Impact of 802.11e EDCA on Mixed TCP-based Applications

Marina Thottan
Center for Networking Research
Bell Laboratories
Murray Hill, NJ 07974
marinat@research.bell-labs.com

Michele C. Weigle
Department of Computer Science
Old Dominion University
Norfolk, VA 23529
mweigle@cs.odu.edu

ABSTRACT

There has been an explosive growth in the use of wireless LANs (WLANs) to support network applications ranging from web- browsing and file-sharing to voice calls. It is difficult to optimally configure WLAN components, such as access points (APs), to meet the quality-of-service requirements of the different applications, as well as ensuring flow-level fairness. Recent work has shown that the widely-deployed IEEE 802.11 MAC Distributed Coordination Function (DCF) is biased against downstream flows. The new IEEE 802.11e standard introduces QoS mechanisms, such as Enhanced Distributed Channel Access (EDCA), that allow this unfairness to be addressed. So far, only limited work has been done to evaluate the impact of these MAC protocols on TCP-based applications. In this paper, through *ns*-2 simulations, we evaluate the impact of EDCA on TCP application traffic consisting of both long and short-lived TCP flows. We find that the performance of TCP applications is very dependent upon the settings of the EDCA parameters and buffer lengths at the AP. We also show that the performance of the admission control strategy employed depends on the buffer lengths at the AP and the traffic intensity.

Categories and Subject Descriptors

C.2.1 [**Computer - Communication Networks**]: Network Architecture and Design—*Wireless communication*; C.2.6 [**Computer - Communication Networks**]: Internetworking—*Standards*

General Terms

Performance, Experimentation

Keywords

IEEE 802.11e, EDCA, TCP, admission control

1. INTRODUCTION

Wireless Internet access is *almost* ubiquitous – present everywhere from universities to airports to private homes. Recent measurement studies [3, 14, 11, 12] have shown that users on a wireless LAN (WLAN) behave similarly to those on a traditional wired LAN, using applications such as web browsing, email, streaming media, and peer-to-peer file sharing. For most of these applications, users download more data than they upload, yet research has shown that there can be significant unfairness present for wireless users who download data [21]. A simple WLAN consists of an access point (AP) that acts as a bridge between the wireless terminals and the wired Internet. The AP is connected to the wired network via a backhaul link that typically has a higher capacity than the system capacity of the WLAN. The most widely-used MAC layer protocol for WLANs, Distributed Coordination Function (DCF), does not provide any mechanism to differentiate between an AP and the user terminals, or between different traffic types. In DCF, both the AP and the user terminals have equal access to the wireless medium. Since users typically download more data than they upload, the AP requires a larger share of the channel than the terminals. With equal access between the AP and terminals, downloads see poorer performance than uploads.

The IEEE 802.11e standard [2] attempts to address some of these quality of service (QoS) concerns by introducing new MAC layer enhancements. The goal of this work is to explore the features of the 802.11e MAC standard with the purpose of aiding the design of a fair, application-aware WLAN that will enable support for diverse Internet applications. The 802.11e MAC standard is comprised of two different MAC enhancements: enhanced distributed channel access (EDCA) and hybrid coordination function (HCF)-controlled channel access (HCCA). In this work we focus on EDCA. Through extensive simulations using realistic traffic traces we will evaluate the impact of the various MAC layer parameters on the performance of TCP applications. We will analyze the upstream / downstream fairness issue along with user-perceived performance for WLAN users.

EDCA provides QoS by controlling medium access and classifying traffic. As will be discussed in Section 2, there are a number of parameters that need to be specified for each of these mechanisms. In this work we will show the impact of various MAC settings on the behavior of TCP application traffic. The 802.11e standard also calls for mandatory admission control policies that are to be implemented by the AP. However, there are no detailed guidelines provided on how the traffic should be classified or how the different MAC settings should be optimized to meet the QoS requirements of the different applications. We will provide working guidelines on setting these parameters by jointly evaluating the impact of the MAC parameters, traffic classification, and admission control. Our larger goal is to make configuring the APs easier, thus paving the way to zero-config AP and leading to faster adoption of QoS-assured services such as voice over the WLAN.

The challenges involved in studying this problem are twofold: (1) the generation of realistic traffic traces that model a real WLAN

WiCon'06, August 2-5, 2006, Boston, MA, United States.

network, and (2) working with a newly-defined standard that has limited real-world implementations. Recent measurements [11] have shown that the WLAN traffic is increasingly comprised of Peer-to-Peer (P2P) flows in addition to the standard web browsing flows. To tackle the challenge of generating realistic traffic patterns, we use a combination of WLAN traffic models developed by Meng *et al.* [17] along with long-lived TCP flows that mimic P2P file transfers. Given the limited real-world implementations of the 802.11e EDCA MAC schemes, we use the *ns-2* [16] simulator with recently-added support for EDCA. In this work, since we are primarily concerned with parameter optimization, we do not consider mobility-specific issues such as hand-off or channel loss. We focus only on system-induced impairments such as queuing delays and system resource requirements such as buffer sizing and ease of implementation.

Section 2 describes DCF in 802.11, provides an overview of 802.11e EDCA, and discusses related work. We describe our approach in Section 3, including our system model, the traffic models used, and the EDCA parameters tested. In Section 4 we describe our experimental setup, and in Sections 5 and 6, we discuss our main results.

2. BACKGROUND

The IEEE 802.11 standard specification [1] defines the mechanisms used for wireless media access and physical-layer encodings. In this work, we focus on the MAC layer and the functions provided by 802.11. In the original 802.11 (now known as "802.11 legacy"), there are two main functions for media access: the distribution coordination function (DCF) and the point coordination function (PCF). IEEE 802.11e [2] is an extension to 802.11 that can provide quality-of-service (QoS) to WLANs. Like legacy 802.11, 802.11e also provides two types of media access: enhanced distributed channel access (EDCA) and hybrid coordination function (HCF) controlled channel access (HCCA). Today, all 802.11 wireless devices support DCF, but only a very few also support PCF. For 802.11e, most devices support some form of EDCA, but none yet support HCCA [4].

2.1 Legacy 802.11

Legacy 802.11 supports DCF and PCF. DCF provides contention-based access to the wireless medium, while PCF provides contention-free access, with the AP choosing which wireless station can transmit at any time.

