

TCP FAIRNESS ISSUES IN IEEE 802.11 WIRELESS LANS

Qian Wu Mingwei Gong Carey Williamson
Department of Computer Science, University of Calgary
Calgary, AB, Canada T2N 1N4
Email: {qianwu, gongm, carey}@cpsc.ucalgary.ca

ABSTRACT

In this paper, we study several TCP fairness issues in wireless local networks, both experimentally and via simulation. We illustrate the unfairness problems caused by the 802.11 MAC protocol, as well as by cross-layer interactions. Two possible solutions, namely SPM-AF queue management and LAS queue management, are proposed in the paper. These two queue management approaches are designed to alleviate the unfairness among competing TCP flows. The evaluation of these approaches uses the *ns-2* network simulator. The simulation results show that the proposed approaches provide better sharing of the limited wireless network bandwidth among TCP flows.

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

1 Introduction

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

WLANs can operate in either *infrastructure* mode or *ad hoc* mode. Our work focuses solely 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. In mixed wired/wireless networks, the wireless LAN is typically the throughput bottleneck [16]. Hence, the limited wireless bandwidth needs to be allocated fairly and efficiently amongst the mobile stations and the AP.

In this paper, we focus on the fairness of bandwidth allocation in wireless LANs. The IEEE 802.11 standard defines a 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 Transmission Control Protocol (TCP), can lead to flow-level unfairness [23].

The contributions of this paper are two-fold. First, we demonstrate, both experimentally and via simulation, the presence of TCP unfairness in typical WLAN scenarios. Second, we propose two different solutions to the TCP fairness problems, and evaluate them using simulation.

In particular, we evaluate two separate queue-based strategies to address the foregoing unfairness. The queue management schemes, called Selective Packet Marking with ACK Filtering (SPM-AF) and Least Attained Service (LAS) scheduling, are designed to solve the unfairness problems caused by interactions between TCP and the WLAN Medium Access Control (MAC) protocol.

The remainder of this paper is organized as follows. Section 2 reviews prior related work. Section 3 discusses fairness issues in 802.11 WLANs. The two subsequent sections, Section 4 and Section 5, present experimental and simulation results, respectively, illustrating the presence of these unfairness problems. Section 6 describes our two proposed solutions, while Section 7 presents the evaluation methodology and the simulation results. Section 8 concludes the paper.

2 Related Work

Fairness issues in wireless LANs have been studied extensively in the literature [3, 9, 17, 18, 19, 21].

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 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, 9, 19]. These solutions achieve fair allocation of the wireless channel by adjusting the contention window of each station dynamically.

Fairness issues in TCP and MAC protocol interaction are discussed in [2, 5, 13, 17, 18, 21, 23]. Pilosof *et al.* [23] observed unfair sharing of the network bandwidth between upstream flows and downstream flows through network measurements. A comprehensive simulation study was also 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 window for upstream flows to limit the upstream traffic. As a result, the downstream flows can obtain their fair share of the bandwidth.

The study on fair bandwidth allocation between short-lived flows and long-lived flows (i.e., mice and elephants) in wireless LANs is limited, though the research in wired networks is extensive [6, 11, 14, 20, 26, 30]. The unfairness

problem between short-lived flows and long-lived flows in wireless LANs can be more severe, since the packet loss rate and the bit error rate are usually higher due to the nature of the wireless channel.

Our proposed SPM-AF and LAS queue management approaches operate at the network and transport layers, in contrast to the solutions proposed in [17, 21], which focus on the MAC layer. The closest work to ours is that 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.

3 WLAN Fairness Issues

Unfairness problems in wireless LANs can be caused either by the 802.11 MAC protocol or by cross-layer interactions with other protocols. In this section, we discuss these issues, which have been well-documented in the literature.

3.1 Unfairness from MAC-Layer Protocol

A wireless channel is a shared scarce resource. When there are multiple mobile hosts present in a WLAN, they compete with each other for access to the wireless channel. To be fair, every mobile host should have equal opportunity to win the contention, and gain an equal share of the wireless channel. However, the two goals in the design of a MAC protocol, namely maximizing wireless channel utilization and fair allocation of wireless channel bandwidth, are not always compatible with each other.

There is always a trade-off between these two goals. For example, consider two mobile hosts with different wireless network conditions. To achieve higher network utilization, more bandwidth should be allocated to the station with better channel conditions, which leads to unfairness for the mobile host with worse channel conditions. On the other hand, assigning equal bandwidth to all the stations independently of channel conditions may degrade the overall network utilization.

