

On the Performance Characteristics of WLANs: Revisited

Sunwoong Choi
School of Computer Science
and Engineering
Seoul National University
Seoul, Korea 151-744
schoi@popeye.snu.ac.kr

Kihong Park
Dept. of Computer Sciences
Purdue University
West Lafayette, IN 47907
U.S.A.
park@cs.purdue.edu

Chong-kwon Kim
School of Computer Science
and Engineering
Seoul National University
Seoul, Korea 151-744
ckim@popeye.snu.ac.kr

ABSTRACT

Wide-spread deployment of infrastructure WLANs has made Wi-Fi an integral part of today's Internet access technology. Despite its crucial role in affecting end-to-end performance, past research has focused on MAC protocol enhancement, analysis and simulation-based performance evaluation without sufficient consideration for modeling inaccuracies stemming from inter-layer dependencies, including physical layer diversity, that significantly impact performance. We take a fresh look at IEEE 802.11 WLANs, and using a combination of experiment, simulation, and analysis demonstrate its surprisingly agile performance traits. Our main findings are two-fold. First, contention-based MAC throughput degrades gracefully under congested conditions, enabled by physical layer channel diversity that reduces the effective level of MAC contention. In contrast, fairness and jitter significantly degrade at a critical offered load. This duality obviates the need for link layer flow control for throughput improvement but necessitates traffic control for fairness and QoS. Second, TCP-over-WLAN achieves high throughput commensurate with that of wireline TCP under saturated conditions, challenging the widely held perception that TCP throughput fares poorly over WLANs when subject to heavy contention. We show that TCP-over-WLAN prowess is facilitated by the self-regulating actions of DCF and TCP congestion control that jointly drive the shared physical channel at an effective load of 2–3 wireless stations, even when the number of active stations is very large. Our results highlight subtle inter-layer dependencies including the mitigating influence of TCP-over-WLAN on dynamic rate shifting.

Categories and Subject Descriptors

C.2.1 [Computer Communication Networks]: Network Architecture and Design—*Wireless communication*

General Terms

Performance, Experimentation

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGMETRICS'05, June 6–10, 2005, Banff, Alberta, Canada.
Copyright 2005 ACM 1-59593-022-1/05/0006 ...\$5.00.

Keywords

DCF performance, physical layer diversity, TCP-over-WLAN performance, inter-layer dependence, rate control

1. INTRODUCTION

1.1 Background and Motivation

Wide-spread deployment of wireless local area networks (WLANs) at campuses, enterprises, residential areas, and commercial hotspots has made IEEE 802.11 WLANs [3], also referred to as Wi-Fi when complying with industry interoperability standards, a key component of today's integrated wireline/wireless Internet. As with ALOHA [4] and Ethernet [24] in the 1970s, the advent of WLANs in the late 1990s has spurred research aimed at understanding its properties, including performance analyses of IEEE 802.11 CSMA/CA [7, 10], fair scheduling [23, 33], capacity analyses of ad hoc WLANs [12, 13] (and hybrid extensions [21, 22]), and measurement-oriented studies [5, 20, 30]. The capacity of CSMA/CD Ethernet was not without controversy—Boggs et al. [8] discuss the wide ranging performance perceptions of the DIX and IEEE 802.3 10Base5 Ethernets deployed in the 1980s—but the subsequent evolution to switched technology, with CSMA/CD preserved for backward compatibility, rendered many of the earlier issues moot. This is not the case for IEEE 802.3 WLANs as physical shielding of the wireless transmission medium from interference and localization of contention at a switch are infeasible. WLAN performance is influenced by three main variables: mobility, channel noise, and multiple access contention. Mobility across access points (APs) in an extended service set is facilitated by reassociation-based handoff augmented by legacy link layer forwarding over wireline distribution networks. Mobility support across extended service sets through Mobile IP, at the present, is rarely instituted—even within a single organization—due to security concerns and insufficient application demand. Mobility support across domains is hindered by administrative boundaries and policy barriers. Channel noise in the unlicensed 2.4/5 GHz ISM/U-NNI bands can be a problem, especially when the distance between wireless station and AP is far. Current remedies include multi-rate support in physical layer modulation, with adaptive rate selection left to vendor discretion.

The focus of this paper is on WLAN performance in hot spots—the dominant mode of WLAN usage today—where performance degradation from contention-based multiple access is the key concern. In several respects, there is ambiguity about the performance of WLANs today as there was

about the capacity of Ethernet in the 1980s. For example, in [7] it is indicated that IEEE 802.11 DCF system throughput significantly declines as offered load is increased. The analysis and simulation, however, do not consider the effect of physical layer channel diversity (also called multiuser diversity [19, 32]) that mitigates throughput degradation. In [34] it is shown that 802.11b throughput decreases from above 6 Mbps to below 2 Mbps as the number of stations is increased from 1 to 14. The performance results are stated as validating the 802.11 capacity analysis in [9]. The results in [34] are difficult to replicate unless dynamic rate shifting is activated which easily confuses collision with channel noise. In [35, 36] it is indicated that forward TCP data traffic can collide with reverse acknowledgment traffic, “dramatically” increasing the frame error rate. In [14], an enhanced DCF protocol called DCF+ is proposed to address this problem. We show that this potential TCP data/ack collision problem does not materialize due to the self-regulating actions of DCF and TCP congestion control. One of the goals of this paper is to help clarify some of the ambiguities surrounding WLAN performance through delineation of wireline and wireless features, incorporation of the impact of inter-layer dependencies, including those stemming from physical layer channel diversity, that impact performance.

1.2 New Contribution

The contribution of this paper is two-fold, discussed in two parts: DCF MAC layer performance and TCP-over-WLAN performance. First, we show that contention-based DCF throughput degrades gracefully as offered load or the number of wireless stations is increased. This is enabled by physical layer channel diversity that reduces the effective level of MAC layer contention, a form of multiuser diversity whose persistent manifestation in WLANs DCF is able to exploit. This obviates the need for recently proposed link layer flow control schemes [16, 37] aimed at preventing MAC throughput degradation. We also show that evidence of drastic throughput degradation of IEEE 802.11b WLANs under moderate load [15, 34] is likely to stem from the influence of dynamic rate shifting implemented in most WLAN cards, which suffer from the problem of not being able to effectively distinguish collision from channel noise. Downshifting to fallback rates (e.g., 5.5, 2 and 1 Mbps in 802.11b) helps when increased distance or channel noise results in small SNR, but is ineffective—in fact, detrimental—when collisions are the primary cause of frame errors. In contrast to throughput, MAC layer fairness and jitter significantly degrade at a critical offered load, which can benefit from traffic controls aimed at operating the system outside the saturation region.