In 802.11 DCF, wireless stations use a carrier sense multiple access/collision avoidance (CSMA/CA) mechanism. When a station wants to send data, it must first listen on the channel. If the channel is idle (*i.e.,* no other station is transmitting) for the duration of a DCF inter-frame spacing (DIFS), then the station can send its data. If the channel is busy, the station must defer until the transmission has ended and the channel has been idle for a DIFS. The station then sets a backoff timer, where the duration of the timer is a function dependent on the physical layer and the contention window parameter (CW). Once the backoff timer has reached zero, the station can transmit. After a station has completed transmission of a frame, it must wait at least as long as its contention window before attempting to transmit again. This backoff method results in the station that picks the shortest random backoff time gaining access to the medium.

2.2 802.11e

The 802.11e standard introduces two QoS-based media access protocols, which offer traffic prioritization and improved channel access over DCF. EDCA is similar to a priority-based DCF scheme, while HCCA is a polled protocol, like PCF, that provides parameterized service.

There are four basic access categories (ACs) used to assign priorities in EDCA. These roughly represent, in order from highest to lowest priority, voice, video, best effort, and background traffic. The priorities are implemented through the use of a separate transmit queue and different EDCA parameter values for each of the four access categories. The adjusted EDCA parameters are the contention window range ($CW_{min} \leq CW \leq CW_{max}$), the arbitration inter-frame spacing (AIFS), and the transmission opportunity (TXOP) limit. EDCA is much like DCF, so there is a backoff period when the channel is found to be busy. The contention window parameter is used in EDCA, as in DCF, to determine how long the backoff period should be. The AIFS is similar to the DIFS in DCF. A station must observe the channel idle for an AIFS before transmitting. The TXOP limit defines how long a station can continue to transmit once it has gained access to the channel. The highest priority access category has the lowest values of the CW and AIFS parameters. The smaller these values, the shorter the amount of time the station has to wait to transmit data. High-bandwidth, high-priority flows should have larger values of the TXOP limit so that they can send multiple frames back-to-back without having to contend for the medium again. The AP maintains its own set of EDCA parameters to govern its access to the medium, and it assigns (potentially different) EDCA parameters to the wireless stations when they associate with the AP. EDCA does not provide explicit QoS guarantees, rather it simply provides a very high statistical likelihood that higher levels of bandwidth are allocated to higher-priority traffic.

2.3 Related Work

Early work on fairness in 802.11 networks [18, 23] focused on MAC-layer, or node-level, fairness. With this notion of fairness, each wireless node should receive equal access to the wireless channel. Much of this work, including investigations of 802.11e fairness [10, 4], have focused on UDP traffic, such as streaming media and VoIP. In contrast to UDP-based applications, applications that use TCP, such as web or file transfer, are congestion-controlled and drastically reduce their sending rates in response to packet loss. Most investigations on how the design of the 802.11 MAC-layer affects TCP flows have focused solely on either upstream traffic [8, 9, 25] or downstream traffic [6, 19]. The focus of these efforts has been on TCP flow-level fairness, with the goal of each TCP flow within the group of upstream or downstream TCP connections obtaining similar performance.

Only recently has the fairness between upstream and downstream TCP connections been investigated. Pilsof *et al.* [21] considered TCP fairness when there are both uploads and downloads occurring on the same wireless channel. If the AP receives data from the wired network at a faster rate than it gains access to the wireless medium, this data is queued at the AP. With DCF, the AP has no higher priority to the wireless channel than any other station, potentially resulting in large queues building up at the AP. Pilosof *et al.* showed that with the asymmetry of channel access, buffer availability was a scarce resource because TCP acknowledgments (ACKs) were consuming precious buffer space in the AP queue. The authors demonstrated that unless buffer sizes at the AP were sufficiently large, upstream flows could obtain much more bandwidth than downstream flows because the downstream data packets were being dropped in large numbers at the AP.

There have been several proposals aimed at addressing the upstream / downstream asymmetry that causes queue overflow at the AP. These include giving the AP higher priority access to the channel [6], adjusting the TCP advertised window to slow down TCP senders so that they do not overflow the AP queue [21], having TCP receivers skip sending some number of ACKs [8], using per-node queues at the AP [5], and employing active queue management techniques in the wireless network [26, 27, 20].

However, there have been far fewer investigations of the effects of 802.11e QoS techniques on TCP flows. Leith and Clifford [15] investigated TCP upload performance with 802.11e EDCA and proposed a method to ensure fairness by grouping the TCP ACKs into one AC and adjusting the *AIFS* and *CW* parameters for that AC at the AP. Tinnirello *et al.* [22] studied the impact of EDCA parameter settings by changing the values at the mobile stations and keeping the default values for the AP. Casetti and Chiasserini [7] evaluated the performance of the 802.11e EDCA protocol for voice traffic in the presence of on-off TCP traffic.

There has been very little work investigating TCP fairness with a realistic traffic mix consisting of many short-lived flows and a few long-lived flows over WLANs. The work of Bottigliengo *et al.* [5] is one of the few. They investigate 802.11b and propose the addition of an LLC-layer algorithm at the AP and wireless stations to give priority access to stations that have experienced channel failures. Our research is different in that we wish to investigate TCP fairness of WLANs given a traffic model [17] derived from WLAN traffic measurements over both 802.11 and 802.11e MAC protocols without operating outside the parameters of the standard protocols (*i.e.*, maintaining the current four access category queues for EDCA).

There have been several recent measurement studies of WLAN traffic [3, 14, 11, 12, 13]. Kotz *et al.* performed two measurement studies of the campus WLAN at Dartmouth College. The first study [14], performed during Fall 2001, showed that users were generally stationary and over 90% of the traffic was TCP. In particular, HTTP accounted for 53% of all traffic. In a follow-up study [11], performed in 2003-2004, the authors found overall increased usage of the wireless network, with marked increases in the usage of peer-to-peer (P2P) file sharing and streaming media. Meng *et al.* [17] used data from the first Dartmouth study [14] to develop traffic models. The authors fit well-known random distributions (Weibull and Lognormal) to empirical distributions of bytes transferred per connection and connection start intervals for each access point in one hour time intervals. We have used this model of WLAN traffic in our studies.

