The 18th Annual IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC'07)

OPTIMIZING AGGREGATE THROUGHPUT OF UPSTREAM TCP FLOWS OVER IEEE 802.11 WIRELESS LANS

Nakjung Choi, Jiho Ryu, Yongho Seok, Taekyoung Kwon and Yanghee Choi
Computer Science and Engineering, Seoul National University, Korea
Email: {fomula, jhryu, yhseok, tk, yhchoi}@mmlab.snu.ac.kr

ABSTRACT

This paper via analysis and simulation revisits the interaction between MAC contention and TCP congestion control over IEEE 802.11 WLANs, misled in the previous efforts. The results reveal that the *effective* number of contending wireless stations is not proportional to the number of wireless stations with an upstream TCP flow in IEEE 802.11 wireless LANs. Thus, we propose a new scheme called TCP ACK Priority (TAP) in which, by allowing an access point to transmit TCP ACKs at the highest priority, the optimal number of competing stations are allowed to contend for media access to utilize link bandwidth efficiently. We use an ns-2 simulator to evaluate the performance of TAP with the IEEE 802.11 DCF. The results show that there is an improvement in network performance without the loss of fairness between upstream TCP flows. The extensions for IEEE 802.11e/n are also considered.

I. INTRODUCTION

Most Internet applications use TCP, but it performs poorly in wireless environments. Thus, there have been several efforts to improve TCP performance in IEEE 802.11 wireless LANs [1]. However, previous authors have blindly assumed that the number of stations (STAs) competing at the MAC layer is proportional to the number of STAs with a TCP flow, which is not true under congested network conditions. This is because TCP has a congestion control mechanism [4] consisting of two phases: *slow start* and *congestion avoidance*. In the slow start phase, a congestion window ($cwnd$) is initialized to one segment when a new TCP connection is established. Each time a TCP ACK is received, $cwnd$ is increased by one segment, resulting in an exponential increase in the window size for every round trip time (RTT). After $cwnd$ reaches the slow start threshold ($ssthresh$), the congestion avoidance phase starts, in which $cwnd$ is incremented by 1/segment after each RTT, resulting in a linear increase in the window size. Note that when a STA has sent $cwnd$ of TCP data, it cannot send more data until it receives a TCP ACK from the recipient. Therefore, a STA waiting for a TCP ACK cannot contend for channel access at the MAC layer. In a congested wireless LANs, the result is that only a small number of wireless STAs that wish to send TCP data will effectively contend for wireless channel access, while most of the STAs cannot transmit TCP packets due to $cwnd$ limitations.

It is normally assumed that each STA and an access point (AP) will get a fair share of the wireless channel capacity in the long term. After a STA transmits TCP packets of $cwnd$ size, it cannot attempt to transmit the next TCP packet until it receives a TCP ACK. That is, after the AP transmits a TCP ACK to the STA, the STA will compete with the AP (and other STAs, if any) for wireless channel access. Our empirical study in Section A. reveals that the number of competing STAs at the MAC layer is mostly around two regardless of the number of STAs with an upstream TCP flow. However, the number of competing STAs should be increased to achieve an optimal aggregate throughput in IEEE 802.11 wireless LANs.

The remainder of this paper is organized as follows. In Section II, we introduce previous research on improving the TCP performance over wireless networks. Then, we analyze why the performance problem takes place in case of TCP over IEEE 802.11 wireless LANs and propose our mechanism called *TCP ACK Priority* (TAP) to enhance the aggregate TCP throughput in Section III. Section IV presents simulation results in terms of the effective number of competing STAs, aggregate throughput, fairness index and delay. Finally, we conclude our work in Section V.

II. RELATED WORK

To improve the performance of TCP over wireless networks, I-TCP [5] uses a split-connection approach. It divides a TCP connection into two distinct connections: a wired connection between a fixed host and an AP, and a wireless connection between an AP and a wireless STA. When the AP receives a TCP data packet through the wired connection, it sends an TCP ACK corresponding to the TCP data packet back to the fixed host and transfers the TCP data packet to the wireless connection. The main advantage of this approach is that transmission errors over wireless links can be hidden from the TCP sender, which is in the wired part of the network, and TCP over the wireless link can be optimized independent of the wired connection. However, it has the disadvantage that it violates the end-to-end semantics of TCP. For example, a TCP ACK may be delivered to a TCP sender before the associated TCP data is actually delivered to the TCP recipient.

