Fairness comparison of FAST TCP and TCP Reno

Liansheng Tan ^a, Lina Dong ^b, Cao Yuan ^b, Moshe Zukerman ^{c,*}

^a *The Research School of Information Sciences and Engineering, The Australian National University, Canberra, ACT 0200, Australia*

^b *Department of Computer Science, Central China Normal University, Wuhan 430079, PR China*

^c *The ARC Special Research Center for Ultra-Broadband Information Networks (CUBIN), EEE Department, The University of Melbourne, Vic. 3010, Australia*

Received 25 May 2006; received in revised form 24 December 2006; accepted 29 December 2006

Available online 9 January 2007

Abstract

In networking, it is often required to quantify by how much one protocol is fairer than another and how certain parameter setting and/or protocol enhancements improve fairness. This paper provides a framework to evaluate the fairness of various protocols in a general telecommunications network. Within this framework, there are two key components: (1) a benchmark and (2) a single dimension metric. We suggest to use the max–min fairness bandwidth allocation as the benchmark and the Euclidean distance between any bandwidth allocation under any protocol and the max–min bandwidth allocation as the metric. Explicitly, we provide a method to compare the fairness of two sets of bandwidth allocation under two different protocols for a given network by using this metric. On the basis of this new framework, we evaluate the fairness of FAST TCP and TCP Reno relative to the max–min fairness criteria. The distance between the max–min fair allocation and allocations based on each of the two protocols is measured using the Euclidean norm. We derive explicit expressions for these distances for a general network and compare the fairness of these two protocols by using their corresponding utility functions. Finally, we numerically demonstrate how this method can be applied to compare the fairness of FAST TCP and TCP Reno for a “Parking Lot” linear network and for the NSFNET Backbone network. In addition to merely a comparison between protocols, such numerical results can provide guidelines for better choice of parameters to make a protocol fairer in a given scenario. © 2007 Elsevier B.V. All rights reserved.

Keywords: FAST TCP; TCP Reno; Max–min fairness; Fairness comparison metric; Utility function

1. Introduction

In the last several decades, fairness has been considered an important criterion in design of computer and telecommunications networks. Its importance have given rise to many discussions in the networking literature on how to define and provide fairness [1–19]. Given the importance of fairness in networking, it is often required to quantify by how much one protocol is fairer than another and by how much certain parameter setting and/or protocol enhancements improve fairness.

The current TCP congestion control algorithm, TCP Reno and its variants, have experienced scalability problems [20,21]. FAST TCP [21] has recently gained popularity as a viable TCP alternative and is considered promising in terms of system stability, throughput, responsiveness and fairness [22–24]. The question of fairness comparing FAST TCP and TCP Reno has therefore been an important and popular research topic [25–31].

In [25], the authors investigate goodput, fairness and friendliness properties presented in Westwood+, New Reno and Vegas TCP by using both NS-2 simulations and live Internet measurements. The work in [26] studies the equilibrium and fairness properties of networks where both TCP Reno and FAST TCP are used. In particular, they show that the fairness property of FAST TCP and TCP Reno is preserved “locally” within each protocol

* Corresponding author. Tel.: +61398879028.

E-mail addresses: liansheng.tan@rsise.anu.edu.au (L. Tan), donglina.ccn@gmail.com (L. Dong), yuancao1980@gmail.com (C. Yuan), mzu@unimelb.edu.au (M. Zukerman).

when TCP Reno flows and FAST TCP flows share the same network bandwidth. In [27], the authors compare the fairness of FAST TCP and TCP Vegas for a single-bottleneck link model in the persistent congestion situation by using a queue equilibrium formulation proposed in [28,29]. For a single bottleneck link, the earlier works by [30] analyzes the performance of TCP Vegas in comparison with TCP Reno in terms of fairness using a simple closed fluid model. In [31], the authors investigate the fairness issue in the situation where TCP Reno and TCP Vegas flows share a single bottleneck link. However, their results are only related to the fairness of networks where heterogeneous protocols co-exist.

This paper provides a framework to evaluate the fairness of various protocols in a general telecommunications network. We promote the use of a “measure” of fairness because it is important to be able to quantify by how much one protocol is fairer than another and how certain parameter setting and/or protocol enhancements improve fairness. Such a measure requires (1) a benchmark and (2) a single dimension metric. We suggest to use the max–min fairness bandwidth allocation as the benchmark and the Euclidean distance between any bandwidth allocation under any protocol and the max–min bandwidth allocation as the metric. Explicitly, we provide explanations on how to compare the fairness of two sets of bandwidth allocation under two different protocols for a given network by using this metric. We illustrate the implication and intuition of this metric by using simple examples. These examples also show when the bandwidth allocation satisfies the max–min fairness, the Euclidean distance is minimized and equal to zero indicating the point of maximum fairness. We apply these new concepts of benchmark and measure to compare between FAST TCP and TCP Reno both for a “Parking Lot” network and the NSFNET backbone network. We also demonstrate the importance of parameter setting by showing that one protocol can be made closer to max–min than the other by choosing better parameter settings.