3. APPROACH

In this section, we present our approach to investigating the impact of EDCA parameters on TCP applications. We describe our overall system model, the traffic models we use, the EDCA parameters tested, and our admission control policy. In this work, our main goal is to optimize the MAC parameters (*AIFS* and *CW* range) with the purpose of improving the end user perception of the quality of the application. Using realistic traffic traces we expose the dependence between the EDCA parameters under different mix of application traffic and traffic intensity.

3.1 System Model

Figure 1 shows our system model. In our model we consider a single AP that is connected to a single wired node by a 100 Mbps link with a 25 ms propagation delay. There are eight stationary wireless nodes that are equi-distant from the AP. We assume that the nodes are close enough to the AP to avoid channel errors. The wireless network is 802.11b with a capacity of 11 Mbps. Our model is aimed at studying queuing-related traffic impairments, so our model does not consider any other impairments such as channel loss or mobility. Furthermore, since our focus is on the service provider perspective, we only consider the configuration and tuning of the APs. We assume that the mobile nodes use the default setting of the MAC parameters.

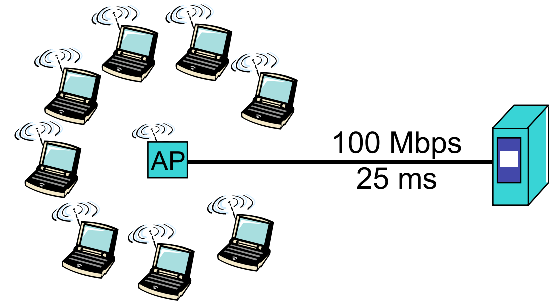

Figure 1: Topology Model. The AP is connected to the wired node. The eight stationary wireless nodes are equi-distant from the AP.

3.2 Traffic Models

Our experiments used three different traffic mixes: long-lived P2P-like (P2P), model-based (AP12), and a mix of model-based and P2P (AP12-P2P). For the P2P traffic, we ran 1 or 2 concurrent 4 MB TCP transfers (1P2P, 2P2P) in each direction (upstream and downstream). Once a file transfer completed, a new transfer was begun after waiting 10 seconds. The model-based traffic was derived from a model of the traffic on individual APs [17] based on traces from Dartmouth College in 2001 [14]. In particular, we used the traffic observed during the 11am-noon hour on AP 12, as it was one of the busiest times at one of the busiest APs. In order to achieve higher load (similar to that described in a later measurement study [11]), we used the traffic from AP 12 and AP 14 together (AP12-14). These TCP flows were generally short-lived, thus emulating web-like traffic, where the wireless node sends a request and the wired node sends back a response. This results in data flowing in both directions (upstream and downstream). The traffic model only specifies the total bytes sent, but does not specify how the bytes were divided between upstream and downstream. We take the total bytes transfered by the APs and split them equally between inbound and outbound traffic at the AP. This enables us to evaluate a more balanced loading condition between upstream and downstream. For the mix of model-based and P2P traffic, we used the traffic from AP 12 along with 1 or 2 P2P flows in each direction (AP12-1P2P, AP12-2P2P), and we also used a mix of AP 12 and AP 14 with 2 P2P flows (AP12-14-2P2P).

3.3 EDCA Parameters

In EDCA, the four standard ACs are assigned priorities based on an input vector consisting of the following parameters: $AIFSN$, CW_{min}, CW_{max}, and $TXOP\_Limit$. We use the default $TXOP\_Limit$ parameters for an 802.11b physical layer. The EDCA AIFS parameter is computed as $AIFS[AC] = AIFSN[AC] \times SlotTime + SIFSTime$, where $SIFSTime$ is the short Inter-Frame Spacing interval. All mobile stations use the default EDCA parameters outlined in the 802.11e standard and shown in Table 1. Note that CW_{min} and CW_{max} dictate the back-off period and are smaller for the higher priority queues (q_0 and q_1). Between the two high priority queues, the waiting period after the channel is idle (AIFS) is the same, but since the CW values are smaller for q_0, the traffic in q_0 will experience the least amount of queuing delay. The two low priority queues both have similar parameters for determining the back-off period,

AC	CW_{min}	CW_{max}	$AIFSN$	$TXOP\_Limit$ (ms)
q_0	7	15	2	3.264
q_1	15	31	2	6.016
q_2	31	1023	3	0
q_3	31	1023	7	0

Table 1: Default EDCA MAC Parameter Settings (EDCA-1)

AC	CW_{min}	CW_{max}	$AIFSN$	$TXOP\_Limit$ (ms)
q_0	7	15	**1**	3.264
q_1	15	31	**1**	6.016
q_2	31	1023	**2**	0
q_3	31	1023	**6**	0

Table 2: EDCA-2 Parameter Settings at the AP. $AIFSN$ **is set to the default** $AIFSN$ **- 1.**

but they differ significantly in the wait period after channel is idle. Also note that q_1 has the longest transmit time once the channel is gained. When $TXOP\_Limit$ is 0, the queue can transmit only one frame at a time.

We evaluate the performance of TCP traffic over five different settings of the main EDCA parameters. The first set of parameters, which we call EDCA-1, are the same as the default parameters.

The second set of parameters, EDCA-2, was designed to give a higher priority to the AP. Since a lower *AIFS* value gives priority when contention arises, we set *AIFSN* on each AP queue to be 1 less than the default. Since the mobile nodes keep the default value, the AP will have a higher priority. These values are shown in Table 2.

Table 3 shows the third set of parameters (EDCA-3). In this case we set the $AIFSN$ to 2 less than the default, giving a yet higher priority to the AP than with EDCA-2.

The IEEE 802.11e standard recommends that the minimum value of $AIFSN$ be set to 2, but with both EDCA-2 and EDCA-3, our formula causes the $AIFSN$ of some queues to fall below the recommended minimum. So, we also evaluated reducing the $AIFSN$ values by 2 (as in EDCA-3), but requiring the minimum to be 2. These parameters, called EDCA-4, are shown in Table 4.

Our final set of parameters modifies only the default CW_{max} parameter of q_2. With EDCA-5, we fix the $AIFSN$ values to the default settings and reduce CW_{max} of q_2 by half. The parameter specifications are shown in Table 5. The goal of this setting is to improve upstream / downstream fairness by giving priority to the AP.