There are many ways to define fairness on a wireless channel. Per-node fairness is defined as an equal channel share for every node. Each node has a per-node queue and a per-node backoff procedure. On the other hand, per-flow fairness is defined as equal channel share for every data flow. Each flow has a queue and a separate independent backoff procedure. A node with multiple flows can be viewed as multiple nodes each with one flow.

The 802.11 WLAN MAC is per-node based. With asymmetric contention and collisions in a wireless network, contention is per-flow rather than per-node. In the scenario of two flows, with one flow to a destination with better channel conditions than another flow, the flow with worse wireless channel conditions may impede the flow

with better conditions, degrading the overall performance of the wireless network.

Consider another WLAN scenario with two downstream flows (from AP to mobile host) and one upstream flow (from mobile host to AP). Suppose that the AP and mobile host share the wireless channel equally. Half of the wireless channel bandwidth is allocated to the two downstream flows, while the remaining half of the wireless channel bandwidth is allocated to the upstream flow. The upstream flow receives twice as much bandwidth as each downstream flow. The per-node fairness scheme in this AP scenario results in per-flow unfairness.

Another notion of fairness is time-based fairness. Fairness over short time scales is essential for delay-sensitive applications like audio or video. Long-term fairness is necessary for system stability. Short-term fairness leads to long term fairness, but long-term fairness does not imply short-term fairness. Some research shows that the Binary Exponential Backoff in IEEE 802.11 can exhibit both short-term and long-term unfairness.

Short-term unfairness may degrade the performance of higher-layer protocols, such as TCP. In the presence of short-term unfairness on WLANs, TCP acknowledgments and data packets may be clustered in time, resulting in bursty transmissions and ACK compression [31]. Bursty traffic can cause router queue overflows, leading to packet losses and performance degradation.

3.2 Cross-layer Interactions

TCP is widely used on the Internet for reliable end-to-end data transfer. The two main objectives of TCP congestion control mechanisms are efficient and fair allocation of the available network bandwidth among TCP flows [29]. However, the interactions between the 802.11 MAC protocol and TCP can result in unfair sharing of the wireless channel among competing TCP flows. In particular, we discuss competition between upstream flows and downstream flows, as well as unfairness between short-lived flows and long-lived flows in the following two subsections.

3.2.1 Upstream versus Downstream Flows

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 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 mobile hosts, generating a total of N upstream TCP flows and N downstream TCP flows. The upstream flows send large TCP data packets to the Internet via the AP, and receive small TCP ACK packets from the Internet via the AP, while the downstream flows receive large TCP data packets from

the Internet via the AP, and send small TCP ACK packets to the Internet via the AP.

Since the AP and the mobile stations share the wireless channel, each of them (on average) obtains a fraction $\frac{1}{N+1}$ of the opportunities for packet transmission. However, the AP has to transmit *half* of the total packets onto the WLAN, and is thus a bottleneck. Furthermore, the $\frac{1}{N+1}$ share that the AP obtains for channel access is 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 [7]. 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 bottleneck queue overflows, packet losses occur.

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 data packet 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.

3.2.2 Short-lived versus Long-lived Flows

TCP’s flow and congestion control algorithms are arguably designed to optimize throughput for long-lived bulk data transfers. The TCP *slow-start* phase is used initially to probe the network capacity, and steady-state is typically reached in the *congestion avoidance* phase. For small files, however, the steady-state congestion avoidance phase is not always reached before the transfer terminates. As a result, short-lived TCP flows are more conservative than long-lived TCP flows in their use of network bandwidth.

It is understandable that in an infrastructure-based WLAN, when long-lived and short-lived TCP flows compete for the downstream bottleneck queue at the access point, packets from the long-lived flows can occupy more buffer space than the short-lived flows, causing higher loss for the short-lived flows. Moreover, unlike wired networks, transmissions on wireless links are prone to bursty bit errors due to the nature of the wireless medium.

For small file transfers, packet losses can be relatively costly, since the congestion window size for these TCP connections may be too small to trigger the efficient “triple duplicate ACK” strategy. This can result in poor performance for small file transfers, in terms of variability of transfer times.

4 Experimental Results

We begin our investigation of TCP unfairness issues in WLANs with empirical experiments using off-the-shelf wireless networking equipment. The purpose of the experiments is to ascertain the existence of the TCP unfairness problems discussed earlier, and assess their magnitude.