We acknowledge that our particular choices of measure and metric may be controversial. We therefore first clarify that the choice of max–min fairness as benchmark is without loss of generality, and the results of this paper are readily applied to any other fairness benchmark. The choice of the Euclidean norm may remain controversial as the measure of fairness depends on users’ utility functions. Nevertheless, the Euclidean distance function measures the ‘as-the-crow-flies’ distance, which means the shortest distance between two points. It has been used in [16,32] to evaluate the closeness of two rate allocations, it is consistent with the fairness index proposed by Jain et al. [1] for the special case of a single bottleneck, and extends the fairness index of [1] to a general network with multiple bottleneck links.

The remainder of the paper is organized as follows. In Section 2, we define the new metric to compare the deviation of the two protocols from max–min. In Section 3, based on the optimal models of FAST TCP and TCP Reno, we derive their deviation from max–min using the

metric. In Section 4, we numerically compare FAST TCP and TCP Reno relative to max–min. Finally, the conclusions are made in Section 5.

2. The metric

A series of recent works, e.g. [16–19,21,22,26,27,32–43] have considered the bandwidth allocation problem as an optimization problem aiming to find an allocation policy that maximizes the sum of all sources utilities subject to link capacity constraints.

Consider a network with l links shared by n flows. Link j ($j = 1, 2, \dots, l$) has capacity c_j , and x_i is the rate of flow i ($i = 1, 2, \dots, n$). We define the routing matrix $R = (R_{ij})$, $i = 1, \dots, n, j = 1, \dots, l$ where $R_{ij} = 1$ if flow i use the link j and $R_{ij} = 0$, otherwise. Consider the following utility optimization problem:

Problem P

$$\begin{cases} \max_{x \geq 0} \sum_i U_i(x_i) \\ \text{subject to } Rx \leq c, \end{cases} \quad (1)$$

where U_i is the utility function of flow i and c is the capacity vector.

In [17], the following class of utility functions is proposed

$$U(x, \alpha) = \begin{cases} (1 - \alpha)x_i^{1-\alpha} & \text{if } \alpha \neq 1 \\ \log x_i & \text{if } \alpha = 1, \end{cases}$$

where $\alpha \geq 0$. This includes four special cases: maximum throughput ($\alpha = 0$), proportional fairness ($\alpha = 1$), minimum potential delay ($\alpha = 2$) and max–min fairness ($\alpha = \infty$) [18,19,40]. The papers [19,40] point out that one allocation policy with larger α is fairer. However, an exact measure of fairness is not considered there.

Jain et al. [1] have proposed a fairness index applicable to any system where all users are sharing the same resource. In other words, it can be used for a quantitative evaluation of the fairness of an allocation policy in a single-bottleneck link network. In this paper, we apply the fairness index of [1] to a general network by comparing the fairness of any two protocols with concave utility functions by considering their deviations from a pre-chosen rate allocation benchmark.

In this paper, the rate allocation benchmark we choose is max–min fairness [2]. The choice is without loss of generality as the method we use can apply to any other rate allocation benchmark. The choice of max–min fairness is based on its popularity and importance in computer networks. Note that for a general network, the max–min fair rate allocation can be easily obtained using the well-known progressive filling algorithm [16,32].

Definition 1. Given a general network, for two protocols (λ and μ) with different utility functions, their rate allocations are solutions to the optimization problem P, which are denoted by $X^\lambda = \{x_i^\lambda, i = 1, \dots, n\}$, and $X^\mu = \{x_i^\mu, i = 1, \dots, n\}$, respectively. The max–min rate allocation

tion for this network is $X = \{x_i, i = 1, \dots, n\}$. If $\|X^\lambda - X\|_2 < \|X^\mu - X\|_2$, we say that the protocol λ is closer to max-min than protocol μ . The above norm is the standard Euclidean norm defined by

$$\|X^\mu\|_2 = \sqrt{\sum_{i=1}^n (x_i^\mu)^2}. \quad (2)$$

Here, we use Euclidean norm to measure the deviation from the rate allocation benchmark because Euclidean norm reflects the distance between any two vectors. It has been used to evaluate the closeness of two rate allocations in [16,32].