Second, we show that TCP-over-WLAN achieves high throughput commensurate with that of wireline TCP, even under saturated conditions where the number of wireless stations is very large. We show that TCP-over-WLAN prowess is effected by the self-regulating actions of DCF and TCP congestion control, which jointly drive the shared baseband medium at an effective load of 2–3 wireless stations. The TCP data/ack collision problem studied in [14, 35, 36] does not manifest as a real performance concern. We evaluate TCP-over-WLAN performance by analyzing a Markov chain over the state space counting the number of backlogged stations. We show a strong negative drift in the Markov chain that is established by inferring transition rates from mea-

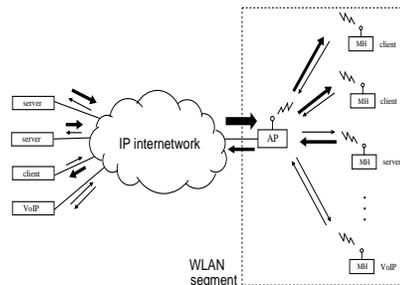
surement traces, both from experiment and simulation. We identify situations—analogue to wireline TCP with mismatched low bandwidth access links [28]—where buffer overflow at the gateway or AP may become a dominant factor.

The remainder of the paper is organized as follows. In the next section, we give the technical set-up. This is followed by Section 3 which discusses the performance properties of IEEE 802.11b DCF. Section 4 studies TCP-over-WLAN performance. We conclude with a discussion of our results.

2. SYSTEM MODEL

2.1 Infrastructure WLAN Environment

We consider an integrated wireless/wireline IP internetwork where access networks are comprised of IEEE 802.11 infrastructure WLANs that service wireless hosts. An access point (AP) may be connected to other APs by a wireline distribution network—typically one or more Ethernet switches—that is then connected to the Internet via a wireline access link. Figure 1(a) depicts the components of such a network system. Our focus is on the wireless access seg-



(a)



(b)

Figure 1: (a) Wireless/wireline IP network with IEEE 802.11 infrastructure WLAN segment. (b) iPAQ pocket PCs and Enterasys RoamAbout R2 AP forming BSS.

ment representative of hot spots deployed at coffee houses, airport lounges, SOHOs, and residential home networks. Typical last mile access technology include broadband cable/DSL and T1/T3 lines for small-to-medium enterprises. We model the wireless/wireline network system as a dumb-bell network where n wireless hosts access the wireline network through an AP—collectively forming a basic service

set (BSS)—that is connected to a router that, in turn, connects n wireline hosts. A canonical configuration treats the wireless hosts as clients that communicate with servers on the wireline side. In *ns-2* simulations, the dumbbell configuration also serves as the physical topology.

2.2 Experimental Set-up

We connect an Enterasys RoamAbout R2 AP supporting IEEE 802.11b through a 100 Mbps Ethernet crossover cable to a Dell Inspiron 8100 laptop with a 1 GHz Mobile Pentium III processor and 512 MB memory running Linux 2.4.7-10. We also performed benchmark experiments on a private IP-over-SONET testbed comprised of 9 Cisco 7206VXR routers—the wireline backbone segment—where the AP is directly connected to one of the routers, with PCs connected to the other 8 routers acting as servers. The two set-ups yield similar results and we report benchmark suites from the laptop-to-AP wireline configuration.

The BSS wireless segment is populated by 16 Compaq/HP 3800 and two 3600 series iPAQ pocket PCs with 64 MB memory (Figure 1(b)) running Familiar Linux v0.7.2. Each pocket PC is plugged into an external PCMCIA jacket which connects an Enterasys RoamAbout 802.11b PC Card. The data rate on the AP and pocket PCs is set to 11 Mbps, RTS/CTS and power control are disabled, and the channel is fixed at an unoccupied frequency band with minimal cross-channel interference. We use WildPackets AiroPeek NX [1], a state-of-the-art commercial WLAN sniffing and analysis tool, to determine the state of the wireless medium for channel selection and monitoring of MAC-level transmission activity. The latter, in conjunction with tcpdump traces collected at the laptop and pocket PCs, is used to infer the backlog Markov chain associated with TCP-over-WLAN dynamics.

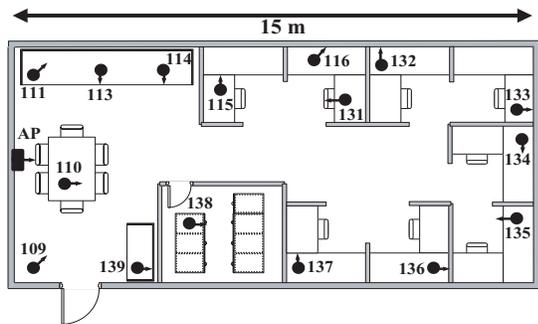


Figure 2: Basement indoor office environment showing locations of AP and wireless stations.

The experiments in this paper focus on an indoor office environment (Figure 2) in the basement of the Computer Science Building at Purdue University. We also performed benchmarks in a table top configuration where all pocket PCs are co-located on a conference table, which minimizes the influence of distance on signal strength and resultant throughput. We use the basement corridor configuration for distance-based experiments in an indoor setting.