In another study [6], a TCP snoop protocol makes changes to the network layer by introducing a new module called a snoop agent at an AP. This agent buffers TCP data packets which are destined for wireless STAs and which have not yet been acknowledged by them. When the loss of a packet is detected by the arrival of duplicated TCP ACKs, or by TCP timeout, the agent performs local retransmission across the wireless link. This scheme also prevents the TCP sender from invoking unnecessary fast-retransmission by dropping the duplicated TCP ACKs at the AP. It improves on the split-connection approach by preserving the end-to-end semantics and using the soft-state at an AP.

DCF+ [7] is a recent approach to enhancing TCP performance over wireless LANs. DCF+ is based on the WLAN MAC layer, and extends IEEE 802.11 DCF while remaining compatible with it. DCF+ assumes that a recipient has a data frame ready to be transmitted back to the sender. Standard DCF requires that a source STA transmitting a data frame must receive a MAC ACK frame from the destination STA. DCF+ makes the ACK frame act as an RTS frame sent by the destination. Then, the source replies with a CTS and the destination can immediately transfer data frames which are ready to send to the source STA. This procedure ensures that the next transmission in reverse direction follows without contention, and hence reduces the time required for successful transmission.

When several mobile hosts upload packets using TCP, flows with only a small number of packets in flight (e.g. newly started TCP flows) are more susceptible to timeouts than flows with a large number of packets in flight. Small flows can therefore be starved for a long time due to timeouts induced by packet losses. This unfairness of upstream TCP flows can be resolved by using additional flexibilities of the IEEE 802.11e MAC, which allows the values of interframe space (IFS) called Arbitrary IFS (AIFS) and CW_{min} to be set on a per-class basis for each wireless STA. STAs with small values of IFS and CW_{min} have more opportunities for packet transmissions. By reducing IFS and CW_{min} for an AP, and increasing them for other wireless STAs, the AP gets a more equitable share of opportunities to access the wireless channel [8].

III. Optimizing the Effective Number of Competing Stations

A. TCP over IEEE 802.11 wireless LANs

In order to understand why TCP performs poorly in IEEE 802.11 wireless LANs, we need to know that TCP flows in IEEE 802.11 wireless LANs have little to do with network congestion. According to [10] that analyzes the saturation throughput of the IEEE 802.11 DCF, as the number of active STAs[1] increases, the saturation throughput of IEEE 802.11 wireless LANs decreases. However, [9] presented different experimental results that the aggregated throughput of TCP flows is stable even though the number of STAs increases. The reason is that the number of active STAs is nearly the same independently of the number of STAs with a TCP flow. Since TCP adopts a window-based flow control mechanism, a TCP sender does not have any TCP DATA packet to send until it receives TCP ACK packets for transmitted TCP DATA packets. In this situation, the number of active STAs is much smaller than the number of STAs.

To simplify an analysis of this window-based flow control mechanism of TCP in IEEE 802.11 wireless LANs, we assume that all STAs and an AP get a chance to transmit a frame in a round-robin manner due to the fairness of the IEEE 802.11 DCF. Fig. 1 depicts a simple n-state Markov chain that models upstream TCP connections in the IEEE 802.11 infrastructure mode. N STAs are trying to upload TCP DATA packets to N fixed hosts via an AP. We assume that each STA has an average of W TCP packets to upload. We devote the state of the network as $b(L1,\ L2)$, where $L1$ is the sum of the lengths of the queues at all STAs and $L2$ is the length of the queue at the AP. After a STA has transmitted W TCP packets, it cannot contend to transmit another TCP packet until it receives a TCP ACK packet from the AP. By running MATLAB with this model until it reaches a steady state, we determined the probability density function of the number of active STAs participating in the wireless channel contention, and these results are plotted in Fig. 2(a).

This graph shows that the number of active STAs is mostly less than five regardless of the number of STAs that are trying to upload TCP DATA packets, so the network is never saturated. These analytic results explain why the number of STAs does not affect the aggregate throughput of TCP flows in IEEE 802.11 wireless LANs. Our empirical study using an ns-2 simulator [11] verifies this analysis, as shown in Fig. 2(b). However, the number of competing STAs should be increased to achieve an optimal aggregate throughput in IEEE 802.11 wireless LANs.