To illustrate this concept, let us consider the following two simple networks. The first network is shown in Fig. 1. It comprises only a single link with two flows passing through it. The capacity of this link is assumed to be one unit. Clearly, the max-min fair bandwidth allocation for this network is $(1/2, 1/2)$. For any bandwidth allocation of this network under any protocol, denoted by (x_1, x_2) , we must have $x_1 + x_2 = 1$. The standard Euclidean norm, that measures the distance of the bandwidth allocation (x_1, x_2) from the max-min bandwidth allocation is

$$\begin{aligned} Eu &= \sqrt{(x_1 - 1/2)^2 + (x_2 - 1/2)^2} = \sqrt{2(x_1 - 1/2)^2} \\ &= \sqrt{2} |x_1 - 1/2|. \end{aligned}$$

In Fig. 2, we plot this standard Euclidean norm as a function of x_1 . Fig. 2 illustrates that the Euclidean distance is consistent with our notion of fairness, obtains higher values as x_1 is further from its max-min fair value at $x_1 = 1/2$, and that this distance is minimized when $x_1 = 1/2$, i.e., when the bandwidth allocation is of max-min fair.

The second network is described by Fig. 3. In this linear network, there are three flows passing three links. The capacities on each of these three links is assumed to be one unit. One easily finds that the max-min fair bandwidth allocation for this network is $(1/3, 1/3, 1/3)$. Let us for simplicity consider the set of rate allocations, denoted by (x_1, x_2, x_3) , that obey the following two conditions: (1) $x_1 = x_2$; and (2) Link 2 is fully utilized. From the second assumption, we have $x_1 + x_2 + x_3 = 1$. Subsequently, the standard Euclidean norm, that measures the distance of the bandwidth allocation (x_1, x_2, x_3) from the max-min bandwidth allocation should be

Fig. 1. The single link two-flows scenario.

Fig. 2. The standard Euclidean norm Eu as a function of x_1 .

Fig. 3. A case of three links and three flows.

$$\begin{aligned} Eu &= \sqrt{(x_1 - 1/3)^2 + (x_2 - 1/3)^2 + (x_3 - 1/3)^2} \\ &= \sqrt{2(x_1 - 1/3)^2 + (1 - 2x_1 - 1/3)^2} \\ &= \sqrt{6} |x_1 - 1/3|. \end{aligned}$$

We plot the above standard Euclidean norm as a function of x_1 in Fig. 4. Again, Fig. 4 shows that the standard Euclidean distance increases as x_1 is further away from its max-min fair value at $x_1 = 1/3$, and that this distance

Fig. 4. The standard Euclidean norm Eu as a function of x_1 .

is minimized when $x_1 = 1/3$, i.e., when the bandwidth allocation is max–min fair.

When there is only a single bottleneck link, the fairness comparison result of two rate allocations by using our Euclidean distance metric is the same as the one obtained using the fairness index of [1]. We clarify this issue in the following. Assume that there is one single bottleneck link; the capacity of bottleneck link is C . We do two separate runs for two different protocols μ and λ . In the first run $n \lambda$ flows pass through the bottleneck link and in the second $n \mu$ flows pass through it. Let the sending rate of each λ flow be denoted by x_i^λ , where $i = 1, 2, \dots, n$; and the sending rate of each μ flow is x_i^μ , where $i = 1, 2, \dots, n$.

The max–min rate allocation for this network is x_i , where $i = 1, 2, \dots, n$. If the protocol λ is fairer than μ , using the fairness index, then

$$\frac{\left(\sum_{i=1}^n x_i^\lambda\right)^2}{n \sum_{i=1}^n (x_i^\lambda)^2} > \frac{\left(\sum_{i=1}^n x_i^\mu\right)^2}{n \sum_{i=1}^n (x_i^\mu)^2}. \quad (3)$$

For the bottleneck link we have

$$\sum_{i=1}^n x_i^\lambda = \sum_{i=1}^n x_i^\mu = C. \quad (4)$$

The inequality (3) is equivalent to

$$\sum_{i=1}^n (x_i^\lambda)^2 < \sum_{i=1}^n (x_i^\mu)^2. \quad (5)$$

Using the new metric proposed by Definition 1, that is

$$\sqrt{\sum_{i=1}^n (x_i^\lambda - x_i)^2} < \sqrt{\sum_{i=1}^n (x_i^\mu - x_i)^2}, \quad (6)$$

we obtain

$$\sum_{i=1}^n (x_i^\lambda)^2 - 2 \sum_{i=1}^n x_i^\lambda \cdot x_i + x_i^2 < \sum_{i=1}^n (x_i^\mu)^2 - 2 \sum_{i=1}^n x_i^\mu \cdot x_i + x_i^2. \quad (7)$$

From (7) and (4), we obtain (5). This shows that using these two metrics to compare fairness gives the same result.