3.4 Admission Control

The goal of our admission control policy is to reduce the response times for short web transfers, which might have real-time delay requirements. For our admission control policy we assume the following traffic classification and prioritization. A typical WLAN traffic composition can consist of delay-sensitive voice traffic, interactive web transfers with real-time constraints, large file transfers that are sensitive to loss, and best-effort traffic. To guarantee the best delay objectives we assume that the voice traffic will be mapped to the access category of q_0 with a small buffer size of 10 packets. The smaller buffer helps to bound the delay jitter in the voice traffic. For the real-time web transfers where the loss is not as critical as the download times, we use the second access category q_1 with a small buffer size of 10 packets. Access category q_2 typically consists of large file transfers, and category q_3 contains the best effort traffic. With these initial priority settings, we employ the following admission control policy:

AC	CW_{min}	CW_{max}	$AIFSN$	$TXOP\_Limit$ (ms)
q_0	7	15	**0**	3.264
q_1	15	31	**0**	6.016
q_2	31	1023	**1**	0
q_3	31	1023	**5**	0

Table 3: EDCA-3 Parameter Settings at the AP. $AIFSN$ **is set to the default** $AIFSN$ **- 2.**

AC	CW_{min}	CW_{max}	$AIFSN$	$TXOP\_Limit$ (ms)
q_0	7	15	**2**	3.264
q_1	15	31	**2**	6.016
q_2	31	1023	**2**	0
q_3	31	1023	**5**	0

Table 4: EDCA-4 Parameter Settings at the AP. $AIFSN$ **is set to the default** $AIFSN$ **- 2 with a minimum value of 2.**

AC	CW_{min}	CW_{max}	$AIFSN$	$TXOP\_Limit$ (ms)
q_0	7	15	2	3.264
q_1	15	31	2	6.016
q_2	31	**511**	3	0
q_3	31	1023	7	0

Table 5: EDCA-5 Parameter Settings at the AP. At q_2, CW_{max} **is set to half of its default value.**

- q_0: *Voice and other delay-sensitive traffic*
  When the queue overflows, packets are dropped.
- q_1: *Real-time web downloads and file transfers*
  When the queue reaches 75% capacity, new flows that would have been classified in q_1 are classified in q_2 instead.
- q_2: *Large file transfers*
  When the queue reaches 75% capacity, new flows that would have been classified in q_2 are classified in q_3 instead.
- q_3: *"Best-effort" traffic*
  When the queue overflows, the packets are dropped.

With our model-based and P2P traffic mix, described in Section 3.2, model-based flows smaller than 10 KB are assigned to $q1$, model-based flows larger than 10 KB are assigned to $q2$ and are subject to the admission control rules described above. P2P flows are assigned to $q2$ and are not subject to the admission control rules.

4. EXPERIMENTAL SETUP

We ran *ns*-2 simulations with both long-lived and model-based (mostly short-lived) dynamic TCP traffic. We evaluate the performance of this mixed traffic using two different settings of the MAC parameters. Upstream and downstream performance is evaluated as function of the variable queue sizes in the transmit queues at the AP.

The wireless topology used in our simulations was shown in Figure 1. The AP is connected to the wired node by a 100 Mbps link with a 25 ms propagation delay (RTT of 50 ms). The eight stationary wireless nodes are equi-distant from the AP. The wireless network is 802.11b with a capacity of 11 Mbps. The maximum TCP segment size is 1420 bytes, and the maximum TCP window is 47 1420-byte packets (approximately 64 KB). We used the Full-TCP model in *ns*-2 with TCP SACK and delayed ACKs. All experiments were run for 30 minutes (1800 seconds) after a 5-minute warmup interval.

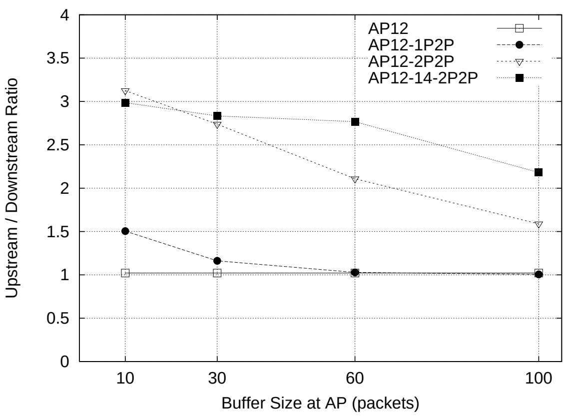

Figure 2: DCF: Ratio of upstream throughput to downstream throughput for various buffer sizes and levels of traffic intensity

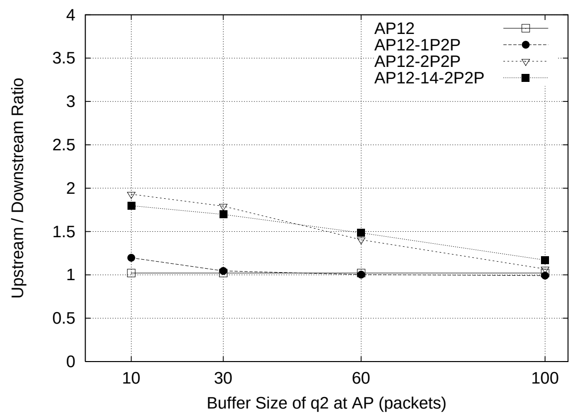

Figure 3: EDCA: Ratio of upstream throughput to downstream throughput for various buffer sizes and levels of traffic intensity

For EDCA, the four different transmit queue sizes at the AP were set up as follows: $q_0 = 10$, $q_1 = 10$, $q_2 = [10, 100]$ and $q_3 = 100$. The queue sizes were chosen keeping in mind the needs of different application traffic. Real-time traffic such as voice is required to meet stringent delay requirements. This requirement is met by q_0 which is the highest priority queue and has a small buffer to reduce inherent queuing delays. Similarly, short interactive web transfers are also sensitive to large queuing delays and hence are assigned to q_1 which also has a small buffer. Larger file transfers, such as P2P traffic and large web downloads, are handled by the lower priority queue q_2 which has a larger buffer. In our experiments we vary the q_2 buffer size from 10 packets to 100 packets ({10, 30, 60, 100}). q_3 has a fixed buffer size of 100 packets and is used to handle the overflow of queues q_1 and q_2 and to handle best-effort traffic. Each wireless terminal node has a transmit queue of 50 packets.