B. TAP: TCP ACK Priority

We propose to modify the mechanism of MAC-layers contention so as to increase the number of competing STAs in order to achieve a higher aggregate TCP throughput. Suppose that the maximum aggregate throughput of a wireless LAN can be achieved when n^* STAs (including an AP) contend to access the wireless channel. IEEE 802.11 DCF specifies that a STA should perform random backoff for collision avoidance even if there is only one STA in the network. The waste of link bandwidth due to this backoff time is reduced when several STAs are competing for wireless channel access. However, as the number of STAs increases, the increasing frequency of frame collisions degrades the aggregate network throughput. Between there extremes is the optimal number of competing STAs, and we will propose a method of ensuing that this number is achieved in practice.

In our proposed mechanism called TCP ACK Priority (TAP), the AP transmits TCP ACKs to $n^* - 1$ STAs in a way that preempts all other transmissions. This can be achieved if the AP intentionally transmits at the 0^{th} slot after a DIFS, with no backoff time. In order to choose an appropriate value of n^*, we have used the throughput analysis of 802.11 DCF wireless LANs by Bianchi [10].

$$\begin{cases} T_c = \begin{cases} DIFS + SIFS + \frac{PHY\ Header}{Basic\ Rate} + \frac{DATA}{Data\ Rate} \\ DIFS + \frac{RTS}{Basic\ Rate} \end{cases}^{*} \\ \tau = \frac{2(1-2p)}{(1-2p)(W+1)+pW(1-(2p)^m)} \\ p = 1-(1-\tau)^{n-1} \\ \tau = \frac{\sqrt{[n+2(n-1)(T_c^*-1)]/(n-1)}}{(n-1)(T_C^*-1)} \approx \frac{1}{\sqrt[n]{T_c^*/2}}, \end{cases}$$

[1]A STA can be classified into an *active* or *inactive* one. An active STA is defined as a STA having at least one frame to send and participating in the wireless channel contention, while an inactive one does not try to send any frame.

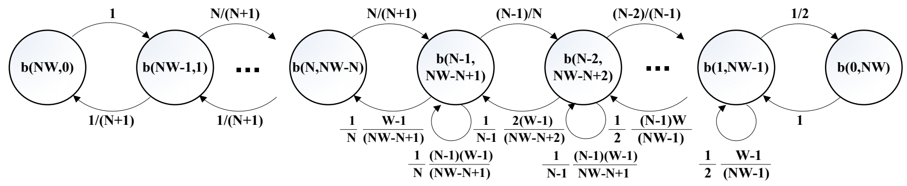

Figure 1: An n-state Markov chain for upstream TCP flows in IEEE 802.11 wireless LANs.

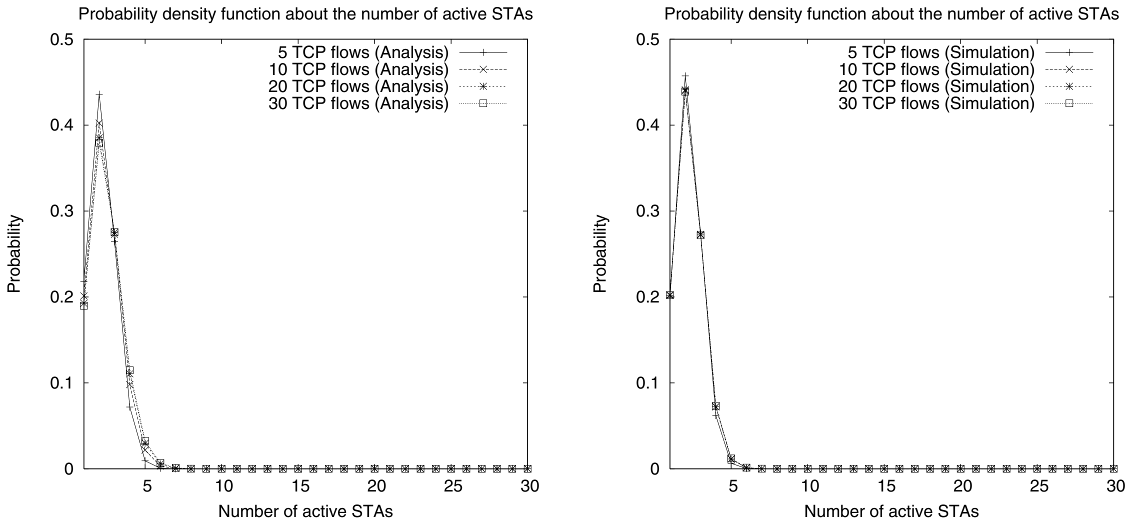

(a) Analysis (b) Simulation

Figure 2: Ratio of competing active STAs.

