by the mean of the total queueing delay of its packets in the path, as in TCP Vegas/Drop-Tail network. The price can also be communicated explicitly to the user by using REM active queue management scheme (with ECN) and here we have a Vegas/REM network. In the first case, we would like to mention that we *implicitly* use the PASTA property for Poisson arrivals with FIFO scheduling principle to derive the proportional relationship between the total rate arrive at the link and the queueing delay. We can also suggest here the Little's formula for this relationship. This is assumed frequently in the literature with or without any mention. In any case, if the aggregate arrival flow is not Poisson (e.g. self-similar traffic), then queue length (queueing delay) is generally larger than the Poisson one. Furthermore, the expression of queueing delay in our model is assumed to be *additive* among links. This is true for a Norton network with Poisson arrivals. So, strictly speaking, our analysis can be considered as a *worst case* analysis for TCP Vegas/Drop-Tail network. For TCP Vegas/REM network, the additive assumption is justified when the mark rates are small. Indeed, let $\pi_l(t)$ be the marking probability at link l at time t and the end-to-end marking probability $q_i(t)$ that the end-point i observes (and to which source algorithm reacts). For small $\pi_l(t)$, $q_i(t) = 1 - \prod_{l \in \mathcal{L}_i} (1 - \pi_l(t)) \approx \sum_{l \in \mathcal{L}_i} \pi_l(t)$.

Under the assumptions mentioned above, our problem can be modelled as a non-cooperative game. The strategy space for a player is its sending rate and is determined by the capacity of the links. The strategy for player i can be defined as $S^{(i)} = \{x_i | 0 < x_i \leq c_{max}^i\}$, where $c_{max}^i = \max\{c_l | l \in \mathcal{L}_i\}$. The strategy space for the game is defined as the Cartesian product $S = \bigotimes_{i=1}^N S^{(i)}$, which is equivalent to the feasible set Λ . Strategy $\mathbf{x} = (x_1, x_2, \dots, x_N) \in \Lambda$ is called a strategy profile. Each player (e.g. player i) chooses the sending rate (x_i) in the feasible set in order to maximize its *own* payoff function $B_i(x_i)$ in a selfish way. By "in a selfish way" we mean that the player does not care about other players' payoff, as far as the rate vector is in the feasible set.

One of the key questions in a non-cooperative flow control game in general, and our game in particular, is whether the network converges to (or settles at) an equilibrium point, such that no player can increase its payoff by adjusting its strategy unilaterally. In the game-theory terminology such a point is called a *Nash equilibrium*. The Nash equilibrium in our game also reflects the *balance* of the gain and the cost for each player as well as for the network as a whole. A non-cooperative game may have no Nash equilibrium (in its pure strategy space), multiple equilibria, or a unique equilibrium. As for the TCP Vegas game, we can prove the following theorem:

Theorem 1. *There exists a unique Nash equilibrium (in its pure strategy space) for the TCP Vegas game described above.*

We follow the proof methodologies provided in [5] as well as in a recent paper [9].

Proof. First, let's consider the existence of the Nash equilibrium for the TCP Vegas game. Notice that the feasible set $\Lambda = \{\mathbf{x} | \mathbf{R}\mathbf{x} \leq \mathbf{c}, \mathbf{x} \geq \mathbf{0}\}$ is a nonempty,

convex and compact set. It is nonempty because $\mathbf{x} = (\epsilon, \epsilon, \dots, \epsilon) \in A$, where $0 < \epsilon < \frac{c_{min}}{N}$, $c_{min} = \mathbf{min}\{c_l | l \in \mathcal{L}\}$. It is bounded because $x_i \leq c_{max}, i \in \mathcal{N}$, where $c_{max} = \mathbf{max}\{c_l | l \in \mathcal{L}\}$. Assume that $\mathbf{x}_1, \mathbf{x}_2 \in A$ and $0 < \rho < 1$, we have:

$$\rho \mathbf{x}_1 + (1 - \rho) \mathbf{x}_2 \leq \mathbf{R}(\rho \mathbf{x}_1 + (1 - \rho) \mathbf{x}_2) \leq \mathbf{c}$$

This result implies the convexity of A .