Our experimental work makes use of both the DCF and EDCA MAC models in *ns-2*. For DCF, we used the 802.11 model present in *ns*-2 with bug fixes provided by Wiethölter and Hoene [24], and for EDCA, we use the 802.11e model written and verified by Wiethölter and Hoene [24].

5. RESULTS

All results shown are of the average metric obtained from five independent replications of each experimental set up along with their 95% confidence interval (though in most cases, the confidence interval is so small that it is not visible on the plots).

5.1 DCF Unfairness

Figure 2 shows the upstream / downstream unfairness problem as a function of queue size when DCF is used. Recall that DCF has only one queue, so there is no QoS employed. The unfairness is due to the fact that the AP has the same priority to send as any wireless node. When there is a large amount of data flowing downstream, the transmit queue at the AP builds up and eventually overflows, resulting in loss and poor performance for the downstream flows. We see that, in general, as the queues grow larger unfairness is mitigated. Note that as additional P2P flows are added, larger buffers are required for fairness (larger than the 100 packets shown here). This finding is consistent with previous results [21]. However increasing the queue size to larger buffers will deplete the system memory resources. Furthermore, with larger buffers there is an inherent delay introduced between the packets of the same flow thus leading to large delay variances. Large increases in the delay variance can cause a significant impairment to the user-perceived performance on certain applications such as voice and video.

5.2 Impact of EDCA

We evaluate the impact that EDCA has on upstream / downstream fairness by using the default parameters (Table 1) both at the AP and mobile stations. The model-based traffic (AP12) is assigned to $q1$, and the P2P traffic is assigned to $q2$. As described in Section 4, $q1$ has a maximum size of 10 packets and $q2$ varies from 10 packets to 100 packets. The ratio of upstream throughput to downstream throughput is shown in Figure 3.

We see that there is no unfairness when only the model-based traffic (AP12) is used with either DCF or EDCA. When P2P flows are added, although some amount of upstream / downstream unfairness is still present, we find that the queuing requirements with EDCA are smaller than for DCF to achieve fairness. Note that the trivial solution of increasing buffer sizes to achieve fairness may prevent real-time flows from meeting their minimum delay constraints. Thus as described in Section 2, the different EDCA parameter settings need to be tuned and the queue sizes optimized for the different traffic classifications.

When EDCA is used, we can separate traffic types into different transmit queues, which helps to keep overall queue sizes low. We can also give priority to downstream flows, which allows the AP to access the wireless medium more often.

5.3 Impact of EDCA Parameter Settings and Buffer Size at $q2$

Here we compare the performance of the different EDCA parameter settings described in Section 3.3 according to user-level metrics. For the model-based flows, we will show the median response time, which is the time between a "request" being sent by a wireless client to the time that the "response" is received by the wireless client. This includes time for both upstream and downstream transfers. We consider the median statistic since the flow sizes vary and the average response time value will be skewed by the heavier flows. For the P2P flows, we will show the mean flow completion time (time to transfer the 4 MB file) for downstream flows. Since the AP12-2P2P and AP12-14-2P2P traffic intensities are the most interesting in terms of fairness, we will focus on these for the remainder of the discussion.

Figures 4 and 5 show the median response times of the model-

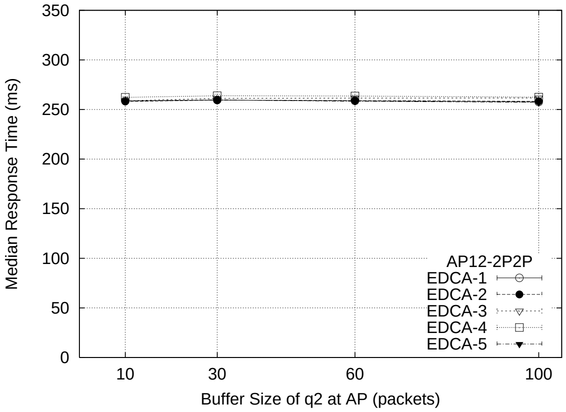

Figure 4: AP12-2P2P: Median response times of model-based flows

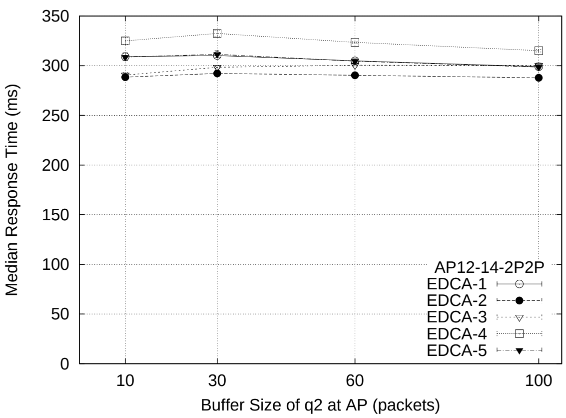

Figure 5: AP12-14-2P2P: Median response times of model-based flows

based flows for AP12-2P2P and AP12-14-2P2P, respectively, for various buffer sizes and EDCA parameter settings. At the lower intensity of model-based traffic (AP12-2P2P), we can see that neither the queue size at $q2$ nor the EDCA parameter settings affect the performance of the model-based flows. The buffer size at q_2 should have little effect on the response times of model-based traffic because all of the model-based traffic is classified in q_1 at a higher priority. At the higher intensity (AP12-14-2P2P), we see that using EDCA-2 results in the best performance over all $q2$ buffer sizes. The EDCA-2 setting decreases $AIFSN$ in each AC at the AP by 1. This gives downstream traffic in each AC slightly higher priority than the upstream traffic in the corresponding AC. We note that the poorest performer is EDCA-4, in which we reduced $AIFSN$ by 2 at the AP, but maintained the standard-recommended minimum of 2. Because of this, both $q1$ and $q2$ have the same $AIFSN$, which is not the case with any of the other settings. The only difference between them is the CW range. Since there is a noticeable effect on the model-based response times, we can say that $AIFSN$ affects response times for flows in $q1$ more than changing the CW range. This is reinforced by the performance of EDCA-5, in which we reduced the value of CW_{max} at $q2$ at the AP. EDCA-5 performs very similarly to EDCA-1, which has the default EDCA parameter settings at the AP. We do note that incorrect settings of the CW