The above also indicates that our new metric is an extension of the fairness index [1] to a general network with multiple bottleneck links. We will further illustrate this extension in the following sections using two popular TCP protocols, TCP Reno (a widely used protocol) and FAST TCP (a leading future option), in a general network settings. The choice of these two protocols

3. Fairness comparison of FAST TCP versus TCP Reno for a general network

We now compare the fairness of FAST TCP and TCP Reno using the proposed metric. We assume that, for a FAST TCP flow i , if it passes multiple links, the summation

of the total queue length in equilibrium at these links is α_i . Consider the utility optimization problem (2), for FAST TCP we use the following utility function as suggested by [21,22,26,27]

$$U_i^f(x_i^f) = \alpha_i \log x_i^f, \quad (8)$$

where x_i^f is the rate of FAST TCP flow i in equilibrium. We take the utility function of TCP Reno [37]

$$U_i^r(x_i^r) = \frac{1.5}{D_i^2} \cdot \left(-\frac{1}{x_i^r}\right), \quad (9)$$

where D_i is the equilibrium round trip time (RTT), i.e., propagation plus queuing delay, that TCP Reno flow i experienced and x_i^r is the rate of TCP Reno flow i . The choice of this utility function is without loss of generality and it is made for the sake of simplicity; a similar analysis can also be performed choosing an alternative utility function.

Solving Problem P may require coordination among possibly all flows; sometimes we solve it by solving its dual problem. By solving its dual problem, the rate of the flow i is obtained to be given by (see [38])

$$x_i = (U_i'(p))^{-1}, \quad (10)$$

where $(U_i'(p))^{-1}$ is the inverse of $U_i'(p)$, and p is the path bandwidth price for flow i . A procedure has been given in [38] on how to calculate the bandwidth allocation for a certain protocol that optimizes the aggregate network utility. We hereby follow that procedure. Using (8), (9) and (10), we are able to obtain the bandwidth allocation for FAST TCP and TCP Reno that optimizes the aggregate utility. They are

$$x_i^f = \frac{\alpha_i}{p_i^f}, \quad (11)$$

$$x_i^r = \sqrt{\frac{1.5}{p_i^r}} / D_i, \quad (12)$$

respectively, where p_i^f, p_i^r are the path bandwidth price of FAST TCP flow i and TCP Reno flow i . From Definition 1, if

$$\sqrt{\sum_{i=1}^n (x_i^f - x_i)^2} \leq \sqrt{\sum_{i=1}^n (x_i^r - x_i)^2}, \quad (13)$$

FAST TCP is closer to max–min than TCP Reno. Eq. (13) requires a solution of Problem P by global optimization. However, combining (11), (12), (13) leads to the following condition

$$\sqrt{\sum_{i=1}^n \left(\frac{\alpha_i}{p_i^f} - x_i\right)^2} \leq \sqrt{\sum_{i=1}^n \left(\sqrt{\frac{1.5}{p_i^r}} / D_i - x_i\right)^2}, \quad (14)$$

which can be solved in a distributed manner by having each flow compute its own price as suggested by [38] and provides condition for fairness comparison between FAST TCP and TCP Reno.

4. Two numerical examples

In this section, we present two examples to demonstrate how to compare FAST TCP and TCP Reno relative to max–min.

Example 1 (*The “Parking Lot” Network*).

Consider the “Parking Lot” network shown in Fig. 5. Assume that there are n links each of which has unit capacity, flow i only passes through the i th link, for $i = 1, 2, \dots, n$, and flow $n+1$ passes through all the links. Let the sending rate of each FAST TCP flow be denoted by x_i^f , where $i = 1, 2, \dots, n+1$, when all the flows are FAST TCP flows, and the sending rate of each TCP Reno flow is x_i^r , where $i = 1, 2, \dots, n+1$, when all the flows are TCP Reno flows.