Now let's consider the payoff functions of the players. Notice that $B_i(x_i)$ is a concave function of x_i . Indeed:

$$B_i''(x_i) = -\frac{\alpha_i}{x_i^2} - \sum_{l \in \mathcal{L}_i} \pi_l' < 0 \quad (7)$$

From what have been discussed so far, our game has the following properties:

1. The joint strategy space is nonempty, convex and compact.
2. The payoff function of each player is concave in its own strategy space.

According to Theorem 1 in [5], there exists a Nash equilibrium in its pure strategy space.

For the uniqueness of the Nash equilibrium, let's consider the (nonnegative) weighted sum of the payoff functions:

$$\sigma(\mathbf{x}, \mathbf{w}) = \sum_{i=1}^N w_i B_i(\mathbf{x}), \quad w_i \geq 0 \quad (8)$$

Denote $g(\mathbf{x}, \mathbf{w})$ the pseudo-gradient of $\sigma(\mathbf{x}, \mathbf{w})$, then the Jacobian of $g(\mathbf{x}, \mathbf{w})$ with respect to \mathbf{x} can be computed as follows:

$$\mathbf{G} = \begin{pmatrix} B_{11} & B_{12} & \dots & B_{1N} \\ B_{21} & B_{22} & \dots & B_{2N} \\ \vdots & \ddots & \ddots & \vdots \\ B_{N1} & B_{N2} & \dots & B_{NN} \end{pmatrix}$$

where

$$B_{ij} = \begin{cases} w_i \left(-\frac{\alpha_i}{x_i^2} - \sum_{l \in \mathcal{L}_i} \pi_l' \right) < 0 & j = i \\ -w_i \sum_{l \in \mathcal{L}_{(i,j)}} \frac{\partial \pi_l}{\partial x_j} < 0 & j \neq i, \mathcal{L}_{(i,j)} \neq \emptyset \\ 0 & j \neq i, \mathcal{L}_{(i,j)} \equiv \emptyset. \end{cases}$$

where $\mathcal{L}_{(i,j)} = \mathcal{L}_i \cap \mathcal{L}_j$. The matrix \mathbf{G} defined above is thus negative definite. As a result, according to Theorem 6 in [5], $\sigma(\mathbf{x}, \mathbf{w})$ is diagonally strictly concave. According to Theorem 2 in [5], the equilibrium point of the TCP Vegas game is unique.

Remark 1. To reach this equilibrium, [5] shows that each player can change its own strategy at a rate proportional to the gradient of its payoff function with respect to its strategy and subject to constraints. This method is in fact equivalent to the gradient projection algorithms described in [10].

Remark 2. The authors of [10], using optimization framework, also showed that, under certain assumptions on the step size, these algorithms converge to a system wide optimal point (which is also proved to be unique). Furthermore, it is proved in [11], [12] that the rate control of TCP Vegas/Drop Tail and TCP Vegas/REM is indeed based on these algorithms. This implies that *the TCP Vegas game described above converges to a unique Nash equilibrium that is system wide optimal.*

3.2 Game 2: Parameter Setting of TCP Vegas

In this game, we consider the parameter setting of TCP Vegas. As described in [16], TCP Vegas tries to maintain the number of backlogged packets in the network between α and β . We examine here the situation when a selfish (and greedy) user tries to increase the number of its backlogged packets in the network in order to grab more bandwidth in the network. If all other players do the same thing (i.e. they are also selfish and greedy), the total number of packets in the network would increase without bound. However, the size of the buffers at routers are bounded and packet loss would occur, reducing the throughput of the connection. We are interested in a situation (i.e. a parameter setting, if at all exists) from where no player would deviate.

We consider a simple topology of N TCP Vegas sources sharing a *single* bottleneck link with a buffer size of B packets. Source i is associated with a set (α_i, β_i) . In this paper, we deal with the case when $\alpha_i = \beta_i$. The case when $\alpha_i \neq \beta_i$ is left for future work.

