Institut Eurécom
Department of Corporate Communications
2229, route des Crètes
B.P. 193
06904 Sophia-Antipolis
FRANCE

Research Report RR-06-183

Size-based scheduling to improve fairness and performance in 802.11 networks

Diego Ferrero, Guillaume Urvoy-Keller

Tel : (+33) 4 93 00 26 26
Fax : (+33) 4 93 00 26 27
Email : {diego.ferrero, urvoy}@eurecom.fr

[1]Institut Eurécom's research is partially supported by its industrial members: Bouygues Téécom, France Téécom, Hitachi Europe, SFR, Sharp, ST Microelectronics, Swisscom, Texas Instruments, Thales.

Abstract

Using wireless technologies to access the Internet (e.g., public hotspots) and setting small ad hoc networks (e.g., during a meeting), are common experiences for end users. Those environments are challenging for TCP because of the variability of the available bandwidth and the competition among flows that span over paths of different lengths and different types (wired or wireless). In this paper, we advocate that part of the problem comes from the use of the FIFO scheduling policy. We propose to replace FIFO with a size-based scheduling policy called LAS (Least Attained Service). We demonstrate that LAS solves most of the fairness issues encountered in those environments and makes a better use of the scarce bandwidth resources.

1 Introduction

TCP is the dominant transport protocol in the Internet, carrying typically over 80% of the bytes of a given link. The vast majority of new applications, e.g., peer-to-peer file sharing/replication, podcasting have adopted TCP as their transport layer. The widespread use of wireless communications, with home wireless networks, wireless LANs [1] or ad hoc networks, should not change the current situation and TCP is likely to remain the transport layer of most applications in those environments.

A number of problems due to the use of TCP in 802.11 wireless networks have been identified over the years; the reason is that TCP has been optimized for a wired network. Our focus in this paper is on TCP in static 802.11 multihop wireless networks. This category encompasses many practical cases, such as single hop networks like wireless LANs and multi-hop ad hoc networks. One of the biggest concerns with TCP in those environments is fairness. We advocate that the fairness problem is exacerbated by the interaction between TCP and the FIFO scheduling policy. For these networks, we propose to replace the FIFO scheduling policy by a size-based scheduling policy called LAS (Least Attained Service - [8], page 172) to improve fairness among TCP flows and the interactivity perceived by the user.

LAS has been shown to solve most of the fairness issues faced by TCP in a wired environment [11] that arise due to the competition between flows with different RTTs or between TCP and UDP flows. Another key feature of LAS is its ability to protect TCP flows in their early phase (typically their slow-start phase) by reducing their loss rate and RTT [10]. In addition, since small TCP flows are given a higher priority under LAS, the interactivity perceived by the end user is improved, e.g., when simultaneously running a peer-to-peer file sharing application and browsing the Web. In the present work, we investigate if those appealing properties of LAS persist in the wireless environments we focus on. Note that while in a wired context, the use of LAS is somehow restricted to specific locations of the network where a bottleneck is observed, LAS has the potential to be used in all wireless networks where bandwidth is in general a scarce resource.

We have compared LAS to FIFO for different network topologies, namely a chain, a grid and a wireless LAN topology and different workloads, namely long lived flows or flow size distributions with various levels of skweness[2]. Using a simulation approach, we show that:

- LAS solves all known fairness issues in wireless networks: short vs. long TCP flows, TCP vs. UDP flows, dowloading vs. uploading TCP flows.

- Fairness is not obtained at the expense of performance in a wireless LAN

[1] We use the term wireless LAN in this paper to refer to a single cell infrastructure 802.11 network where all wireless clients are in the data range of the access point.

[2] Skweness refers here to the ratio of small to large flows and the fraction of bytes carried by those two groups. Refer to Section 4.

scenario. In a multihop network scenario, the price to pay to enforce fairness with LAS is a smaller aggregate throughput as compared to FIFO.

- Fairness and to a lesser extent performance are significantly less sensitive to the choice of the advertised window and buffer size under LAS than under FIFO.

- For grid and wireless LAN topologies and for realistic workload models, LAS is able to increase interactivity with a reasonable penalty for large flows.

