Queue Management Strategies to Improve TCP Fairness in IEEE 802.11 Wireless LANs

Mingwei Gong Qian Wu Carey Williamson
Department of Computer Science
University of Calgary
Email: {gongm,qianwu,carey}@cpsc.ucalgary.ca

***Abstract*— Wireless Local Area Networks (WLANs) based on the IEEE 802.11 technology have become increasingly popular and ubiquitous. The 802.11 standard allows each station in a WLAN equal opportunity to access the wireless channel, which can result in unfair sharing of network bandwidth between upstream and downstream TCP flows at an AP. In this paper, we propose two different queue management techniques to alleviate the unfairness problem, with one based on Selective Packet Marking (SPM), and the other based on Least Attained Service (LAS) scheduling. We evaluate these proposed solutions using the *ns-2* network simulator. The simulation results show that, compared to a conventional DropTail queue mechanism for NewReno TCP sources, the proposed solutions improve the fairness index by 20-40%, while achieving comparable aggregate throughput.**

Keywords: IEEE 802.11 WLAN, TCP fairness, ns-2

I. INTRODUCTION

In recent years, IEEE 802.11 Wireless Local Area Networks (WLANs) have become increasingly prevalent. WLANs provide wireless access to the Internet, providing a viable alternative to Ethernet LAN connectivity.

WLANs can operate in either *infrastructure* mode or *ad hoc* mode, though our work focuses only on infrastructure mode. In a typical deployment, a mobile station equipped with an 802.11 interface communicates with an Access Point (AP) on a wireless channel, and the AP relays traffic to and from the Internet backbone. The WLAN is typically the throughput bottleneck in these scenarios [15]. Hence, the limited wireless bandwidth needs to be allocated fairly and efficiently amongst the mobile stations and the AP.

The IEEE 802.11 standard defines a mandatory contention-based channel access protocol called Distributed Coordination Function (DCF). The DCF protocol provides all WLAN stations equal opportunity to access the transmission medium.

Unfortunately, this station-based fairness mechanism, when coupled with the Internet's Transmission Control Protocol (TCP) for reliable data delivery, can lead to flow-level unfairness [22]. In particular, an upstream flow can attain higher throughput than a downstream flow.

In this paper, we focus on fair bandwidth allocation between upstream TCP flows generated by mobile stations and downstream TCP flows destined to mobile stations. Two independent queue-based strategies are proposed in this paper to improve fairness while maintaining aggregate system throughput. The first approach is called Selective Packet Marking with ACK Filtering (SPM-AF). This approach sets a priority level for each outgoing packet, and manages packets differently based on their priority levels. The second approach is based on Least Attained Service (LAS) scheduling, which gives precedence to flows that have received less service. Hence, fairness can be achieved amongst the flows sharing the AP.

The remainder of this paper is organized as follows. Section II illustrates the problem of TCP fairness over 802.11 WLANs. Section III describes two proposed approaches for solving the fairness problem. Section IV presents the evaluation methodology and the simulation results. Section V reviews prior related work. Section VI concludes the paper.

II. PROBLEM ILLUSTRATION

In an infrastructure-based WLAN, the AP acts as a bridge to relay traffic between the wireless local network and the wired network. However, the DCF mechanism used by WLAN cards does not provide higher priority for the AP to access the wireless transmission medium, which causes a bottleneck at the AP. This issue can lead to pronounced unfairness between upstream and downstream TCP flows.

Consider a scenario in which there are N upstream TCP flows and N downstream TCP flows, where $1 < N < 30$. For simplicity, assume that all flows are bidirectional, and send an equal number of TCP packets in each direction. However, the upstream flows send large TCP data packets to the AP and receive small TCP ACK packets from the AP, while the downstream flows receive large TCP data packets from the AP and send small TCP ACK packets to the AP. For the rest of this discussion, we focus only on the large TCP data packets, since they have the dominant influence on network throughput.

Since the AP and the mobile stations share the wireless channel, each of them transmits $\frac{1}{N+1}$ of the total data packets, on average. However, the $\frac{1}{N+1}$ share that the AP obtains for channel access is further partitioned across the N downstream flows, making each downstream flow progress slowly.

The equal opportunity nature of DCF makes the downstream queue at the AP a bottleneck [5]. When upstream and downstream TCP flows are both present in the network, the arrivals to the bottleneck queue include the data packets for the downstream flows and the TCP ACK packets for the upstream flows. If the arrival rate to the queue exceeds the service rate at which the AP transmits packets onto the WLAN, then the bottleneck queue will fill and overflow, causing packet losses.

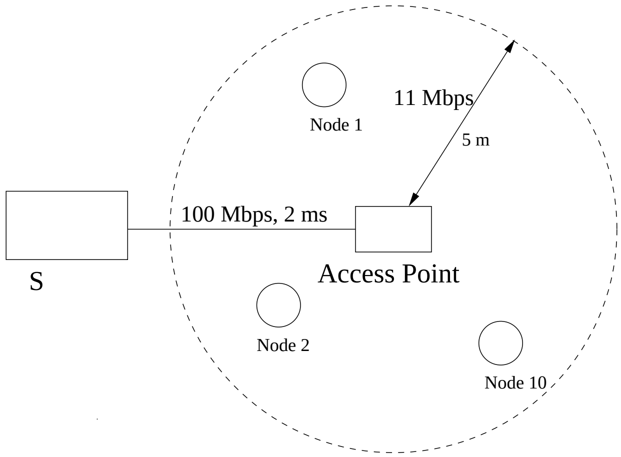

Fig. 1. Network Model for ns-2 Simulations

For TCP flows, loss of a data packet and loss of an ACK have different impacts on performance. The loss of a data packet can be detected either by a timeout or by the more efficient "triple duplicate ACK" mechanism. In either case, the loss typically triggers reduction of the congestion window at the TCP sender, resulting in lower throughput. The loss of an ACK, on the other hand, is largely irrelevant: a subsequent ACK conveys the lost information, because of the cumulative nature of TCP acknowledgments. Therefore, the loss of an ACK typically has negligible impact on the throughput.