Where, τ is the probability that a STA transmits in a randomly chosen time-slot and p the conditional collision probability. A frame being transmitted on the channel has a probability p of experiencing a collision. T_c is the time wasted due to a packet collision in a basic access (DATA/ACK) or optional access (RTS/CTS/DATA/ACK) mechanism. For more details about these equations, please refer to Equations (7), (9) and (28) in [10]. From these equations we can determine compute the number of STAs which achieves the maximum throughput.

Fig. 3 shows the optimal number of competing STAs as the sizes of transmitted frames (the sizes of MAC service data units) changes. There are four different bit rates and two access mechanisms. For the purpose of this analysis, all STAs using the basic access mechanism are assumed to have the same data rate. In the basic wireless LAN access mechanism, the collision time is determined by the longest transmission time $(\frac{DATA}{Data\ Rate}^*)$ of all STAs involved in a collision; and we can also see that the optimal number of STAs varies depending on the data rate and the frame size. As the data rate increases and the frame size decreases, the optimal number of STAs becomes larger because the frame transmission time is shorter and so the time wasted due to frame collisions is reduced. However, when the RTS-CTS access mechanism is used, the optimal number of STAs is fixed at 7.526. In this case collision times do not depend on the data rate or on the size of the data frame, since collisions can only occur during RTS-CTS handshaking when

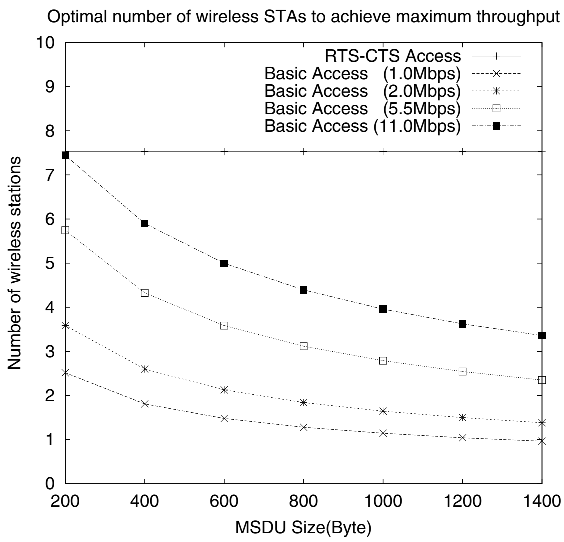

Figure 3: The optimal number of STAs for the maximum throughput.

a fixed frame size and the basic rate are being used. From Fig. 3, we can conclude that n^* should be let adaptively to a value between 1 and 8, depending on the data rate of the STAs, the average size of the frames being transmitted and the type of the