Let the FAST TCP and TCP Reno Problem P utility functions be according to (8) and (9), respectively, the rates of FAST TCP and TCP Reno flows are obtained as follows. To calculate the FAST TCP flows’ rate, let

$$\gamma = \frac{1}{1 + \frac{\sum_{i=1}^n \alpha_i}{\alpha_{n+1}}},$$

by using (8) and (10), we have the FAST TCP bandwidth allocation

$$x_i^f = 1 - \gamma, \quad (i = 1, 2, \dots, n), \quad x_{n+1}^f = \gamma. \quad (15)$$

To calculate the TCP Reno flows’ rate, let

$$\omega = \frac{1}{1 + D_{n+1} \cdot \sqrt{\sum_{i=1}^n \frac{1}{D_i^2}}},$$

by using (9) and (10), we have the TCP Reno bandwidth allocation

$$x_i^r = 1 - \omega, \quad (i = 1, 2, \dots, n), \quad x_{n+1}^r = \omega. \quad (16)$$

The max–min fairness rates are computed to be

$$x_i = 1/2 \quad (i = 1, 2, \dots, n+1). \quad (17)$$

Here, we use (13) for fairness comparison. Let

$$\begin{aligned} \Gamma &= \sqrt{\sum_{i=1}^n (1 - \gamma - 1/2)^2 + (\gamma - 1/2)^2} \\ &= \sqrt{n+1} \cdot |1/2 - \gamma|, \end{aligned} \quad (18)$$

Fig. 5. The “Parking Lot” example.

$$\begin{aligned} \Omega &= \sqrt{\sum_{i=1}^n (1 - \omega - 1/2)^2 + (\omega - 1/2)^2} \\ &= \sqrt{n+1} \cdot |1/2 - \omega|. \end{aligned} \quad (19)$$

For $D_{n+1} > D_i$ ($i = 1, 2, \dots, n$), then

$$D_{n+1} \cdot \sqrt{\sum_{i=1}^n \frac{1}{D_i^2}} = \sum_{i=1}^n \left(\frac{D_{n+1}}{D_i} \right)^2 > 1,$$

so $\omega < 1/2$. Then (19) is reduced to

$$\Omega = \sqrt{n+1} \cdot (1/2 - \omega). \quad (20)$$

If FAST TCP is closer to max–min than TCP Reno, that is $\Gamma < \Omega$, which is equivalent to

$$\sqrt{n+1} \cdot |1/2 - \gamma| < \sqrt{n+1} \cdot (1/2 - \omega), \quad (21)$$

that is

$$\omega < \gamma < 1 - \omega, \quad (22)$$

from it, we further obtain

$$\frac{1}{D_{n+1} \cdot \sqrt{\sum_{i=1}^n \frac{1}{D_i^2}}} < \frac{\sum_{i=1}^n \alpha_i}{\alpha_{n+1}} < D_{n+1} \cdot \sqrt{\sum_{i=1}^n \frac{1}{D_i^2}}. \quad (23)$$

By using (23), we can set the FAST TCP parameter α to make FAST TCP closer to max–min fairness than TCP Reno.

Example 2 (*The NSFNET Backbone Network Example*).

Let us consider the NSFNET Backbone network shown in Fig. 6. Again we set all link capacities to unity, and consider 10 flows. The routing information is given in Table 1. We run the progressive filling algorithm to compute the max–min fair allocation.

Considering a general utility function, the relevant optimization problem is thus formulated as follows:

Maximize $U(x)$

subject to :

$$x_1 + x_2 + x_5 \leq 1,$$

$$x_1 + x_2 + x_3 + x_4 \leq 1,$$

$$x_1 + x_3 + x_4 \leq 1,$$

$$x_1 + x_4 \leq 1,$$

$$x_6 + x_7 + x_8 + x_9 \leq 1,$$

$$x_7 + x_8 + x_9 + x_{10} \leq 1,$$

$$x_8 + x_9 + x_{10} \leq 1,$$

$$x_9 + x_{10} \leq 1.$$

Assume all α are equal for all FAST TCP flows. We use the tools provided by AMP Raw Data Query [44] to measure the RTTs between the source and destination pairs of the 10 flows. We set the source and destination nodes as the university machines in those cities. We use FAST TCP utility functions given by (8) and TCP Reno utility functions

Fig. 6. The NSFNET Backbone network.