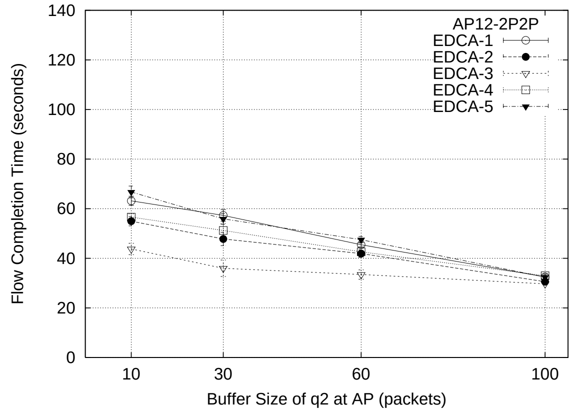

Figure 6: AP12-2P2P: Mean flow completion times of downstream P2P flows

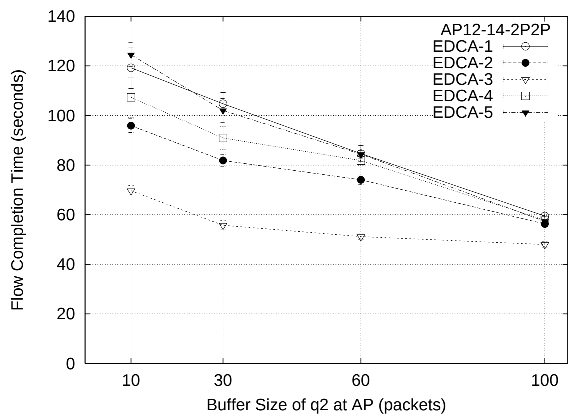

Figure 7: AP12-14-2P2P: Mean flow completion times of downstream P2P flows

parameters can have a detrimental effect. We initially ran a set of experiments where the CW range of $q2$ overlapped that of $q1$. This gave $q2$ priority over $q1$ in some instances, causing poor performance for the model-based flows.

Figures 6 and 7 show the mean flow completion times of the downstream P2P flows for AP12-2P2P and AP12-14-2P2P, respectively. For all EDCA parameter settings and $q2$ buffer sizes, the introduction of additional model-based flows (AP12-14-2P2P) has a detrimental impact on the P2P flows. Since the P2P flows are not high-priority, this is not a great concern, but we would like to have settings that minimize the degradation of the P2P performance. As the $q2$ buffer size increases towards 100 packets, the flow completion times decrease. Since the P2P flows are at a lower priority, they will have less opportunity to access the wireless channel and packets will have to be queued in $q2$ at the AP. With a larger queue buffer, the AP can handle larger bursts of arriving packets without overflowing. An overflow would result in packet loss, which greatly reduces TCP throughput. Again, we want to strike a balance between the benefit of increased throughput and the cost of increased memory requirements that larger buffer sizes bring. At $q2$ of 100 packets and EDCA-3, which is by far the best EDCA setting for the P2P flows, the increase in flow completion time with additional model-based traffic is less than 20 seconds. Using EDCA-3

Traffic	EDCA Setting	Upstream/ Downstream
AP12-2P2P	EDCA-2	1.00
AP12-2P2P	EDCA-3	1.00
AP12-14-2P2P	EDCA-2	1.14
AP12-14-2P2P	EDCA-3	0.98

Table 6: Upstream/downstream throughput ratio for AP12-2P2P and AP12-14-2P2P with EDCA-2 and EDCA-3

gives the downstream traffic more priority over the upstream traffic because *AIFS* at the AP is 2 less than the default, which is used at the mobile stations. This allows $q2$ at the AP to send packets out more often and causes the queue to be lower, reducing the time that packets spend waiting to be transmitted. All of this leads to lower flow completion times for the P2P flows.

Figures 4-7 show that $q2$ with a size of 100 packets gives the best performance for both model-based and P2P flows. Also, the figures show that the EDCA-2 parameter setting results in the best performance for model-based flows, and the EDCA-3 parameter setting results in the best performance for P2P flows. We will continue the discussion of the results focusing on EDCA-2 and EDCA-3 with a $q2$ setting of 100 packets.

6. DISCUSSION

In this section, we will explore in greater detail the tradeoffs and impact of the EDCA-2 and EDCA-3 parameter settings with $q2$ = 100 packets. We will also discuss the impact of the queue size at $q1$ and admission control.

6.1 Upstream/Downstream Fairness

Table 6 shows the upstream / downstream throughput ratio for AP12-2P2P and AP12-14-2P2P using EDCA-2 and EDCA-3. Both EDCA settings are fair for AP12-2P2P, but at AP12-14-2P2P with EDCA-2, the upstream throughput is slightly higher than the downstream throughput. Recall that the EDCA-2 setting modified the *AIFSN* parameter at the AP by 1 and the EDCA-3 setting modified the *AIFSN* parameter at the AP by 2. So, EDCA-3 gives slightly higher priority to the AP, allowing the downstream flows to achieve fairness even with additional model-based traffic.

6.2 Packet Loss

Figure 8 shows packet loss for model-based and P2P flows in both upstream and downstream directions. We would expect the P2P loss rate to increase with an increase in model-based traffic because the queue builds up as $q1$ gets priority access to the channel. But, the increase in loss for the model-based flows shows that the $q1$ buffer size of 10 packets is insufficient for the increased traffic with AP12-14.

With the increase of the size of $q1$ to 20 packets, model-based packet loss is greatly reduced even with increasing model-based traffic, as shown in Figure 9. We note that this result will be dependent upon the traffic intensity of these types of flows.

6.3 Impact of Admission Control

The main goal of our admission control policy is to protect small model-based flows, which represent interactive web flows that have real-time quality requirements. The expected outcome of admission control should be reduced response times for these small flows.

We first show the packet loss rates when using admission control with a $q1$ size of 10 packets. Figure 10 shows that employing admission control can greatly reduce loss for model-based flows.

