Computer Networks 52 (2008) 2728–2744

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

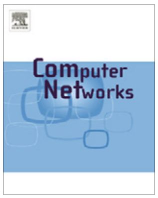

An experimental study of inter-cell interference effects on system performance in unplanned wireless LAN deployments

Mesut Ali Ergin *, Kishore Ramachandran, Marco Gruteser

WINLAB, Electrical and Computer Engineering Department, Rutgers, The State University of New Jersey, 671 Route 1 South, North Brunswick, NJ 08902-3390, USA

A R T I C L E   I N F O

Article history:
Received 4 October 2007
Received in revised form 10 March 2008
Accepted 28 May 2008
Available online 13 June 2008

Responsible Editor: E. Gregori

Keywords:
Unplanned WLAN Deployments
High-density WLANs
Inter-cell interference
Wireless testbeds
Real-world evaluation
Wireless experimentation

ABSTRACT

In this paper, we report on our experimental study of the effects of inter-cell interference on IEEE 802.11 performance. Due to growing use of wireless LANs (WLANs) in residential areas and settings supporting flash crowds, chaotic unplanned deployments are becoming the norm rather than an exception. Environments in which these WLANs are deployed, have many nearby access points and stations on the same channel, either due to lack of coordination or insufficient available channels. Thus, inter-cell interference is common but not well-understood. According to conventional wisdom, the efficiency of an IEEE 802.11 network is determined by the number of active clients. However, we find that with a typical TCP-dominant workload, cumulative system throughput is characterized by the number of actively interfering access points rather than the number of clients. We verify that due to TCP flow control, the number of backlogged stations in such a network equals twice the number of active access points. Thus, a single access point network proves very robust even with over one hundred clients, while multiple interfering access points lead to a significant increase in collisions that reduces throughput and affects media traffic. Only two congested interfering cells prevent high quality VoIP calls. Based on these findings, we suggest a practical contention window adaptation technique using information on the number of nearby access points rather than clients. We also point out the need for collision-resilient rate adaptation in such a setting. Together these techniques can largely recover the 50% loss in cumulative throughput in a setting with four strongly interfering access points.

1. Introduction

Penetration of IEEE 802.11 Wireless LANs (WLANs)[1] into everyday life has started with data communications for laptop PCs and is now continuing with the integration of WLAN capabilities into consumer devices such as printers, music players and mobile phones. Recent projections estimate that by 2011, 340 million WiFi enabled phones will be shipping yearly [1]. Such a usage growth likely leads to increased inter-cell interference. Conventional inter-cell interference mitigation requires careful site surveys to prevent neighboring cells from operating on the same channel, as well as transmission power control and other receiver parameter optimizations (such as carrier sensitivity adjustment) to increase spatial reuse. Dynamic forms of these approaches are also proposed to ease their adoption, such as distributed channel selection.

Even with such mitigation techniques, however, several factors make complete elimination of inter-cell interference non-trivial in WLANs. First, many WLANs, especially in residential or small business areas, are deployed in an unplanned and chaotic manner. Lack of coordination between neighbors and the restrictions of many ISP contracts

---

* Corresponding author. Tel.: +1 732 932 6857; fax: +1 732 932 6882.
E-mail addresses: ergin@ieee.org, ergin@winlab.rutgers.edu (M.A. Ergin), kishore@winlab.rutgers.edu (K. Ramachandran), gruteser@winlab.rutgers.edu (M. Gruteser).

[1] In this study, we are interested in the performance of IEEE 802.11a/b/g Wireless LANs. Therefore the term ‘WLAN’ has been used throughout the text to refer to IEEE 802.11a/b/g Wireless LANs.

1389-1286/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2008.05.013

lead to over-deployment of access points. War-driving studies of these networks have found up to 80 access points in communication range, with 40% of them configured on the industrial, scientific and medical (ISM) radio bands at 2450 MHz, channel number 6 [2]. Second, the increased usage of wireless multimedia including voice-over-WiFi has motivated WLAN system designs where neighboring light-weight access points operate on the same channel to minimize handoff time [3]. Third, the number of available non-overlapping channels in the unlicensed spectrum for IEEE 802.11a/b/g is still low, even with the 2006 US FCC allocation of 11 additional channels in 5 GHz spectrum. Hence, even frequency planned networks that provide blanket coverage over larger campuses cannot fully avoid inter-cell interference, especially with the trend towards high capacity radio switches that require multiple channels per cell [3].

Unfortunately, in these unplanned environments, little is known about the effects of inter-cell interference on IEEE 802.11 system performance. Detailed analytical and simulation models exist for the MAC protocol scalability [4,5] and experimental studies have characterized scalability under TCP and UDP workloads [6] in the single-cell case. However, system scalability in the common *unplanned* multi-cell case remains largely unexplored. Multi-cell networks have been studied through measurement campaigns in real-world campus [7] or conference settings [8,9] and recent measurements in a dense conference deployment have detected performance anomalies [9,10], but the data does not allow a detailed analysis of root causes.

New contributions: This paper presents a systematic analysis of the effect of inter-cell interference in such unplanned, high-density WLANs through detailed experiments and simulations. Our work complements previous real-world measurements through experimentation with over hundred IEEE 802.11 enabled nodes in a repeatable laboratory setting with controlled interference. Thus, it allows in-depth analysis through simulations and repeatable experiments, with precisely known configurations.

We expand our initial findings presented in [11] by conducting elaborate large-scale realistic experiments on the ORBIT testbed [12] as well as evaluating multimedia performance (esp. VoIP) through experiments and simulations. We also suggest a practical solution that can recover significant part of the losses demonstrated in this study. In summary, our contributions with this paper include:

- Analyzing system performance using a realistic TCP dominated workload in unplanned multi-cell WLANs by conducting experiments on the ORBIT testbed as well as QualNet simulator [13]. Results show that a single-cell network remains remarkably robust even with 125+ clients; the collision rate remains low. This extends Choi et al.'s empirical results [6] for 16 clients to a much larger network, with realistic client association patterns, and bursty traffic mixes. We also show that, in an unplanned multi-cell network, however, the collision rate increases significantly.
- Providing novel insights into the behavior of TCP in multi-cell WLANs. Due to TCP flow control, the number of backlogged stations equals twice the number of active access points, meaning that network efficiency is determined by *the number of interfering access points*, not *the number of clients*. In addition, we show that TCP cannot regulate the flows in the IEEE 802.11 network for optimal system operating point (i.e. max. throughput) across different contention window settings.
- Quantifying the effect of inter-cell interference on multimedia traffic and on throughput loss due to inefficient rate adaptation. Even with Wireless Multimedia Extensions, based on the IEEE 802.11e standard [14], VoIP users can still experience substantial performance degradation in unplanned deployments. This deterioration starts to occur even when the number of interfering APs is *relatively small* (three). Video streaming in the network makes the system performance worse for VoIP users.
- Identifying a practical distributed interference mitigation technique: *contention window adaptation based on the number of active access points*, not the number of clients in the network. We also show that an additional 20% gain would be possible with collision-resilient rate adaptation.

Road map: The remainder of this paper is organized as follows. The next section explains our evaluation methodology and the details of the realistic traffic workload we have used. In Section 3, we characterize the effect of node-density (both client and AP density), traffic variability, and client arrival pattern, on system performance. This is followed by Section 4, which details TCP performance in single-cell and unplanned multi-cell WLANs. Section 5 goes on to describe the effects of inter-cell interference on multimedia applications. In Section 6, we examine performance restoration techniques and the gains that could be obtained by accurate implementations. We discuss related work in Section 7 and conclude the paper in Section 8.

2. Experiment setup and traffic models

We leverage the publicly accessible ORBIT testbed [12] to carry out systematic and controlled experiments. In our evaluations, we use the network topology shown in Fig. 1. The main components of this integrated wired/wireless IP network are wired nodes hosting application servers, wireless access points (AP) and stations (STA). We focus on application behavior in the wireless access segment, which consists of multiple, interfering basic service sets (BSS) on the same channel, in close proximity. All nodes remain in communication range emulating future very high-density deployments. Evaluating the effect of hidden nodes is beyond the scope of this paper. Pairwise SNR profile of all nodes in our testbed allows communication using IEEE 802.11a at 54 Mbits/s rate with less that 1% packet error rate.

2.1. Experimental and simulation setup

Table 1 provides a list of the parameters we used on the ORBIT testbed. To ensure that our results are representative

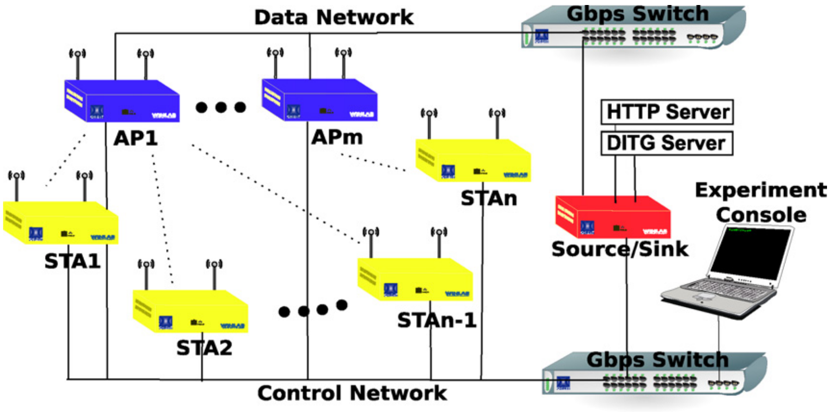

Fig. 1. ORBIT testbed experiment setup with configurable number of access points and up to 400 nodes.

Table 1
Attribute summary for ORBIT experiments

Attribute	Value
Radio nodes	1 GHz VIA C3 Processor, 512 MB RAM, 20 GB HDD
Wireless interfaces	2 X Atheros AR5212 based mini-PCI 802.11a/g
Wireless output power	18 dBm
PHY/MAC/Freq. used	IEEE 802.11a/ Operating at UNII Band Channel 52
PHY link speed (Fixed)	Up to 54 Mbits/s
MAC payload size	1300 bytes
MAC retries	10
O/S used	Linux 2.6.18
Driver software	MadWifi svn.21XX [15]

of real-world behavior, we first carried out calibration tests comparing throughput of the ORBIT machines configured as APs (that use the MadWifi driver [15]) with commercial Cisco (1200 series) and D-Link APs. We did not observe a significant difference (less than 5%). We also configured the APs as bridges and used traffic sniffers (via tcpdump) on both the wired and wireless segments. To allow other researchers to replicate, and build on our experiments in the ORBIT environment, we have published relevant traffic scripts, tools, and frame dump traces through the CRAWDAD archive [16] as *rutgers\ap_density* traceset [17].

