

Experiences with a Centralized Scheduling Approach for Performance Management of IEEE 802.11 Wireless LANs

Malati Hegde*, Pavan Kumar*, K. R. Vasudev*, N. N. Sowmya*, S. V. R. Anand*,
Anurag Kumar*, Joy Kuri†

*Dept. of Electrical Communication Engineering, †Department of Electronic Systems Engineering,
Indian Institute of Science, Bangalore, India

Email: malati, anand, anurag@ece.iisc.ernet.in, pavan.k.pujari, vasudevkr@gmail.com, kuri@cedt.iisc.ernet.in

Abstract—We present a centralized integrated approach for (i) enhancing the performance of an IEEE 802.11 infrastructure WLAN, and for (ii) managing the access link that connects the WLAN to the Internet. Our approach, which is implemented on a standard Linux platform, and which we call ADWISER (ADvanced Wi-fi Internet Service EnhanceR), is an extension of our previous system WM (WLAN Manager [1], [2]).

ADWISER addresses several infrastructure WLAN performance anomalies such as mixed-rate inefficiency, unfair medium sharing between simultaneous TCP uploads and downloads, and inefficient utilisation of the Internet access bandwidth when Internet transfers compete with LAN-WLAN transfers, etc. The approach is via centralized queueing and scheduling, using a novel, configurable, cascaded packet queueing and scheduling architecture, with an adaptive service rate. In this paper, we describe the design of ADWISER, and report results of extensive experimentation conducted on a hybrid test-bed consisting of real end systems and an emulated WLAN on Qualnet. We also present results from a physical test-bed consisting of one AP and a few end systems.

Index Terms—WLAN QoS management, WLAN controllers, split-MAC WLAN controllers, IEEE 802.11 wireless networks

I. INTRODUCTION

Enterprise and home wireless local area networks (WLANs) are based on the CSMA/CA based Distributed Coordination Function (DCF) MAC (medium access control) standardized in IEEE 802.11 [3], and commercially referred to as “WiFi”. It is well known that, due to the intrinsic properties of the DCF MAC protocol, there are several limitations to the QoS (quality of service) offered by WiFi WLANs. For example, from Figure 1 we see that when out of 3, 6, 9, 12, etc., stations (STAs), $\frac{1}{3}$ of the STAs are associated with an Access Point (AP) at each of the rates 54 Mbps, 24 Mbps, and 6 Mbps, then the aggregate bulk, TCP controlled, download throughputs of the STAs associated at each rate is *the same* (just over 3 Mbps: bottom three curves in Figure 1) and the total download throughput is about 9.5 Mbps (third curve from the top), whereas when all the STAs are associated at 54 Mbps then the aggregate throughput is about 22.5 Mbps (the top curve); see also [4] for the same observation in the *saturated*

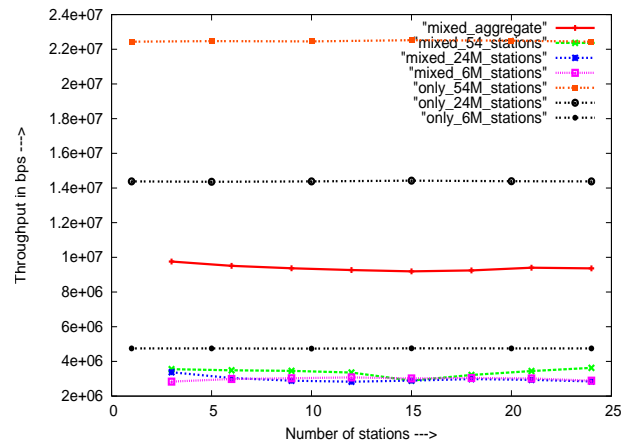


Fig. 1. Simulation results showing TCP download throughputs for n STAs when *all* n STAs are associated at the PHY rates of 54 Mbps (plots labeled “only”), or at 24 Mbps, or at 6 Mbps, and also when $\frac{1}{3}$ of the STAs are at each of the three rates (plots labeled “mixed”).

node setting. In practical WLAN deployments, such mixed rate associations can be expected, thus leading to a reduction in the overall network throughput. For an analytical understanding of this phenomenon, in the context of TCP controlled bulk transfers, see [5] and [6].

Another QoS problem reported (see [7]) is that of unfairness between upload and download TCP controlled bulk transfers for practical sizes of the AP buffer (which stores the packets contending for access to the wireless medium), see [8] for an analytical model that explains the observations in [7]. Real time interactive traffic such as that due to voice over IP (VoIP) telephony is known to perform poorly in the presence of TCP traffic in WLANs. To solve this problem, the IEEE 802.11e standard provides an Enhanced DCF (EDCF) in which STAs carrying voice transfers can obtain prioritized service. While these mechanisms serve to isolate VoIP traffic from TCP transfers, admission control of VoIP connections is needed, specially if some rate guarantees are to be provided to TCP transfers (see [9]).

We earlier proposed WM ([1], [2]), a device through which all traffic between the WLAN and the wireline LAN passes (see Figure 2). WM can be viewed as implementing a Split-MAC based architecture (see the RFCs [10] and [11]) that

This work was supported by research grants from the Department of Information Technology (DIT), and the Department of Science and Technology (DST), India.

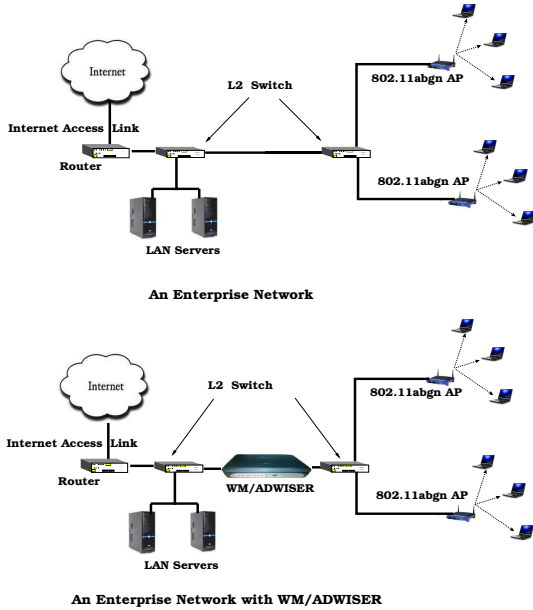


Fig. 2. Top: An enterprise network in which all traffic between the WLAN and the wireline local network passes through one Ethernet link. Bottom: The same network with WM or ADWISER in place.

works with off-the-shelf commercial products without requiring any changes in the APs. While retaining several design principles employed in WM [1], [2], in the present paper, we extend the functionality to address the unfair Internet bandwidth sharing that arises when there is a mix of Internet-WLAN and LAN-WLAN TCP transfers. This unfairness can be explained as follows. The large difference in the RTTs seen by these flows causes unequal buffer sharing at the AP in favour of flows with smaller RTTs, which results in the AP more often “serving” LAN-WLAN packets as compared to WAN-WLAN packets (see Section II-D for an experimental demonstration of this problem). We propose a cascaded scheduling mechanism which eliminates this unfairness problem. This additional new feature thus enables us to manage both the Internet access bandwidth and enhance WLAN performance of an enterprise having LAN and WLAN users, all within a single box, *without any modifications to the AP or the STAs*. We have appropriately named this evolution of WM as ADWISER (ADvanced Wi-fi Internet Service Enhancer).

Contributions of this Paper: This paper is about implementation of ADWISER and its experimental evaluation. Our experimental work has been performed on a hybrid test-bed (in which the PHY and MAC of the WLAN are modelled on the Qualnet simulator, whereas the servers, STAs and the wireline LAN are physical entities), and also on a real test-bed with a single AP and several STAs. We report our experiences in implementing ADWISER on an off-the-shelf Linux platform, and we provide the implementation details. We then report the highly encouraging experimental results obtained, and the insights gained; we focus here exclusively on QoS management of TCP transfers.

In Section II we provide an experimental review of the TCP QoS problems that ADWISER aims to solve. This review makes the paper self-contained, and also helps to understand

the results on the experimental evaluation of ADWISER. Further, to the best of our knowledge, the experimental observations of starvation of the Internet access link during simultaneous WAN-WLAN and LAN-WLAN transfers, and the tradeoff between VoIP calls and TCP throughput, under IEEE 802.11e EDCF, are novel in the IEEE 802.11 WLAN setting. Section III describes the design principles behind ADWISER, some details of the system architecture, and the various queueing, scheduling and service rate adaptation algorithms. In Section IV we provide an extensive experimental study that illustrates the efficacy of ADWISER in solving the TCP QoS problems.

Related Literature: Since the publication of our proposal [1], another research group has independently reported a similar idea (see [12] and the references therein). While the general motivating concerns in the two proposals are similar, we have taken some specific TCP performance issues, namely (i) fairness and efficiency due to different PHY rates, (ii) fairness for arbitrary combinations of upload and download connections, (iii) prevention of starvation of transfers from the Internet in the presence of transfers from the wired LAN, and (iv) policy based TCP throughput differentiation, to the stage of complete implementation as an ADWISER prototype. Since our implementation is not in the Access Point, in order to achieve these objectives, ADWISER has to adaptively learn the *effective service capacity* of the WLAN.

TCP-related QoS problems in WLANs have been addressed earlier by many authors. Multirate anomaly was addressed in [13], where time-fairness was ensured using a Leaky Bucket. The work in [14] proposed the control of the “channel occupancy time” of a TCP connection by TCP window adjustments, based on link-layer measurements.

When STAs are associated with an AP at the *same* rate, fairness issues arise during simultaneous uploads and downloads. This was addressed in [15], where, to provide downloads adequate medium access opportunities, the value of the uplink TCP receiver’s advertised window was reduced as the ACKs passed the AP. The proposal in [16] was to restrict the rate at which uplink data packets can reach the corresponding TCP receivers by incorporating a rate limiter inside the AP. On the other hand, the approach in [17] was to introduce separate queues for storing TCP data packets and ACKs, and serving the latter with higher probability. The idea in [18] is similar, except that a Weighted Round Robin (WRR) scheduler is used to serve the queued data and ACK packets. Yet another approach was followed in [19], where uplink packets were dropped deliberately at the AP, to trigger congestion control mechanisms at the TCP sender. The approach in [20] is based on per-connection queueing and fair scheduling at the AP, with the purpose of releasing TCP ACKs in a controlled manner to slow down an uplink TCP connection’s window growth. The Selective Packet Marking (SPM) scheme in [7] aims to protect the first few packets in a downlink transfer (crucial for window growth) by marking them as high-priority and requiring the AP to buffer them preferentially. Unfairness between simultaneous uplink transfers has been reported earlier in [21]. This paper provides one solution by invoking IEEE 802.11e service differentiation to handle downlink TCP ACKs in the AP;

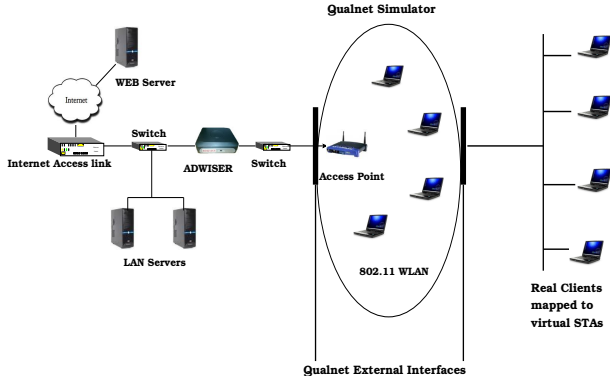


Fig. 3. ADWISER in a hybrid test-bed consisting of real end systems and a WLAN emulated in the Qualnet simulator.

another solution is provided in [22], where the idea is to delay TCP ACKs by per-flow queuing.