The remaining of this paper is organized as follows. In Section 2, we survey the related work concerning TCP in 802.11 networks. In Section 3, we present LAS. In Section 4, we detail our simulation scenarios, in terms of topologies and workload models. In Section 5, we describe our performance and fairness metrics. In Section 6, we present numerical results for the comparison between FIFO and LAS. In Section 7, we provide some conclusions and directions for future work.

2 Related Work

Fairness and performance of TCP flows in ad hoc networks and wireless LANs have received a lot of attention. This stems from the poor and unpredictable performance observed in those networks, as confirmed by many experimental and measurement studies [13, 4, 7].

Many studies on ad hoc networks have underlined the crucial impact of the advertised window parameter [1, 5]. The main conclusion from those studies is twofold. First, the optimal value of the advertised window is in general small, between 2 and 10 packets. Second, TCP generally overshoots this optimal value which results in a higher level of contention and a decrease of fairness. To force TCP to operate around its optimal value, a RED-like mechanism is proposed in [5], where the decision to discard a packet is based on the level of contention experienced by a given node. Note however, that the above results hold for networks with hidden nodes. In contrast, when there is no hidden node, as it is the case in our paper (except in Section 6.1.2), performance increases monotically with an increasing advertised window for the two policies we consider, namely FIFO and LAS.

The case of a wired-cum-wireless network has to be treated individually. Indeed, as TCP connections now span over both a wired and a wireless network, the advertised window can no longer be clamped to small values as in pure ad hoc networks [1, 5]. In this context, fairness among TCP flows or UDP and TCP flows is hardly obtained and performance can quickly degrade as confirmed experimentally in [16]. In [17], the authors propose to replace the FIFO scheduling policy by a non work-conserving policy, with 8 parameters to tune, that aims at minimizing the contention on the air interface. In contrast, LAS is a work-conserving policy with no parameter to tune. However, the focus of [17] is on networks with hidden nodes, which we do not address in the present paper, except in Section 6.1.2.

In [15], the authors focus on the case of an 802.11 wireless LAN with simultaneous uploads and downloads to and from the Internet and relate the observed unfairness (the uploading flow taking advantage over all downloads) to the buffer size at the base station. They propose to enforce fairness by adjusting the advertised window of TCP connections. The authors however acknowledge that their solution is not scalable as it requires to passively estimate the RTT of each connections to adjust the advertised window, which is a non trivial task.

In the present work, we propose to use LAS, a size-based scheduling policy, in a wireless context. Size-based scheduling has already been proposed in a wireless context in [9]. The considered size-based policy is not LAS but the shortest job first (SJF) policy. The objective of the authors in [9] is to apply this policy at the packet level to ensure fairness between short and long packets and more generally to favor multimedia applications. In contrast, we use LAS to solve fairness issues at the TCP level and we thus apply it at the flow level.

3 The Least Attained Service Policy

LAS is a size-based scheduling policy. It has been initially proposed and studied in the context of time-sharing computers in the late 60s [14]. Under LAS priority is given to the job that has received the least amount of service. In case of ties, jobs share the server in a round-robin manner. A salient feature of LAS is that it has no internal parameter to tune.

In [12], LAS has been studied in the context of a packet network like the Internet, where it is extended to incorporate a buffer management policy. The resulting policy that we refer to as LAS in the remaining of the paper works as follows. Upon reception, a packet is assigned a priority which is inversely proportional to the number of bytes sent so far by the corresponding connection (the first packet of a new connection thus has maximum priority). If ever the queue is full upon the arrival of a new packet, this packet is assigned its priority, inserted in the queue and the packet with the lowest priority is discarded.

The results on LAS relevant in the context of this paper are the following. In [12], LAS has been proved to improve the performance of most flows except a small fraction of the largest ones as compared to FIFO if ever the flow size distribution is highly skewed, i.e., a clear minority of the largest flows (say less than 5%) convey most of the bytes. Internet traffic has often been observed to exhibit highly skewed distributions [2]. LAS has also been proved to interact nicely with TCP by protecting flows in their slow-start phase in [10] and to solve classical unfairness situations: UDP vs. TCP or TCP flows with different RTTs [11].