For our simulations, we chose QualNet [13] (the commercial successor of GloMoSim) due to its accurate physical layer interference model, which can affect higher layer performance comparisons [18]. In particular, QualNet's SINR calculation taking into account the cumulative interference power from all concurrent senders is very important for measuring the effect of MAC collisions in our high-density simulations. Note that the default ns-2 model may underestimate collisions, since it only keeps track of the strongest interferer, not the sum of all interference signals.

Before running our simulations, we made sure that two bugs in the particular version of the simulator: (i) losing slot synchronization when resuming backoff (already described in [19]), and (ii) improper resetting of CW values for IEEE 802.11e EDCA [14] access queues, were fixed. The parameters that we used in our simulations are outlined in Table 2[2].

Metrics: For performance evaluation, we choose to use application-specific metrics. Thus, for the performance of TCP-based applications, we focus on system throughput and throughput fairness (Jain's fairness index [20]). For multimedia applications such as VoIP, we utilize both quantitative metrics such as application-level packet drop rate, latency and jitter as well as a standard subjective quality metric, namely mean opinion score (MOS).[3]

The rationale behind conducting our experiments in a controlled laboratory setting (such as ORBIT) rather than a real deployed WLAN is as follows. First, it allows detailed instrumentation to understand MAC-level behavior without the use of large numbers of sniffers. Second, experiments are repeatable that is they are not dependent on time-varying shadowing and interference patterns. These allow both easier investigation of root causes and directly comparing alternative solutions. Finally, the high-density placement of 400 nodes allows us to experiment with densities that may be expected in future years rather than focussing solely on today's system performance.

2.2. Validation experiment

Throughout the paper, we will use collision-dominated packet error rate (PER) and system throughput measurements obtained from Bianchi's IEEE 802.11 model that assumes that all stations are backlogged [4], the QualNet simulator [13], and the ORBIT wireless testbed [12], to compare the performance of the wireless networks of interest.[4] To be able to validate such a comparison, we conduct a saturation data transfer experiment on an IEEE 802.11a network of two active nodes using 36 Mbits/s fixed

---

[2] Note that we used a modified MAC retry value of 10 to match those used in our ORBIT experiments.

[3] MOS is a subjective score (ranging from 0 to 5) used to evaluate voice quality as perceived by an average user of the system. Details can be found in ITU G.107 and G.113.

[4] Note that whenever Bianchi's saturation model is used for comparisons, corresponding experiments on simulator and the testbed will incur bulk transfer workloads to satisfy backlogged station requirements of the model.

Table 2
Attribute summary for simulations

Attribute	Value
PHY	5.2 GHz, Two Ray Ground, Tx. power of 18 dBm, Rx. Sensitivity of −78 dBm (at 36 Mbps)
MAC	Basic access (RTS off), Variable CW_{min} (mostly 15), CW_{max} of 1023, 16 µs SIFS, 34 µs DIFS, 9 µs Slot time, 28 bytes MAC header, 14 bytes ACK frame size, 1300 bytes MAC payload, 10 retries
NET	IPv4, IP Queue size of 75,000 bytes
TCP	NewReno with RFC 1323, Max. segment size of 1300 bytes, Send/Rcv buffer size of 110 KB, Delayed ACK disabled

Table 3
Validation experiment results for Bianchi's model, QualNet simulator and the ORBIT testbed

	Avg. PER (%)	Avg. throughput (Mbits/ s)
Bianchi's model	10.17	23.14
QualNet simulator	11.02	22.28
ORBIT testbed	9.86	24.67

rate, sending 1500-byte UDP packets, for a duration of five minutes. Average PER and throughput results obtained from MATLAB simulations of Bianchi's model, the QualNet simulator and the ORBIT testbed for this configuration are reported in Table 3. From this and other similar experiments, for the configurations tested, we have repeatedly observed that results from all three methods are very close, validating their potential comparability.

2.3. Traffic model

Workloads: Since inter-cell interference patterns are affected by end-user workloads, we designed a synthetic office workload in addition to bulk TCP-only workloads. The office workload is based on several hours of sniffer traces obtained in our academic office/lab environment from a single access point serving up to 50 students and faculty. These measurements indicate that 97% of packets use the TCP protocol and about 75% of traffic is generated by web traffic, as illustrated in Fig. 2a. In the figure, all percentages are based on the number of bytes communicated in the WLAN and only application protocols with >2% contribution are individually referred. These measurements are reasonably consistent with, except for a 20% increase in web traffic, with an earlier analysis of SIGCOMM 2001 conference traces covering 4 APs and 195 stations [8], which is also shown in the figure.

Thus, 75% of the synthetic workload consists of bursty web traffic, following the self-similar ON–OFF traffic model described in [21]. We directly emulate the HTML transfer, browser processing, and HTML object retrieval phases using the HTTP 1.1 compliant *GNU wget* page retrieval tool [22] to access a local webserver serving web pages and objects obtained from an academic web server. The model we used for this emulation is illustrated in Fig. 2b. The user's thinking time X (inactive time) between page accesses follows a Pareto distribution:

$$Pr(X > x) = \left(\frac{x}{x_m}\right)^{-k} \qquad (1)$$

with shape parameter $k = 1.5$ and lower bound $x_m = 1$, as suggested in [21]. We concentrate on TCP downlink traffic, since it represents typical access point usage for web browsing. Also, earlier results [23] showed that the direction of TCP traffic is not significant. This is mostly due to the equal frame flow rate requirement of TCP in data segment and ACK directions.

The remaining share of the workload comprises a mix of VoIP traffic (over UDP/IP) using the G.711 codec with H.323 signalling (3% of overall volume), and TCP packet transfers with exponentially distributed interarrival times (21% of the overall volume on average) as background traffic. These flows are emulated through the D-ITG traffic generator v.2.4.3 [24]. In the experiments, each station is assigned a traffic generation profile to satisfy the workload distribution outlined above, and keeps this profile until the end of the experiment.

User arrival pattern: Another factor that might potentially affect system performance is client association dynamics. To measure performance in a more realistic manner, we extracted the user arrival patterns from WLAN traces of the 62nd IETF meeting [9]. In particular, we use the arrival pattern of the users returning from lunch between 12:30 p.m. to 1:00 p.m., as illustrated in Fig. 3. The IETF WLAN comprised over 150 APs and more than 700 users. Note the significant variance in user associations, for example at 43 min into the trace, the number of users on channel 11 quickly increases from approximately 50 to over 250 within a two minute window.

3. Analysis of the system performance

In this section, we study the system performance of a multi-cell wireless network deployment, where the cells interfere with each other. We systematically examine the effect of access point density, traffic variability, user arrival pattern and station density to understand root causes of performance problems. We begin with an experiment that emulates the characteristics of short-term conference deployments [9,10], with multiple APs, a web-dominated traffic workload, and a dynamic station arrival pattern. Emulating such a scenario in a controlled environment will allow us to isolate the effect of these factors on network throughput.

3.1. Multi-cell network with realistic workload

This 11 min experiment comprises four APs and 75 STAs, selected randomly from the 400-node ORBIT main radio grid. Once APs are operational on UNII 5 GHz Band Channel 52, using IEEE 802.11a, stations start to associate with the network by following the dynamic client arrival pattern described in Section 2.3. This arrival pattern is

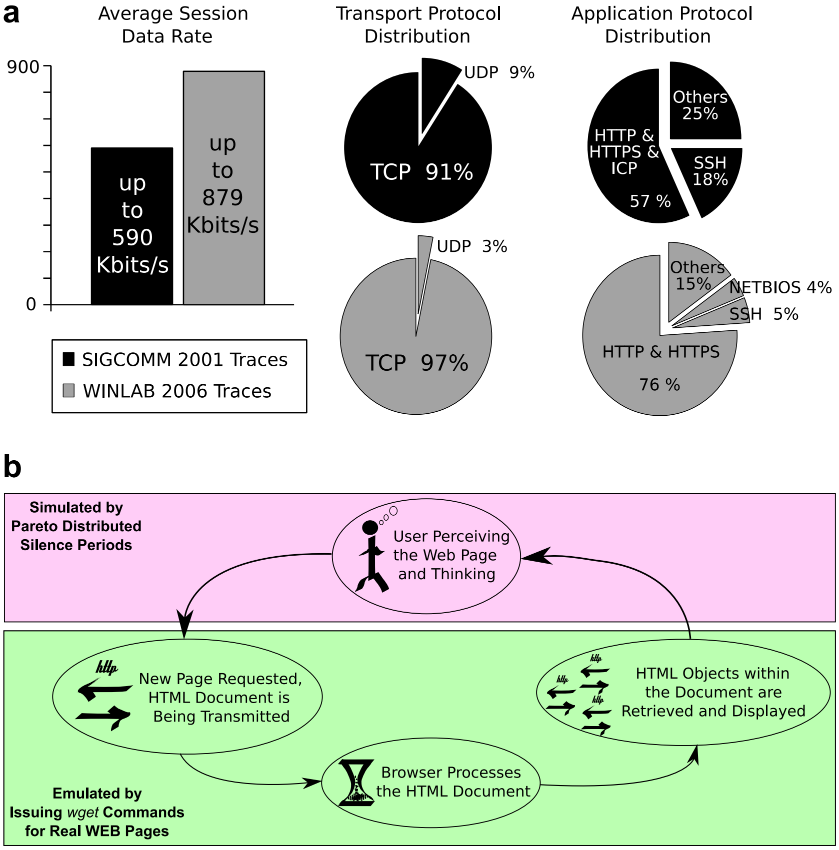

Fig. 2. (a) WLAN application workload characterization. (b) ON–OFF model for emulating realistic WWW access [21].