2.3 Simulation Set-up

We use *ns* [2] (version 2.27) with CMU’s WLAN extension to simulate the dumbbell wireless/wireline WLAN topology corresponding to Figure 1(a), where the IP network

cloud is replaced by a single router. Wireline bandwidth is set to 100 Mbps, with link latency between router and AP set to 1 msec and the link latency between wireline hosts and router set to 10 msec. The IEEE 802.11 BSS data rate is 11 Mbps and ACK frames are transmitted at 2 Mbps. Default channel noise in the form of BER is set at 10^{-6} . Other IEEE 802.11b related parameter specifications are summarized in Table 1. An upper bound on the

slot time	20 μ s
SIFS	10 μ s
DIFS	50 μ s
CWmin	31
CWmax	1023
physical preamble	144 bits
physical header	48 bits
ACK frame	112 bits
MAC header + CRC	224 bits

Table 1: IEEE 802.11 DCF MAC parameters

maximum achievable throughput by a single wireless station may be computed based on these parameters, which yields a frame completion time of $50 + T_{BO} + 192 + (224 + 8\ell)/11 + 10 + 192 + 112/2$ (μ s) where ℓ is the frame size and T_{BO} is the length of the backoff interval. If the backoff interval is ignored, it takes about 1.25 msec to transmit a frame of size $\ell = 1000$ bytes, which implies a MAC throughput bound of 6.4 Mbps. The bound varies as a function of frame size ℓ . The theoretically expected single station throughput, considering the backoff counter, is 5.136 Mbps. Simulated throughput is 5.117 Mbps, and in iPAQ experiments the single station throughput is 5.155 Mbps. When 2–4 wireless stations are present, the total throughput, despite collisions, can be higher than that of a single station due to a concurrent countdown effect of DCF that leaves the channel less idle on average.

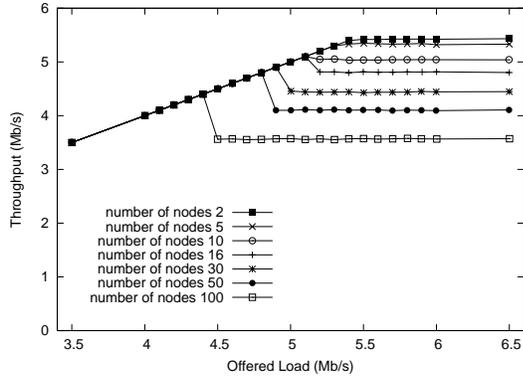
3. CHARACTERISTICS OF IEEE 802.11 DCF PERFORMANCE

3.1 DCF Throughput: Agility

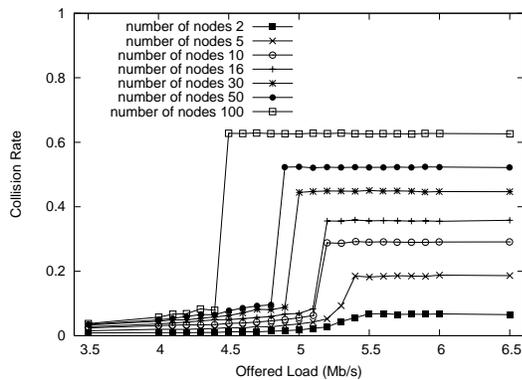
The first issue concerns the throughput of DCF which has been observed to degrade significantly as the contention level—offered load or number of stations—is increased. We will start with a standard simulation-based performance evaluation, then incorporate the impact of physical layer channel diversity and dynamic code rate shifting.

3.1.1 Equidistant DCF Throughput: Simulation

We consider a BSS configuration where wireless nodes are symmetrically placed on a circle of radius 10 m with the AP located at the center. Equidistance and symmetry affect maximal DCF contention. Single point location—a logical configuration commonly used in contention-based MAC studies where all stations, including the AP, are co-located at a single point, random location, and other layout configurations are considered in Section 3.1.3. Figure 3(a) shows simulated DCF throughput as a function of offered load for 2–100 wireless stations. Offered load is constant bit rate (CBR) traffic with a small uniformly random inter-packet noise added to break up deterministic synchronization effects. First, we note that the 2 and 5 station throughput



(a)



(b)

Figure 3: (a) Simulated IEEE 802.11 DCF throughput as a function of offered load for 2–100 wireless stations. (b) Corresponding DCF collision rate.

is higher than the single station throughput of 5.117 Mbps due to the concurrent countdown effect. As offered load is increased, throughput grows linearly until a saturation point at which throughput ceases to increase; in fact, it may decline. Onset of saturation throughput occurs after a critical offered load where the collision rate increases sharply as shown in Figure 3(b). When the number of stations is large, both peak and saturation throughput decrease and their relative gap widens. The drop in peak and saturation throughput as a function of the number of wireless stations is shown in Figure 4. Throughput decline is overall gradual with saturation throughput exhibiting a faster drop. At a moderate load of 16 wireless nodes, peak throughput decreases 5.9% from the throughput level of 2 wireless stations and saturation throughput decreases by 11.4%. At a heavy load of 50 stations, the reductions are 11.4% and 24.2%, respectively.

Peak throughput—vis-à-vis saturation throughput—is relevant for two reasons. In [7] it is argued that saturation throughput should be considered as the attainable equilibrium performance measure. We find that under physical layer channel diversity, the gap between peak and satura-

tion throughput shrinks markedly, and when TCP operates over WLAN the joint actions of TCP congestion control and DCF drive the baseband channel near the peak throughput level. The first point is discussed next.

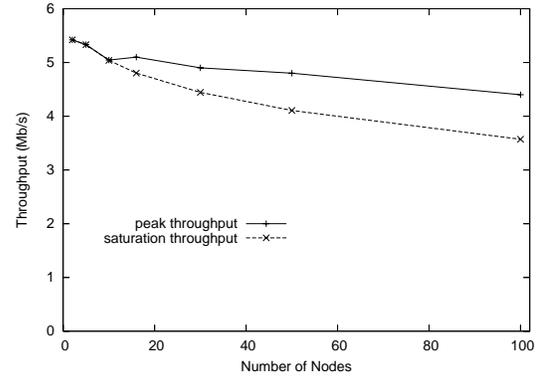


Figure 4: Decrease in DCF peak and saturation throughput as the number of nodes is increased.

3.1.2 Physical Layer Diversity: Experiment

Physical layer channel diversity—a form of multiuser diversity [19, 32]—plays a significant role in mitigating DCF throughput degradation under heavily contented conditions. Figure 5 shows measured throughput of 802.11 DCF in the indoor office environment depicted in Figure 2, as offered load is increased from 4 Mbps to 7 Mbps with 2–16 pocket PCs competing for the shared channel. We observe two differences when compared to equidistant simulated DCF throughput in Figure 3(a): the overall throughput for 2, 5, 10 and 16 wireless stations is higher—5.7–5.876 Mbps (experiment) vs. 4.816–5.432 Mbps (simulation)—and the gap between peak and saturation throughput is much narrower. For 16 stations, the gap is 5.9% in simulation versus essentially 0% in experiment.