Some authors have taken the approach of modifying MAC layer characteristics to ensure improved fairness at higher layers. The proposal in [23] is to set the MTU of a station proportional to its rate of association; effectively, this ensures time-fairness. In [24], the same effect is achieved by setting the initial congestion window of a station CW_{min} to a value inversely proportional to the station’s rate of association. The uplink-downlink fairness issue is addressed in [25] by utilizing the shortest interframe spacing PIFS for downlink access by APs. The work reported in [26] explores the use of the parameters CW_{min} as well AIFS of the 802.11e standard to alleviate unfairness between competing uplink TCP transfers.

Thus, these approaches require either changes to MAC parameters and/or modifications to the firmware running on the AP. On the other hand, ADWISER does not require any modification to the MAC parameters or the AP firmware, and works with any existing IEEE 802.11 based infrastructure WLAN.

II. TCP QoS PROBLEMS: EXPERIMENTAL ILLUSTRATIONS

In the process of understanding the QoS problems and developing ADWISER, we have used a *hybrid test-bed*, which combines physical network devices and a QualNet simulator [27]. In our experimental evaluation section, we also report results from a test-bed with all physical network devices.

The hybrid test-bed is shown in Figure 3, and permits us to experiment with a larger WLAN, in a controlled environment, than is practical to set up in our laboratory. The QualNet simulator (v 4.5) is used to emulate the MAC and PHY of the WLAN devices. QualNet provides a mechanism to inject actual packet traffic into the simulated wireless nodes through its external interface APIs. As shown in the figure, the hybrid test-bed comprises a physical server, an Internet link emulator (8Mbps, 300ms RTT emulated using Linux *netem*), a router, an Ethernet LAN switch, the Linux machine on which ADWISER is implemented, and laptops on which the client applications run. The QualNet simulator models the AP, the STAs (each corresponding to one of the laptops), and the wireless medium. Although Qualnet implements a rate-adaptation algorithm, this is turned off in our experiments, so that each STA is always associated at a fixed PHY rate with an AP. This facilitates

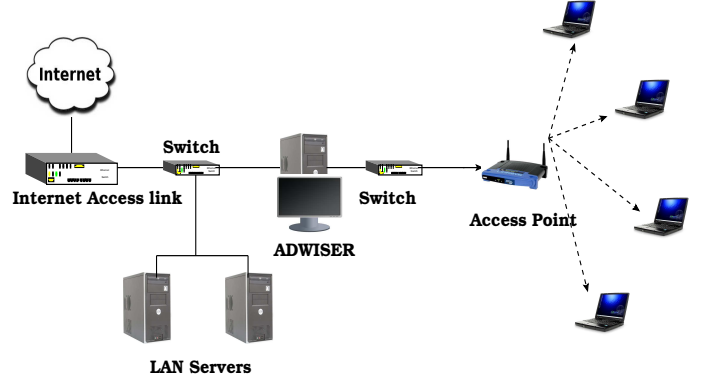


Fig. 4. Physical test-bed for WAN-WLAN and LAN-WLAN transfers. By a software command, ADWISER can be bypassed or inserted.

comparison of experimental results with those predicted from analysis. Our TCP transfers are initiated from the end systems; then the packets of the TCP flows pass through the simulated MAC in the QualNet simulator, thus experiencing the WLAN contention and delays as they would have in a real network.

Our experiments with actual access points were conducted on the *physical test-bed* shown in Figure 4. The test-bed comprised a single WME (WiFi Multimedia Extensions) compliant Cisco Aironet AP (1240 series), and standard laptops running Linux (kernel release 2.6.23, Fedora 8 distribution). In our experiments, the AP is configured to associate STAs at a single rate, thus effectively disabling the rate adaptation algorithm. An 8 Mbps Internet access link with a propagation delay of 300ms was configured using the Linux *netem* utility.

Where we report experiments with web traffic, we have used scripts that generate *web-like* traffic. A “web page” is modeled as containing multiple objects composed of text and images. Every page downloaded has a main file plus a random number of auxiliary files; this random number takes values between between 1 and 6. After each page download, there is an inactive think time. The files in the web page are sampled from a Pareto distribution (the complementary c.d.f. has tail exponent = 1.1) with an average file size of 20KB. On an average, every page has 4 files with an average size of 20KB and these files are downloaded in parallel over separate TCP connections from the web server.

A. Multi-Rate Unfairness

This problem was already mentioned in the Introduction and illustrated via Figure 1, which was obtained from the hybrid test-bed. The mixed PHY rate scenario leads to a drastic reduction in the overall WLAN throughput, while equalising connection throughputs, since the DCF MAC is *packet fair* and hence gives an equal chance for all backlogged nodes to contend, irrespective of the PHY rates at which they can send.

B. Upload-Download Unfairness

In addition to the literature discussed earlier, we have made observations about TCP upload-download unfairness that we briefly present here. Unfairness between competing TCP uploads and downloads has been addressed in [28] as well. However, apart from the forward/reverse path asymmetry

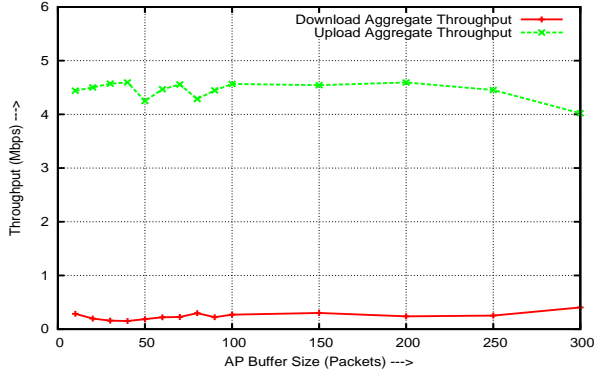


Fig. 5. Aggregate upload and download throughputs of 10 STAs (5 performing uploads and 5 performing downloads) at 11 Mbps with varying input AP buffer size. Hybrid test-bed; TCP window scaling is enabled.

mentioned in [28], the issue of AP buffer space is important, as the discussion below shows.

Let us first recall the reason for the upload-download unfairness reported in [7]. An analytical model for the observations in [7], can be found in [8]. In an infrastructure WLAN, all traffic between the STAs and a server on the wireline network must pass through the AP. We recall that the AP is one device that contends to send packets for all the connections, whereas each STA contends only for its own connection. The AP is thus a bottleneck in the infrastructure WLAN setting (see [5] and [6]). The packets from upload and download TCP connections (respectively, TCP ACKs, and TCP data packets) fill up the AP queue, which is stored in a finite buffer. When the TCP windows grow, packet drops take place at the tail of the AP buffer. For the downlink transfers, data packet drops at the AP buffer affect the growth of the TCP congestion windows. Since ACK packets (corresponding to uplink transfers) are smaller, they suffer a smaller drop rate. Further, TCP ACKs are cumulative, i.e., each ACK acknowledges all the packets up to the packet that it acknowledges. Hence, uplink transfers are not as severely affected by tail-drop packet losses at the AP buffer. In our experiments, we have found that with the recent versions of TCP that implement window scaling and the availability of large TCP receive buffers (8 MB was the default in our experiments), this reported upload-download unfairness can be very severe, and is not relieved by increasing the AP buffer space; see the results reported later in this section.

To illustrate the unfairness problem between uplink and downlink traffic flows, we ran experiments on the hybrid test-bed. Ten wireless STAs are connected at the 11 Mbps PHY rate to an AP. STA 1 to STA 5 each upload a large file to a local server (on the wired LAN) via the AP, while STA 6 to STA 10 download a file of the same size from the server. All the transfers are initiated simultaneously using the Linux *wget* file transfer utility. The TCP receivers employ the delayed ACK mechanism. The buffer size in the AP is varied from 10 to 350 packets (15 KBytes to 525 KBytes buffer at AP) and the TCP maximum segment size (MSS) is set to 1500 bytes. The 10 transfers are allowed to run for a period of 300 seconds. Figure 5 shows the aggregate upload and download throughputs for the two groups of 5 STAs each, for varying AP buffer size. We see that the 5 STAs uploading files to the server obtain by far the larger part of the network throughput.




Fig. 6. Congestion windows (vs. time) for the 5 upload TCP transfers, when 5 upload transfers and 5 download transfers are in progress. Hybrid test-bed; AP buffer size = 100 packets.

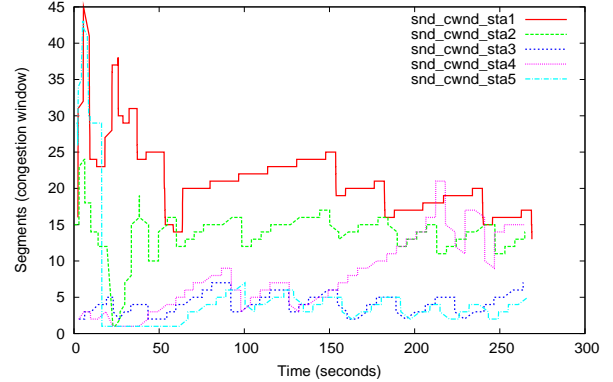


Fig. 7. Congestion windows (vs. time) for the 5 download TCP transfers, when 5 upload transfers and 5 download transfers are in progress. Hybrid test-bed; AP buffer size = 100 packets.

We observe that an increase in the AP buffer size has little impact on the aggregate throughput seen by both uplink and downlink transfers.

Figure 6 and Figure 7 show the evolution of the TCP congestion windows for the uplink and downlink transfers of these 10 STAs for an input AP buffer size of 100 packets (150 KBytes)¹. Each DATA packet is 1500 bytes, whereas a TCP ACK packet is 40 bytes, plus fields for options such as time-stamp and SACK (if any). Thus, the AP buffer can accommodate about 100 DATA packets, but 1000s of TCP ACK packets. The round trip time between the AP and the server on the LAN is negligible, hence, as explained earlier in this section, the congestion windows of all connections, reside almost entirely in the AP buffer, and the windows of the uplink transfers are able to grow to large values (300 to 700 segments), with the TCP ACKs then occupying more than 50% of the AP buffer (see Figure 6). On the other hand the congestion windows of the downlink transfers grow to just 5 to 15 segments, and the remaining AP buffer is occupied by these packets (see Figure 7). Hence, the AP more often serves ACKs for the uplink connections, which therefore obtain a very high throughput as compared to the downlink transfers.

¹Although the buffer size is specified in number of packets, it is maintained internally as an equivalent number of contiguous bytes.




Fig. 8. Upload throughputs; 10 STAs associated at 11 Mbps; hybrid test-bed.