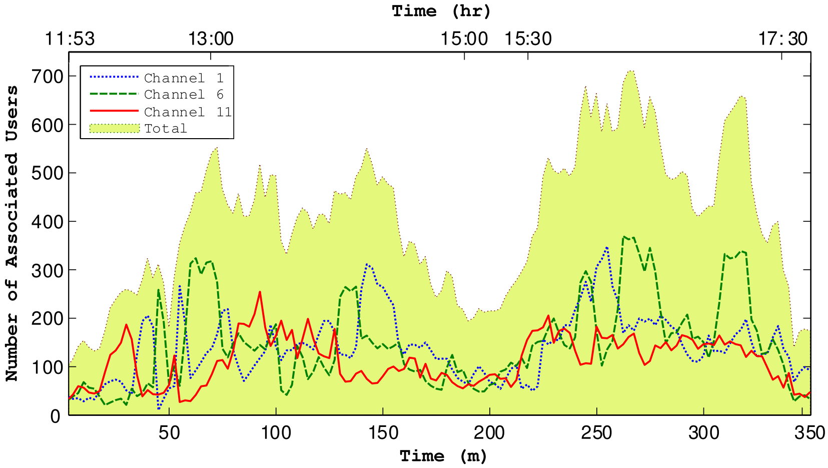

Fig. 3. Dynamic client arrival patterns from traces of the 62th IETF meeting [16] showing rapid changes in number of associated users.

illustrated in Fig. 4a. All APs use the same SSID, thus STAs select the AP with the highest Received Signal Strength Indication (RSSI) at their position. In this experiment, the four APs have 32, 13, 13, and 17 stations associated, respectively.

Fig. 4b summarizes the system performance in terms of the cumulative system throughput for the network. We attribute the throughput spikes, especially with about 30 STAs, to our user model for web browsing traffic (i.e., thinking/perceiving phase vs. fetching/downloading phase). As more STAs arrive, communication demands increase, and an averaging effect is observed in the overall system throughput. Our observations from this experiment are:

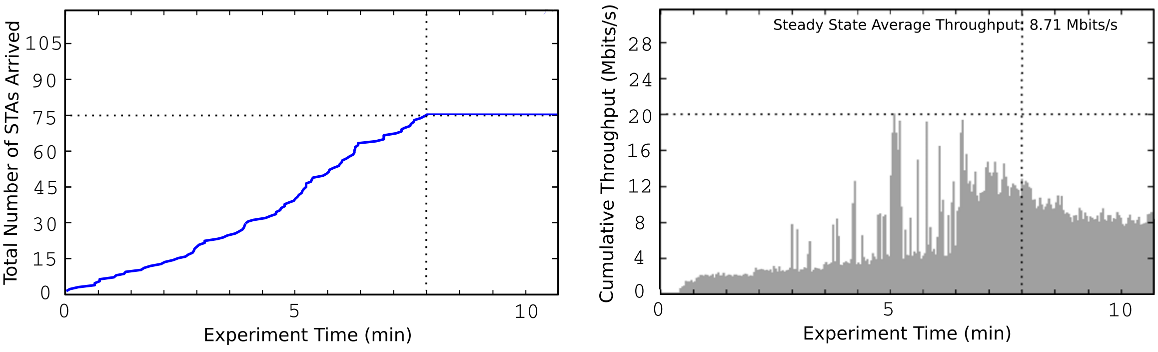

(a) Station (STA) arrival dynamics (b) ORBIT experiment for 4 APs, 75 STAs

Fig. 4. System performance for the most general case: multi-cell with many clients using realistic traffic and arrival patterns.

- The steady state average system throughput, *8.71 Mbits/s* (calculated over the last three minutes of the experiment after all 75 STAs arrived) is below the capacity of a 54 Mbits/s IEEE 802.11a network [25]. In fact, a one AP/one-STA baseline experiment using TCP bulk data transfer, we conducted, yielded *24.02 Mbits/s* average steady state throughput in the same experiment environment.
- Distribution of clients across access points is uneven. One of the APs in the experiment serves more than twice the number of STAs of another AP. To investigate this further, we conducted an experiment with 320 stations (STA) and 12 co-located (to the extend ORBIT testbed's fixed node placement allows) APs. We observe a similar uneven distribution as shown in Fig. 5. At the stations, we have logged RSSI measurements for the co-located APs and found out that they differed likely due to a combination of multipath, antenna gains, RF frontend dynamics, connectors, and cabling. In summary, even a high-density distribution of client positions with line-of-sight propagation does not necessarily lead to an even distribution of associations over co-located access points. This observation motivates the need for association control techniques in multi-channel/multi-AP WLAN installations (e.g., [10]) aiming to create even load distribution within the network.

To understand the root cause of threefold throughput change, we will study the effect of access point density, traffic workload, dynamic user arrivals, and station density on this result.

3.2. The Effect of the Number of APs

To measure the effect of *the number of APs*, we repeat the previous experiment with one, two, and three access points, while keeping the total number of STAs and their offered load the same. Results from this experiment are

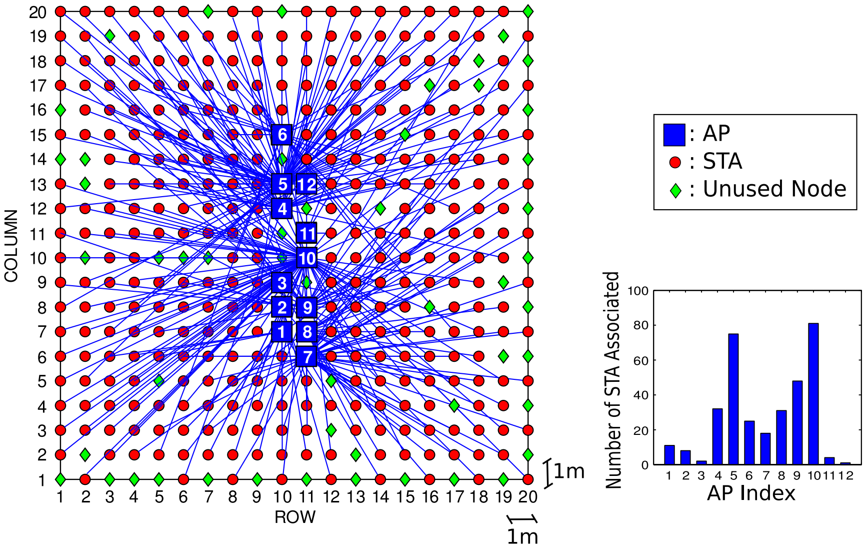

(a) STA and AP Distributions

Fig. 5. Experiment with 320 stations and 12 APs showing association patterns and distribution to access points.

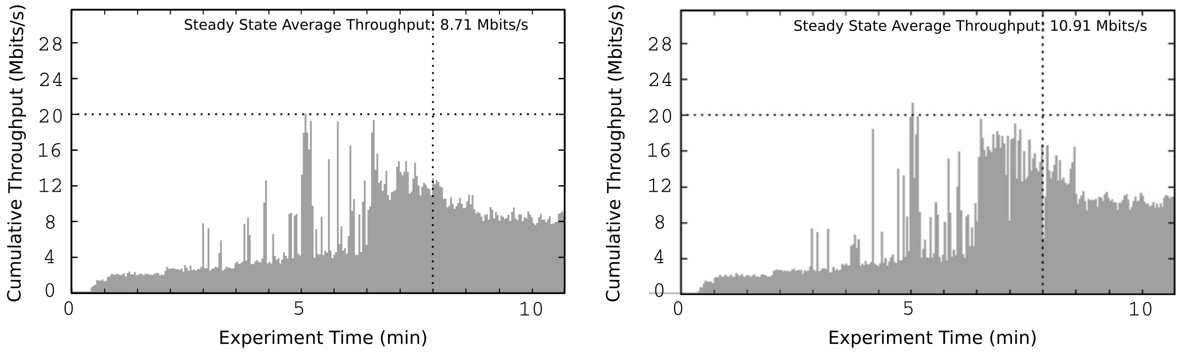

(a) ORBIT Experiment for 4 AP, 75 STAs (b) ORBIT Experiment for 3 AP, 75 STAs

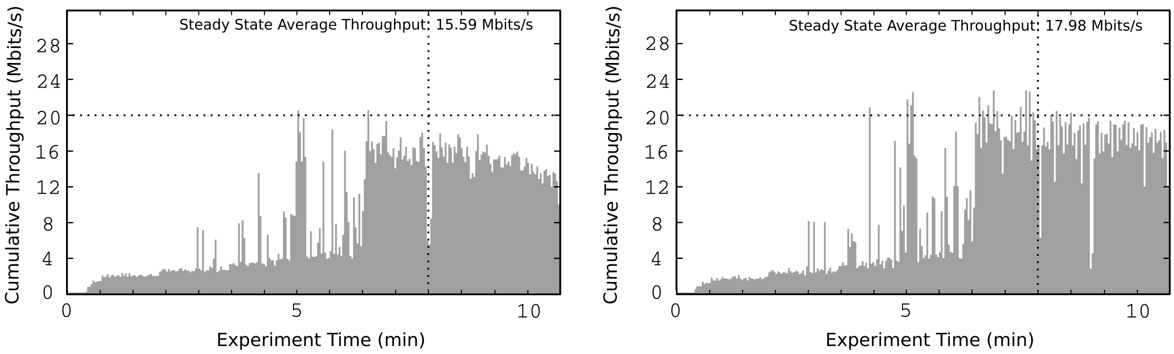

(c) ORBIT Experiment for 2 AP, 75 STAs (d) ORBIT Experiment for 1 AP, 75 STAs

Fig. 6. Investigating the effect of the number of APs to system performance. Everything else is kept the same as the previous experiment in Section 3.1.

illustrated in Fig. 6, shown together with the prior four AP experiment result to facilitate comparisons. Our observations follow:

- All four experiments show two distinct phases. During the first phase, lasting until about 75 STAs are associated, system behavior is comparable across all four cases. We attribute this to the very spiky offered load staying mostly below the capacity provided by the system – throughput is *offered-load limited*.
- In the following *capacity-limited* steady state phase (after all 75 STAs join), configurations with fewer APs result in a significant increase in throughput. In particular, using three APs increases the average throughput by 25.2% compared to four APs. Similarly, reduction to two APs increases the average throughput an additional 53.8% compared to three APs. Finally, a single AP network achieves an average system throughput of *17.98 Mbits/s*, which is about a 106.4% *performance improvement* compared to the four AP network. Despite this improvement, the average system performance is still below our 24.02 Mbits/s baseline result.