The experiments are carried out in the Wireless Internet Performance Laboratory at the University of Calgary. Figure 1 illustrates the experimental setup. An IBM ThinkPad laptop with Linksys WPC11 version 3 wireless network PC card is configured as an AP. The IBM ThinkPad has an Intel Pentium 4-M 2.2 GHz processor with 256 MB RAM and 40 GB portable IDE hard drive. Installing the HostAP freeware on the IBM ThinkPad makes the laptop function as an AP. With bridging and IP-forwarding enabled on the AP laptop, all the frames received by the AP wireless network interface card are forwarded to the wired network through the AP Ethernet network interface card, and vice versa. Mobile hosts can communicate with the Internet through the AP via the wireless channel.

The WLAN tests transfer data to and from a wired server on the University of Calgary network. A similar configuration was used for WAN experiments, with the wired server located remotely at the University of Saskatchewan. In both tests, the maximum achievable throughput for a single TCP flow is approximately 5.1 Mbps.

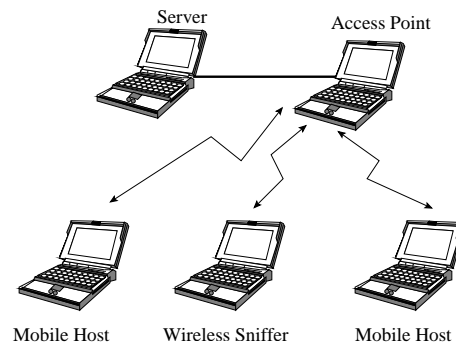


Figure 1. Experimental Setup: LAN Environment

TCP traffic is generated using netperf. The laptops were running Linux kernel version 2.4. Special instrumentation was added to record the transmission queue dynamics both in the mobile hosts and in the AP. The maximum queue size was 100 packets. Network traffic traces were collected using the SnifferPro 4.6 wireless network analyzer.

4.1 Head-of-Line (HOL) Blocking

Our first set of experiments identified a Head-of-Line (HOL) blocking problem that pervades wireless networks. This experiment featured two mobile hosts, with one laptop

located next to the AP, and the other laptop farther away. Each mobile host initiated a bulk transfer downstream TCP flow. The mobile host near the AP achieved a throughput of 3.64 Mbps, while the throughput of the mobile host 10 meters from the AP was 1.44 Mbps.

Figures 2(a) and (b) show the instantaneous throughput observed for the two flows. Flow 1, which started 8 seconds earlier, always achieves higher throughput than Flow 2. However, there is a strong synchronization between the flows, as indicated by the vertical time markers on the two graphs. This synchronization behaviour is caused by MAC-layer retransmissions of TCP packets destined to the (far) mobile host with poorer wireless channel conditions. When the AP spends time retransmitting to Flow 2, the instantaneous throughput of Flow 1 suffers. This phenomenon is further illustrated in Figures 2(c) and (d), which show the RTT values for each flow. There are distinct spikes in the RTT of Flow 1 as a consequence of repeated retransmissions to Flow 2.

4.2 Upstream versus DownStream Flows


The second set of experiments studies the influence of TCP flow directionality, as well as the number of upstream and downstream flows. These experiments were conducted in both LAN and WAN scenarios; the results were similar, so only the LAN scenario is discussed here.

For a single upstream flow and a single downstream flow, we observed flow unfairness in favour of the *downstream* flow. This result actually contradicts the argument presented in Section 3, so we investigated it further.


In this scenario, there are two mobile hosts. One is the sender of the upstream flow, and the other is the receiver of the downstream flow. The throughput achieved by the upstream flow is 2.41 Mbps, and the throughput for the downstream flow is 2.86 Mbps. The downstream flow achieves higher throughput than the upstream flow, as shown in Figure 3(a) and (b).

The explanation for this behavior is the transmission queue delays experienced by the flows. Both flows share a common bottleneck at the AP transmission queue: the downstream flow for its data packets, and the upstream flow for its ACKs. However, the upstream flow also experiences delays at another queue, namely the outbound transmission queue at the mobile host itself. The presence of this large queue of full-sized TCP packets adds significant delay to the RTT for the upstream flow (see Figures 3(c) and (d)), hindering its TCP throughput.