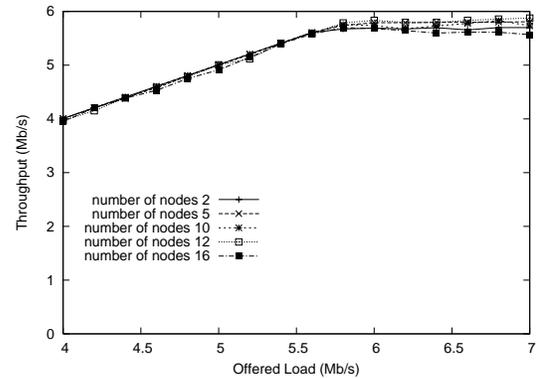


Figure 5: Empirical IEEE 802.11 DCF throughput as a function of offered load for 2, 5, 10, 12, 16 wireless stations in indoor office environment.

Figure 6 compares the saturation throughput of DCF as the number of wireless stations is increased for the indoor office experiment, indoor table top experiment, and equidistant circle simulation. For 10 or 16 stations, we observe a 1 Mbps (20%) difference in achieved throughput between

simulation and indoor office experiment, indicating the influence of physical layer channel diversity. We note that when only a single wireless station is accessing the channel, saturation throughput, for both simulation and experiment, is in the 5.1–5.2 Mbps range.

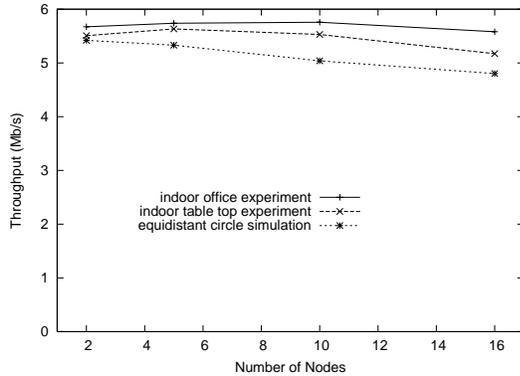


Figure 6: Comparison of 802.11 DCF saturation throughput as a function of the number of wireless stations for indoor office experiment, indoor table top experiment, and equidistant circle simulation.

In general, multiuser diversity may result from mobility and time-varying fading effects, and a scheduler that wishes to harness channel diversity for throughput maximization would try to allocate the channel at a given moment in time to the user with the best channel condition. In infrastructure WLAN hot spots where users, for the most part, change their location infrequently, physical layer channel diversity manifests in a persistent, location dependent manner. That is, “advantaged” stations continue to be advantaged, and “disadvantaged” stations remain disadvantaged. Figure 7 shows the throughput share of 8 wireless stations under different offered loads in the indoor office environment. We observe that biases in channel diversity are preserved which translates to corresponding biases in individual throughput share. The bias is solely location dependent—e.g., switching the pocket PCs makes little difference—determined by the geometry of radio propagation in indoor environments. It is well-known that signal strength distribution in closed spaces

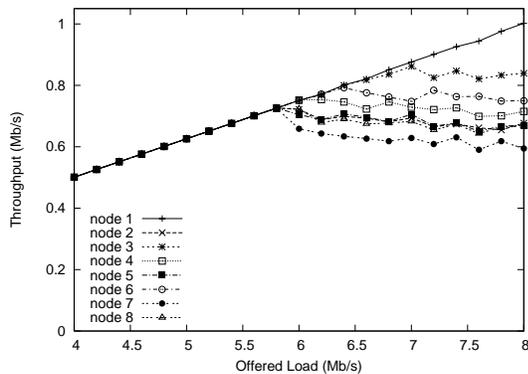


Figure 7: Persistent stratification of throughput share in indoor office WLAN experiment with 8 iPAQ pocket PCs.

is varied [11, 17, 31], whose theoretical foundation with respect to persistence, sensitivity, and high variability may be found in the chaotic nature of radio wave propagation in cavities [29]. DCF exploits this physical layer channel diversity in two ways: one, by a simple capture effect where a collision between two frames of different strengths may result in a successful decoding of the dominant frame due to the signal differential, and two, by subsequent exponential backoff of the weaker station which amplifies the access priority that the stronger station receives. In Figure 7, we observe that as collision rate increases with increased offered load, the throughput share of the strongest station increases due to the reinforcing action of DCF. The DCF amplification effect is reminiscent of “the rich get richer and the poor get poorer” dynamics.

From a system throughput perspective, as DCF “schedules” stations with stronger signal strength more frequently, system throughput improves by the maxim of multiuser diversity scheduling [12]. Figure 8 shows DCF throughput as a function of offered load from the indoor table top experiment. The less varied physical layer channel diversity resulting from close station proximity effects a smaller throughput

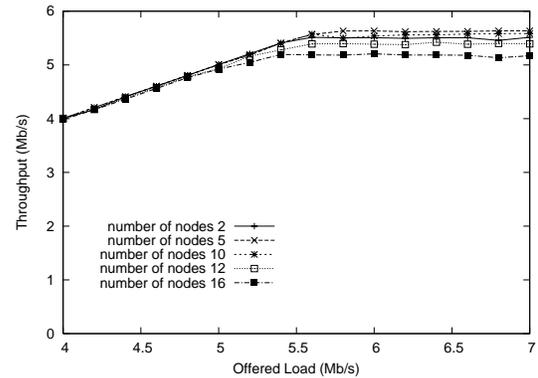


Figure 8: Empirical IEEE 802.11 DCF throughput as a function of offered load for 2, 5, 10, 12, 16 wireless stations in indoor table top environment.

gain vis-à-vis equidistance circle simulation than the indoor office configuration. The quantitative influence of multiuser diversity on DCF throughput gain as a function of the number of wireless stations is indicated in Figure 6, where table top throughput is sandwiched between equidistance circle simulation and indoor office experiment.