The implementation of LAS can be performed at the flow or at the connection level. In order not to penalize applications that generate long connections but send at low rate, e.g, a routing process between two machines that sends keep-alive messages at a low rate most of the time, it seems wiser to implement LAS at a flow level. This is the option we take in the rest of this paper. In addition, to account for connections sending packets of different sizes, we chose to use as attained service the amount of bytes of this flow serviced so far. TCP level ACK packets are processed in a specific way. The size of one ACK packet is in general an order of magnitude smaller than a TCP packet. In order to avoid to give full priority to the ACK flow, the priority assigned to an ACK packet is proportional to the amount of data acknowledged so far for the corresponding flow (see Section 6.1.3). The overhead of LAS as compared to FIFO is that per flow statistics must be kept. It should however be an affordable task for an ad hoc node or an access point that should not service a large number of simultaneous connections.

4 Simulation Scenarios

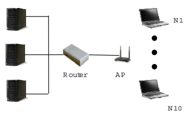

Figure 1: Wireless LAN scenario

In this section, we outline the scenarios that we investigate in Section 6. For the case of ad hoc networks, we consider two basic topologies: a chain and a grid.

For the chain topology, we consider either long-lived FTP flows or a mix of FTP and constant bit rate UDP flows. This type of workload allows us to investigate classical unfairness situations and to understand the role of two key parameters, namely the buffer size at each node and the maximum advertised window of TCP. For the grid topology, we consider a classical workload model where FTP connections arrive according to a Poisson process, the duration of connections follows an exponential distribution and the source and destination of a connection are chosen randomly.

For the wireless LAN case (see Figure 1), we first consider FTP connections with a mix of uploads and downloads to and from the Internet. The idea here is to investigate the unfairness that occurs between uploads and downloads and also between uploads only or downloads only when the latencies of the paths differ. We next consider a Web-like workload. We aim at investigating the ability of LAS to offer a good interactivity with this scenario. Our Web-like model is based on the work in [3]. There is a pool of clients and a pool of servers. Clients request pages to servers according to a Poisson process. Each Web page contains a certain number of objects. Each time an object is requested, a new TCP connection is established. Distribution of the objects size follows either an exponential or a Pareto distribution. We set the Web parameters as shown in Table 1. We have used different coefficient of variation [3] CoV values for the page size distribution in order

[3]The coefficient of variation of a distribution is the ratio of the standard deviation to the mean of

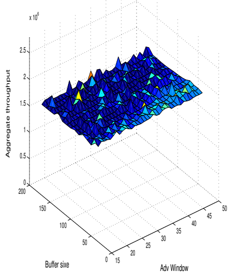

Figure 2: LAS - Aggregate throughput

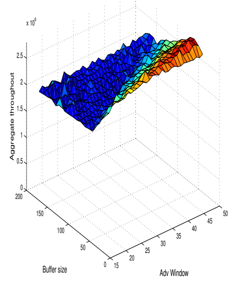

Figure 3: FIFO - Aggregate throughput

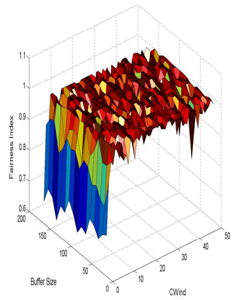

Figure 4: Fairness index of the throughput obtained under LAS

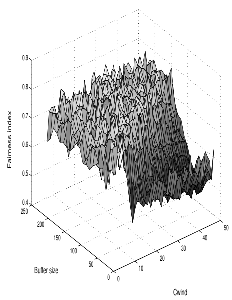

Figure 5: Fairness index of the throughput obtained under FIFO

to investigate how LAS and FIFO behave for various amounts of skewness in the flow size distribution. Note that in a wired context, it has been shown that the more skewed the distribution, the better the performance of LAS as compared to FIFO [12].

The simulation software we use is ns version 2.29. We have chosen 802.11b as a MAC protocol. We report in Table 2 the basic configuration used for all simulations. The version of TCP is NewReno. The Maximum Segment Size is equal to 1000 bytes. UDP segments are also of size 1000 bytes.