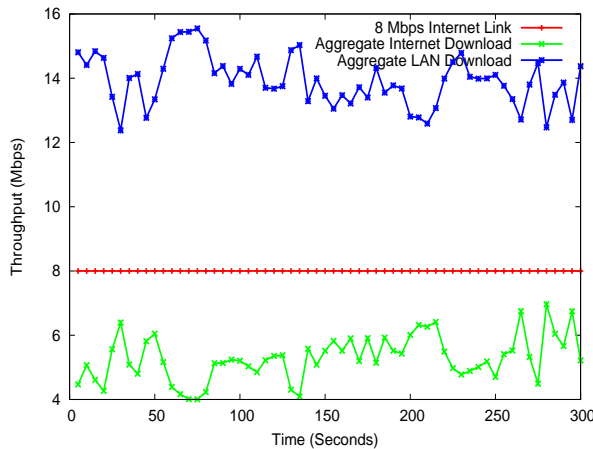


Fig. 9. Aggregate throughputs for WAN-WLAN and LAN-WLAN transfers; physical test-bed of Figure 4

C. Unfairness Among Multiple Uploads

When there are only several upload transfers, a newly initiated TCP flow can lose ACK packets associated with its first few data transmissions, thus inducing a timeout. The ACK packets for the retransmitted data packets can also be lost, leading to further timeouts (with associated doubling of the retransmit timer), and so the flow can be completely starved for long periods. We conducted an experiment on the hybrid test-bed with 10 STAs, all associated at 11 Mbps rate with one AP. Upload transfers are initiated simultaneously from all the STAs and throughput observations are made for 300 seconds. The bar plot in Figure 8 shows the average throughputs of the 10 upload connections. Only some of the STAs (6 to 10) obtain considerable bulk transfer throughputs, whereas some STAs (3, 4 and 5) obtain much smaller throughputs. Thus, even among upload transfers there can be random unfairness depending on loss events, and the consequent unfairness in TCP window evolution.

D. WAN-WLAN and LAN-WLAN Unfairness

Unfairness among competing TCP flows at a bottleneck router with different RTTs is common in wired networks. The large difference in the RTTs causes unequal sharing of the bottleneck link, favouring the TCP connections with

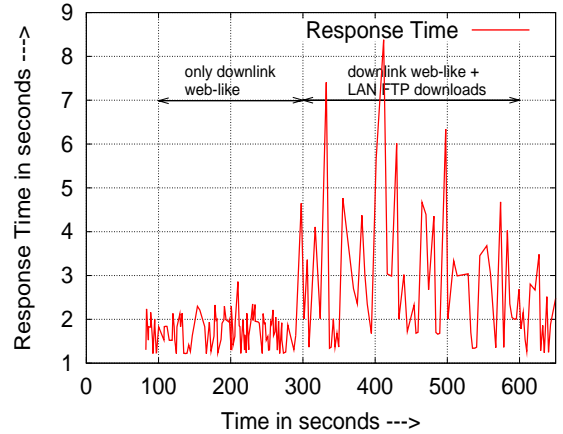


Fig. 10. Response times of web-like downloads from the server on the WAN, in the presence of bulk download from the wireline LAN; physical test-bed. The completion times of successive downloads are plotted as points joined by lines.

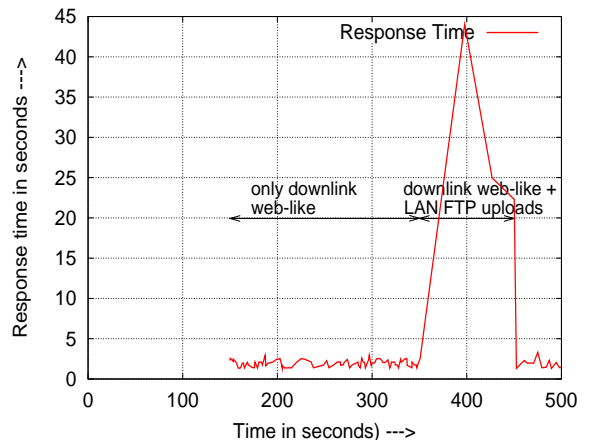


Fig. 11. Response times of web-like downloads from the server on the WAN, in the presence of bulk uploads to the wireline LAN; physical test-bed. The completion times of successive downloads are plotted as points joined by lines.

smaller RTT. The same phenomenon can be observed in the infrastructure WLAN network with STAs initiating transfers from Internet servers and local LAN servers. This causes unfair wireless channel access times between the competing flows from LAN-originated traffic and the Internet traffic destined to the WLAN (see also [29]). This unfairness can be explained as follows.

Consider the situation when both the transfers are downlink. The WAN-WLAN connection has a large round trip propagation delay, and hence requires a large number of packets in flight in order to sustain its throughput. Losses at the AP buffer result in its window being reduced, much of the window is “in flight,” and little remains in the AP buffer. On the other hand the LAN-WLAN connection requires a very small number of packets in flight, can rapidly recover from losses, and ends up occupying most of the AP buffer. Thus, since the AP renders FIFO service to packets in its buffer, it finds a LAN-WLAN packet to serve with a much higher probability than that of a WAN-WLAN connection. This results in the throughput unfairness that we have found below in our experiments.

We conducted an experiment on the physical test-bed shown

in Figure 4, with 4 STAs all associated at 54 Mbps. TCP transfers were initiated on the 2 STAs downloading files from the Internet server. Simultaneously, from the other 2 STAs, TCP download transfers from the local LAN servers were started. Figure 9 shows the aggregate TCP throughputs for Internet and LAN transfers. It is observed in Fig 9, that the LAN-WLAN transfers see higher throughputs, and the Internet access link gets starved. This is undesirable since the Internet access link is often an expensive resource for a small enterprise.

As a consequence of the above unfairness, we observed that when we introduced web-like transfers from the Internet server, the web response times (the time taken to download a complete page from the web server) increased substantially in the presence of LAN transfers.

The test-bed shown in Figure 4 was used for experiments with web-like transfers. One station is performing web-like downloads from the Internet server. In Figure 10, from 100s to 300s there are only web-like downloads, and during this period the average response time is 1.76s for an RTT of 300ms. Even though the Internet access link is 8Mbps, this response time is mainly dictated by the RTT, as the TCP connection setup followed by a web download of 20KB file will take approximately 5 RTTs. Then, from 300s to 600s, we initiated a TCP bulk download from the LAN server along with the web-like downlink transfers. We observe (again, from Figure 10) that the response times shot up to an average value of 2.60s, and became highly variable with large peaks.

We used the physical test-bed to demonstrate the response time variation for web-like downlink transfers in the presence of a bulk upload to a LAN server. In Figure 11, from 150s to 350s, one station downloads web-like traffic from an Internet server. Here we observe the average response time to be 1.85s. From 350s until 450s we initiated a TCP bulk upload to the LAN server along with the web-like downloads. We observe that the response time shot up to an average value of 30.43s. This large increase of response time of web transfers in the presence of a LAN upload as compared to a LAN download is due to the upload-download unfairness (Section II-B) which is exacerbated by the unfairness under discussion in this section.

E. Effect of VoIP calls on TCP throughput

It is well known that unless there is service differentiation between packet voice and TCP transfers, VoIP traffic obtains very poor performance. When service differentiation is provided, however, connection admission control (CAC) may need to be exercised for VoIP calls so as to guarantee some aggregate throughput for TCP transfers.

In IEEE 802.11 WLANs, service differentiation between voice and TCP packets is implemented by the Enhanced DCF (EDCF) mechanisms (standardized in IEEE 802.11e) which permits the voice access category to use more aggressive backoff and reattempt parameters. These mechanisms provide the voice traffic with an approximation to preemptive and rate priority. Harsha et al. [9] have provided an analytical model for the decrease in TCP throughput as the number of VoIP calls is increased. It was shown that, as the number of voice calls




Fig. 12. Download throughput obtained by a single TCP download with increasing number of VoIP calls. Results are from the hybrid test-bed. The MAC protocol is IEEE 802.11e, over 54 Mbps PHY; VoIP: G.711 coding, 200 byte packets; TCP DATA packets: 1500 bytes.

is increased up to the voice call capacity, the TCP throughput drops linearly with the number of accepted voice calls. We have corroborated the reported behavior on our hybrid test-bed shown in Figure 3. In Figure 12 we have plotted the TCP throughput vs. the number of VoIP calls for 802.11g/e, with the VoIP traffic being carried on AC 3. The STAs are associated at the PHY rate of 54 Mbps, the VoIP codec is G.711 with a packet size of 200B, and TCP data packet size is 1500 bytes. We have considered a single TCP download. As seen in Figure 12, the TCP throughput is close to 2.5 Mbps when the number of admitted VoIP calls reaches 40 (up to which point the target maximum delay of 20 ms for voice packets is respected). It can be seen that the TCP throughput drops roughly linearly at the rate of 0.45 Mbps for each admitted voice call. Thus, from Figure 12, we see that limiting the number of VoIP calls (of the above type) to, say 10, guarantees that the TCP transfers are guaranteed an aggregate throughput of about 17 Mbps.

III. ADWISER: APPROACH, ALGORITHMS, AND IMPLEMENTATION

A. The ADWISER Approach [1], [2]

We have implemented ADWISER on a standard personal computer platform on Linux. The device is located (see Figure 2) so that all packets that pass through the AP also pass through the device. In this subsection we present a conceptual overview of the ADWISER approach. Details are provided in the subsequent subsections.

In this paper we assume that the user devices are all on the WLAN. These devices could have TCP connections to servers on the wireline LAN or on the Internet (see Figures 3 and 4).