3.1.3 Non-equidistant DCF Throughput: Simulation

In the CMU wireless extension of *ns-2*, capture effect may be considered by an application of SIR thresholding. We consider four additional BSS configurations where distance-based capture effect may be isolated and discerned: “line” where all stations are aligned 1 m apart on a line with the AP at one end, “near random” where stations are placed uniformly randomly inside a disk of radius 10 m with the AP at the center, “far random” where the AP is not at the center of the random disk but outside at a distance of 20 m from the center, and “single point” where all nodes, including the AP, are co-located at a single point. Figure 9 shows simulated DCF throughput as a function of the number of stations for the five configurations. We observe that the line configuration achieves the highest throughput, closely followed by the

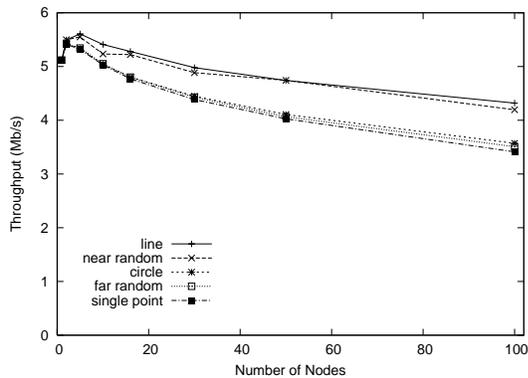


Figure 9: Simulated DCF throughput as a function of the number of stations for line, near random, far random, single point, and circle configurations.

near random configuration. This is expected since channel diversity—due to path loss modeled by distance—is highest for the line configuration followed by the near random configuration. In the far random configuration, which is in distant fourth place, relative channel diversity inside the disk is dampened by the far distance of the wireless stations from the AP. The circle and single point configurations achieve a similar throughput as the far random configuration due to their lack of diversity and resultant unbiased DCF multiple access competition.

3.1.4 Dynamic Rate Shifting: Experiment

In [34] it is shown that empirical IEEE 802.11b throughput in an infrastructure WLAN decreases from above 6 Mbps to below 2 Mbps as the number of stations is increased from 1 to 14. Perception of significant throughput degradation at moderate load is wide-spread, for example, in [15] it is claimed that “as the number of contending stations increases, aggregate capacity drops precipitously (to less than 1 Mb/s with 10 contending stations)” which is blamed on the multiple access nature of IEEE 802.11 CSMA/CA. We do not doubt the validity of the empirical observations—we believe they are factually accurate—but we question the conclusions drawn from the observations, in particular, the assignment of responsibility to DCF.

Figure 10 shows IEEE 802.11 DCF throughput as a function of the number of stations from an indoor table top experiment for different data rates of 802.11b—11 Mbps, 5.5 Mbps, 2 Mbps and 1 Mbps—and dynamic rate shifting (“auto rate”) implemented in Enterasys WLAN cards. At fixed data rates, the throughput curves are relatively flat as the number of pocket PCs is increased from 2 to 18, consistent with the graded throughput decline seen in Figure 6. When auto rate is enabled—the default mode in most WLAN cards—then aggregate throughput declines drastically as seen in Figure 10, reaching 1 Mbps at 16 wireless stations. The prominent throughput decline, however, is not due to DCF. Testing of several IEEE 802.11b vendor cards reveals that the dynamic rate shifting algorithms implemented in WLAN cards—most are realized in firmware—have difficulty distinguishing collision from channel noise. For example, in the Enterasys RoamAbout 802.11 DS High Rate card (Orinoco chipset), down shifting is triggered by 2 consecutive failures to receive an 802.11 ACK frame. The

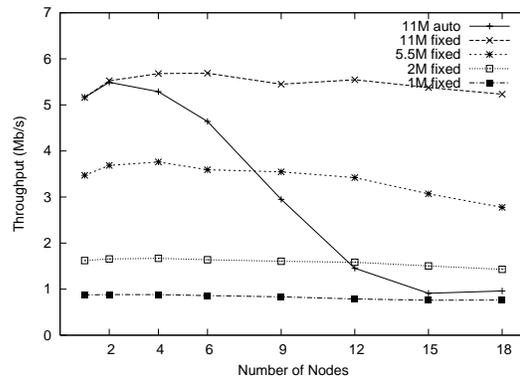


Figure 10: Empirical IEEE 802.11 DCF throughput as a function of the number of pocket PCs for auto rate and fixed data rates 11 Mbps, 5.5 Mbps, 2 Mbps and 1 Mbps in indoor table top environment.

transmitter, therefore, cannot know if the data frame was corrupted by channel noise or collision. The rudimentary nature of vendor-specific dynamic rate shifting schemes—rate control is not part of the IEEE 802.11 standard—has the detrimental effect of significantly reducing throughput under moderate contention, even when channel noise is small.

Given that dynamic rate shifting has an undesirable side effect of significantly reducing WLAN throughput, how well does it counter channel noise? Figure 11 shows the throughput of a single pocket PC—the influence of collision is factored out—benchmarked at different locations in the basement of the CS Building along a rectangular corridor for fixed data rates 11 Mbps, 5.5 Mbps, 2 Mbps, 1 Mbps, and auto rate. The spacings are approximately 7 m, with PT1

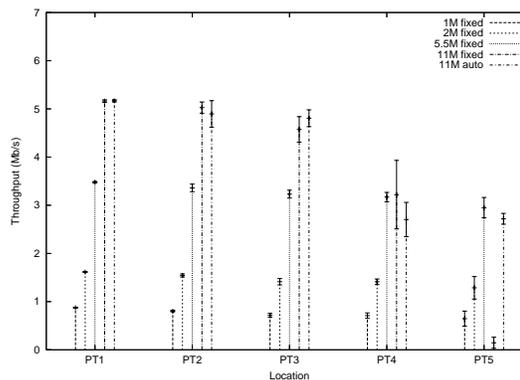


Figure 11: Empirical 802.11 DCF throughput of a single pocket PC at different locations along a rectangular corridor in the basement of the CS Building.

being the closest location to the AP and PT5 the farthest, following two 90 degree bends in the corridor (Figure 12). We observe that until location PT4, 11 Mbps data rate yields the highest throughput, matched closely by 5.5 Mbps data rate at PT4. At location PT5 inversion takes place and 11 Mbps data rate attains the least throughput. Throughput, auto rate achieves throughput comparable to that of the highest fixed data rate indicating that it is able to adapt to varying channel noise in the absence of collision. In the presence of

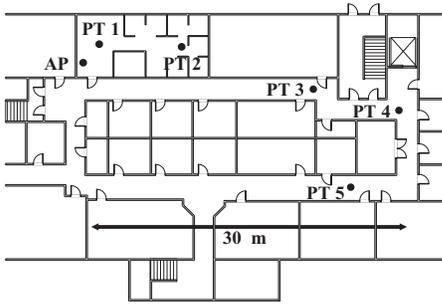


Figure 12: Basement corridor environment with single pocket PC positioned at locations PT1–PT5.

collision, even at moderate loads of 12–18 stations, we find that dynamic rate shifting may do more harm than good in indoor environments.