Page size (in objects)	3
Average Object size	30 pkts
CoV of Object size	1/2/5
Inter-Page time	0.1s
Inter-object time	0.1s

Table 1: Main parameters of the Web traffic model

Data rate	11Mb/s
RTC/CTS option	off
TCP delay-ACK option	on
Data range	250m
Interference range	550m
Routing protocol	DSDV

Table 2: Main parameters of the simulated scenarios

the distribution. It allows to compare the variability of different distributions irrespectively of their actual mean value.

5 Performance and fairness metrics

We use the following metrics to quantify performance and fairness:

- *Aggregate throughput*: the aggregate throughput is defined as the sum of the bytes correctly received by the TCP layer at the destination. It might include duplicated packets.
- *Total number of TCP data packets transmitted*: The total number of TCP data packets transmitted over the network is the sum of all such packets exchanged between any two nodes, that might be either intermediate nodes or the final destination and that were correctly received by the final destination. Practically, this means that if a data packet has been transmitted three times before reaching its final destination, it will be counted three times in the metric.
- *Fairness index*: The Jain's fairness index [6] allows to compare how the available bandwidth is shared among the competing connections. For a set of n connections with average throughputs $x_i,\ i \in \{1, n\}$, the Jain's fairness index I is defined as:

$$I = \frac{\left(\sum_{i=1}^{n} x_i\right)^2}{n \sum_{i=1}^{n} x_i^2} \quad (1)$$

The closer I is from 1, the more fair is the share of the available bandwidth among the connections.

6 Numerical Results

In this section, we first investigate the fairness of LAS as compared to FIFO for an ad hoc chain and a wireless LAN topology. This study is performed for long lived flows that allow to clearly reveal unfair situations. We next investigate the ability of LAS to favor short flows (at the expense of a reasonable penalty for the largest flows) for a grid and a wireless LAN topology.

6.1 Fairness

Our main focus in this paper is on networks without hidden node. We however present in Section 6.1.2 some preliminary results that show that LAS can solve some of the issues raised by the presence of hidden nodes.

6.1.1 Chain topology without hidden node

Long-lived TCP connections We consider a chain topology with 3 FTP sessions between node 0 and nodes 1, 2, 3 respectively (nodes are numbered from left to right in the chain). We term them as connection 1, 2 and 3.

We study the performance and the fairness of FIFO and LAS, both as a function of the advertised window and buffer size values. As small values of the advertised window have often been observed to lead to the best performance in ad hoc networks, we investigated values between 1 and 50 kbytes. We considered buffer sizes between 20 and 200 kbytes. Larger values do not change the results as with a buffer of 200 kbytes, each node can store all the outstanding packets (maximum of 50 kbytes) of all the 3 FTP connections.

Figures 2 and 3 represent the aggregate throughputs of LAS and FIFO respectively. Those figures reveal that the aggregate throughput under FIFO is in general 50% higher than under LAS. However, since the total numbers of TCP data packets sent under LAS and under FIFO are similar (we do present the figures here due to space limitation), the lower aggregate throughput under LAS should not be interpreted as a lower efficiency at the MAC level, but as a different strategy to allocate the bandwidth to the connections.

The fairness indexes in Figures 4 and 5 for LAS and FIFO show that:

- The fairness under LAS is consistently higher than under FIFO and the sharing of the available bandwidth under LAS is often almost perfect, i.e., close to 1.
- LAS is clearly less sensitive than FIFO to the choice of the advertised window or the buffer size values.

TCP vs. UDP Our objective here is to show that LAS enforces fairness among TCP and constant bit rate UDP connections. To this end, we considered a worst case scenario with one UDP connection between nodes 0 and 1 and another one between nodes 0 and 2. We next set up an FTP connection between the two ends of the chain, nodes 0 and 3. The main result from those simulations is that under FIFO, the FTP connection completely starves, while under LAS the three connections share the channel perfectly (we do not present any curve due to space constraints). As a conclusion, even in a worst case scenario like the one considered here, LAS is able to protect TCP flows and enforce a fair share of the channel between TCP and UDP flows.