Figure 4 shows the results for *multiple* upstream and downstream flows. When there are two (or more) upstream flows, as well as two (or more) downstream flows, unfairness is very pronounced in favour of the *upstream* flows. For example, Figure 4 shows an average of 1.5 Mbps for the upstream flows, and an average of 1.0 Mbps for the downstream flows. As mentioned in Section 3, this phenomenon happens because the AP transmission queue saturates, losing packets for all flows. However, the packet losses have




(a) Downstream Flow 1 Throughput



(b) Downstream Flow 2 Throughput




(c) Downstream Flow 1 RTT




(d) Downstream Flow 2 RTT


Figure 2. WLAN Experimental Scenario with 2 Downstream Flows




(a) Upstream Flow Throughput



(b) Downstream Flow Throughput



(c) Upstream Flow RTT



(d) Downstream Flow RTT

Figure 3. WLAN Experimental Scenario with 1 Upstream Flow and 1 Downstream Flow

greater impact on the downstream flows (losing data packets) than the upstream flows (losing TCP ACKs).

4.3 Summary

The foregoing experiments demonstrate the presence of unfairness for TCP flows in an infrastructure-based IEEE 802.11b WLAN. These unfairness problems arise from unexpected interactions between TCP and the MAC-layer protocols.

5 Simulation Results

To further explore WLAN unfairness problems, we present experiments using the *ns-2* network simulator. The simulation approach allows us to explore larger scenarios, and to generalize our experimental results to other network and traffic flow configurations.

5.1 Upstream versus Downstream Flows

To illustrate the unfairness problem between upstream and downstream flows, we present the simulation model shown in Figure 5. 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 the wired station *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.


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 application layers for both downstream and upstream transfers have infinite data to send (i.e., bulk FTP 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 to determine steady-state throughput.


Figure 6 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.


The main 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 diminish. 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




(a) Upstream Flow 1 Throughput



(b) Upstream Flow 2 Throughput



(c) Downstream Flow 1 Throughput



(d) Downstream Flow 2 Throughput

Figure 4. WLAN Experimental Scenario with 2 Upstream Flows and 2 Downstream Flows




Figure 5. Network Model for ns-2 Simulations

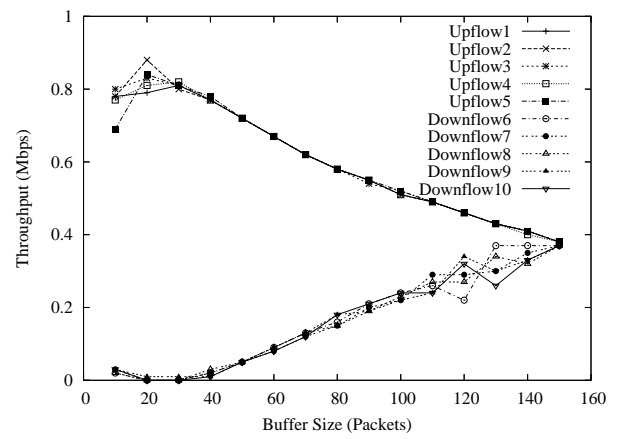


Figure 6. Throughput Results for Upstream and Downstream Flows

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.

5.2 Short-lived versus Long-lived Flows

We investigated the impact of long-lived TCP flows on short-lived flows by simulation. In the study, the number of mobile stations is varied from 2 to 40. Half of the mobile stations transfer large files with a fixed file size of 1000 packets. Each of these hosts downloads 100 such large files. The other half of the mobile stations transfer small files with a fixed file size of 10 packets. The small file transfers are repeated over and over until all large file transfers are finished. The TCP NewReno model and DropTail queue management are used. The data transmission rate used between the AP and the mobile stations is fixed at 11 Mbps. The packet loss rates and the Coefficient of Variation (COV) of the transfer times across the multiple transfers are calculated. If the long-lived flows and the short-lived flows are treated equally by the system, they should have similar packet loss rates and comparable COV values, though their mean transfer times may differ greatly.

Figure 7 shows the COV of the transfer times for the small files (10 packets) and the large files (1000 packets), respectively. The vertical axis shows the COV value, and the horizontal axis shows the number of mobile hosts. There are two lines on each graph. The line labeled "mixed traffic" represents a mix of small and large file transfers in the scenario described above. The line labeled "small only" or "large only" represents the scenario in which all mobile stations transfer either small files or large files. This scenario is used as the baseline for comparison.

Figure 7 (a) clearly indicates that the presence of long-lived TCP flows affects the performance of small file transfers. Compared to the baseline case, the COV values increase significantly. The performance degradation is more pronounced when there are more hosts in the wireless LAN. On the other hand, Figure 7 (b) shows that adding short-lived flows to the network has little impact on the large file transfers. Note that the COV values of the small file transfers 2 to 10 times larger than those of the large file transfers.