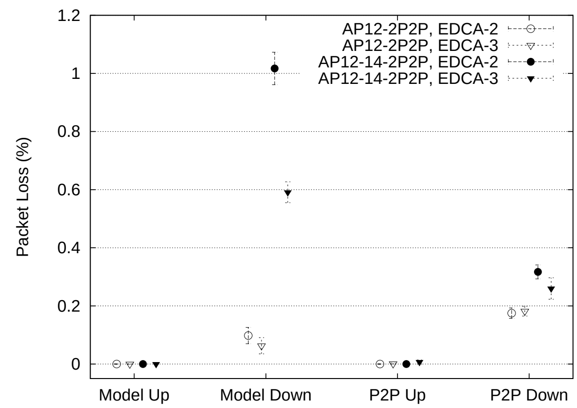

Figure 8: Packet loss for model-based and P2P flows in both upstream and downstream directions with $q1$ = 10 packets

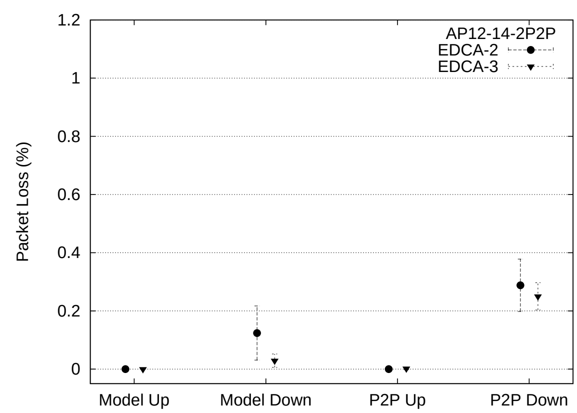

Figure 9: Packet Loss for model-based and P2P flows in both upstream and downstream directions with $q1$ = 20 packets

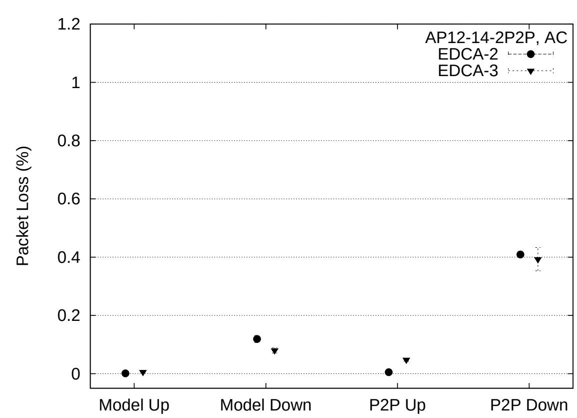

Figure 10: Packet loss for model-based and P2P flows in both upstream and downstream directions with $q1$ = 10 packets and admission control

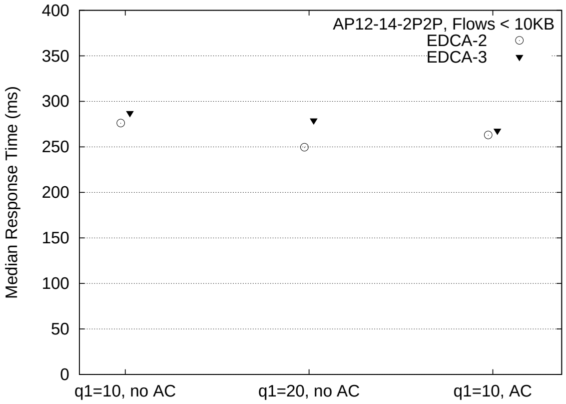

Figure 11: Median response times for model-based flows less than 10 KB

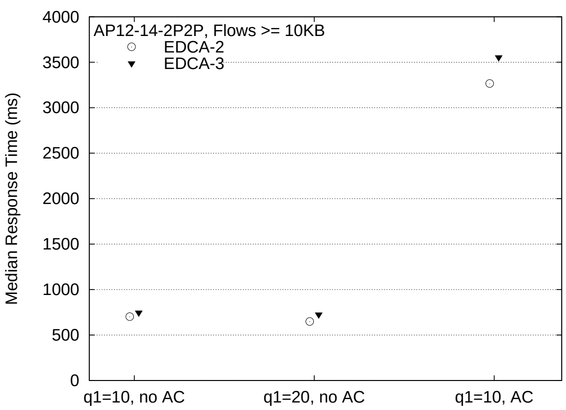

Figure 12: Median response times for model-based flows greater than 10 KB

Figure 11 shows the median response times for model-based flows less than 10 KB with AP12-14-2P2P traffic. We show the performance for $q1$ with 10 packets and no admission control, $q1$ with 10 packets and admission control, and $q1$ with 20 packets and no admission control. Admission control does reduce the response times slightly, but increasing the $q1$ buffer also has a positive impact on response times.

Figure 12 shows the median response times for larger model-based flows. Using admission control greatly increases the median response time for these flows, but that was expected with the admission control policy set forth. Flows larger than 10 KB were put into $q2$ along with the P2P flows. The buffer size at $q2$ was 100 packets, so these longer flows had a larger queuing delay as well as a lower priority than the shorter flows. In addition, if $q2$ was near capacity, these flows would have been passed to $q3$, which has an even lower priority.

Although using admission control does not reduce response times for small flows much beyond the gains obtained by increasing the buffer at $q1$, we anticipate that with increasing loads, increasing the size of the buffer will no longer be an ideal option. The larger the queue buffer, the larger queuing delays will be, thus increasing response times. If we apply admission control policies to keep the queue size small, we can protect smaller interactive flows, while still admitting larger flows and avoiding packet loss.

7. CONCLUSIONS

The support of Internet applications over the WLAN is dependent on the QoS assurance for these applications. This requires the fair allocation of system bandwidth to traffic in both upstream and downstream directions. In this work we have shown that by using the 802.11e EDCA MAC the upstream / downstream bias can be tuned. Furthermore the assignment of traffic to different access categories helps in meeting the varied QoS objectives of the different applications. However the tuning of the MAC parameters to achieve the required QoS criteria is non-trivial.

Based on the extensive set of simulations that we have performed, we conclude that for optimal QoS assurance for different traffic classes, it is necessary to jointly optimize the following parameters: (1) the AIFS interval, (2) CW interval, (3) the queue sizes of the different access categories, and (4) the admission control policy. It is important to carefully set both the AIFS and CW intervals, keeping in mind the traffic types in the different access categories. The optimal allocation of queue sizes to the different access categories considers the limitations imposed by the system resources. The choice of the admission control scheme can play a role in reducing the response times for high-priority flows while keeping small queue buffers to minimize delay.