6.1.2 Chain topology with one hidden node

We have analyzed the case of a chain topology with 5 nodes, where the last one is a hidden node. We set up 4 FTP sessions between node 0 and nodes 1, 2, 3, 4 respectively. We term them as connection 1, 2, 3 and 4. We ran simulations for the same range of parameters for the advertised window and the buffer sizes as in the case without hidden nodes. We do not report the graphs here due to space constraints. We observed from those simulations that under FIFO, connection 4 always starves, while under LAS, the four connections almost always obtains the same throughput. In addition, the total number of packets sent under FIFO is the same as under LAS, just like in the case without hidden node. Those preliminary results are encouraging. We leave for future work a more in depth study of the performance of LAS in networks with hidden nodes.

6.1.3 Wireless LAN scenario

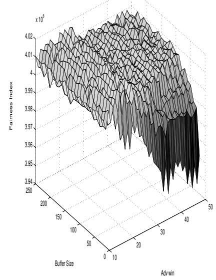

Figure 6: LAS - Wireless LAN - Aggregate throughput

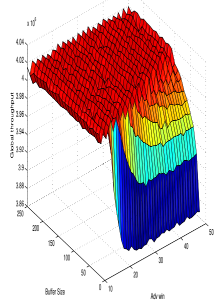

Figure 7: FIFO - Wireless LAN - Aggregate throughput

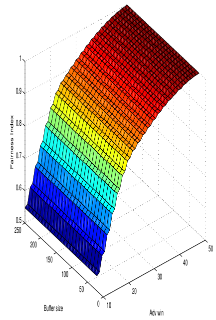

Figure 8: LAS - Wireless LAN - Fairness Index

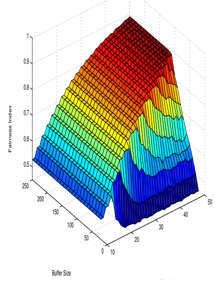

Figure 9: FIFO - Wireless LAN - Fairness Index

We compare in this section LAS to FIFO for a wireless LAN scenario as the one of Figure 1, with 3 hosts on the wired part and 3 active wireless hosts. Links are provisioned such that only the queue of the access point can build up.

Downloads only scenario We first investigated the case of a wireless LAN with 3 FTP servers in the wired part and 3 FTP clients in the wireless part. Latencies on the wired part of the path for the 3 connections are respectively 2, 50 and 150 ms. The advertised window value and the buffer size at the access point are key parameters that influence the performance. We consider advertised window values between 10 and 50 kbytes. We do not consider smaller values (less than 10 kbytes) as in Section 6.1.1 because of the bad utilization of the wired part that results from such a choice.

We first observe from Figures 6 and 7 that LAS and FIFO achieve in general the same aggregate throughput (aggregate throughput and total number of packets coincide in a single hop scenario). Fairness indexes for both policies are presented in Figures 9 and 8. Those figures reveal that LAS becomes fair (a fairness index above 0.9) as soon as the congestion window values are above 25 kbytes and irrespectively of the buffer size. This is not the case for FIFO that tends to be fair (its fairness index remains below 0.9) for larger buffer sizes only but whose performance degrades for decreasing buffer sizes. Those results highlight the detrimental effect that a non adequate choice of the buffer size can have on TCP when used in conjunction with FIFO. Such a trade-off does not exist with LAS for which fairness depends only on the advertised window with a simple rule: the larger the advertised window, the better the fairness.

Downloads and uploads scenario We investigate here a different flavor of the previous scenario where one FTP server is in the wireless part while the two other ones remain in the wired part. Such a scenario, with simultaneous downloads and uploads, has been observed to lead to highly unfair situations where the uploading connection obtains the highest throughput [15]. The reason behind unfairness is the competition that takes place between the (TCP level) ACK stream of the uploading flows and the data streams of the downloading flows at the buffer of the access point. We have performed for this scenario the same analysis as for the previous case. Simulation results are similar to the downloads only case, i.e., LAS is able to enforce fairness while FIFO is not (we do not report the graphs here due to space limitation). However, for LAS to enforce fairness, one must use the implementation technique described in Section 3: the priority of an (TCP level) ACK packet must be set equal to the amount of bytes sent by the corresponding data streams[4]. If it is not the case, an ACK packet being 25 times smaller than a data packet, the ack stream of the upload gain full priority at the buffer of the access point and the fairness between uploads and downloads can no longer be maintained.