Figure 8 shows the packet loss ratio results. The statistics include the packet loss rate for the short-lived flows, the packet loss rate for the long-lived flows, and the overall average packet loss rate. The short-lived flows have higher packet loss ratio than the long-lived flows in all cases studied.

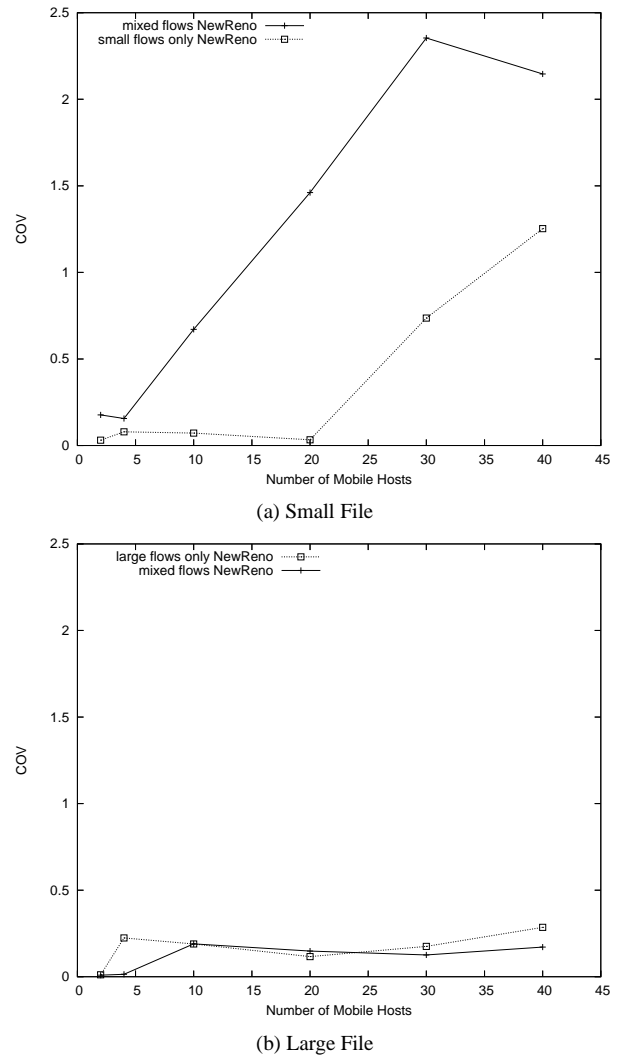


Figure 7. COV Results for Short-lived and Long-lived Flows




Figure 8. Packet Loss Ratio Results

5.3 Summary

The foregoing simulation experiments further demonstrate the presence of unfairness for TCP flows in wireless local networks. Solutions to these problems are explored in the next section.

6 Proposed Solutions

As discussed in Section 3.2, the interaction between the 802.11 MAC protocol and TCP's congestion control mechanism can result in unfair sharing of the wireless channel among competing TCP flows. These observations motivate 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.

6.1 SPM-AF Queue Management Scheme

The SPM-AF algorithm borrows ideas from Selective Packet Marking (SPM) proposed in [27] 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 [27], some packets in a TCP connection (i.e., any packets sent when the congestion window is smaller than 4 packets) are more crucial than others, since losses of these packets lead to a coarse-grained 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, one or more buffers are freed for use by data packets of downstream flows.

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), 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.

One side effect of ACK filtering is that the reduced frequency of ACK feedback slows 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.

6.2 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 [4].

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 network routers [24, 25]. 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 [24, 25].

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.

For LAS scheduling, one possible drawback is that one newly arriving long-lived flow can block other existing long-lived flows. That is, the packets belonging to the new flow are more likely inserted in front of the packets belonging to the other long-lived flows until all flows have received the same amount of service. This nature of LAS scheduling causes severe unfairness among long-lived flows.

Threshold-based LAS scheduling is proposed to solve this unfairness. The idea is to give the newly arriving flow service priority up to a certain threshold (e.g., 50 pack-

ets). Once the threshold is reached, FIFO scheduling is employed on this flow. The Threshold-based LAS, or T-LAS for short, can provide both fairness to short-lived flows as well as long-lived flows.

7 Simulation Evaluation

This section evaluates our queue management schemes. As introduced in Section 6, SPM-AF and LAS queue management strategies were designed to tackle unfairness among competing TCP flows.

7.1 Upstream versus Downstream Flows

The network model used in the simulation evaluation is the same as that described in Section 5. 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 fairness index [15]. Higher values of both metrics are desirable. We expect our proposed solutions to show similar aggregate TCP throughput, but improved fairness.