Players: N TCP Vegas flows

Actions: Each player can set its parameter (α_i) in order to control the number of its backlogged packets in the queue of the bottleneck link (with capacity μ and delay τ). The router is assumed to use Drop-Tail mechanism (FIFO principle)

Payoff: $f(\alpha_i) = \lambda_i$ (the average throughput)

If the total number of backlogged packets is smaller than the buffer size at the bottleneck router (i.e. $\sum_{j=1}^N \alpha_j < B$) then the payoff function of player i can be expressed as follows:

$$f(\alpha_i) = \lambda_i = \frac{\alpha_i}{\frac{\sum_{j=1}^N \alpha_j}{\mu}} = \frac{\mu \alpha_i}{\sum_{j=1}^N \alpha_j} = \frac{\mu \alpha_i}{\alpha_i + \sum_{j \neq i} \alpha_j} \quad (9)$$

From Equation 9 we have:

$$\frac{\partial f}{\partial \alpha_i} = \frac{\mu \sum_{j \neq i} \alpha_j}{(\alpha_i + \sum_{j \neq i} \alpha_j)^2} > 0, \quad i = \overline{1 \dots N} \quad (10)$$

Since $\sum_{j \neq i} \alpha_j$ is always positive, it follows from Equation 9 that $\frac{\partial f}{\partial \alpha_i} > 0, \forall i$. This implies that given other players' strategies, player i will set α_i as high as possible in order to maximize its payoff. Notice that Equation 9 is valid only

if $\sum_{j=1}^N \alpha_j < B$. Otherwise, TCP Vegas, according to [17], behaves exactly like TCP Reno. In this case, there are two possibilities [17]:

$$\lambda_i^{Reno} = \begin{cases} \frac{(1-\gamma^2)(\omega+1)^2}{2(1-\gamma)(\omega^2+\omega)+1} \frac{\mu}{N} < \frac{\mu}{N} & \text{if } \omega \geq \frac{\gamma}{1-\gamma} \\ \frac{\mu}{N} & \text{otherwise.} \end{cases} \quad (11)$$

Thus, we have two cases:

Case 1: $w < \frac{\gamma}{1-\gamma}$

It is important to note that in this case, the link is fully utilized both for TCP Vegas and TCP Reno. Furthermore, in TCP Reno style performance, the bandwidth is fairly (equally) shared between flows (because they have the same RTT). Denote $\alpha^* = (\alpha_1^*, \alpha_2^*, \dots, \alpha_N^*)$ be the Nash equilibrium of the game in this case. Without losing generality, we can assume that $\alpha_1^* \leq \alpha_2^* \leq \dots \leq \alpha_N^*$. Notice that in Nash equilibrium, we must have $\alpha_1^* = \alpha_2^* = \dots = \alpha_N^*$. Otherwise, player 1 has the incentive to deviate (i.e. to increase its number of backlogged packets - α_1) in order to get higher throughput, because in Reno style performance, it would get a fairer share of the total bandwidth (i.e. $\frac{\mu}{N}$). As a result, we have the Nash equilibria for this game: $\alpha^* = (\alpha_1^*, \dots, \alpha_N^*)$ where $\alpha_i^* \geq \lfloor \frac{B}{N} \rfloor$, $\forall i$. This means that, in this case, in Nash equilibrium, the parameter α can be arbitrarily large.

Case 2: $w \geq \frac{\gamma}{1-\gamma}$

In this case, the link is not fully utilized. Following similar reasoning as in Case 1, we have a set of Nash equilibria defined as follows: $\Omega = \{\alpha = (\alpha_1, \dots, \alpha_N) \mid \alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_N\}$ with the conditions that $\sum_{i=1}^N \alpha_i = B-1$ and $\alpha_1 \geq \frac{(1-\gamma^2)(\omega+1)^2}{2(1-\gamma)(\omega^2+\omega)+1} \frac{B-1}{N}$. The latter expression simply means that even player 1 (who gets the smallest bandwidth) would not deviate, so no other player would deviate. If this condition does not hold, player 1 would deviate to get higher bandwidth share.