3.2 DCF Fairness and Jitter

In this section, we discuss the sensitive dependence of DCF fairness and jitter on offered load, and issues related to its control.

3.2.1 DCF Fairness: Diversity Amplification

Figure 13 shows simulated fairness, presented as the ratio of minimum throughput over maximum throughput across participating stations, as a function of offered load for 2–100 stations in the equidistant circle configuration. We observe a sharp degradation of fairness at a critical offered load commensurate with a sharp increase in the collision rate (cf. Figure 3(b)). Unlike throughput, however, which

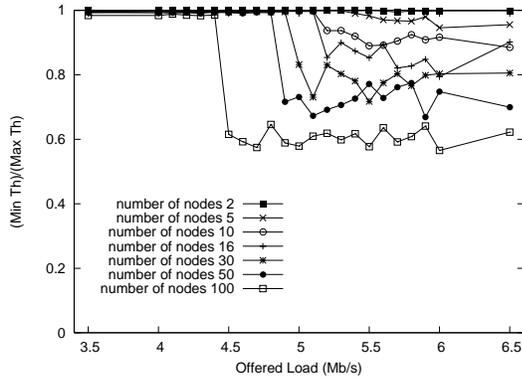


Figure 13: Simulated DCF fairness as a function of offered load for 2–100 wireless stations in equidistant circle configuration.

is mitigated by physical layer channel diversity, unfairness is amplified. This is clearly discernible on a per-flow basis in Figure 14 where the individual throughput share of 16 pocket PCs in the indoor office environment is shown as a function of offered load. Up to a critical offered load near 5.5 Mbps, throughput share is overall equitable. After the critical load where collision rate sharply increases, individual throughput share stratifies into persistent levels (e.g., five for 8 Mbps) whose total width—the minimum throughput over maximum throughput ratio—continues to expand. DCF exerts a self-reinforcing “the rich get richer the poor

get poorer” effect that amplifies unfairness triggered by multi-user diversity and capture effect at high loads.

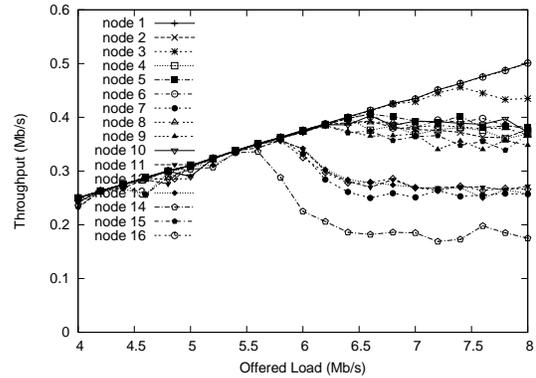


Figure 14: Empirical DCF fairness with respect to individual throughput share as a function of offered load for 16 iPAQs in indoor office environment.

3.2.2 DCF Jitter: Sensitivity

We consider jitter—in the sense of throughput variation over time—at different offered loads. We use standard deviation to capture throughput variability in time. Figure 15 shows average standard deviation as a function of offered load for 2–100 stations in the equidistant circle configuration. We observe that at the critical load where collision rate sharply increases, jitter exhibits a sudden jump. Operating an infrastructure WLAN in the saturation regime need not significantly degrade throughput, however, it significantly increases throughput variability which has direct bearing on VoIP (voice-over-IP) and multimedia streaming applications. This is confirmed in experiments in the indoor office environment.

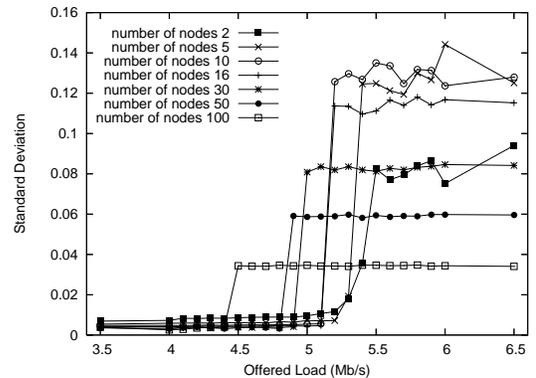


Figure 15: Simulated DCF jitter performance—captured as standard deviation of throughput—for the equidistant circle configuration.

3.2.3 Link Layer Rate Control

In [16, 37] link layer rate control is proposed to mitigate CSMA/CA throughput decline under excessive offered load. Our results show that unless the number of wireless stations in a BSS is very large, DCF throughput degrades gracefully,

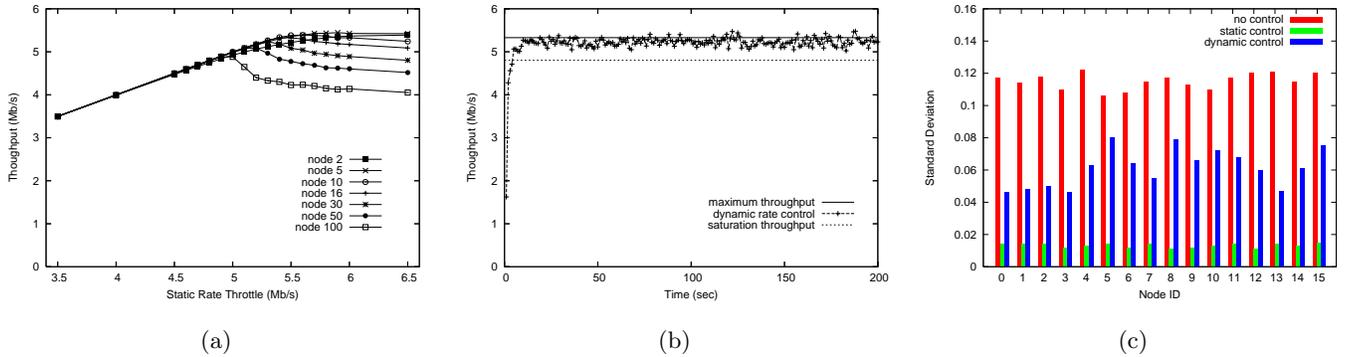


Figure 16: (a) Throughput profile under static rate throttling. (b) Throughput evolution under dynamic rate control. (c) Throughput jitter: no control, dynamic control, and static control.

aided by DCF’s ability to exploit physical layer channel diversity. Fairness and jitter, however, suffer significantly under excessive offered load, and rate controls that limit traffic impinging on the 802.11 MAC can improve fairness and jitter. Viewing the “dome shaped” static input-output profile of IEEE 802.11 DCF (cf. Figure 3(a)) as a feedback control problem, it is not difficult to design a PI controller that is asymptotically stable and efficient. The well-known instability problem of ALOHA (and ALOHA-like) contention-based protocols [18, 27] applies only in the queuing sense—backlog growing unboundedly if the application layer arrival rate exceeds the link layer service rate—but not for maximizing the link layer’s service rate which is a function of the controlled link layer arrival rate.