The basic ideas implemented in ADWISER [1], [2] are the following (see Figure 13):

a) Virtual servers and queueing: If packets are allowed to accumulate in the AP and the STAs, then the usual behavior of IEEE 802.11 DCF shows up; see Figure 1. ADWISER works by attempting to draw into itself the queues from the two bottleneck resources that it manages: the WLAN medium and the Internet access link. This is achieved by creating virtual servers inside ADWISER. Corresponding to the WLAN

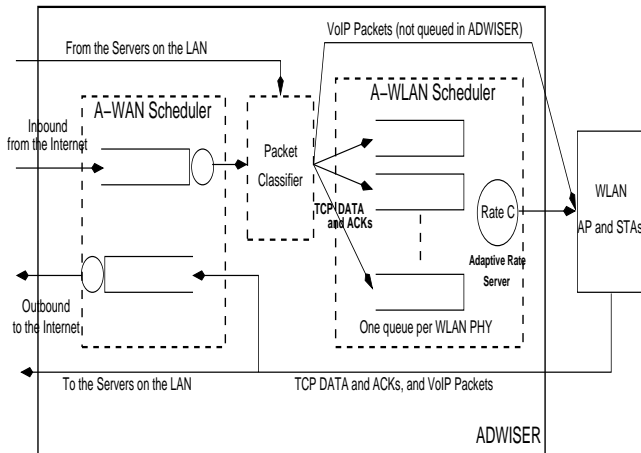


Fig. 13. ADWISER: A schematic of the scheduler that controls the sharing of the WLAN medium. All queues shown are hierarchical queues, so that service differentiation within each queue can be configured by the user.

medium, there is one virtual server (with its associated queues) since the IEEE 802.11 WLAN presents itself as a half-duplex wireless link; we call this the ADWISER WLAN (A-WLAN) scheduler. On the other hand, two virtual servers (with their associated queues) correspond to the inbound and outbound directions of the full-duplex Internet WAN access link; we call this the ADWISER WAN (A-WAN) scheduler. Queueing is forced to take place in ADWISER rather than at the devices physically connected to the bottleneck links (i.e., the AP in the case of the WLAN medium, and the access routers in the case of the Internet access link) by setting the service rates of the virtual servers a little *less* than the effective service rates of the corresponding physical links. This approach also prevents packet drops from occurring due to buffer overflow in the AP. Ideally, with ADWISER, the queue build-up at the APs should be minimal, without starving the wireless medium. All queues shown in Figure 13 are hierarchical in order to facilitate additional user-configured service differentiation. Coordination between the A-WAN and the A-WLAN schedulers is crucial for addressing the throughput unfairness reported in Section II-D and is described in Section III-C2.

b) Flow of and scheduling of TCP and VoIP traffic:

For TCP-controlled traffic, all packets flowing towards the WLAN (TCP DATA or TCP ACKs arriving from servers on the Internet or on the wired LAN) are classified and queued into per rate class queues behind the A-WLAN server². These queues are then served by a virtual server using a fair queueing algorithm, which can enforce user-configured rate differentiation between the TCP transfers (see Figure 13). Control of TCP download traffic (from the servers to the STAs) and of TCP upload traffic (from the STAs to the servers)

²Even in the presence of rate adaptation, we have found experimentally (with the Minstrel algorithm; see [30]) that, for various ranges of received signal strength (RSSI) values, there are dominant rates at which packets are sent between the AP and an associated STA. SNMP polls to the AP elicit the RSSI values at which packets are received from an STA. These RSSI values can be used to map the STAs into, say, 3 or 4 rate classes. In our experiments, however, rate adaptation was effectively disabled: in the hybrid experiments by turning rate adaptation off in the Qualnet simulator, and in the physical AP experiments by permitting the AP to associate STAs only at a particular rate, e.g., 54 Mbps or 11 Mbps.

is achieved by serving the TCP DATA packets and the TCP ACKs appropriately (see Section III-C3). VoIP packets flowing towards the WLAN are placed in a strict priority queue in the A-WLAN server and essentially pass through freely into the AP. All packets flowing out of the WLAN towards the servers on the wired LAN are passed through ADWISER without any further queueing. All packets arriving into ADWISER from the WAN access link are queued in the A-WAN scheduler at the virtual server corresponding to the inbound direction of the Internet access link. Similarly, packets arriving from the WLAN into ADWISER, and destined for the Internet access link are queued in the A-WAN scheduler behind the virtual server corresponding to the outbound direction of the Internet access link.

c) *Rates of the virtual servers:* Whereas there is a well defined bit rate of the Internet access link, e.g., 8 Mbps in our experiments, the effective service rate of the WLAN medium depends on the medium access overheads and the channel conditions prevailing on the wireless links between the AP and the STAs. Further, since IEEE 802.11e service differentiation is used between VoIP connections and TCP connections, the effective service rate that the WLAN medium provides to the TCP connections depends on the number of (and the coding used by) the VoIP connections. A key requirement in ADWISER is that the rate of the virtual server in the A-WLAN scheduler is dynamically adjusted by an on-line rate adaptation algorithm. This service rate adaptation is crucial to the proper working of ADWISER. If the service rate is too high, packet accumulation occurs in the AP and not in ADWISER, and thus ADWISER loses the ability to enforce the desired service ratio between the packets of the various connections. On the other hand, if the A-WLAN server's rate is too small, then the WLAN starves and the system is inefficient. We have implemented a simple service rate adaptation algorithm that dynamically adapts the service rate, even as the number of STAs associated at each rate varies; see Section III-C1c for details.

d) *VoIP CAC:* ADWISER can perform VoIP connection admission control in order to provide desired aggregate throughput to the TCP transfers. Having admitted a VoIP call, ADWISER exercises no additional control on VoIP packets, letting the IEEE 802.11e mechanisms take care of providing the necessary service differentiation between VoIP and TCP transfers.

B. ADWISER Software Architecture

A schematic of the high level ADWISER software architecture is depicted in Figure 14. In the *packet capture module*, we capture the incoming traffic from the wired link towards the AP. The idea is to buffer only the packets coming from the wired link, going towards the WLAN, in order to achieve the desired fairness in both directions. Note that this is possible for TCP upload connections since the flow of ACKs in the downlink direction can be controlled, thereby controlling the TCP throughput. The *packet classifier module* places an arriving packet into one of the physical queues. The information in the IP and TCP headers is used to form

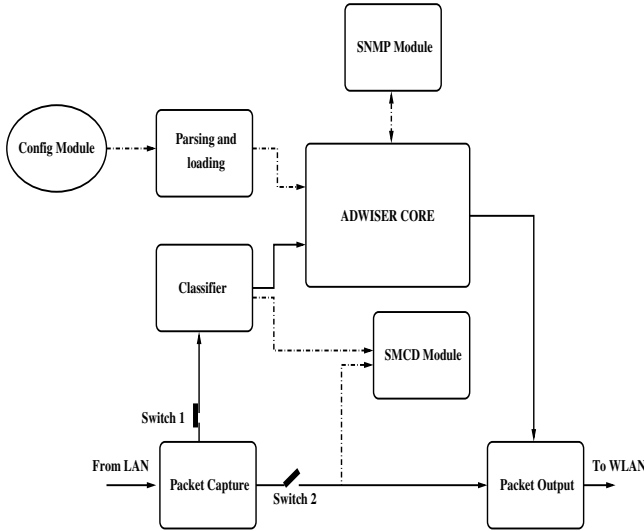


Fig. 14. ADWISER: High level software architecture. When the positions of the two “switches” are as shown, then ADWISER is in the path of the traffic flow; when both positions are reversed, ADWISER is bypassed.

a 4-tuple (source address, destination address, source port and destination port) in identifying the direction and the leaf node in the queuing hierarchy. The voice packets are given strict priority and appropriate DSCP marking is done in ADWISER before forwarding to the AP. These marked packets get mapped to the voice access category, AC3, within an 802.11e AP. ADWISER therefore takes advantage of the QoS offered by 802.11e. The *SNMP module* periodically polls the APs to obtain information on their associated STAs; this relationship is explained later in Section III-C1c. The *SMCD module* (where SMCD expands to Statistics and Measurements Capture Daemon) captures and stores various statistics of the traffic passing through ADWISER.

C. The Algorithms in ADWISER Core

The ADWISER Core (see Figure 13) is where all the QoS algorithms are implemented. In this section we describe these algorithms in detail.

1) Service Rate in the A-WLAN Scheduler:

a) *The Effective Service Rate for a PHY:* ADWISER uses SNMP polls to the AP to determine which STAs are associated and the PHY rates at which they are associated. Then for an STA associated at PHY rate, say R_i , the value of C_i , the effective bit rate at which TCP controlled transfers take place between the AP and this STA, if this was the only connection, is calculated as follows:

$$C_i = \frac{(d \times L_{DATA}) + L_{ACK}}{(d \times (\frac{L_{DATA}}{R_i} + T_{OVH})) + (\frac{L_{ACK}}{R_i} + T_{OVH})}$$

where $d = 2$ to take care of delayed TCP ACKs, and T_{OVH} is the overhead time due to the medium access (RTS/CTS for TCP DATA and Basic Access for TCP ACK), the various inter-frame spaces, and the MAC ACK. Since ADWISER’s attempt is to draw the queues into itself, thereby minimizing contention on the WLAN, when computing T_{OVH} just one back-off duration is taken and it is assumed that collisions do not take place. As will be seen below, any errors in the

computation of the C_i values only affect the weights, ϕ_i , that are used in the A-WLAN scheduler.

b) *The Need for Service Rate Adaptation:* The methodology to achieve the desired throughput fairness requires the ADWISER schedulers to serve the packets at certain desirable service rates; recall the discussion in Section III-A. We need to obtain the A-WLAN scheduler service rate that permits the queuing to take place in ADWISER so that it has control over the TCP throughputs, while keeping the WLAN maximally utilized.

Here is how we can think of the service rate. Suppose there are m STAs associated with the AP, at the PHY rates $r_i, 1 \leq i \leq m$. Each is carrying out a large file download from a server on the wired LAN. Let us imagine m queues in ADWISER, one for each station. The queues contain all the application level bits that need to be transmitted on the wireless medium for that connection. Thus, we imagine that, for each downlink data packet, even its uplink TCP ACK is queued in ADWISER; and for each downlink ACK packet, the corresponding TCP data packets that will be generated are also queued in ADWISER.

Suppose that ideal (bit level) fair queuing is used to serve these queues using the weights $\phi_i = (\phi_1, \phi_2, \dots, \phi_m)$, with $\sum_{j=1}^m \phi_j = 1$. Suppose all the queues are backlogged. Then, out of b bits (where, b is a large number) sent by ADWISER, $\phi_i b$ bits belong to connection i . With C_i as defined earlier in this section, the time occupied on the medium by these bits from connection i is $\frac{\phi_i b}{C_i}$. The total time taken to transmit all the b bits from the m connections is then given by $\sum_{i=1}^m \frac{\phi_i b}{C_i}$. Dividing b by this expression yields the effective rate at which the medium will carry bits above the MAC layer, i.e.,

$$C^* = \frac{1}{\frac{\phi_1}{C_1} + \frac{\phi_2}{C_2} + \dots + \frac{\phi_m}{C_m}} \quad (1)$$

Thus, C^* depends on the weights, $\phi_i, 1 \leq i \leq m$, and the PHY rates at which the STAs are connected.

Now the idea of serving the queues of actual packets in ADWISER is the following:

- 1) Virtually replace each packet in the ADWISER queue by the number of higher layer bits (above MAC and PHY) that will need to be sent on the medium if the packet is released into the WLAN (details in Section III-C3).
- 2) Adapt the service rate of these queues (of virtual bits) so that the rate is a little less than C^* . Adaptation of C^* is needed since the population of active STAs, and the PHY rates at which they are associated will keep varying over time. In addition, the C_i are affected by VoIP calls.