To illustrate the unfairness problem, we present an experiment using the *ns-2* network simulator [21]. All simulations use the TCP NewReno model, with delayed ACKs enabled. In steady state, each flow sends one ACK for every two TCP data segments received.

The network model used in the experiment is shown in Figure 1. A wired station S is connected to an AP via 100 Mbps Ethernet. Ten mobile stations are placed randomly within 5 meters of the AP. Five mobile stations (Node 1 to Node 5) are TCP senders. They upload data to S using fixed-size 1500-byte packets. The other five mobile stations (Node 6 to Node 10) are TCP clients, downloading data from S using the same TCP data packet size as the upstream flows.

The application layers for both downstream and upstream transfers have infinite data to send (i.e., FTP-like bulk transfers). The data transmission rate used in communications between the AP and the mobile stations is fixed at 11 Mbps. The buffer size at the AP's downstream bottleneck link is varied from 10 packets to 150 packets, in steps of 10. Each experiment runs for 300 seconds of simulated time.

Figure 2 shows the throughput results for each of the upstream and downstream flows. The average throughput for the upstream flows is much higher than that for the downstream flows, especially for small buffer sizes.

The unfairness phenomenon is alleviated as the buffer size increases. In fact, upstream flows and downstream flows achieve comparable throughput when the buffer size reaches 150 packets.

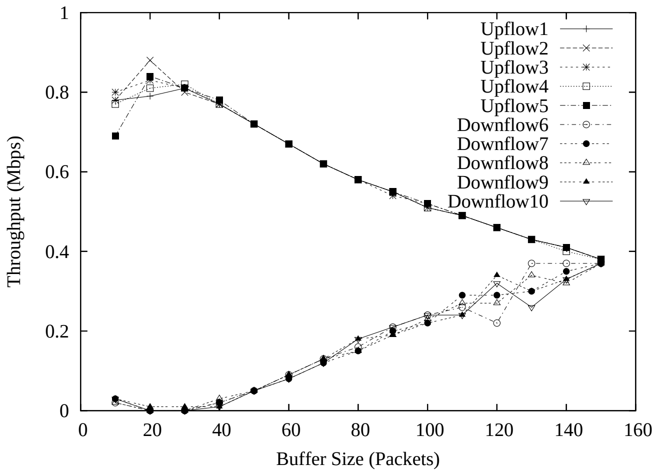

Fig. 2. Throughput Results for Upstream and Downstream Flows

One of the reasons for the improved fairness is explained by Bruno *et al.* [5]. In particular, because TCP uses a closed-loop flow control strategy, the presence of the AP bottleneck implicitly regulates the number of contending stations on the WLAN. These interactions between TCP flow control and the MAC channel access mechanism lead to improved fairness as the AP buffer size increases [5], [6].

A second reason for the improved fairness is the lower packet loss rate. In our experiment, fair sharing occurs when the buffer size is large because packet losses no longer occur. For each downstream flow, the maximum number of outstanding packets is bounded by the receiver's advertised window, which is 20 packets by default in *ns-2*. For 5 flows, this total is at most 100 data packets. Since delayed ACKs are used, the maximum number of ACKs at the AP is 10 per upstream flow. For 5 flows, this total is at most 50 ACK packets. Hence, the total number of packets present simultaneously at the AP's downstream queue is at most 150. In other words, when the buffer size is 150 packets or more, all TCP senders can reach their maximum window size without losing any packets at the bottleneck queue. Thus, the unfairness problem is solved. However, setting the buffer size this large can lead to a large queueing delay (e.g., several hundred milliseconds for an 802.11b WLAN).

Figure 3 shows the congestion window evolution for one upstream flow (from Node 1) and one downstream flow (to Node 6). The buffer size used for this experiment is 60 packets. The asymmetric growth for the two flows is evident. The congestion window grows quickly for the upstream flow, showing that the losses of ACKs have negligible effect on the congestion window of the TCP sender (Node 1). As the congestion window size of the upstream flow grows, the downstream flow is disadvantaged. The congestion window size for the downstream flow stays around 2 packets because of many packet losses.

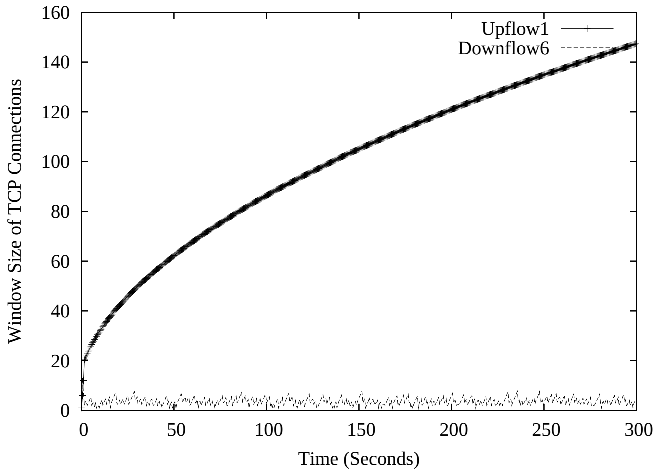

Fig. 3. Congestion Window for Upstream and Downstream Flow

III. QUEUE MANAGEMENT STRATEGIES

As shown in Section II, the root cause for the unfairness is that the loss of a data packet has more adverse impacts on TCP performance than the loss of an ACK packet. This motivates our two separate approaches to queue management. The first approach, called Selective Packet Marking with ACK Filtering (SPM-AF), differentiates TCP data packets from TCP ACKs, giving the data packets higher priority to enter the bottleneck queue. The second approach, called LAS, gives higher priority to flows that have received less service relative to other flows. The following two subsections describe these two approaches.