Figure 9 shows the simulation results. 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 9 (a) shows the aggregate throughput of 10 TCP flows (upstream and downstream). Figure 9 (b) shows the fairness index results.

Four important observations are evident in Figure 9. First, as expected, the proposed approaches achieve similar aggregate throughput. Figure 9 (a) shows that the total throughput for all tested strategies is roughly the same, though the SPM-AF scheme achieves 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 32% for the SPM-AF scheme and 44% for LAS scheduling. Fourth, for the two proposed queue management approaches, LAS scheduling provides the best fairness.

The fairness index generally improves for all strategies when the buffer size is increased, because there are fewer packet losses. Figure 10 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 10 also shows that the SPM-AF scheme has higher loss than the other two schemes, since SPM-AF can proactively remove ACKs even before the buffer overflows.

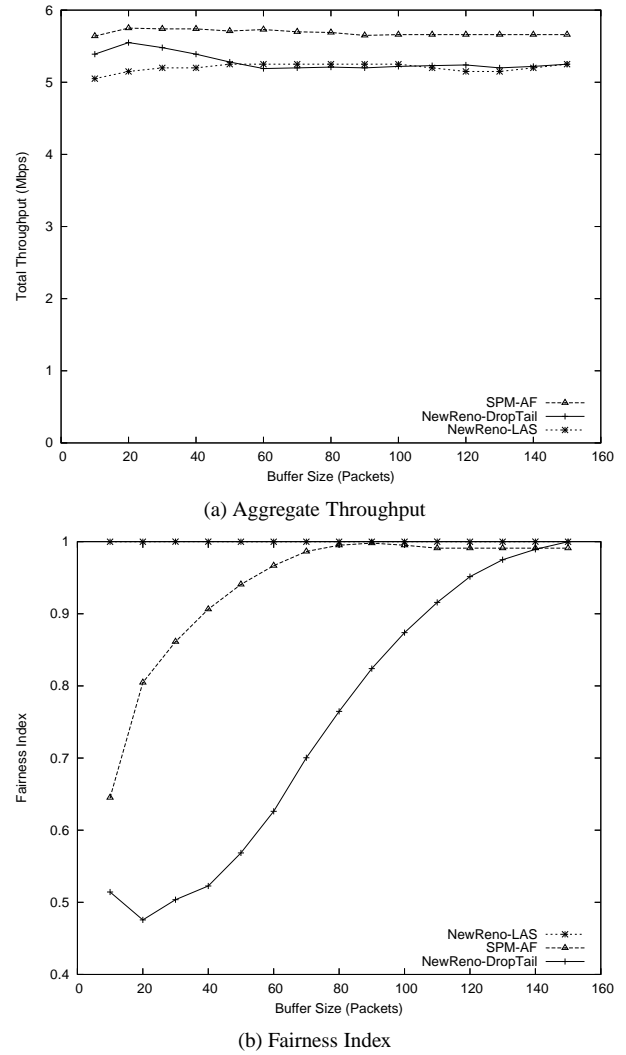


Figure 9. Simulation Results for Throughput and Fairness




Figure 10. Packet Loss Ratio Results

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-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.

7.2 Short-lived versus Long-lived Flows

The SPM-AF and improved T-LAS solutions are compared to the conventional TCP NewReno model and DropTail queue management using the same simulation model as described previously. We consider the packet loss ratio and the COV of the transfer times as the metrics of interest. We expect our proposed solutions to reduce COV values for the small file transfers, and to show comparable COV values for small and large file transfers.

Figure 11 shows the simulation results for the small files (10 packets) and the large files (1000 packets), respectively. On each graph, the number of mobile hosts is represented on the horizontal axis, with the COV of transfer times shown on the vertical axis.

There are four lines shown on each graph in Figure 11. Three of them represent SPM-AF, NewReno-T-LAS, and NewReno-DropTail mechanisms, respectively. The fourth line, labeled “10 packets only” in Figure 11 (a) and “1000 packets only” in Figure 11 (b), is for the scenario in which all mobile hosts transfer files of the same size (10 packets or 1000 packets) using the TCP NewReno model with DropTail queue management.

Figure 11 (a) shows that the proposed SPM-AF scheme reduces the COV for small file transfers. Compared to the NewReno-DropTail scenario, there is a dramatic reduction in the variability of transfer times. For example, the COV values for SPM-AF are 52.3% (10 mobile hosts) to 84.8% (40 mobile hosts) lower than for the NewReno-DropTail combination. SPM-AF also improves the performance of large file transfers as shown in Figure 11 (b), though the performance advantage is relatively small. More importantly, as expected, SPM-AF makes the COV values comparable for small and large transfers.