As an example, if m_j STAs are associated at rate $r_j, 1 \leq j \leq n$, then we have n queues in ADWISER. Consider the following weights, for $1 \leq i \leq n$,

$$\phi_i = \frac{m_i C_i}{\sum_{j=1}^n m_j C_j} \quad (2)$$

i.e., *proportional fairness* or *time fairness* (i.e., the target throughputs are proportional to the PHY rates of the connec-

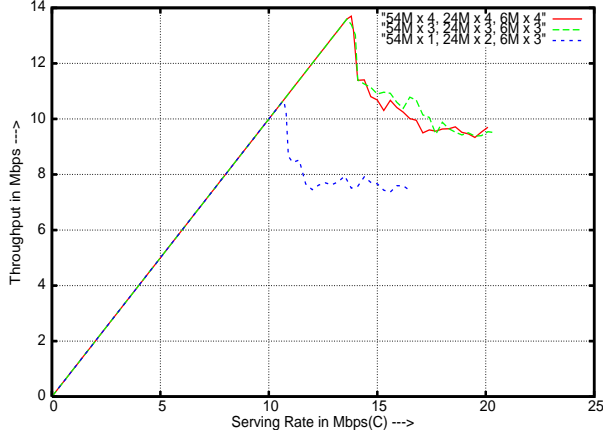


Fig. 15. Aggregate throughput for bulk file downloads, with ADWISER, for various numbers of STAs associated at 54, 24 and 6 Mbps (IEEE 802.11a) plotted vs. the ADWISER service rate C ; results from the hybrid test-bed.

tions), then

$$C^* = \frac{\sum_{i=1}^n m_i C_i}{\sum_{j=1}^n m_j} \quad (3)$$

The above discussion is illustrated by the hybrid experiment (recall Figure 3) results shown in Figure 15, where the scenario is that m_1, m_2 and m_3 STAs are associated, respectively, at the PHY rates of 54 Mbps, 24 Mbps, and 6 Mbps. In ADWISER there is one queue for each rate class. The service weights are given by the expression in Equation (2). The ADWISER service rate C is then varied and the aggregate download throughput is plotted vs. C . There are three plots, corresponding to $m_1 = m_2 = m_3 = 4$, $m_1 = m_2 = m_3 = 3$, and $m_1 = 1, m_2 = 2, m_3 = 3$. The calculated values of C^* , from Equation (3) (where the various parameter values from IEEE 802.11a are used), are 14.14 Mbps, 14.14 Mbps, and 11.00 Mbps. We notice that as C increases from a small value, the aggregate throughput is equal to C , showing that ADWISER is the bottleneck and the wireless medium is underutilized. However, when C approaches C^* , the aggregate throughput drops sharply and stabilises at a value significantly smaller than C^* , as C further increases. Basically, the queues move into the AP and the default behavior takes over, i.e., in the mixed PHY rate situation the aggregate throughput drops due to the prevalence of packet-fairness rather than time-fairness.

c) Service Rate Adaptation Algorithm: Above, we arrived at the service rate C^* analytically for a set of weights $\phi_i, 1 \leq i \leq m$. The experimental result in Figure 15 illustrates the necessity of adapting the ADWISER service rate as the number of STAs associated at each rate changes over time. There is another important reason why the ADWISER service rate needs to be adaptively learnt, and cannot simply be computed from Equation (1). Note that the rates C_i may be achieved only if there are no wireless channel losses. When there are losses, due to the SINR for connection i not being the best possible for the PHY rate at which the STA is associated, then more packets will be sent on the medium (due to MAC level retransmissions) than are accounted for by introducing virtual bits in the ADWISER queues. These losses will result in a smaller value of C_i , which will show up if the ADWISER

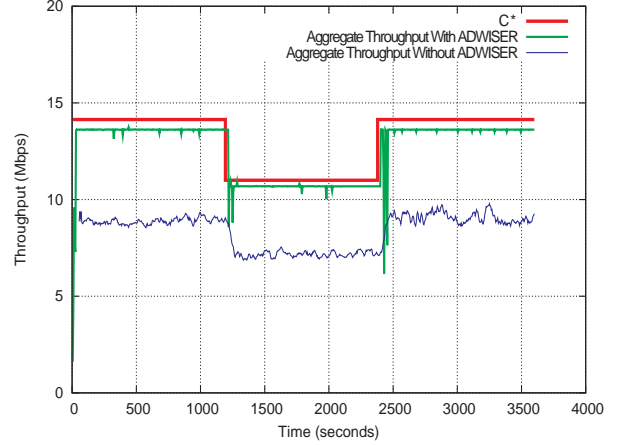


Fig. 16. ADWISER service rate adaptation: aggregate download throughput with and without ADWISER when the number of STAs associated at each rate varies over time; results from the hybrid test-bed.

service rate is adapted based on on-line measurements.

Based on the insights gained from Figure 15 we have implemented the following simple adaptation algorithm. The service rate, C , of the A-WLAN scheduler in ADWISER is iteratively obtained as follows.

- 1) Initialize C to a value less than C^* .
- 2) Carry out the following update after each measurement interval (taken to be 1 second in our implementation)
 - a) If the measured TCP throughput over the WLAN is equal to C , then

$$C \leftarrow \min\{C^*, C + \Delta\},$$

where Δ is a tunable parameter. We found 0.1 Mbps to be good value for Δ .

- b) If the measured TCP throughput is less than C , then

$$C \leftarrow C - \Delta.$$

- 3) In parallel with the above steps, the AP continues to be polled in order to obtain the most current information about STA association. If there is any change, then C^* is recomputed as above.

To illustrate the performance of our approach, we provide a sample result from the hybrid test-bed; see Figure 16. The setup is the same as that for the “open-loop” experiment conducted for Figure 15. Now, the number of STAs in each rate class is changed at an interval of 20 minutes in the following order:

- 3 STAs in each rate class for first 20 minutes,
- 1, 2, and 3 STAs in 54 Mbps, 24 Mbps and 6 Mbps rate classes, respectively, for the next 20 minutes, and
- 4 STAs in each rate class for the last 20 minutes

In Figure 16, we show the values of C^* for each interval, and the service rate, C , to which ADWISER adapts. We also show the aggregate throughput without ADWISER; this is the bottom plot in the figure. We see that, with ADWISER, the aggregate throughput tracks (but falls a little short of) the computed value of C^* . When C^* changes, the adaptation takes 50 to 60 seconds.

2) *A-WAN and A-WLAN Scheduler Coordination*: We have assumed that the Internet access link is an expensive resource and therefore would like it to be fully utilized whenever possible. Thus, WAN-WLAN transfers are allowed to obtain whatever aggregate throughput they can get. This aggregate throughput is measured and subtracted from C , the service rate of the A-WLAN scheduler, which is obtained as explained in Section III-C1c. The LAN-WLAN transfers are then allowed to share the residual rate. The importance of setting C correctly is, thus, further underlined. For, if C is set too high then the queueing will shift from ADWISER back to the AP, resulting in the phenomenon observed in Section II-D. On the other hand, if C is set too low then the WLAN medium will be underutilized.

3) *The A-WLAN Packet Scheduler*: ADWISER maintains separate queues for TCP DATA packets for download connections and TCP ACK packets for upload connections. The ADWISER packet scheduler releases these packets into the WLAN. A TCP ACK packet transmitted by the AP results in a certain number of uplink DATA packets towards the AP. One issue we need to address when trying to use an existing fair queueing scheduler, meant for a wired full-duplex link, is the half-duplex nature of the shared wireless link. The A-WLAN packet scheduler therefore has to release DATA and ACK packets in such a way that it accounts for the time required on the medium for the packets that will be triggered by the reception of these packets by the STAs.

When an ACK packet corresponding to the uplink traffic is served, in the TCP steady state, this results in two uplink DATA packets at the STA because of TCP's delayed ACK mechanism. The scheduler replaces each packet queued in the ADWISER buffer with a number of "virtual bits" corresponding to the number of bits that will actually be carried by the WLAN MAC and PHY as a result of the packet being released into the WLAN. Thus, for each ACK, the scheduler assumes a number of virtual bits equal to an ACK and two DATA packets. For a DATA packet the scheduler assumes a number of virtual bits equal to a DATA packet and half an ACK (to roughly account for delayed ACKs). These are, of course, approximations; in practice, the ratio of ACKs and DATA packets can vary depending on the TCP state. In practice, we have found our simple approach to work quite well.

Packet scheduling is performed using the Start Time Fair Queuing (STFQ) [31] scheduling policy by suitably tagging start and finish numbers to each packet, and scheduling the packet transmissions appropriately. Let us consider an arrival of a packet, $k + 1$, of virtual length $l_{k+1}^{(j)}$ (see above) into a queue j , at arrival instant $a_{k+1}^{(j)}$. Let $F_k^{(j)}$ denote the finish number of packet k in queue j . $V(t)$ denotes the (global) virtual time at time t . The start number $S_{k+1}^{(j)}$ is computed as

$$S_{k+1}^{(j)} = \max\{F_k^{(j)}, V(a_{k+1}^{(j)})\} \quad (4)$$

Then

$$F_{k+1}^{(j)} = S_{k+1}^{(j)} + \frac{l_{k+1}^{(j)}}{\phi_j} \quad (5)$$

where $F_0^{(j)} = 0$. In STFQ, instead of computing the virtual time from a simulation of the corresponding GPS system,

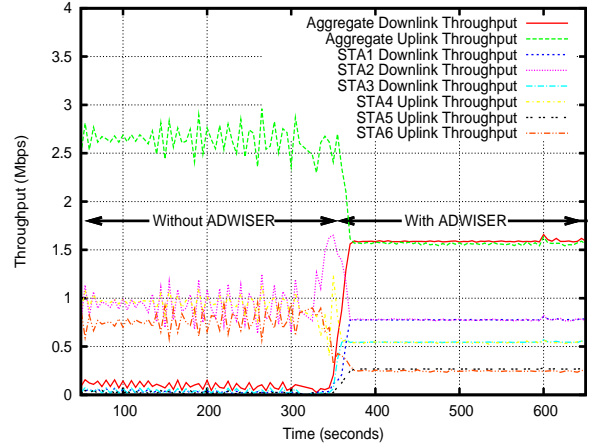


Fig. 17. Hybrid test-bed: Uplink and downlink throughputs vs. time for 6 STAs associated at 11, 5.5 and 2 Mbps PHY rates, with and without ADWISER.

the following approximations are made. The virtual time is initialized to 0, and increases in jumps as follows. When a packet arrives (say, at t), and there is a packet in service, then STFQ approximates $V(t)$ as the start time of the packet in service. If the packet arrives, at time t to an idle system, then the value of $V(t)$ is taken to be the finish time of the last packet in the previous busy period. See [31] for details.

Packets are released into the WLAN in the order of their *start numbers*. This is because the actual server is the wireless medium itself. After releasing a packet into the WLAN from queue i , the ADWISER scheduler allows time for the virtual bits corresponding to the packet to be served at the current ADWISER service rate C .

IV. EXPERIMENTAL RESULTS

A. Results from the Hybrid Test-Bed

Recall Figure 3 which showed a schematic of the hybrid test-bed. The wireless PHY and MAC are simulated in Qualnet. The WAN between the web servers and the ADWISER is emulated by a computer running the Linux *netem* link emulator, in which a propagation delay of 300 ms is configured. The AP buffer is set to the default size of 100 packets. The wireless STAs and the web server run on Linux. We have used the default Linux kernel settings. Thus, TCP window scaling is enabled, and delayed acknowledgment is on. Bulk file transfers (upload and downloads) are started between the STAs and the web server. ADWISER is configured to provide equal aggregate throughputs to uploads and downloads. In the experiments described below, initially ADWISER is bypassed (cases labeled "Without ADWISER"), and later ADWISER is introduced (cases labeled "With ADWISER").