A. SPM-AF Queue Management Scheme

The SPM-AF algorithm borrows ideas from Selective Packet Marking (SPM) proposed in [26] and from ACK Filtering (AF) proposed in [1]. It requires changes to the TCP source and the queue management.

The key idea for the SPM-AF approach is to give data packets greater opportunity to enter the bottleneck queue. More specifically, the TCP source implementing the packet marking scheme sets the priority level for each output data packet based on the congestion window size. As presented in [26], in a TCP connection, some packets (i.e., any packets sent when the congestion window is smaller than 4 packets) are more crucial than others, since losses of crucial packets incur a coarse-grain timeout, while losses of other packets trigger a *fast retransmit* to recover the missing packet. Within each flow, the TCP SPM model marks the crucial data packets “high” priority. All other packets are marked with default “low” priority.

The AF queue management mechanism is used at the AP’s downstream bottleneck queue. AF treats TCP data packets and TCP ACKs differently. When an ACK arrives at the bottleneck queue, the AF policy scans the queue to check if there are any other ACKs for the same connection. If so, all of these previous ACKs are removed from the queue, relying on the cumulative acknowledgment nature of the newly queued ACK to supersede the information in the previously queued (but now removed) ACKs. As a result, some buffer space is freed for use by data packets of downstream flows.

The placement of the consolidated ACK in the queue is configurable in AF. It can be placed either at the tail of the queue (“Tail Filtering”) or at the location of the oldest ACK found (“Head Filtering”).

AF also differentiates the high priority data packets from the low priority data packets. A low-priority arriving packet that encounters a full queue is dropped. A high-priority arriving packet that encounters a full queue may or may not be dropped, depending on the state of the queue. In particular, if there is at least one ACK or low-priority data packet in the queue, then one such packet is removed (if both types are present, then an ACK is chosen for removal), making room for the high-priority packet to be added at the tail of the queue. If the full queue contains no ACKs and no low-priority data packets, then the arriving packet is always dropped.

The reduced ACK frequency could slow down the growth of the sender’s congestion window. To counteract this problem, the TCP sender increases the congestion window by counting how many segments are acknowledged in an ACK, rather than counting the number of ACKs received.

B. Least Attained Service Scheduling

The second proposed queue management mechanism is based on Least Attained Service (LAS) scheduling. LAS scheduling is a size-based scheduling policy that has been studied extensively in recent years.

LAS is a multi-level scheduling policy that always gives service to the job that has received the least service so far. LAS is also known as *Foreground-Background* (FB) or *Shortest Elapsed Time* (SET) first scheduling.

Recently, the LAS scheduling policy is proposed at a network router [23], [24]. A network router can easily identify a network flow by its source and destination addresses and ports. The router uses a counter to keep track of the amount of service attained by a flow, and inserts the newly arriving packet into the priority queue. Studies of LAS scheduling show that LAS favours short jobs, while negligibly penalizing large jobs [23], [24].

The implementation of LAS is quite straightforward. When a new packet arrives at the queue, it is inserted into the appropriate position according to its sequence number. For a particular flow that has received less service than its peers, the incoming packets of the flow are inserted closer to the head of the queue. Moreover, if the queue is full when a packet arrives, a drop occurs. The DropTail queue simply discards the arriving packet, while the LAS queue discards the lowest priority one, which may be the arriving packet or the packet at the tail of the queue. Thus, a flow that has received less service so far avoids losing any packets until it achieves its fair share. Eventually, the flow will catch up with others.

IV. EVALUATION METHODOLOGY

A. Simulation Design

We evaluate our proposed solutions using simulation, with the *ns-2* network simulator (version 2.28). The network model used in the evaluation is the same as that shown in Figure 1. The conventional TCP NewReno model and DropTail queue management are used as the baseline in our simulation study.

The performance metrics of interest are the aggregate TCP throughput and the TCP fairness index [14]. Higher values of both metrics are desirable. We expect our proposed solutions to show similar aggregate TCP throughput, but improved fairness.

B. Simulation Results

Figure 4 shows the simulation results for the network model described in Section II. The horizontal axis in each graph shows the buffer size (in packets) at the AP's downstream bottleneck queue. The buffer size is varied from 10 packets to 150 packets. The vertical axis in Figure 4(a) shows the aggregate throughput of 10 TCP flows (upstream and downstream). Figure 4(b) shows the fairness index results.

There are four different lines represented on each graph in Figure 4. The line labeled "NewReno-DropTail" is for the conventional TCP NewReno model with DropTail queue management. The "NewReno-LAS" line is for TCP NewReno with LAS queue management. The line labeled "SPM-AF Head" represents the "Head Filtering" variation of the SPM-AF queue management described in Section III-A. The "SPM-AF Tail" line is for the "Tail Filtering" version of SPM-AF.

Four important observations are evident in Figure 4. First, as expected, the proposed approaches achieve similar aggregate throughput. Figure 4(a) shows that the total throughput for all tested strategies is roughly the same, though the SPM-AF schemes achieve slightly higher throughput. Second, the fairness index improves as the buffer size is increased. This result is consistent for most cases studied. Third, compared to the NewReno-DropTail combination, the proposed approaches improve the fairness index significantly, especially when the buffer size is small. The average performance gains are 24% for the "Head Filtering" version of SPM-AF, 32% for the "Tail Filtering" version of the SPM-AF scheme, and 44% for LAS scheduling. Fourth, for the two proposed queue management approaches, LAS scheduling provides the best fairness, and the "Tail Filtering" version of SPM-AF outperforms the "Head Filtering" version with respect to the fairness index. The performance difference for the latter is evident when the buffer size is smaller than 80 packets.