The low variability of transfer times for small files is achieved because the crucial packets are marked with high priority, and thus are less likely to be dropped on the bottleneck link. The “outlier” transfer times are reduced or eliminated with the loss priority mechanism employed by SPM-AF.

As for the T-LAS scheme, it inserts the packets from the small files close to the head of the queue. Thus, all small file transfers experience comparable transfer time, achieving smaller COV values compared with the conventional NewReno-DropTail scheme. For large file transfers,

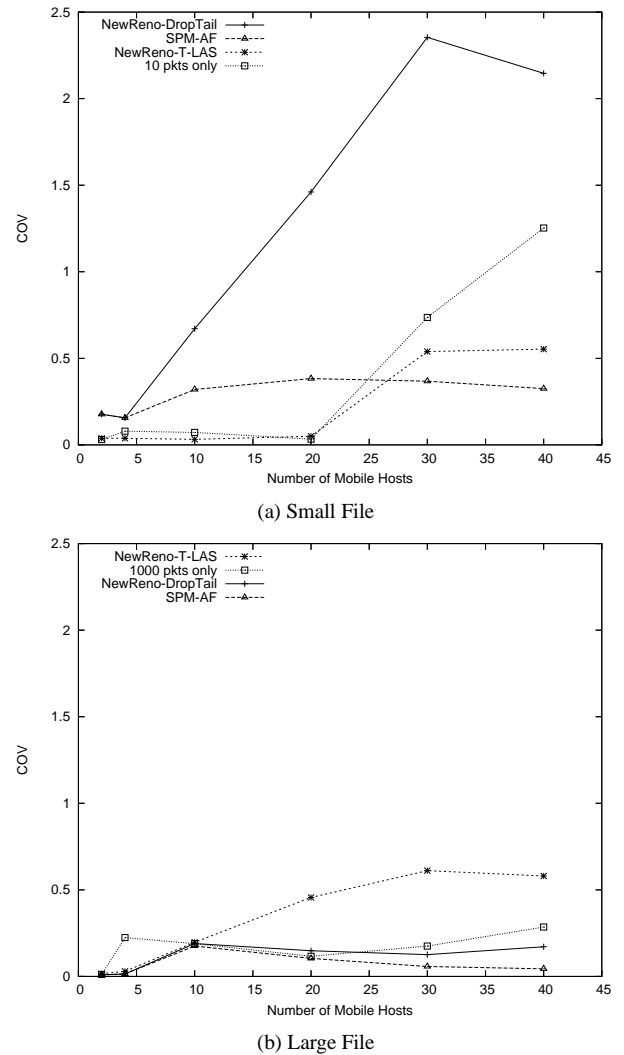




Figure 11. Simulation results for COV of Transfer Times



(a) SPM-AF



(b) T-LAS

Figure 12. Packet Loss Ratio Results for SPM-AF and T-LAS




Figure 13. COV Results for Different File Sizes

the T-LAS scheme assigns high service priority to packets up to a certain threshold. As shown in Figure 11, T-LAS achieves fairness for both small file transfers and large file transfers.

Figure 12 shows the packet loss ratio results for SPM-AF and T-LAS scheduling. There are three lines on each graph, representing the packet loss ratio for short-lived flows, the packet loss ratio for long-lived flows, and the overall average packet loss ratio.

Compared to Figure 8, Figure 12 clearly indicates that SPM-AF and T-LAS queue management dramatically transform the relative packet loss rates for the short-lived flows and long-lived flows. Moreover, Figure 12 (b) shows that when using T-LAS scheduling, the packet loss rate for small file transfers is near zero for most cases studied.

The SPM-AF scheme and T-LAS scheduling are also evaluated in a more complicated scenario. There are 20 mobile hosts in this experiment. Ten of them start FTP bulk data transfers in the background at the beginning of the simulation, and stop the transfers 500 seconds later. During the 500 seconds, the other ten hosts repeatedly transfer small Web documents using HTTP. The Web document (file) size is randomly generated between 5 packets and 25 packets (inclusive).

Figure 13 shows the COV of transfer times for the small file transfers. The horizontal axis shows the file size, while the vertical axis shows the COV of transfer times. Again, there are four different lines in the figure. Three of them are for SPM-AF, NewReno-T-LAS, and NewReno-DropTail schemes respectively. The line labeled “HTTP only” is for the case with no background FTP transfers; all hosts merely transfer small Web documents for 500 seconds.