3.3. The effect of traffic workload, dynamic station arrival, and number of stations

To analyze the remaining difference, we first repeat the previous single access point experiment (using short-lived TCP sessions generated according to a web-dominated workload) with a bulk TCP transfer workload. The bulk TCP workload is constructed by initiating a TCP session from each STA upon arrival and downloading a large file (large enough to not complete during the experiment duration) from a server on the wired network. Results of this experiment are illustrated in Fig. 7a. We observe that the bulk TCP workload results in an average system throughput 23.2 Mbits/s, an increase of 29% compared with the web-dominated workload. We believe that the main cause for this throughput difference is TCP's rate control mechanism not adjusting quickly enough to the optimal TCP congestion window size when short lived TCP sessions are dominant in the network.

While the throughput with bulk TCP is close to the baseline experiment, we will also show the effect of dynamic station arrivals for the purpose of completeness. Fig. 7b presents results from the 75 STA experiment with bulk TCP workload and all stations joining the network simultaneously (i.e., without the dynamic arrival pattern). In steady state, only a negligible difference of 0.2 Mbits/s can be observed between the two experiments. We also observed high throughput fairness at steady state for both arrival patterns. Jain's fairness index [20] is 0.96 with realistic arrivals and 0.94 otherwise.

We have also investigated the STA association performance, since it is the other major system functionality likely to be affected by the type of the arrival pattern being used. In the dynamic arrival case, new clients come into an

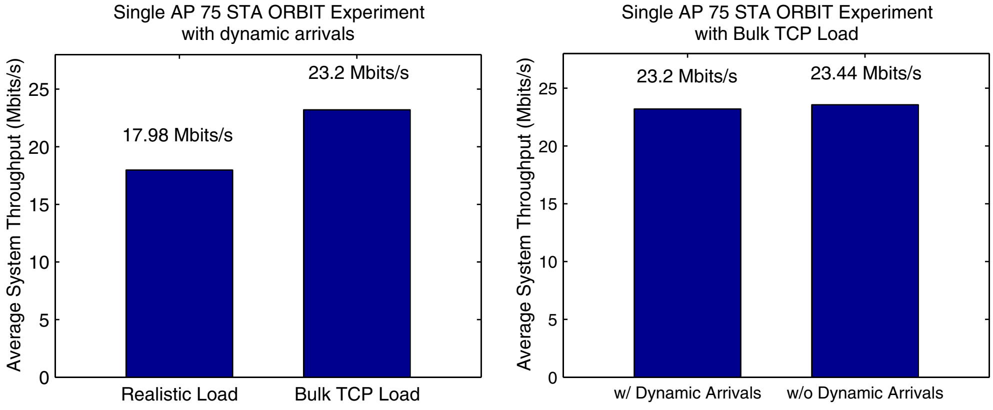

(a)Effect of the Traffic Workload (b)Effect of the Arrival Pattern

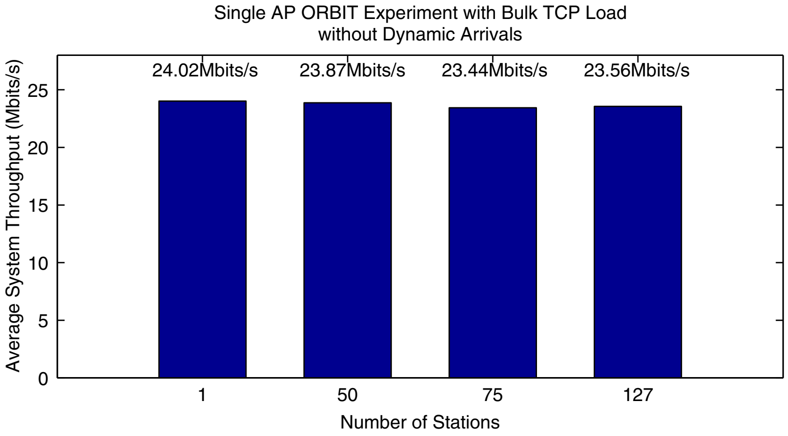

(c)Effect of the number of STAs

Fig. 7. Results from the ORBIT experiments investigating the effects of three different factors on system performance.

already loaded WLAN, and associations are frequently delayed. We observed that with dynamic arrivals, about 40% of the STA associations take e 270 ms, about 50% of them take 5.5 s and a further 5% take up to 17 s. In the case where all STAs appear at the same time, we observed that 80% of the associations complete within 270 ms while the remaining 20% takes up to 5.5 s. This stepping behavior can be explained by association disruptions, whereupon the driver enters an active scan cycle of all 23 IEEE 802.11a/g channels. Scanning takes approximately five seconds in the Madwifi [15] driver implementation for Linux.

The final minor difference in throughput is due to the number of stations. Fig. 7c shows the same experiment repeated with 1, 50, and 127 stations.[5] Results show little dependence on the number of stations.

From the experiments we have conducted so far, it is empirically observed that system performance has the strongest dependence on the *number of interfering APs*. We will continue with the investigation of this dependence in the following sections.

3.4. Discussion

The foregoing experiments show that (i) a single AP network performs efficiently under TCP workload irrespective of the number of stations it serves (up to 127 stations in our setup) and that (ii) a multi-AP network serving the same number of clients on the same channel leads to significant throughput degradation.

As a byproduct, secondary effects that have been observed in practice such as inefficient bit-rate adaptation [9] do not manifest themselves in *single-cell deployments* with such traffic characteristics no matter how many users actively use the system.

These results show that congested WLAN systems cannot be fully understood through traditional MAC layer analysis. Models from neither Bianchi [4] nor Kumar et al. [5] explain these results. According to Bianchi's model the efficiency of a WLAN depends primarily on the number of active stations, regardless of their role as access points or clients.

The single AP result confirms more recent theoretical and experimental work with a smaller network setup [6,23] which have suggested that TCP flow control, when used in a single-cell network, operates the network effi-

---

[5] 127 STA limitation comes from the particular Madwifi driver version we have used on our APs.

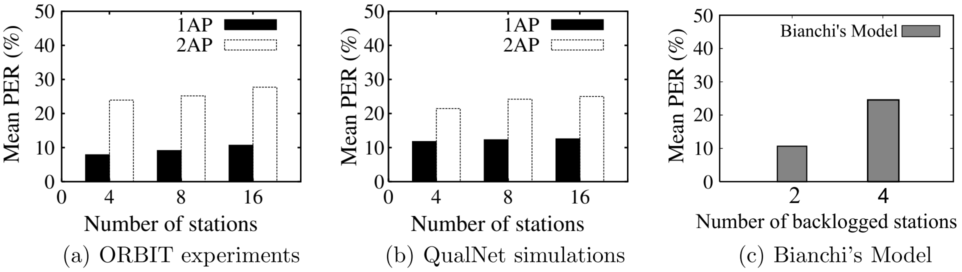

(a) ORBIT experiments (b) QualNet simulations (c) Bianchi's Model

Fig. 8. (a) Empirical and (b) simulated collision-caused packet error rate (PER) with up to 16 clients and 2 APs. (c) Collision-PER predicted by Bianchi's Model.

ciently and maintains robust system throughput regardless of the number of STAs. The multi-AP result, to our knowledge, has not yet been reported and we analyze it further in the following section.

4. TCP analysis in multi-cell networks

The reduction in cumulative throughput for multi-cell networks raises the following inter-related questions:

- Why is TCP Reno over WLAN robust against intra-cell congestion but not against inter-cell congestion and interference?
- Does TCP Reno adjust the flow rate to minimize collisions?

In this section, we will answer the above questions by identifying the applicability of Gong and Marbach's TCP [26] model to multi-cell networks, validate it through experiments and simulation, and finalize by discussing TCP's flow control's ability to identify the optimal operating point.

4.1. Gong and Marbach's model for multi-cell networks

According to Gong and Marbach's model [26] of TCP, in a single BSS case, on average, two stations will be backlogged, irrespective of the number of clients. Also, if n additional flows, in the form of Independent basic service set (IBSS, a.k.a. an ad hoc network) are added as interferers, the expected number of active (i.e., backlogged) nodes in the network would be $2(1+n)$.

While Gong and Marbach do not comment on the multi-AP case, all assumptions made for IBSS flows also hold for BSS networks. Following the same steps, one can therefore also derive the following proposition:

Consider multiple IEEE 802.11 infrastructure networks (BSS) with the following two characteristics: (1) Each BSS consists of at least one station and a single AP. All BSS are within transmission range of each other. (2) There is a single TCP connection per client and applications using TCP connections always have data to send. For a network topology consisting of i BSSs, in steady state, the expected number of backlogged nodes then equals $2i$.

4.2. Verification via experiments and simulations

To verify the proposition in Section 4.1, we study the collision rate observed in the ORBIT experiments,[6] where we use bulk TCP workload on six different network configurations (one and two APs, four, eight and sixteen STAs). Using Bianchi's model, we can then calculate the collision rates for the number of backlogged stations predicted by Gong and Marbach's model and compare it with the measurement result. We use this indirect approach, since MAC queues are maintained in hardware and we cannot directly determine whether a station is backlogged.

We further approximate the collision rate with the overall packet error rate (PER) from these experiments, since WLAN devices cannot distinguish collisions from other transmission errors and thus do not allow direct collision measurements. This approximation is accurate, since we operate in a high SNR environment where frame-errors due to pure channel bit-errors are negligible (<1% in our tests). As additional validation, we also simulate the same configuration in QualNet where we can extract the exact collision rate.

Fig. 8a and b shows the mean collision rate as the number of STAs is increased for one and two AP networks in the ORBIT testbed experiments and simulations, respectively. In both cases, we observe that PER due to collisions is marginally affected by an increase in the number of stations but is significantly affected by an increase in the number of APs – an increase from 1 to 2 APs more than doubles the average PER from 11% to nearly 28%. Near identical results from the simulations and the ORBIT experiments also indicate that the collision rate approximation we used was reasonable.

The empirical collision rate with a single AP (for 4, 8, and 16 STAs) matches the PER value predicted by Bianchi's saturation model [4] with two backlogged stations. Similarly, the empirical collision rate with two APs (for 4, 8, and 16 STAs) matches the PER value predicted by Bianchi's model with four backlogged stations. Predictions of Bianchi's model, obtained from MATLAB simulations, for the same experiment configuration are provided in Fig. 8c. To further verify the accuracy of this model, we conduct additional

---

[6] In these experiments, we do not consider dynamic STA arrivals and bit-rate adaptation in order to focus on MAC contention in a baseline scenario.