Table 1
Bandwidth allocation results for the NSFNET example

Flows	Route	RTT (ms)	Max-min rate	FAST TCP rate	TCP Reno rate
1:1-16	A-B-C-D-E-G	75	0.2500	0.2	0.2712
2:1-8	A-B-C	50	0.2500	0.2	0.1808
3:5-15	B-C-D-E	72.80	0.2500	0.3	0.2708
4:5-16	B-C-D-E-G	74.51	0.2500	0.3	0.2772
5:1-4	A	36.18	0.5000	0.6	0.5480
6:1-2	H	34.42	0.2500	0.4	0.2961
7:1-3	H-I	27.66	0.2500	0.2	0.1783
8:1-7	H-I-J	60.06	0.2500	0.2	0.2590
9:1-11	H-I-J-K	57.03	0.2500	0.2	0.2666
10:2-11	I-J-K	61.11	0.2500	0.4	0.2961

given by (9). These two utility functions will lead to two rate allocation results. We list the flows' route information, max-min rate allocations, the FAST TCP rate allocation, and the TCP Reno rate allocations in Table 1.

We calculate $(\sum_{i=1}^{10} (x_i^f - x_i)^2)^{1/2} = 0.185$, and $(\sum_{i=1}^{10} (x_i^r - x_i)^2)^{1/2} = 0.1050$, so TCP Reno is fairer than FAST TCP in this case. However, when parameters change, different outcomes may arise.

5. Conclusions

We have provided further insights into the issue of how to quantify by how much one TCP protocol is fairer than another and how parameter setting improves the fairness of one protocol over another. Using max-min fairness as a benchmark and the Euclidean norm as a measure of fairness, we have compared FAST TCP and TCP Reno for general networks. Two examples for the use of this measure have been provided.

Acknowledgements

This research has been supported by National Natural Science Foundation of China under Grant Nos. 60473085, 60573118, by the Program NCET-05-0673 from Chinese Ministry of Education and partially by the Australian Research Council.

References

- [1] R. Jain, D. Chiu, W. Hawe, A quantitative measure of fairness and discrimination for resource allocation in shared computer systems, Technical Report DEC-TR-301, Digital Equipment Corporation, September 1984.
- [2] D. Bertsekas, R. Gallager, Data Networks, Prentice Hall, 1987.
- [3] A. Charny, An algorithm for rate allocation in a packet-switching network with feedback, Technical Report MIT/TR-601, MIT, Cambridge, 1994.
- [4] R. Mazumdar, L.G. Mason, C. Douligeris, Fairness in network optimal flow control: optimality of product forms, IEEE Trans. Commun. 39 (5) (1991) 775–782.