The "Head Filtering" approach performs worse than "Tail Filtering" because it exacerbates the unfairness between upstream and downstream flows. Because "Head Filtering" replaces the oldest ACK with the newly arrived ACK for the same flow, the congestion window for that flow is increased sooner than with the "Tail Filtering" version. Therefore, the upstream flows are more aggressive. They achieve higher throughput, resulting in worse fairness. To illustrate this phenomenon, Table I shows the detailed throughput results for each flow when the buffer size is 60 packets. Head Filtering increases the throughput of the upstream flows.

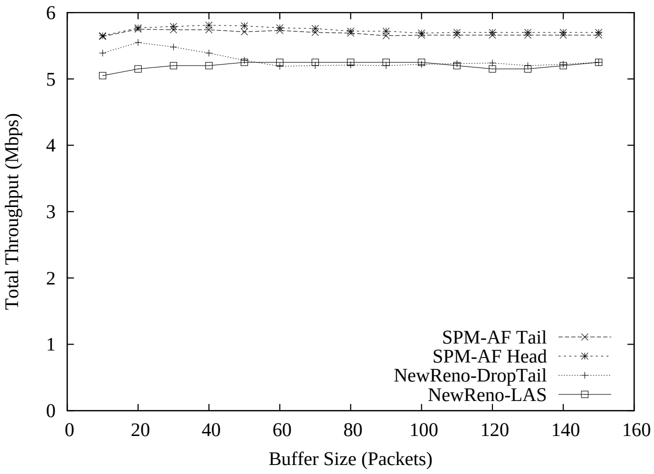

(a) Aggregate Throughput

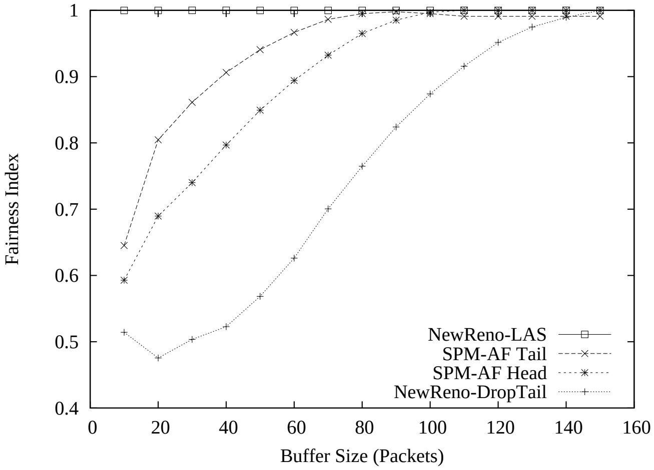

(b) Fairness Index

Fig. 4. Simulation Results for Throughput and Fairness

The fairness index generally improves for all strategies when the buffer size is increased, because there are fewer packet losses. Figure 5 shows the packet loss ratio at the bottleneck buffer for the four schemes in the study. With more space in the buffer, fewer packets are dropped. Hence, the throughput for downstream TCP traffic is increased, leading to better sharing of the network bandwidth. Figure 5 also shows that the SPM-AF scheme has greater loss than the other two schemes, since SPM-AF can remove ACKs even before the buffer overflows.

The fairness advantages for SPM-AF come from the ACK filtering, as well as the priority marking of data packets. With some ACKs removed from the queue, the buffer space is freed to accommodate the data packets. Moreover, protecting crucial data packets from dropping can avoid the "expensive" time-

TABLE I

SPM-AF THROUGHPUT COMPARISON: TAIL VS. HEAD

		Tail Filtering	Head Filtering
Upstream Flow Throughput (Mbps)	Flow 1	0.68	0.77
	Flow 2	0.67	0.78
	Flow 3	0.68	0.78
	Flow 4	0.68	0.76
	Flow 5	0.68	0.78
Downstream Flow Throughput (Mbps)	Flow 6	0.44	0.34
	Flow 7	0.45	0.42
	Flow 8	0.47	0.39
	Flow 9	0.51	0.34
	Flow 10	0.47	0.41

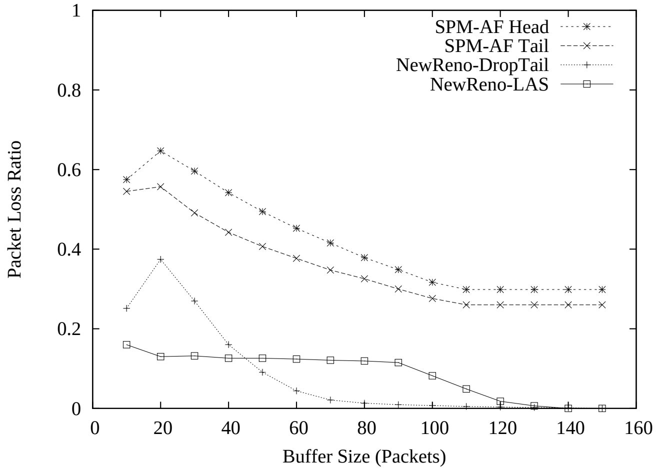

Fig. 5. Packet Loss Ratio Results

out recovery. As a result, the throughput of the downstream flows are improved. As for the LAS scheme, it guarantees that the flow with the least service is served first. Thus, all flows can be treated fairly. That is why LAS performs the best in terms of fairness among the schemes evaluated.