1) *STAs associated at multiple PHY rates*: We consider 6 wireless STAs with half of them downloading files from the web server and the other half uploading files. The 6 wireless STAs are grouped into three sets of 2 STAs, each set associated at 11 Mbps, 5.5 Mbps and 2 Mbps, respectively. The plot in Figure 17 shows the throughputs. During the initial 300 seconds, ADWISER has been bypassed. We notice that the download transfers obtain very small throughputs (the cluster

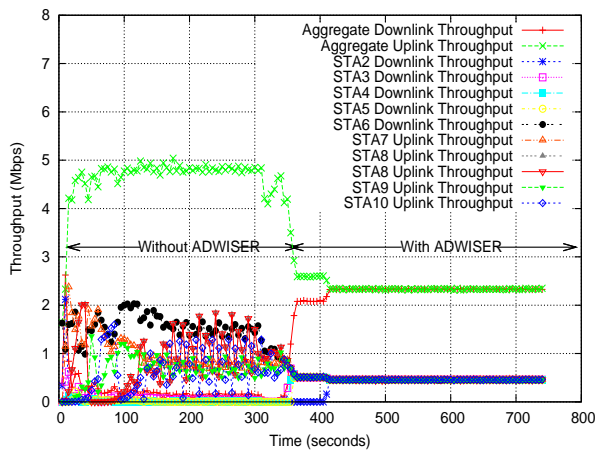


Fig. 18. Hybrid test-bed: Throughputs of 10 STAs, 5 of which are performing uploads and 5 performing downloads. The STAs are associated at 11 Mbps with an AP. Results with and without ADWISER are shown.

of 4 plots near the bottom, including the solid plot for aggregate throughput). Recall the discussion in Section II-B. The upload transfers obtain equal throughputs of 0.8–0.9 Mbps (the three jagged plots in the middle), with an aggregate throughput of about 2.6 Mbps (top-most plot). After ADWISER is introduced, the upload and download throughputs for each PHY rate become equal; the aggregate upload and download throughputs become about 1.578 Mbps each, totaling to 3.156 Mbps, which is just less than $C^* = 3.26$ Mbps obtained from the parameters of this experiment. Further, proportional service differentiation leads to the 11 Mbps, 5.5 Mbps, and 2 Mbps transfers obtaining aggregate throughputs of, respectively, 2×0.778 Mbps = 1.556 Mbps, 2×0.545 Mbps = 1.09 Mbps, and 2×0.255 Mbps = .51 Mbps, which are roughly in the ratio of the $C_1 = 4.89$ Mbps, $C_2 = 3.33$ Mbps, and $C_3 = 1.58$, the ideal contention free TCP throughputs at 11, 5.5 and 2 Mbps. Thus, ADWISER provides a higher aggregate throughput for the network, while providing upload-download fairness, and proportional throughput differentiation across the rate classes.

Relation to the literature: Time-fairness was achieved in [13] by a Leaky Bucket based algorithm, whereas the proposal in [14] was to control a station’s channel occupancy time by a TCP window size adjustment algorithm, using channel occupancy measurements at the MAC layer. Both require modifications to the AP software. In contrast, our solution is completely transparent.

2) *STAs associated at a single PHY rate:* In the setup shown in Figure 3, 10 STAs are now associated with the AP at the 11 Mbps PHY rate, keeping the rest of the setup the same. The following experiments are concerned with fairness between upload and download TCP transfers, and between just upload TCP transfers.

Uploads and Downloads: In Figure 18, we have plotted the throughputs of 10 TCP transfers one from each of 10 STAs, half of which perform uploads and the other half perform downloads, of large files to and from the web server. We observe from the “Without ADWISER” segment that the download throughputs are very small; these are the plots near the bottom, before 300 sec. On the other hand the aggregate

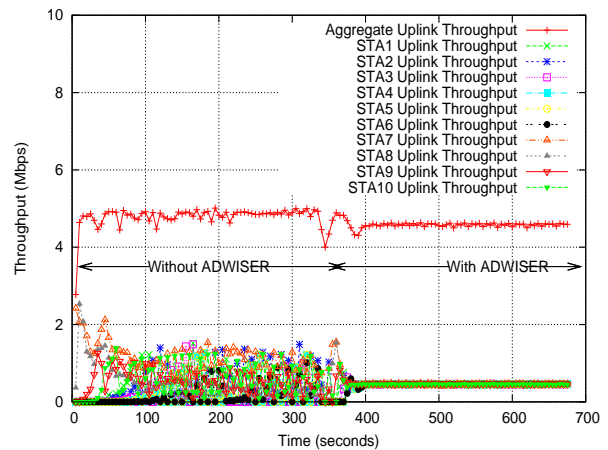


Fig. 19. Hybrid test-bed: Upload throughputs obtained by 10 STAs associated at 11 Mbps with and without ADWISER. The uppermost plot (+ points) shows the aggregate throughput.

upload throughput is about 4.8 Mbps. ADWISER is introduced at 300 sec, and after a transient period, the aggregate upload and download throughputs become equal (about 2.4 Mbps each), and so do the individual throughputs of the 10 file transfers (about 0.48 Mbps each). The transient period can be explained from the observations we had made in Section II-B. With ADWISER bypassed, 1000s of ACKs of the upload connections accumulate in the AP buffer. It takes time to drain out these ACKs. Further, ADWISER does not take control of the queues until the service rate C adapts. Of course, such a large transient is not expected to occur during normal operation with ADWISER in place.

Relation to the literature: The key idea in [15] was to prevent uplink senders from sending too much data, by presenting them with an artificially reduced receiver advertised window. The approach in [17] was to store data and ACK packets in two different queues, and serving the latter with lower probability. Yet another proposal ([19]) was to trigger TCP congestion control actions by deliberately dropping uplink packets. “Virtual Flow Queueing” in [20] is based on adapting the well-known Weighted Fair Queueing (WFQ) policy to the wireless context. Packet lengths are inflated by adding “virtual bits” that account for MAC/PHY and interframe space overheads, as well as backoff intervals. The high-level approach outlined in [12] is also similar to ours.

Nearly all the proposals mentioned above are meant to be implemented on APs. Our transparent solution does not require per-connection TCP window manipulation, or TCP ACK policing using Leaky Buckets, and, unlike [20], uses STFQ scheduling to handle a link with varying service rate.

Uploads only: There are 10 upload file transfers from 10 STAs. Figure 19 shows the results from the hybrid test-bed. Before ADWISER is introduced, there is considerable unfairness between the upload throughputs. Recall the discussion in Section II-C. On the other hand, with ADWISER, the upload throughputs become exactly equal.

Relation to the literature: To address upload-upload unfairness, [22] uses a TCP ACK control mechanism at the AP which regulates the number of ACKs sent out according to the estimated average bandwidth each flow gets. The approach in

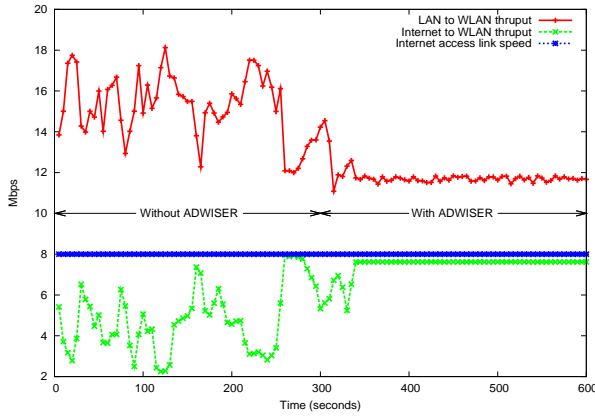


Fig. 20. Physical test-bed: Aggregate TCP download throughputs for WAN-WLAN and LAN-WLAN transfers with and without ADWISER. The Internet access link speed was 8 Mbps.

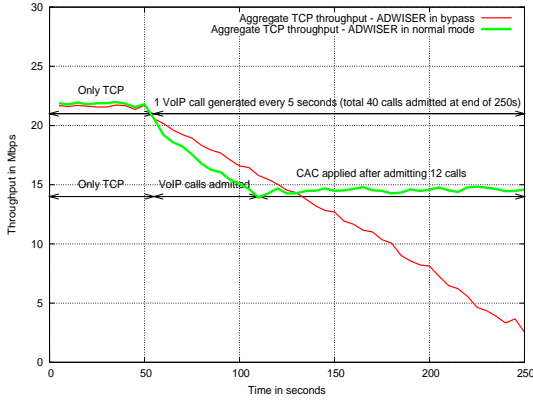


Fig. 21. Hybrid test-Bed: TCP throughput in the presence of VoIP calls with and without ADWISER.

[19] is to regulate TCP traffic by inducing TCP congestion through packet drops. Malone et al. [21] aim to solve the problem by invoking IEEE 802.11e service differentiation to prioritise TCP ACKs in the AP. We use a queueing based approach for controlled release of TCP acks.

3) *WAN-WLAN and LAN-WLAN downloads*: An 8 Mbps Internet link between the web servers and the gateway router, with a propagation delay of 300ms, is emulated using the Linux *netem* utility. We consider a group of two STAs downloading bulk files from the Internet web server and another group of two STAs downloading bulk files from the LAN server. The plot in the Figure 20 shows the bulk file throughputs “Without ADWISER” and “With ADWISER”. We notice that without ADWISER, the Internet access link is poorly utilised in the presence of TCP transfers from the local LAN servers. In contrast, ADWISER eliminates the unfair wireless channel access times between the competing flows from LAN traffic and the Internet traffic destined to the WLAN. The Internet transfer rates are now extremely stable, with higher aggregate throughputs that are close to 8 Mbps. The LAN-WLAN transfers see stable but reduced throughputs; their throughput is reduced since the WAN-WLAN transfers need to obtain more medium time on the WLAN.

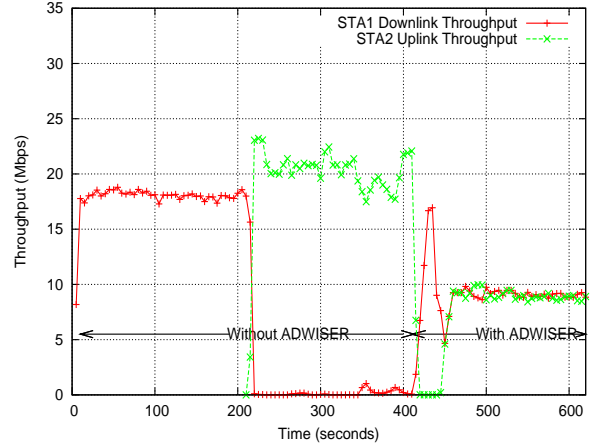


Fig. 22. Physical test-bed: Upload and download throughputs for 2 STAs associated at 54 Mbps (IEEE 802.11g) with and without ADWISER.

Relation to the literature: To the best of our knowledge, existing solutions do not address this problem.

4) *Connection Admission Control for VoIP Calls*: We have implemented VoIP Connection Admission Control (CAC) in ADWISER to provide aggregate throughput guarantees to TCP transfers in the presence of VoIP traffic.