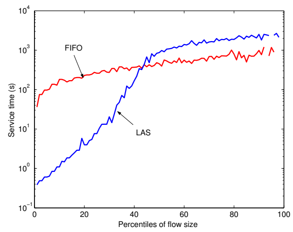

Figure 10: Grid - Service time with LAS

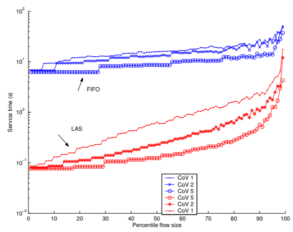

Figure 11: Percentile of service time under LAS and FIFO with different CoV of the page size

[4]As the forward and return path might not be the same, we use in our current implementation the highest byte number acknowledged by the ACK. We leave for future work the problem of losses and reordering.

6.2 Interactivity

6.2.1 Grid Topology

We compared LAS to FIFO for the case of an ad hoc network with a grid topology and a workload model where FTP connections ”arrive” according to a Poisson process, are established between two nodes chosen at random and have a exponentially distributed size (see Section 4). The grid consists of 16 nodes spaced by a distance of 128 m such that a node is in the data range of its direct neighbors only and all nodes are in the same interference range. We used buffer sizes of 10^3 kbytes and maximum advertised windows of 60 kbytes for all nodes and all connections. The objective of this experiment is to show that LAS favors short interactive flows with a negligible penalty for large flows in a realistic workload scenario.

We consider as a metric the conditional response times of the connections i.e., the time required to complete the transfer for each possible connection size. Conditional response times are plotted in Figure 10 against the percentiles of the connection sizes so as to account for the relative proportion of connections of different sizes.

We observe from Figure 10that the response time of the 45% smallest flows is significantly smaller under LAS than under FIFO while the increase in response time of the largest flows remains reasonnable. By reasonnable, we mean that we expect that the benefit due to a better interactivity will outweigh the slowdown of the longest transfers from the user perspective.

The scenario used in this section for the grid topology is based on a flow size distribution which is not highly skewed as the CoV of the exponential distribution is equal to 1. We made this choice because of a lack of evidence that the flow size distribution can be highly skewed in a pure ad hoc network. In contrast, for a wireless network connected to the Internet as it is the case for the wireless LAN scenario in the next section, the flow size distribution can be clearly more skewed. In this case, the advantage of LAS over FIFO will become even more evident.

6.2.2 Wireless LAN Scenario

Wireless LANs are typically installed in order to provide Internet connectivity to mobile users. We consider the scenario of Figure 1 with 3 web servers in the wired part and 10 wireless clients. The workload model is the one depicted in Section 4. We considered a Poisson arrival process of sessions, with 3 different CoV values of 1,2 and 5 for the page size distribution, in order to show the performance of LAS under different conditions.

As one can see from Figure 11, the service time under LAS is smaller than the one under FIFO for 99% of the (smallest) connections. In addition, when the CoV of the page size distributions increases, the advantage of LAS becomes more pronounced, in line with what has been observed in a wired network [12]. The comparison between the grid case and the wireless LAN case with an exponential distribution for the sizes of the connections (CoV=1), also reveals that LAS provides better results in terms of interactivity as compared to FIFO when the number of bottlenecks is smaller (there is a single bottleneck in the Wireless LAN scenario and potentially 16 in the grid scenario).

7 Conclusion and future work

In this work, we proposed to replace FIFO by LAS, a size-based scheduling policy, to improve fairness and interactivity in wireless networks. Specifically, we investigated three topologies: a chain, a grid and a wireless LAN. For those scenarios, we demonstrated that LAS is less sensitive than FIFO to the choice of two key parameters that greatly impact the performance of TCP in wireless environments, namely the advertised window and the buffer size values. In addition, LAS enforces fairness without impacting the network utilization,as the total numbers of packets transmitted are the same under LAS and under FIFO. More generally, LAS makes a better usage than FIFO of the available bandwidth by favoring TCP flows in their early phase, typically their initial slow start phase. This is a very important feature in wireless environments where the available bandwidth can heavily fluctuate due to the shared nature of the medium as it allows to preserve the interactivity perceived by the end user. The next steps for us are to investigate LAS in wireless networks with hidden nodes and also in wireless mesh networks that constitute a multihop extension of the Wireless LAN scenario.