Our final comment on this unique Nash equilibrium is that each TCP Vegas flow (player) maintains the number of its own backlogged packets as many as possible. As a result, the buffer is nearly full and the queueing delay is unnecessarily high. A nearly full buffer may cause many difficulties for TCP Vegas (e.g. the estimation of *baseRTT* might be inaccurate if there are already many packets in the queue when the connection starts)

4 Conclusion

We have demonstrated, by using game-theoretic approach, how TCP Vegas' inherent pricing schemes as well as the parameter setting impact on its performance. Our analysis shows that these inherent pricing schemes result in a rate control equilibrium state that is a Nash equilibrium in game-theoretic terms which is also a global optimum of the all-Vegas networks. We also proved that the parameter setting of TCP Vegas is very vulnerable to selfish actions of the users. This poses a serious threat to the possible deployment of FAST TCP in the future Internet.

References

1. A. Akella, R. Karp, C. Papadimitrou, S. Seshan, and S. Schenker, *Selfish behavior and stability of the Internet: A game-theoretic analysis of TCP* ACM SIGCOMM, 2002.
2. Scott Schenker, *Making Greed Work in Networks: A game-theoretic analysis of switch service disciplines*, IEEE/ACM Transactions on Networking, vol. 3, 1995.
3. Y. A. Korilis and A. A. Lazar, *On the existence of equilibria in noncooperative optimal flow control*, Journal of the ACM, vol. 42, no 3 pp. 584-613, 1995.
4. M. J. Osborne and A. Rubenstein, *A course in game theory*, Cambridge, Massachusetts: The MIT Press, 1994.
5. J. B. Rosen, *Existence and uniqueness of equilibrium points for concave n-person games*, Econometrica, vol. 33, pp. 520-534, Jul. 1965.
6. P. Reichl, B. Stiller, *Pricing model for Internet services*, <http://www.tik.ee.ethz.ch/cati>
7. F. Kelly et al, *Rate control for communication networks: Shadow prices, proportional fairness and stability*, Journal of Operation Research Society, vol 49, no. 3, pp. 237-252, March 1998.
8. Y. Korilis et al, *Incentive-compatible pricing strategies in noncooperative networks*, in Proc. of IEEE Infocom 1998.
9. T. Alpcan, T. Basar, *Distributed algorithms for Nash equilibria of flow control games*, to appear in Annals of Dynamic Games, 2004
10. S. Low and D. Lapsley, *Optimization flow control, I: basic algorithm and convergence*, IEEE/ACM Transactions on Networking, 7(6):861-874, December 1999.
11. S. Low, L. Peterson, and L. Wang, *Understanding Vegas: a duality model*, Journal of ACM, 49(2):207-235, March 2002.
12. S. Low, *A duality model of TCP and queue management algorithms*, IEEE/ACM Transactions on Networking, October 2003.
13. J. MacKie-Mason, H. Varian, *Pricing the Internet*, in B. Kahin and J. Keller, eds., Public Access to the Internet, MIT Press, Cambridge, MA, 1995.
14. J. MacKie-Mason, H. Varian, *Pricing congestible network resources*, Journal on Selected Areas of Communications, pp. 1141-1149, 13(1995).
15. S. Athuraliya, V. Li, S. Low, Q. Yin, *REM: active queue management*, IEEE Network, June 2001.
16. L. Brakmo, S. O'Malley, and L. Peterson, *TCP Vegas: new techniques for congestion detection and avoidance*, IEEE/ACM SIGCOMM 94, London, UK, Sept. 1994.
17. T. Bonald, *Comparison of TCP Reno and TCP Vegas*, Workshop on the modeling of TCP, 1998.
18. Cheng Jin, David X. Wei and Steven H. Low *FAST TCP: motivation, architecture, algorithms, performance*, IEEE Infocom, March 2004.