Figure 6 shows the queue status in terms of the number of ACKs and the number of data packets, as sampled every 1 second in the bottleneck queue when the buffer size is 60 packets. The queueing dynamics are quite different for the studied schemes:

- For the conventional NewReno-DropTail scheme, due to the cumulative ACK property, the upstream flows see little effect of congestion and keep sending as if nothing happened. Thus, the ACK packets occupy more than two thirds of the buffer space (see Figure 6(a)).
- For the LAS scheme, the upstream senders are implicitly regulated. That is, the more ACK packets there are in the queue, the more likely they are inserted at the back of the queue. This protects the more sensitive data packets from loss. In steady state, upstream flows and downstream flows receive fair service. There are usually twice as many data packets in the queue as ACK packets (see Figure 6(b)), because TCP delayed ACKs are used.
- For the SPM-AF schemes in Figure 6(c) and (d), the ACKs occupy about 10% of the buffer during the simulation. Note that the reduced ACK frequency does not hinder congestion window growth for the upstream flows, since the TCP SPM model increases the congestion window by counting how many segments are acknowledged in an ACK, as described in Section III-A.

Figure 7 shows the evolution of the TCP congestion window size versus time. Figure 7(a) shows the simulation results for LAS, while Figure 7(b)). shows the results for the Tail Filtering version of SPM-AF. The buffer size limit is 60 packets.

The TCP congestion window size is strongly affected by the queue management strategy. Recall that in the NewReno-DropTail scenario (Figure 3), the congestion window size for the representative upstream flow grows in an unbounded fashion, to the detriment of the downstream flow. In contrast, Figure 7(a) shows that LAS scheduling constrains the window size growth for the upstream flow, allowing the downstream flow to achieve reasonable throughput. A similar observation holds for SPM-AF in Figure 7(b), although the congestion window size for the upstream flow can still grow quite large.

Since downstream flows are more prevalent in most WLANs, we also test our proposed schemes in a scenario with one upstream flow and multiple (from 5 to 25) downstream flows. The simulation results are shown in Figure 8. These results show that unfairness occurs even with a single upstream flow. The more downstream flows there are, the worse the unfairness problem is. The SPM-AF scheme improves both aggregate throughput and fairness, compared to the conventional DropTail scheme. LAS provides a tradeoff between fairness and throughput. LAS has perfect fairness performance under all situations, while its aggregate throughput drops slightly.

V. RELATED WORK AND DISCUSSION

Fairness issues in wireless LANs have been studied extensively [3], [8], [11], [12], [17], [18], [19], [20].

Fairness problems can arise because of MAC protocol mechanisms. For example, the backoff procedure used by the 802.11 DCF protocol doubles the contention window (CW) whenever an attempted transmission fails. A station with a larger contention window has a lower probability to access the medium. Several enhancements to 802.11 DCF have been proposed [3], [8], [12], [19]. These solutions achieve fair allocation of the wireless channel bandwidth by adjusting the contention window of each station dynamically.

The authors in [17], [20] conducted network measurements in an IEEE 802.11e WLAN environment. They proposed an approach that gives the AP higher priority to access the medium by setting the TXOP, AIFS, and CWmin parameters to alleviate the unfairness between upstream flows and downstream flows.

Fairness issues in TCP and MAC protocol interaction are discussed in [2], [4], [10], [17], [18], [20], [22]. Pilosof *et al.* [22] observed unfair sharing of the network bandwidth between upstream flows and downstream flows through network measurements. A comprehensive simulation study was also