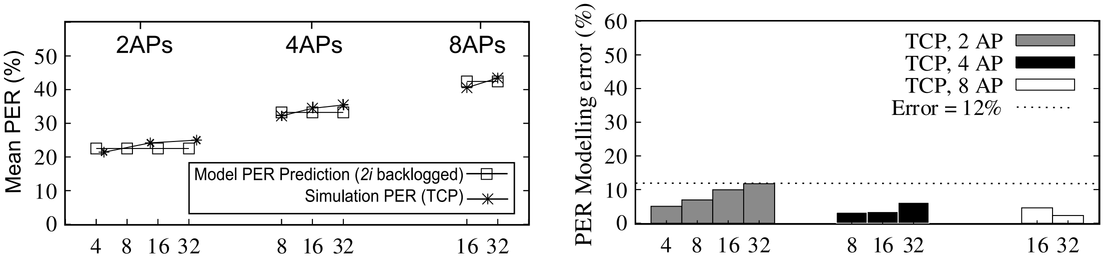

(a) PER (multiple APs)   (b) PER modelling error

Fig. 9. Accuracy of the model relative to simulation results. All simulation results exhibited very small variance characteristics, hence errorbars are not shown here.

simulations with four and eight APs connecting up to 32 STAs and report the modeling error. Observed PER from model predictions and simulations closely follow each other. The percentage modelling error is illustrated in Fig. 9b – the worst case error is only 11.76%. Moreover, modeling error reduces with an increasing number of interfering APs. We also observed that these simulation results exhibited negligible variance (a maximum PER variance of 0.12) across ten runs with different random seeds. These results verify that each active cell increases the total number of backlogged stations in the network by *two*, as stated by the proposition in Section 4.1. The network efficiency in a TCP dominated network, for this reason, is primarily a function of the number of interfering APs.

4.3. Discussion

Incidentally, an IEEE 802.11a network running with two backlogged stations maximizes throughput according to Bianchi's model. Fewer backlogged stations lead to too much idle time, while more backlogged stations lead to too much collision overhead. Thus, one may ask whether TCP's flow control algorithm, designed for managing congestion in the Internet, can also control the flow rate to maximize throughput in congested wireless networks? And if so, why does it not achieve this under co-channel interference from other access points?

To address these questions, we conducted simulations in which we change the MAC CW_{min} parameter, while keeping the number of stations constant (32) and using bulk TCP transfers to saturate the channel. If TCP can identify the optimal operating point to maximize throughput in single AP networks, it should respond to the changed MAC CW_{min} settings with a corresponding change in the number of backlogged stations. Recall that to maximize 802.11 network throughput, the number of backlogged stations has to increase with an increase in MAC CW_{min}, to balance collisions against idle overhead [4].

Fig. 10 shows the MAC PER under TCP for the single AP case (with 32 STA) with increasing CW_{min}. It also shows the expected packet error rate when only two stations are backlogged (i.e., no TCP adjustment) labeled as "analytical PER prediction" and an optimal PER curve, which corresponds to the PER at which the cumulative MAC throughput is maximized. Note that the observed (measured) MAC PER curve approaches the optimal value for CW_{min} = 16, but does not follow the curve for higher CW_{min} settings. Instead it tracks the PER for two backlogged stations, indicating that in general, TCP's flow control cannot identify the MAC operating point at which throughput is maximized.

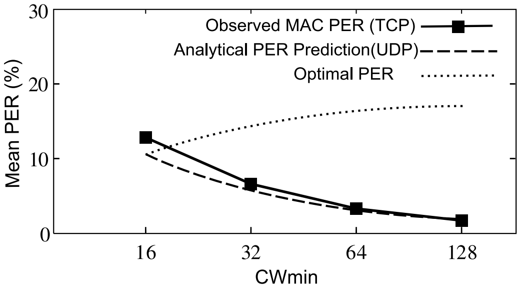

Fig. 10. Collision rate behavior of TCP under varying CW_{min} values.

Instead the measured results can be explained through an interaction between flow control and MAC layer channel access. A TCP workload leads to two backlogged stations for each access point because TCP requires an equal packet flow rate both up- and downlink directions (i.e., DATA and ACK). Regardless of the direction of the data traffic, the flow is limited by the AP's MAC layer channel access probability. Two backlogged stations means that on average the AP itself and one associate station are backlogged. If more associated stations were backlogged, they together would have a higher channel access probability than the access points, and the backlogged queues would empty. Similarly, adding more APs increases the cumulative channel access probability for APs, hence they can excite more client stations (one per AP).

Note that this result is not a function of the downlink dominance of the traffic [23]. We have also independently carried out simulations to compare a TCP-uplink dominated scenario and observed that direction of data traffic did not change the throughput or collision rates significantly (<1%).

5. Effect of AP density on multimedia over WLAN

In this section, we evaluate the effect of inter-cell interference on multimedia traffic, with and without wireless

Table 4
Empirical VoIP performance

No. of APs	VoIP Call #1			VoIP Call #2		
	Packet drop rate (%)	Avg. jitter (ms)	Avg. latency (ms)	Avg. drop rate (%)	Avg. jitter (ms)	Avg. latency (ms)
1	0.14	2.11	54	0.78	2.16	78
2	1.55	4.16	77	1.12	2.50	65
3	2.97	4.82	101	1.84	5.76	138
4	11.40	8.63	304	10.82	7.85	191

multimedia extensions (WME).[7] We place a special emphasis on voice-over-IP (VoIP) application performance since most current video streaming traffic is not interactive, i.e. it can be buffered and transported over HTTP/TCP (e.g., YouTube™ video streaming). The recent emergence of IEEE 802.11 in VoIP handsets and cell phones, however, can be expected to lead to increased VoIP usage on access points.

5.1. Performance over Legacy IEEE 802.11a

We first characterize VoIP performance in environments with legacy IEEE 802.11a stations using our realistic workload mix. For this purpose, we conduct ORBIT experiments configured as in Section 3.1 (IEEE 802.11a, 75 STAs with realistic arrivals, and up to 4 APs). We then designate two of the stations randomly as VoIP devices and allow them to each initiate a VoIP call towards our sink on the wired network of the testbed. VoIP sessions start at uniformly distributed random times during the experiment and last for 270 s. A VoIP call runs over an RTP/UDP session using G.711 with one sample per packet as the codec (voice activity detection disabled). For each VoIP call, we measure packet error rate, mean jitter, and mean latency. Table 4 summarizes results from these experiments.

As per ITU-T Recommendation G.114 [27], we consider 150 ms as the upper latency limit for acceptable VoIP communications. Similarly, we consider 75 ms and 3% the upper limits on jitter and packet loss, respectively [28]. Our key observations follow:

- A single congested access point can support two VoIP calls with acceptable performance, even without the use of WME quality of service differentiation.
- The two VoIP calls can be supported with adequate performance only in a congested environment with no more than two interfering APs. With the addition of the third AP, packet loss rate and latency values for both VoIP calls approach the limits described above. For four APs, packet error rates and latency reach an unacceptable 11% and 300 ms, respectively.

Overall, the addition of two VoIP flows has a negligible effect on the throughput of the other TCP flows, since each VoIP flow only generates an application layer load of 64 Kbits/s. We do not observe unacceptable jitter in any repetition of the experiment, likely because the VoIP receiver was placed only one hop away from the VoIP sources on our wired testbed network.

These results indicate that for WLAN-enabled hybrid phones to work well in interference-limited multi-AP deployments (e.g., apartments with many residents each using their own WLAN) changes to legacy IEEE 802.11a/b/g are inevitable. Next, we investigate whether the recent changes provided by WiFi Alliance's WME address this problem.

5.2. Performance over IEEE 802.11e (WME)

For WME experiments, we relied on QualNet simulations, because MadWifi WME implementation [15] on Atheros WLAN hardware was in its early stages of development while we conducted our experiments. We first conduct WME-enabled version of the VoIP experiments we carried out in the ORBIT testbed with two VoIP devices, for the purposes of facilitating direct comparisons. Then, we also experiment with a higher number of VoIP sessions as well as scenarios involving video streaming sessions. For the latter, we use the mean opinion score (MOS) metric to characterize user perceived audio quality of VoIP calls. For the video streaming traffic, we configured the throughput mean and variance of the offered load to those captured from a five minute long movie segment encoded from DVD using DivX 5.1. The mean data rate of the stream used in the experiments presented here was 382.5 Kbit/s.

We assign application traffic to the WME MAC access queues as follows: video streaming to the *Video* queue (*WME_AC_VI*), and VoIP calls to the *Voice* queue (*WME_AC_VO*) and all other TCP data flows to the *Best Effort* queue (*WME_AC_BE*). EDCA parameter-set values for all four access queues conform to the default values suggested by the IEEE 802.11e standard [14].

Results from the WME-enabled simulations of one and four AP networks are given in Table 5, and they can be compared to the ones reported in Table 4, obtained from the experiments on the ORBIT testbed. Fig. 11 presents the results from the second set of experiments where we report the subjective quality for each of the VoIP calls in terms of MOS. These experiments vary the number of VoIP and video streaming sessions and the number of interfering access points. We observe that:

- Use of WME improves media quality in the multi-cell case. Average latency values of the two VoIP calls in the four AP network remains acceptable, in contrast to the non-WME case. With WME, up to 10 concurrent

---

[7] WME is an interoperability standard from the WiFi Alliance based on the IEEE 802.11e standard [14].