- [5] D.-M. Chiu, R. Jain, Analysis of the increase and decrease algorithms for congestion avoidance in computer networks, *Comput. Netw. ISDN* 17 (1) (1989) 1–14.
- [6] M. Zukerman, L. Tan, H. Wang, I. Ouveysi, Efficiency-fairness tradeoff in telecommunications networks, *IEEE Commun. Lett.* 9 (7) (2005) 643–645.
- [7] L. Tan, X. Zhang, L.L.H. Andrew, S. Chan, M. Zukerman, Price-based max–min fair rate allocation in wireless multi-hop networks, *IEEE Commun. Lett.* 10 (1) (2006) 31–33.
- [8] S. Chan, M. Zukerman, Is max–min fairness achievable in the presence of insubordinate users? *IEEE Commun. Lett.* 6 (3) (2002) 120–122.
- [9] B. Vandalore, S. Fahmy, R. Jain, R. Goyal, M. Goyal, General weighted fairness and its support in explicit rate switch algorithms, *Comput. Commun.* 23 (2000) 149–161.
- [10] I. Cidon, Y. Ofek, MetaRing - a full-duplex ring with fairness and spatial reuse, *IEEE Trans. Commun.* COM-41 (1) (1993) 110–120.
- [11] Y. Ofek, Overview of the metaring architecture, *Comput. Netw. ISDN* 26 (6–8) (1994) 817–830.
- [12] G. Anastasi, L. Lenzini, M. La Porta, Y. Ofek, Dynamic max–min fairness in ring networks, *Cluster Comput.* 3 (3) (2000) 215–230.
- [13] G. Anastasi, L. Lenzini, Y. Ofek, Tradeoff between the cycle complexity and the fairness of ring networks, *Microprocess. Microsyst.* 25 (1) (2001) 41–59.
- [14] E.L. Hahne, A.K. Choudhury, N.F. Maxemchuk, DQDB networks with and without bandwidth balancing, *IEEE Trans. Commun.* 40 (1992) 1192–1204.
- [15] M. Zukerman, P. Potter, The DQDB protocol and its performance under overload traffic conditions, *Comput. Netw. ISDN* 20 (1990) 261–270.
- [16] J.L. Boudec, Rate adaptation, congestion control and fairness: a tutorial. Available from: <http://ica1www.epfl.ch/PS_files/LEB3132.pdf>.
- [17] J. Mo, J. Walrand, Fair end-to-end window-based congestion control, *IEEE/ACM Trans. Network.* 8 (5) (2000) 556–567.
- [18] L. Massoulié, J. Roberts, Bandwidth sharing: objectives and algorithms, *IEEE/ACM Trans. Network.* 10 (3) (2002) 320–328.
- [19] A. Tang, J. Wang, S.H. Low, Is fair allocation always inefficient, *Proceedings of IEEE Infocom, Hong Kong, China, March 2004*.
- [20] L.S. Brakmo, S.W. O'Malley, L.L. Peterson, TCP Vegas: new techniques for congestion detection and avoidance, *Proceedings of ACM SIGCOMM 1994, London, UK, August 31–September 2, 1994*, pp. 24–35.
- [21] C. Jin, D. Wei, S.H. Low, FAST TCP: motivation, architecture, algorithms, and performance, *Proceedings of IEEE INFOCOM 2004, Hong Kong, China, March 2004*.
- [22] J. Wang, D.X. Wei, S.H. Low, Modeling and stability of FAST TCP, *Proceedings of IEEE INFOCOM (2005)* 938–948.
- [23] C. Jin, D.X. Wei, S.H. Low, G. Buhrmaster, J. Bunn, D.H. Choe, R.L.A. Cottrell, J.C. Doyle, W. Feng, O. Martin, H. Newman, F. Paganini, S. Ravot, S. Singh, FAST TCP: from theory to experiments, *IEEE Network* 19 (1) (2005) 4–11.
- [24] S. Hegde, D. Lapsley, B. Wyrowski, J. Lindheim, D. Wei, C. Jin, S.H. Low, H. Newman, FAST TCP in High Speed Networks: An Experimental Study, *Proceedings of GridNets, San Jose, CA, 2004*.
- [25] L.A. Grieco, S. Mascolo, Performance evaluation and comparison of Westwood+, New Reno, and Vegas TCP congestion control, *ACM SIGCOMM Comput. Commun. Rev.* 34 (2) (2004) 25–38.
- [26] A. Tang, J. Wang, S. Hegde, S.H. Low, Equilibrium and fairness of networks shared by TCP Reno and Vegas/FAST, *Telecommunications Systems—Special issue on High Speed Transport Protocols* 30 (4) (2005) 417–439.
- [27] L.L.H. Andrew, L. Tan, T. Cui, M. Zukerman, Fairness comparison of FAST TCP and TCP Vegas, *Proceedings of 19th ITC, Beijing, China, August 2005*, pp. 1375–1384.
- [28] L. Tan, C. Yuan, M. Zukerman, FAST TCP: fairness and queuing issues, *IEEE Commun. Lett.* 9 (8) (2005) 762–764.
- [29] T. Cui, L.L.H. Andrew, M. Zukerman, L. Tan, Improving the fairness of FAST TCP to new flows, *IEEE Commun. Lett.* 10 (5) (2006) 414–416.
- [30] J. Mo, R.J. La, V. Anantharam, J. Walrand, Analysis and comparison of TCP Reno and Vegas, *Proceedings of IEEE INFOCOM 1999*, vol. 3 March 1999 pp. 1556–1563.
- [31] K. Kurata, G. Hasegawa, M. Murata, Fairness comparisons between TCP Reno and TCP Vegas for future deployment of TCP Vegas, in: *Proceedings of INET 2000, Yokohama, Japan, July 18–21, 2000*. Available from: <http://www.isoc.org/inet2000/cdproceedings/2d/2d_2.htm>.
- [32] D.D. Luong, J. Br, Bandwidth sharing scheme of end-to-end congestion control protocols, *Period. Polytech., Ser. El. Eng.* 45 (2) (2001) 79–90.
- [33] F.P. Kelly, Charging and rate control for elastic traffic, *European Trans. Telecommun.* 8 (1) (1997) 33–37.
- [34] F.P. Kelly, A.K. Maulloo, D.K.H. Tan, Rate control for communication networks: shadow prices, proportional fairness and stability, *J. Oper. Res. Soc.* 49 (3) (1998) 237–252.
- [35] S. Kunniyur, R. Srikant, End-to-end congestion control schemes: utility functions, random losses and ECN marks, *IEEE/ACM Trans. Network.* 11 (5) (2003) 689–702.
- [36] R.J. La, V. Anantharam, Utility-based rate control in the Internet for elastic traffic, *IEEE/ACM Trans. Network.* 10 (2) (2002) 272–286.
- [37] S.H. Low, A duality model of TCP and queue management algorithms, *IEEE/ACM Trans. Network.* 11 (4) (2003) 525–536.
- [38] S.H. Low, D.E. Lapsley, Optimization flow control, I: basic algorithm and convergence, *IEEE/ACM Trans. Network.* 7 (6) (1999) 861–874.
- [39] S.H. Low, F. Paganini, J.C. Doyle, Internet congestion control, *Control Syst. Mag.* 22 (1) (2002) 28–43.
- [40] A. Tang, J. Wang, S.H. Low, Counter-intuitive throughput behaviors in networks under end-to-end control, *IEEE/ACM Trans. Network.* 14 (2) (2006) 355–368.
- [41] B. Wyrowski, M. Zukerman, MaxNet: a congestion control architecture, *IEEE Commun. Lett.* 6 (11) (2002) 512–514.
- [42] B. Wyrowski, M. Zukerman, QoS in best-effort networks, *IEEE Commun. Mag.* 40 (12) (2002) 44–49.
- [43] B. Wyrowski, L.L.H. Andrew, M. Zukerman, MaxNet: a congestion control architecture for scalable networks, *IEEE Commun. Lett.* 7 (10) (2003) 511–513.
- [44] Measurement and network analysis: AMP raw data query. Available from: <http://watt.nlanr.net/Active/raw_data/cgi-bin/data_form.cgi>.