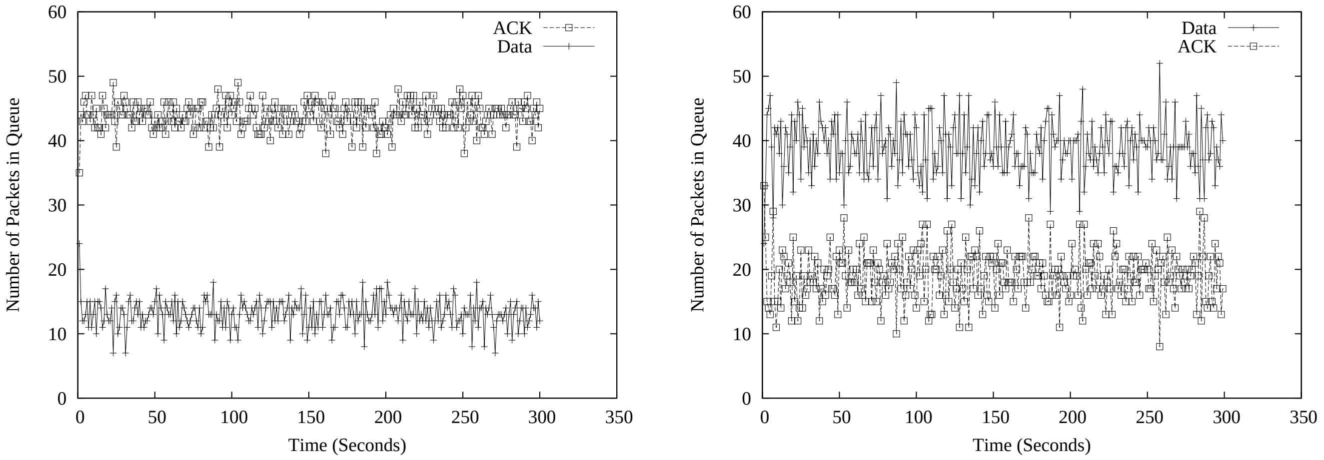

(a) NewReno-DropTail

(b) NewReno-LAS

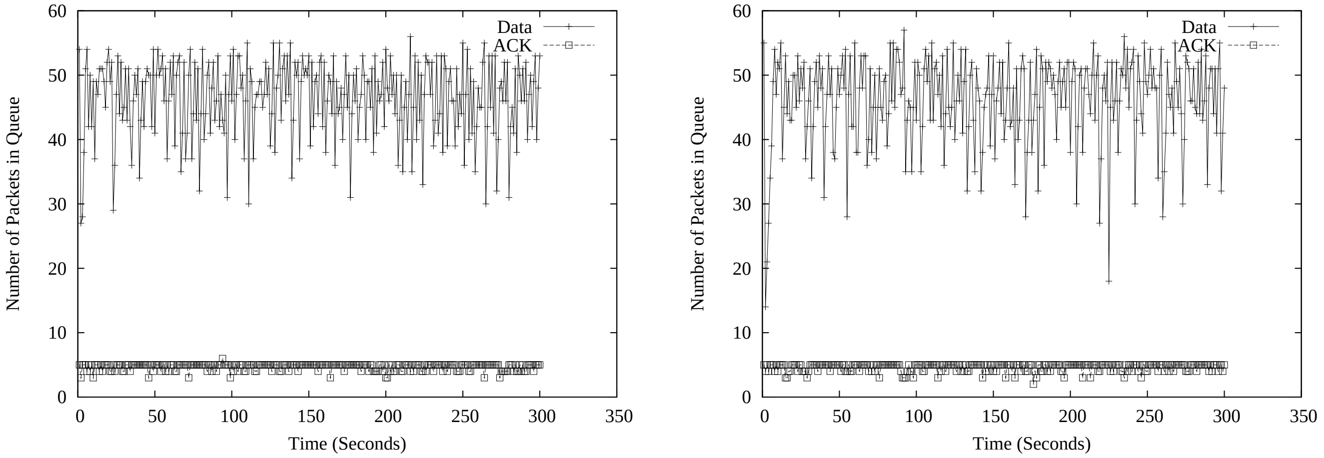

(c) SPM-AF Tail Filtering

(b) SPM-AF Head Filtering

Fig. 6. Queue Behavior for Each Scheme

conducted to identify the causes of unfairness. The authors reported that the buffer size at the AP plays an important role in the wireless channel bandwidth allocation. They also proposed a solution that manipulates the advertised receiver windows of the upstream flows to limit the upstream traffic. As a result, the downstream flows can obtain their fair share of the bandwidth.

In this paper, we propose two separate queue management schemes to address the upstream/downstream unfairness. Our solutions operate at the network and transport layers, in contrast to the solutions proposed in [17], [20], which focus on the MAC layer. The closest work to ours is the one presented in [18]. Their approach is based on virtual queue management. Each flow at the AP corresponds to a virtual queue (VQ). By serving each VQ in a round-robin fashion, fairness can be achieved between upstream flows and downstream flows. However, this scheme requires per-flow state information.

The main challenge for our proposed schemes is the overhead induced at the AP. The SPM-AF approach requires selective marking of packets at the TCP sources, which can be done either at the transport layer or via DiffServ mechanisms [27]. The queue management at the AP must consider the packet markings, requiring extra processing for ACK filtering. The LAS approach requires sequence numbers, either at the transport layer or a lower layer of the protocol stack. With this information, the LAS approach offers greater control than traditional queue management strategies, including the ability to regulate the number of active flows and protect short-lived flows that are prevalent in Web browsing.

Another issue in *multi-rate* WLAN environments is time-based fairness versus throughput-based fairness [7]. As presented in [11], the throughput of a high-rate station suffers from the presence of a low-rate station. Although the high-rate station and the low-rate station achieve similar throughput, the overall network efficiency is degraded. Several researchers