References

[1] K. Chen, Y. Xue, and K. Nahrstedt, "On setting TCP's congestion window limit in mobile ad hoc networks", In *ICC '03*, volume 2, pp. 1080–1084, 2003.

[2] M. Crovella and A. Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes", *IEEE/ACM Transactions on Networking*, pp. 835–846, December 1997.

[3] A. Feldmann, A. Gilbert, P. Huang, and W. Willinger, "Dynamics of IP Traffic: A Study of the Role of Variability and the Impact of Control", In *ACM SIGCOMM*, August 1999.

[4] M. Franceschinis, M. Mellia, M. Meo, and M. Munaf, "Measuring TCP over WiFi: A Real Case", In *1st workshop on Wireless Network Measurements (Winmee)*, Riva Del Garda, Italy, 2005.

[5] Z. Fu, H. Luo, P. Zerfos, S. Lu, L. Zhang, and M. Gerla, "The impact of multihop wireless channel on TCP performance", *Mobile Computing, IEEE Transactions on*, 4(2):209–221, 2005.

[6] R. Jain, D. M. Chiu, and W. Hawe, "A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Computer Systems", Technical report 301, DEC, Littleton, MA, September 1984.

[7] V. Kawadia and P. R. Kumar, "Experimental investigations into TCP performance over wireless multihop networks", In *E-WIND '05: Proceeding of the 2005 ACM SIGCOMM workshop on Experimental approaches to wireless network design and analysis*, pp. 29–34, New York, NY, USA, 2005, ACM Press.

[8] L. Kleinrock, *Queuing Systems, Volume II: Computer Applications*, Wiley, New York, 1976.

[9] S. Pollin, A. Motamedi, A. Bahai, F. Catthoor, and L. Van der Perre, "Delay improvement of IEEE 802.11 distributed coordination function using size-based scheduling", volume 5, pp. 3484–3488 Vol. 5, 2005.

[10] I. Rai, E. W. Biersack, and G. Urvoy-Keller, "Size-based Scheduling to improve the Performance of Short TCP Flows", *IEEE Network Magazine*, 2005.

[11] I. A. Rai, Urvoy-Keller, and E. W. Biersack, "LAS Scheduling Approach to Avoid Bandwidth Hogging in Heterogeneous TCP Networks", In *HSNMC'04*, July 2004.

[12] I. A. Rai, G. Urvoy-Keller, and E. W. Biersack, "Analysis of LAS Scheduling for Job Size Distributions with High Variance", In *Proc. ACM SIGMETRICS*, pp. 218–228, June 2003.

[13] M. Rodrig, C. Reis, R. Mahajan, D. Wetherall, and J. Zahorjan, "Measurement-based characterization of 802.11 in a hotspot setting", In *E-WIND '05: Proceeding of the 2005 ACM SIGCOMM workshop on Experimental approaches to wireless network design and analysis*, pp. 5–10, New York, NY, USA, 2005, ACM Press.

[14] L. E. Schrage, "The Queue M/G/1 with Feedback to Lower Priority Queues", *Management Science*, 13(7):466–474, 1967.

[15] P. Sinha, S. Pilosof, R. Ramjee, D. Raz, and Y. Shavitt, "Understanding TCP fairness over Wireless LAN", In *Proceedings Infocom 2003*, April 2003.

[16] K. Xu, S. Bae, S. Lee, and M. Gerla, "TCP behavior across multihop wireless networks and the wired internet", In *WOWMOM '02*, pp. 41–48, New York, NY, USA, 2002, ACM Press.

[17] L. Yang, W. K. Seah, and Q. Yin, "Improving fairness among TCP flows crossing wireless ad hoc and wired networks", In *MobiHoc '03*, pp. 57–63, 2003.