Using the hybrid test-bed in Figure 3, we reported the effect of VoIP calls on TCP download throughput in Section II-E. To demonstrate the VoIP CAC feature in ADWISER, we had a single STA perform TCP bulk download from the LAN server, and then we started generating G.711 VoIP calls at the rate of 1 call every 5 sec. In Figure 21, up to 50s, there is only a bulk TCP download from a server on the LAN to the STA. After 50s, we started generating VoIP calls to other STAs. It is observed that, without ADWISER, all the VoIP calls are admitted and the TCP throughput drops to as low as 3Mbps as the calls are added. When ADWISER is in place, with a CAC limit of 12 calls, it is observed that after admitting 12 calls (at about 110 sec), no more VoIP calls are admitted. As a result, the TCP download throughput drops only to about 14Mbps and stays there.

Relation to the literature: CAC for VoIP traffic is not new, but exercising it for providing QoS to TCP in the IEEE 802.11e context is a novel feature in ADWISER. We note here that, in practice, the CAC limit for VoIP calls will depend on their coding rates. Thus, ADWISER will need to have a matrix showing the combination of numbers of calls of each coding rate to be admitted. As an example, see [6] for an analytical approach for obtaining CAC limits for multiclass VoIP calls for IEEE 802.11 DCF.

B. Results from a Test-Bed with a Real AP

In these experiments, we have an IEEE 802.11g Cisco Aironet AP with which two Linux based laptops (STA1 and STA2) are associated, each at 54 Mbps. The remaining experimental test-bed and parameters are the same as in Section IV-A. Delayed ACKs were enabled in the TCP receivers; RTS/CTS was used to send data packets, whereas Basic Access was used for TCP ACKs. With this, the maximum possible throughput on the wireless medium (with 54 Mbps PHY rate)

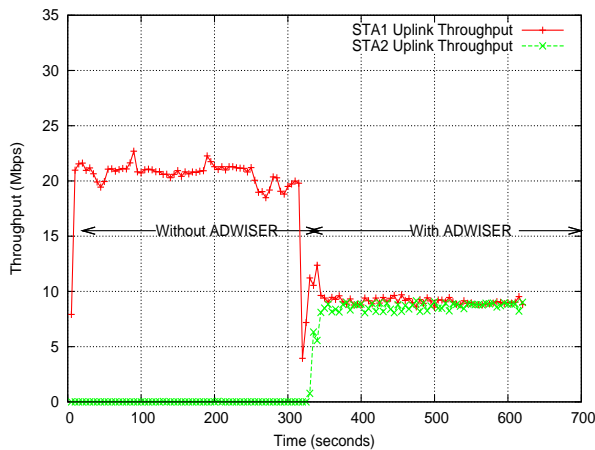


Fig. 23. Physical test-bed: Upload throughputs for 2 STAs associated at 54 Mbps (IEEE 802.11g) with and without ADWISER.

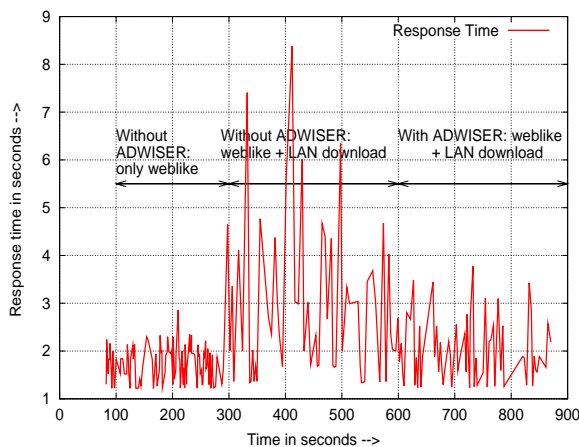


Fig. 24. Physical test-bed: Response times of web-like downloads in the presence of bulk download from the LAN, with and without ADWISER.

can be calculated to be 23.14 Mbps [9]. The actual throughputs achieved will, of course, be less.

Upload and Download: See Figure 22. A large file download from the wireline web server was initiated from STA1, and was allowed to run for 200 seconds; we observe a throughput of about 18.56 Mbps. Then an upload was started from STA2; the upload throughput was 21.2 Mbps, whereas the download throughput fell to 0.34 Mbps. The upload transfer obtains a slightly higher throughput than a download alone because some of the TCP ACKs are lost in the AP, and hence fewer bits are carried on the medium per data packet. Once ADWISER is turned on at 400 seconds, the throughputs equalise, giving a total throughput of about 19 Mbps. Recall, from Section III-C1c, that we adapt the ADWISER service rate to a value a little less than C^* , which is the reason for the lower aggregate throughput than what might ideally be expected.

Two Uploads: This experiment demonstrates unfairness among uploads, and ADWISER’s ability to enforce fairness. In Figure 23, in the “Without ADWISER” period, only the upload from STA1 obtains throughput (about 21.5 Mbps), whereas the one from STA2 is starved completely. This complete starvation of an STA, when performing TCP uploads, was also reported

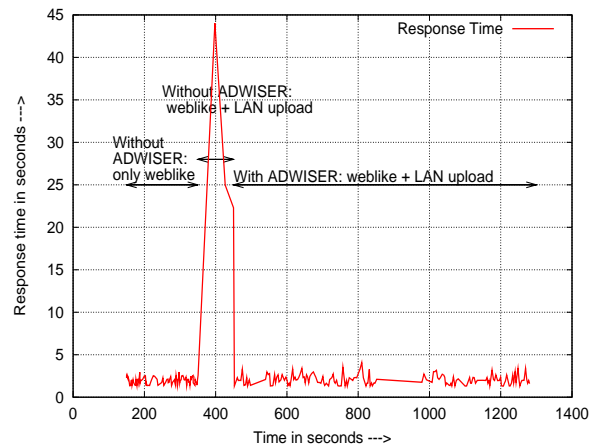


Fig. 25. Physical test-bed: Response times of web-like downloads in the presence of bulk upload to the LAN, with and without ADWISER.

by Malone et al. [21]. Once ADWISER is introduced, both the uploads obtain equal throughputs, with the aggregate being 19 Mbps.

Web-like traffic: The following two experiments were carried out with the set up shown in Figure 4, with ADWISER inserted. The results are shown in Figures 24 and 25

In Figure 24, the 300s to 600s interval shows the response times of web-like downloads from the Internet server in the presence of a bulk *download* from the LAN server, without ADWISER. As can be seen, the average value of response time is 2.60s for this interval. After introducing ADWISER at 600s, we see that the response time improves, achieving a smaller average of about 2sec, and also lower variability.

In Figure 25, the 350s to 450s interval shows the response time of web-like download from Internet server in the presence of a bulk *upload* to the LAN server, without ADWISER. The average value of response time is 30.43s. After introducing ADWISER at 450s, we see that the response time improves with an average value of 2.05s, which is a very significant improvement.

C. Policy Based Service Differentiation

Consider two stations, STA1 and STA2, associated at 11 Mbps, each downloading a large file. STA1 is a privileged user and requires some guaranteed fraction of the total throughput. ADWISER provides for such configurable policies. In Table I we have presented two cases. In the first row ADWISER is configured to give STA1 five times the throughput of STA2. In the second row STA1 gets nine times the throughput of STA2. In ADWISER a queue is created for each STA and appropriate weights (recall the ϕ_i values from Section III-B) are assigned. We see from Table I that ADWISER is able to provide the required service differentiation. The total throughput in each case is the same: 4.655 Mbps.

Figure 26 shows the results of policy-based service differentiation on a test-bed with a real AP. We have classified two STAs into the “gold” category, and two into the “bronze” category. All are associated at 54 Mbps. All the four stations download bulk data from the Internet server. We use Linux

Configured Ratio	STA1 Thpt Mbps	STA2 Thpt Mbps	Total Thpt Mbps
5:1	3.879	0.776	4.655
9:1	4.190	0.465	4.655

TABLE I
SERVICE DIFFERENTIATION USING ADWISER.

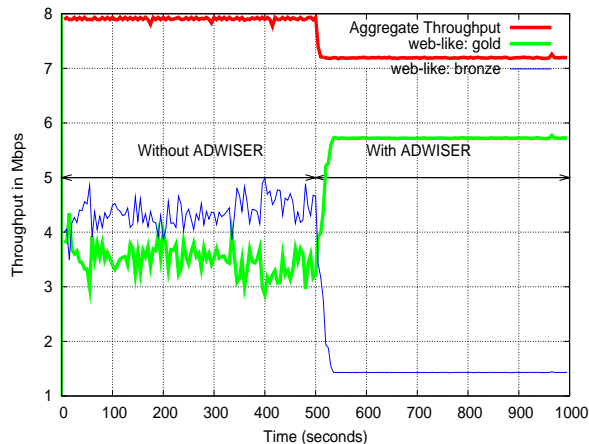


Fig. 26. Physical test-bed: Enforcement of a throughput differentiation policy for Internet downloads using ADWISER.

netem to set a RTT of 300ms for the gold stations, and a slightly lower RTT of 200ms for the bronze stations. In Figure 26, in the interval 0s to 500s, when ADWISER is bypassed, it is observed that the bronze STAs obtain a slightly higher throughput due to the lower RTT. But from 500s to 1000s, we have introduced ADWISER, with a policy 4:1. The virtual server in the WAN scheduler in ADWISER has a service rate that is 90% of the Internet access link speed (recall Section III-A); hence, the total throughput over the WAN is limited to 7.2 Mbps when ADWISER is inserted. We observe that the total throughput of 7.2Mbps is shared in the ratio of 4:1 between the gold STAs and the bronze STAs.

Relation to the literature: ADWISER supports policy based service differentiation. This is a novel aspect that adds an important capability to traffic management in WLANs, and that, to our knowledge, is not available in the open literature.

V. CONCLUSION

ADWISER, an enhanced version of WM ([1], [2]), is a centralized WLAN controller that can manage certain QoS issues in an infrastructure WLAN, without any modification to APs or the clients. These systems can be viewed as controllers in the “Split-MAC” architecture ([10], [11]). We have reported our experiences in implementing ADWISER on a Linux platform, and the performance results we have obtained. All the results were obtained from test-beds in which the actual physical ADWISER was in place, and was carrying TCP traffic between a server and clients. We have reported that several TCP QoS problems in WLANs can be effectively solved by ADWISER. ADWISER can also be used to create any desired service differentiation between STAs.

Another experience we must report was that, in the course of evolving the ADWISER idea and experimenting with it, the insights gained from analytical models (such as [5] and [6], to list only a couple of many such papers in our bibliography)

were invaluable in (i) understanding, explaining and predicting the experimental results, and in (ii) guiding the design of the algorithmic techniques.

In ongoing work, we are extending ADWISER to manage multiple APs, which is possible only because of the centralized nature of our solution, implemented outside the APs.