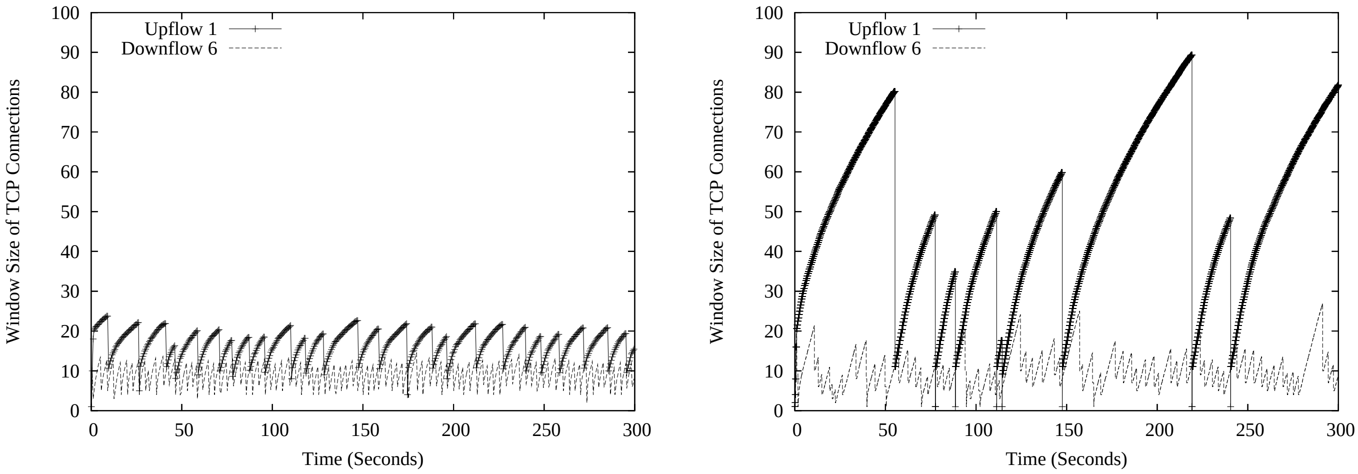

(a) NewReno-LAS

(b) SPM-AF Tail Filtering

Fig. 7. Simulation Results for Congestion Window Size Evolution

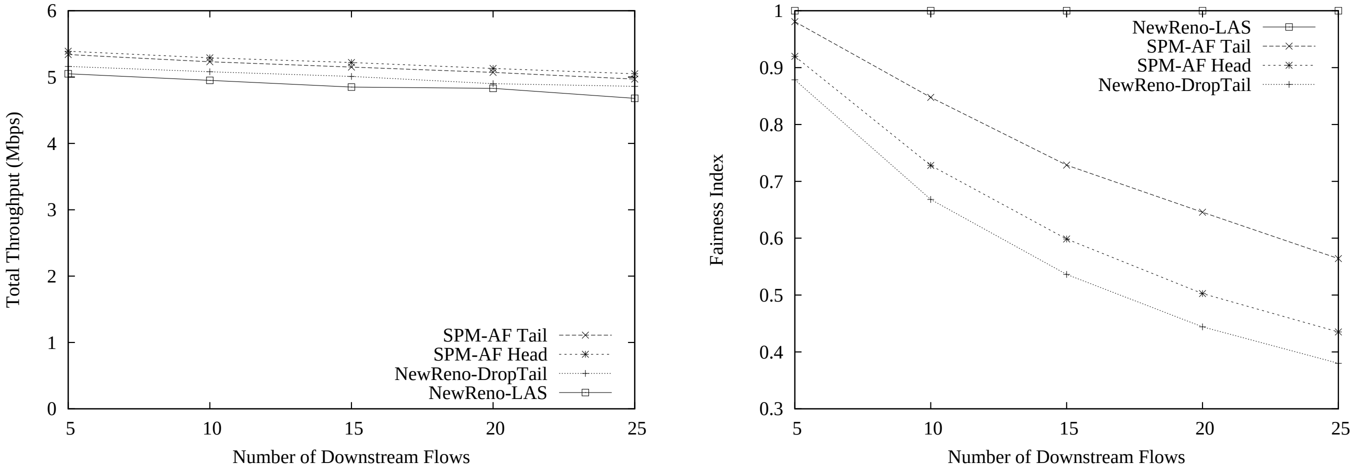

(a) Aggregate Throughput

(b) Fairness Index

Fig. 8. Simulation Results for 1 Upstream, N Downstream Flows

have proposed solutions that use time-based fairness instead of throughput-based fairness in multi-rate WLANs. By allocating more network bandwidth to high-rate stations, the overall system efficiency improves [9], [12], [13], [16], [25].

VI. SUMMARY AND CONCLUSIONS

The 802.11 standard offers each station in a WLAN equal opportunity to access the wireless channel. This can result in unfair bandwidth allocations between upstream TCP flows and downstream TCP flows sharing an AP.

In this paper, we propose two independent queue management approaches to alleviate the unfairness. One approach is based on Selective Packet Marking with ACK Filtering (SPM-AF), while the other is based on Least Attained Service (LAS) scheduling.

We evaluate both proposed solutions using the *ns-2* network simulator. The simulation model assumes a typical infrastructure WLAN configuration, with an AP connected to a wired network and ten mobile hosts in the WLAN.

The simulation results show that the proposed solution improves the fairness index by 20-40% compared to the conventional DropTail queue mechanism, while achieving comparable throughput. Both approaches successfully solve the unfairness problem. The SPM-AF approach produces slightly higher throughput than the LAS approach, while the LAS approach produces fairness superior to SPM-AF. The choice of solutions depends on the performance metric of interest, as well as the network and traffic configuration in the WLAN.

Our future work will study fairness issues for more general network traffic scenarios in WLAN environments. For example, another fairness issue that has been studied recently is short-term fairness for TCP flows. Since short-lived TCP flows are very sensitive to losses during the early stages of TCP congestion window growth, it is important to protect short-lived TCP flows during the critical small window regime. Both SPM and LAS schemes show promise for improving performance of short-lived TCP flows in wired networks. Our ongoing work is studying how well these schemes work in the WLAN context, particularly with dynamic traffic flows.

ACKNOWLEDGEMENTS

Financial support for this research was provided by iCORE (Informatics Circle of Research Excellence) in the Province of Alberta, as well as NSERC (Natural Sciences and Engineering Research Council) in Canada. The authors thank the anonymous RAWNET reviewers for their constructive feedback on the original version of this paper.

REFERENCES

[1] H. Balakrishnan, V. Padmanabhan, and R. Katz, “The Effect of Asymmetry on TCP Performance”, *Mobile Networks and Applications (MONET)*, Vol. 4, No. 3, pp. 219-241, October 1999.

[2] A. Banchs, A. Azcorra, C. Garcia, and R. Cuevas, “Applications and Challenges of the 802.11e EDCA Mechanism: An Experimental Study”, *IEEE Network*, Vol. 19, No. 4, pp. 52-58, July 2005.

[3] G. Bianchi and I. Tinnirello, “Kalman Filter Estimation of the Number of Competing Terminals in an IEEE 802.11 Network”, *Proceedings of IEEE INFOCOM*, San Francisco, CA, pp. 844-852, March 2003.

[4] M. Bottigliengo, C. Casetti, C. Chiasserini, and M. Meo, “Short-term Fairness for TCP Flows in 802.11b WLANs”, *Proceedings of IEEE INFOCOM*, Hong Kong, China, pp. 1383-1392, March 2004.