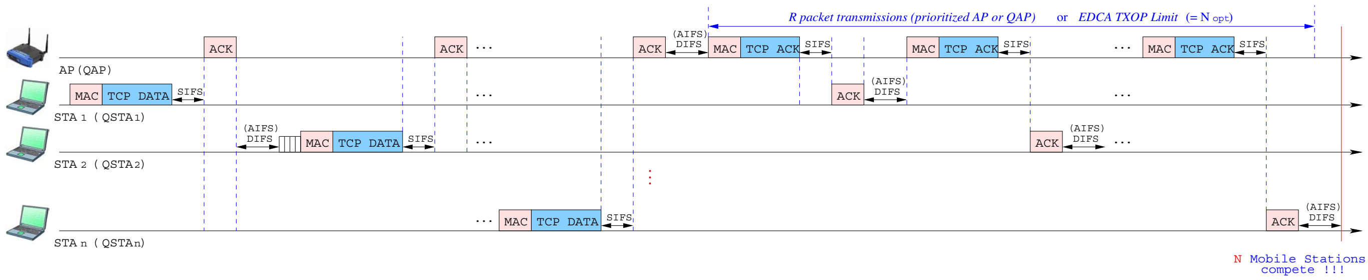

(a) IEEE 802.11 DCF 0^{th}-slot or IEEE 802.11e EDCA TXOP

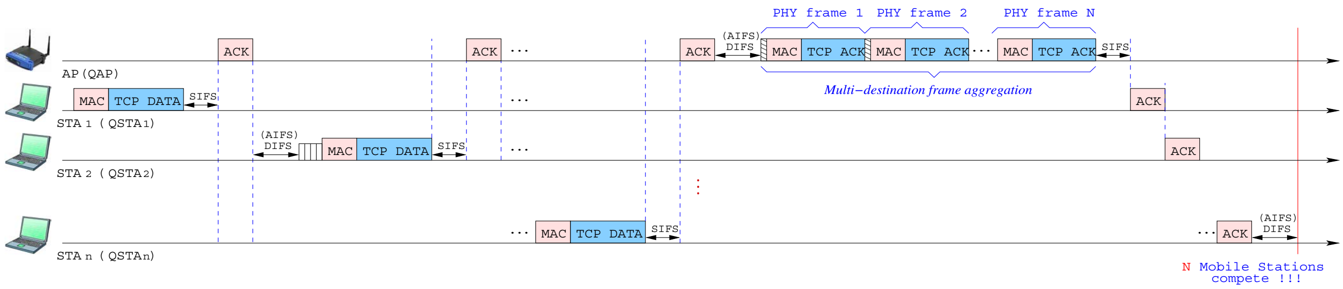

(b) IEEE 802.11n multi-destination frame aggregation

Figure 4: TAP operation for prioritizing TCP ACK packets.

channel access mechanism.

To incorporate TAP into wireless LANs without any modification to IEEE 802.11, we can utilize the new features of IEEE 802.11e/n. The AP can transmit multiple TCP packets successively by using the Transmission Opportunity (TXOP) feature of the IEEE 802.11e standard [2], or the multi-destination frame aggregation facility that is discussed in the IEEE 802.11n draft [3]. After the AP obtains the wireless channel, as described above, sufficient TXOP to transmit $n^* - 1$ TCP packets is given to the AP, as shown in Figure 4(a). Alternatively, several MAC Protocol Data Unit (MPDU) frames involving TCP ACK packets for each STA can be integrated into a single MPDU, which is then transmitted by the AP, as shown in Fig. 4(b). Although multi-destination frame aggregation can avoid overheads such as the need for a SIFS and a MAC header for each TCP packet, several TCP flows will all be interrupted if the aggregated frame is lost.

IV. Performance Evaluation

We evaluated TAP using the ns-2 simulator. The simulated network topology consists of n STAs and an AP in the wireless part of the network and n hosts in the wired part, where n is a variable. Each wireless STA is sending TCP traffic to the corresponding host. The AP and n STAs compete to transmit TCP ACK and TCP DATA packets, respectively. We use the FTP traffic model with 1024-byte TCP DATA packets and every TCP DATA packet is transmitted at 11 Mbps with or without an RTS/CTS mechanism. If an RTS/CTS mechanism is enabled, the RTS/CTS frames are transmitted at 1 Mbps. The total simulation time is 200s in all cases, and we compare IEEE 802.11 DCF and TAP using four performance metrics: the probability density function of the number of active STAs, aggregate throughput, fairness index and delay. The value of n^* is set to 8 for the RTS/CTS mechanism and to 4 for the basic access mechanism (see Fig. 3).

Fig. 5(a) shows how the probability density function of the number of active STAs depends on the number of TCP flows with no RTS/CTS. The average number of active STAs in 802.11 DCF stays around 2 while the number of active STAs in TAP is about 4, which are almost optimal values with these conditions. Fig. 5(b) shows that TAP achieves higher aggregate throughput by making 8 STAs compete for the wireless channel access with the RTS/CTS mechanism and 4 with the basic access mechanism.

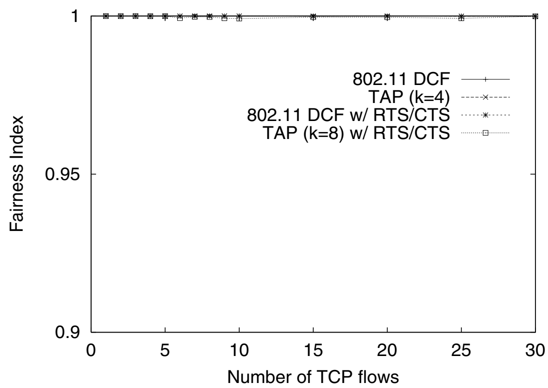

Figure 6: Fairness between upstream TCP flows.