Liansheng Tan is now with Research School of Information Sciences and Engineering, The Australian National University, Canberra ACT 0200, Australia. Dr. Tan received his Ph.D. degree from Loughborough University in the UK in 1999. He has done research in computer communication network in School of Information Technology and Engineering at University of Ottawa, Ontario, Canada as a postdoctoral research fellow and a visiting research scientist in 2001. During 2002 till 2006, he worked at Department of Computer Science, Central China Normal University, China as a professor and Head of Department. During that period, he also held a number of visiting research positions at Loughborough University, University of Tsukuba, City University of Hong Kong and University of Melbourne. Dr. Liansheng Tan is currently an Associate Editor of *Dynamics of Continuous, Discrete & Impulsive Systems (Series B: Applications & Algorithms)*, he is also on the Editorial Board of *International Journal of Communication Systems*. He has published over seventy referred papers in international journals and conferences. His research interests include modeling, congestion control analysis and performance evaluation of computer communication networks, resource allocation and management of wireless and wireline networks and routing and transmission control protocols.

Lina Dong received her BSc and MSc degrees from Centre China Normal University, China, in 2003 and 2006, respectively, all in computer science. Currently she is working in Huawei Technologies. Her research interests include communication networks, queueing theory, and Internet congestion control.

Cao Yuan received his BSc. and MSc. degrees from Central China Normal University, China, in 2002 and 2006, respectively, all in computer science. Currently he is a PhD candidate in Central China Normal University. His research interests include communication networks, queueing theory, Internet protocols and control system.

Moshe Zukerman received his B.Sc. in Industrial Engineering and Management and his M.Sc. in Operation Research from Technion – Israel Institute of Technology and a Ph.D. degree in Electrical Engineering from The University of California Los Angeles in 1985. Dr. Zukerman was an independent consultant with IRI Corporation and a post-doctoral fellow at UCLA during 1985–1986. During 1986–1997 he served in Telstra Research Laboratories (TRL), first as a research engineer and between 1988–1997 as a

project leader managing a team of researchers providing expert advice to Telstra on network design and traffic engineering, and on traffic aspects of evolving telecommunications standards. In 1990 he received the Telstra Research Laboratories Outstanding Achievement Award. In 1997 he joined The University of Melbourne where he is now a professor responsible for promoting and expanding telecommunications research and teaching in the Electrical and Electronic Engineering Department. During 1990–2001, he has also taught and supervised graduate students at Monash University. He has served as a session chair and member of technical and organizing committees of numerous national and international conferences. He gave tutorials in several major international conferences such as IEEE ICC and IEEE GLOBECOM. He has also served on the editorial boards of the Australian Telecommunications Research Journal, during 1991–1996, and the Computer Networks, during 1999–2003. He served as a Series Editor for the IEEE Communications Magazine during 2001–2004 and twice as a Guest Editor of IEEE JSAC. Presently, he is serving on the editorial board of the IEEE/ACM Transactions on Networking, the International Journal of Communication Systems, and as a Senior Editor of IEEE JSAC. Professor Zukerman is an IEEE Fellow and has over 250 publications in scientific journals and conference proceedings and has been awarded several national and international patents.