Our future work will focus on further understanding of the interaction of the different MAC parameters as it applies to the QoS metrics of different applications. Specifically we will study the behavior of voice traffic and other UDP-based best effort services. We will also explore the admission control policies in more detail and provide quantifiable performance bounds that take into consideration both system resources as well as available bandwidth. We will also evaluate the performance of the EDCA scheme under a realistic scenario by taking into consideration both channel conditions and node mobility.

8. REFERENCES

[1] ANSI/IEEE. 802.11: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications. IEEE Std 802.11, 1999 Edition (R2003), 1999.

[2] ANSI/IEEE. 802.11e: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications, Amendment 8: Medium access control (MAC) enhancements for quality of service (QoS). IEEE Std 802.11e-2005, Nov. 2005.

[3] A. Balachandran, G. M. Voelker, P. Bahl, and P. V. Rangan. Characterizing user behavior and network performance in a public wireless LAN. In *Proceedings of ACM SIGMETRICS*, pages 195–205, June 2002.

[4] A. Banchs, A. Azcorra, C. Garcia, and R. Cuevas. Applications and challenges of the 802.11e EDCA mechanism: An experimental study. *IEEE Network*, pages 52–58, July/August 2005.

[5] M. Bottigliengo, C. Casetti, C.-F. Chiasserini, and M. Meo. Short-term fairness for TCP flows in 802.11b WLANs. In *Proceedings of IEEE INFOCOM*, 2004.

[6] R. Bruno, M. Conti, and E. Gregori. Throughput evaluation and enhancement of TCP clients in Wi-Fi hot spots. In *Proceedings of the 1st Working Conference on Wireless On-demand Network Systems (WONS 2004)*, pages 73–86, 2004.

[7] C. Casetti and C.-F. Chiasserini. Improving fairness and throughput for voice traffic in 802.11e EDCA. In *Proceedings of IEEE PIMRC*, 2004.

[8] S. Gopal, S. Paul, and D. Raychaudhuri. Investigation of the TCP simultaneous-send problem in 802.11 wireless local area networks. In *Proceedings of the IEEE International Conference on Communication (ICC)*, Seoul, South Korea, May 2005.

[9] S. Gopal and D. Raychaudhuri. Experimental evaluation of the TCP simultaneous-send problem in 802.11 wireless local area networks. In *Proceeding of the 2005 ACM SIGCOMM Workshop on Experimental Approaches to Wireless Network Design and Analysis*, pages 23–28, 2005.

[10] A. Grillo and M. Nunes. Performance evaluation of IEEE 802.11e. In *Proceedings of IEEE PIMRC*, 2002.

[11] T. Henderson, D. Kotz, and I. Abyzov. The changing usage of a mature campus-wide wireless network. In *Proceedings of ACM MOBICOM*, Philadephia, PA, 2004.

[12] F. Hernandez-Campos and M. Papadopouli. Assessing the real impact of 802.11 WLANs: A large-scale comparison of wired and wireless traffic. In *Proceedings of the 14th IEEE Workshop on Local and Metropolitan Area Networks (LANMAN)*, 2005.

[13] F. Hernandez-Campos and M. Papadopouli. A comparative measurement study of the workload of wireless access points in campus networks. In *Proceedings of IEEE PIMRC*, 2005.

[14] D. Kotz and K. Essien. Analysis of a campus-wide wireless network. In *Proceedings of ACM MOBICOM*, pages 107–118, 2002.

[15] D. J. Leith and P. Clifford. Using the 802.11e EDCF to achieve TCP upload fairness over WLAN links.

[16] S. McCanne and S. Floyd. ns Network Simulator. Software available at `http://www.isi.edu/nsnam/ns/`.

[17] X. G. Meng, S. H. Wong, Y. Yuan, and S. Lu. Characterizing flows in large wireless data networks. In *Proceedings of ACM MOBICOM*, pages 174–186, 2004.

[18] T. Nandagopal, T.-E. Kim, X. Gao, and V. Bharghavan. Achieving MAC layer fairness in wireless packet networks. In *Proceedings of ACM MOBICOM*, Boston, MA, 2000.

[19] C. H. Ng, J. Chow, and L. Trajkovic. Performance evaluation of TCP over WLAN 802.11 with the Snoop performance enhancing proxy. In *Proceedings of OPNETWORK*, 2002.

[20] Q. Pang, S. C. Liew, C. P. Fu, W. Wang, and V. O. Li. Performance study of TCP Veno over WLAN and RED router. In *Proceedings of GLOBECOM*, 2003.

[21] S. Pilosof, R. Ramjee, D. Raz, Y. Shavitt, and P. Sinha. Understanding TCP fairness over wireless LAN. In *Proceedings of IEEE INFOCOM*, 2003.

[22] I. Tinnirello, G. Bianchi, and L. Scalia. Performance evaluation of differentiated access mechanisms effectiveness in 802.11 networks. In *Proceedings of GLOBECOM*, 2004.

[23] N. H. Vaidya, P. Bahl, and S. Gupta. Distributed fair scheduling in a wireless LAN. In *Proceedings of ACM MOBICOM*, 2000.

[24] S. Wiethölter and C. Hoene. Design and verification of an IEEE 802.11e EDCF simulation model in ns-2.26. Technical Report TKN-03-019, Telecommunication Networks Group, Technische Universität Berlin, Dec. 2003.

[25] H. Wu, Y. Peng, K. Long, S. Cheng, and J. Ma. Performance of reliable transport protocol over IEEE 802.11 wireless LAN: Analysis and enhancement. In *Proceedings of IEEE INFOCOM*, 2002.

[26] H. Xu, Q. Xue, and A. Ganz. Adaptive congestion control in infrastructure wireless LANs with bounded medium access delay. In *Proceedings of the Workshop on Mobility and Wireless Access (MobiWAC)*, 2002.

[27] S. Yi, M. Keppes, S. Garg, and X. Deng. Proxy-RED: An AQM scheme for wireless local area networks. In *Proceedings of ICCCN*, Chicago, IL, Oct. 2004.