Fig. 6 shows how the long-term fairness (Jain at el, [12]) changes with the number of TCP flows. In IEEE 802.11 DCF,

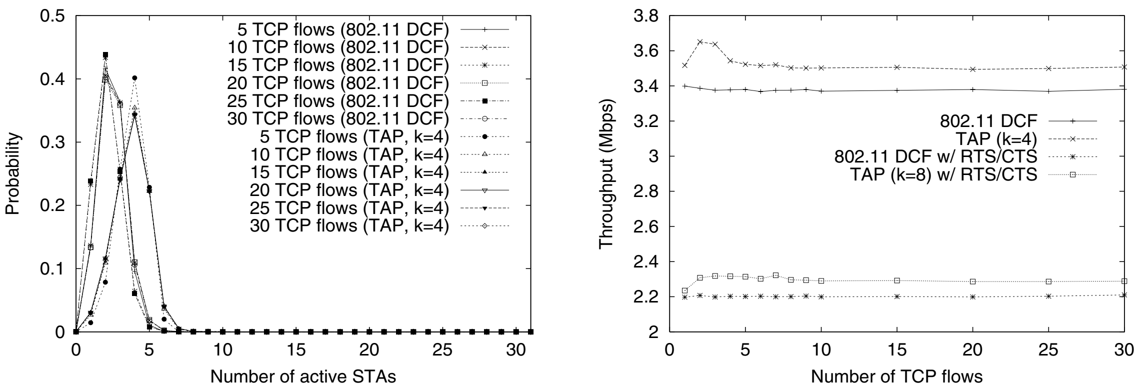

(a) Pdf of the number of active STAs (k=4)

(b) Throughput

Figure 5: Improvement on the performance of TAP.

a TCP sender cannot transmit any TCP DATA packet until it receives a TCP ACK packet, while in TAP, several TCP senders receive TCP ACK packets, which are transmitted with higher priority, so that a short-term fairness problem may occur. However, this result reveals that there are no long-term fairness problem since the fairness index stays around 1 with both the IEEE 802.11 DCF and TAP.

We observed that two behaviors in TAP: (a) more active STAs compete, which results in more collisions and longer delays, and (b) the AP transmits TCP ACK packets with higher priority, which increases TCP DATA packets' delay. Hence, TAP has a little bit increased latency. However, recall that the number of the active STAs is kept as 4 in TAP and 2 in IEEE 802.11 DCF, thus the average delay is maintained almost stable regardless of the number of STAs in the wireless LAN.

V. Conclusion

Our analysis and simulation results suggest that previous models of competing TCP flows in IEEE 802.11 wireless LANs are misleading. We have analyzed the interaction between MAC contention and TCP congestion control. By introducing prioritized access by the access point, the proposed scheme makes the optimal number of competing stations contend for media access. Through ns-2 simulations we have demonstrated that our proposal can achieve a higher aggregate TCP throughput than the original IEEE 802.11 DCF at the cost of slight delay increase.

References

[1] IEEE WG, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer(PHY) Specifications," Standard, IEEE, 1999.

[2] IEEE WG, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: Medium Access Control (MAC) Quality of Service Enhancements," Standard, IEEE, 2005.

[3] IEEE WG, "TGn Sync Proposal Technical Specification," IEEE 802.11-04/889r2, 2005.

[4] M. Allman, V. Paxson and W. Stevens, TCP Congestion Control, RFC 2581, Internet Engineering Task Force (IETF), April 1999.

[5] A.Bakre and B.R.Badrinath, "I-TCP: indirect TCP for mobile hosts," Proc. Distributed Computing Systems, 1995.

[6] H.Balakrishnan, S.Seshan, and R.H.Katz, "Improving reliable transport and handoff performance in cellular wireless networks," ACM Wireless Networks, vol.1, pp.469-481, December 1995.

[7] H.Wu, Y.Peng, K.Long, S.Ceng, and J.Ma, "Performance of reliable transport protocol over IEEE 802.11 wireless LAN: analysis and enhancement," Proc. INFOCOM, 2002.

[8] D.J. Leithm and P. Clifford, "Using the 802.11e EDCF to achieve TCP upload fairness over WLAN links," WIOPT, 2005.

[9] S. Choi, K. Park, and C. Kim, "On the performance characteristics of WLANs: revisited," Proc. ACM SIGMETRICS, 2005.

[10] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," IEEE Journal on Selected Areas in Communications, Vol. 18, No. 3, March 2000.

[11] VINT group, UCB/LBNL/VINT Network Simulator : ns (version 2), http://www.isi.edu/nsnam/ns/.

[12] R. Jain, G. Babic, B. Nagendra, and C. Lam, "Fairness, call establishment latency and other performance metrics," Tech. Rep. ATM Forum/96-1173, ATM Forum Document, August 1996.