REFERENCES

- [1] M. Hegde, S. V. R. Anand, A. Kumar, and J. Kuri, “WLAN Manager (WM): a Device for Performance Management of a WLAN,” *International Journal of Network Management*, vol. 17, pp. 155–170, January 2007.
- [2] M. Hegde, P. Kumar, K. R. Vasudev, S. V. R. Anand, A. Kumar, and J. Kuri, “Experiences with WM: A Centralised Scheduling Approach for Performance Management of IEEE 802.11 Wireless LANs,” in *Proceedings of WISARD (Workshop on Wireless Systems: Advanced Research and Development)*, Bangalore, January 2010.
- [3] “Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, ANSI/IEEE Std 802.11, 1999 Edition,” 1999, IEEE Std 802.11.
- [4] G. Berger-Sabbatel, F. Rousseau, M. Heusse, and A. Duda, “Performance Anomaly of 802.11b,” in *Proceedings of IEEE Infocom 2003*, IEEE, 2003.
- [5] R. Bruno, M. Conti, and E. Gregori, “An Accurate Closed-form Formula for the Throughput of Long-lived TCP Connections in IEEE 802.11 WLANs,” *Computer Networks*, vol. 52, pp. 199–212, 2008.
- [6] G. Kuriakose, S. Harsha, A. Kumar, and V. Sharma, “Analytical Models for Capacity Estimation of IEEE 802.11 WLANs using DCF for Internet Applications,” *Wireless Networks*, vol. 15, no. 2, pp. 259–277, February 2009.
- [7] Q. Wu, M. Gong, and C. Williamson, “TCP Fairness Issues in IEEE 802.11 Wireless LANs,” *Computer Communications*, vol. 31, no. 10, June 2008.
- [8] O. Bhardwaj, G. Sharma, M. Panda, and A. Kumar, “Modeling Finite Buffer Effects on TCP Traffic over an IEEE 802.11 Infrastructure WLAN,” *Computer Networks (an Elsevier journal)*, vol. 53, pp. 2855–2869, 2009.
- [9] S. Harsha, A. Kumar, and V. Sharma, “An Analytical Model for the Performance Evaluation of Multimedia Applications EDCA in an IEEE 802.11e WLAN,” *Wireless Networks*, vol. 16, pp. 367–385, 2010.
- [10] L. Yang, P. Zerfos, and E. Sadot, “Architecture Taxonomy for Control and Provisioning of Wireless Access Points (CAPWAP),” June 2005, RFC 4118.
- [11] P. Calhoun, M. Montemurro, and D. Stanley, “Control and Provisioning of Wireless Access Points (CAPWAP) Protocol Specification,” March 2009, RFC 5415.
- [12] K. Cai, R. Lotun, M. Blackstock, J. Wang, M. Feeley, and C. Krasic, “A Wired Router Can Eliminate 802.11 Unfairness, But It’s Hard,” in *ACM HotMobile 2008: The Ninth Workshop on Mobile Computing Systems and Applications*, 2008.
- [13] G. Tan and J. Guttag, “Time-based Fairness Improves Performance in Multi-rate WLANs,” in *Proc. USENIX Annual Tech. Conf., Berkeley, CA, USA*, May 2004, pp. 269–282.
- [14] A. J. Kenichi Kashibuchi and N. Kato, “Channel Occupancy Time Based TCP Rate Control for Improving Fairness in IEEE 802.11 DCF,” *IEEE Transactions ON Vehicular Technology*, vol. 59, no. 6, July 2010.
- [15] S. Pilosof, R. Ramjee, D. Raz, Y. Shavitt, and P. Sinha, “Understanding TCP Fairness over Wireless LANs,” in *Proceedings of the IEEE INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies. IEEE*, vol. 2, no. ISSN: 0743-166X. IEEE, March 2003, pp. 863–872.
- [16] A. Detti, E. Graziosi, V. Minichiello, S. Salsano, and V. Sangregorio, “TCP Fairness Issues in IEEE 802.11 based Access Networks.” [Online]. Available: http://www.eln.uniroma2.it/Stefano_Salsano/papers/salsano-tcp-fair-wlan.pdf
- [17] J. Ha and C.-H. Choi, “TCP Fairness for Uplink and Downlink Flows in WLANs,” in *ACM HotMobile 2008: The Ninth Workshop on Mobile Computing Systems and Applications*, 2006.
- [18] F. Vacirca and F. Cuomo, “Experimental Results on the Support of TCP over 802.11b: An Insight into Fairness Issues,” in *Proceedings of Wireless On-Demand Network Systems and Services*, January 2006.
- [19] N. Blefari-Melazzi, A. Detti, I. Habib, A. Ordine, and S. Salsano, “TCP Fairness Issues in IEEE 802.11 Networks: Problem Analysis and Solutions Based on Rate Control,” *IEEE Transactions on Wireless Communications*, vol. 6, no. 4, pp. 1346–1355, 2007.

- [20] F. R. Paul Starzetz, Martin Heusse and A. Duda, "Virtual Flow Queueing for Improving TCP Performance over IEEE 802.11 WLANs," in *WCNC*, 2008.
- [21] D. Malone, D. J. Leith, A. Aggarwal, and I. Dangerfield, "Spurious TCP timeouts in 802.11 networks," in *Proc. 6th Intl. Symposium on Modeling and Optimization (WiOpt), Berlin, Germany*, April 2008.
- [22] F. Keceli, I. Inan, and E. Ayanoglu, "TCP ACK Congestion Control and Filtering for Fairness Provision in the Uplink of IEEE 802.11 Infrastructure Basic Service Set," in *Proceedings of the ICC*, 2007, pp. 4512–4517.
- [23] S.-H. Yoo, J.-H. Choi, J.-H. Hwang, and C. Yoo, "Eliminating the Performance Anomaly of 802.11b," *Lecture Notes in Computer Science*, 2005.
- [24] H. Kim, S. Yun, I. Kang, and S. Bahk, "Resolving 802.11 Performance Anomalies through QoS Differentiation," *IEEE Communications Letters*, pp. 655–657, 2005.
- [25] S. W. Kim, B.-S. Kim, and Y. Fang, "Downlink and Uplink Resource Allocation in IEEE 802.11 Wireless LANs," *IEEE Transactions on Vehicular Technology*, vol. 54, no. ISSN: 0018-9545, pp. 320–327, January 2005.
- [26] D. J. Leith and P. Clifford, "Using the 802.11e EDCF to Achieve TCP Upload Fairness over WLAN Links," in *Proceedings of the Third International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks*. IEEE Computer Society, 2005, pp. 109–118.
- [27] "Scalable Networks," 2010, www.scalable-networks.com.
- [28] D. J. Leith, P. Clifford, D. Malone, and A. Ng, "TCP Fairness in 802.11 WLANs," *IEEE Communications Letters*, vol. 9, no. 11, pp. 964–966, November 2005.
- [29] D. J. Leith and P. Clifford, "Modelling TCP Dynamics in Wireless Networks," in *Proceedings of the International Conference on Wireless Networks, Communications and Mobile Computing*, vol. 2, June 2005, pp. 906–911.
- [30] <http://linuxwireless.org/en/developers/Documentation/mac80211/RateControl>
- [31] P. Goyal, H. M. Vin, and H. Cheng, "Start-Time Fair Queueing: A Scheduling Algorithm for Integrated Services Packet Switching Networks," *IEEE/ACM Transactions on Networking*, vol. 5, no. 5, pp. 690–704, May 1997.



Malati Hegde received her PhD from the Indian Institute of Technology, Kanpur, in 1978. She is working as a Principal Scientific Officer in the ECE Department, Indian Institute of Science (IISc), Bangalore. Her areas of interest are network management and performance management in wired, wireless LANs and wireless sensor networks. She is actively involved the Education and Research Network (ERNET) project, a nation-wide academic and research network, and shares the responsibility of the ERNET NOC at IISc, Bangalore, in network

planning and providing Internet services.



Pavan Kumar received the B.Tech. degree in Computer Science from National Institute of Technology Rourkela, India, in 2005. He joined the ECE Department, IISc, Bangalore as a Project Assistant in 2007. Since 2010, he has been with Brocade Communications Systems, Bangalore, developing network protocol software for Ethernet switches.



K. R. Vasudev received the B.E. in Electronics and Communications from the Visvesvaraya Technological University, Belgaum, Karnataka. He was with the ECE Department, IISc, Bangalore, from 2008 to 2009 as a project assistant. He then obtained a Masters degree in Digital Communications at University of Kiel, Germany. His Masters thesis was on the topic of multi-channel speech quality improvement for automotive systems. Currently he is pursuing his PhD in the field of Speech and Audio signal processing at the Digital Signal Processing and System Theory (DSS) chair at University of Kiel, Germany.



N. N. Sowmya was born on 30th November, 1986 in Kerala, India. She completed her B.E. (Hons) EEE from the Birla Institute of Technology and Science, Pilani - Goa Campus in 2008. She worked with Honeywell Technology Solutions, Bangalore, as a Software Engineer for 2 years. She was then a project assistant in the ECE Department, Indian Institute of Science, for one year, during which period she worked on the WLAN Manager/ADWISER project.



S. V. R. Anand received the B.E. from the ECE Department, Indian Institute of Science in 1987. Working as a technical consultant, he is involved in the implementation aspects of various sponsored R&D projects undertaken by the department. His areas of interests include network system software development, and operating systems. He is also associated with the ERNET Network Operations Centre at the Indian Institute of Science.



Anurag Kumar obtained his B.Tech. degree in Electrical Engineering from the Indian Institute of Technology at Kanpur. He obtained the PhD degree from Cornell University. He was then with Bell Laboratories, Holmdel, N.J., for over 6 years. Since 1988 he has been with the Indian Institute of Science (IISc), Bangalore, in the Dept. of Electrical Communication Engineering, where he is now a Professor, and also Chair of the Electrical Sciences Division. From 1988 to 2003 he was the Coordinator at IISc of the Education and Research Network Project (ERNET), India's first wide-area packet switching network. His area of research is Communication Networking; specifically, modeling, analysis, control and optimisation problems arising in communication networks and distributed systems. He has been elected Fellow of the IEEE, the Indian National Science Academy (INSA), the Indian National Academy of Engineering (INAE), and the Indian Academy of Sciences (IASc). He was an associate editor of *IEEE Transactions on Networking*, and an area editor of *IEEE Communications Surveys and Tutorials*. He is a coauthor of the graduate text books "Communication Networking: An Analytical Approach," and "Wireless Networking," by Kumar, Majunath and Kuri, published by Morgan-Kaufman/Elsevier.

Project (ERNET), India's first wide-area packet switching network. His area of research is Communication Networking; specifically, modeling, analysis, control and optimisation problems arising in communication networks and distributed systems. He has been elected Fellow of the IEEE, the Indian National Science Academy (INSA), the Indian National Academy of Engineering (INAE), and the Indian Academy of Sciences (IASc). He was an associate editor of *IEEE Transactions on Networking*, and an area editor of *IEEE Communications Surveys and Tutorials*. He is a coauthor of the graduate text books "Communication Networking: An Analytical Approach," and "Wireless Networking," by Kumar, Majunath and Kuri, published by Morgan-Kaufman/Elsevier.



Joy Kuri has a B.E. in Electronics and Telecommunication Engineering from Jadavpur University, Kolkata. After a year in the industry, he joined the M.E. program in the Department of Electrical Communication Engineering at the Indian Institute of Science. He went on to receive a PhD from the same department at IISc in 1995. Subsequently, he spent two years at Ecole Polytechnique, University of Montreal, Canada, and one and a half years in INRS-Telecommunications, University of Quebec, Canada as a Research Associate. Since 1999 he has

been with the Centre for Electronics Design and Technology, Indian Institute of Science, where he is now a Professor. His research and teaching interests are in the areas of modeling, analysis and control of communication networks and stochastic systems.