Table 5
Simulated VoIP performance

No. of APs	VoIP Call #1			VoIP Call #2		
	Packet drop rate (%)	Avg. jitter (ms)	Avg. latency (ms)	Avg. drop rate (%)	Avg. jitter (ms)	Avg. latency (ms)
1	0	1.01	30	0	1.21	34
4	4.11	6.86	117	6.20	7.24	126

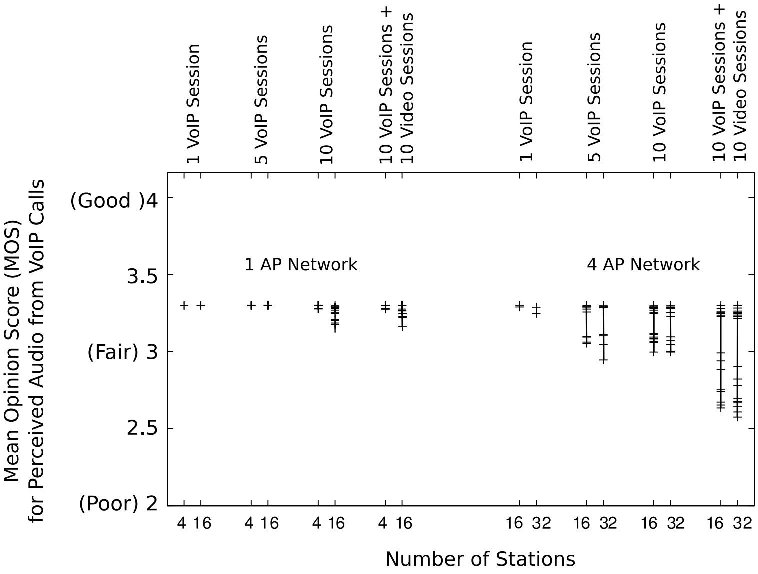

Fig. 11. Simulated VoIP MOS results for varying intensities of multimedia traffic carried over one and four AP networks using IEEE 802.11e (WME).

VoIP sessions can now be supported with a MOS at the fair level, compared to 2 VoIP sessions with 2 interfering access points without WME.

- With increasing amounts of media traffic, the MOS for VoIP sessions still indicate significant quality degradation. At 10 video and 10 concurrent VoIP sessions, average quality approaches the poor rating.
- While not shown in the figure, we also observed a tenfold increase in video streaming jitter for the four AP scenario. For non-interactive videos this can likely be addressed through application buffering. For video conferencing applications, however, this jitter may be unacceptable.

As expected, we observed a throughput reduction for best effort traffic when more media streams are added, demonstrating effective MAC layer prioritization of media traffic. Also note that switching to WME reduces capacity even without media traffic, because the best effort queue to which all regular traffic is assigned uses a larger interframe space than default IEEE 802.11a/b/g (i.e., an AIFSN of 3). This increase in the interframe space by one slot (9 μs for IEEE 802.11a, and 20 μs for IEEE 802.11b) results in an 14% reduction of best effort throughput with 32 stations and 4 APs in IEEE 802.11a. We observed that up to 14% drop in throughput was likely for 32 station 4 AP WME simulations, when compared to the results from legacy IEEE 802.11a (see Fig. 12).

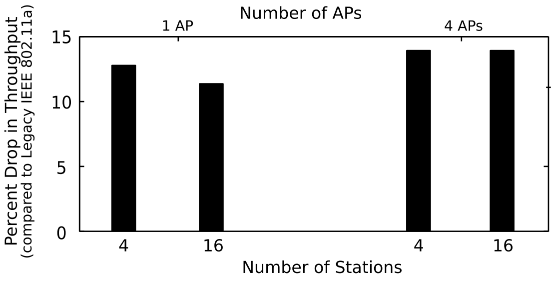

Fig. 12. Default value for AIFSN parameter of IEEE 802.11e *Best Effort* queue results in decreasing throughput for bulk data traffic, even in the absence of any multimedia traffic.

Overall, WME provides an improvement in interference tolerance but not a complete solution. Additional measures will be needed to make media traffic resilient against inter-cell interference. We investigate such additional techniques next.

6. Improving performance of unplanned dense AP deployments

We have seen that inter-cell interference reduces cumulative throughput in WLAN systems with TCP-dominant workloads much more significantly than intra-cell contention. Also, the detrimental effects of inter-cell interference are more severe on multimedia traffic.

We propose a contention window adaptation solution that addresses this challenge at the MAC layer. We propose a MAC layer solution because the root cause of the problems we observed was increasing contention due to added interference, and it could be best addressed with a local solution at the wireless last hop. End-to-end approaches (e.g., TCP tweaks) or network-wide solutions (e.g., IP tweaks) are undesirable due to changes that will be required at millions of hosts all over the world. The other PHY/MAC layer techniques are complimentary to our proposal, they can increase overall capacity, but in chaotic dense deployments several interfering access points might still remain after applying these techniques. Most relevant ones are visited in Section 7.

We first describe and evaluate the contention window adaptation approach, and then end with a discussion that highlights other challenges that should be addressed for a complete system solution.

6.1. CW_{min} adaptation using active AP count

It is well-known that IEEE 802.11 MAC performance can be improved by selecting a $\mathrm{CW_{min}}$ appropriate for the current number of active clients in the network [29–31]. We propose instead that *the selection of CW_{min} for a typical TCP dominated multi-cell WLAN system should be based on the number of interfering access points, not the number of clients.* This method provides the following advantages:

- Knowing the number of access points allows more accurate selection of $\mathrm{CW_{min}}$ under a TCP workload than knowing the number of transmitting clients. As our earlier analysis has shown the collision rate under TCP is determined by the number of access points, since TCP flow control regulates client activity.
- The number of active APs is easier to obtain than the number of clients, since APs announce their presence through beacon packets and tend to remain active and stationary over longer durations. Clients might also help determining the number of access points by listening for AP beacon packets.

Given the number of active APs N_b, the optimal contention window can be derived by combining our insight regarding the number of backlogged stations under TCP workloads with earlier contention window adaptation work. According to [31] bandwidth can be maximized with a contention window $\mathrm{CW}_j^* = \frac{\sqrt{2\beta T'_D}}{r_j} + 1$, where $\beta = \left(\Sigma_{j=1}^M N_j r_j\right)^2 - \Sigma_{j=1}^M N_j r_j^2$ and N_j is the number of active users for each priority queue. Since we do not consider multiple priority queues here, β reduces to $N(N-1)$. Also substituting $T'_D = \frac{T}{T_s}$ yields

$$\mathrm{CW_{min}} = \sqrt{2N(N-1)\frac{T}{T_s}} + 1, \qquad (2)$$

where T is the time required for the transmission of a MAC frame (excluding DIFS and backoff, including ACK reception), and T_s is the duration of a MAC time-slot. Using the insight from Section 4 that the number of backlogged stations equals twice the number of active access points, we can substitute $N = 2N_b$, which formulates $\mathrm{CW_{min}}$ in terms of the number of active access points as

$$CW_{min} = 2\sqrt{N_b(2N_b - 1)\frac{T}{T_s}} + 1. \qquad (3)$$

6.2. Implementation of CW_{min} adaptation

This subsection describes an AP-centric algorithm for determining the number of active access points N_b and distributing the $\mathrm{CW_{min}}$ setting. While a client-centric or hybrid approaches are also possible, we have chosen this approach because it minimizes the number of devices that need to be modified to the access points.

The number of active access points N_b should ideally only include access points that actively communicate with at least one client. Since APs transmit beacons even when none of their clients are active, however, the proposed algorithm determines the expected $E[N_b]$, by considering the percentage of the time neighboring APs consume channel resources. The IEEE 802.11h standard [32], which defines spectrum and transmit power management extensions, allows for channel-related measurement-exchange mechanisms that can be extended to support $E[N_b]$ calculation in a standards-compliant way.

Algorithm 1. $\mathrm{CW_{min}}$ update algorithm APs execute. It finds the effective number of active APs in the vicinity and uses this to update $\mathrm{CW_{min}}$ and propagate it further down to its STAs with the next beacon.

```
// Accepts: CurrentChannel, $t_p$, $T$, $T_s$
// Updates: (System Parameter) $\mathbf{CW_{min}}$
for every $t_p$ seconds do
    NeighborAPList[]=doBackgroundScan(CurrentChannel);
    for each $AP_i$ in NeighborAPList[] do
        sendFrame(80211H_MEASR_REQUEST, $AP_i$);
        recvFrame=readFrame(80211H_MEASR_REPORT);
        %_Backlg_i=parseFrame(recvFrame, BACKLOG_FRACTION);
    end
    %_Busy_self = measure(CHANNEL_BUSY_FRACTION, self);
    $E[N_b] = \frac{1}{100} \cdot [\%\_Busy_{self} + \sum_{\forall i} \%\_Backlg_i]$;
    $\mathbf{CW}_{\min} = \lceil 2\sqrt{E[N_b](2E[N_b] - 1)\frac{T}{T_s}} + 1 \rceil$;
    update_EDCA_ParamSet(CW_min);
end
```

The AP-centric approach for this purpose is outlined in Algorithm 1. The algorithm periodically calculates $\mathrm{CW_{min}}$ every t_p *seconds* by requesting neighboring APs (visible through their beacons) to report the percentage of the time their frame queues were not empty (i.e. *BACKLOG_FRACTION*) during the last t_p seconds. This message exchange can be implemented within the IEEE 802.11h measurement request/report framework by using one of the *reserved* measurement type definitions. Querying AP sends an IEEE 802.11h measurement request frame (80211*H_MEASR_REQUEST*) to the AP of interest, and the receiving AP responds to this with a measurement report frame (80211*H_MEASR_REPORT*) including its *BACKLOG_FRACTION* measurement (see lines 2–7 of the algorithm). These reports from neighboring APs, when combined with the measuring access point's own wireless channel-busy percentage measurement (*CHANNEL_BUSY_FRACTION*), allow determining $E[N_b]$ more precisely; a neighboring AP which is backlogged 50% of the time for example, increases $E[N_b]$ by $\frac{1}{2}$. $E[N_b]$ calculated this way is used to update $\mathrm{CW}_{\min}$ (see line 10 of the algorithm), which in turn is included in the next AP beacon to be announced to the STAs of this AP through IEEE 802.11e EDCA parameter set information element [14].

6.3. Simulation results

We conduct fixed-bit-rate simulations with 32 STAs to observe the potential improvement that our $\mathrm{CW}_{\min}$ selection can provide. The results assume that all stations accurately estimate the number of active access points. Fig. 13 shows that for both exact $\mathrm{CW}_{\min}$ values suggested by the Eq. (3) above, and for the nearest power of two (which is a practical restriction in today's WLAN hardware), $\mathrm{CW}_{\min}$ adaptation based on the number of APs reduced the collision-based losses significantly, keeping them close to the residual collision losses of a single AP scenario. Also, the granularity of $\mathrm{CW}_{\min}$ adaptation in powers of two does not have a significant negative effect on performance. Note that reduction in PER not only improves throughput but also reduces delay that affected multimedia streams as described in Section 5.