[5] R. Bruno, M. Conti, and E. Gregori, “Throughput Evaluation and Enhancement of TCP Clients in Wi-Fi Hot Spots”, *Proceedings of the First Working Conference on Wireless On-demand Network Systems* (WONS), Madonna di Campiglio, Italy, pp. 73-86, January 2004.

[6] R. Bruno, M. Conti, and E. Gregori, “Analytical Modeling of TCP Clients in Wi-Fi Hot Spots”, *Proceedings of IFIP Networking Conference*, Athens, Greece, pp. 626-637, May 2004.

[7] R. Bruno, M. Conti, and E. Gregori, “Throughput vs. Temporal Fair MAC Protocols in Multi-Rate WLANs: Analysis and Performance Evaluation”, *Proceedings of IEEE Vehicular Technology Conference* (VTC-Spring), Milan, Italy, May 2004.

[8] F. Cali, M. Conti, and E. Gregori, “Dynamic Tuning of the IEEE 802.11 Protocol”, *IEEE/ACM Transactions on Networking*, Vol. 8, No. 6, pp. 785-799, December 2000.

[9] G. Cantieni, Q. Ni, C. Barakat, and T. Turletti, “Performance Analysis under Finite Load and Improvements for Multirate 802.11”, *Computer Communications*, Vol. 28, No. 10, pp. 1095-1109, June 2005.

[10] A. Grilo and M. Nunes, “Performance Evaluation of IEEE 802.11e”, *Proceedings of the 13th IEEE International Symposium on Personal, Indoor, and Mobile Radio Communications*, Lisbon, Portugal, pp. 511-517, September 2002.

[11] M. Heusse, F. Rousseau, G. Berge-Dabbatel, and A. Duda, “Performance Anomaly of 802.11b”, *Proceedings of IEEE INFOCOM*, San Francisco, CA, pp. 836-843, March 2003.

[12] M. Heusse, F. Rousseau, G. Berge-Dabbatel, and A. Duda, “Idle Sense: An Optimal Access Method for High Throughput and Fairness in Rate Diverse Wireless LANs”, *Proceedings of ACM SIGCOMM*, Philadelphia, PA, pp. 121-132, August 2005.

[13] G. Holland, N. Vaidya, and P. Bahl, “A Rate-Adaptive MAC Protocol for Multi-Hop Wireless Networks”, *Proceedings of ACM/IEEE MOBICOM*, Rome, Italy, pp. 236-251, July 2001.

[14] R. Jain, *The Art of Computer Systems Performance Analysis*, Wiley & Sons, New York, NY, 1991.

[15] J. Jun, P. Peddabachagari, and M. Sichitiu, “Theoretical Maximum Throughput of IEEE 802.11 and its Applications”, *Proceedings of the 2nd IEEE International Symposium on Network Computing and Applications* (NCA’03), Cambridge, MA, pp. 249-256, April 2003.

[16] T. Kuang, Q. Wu, and C. Williamson, “MRMC: A Multi-Rate Multi-Channel MAC Protocol for Multi-Radio Wireless LANs”, *Proceedings of the 2005 Workshop on Wireless Networks and Communication Systems* (WiNCS), Philadelphia, PA, pp. 263-272, July 2005.

[17] D. Leith and P. Clifford, “Using the 802.11e EDCF to Achieve TCP Upload Fairness over WLAN Links”, *Proceedings of the Third International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, Trentino, Italy, pp. 109-118, April 2005.

[18] X. Lin, X. Chang, and J. Muppala, “VQ-RED: An Efficient Virtual Queue Management Approach to Improve Fairness in Infrastructure WLAN”, *Proceedings of IEEE Workshop on Wireless Local Networks (WLN)*, Sydney, Australia, pp. 632-638, November 2005.

[19] H. Ma, H. Li, P. Zhang, S. Luo, C. Yuan, and X. Li, “Dynamic Optimization of IEEE 802.11 CSMA/CA Based on the Number of Competing Stations”, *Proceedings of IEEE International Conference on Communications (ICC)*, Paris, France, pp. 191-195, June 2004.

[20] A. Ng, D. Malone, and D. Leith, “Experimental Evaluation of TCP Performance and Fairness in an 802.11e Test-bed”, *Proceedings of ACM SIGCOMM Workshop on Experimental Approaches to Wireless Network Design and Analysis*, Philadelphia, PA, pp. 17-22, August 2005.

[21] ns2, The Network Simulator ns-2, http://www.isi.edu/nsnam/ns/

[22] S. Pilosof, R. Ramjee, D. Raz, Y. Shavitt, and P. Sinha, “Understanding TCP Fairness over Wireless LAN”, *Proceedings of IEEE INFOCOM*, San Francisco, CA, pp. 863-872, March 2003.

[23] I. Rai, G. Urvoy-Keller, and E. Biersack, “Analysis of LAS Scheduling for Job Size Distributions with High Variance”, *Proceedings of ACM SIGMETRICS*, San Diego, CA, pp. 218-228, June 2003.

[24] I. Rai, G. Urvoy-Keller, M. Vernon, and E. Biersack, “Performance Models for LAS-based Scheduling Disciplines in a Packet Switched Network”, *Proceedings of ACM SIGMETRICS*, New York, NY, pp. 106-117, June 2004.

[25] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, “Opportunistic Media Access for Multirate Ad Hoc Networks”, *Proceedings of ACM/IEEE MOBICOM*, Atlanta, GA, pp. 24-35, September 2002.

[26] C. Williamson and Q. Wu, “A Case for Context-Aware TCP/IP”, *ACM Performance Evaluation Review*, Vol. 29, No. 4, pp. 11-23, March 2002.

[27] Q. Wu, and C. Williamson, “A DiffServ Framework for Context-Aware TCP/IP”, submitted for publication, 2006.