Figure 16(a) shows DCF throughput under static rate throttling for the set-up corresponding to Figure 3(a). Compared to the latter, both the onset of critical offered load and the magnitude of throughput drop are improved which stems from the smoothing effect of the rate throttle. Figure 16(b) shows throughput evolution under dynamic rate control with 16 wireless stations where local control implemented by each station is geared toward throughput maximization: improvement in jitter and fairness may result as an indirect consequence of well-behavedness of jitter and fairness outside the saturation region. We observe that dynamic rate control is able to approximate the maximum statically throttled throughput level. Figure 16(c) compares throughput jitter—captured as throughput standard deviation—for no control, dynamic control, and static control. Throughput jitter is improved by dynamic rate control vis-à-vis no control, however, it falls short of the small jitter level achieved by static rate control. This stems from the adaptivity and throughput-centric nature of the implemented feedback rate control which causes recurrent excursions into the saturation region when reaching for peak throughput. By sacrificing a little throughput, jitter and fairness can be further improved without transitioning to full-fledged jitter- and fairness-specific traffic controls. Also, with additional control messaging overhead, distributed fairness control can affect rate-based WFQ emulation with given service weights.

There are open-loop applications such as VoIP for which link layer rate shaping is directly harmful. For closed-loop, elastic applications such as TCP-based file transfer—the

dominant constituent of Internet traffic which has also been confirmed for WLAN access [5, 20]—link layer rate control may be construed beneficial, especially due to the increased collision rate associated with the saturation regime (cf. Figure 3(b)) that may translate to increased frame error rate, also evident from increased throughput jitter (Figure 16(c)). The VoIP bandwidth multiplexing problem may be addressed using IEEE 802.11e and rate-based WFQ emulation, among other control options. We established that dynamic rate shifting implemented in WLAN cards easily confuse collision with channel noise leading to significant throughput degradation. In the case of TCP, known to be sensitive to spurious packet loss, this may prove to be problematic. Interestingly TCP-over-WLAN performance does not suffer under this problem, achieving performance on par with wireline TCP. The reasons underlying this surprising agility is the subject of the next section.

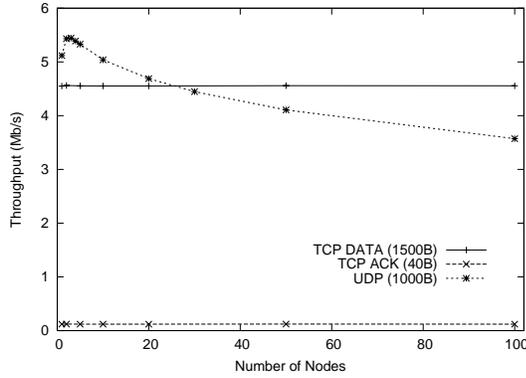
4. TCP-OVER-WLAN PERFORMANCE

We first discuss throughput performance of TCP-over-WLAN, followed by explanation of its prowess and mitigating influence on dynamic rate shifting. We use simulation to study the details of TCP-over-WLAN dynamics, augmented by experimental benchmarks.

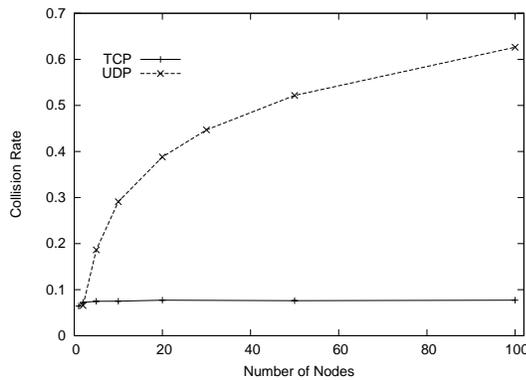
4.1 Scalable Throughput

It is known from steady-state TCP analysis [25] that TCP throughput depends polynomially on packet loss rate (i.e., $\propto p^{-1/2}$) which renders TCP sensitive to spurious packet loss. An important focus of past TCP-over-Wireless research [6, 26] has been distinguishing channel noise from congestion, so that TCP’s congestion control may be desensitized against random packet loss. Pronounced collision, a form of congestion on the wireless segment, can in indoor environments dominate channel noise—the reason why WLAN throughput, under the action of dynamic rate shifting, drastically degrades when subject to moderate multiple access contention (cf. Figure 10)—and significantly diminish TCP throughput.

To evaluate TCP-over-WLAN performance under multiple access contention and resultant collision, we use the single point simulation configuration which suffers under the heaviest collision and throughput drop (Figure 9), providing



(a)



(b)

Figure 17: (a) TCP-over-WLAN throughput and ACK traffic as a function of the number of wireless stations; UDP throughput is shown for comparative purposes. (b) Corresponding collision rate.

maximal challenge to TCP. We consider a client/server environment where wireless stations, acting as clients, fetch files from wireline servers. This incorporates the TCP data/ack collision problem [14, 35]. We use TCP NewReno without selective and delayed ACK, and consider infinite source sessions in steady-state. Figure 17(a) shows TCP throughput as a function of the number of wireless stations. As multiple access contention is increased from 2 to 100, TCP-over-WLAN throughput remains flat. The same goes for TCP ACK traffic which is at the same packet rate (pps) but smaller data rate (bps). In contrast, UDP (i.e., DCF) throughput declines under high contention. Figure 17(b) shows the corresponding collision rate experienced by TCP data and ACK frames at the DCF MAC layer which also remains essentially flat. We use a MSS of 1500 bytes for the above benchmarks but the same qualitative profile holds for other frame sizes. The AP has a buffer size of 200 (in frame units), a capacity in the same range as the Enterasys Roam-About R2 AP.