Note that the optimality of this $\mathrm{CW}_{\min}$ tuning strategy, in terms of achieving proportionally fair bandwidth allocation and maximizing utilization has already been shown in [31]. Note also that, in the trivial case of a single-cell, $N_b = 1$ and $\mathrm{CW}_{\min}$ is reduced to a constant, consistent with the result that, in the case of a single-cell, the number of backlogged stations is constant. Finally, the proposed $\mathrm{CW}_{\min}$ adaptation strategy is optimized for TCP dominated workloads (as compared to the earlier work which provides a solution for UDP-dominated workloads [31]). In mixed traffic environments, where UDP accounts for a significant share of network traffic and is used on many stations, the proposed solution would have to take into account the increase in the number of active nodes due to this additional UDP traffic. Estimating the exact number of active nodes in this case remains an interesting open topic for future work.

6.4. Discussion: relationship to collision-resilient bit-rate adaptation

Collision-aware rate adaptation is a further MAC layer technique to improve performance in congested IEEE 802.11 networks. Prior work [9,33,34] has shown that many bit-rate algorithms choose low bit-rates in congested environments – short-term collision errors are misinterpreted as longer-term changes in path loss. The bit-rate reduction further decreases available capacity, leading to more collisions.

One might believe that $\mathrm{CW}_{\min}$ adaptation eliminates the need for collision-aware rate adaptation techniques because it substantially reduces collisions. From the $\mathrm{CW}_{\min}$ adaptation results in Fig. 13a, we observe, however, that even at the optimal operating point a residual collision rate remains. We also observe from our multi-cell experiments that bit-rate adaptation is responsible for a substantial fraction of throughput loss and that even with the residual collision rate bit-rate adaptation in the MadWifi driver does not choose optimal rates. Fig. 14 shows the bit-rate distribution of SampleRate [35], the default bit-rate adaptation scheme in the MadWifi driver, for up to 10 access points and 50 clients. With more APs, an increasing percentage of frames are transmitted at the lower bit-rates, even though signal-to-noise ratio (SNR) at most receivers in our experiment environment is high enough for communication at 54 Mbits/s and the total number of clients remains constant. For example, 60% of frames use 6 Mbits/s for the 10 AP experiment. Even with the residual collision rate in the one AP case, however, less than 20% of frames are transmitted at the optimal bit-rate of 54 Mbit/s.

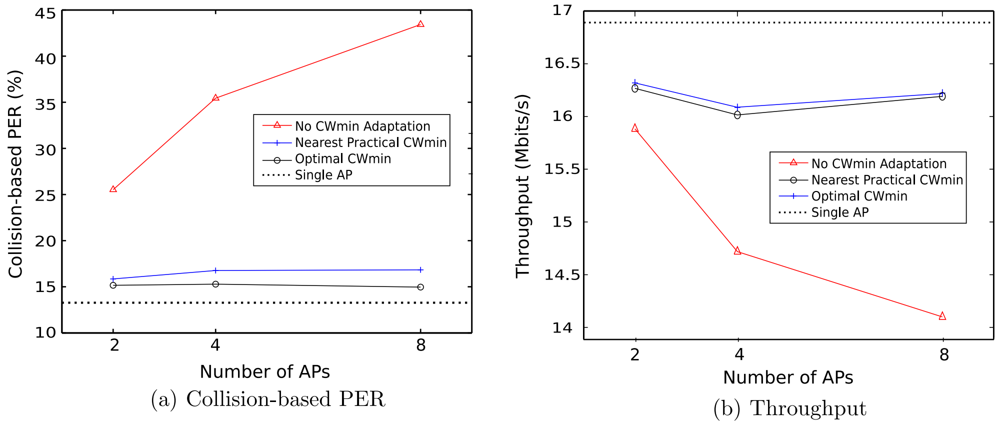

(a) Collision-based PER
(b) Throughput

Fig. 13. Simulations with 32 STAs showing potential gains from incorporating AP Count-based CW adaptation into WLAN with inter-cell interference (in (b), throughput curves for "Optimal $\mathrm{CW}_{\min}$" and "Nearest practical $\mathrm{CW}_{\min}$" almost overlap).

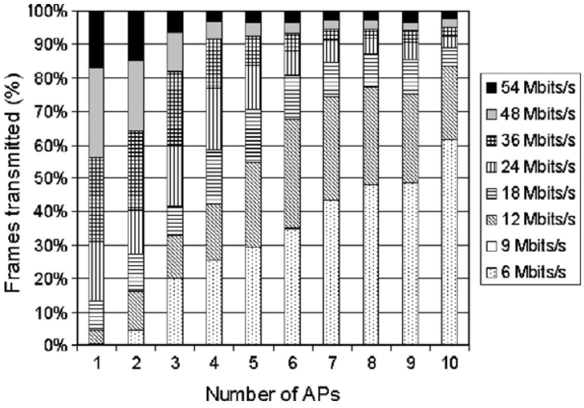

Fig. 14. Empirical IEEE 802.11 bit-rate distribution with the SampleRate algorithm.

Hence, we believe that improving the collision-resiliency of bit-rate adaptation mechanisms must be an integral part of the solution space. Based on the observations in [34,36,33], a potential solution to achieving collision-resiliency could be through the adaptive use of the RTS/CTS mechanism. Rate adaptation algorithms could leverage this backwards-compatible mechanism such that they ignore RTS losses, which could be due to collision, and react only to data packet losses that occur after the channel is reserved. Packet losses that occur when the channel is reserved are most likely because SNR at the receiver is not high enough to support the current bit-rate. Thus, this mechanism could be used to distinguish between losses due to collision and those due to poor channel conditions.

7. Related work

Related work, most relevant to our study, that models the interaction between TCP and IEEE 802.11 is presented in [6,23,26]. Choi et al. [6] and Bruno et al. [23] show that for both downlink and uplink traffic, cumulative TCP throughput does not degrade in single-cell WLANs. In [26], a discrete time model explaining the interaction between TCP and IEEE 802.11 is presented for a network topology consisting of multiple source destination pairs along with an infrastructure-like network. In this paper, we showed that their model can also be extended for multi-cell infrastructure networks. In addition, the authors of [37], study TCP fairness in the presence of simultaneous uplink and downlink traffic and observed significant unfairness among TCP flows. They model the interaction between TCP and IEEE 802.11 but do not consider the effect of MAC congestion. We differ from all these studies since we study the interaction between TCP and IEEE 802.11 for high-density multi-cell WLANs.

Several studies in the past have also proposed tuning the contention window (CW) to maximize utilization [29–31]. In [29], the authors determine that finding a balance between the bandwidth loss associated (i) with collisions and (ii) with the time spent by the nodes backing off (idle period) is possible. Heusse et al. [30] use an AIMD algorithm to tune the CW so as to maintain the idle period at a desirable level. Hu et al. [31] analyze the ability of IEEE 802.11e EDCA [14] to maximize bandwidth utilization and provide service differentiation. However, all these studies focus on the single-cell WLAN and address the situation of multiple competing clients. We believe that we are the first to apply CW_{min} adaptation to mitigate interference from adjacent BSSs in unplanned deployments.

Apart from CW_{min} adjustment, there are a number of complementary techniques to address interference-based performance degradation in IEEE 802.11 WLANs. They can be grouped into the following broad categories:

Transceiver parameter optimization: This category includes transmit power, carrier sense threshold and receiver sensitivity adjustments. In [38], the authors propose tuning the carrier sense threshold at the receiver to mitigate interference effects and a dynamic power management technique to reduce interference in unplanned deployments is proposed in [2]. In [39], a combination of receiver sensitivity and clear-channel-assessment adaptation is proposed. However, the authors themselves point out the suboptimal behavior of their approaches in uncoordinated environments. More recently [40], proposes a distributed algorithm to jointly adjust transmit power and IEEE 802.11 bit-rate to reduce interference while not sacrificing performance.

Channel assignment: This category includes static and dynamic channel assignment techniques to mitigate interference. In [41], the authors show that static channel assignment techniques cause unfairness in unplanned deployments and then describe a decentralized channel hopping scheme that improves fairness by distributing interference evenly among neighboring BSSs. We argue that this solution may not fully be able to mitigate contention in environments where the number of interfering APs are higher than the number of orthogonal channels. The CFAssign-RaC algorithm presented in [42] jointly address the issues of channel assignment and load balancing in centrally administered WLANs. Given the uncoordinated nature of the deployments we consider, this solution is not directly applicable here.

Association control and load balancing: These approaches balance client load across a set of APs [42–44] by changing the point of association of the clients. For that reason, they have an inherent assumption of requiring coordination and orthogonality of channels across APs. Recently, [10] propose an association management solution to prevent WLANs from accepting more clients than they can serve efficiently. However, their solution mandates that clients incur delays on the order of minutes.

Some WLAN equipment vendors [45–47] offer WLAN controller-based products that improve performance in high-density deployments through load balancing and other proprietary algorithms. These products are targeted at planned deployments where all access points are connected to the same controller infrastructure, hence they do not address unplanned deployments.

8. Conclusions

In this work, we have investigated the effect of inter-cell interference on unplanned WLAN performance. While inter-cell interference should ideally be avoided through careful access point placement, frequency selection, and transceiver parameter control, current chaotic wireless deployments and the limited number of available channels make inter-cell interference a reality. Therefore, we have measured the effect of such interference both in a testbed with more than one hundred nodes and through simulations. We have

- Found that cumulative throughput degrades significantly, by 50% with four interfering access points, while it remains remarkably robust with over one hundred clients in the single-cell case.
- Verified that TCP's flow control leads to an average of $2i$ nodes concurrently backlogged in the network, where i is the number of actively interfering access points. Thus, the collision rate increases with the number of access points. TCP does not adjust its flow rate to optimize throughput for different CW_{min} choices.
- Showed that increased collision rate with inter-cell interference also affects media streaming. With only two congested interfering access points, VoIP mean opinion score (MOS) is unsatisfactory.
- Described a novel and practical approach that selects the appropriate CW_{min} based on the number of active access points, not the number of active clients. This approach leads to increased throughput in a network with primarily TCP workloads.