Figure 13 shows the performance advantages. Compared to NewReno-DropTail, the SPM-AF scheme reduces the average COV value by a factor of 4, while T-LAS scheduling reduces the COV by about a factor of 5.

8 Summary and Conclusions

In this paper, we studied several TCP fairness issues in wireless local networks. We illustrated TCP unfairness problems caused by the 802.11 MAC protocol, as well as by cross-layer interactions.

Two solutions, namely SPM-AF queue management and LAS queue management, are evaluated in the paper. These queue management approaches were designed to alleviate the unfairness among competing TCP flows.

Our evaluation uses the *ns-2* network simulator. The simulation model assumes a typical infrastructure WLAN configuration, with an AP connected to a wired network and multiple mobile hosts in the WLAN.

The simulation results for SPM-AF and LAS queue management show that they can fairly and efficiently allocate the wireless channel among the competing TCP flows. In addressing the unfairness between upstream and downstream flows, the proposed solutions improve the fairness index by 20-40% compared to the conventional DropTail queue mechanism, while achieving comparable throughput. In tackling the unfairness between short-lived and long-lived flows, SPM-AF and LAS reduce the variability of transfer times for small files by 50% or more, while maintaining the throughput of large file transfers.

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 are grateful to Shelly Xiao for her efforts with the WLAN measurements.

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] E. Biersack, B. Schroeder, and G. Urvoy-Keller, "Scheduling in Practice", *ACM Performance Evaluation Review*, Vol. 34, No. 4, pp. 21-28, March 2007.
- [5] 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.
- [6] C. Brandauer, G. Iannaccone, C. Diot, and T. Ziegler, "Comparison of Tail Drop and Active Queue Management Performance for Bulk-data and Web-like Internet Traffic", *Proceedings of ISCC*, Hammamet, Tunisia, July 2001.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] J. Dunn, M. Neufeld, A. Sheth, D. Grunwald, and J. Bennett, "A Practical Cross-Layer Mechanism for Fairness in 802.11 Networks", *Mobile Networks and Applications*, Vol. 11, No. 1, pp. 37-45, February 2006.
- [11] S. Ebrahimi-Taghizadeh, A. Helmy, and S. Gupta, "TCP vs. TCP: a Systematic Study of Adverse Impact of Short-lived TCP Flows on Long-lived TCP Flows" *Proceedings of INFOCOM*, Vol. 2, pp. 926-937, Miami, FL, March 2005.
- [12] M. Gong, Q. Wu, and C. Williamson, "Queue Management Strategies to Improve TCP Fairness in IEEE 802.11 Wireless LANs", *Proceedings of RAWNET*, Boston, MA, April 2006.
- [13] 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.
- [14] L. Guo, and I. Matta, "The War Between Mice and Elephants", Technical Report BU-CS-2001-005, Computer Science Department, Boston University, May 2001.
- [15] R. Jain, *The Art of Computer Systems Performance Analysis*, Wiley & Sons, New York, NY, 1991.
- [16] 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.
- [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] I. Matta, and L. Guo, “Differentiated Predictive Fair Service for TCP Flows” Technical Report BU-CS-2000-012, Computer Science Department, Boston University, May 2000.
- [21] 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.
- [22] ns2, The Network Simulator ns-2, <http://www.isi.edu/nsnam/ns/>
- [23] 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.
- [24] 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.
- [25] 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.
- [26] K. Tokuda, G. Hasegawa, and M. Murata, “Analysis and Improvement of the Fairness Between Long-lived and Short-lived TCP Connections”.
- [27] 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.
- [28] Q. Wu, and C. Williamson, “A DiffServ Framework for Context-Aware TCP/IP”, submitted for publication, 2006.
- [29] F. Xiao, “TCP Fairness Issues on Wireless Networks”, M.Sc. Thesis, Department of Computer Science, University of Calgary, August 2004.
- [30] S. Yilmaz, and I. Matta, “On Class-based Isolation of UDP, Short-lived and Long-lived TCP Flows”, in *Proceedings of MASCOTS*, Cincinnati, OH, pp. 415-422, August 2001.
- [31] L. Zhang, S. Shenker, and D. Clark, “Observations on the Dynamics of a Congestion Control Algorithm: The Effects of Two-way Traffic”, *Proceedings of ACM SIGCOMM*, Zurich, Switzerland, pp. 133-147, September 1991.