These findings underline the need to consider system performance in addition to studying the MAC layer in isolation. We also point out the importance of collision-resilient rate adaptation algorithms. Even with improved CW_{min} selection, a further 20% cumulative throughput gain may be possible through collision-resilient bit-rate adaptation.

Acknowledgement

We thank Dipankar Raychaudhuri and Ivan Seskar of WINLAB, Rutgers University for their valuable comments on this study. This material is based upon work supported by the National Science Foundation under Grants Nos. CNS 0335244 and CNS 0627032.

References

[1] ABI research <http://www.abiresearch.com>, 2007.
[2] A. Akella, G. Judd, S. Seshan, P. Steenkiste, Self-management in chaotic wireless deployments, in: Proceedings of ACM MOBICOM, 2005, pp. 185–199.
[3] Extricom, Inc. <http://www.extricom.com>.
[4] G. Bianchi, Performance analysis of the IEEE 802.11 distributed coordination function, IEEE JSAC 18 (3) (2000) 535–547.
[5] P. Gupta, P. Kumar, The capacity of wireless networks, IEEE Trans. Info. Theory 46 (2000) 388–404.
[6] S. Choi, K. Park, C. Kim, On the performance characteristics of WLANs: revisited, SIGMETRICS Perform. Eval. Rev. 33 (1) (2005) 97–108.
[7] T. Henderson, D. Kotz, I. Abyzov, The changing usage of a mature campus-wide wireless network, in: Proceedings of ACM MOBICOM, 2004, pp. 187–201.
[8] A. Balachandran, G.M. Voelker, P. Bahl, P.V. Rangan, Characterizing user behavior and network performance in a public wireless LAN, in: Proceedings of SIGMETRICS, 2002, pp. 195–205. URL <http://doi.acm.org/10.1145/511334.511359>.
[9] A.P. Jardosh, K.N. Ramachandran, K.C. Almeroth, E.M. Belding-Royer, Understanding congestion in IEEE 802.11b wireless networks, in: Proceedings of USENIX IMC, 2005, pp. 279–292.
[10] A.P. Jardosh, K. Mittal, K.N. Ramachandran, E.M. Belding, K.C. Almeroth, IQU: Practical queue-based user association management for WLANs, in: Proceedings of ACM MOBICOM, 2006, pp. 158–169.
[11] M.A. Ergin, K. Ramachandran, M. Gruteser, Understanding the effect of access point density on wireless LAN performance, in: Proceedings of ACM MOBICOM, 2007, pp. 350–353.
[12] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu, M. Singh, Overview of the orbit radio grid testbed for evaluation of next-generation wireless network protocols, in: Proceedings of IEEE WCNC, vol. 3, 2005, pp. 1664–1669.
[13] Scalable Network Technologies, Qualnet v.3.9.5 user's guide <http://www.qualnet.com>, May 2006.
[14] IEEE, Standard Document, IEEE Std 802.11e-2005 (Amendment to IEEE Std 802.11, 1999 Edition (Reaff 2003).
[15] MADWiFi, Multiband Atheros Driver for WiFi <http://madwifi.org>, 2007.
[16] Wireless packet traces of the 62. IETF Meeting, archived at CRAWDAD <http://crawdad.cs.dartmouth.edu/ietf2005>.
[17] Acess Point Density Traces rutgers/ap_density at CRAWDAD <http://crawdad.cs.dartmouth.edu/rutgers/ap_density>, 2007.
[18] M. Takai, J. Martin, R. Bagrodia, Effects of wireless physical layer modeling in mobile ad hoc networks, in: Proceedings of ACM MOBIHOC, 2001, pp. 87–94.
[19] H. Chang, V. Misra, D. Rubenstein, A general model and analysis of physical layer capture in 802.11 networks, in: Proceedings of IEEE INFOCOM, 2006, pp. 1–12.
[20] R. Jain, D. Chiu, W. Hawe, A quantitative measure of fairness and discrimination for resource allocation in shared computer systems, Tech. Rep., DEC TR-301, September 1984.
[21] P. Barford, M. Crovella, Generating representative web workloads for network and server performance evaluation, in: Proceedings of ACM SIGMETRICS, 1998, pp. 151–160.
[22] GNU wget file retrieval tool <http://www.gnu.org/software/wget>.
[23] R. Bruno, M. Conti, E. Gregori, Performance modelling and measurements of tcp transfer throughput in 802.11-based WLAN, in: Proceedings of ACM MSWiM, 2006, pp. 4–11.
[24] A. Botta, A. Dainotti, A. Pescape, Multi-protocol and multi-platform traffic generation and measurement, in: IEEE INFOCOM 2007 Demo Session, 2007.
[25] Y. Xiao, J. Rosdahl, Throughput and delay limits of IEEE 802.11, IEEE Commun. Lett. 6 (8) (2002) 355–357.
[26] Y. Gong, P. Marbach, Interaction of rate and medium access control in wireless networks: the single cell case, in: Proceedings of ACM MOBIHOC, 2006, pp. 178–189.
[27] ITU-T Rec. G.114 One-way Transmission Time <http://tinyurl.com/2z4co5>.
[28] Cisco Systems Inc, White Paper – A clear understanding: what makes voice over IP work? <http://tinyurl.com/3jko5>, 2002.
[29] F. Cali, M. Conti, E. Gregori, Dynamic tuning of the IEEE 802.11 protocol to achieve a theoretical throughput limit, IEEE/ACM Trans. Netw. 8 (6) (2000) 785–799.
[30] M. Heusse, F. Rousseau, R. Guillier, A. Duda, Idle sense: an optimal access method for high throughput and fairness in rate diverse wireless LANs, in: Proceedings of ACM SIGCOMM, 2005, pp. 121–132.
[31] C. Hu, J.C. Hou, A novel approach to contention control in IEEE 802.11e-operated WLANs, Tech. Rep. UIUCDCS-R-2005-2600, Univ. of Illinois at Urbana-Champaign (August 2006).
[32] IEEE, Std. 802.11h: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications: Amendment 5: Spectrum and transmit power management extensions in the 5 GHz band in europe.
[33] J. Kim, S. Kim, S. Choi, D. Qiao, CARA: Collision-aware rate adaptation for IEEE 802.11 WLANs, in: Proceedings of IEEE INFOCOM, Barcelona, Spain, 2006, pp. 1–11.
[34] K. Ramachandran, H. Kremo, M. Gruteser, P. Spasojevic, I. Seskar, Scalability analysis of rate adaptation techniques in congested IEEE 802.11 networks: An orbit testbed comparative study, in: Proceedings of IEEE WoWMoM, 2007, pp. 1–12.
[35] J.C. Bicket, Bit-rate selection in wireless networks, Master's thesis, M.I.T., Cambridge, MA, February 2005.
[36] S.H.Y. Wong, S. Lu, H. Yang, V. Bharghavan, Robust rate adaptation for 802.11 wireless networks, in: Proceedings of ACM MOBICOM, 2006, pp. 146–157.
[37] S. Pilosof, R. Ramjee, D. Raz, Y. Shavitt, P. Sinha, Understanding TCP fairness over Wireless LAN, in: Proceedings of IEEE INFOCOM, 2003, pp. 863–872.

[38] A. Vasan, R. Ramjee, T.Y.C. Woo, ECHOS - enhanced capacity 802.11 hotspots., in: Proceedings of IEEE INFOCOM, 2005, pp. 1562–1572.
[39] J. Zhu, B. Metzler, X. Guo, Y. Liu, Adaptive csma for scalable network capacity in high-density WLAN: A hardware prototyping approach, in: Proceedings of IEEE INFOCOM, 2006, pp. 1–10.
[40] T.-S. Kim, J.C. Hou, H. Lim, Improving spatial reuse through tuning transmit power, carrier sense threshold, and data rate in multihop wireless networks, in: Proceedings of ACM MOBICOM, 2006, pp. 366–377.
[41] A. Mishra, V. Shrivastava, D. Agarwal, S. Banerjee, S. Ganguly, Distributed channel management in uncoordinated wireless environments, in: Proceedings of ACM MOBICOM, 2006, pp. 170–181.
[42] A. Mishra, V. Brik, S. Banerjee, A. Srinivasan, W. Arbaugh, A client-driven approach for channel management in wireless LANs, in: Proceedings of IEEE INFOCOM, 2006, pp. 1–12.
[43] Y. Bejerano, S.-J. Han, L.E. Li, Fairness and load balancing in wireless LANs using association control, in: Proceedings of ACM MOBICOM, 2004, pp. 315–329.
[44] F. Guo, T. cker Chiueh, Scalable and robust WLAN connectivity using access point array, in: Proceedings of DSN 2005, 2005, pp. 288–297.
[45] Aruba Networks, Inc. <http://www.arubanetworks.com>.
[46] Cisco Systems, Inc. <http://www.cisco.com>.
[47] Trapeze Networks <http://www.trapezenetworks.com>.

Mesut Ali Ergin (ergin@winlab.rutgers.edu) is a Ph.D. candidate in the Department of Electrical and Computer Engineering at Rutgers, the State University of New Jersey, USA and works as a graduate assistant in Wireless Information Networks Laboratory (WINLAB) since 2004. He received an MS degree in Computer Engineering from Bogazici University, Turkey, and a BS degree in Control and Computer Engineering from Istanbul Technical University, Turkey, in 2003 and 1999, respectively. His research interests include design and evaluation of wireless and mobile systems for realistic future usage scenarios. He is a member of the IEEE Communication Society.

Kishore Ramachandran (kishore@winlab.rutgers.edu) is a Ph.D. Candidate in the Department of Electrical and Computer Engineering at Rutgers University and a graduate research assistant at WINLAB. His research interests include wireless networks and mobile computing with an emphasis on radio resource management, network scalability and measurement. He received the M.S. degree in Computer and Systems Engineering from Rensselaer Polytechnic Institute (RPI), NY. and the B.E. degree in Telecommunications Engineering from the University of Mumbai, India, in 2002 and 2000, respectively.

Marco Gruteser (gruteser@winlab.rutgers.edu) is an assistant professor at WINLAB, Rutgers University. His research interests include location-aware networking and building reliable, secure, and privacy-aware communication systems for vehicular networks. He received his Ph.D. in computer science from the University of Colorado at Boulder. He is a member of the IEEE Computer Society and the ACM.