Figure 18 shows frame error rate, collision rate, channel error rate, frame discard rate, and AP buffer drop rate as

the number of contending stations is increased. At the given BER level 10^{-6} , collisions make up the bulk of frame errors, however, no frames are discarded by the 802.11 MAC layer due to DCF’s ARQ which performs 7 retransmissions under exponential backoff before giving up. The loss rate at the AP shows that it is downstream buffer overflow at the AP—a classical wireline bottleneck that exists in other pure wireline contexts such as low bandwidth access links [28]—

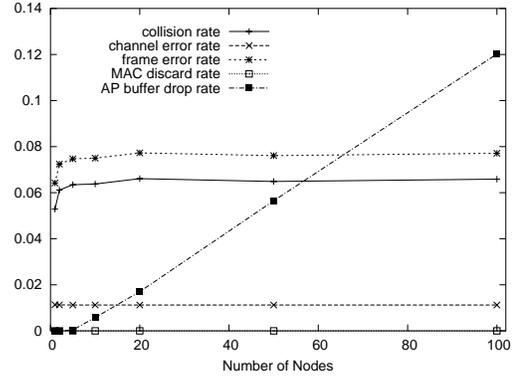


Figure 18: Frame error rate, collision rate, channel error rate, frame discard rate, and AP buffer drop rate as a function of the number of stations.

that throttles the application layer data rate. The flat TCP throughput curve is caused by the flat WLAN collision rate. Collision rate stays flat even if the AP is endowed with infinite buffer capacity. TCP-over-WLAN is able to operate at a small, constant collision rate independent of the number of TCP flows competing for the shared wireless baseband bandwidth. The TCP data/ack collision problem raised in [14, 35] does not materialize. We refer to this as TCP-over-WLAN’s scalable throughput property.

4.2 Efficiency: UDP Emulation

To evaluate how efficiently TCP utilizes WLAN capacity in the sense of “DCF throughput vs. offered load,” we emulate TCP using UDP in a rate-matched, open-loop fashion where n downstream UDP sessions transmit 1500 byte data frames at λ pps that is matched by n upstream UDP flows. Figure 19 shows downstream and upstream UDP transmitting 40 byte packets (mimicking ACK packets) at the same

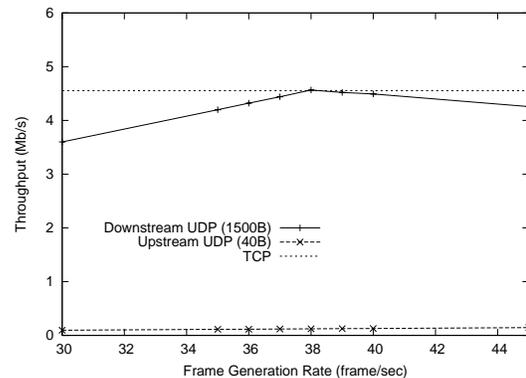


Figure 19: UDP emulation of TCP with 10 wireless stations.

rate. From Section 3 we know that under excessive offered load collision increases sharply and DCF throughput declines, albeit gradually. Our goal is to understand at what offered load TCP drives DCF. throughput as the traffic generation rate λ is varied. We also show the throughput attained by TCP under the same benchmark setting. We observe that UDP throughput increases with increased offered load until it hits a peak—coinciding with TCP’s throughput—after which it declines due to increased collision. The performance result in Figure 19 shows that when the AP-to-wireless station traffic ratio (in pps) is $n : 1$, a typical scenario for WLAN hot spots which are used as Internet access networks, TCP is able to drive DCF at its maximal throughput level.

4.3 TCP-over-WLAN Dynamics: Effective Contention

4.3.1 Markov Chain Analysis: Negative Drift

Our goal is to understand how TCP-over-WLAN is able to achieve a constant collision rate over a very large contention range of 2–100 wireless stations. We define a Markov chain that tracks TCP-over-WLAN dynamics at the WLAN level: at any instant in time, the state of the birth-death Markov chain is given by the number of backlogged wireless stations, including the AP, where backlogged means that the NIC card has received one or more frames from the upper layer that have not been successfully transmitted yet. Thus the backlog chain counts the number of actually active stations participating in the multiple access competition at an instant in time. In simulation, it is straightforward to infer the backlog Markov chain from measurement logs: starting from time 0, we have a point process whose realization or sample path we track. Conditioned on these states, we compute next state transition statistics from which the transition rates can be estimated.

Figure 20 shows the backlog Markov chain inferred from a single point configuration simulation with 20 wireless stations. Starting from counting state 3, we see a strong negative drift that ends at state 6. In the client/server wireless/wireline TCP environment, the AP is a conduit for all downstream traffic from wireline servers to wireless clients,

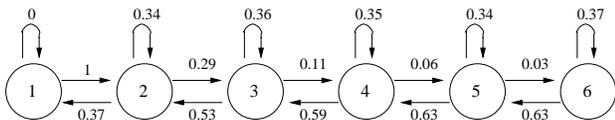


Figure 20: TCP-over-WLAN dynamics: inferred backlog Markov chain from simulation benchmark.

and vice versa for the resultant upstream ACK traffic. If the AP successfully grabs the shared channel, it may send a TCP data frame to a wireless station with existing backlog—in which case the counting state remains put (self-loop)—or it may send the data frame to a non-backlogged station in which case the counting state increases by 1. The full Markov chain contains micro-states that count the number of backlogged packets per station. A client station, upon grabbing the channel, may have more than one frame in its backlog, in which case the counting state stays put; otherwise, it is decremented by 1.

DCF’s contribution to the negative drift is as follows: the

AP in a BSS, by virtue of its conduit role with a $n : 1$ traffic ratio, is in a special forwarding position. However, under DCF all wireless stations, including the AP, are treated equal with respect to channel access. For the counting chain to grow, it requires that the AP transmit a TCP data frame. But the larger the counting state k , the smaller the probability ($\approx 1/k$) that the AP will win in the channel grabbing competition. Conversely, the larger the likelihood ($\approx 1 - 1/k$) that one of the wireless stations will win the competition. Hence the strong negative drift that pulls the Markov chain toward 2–3 counting states. Figure 21 shows the av-

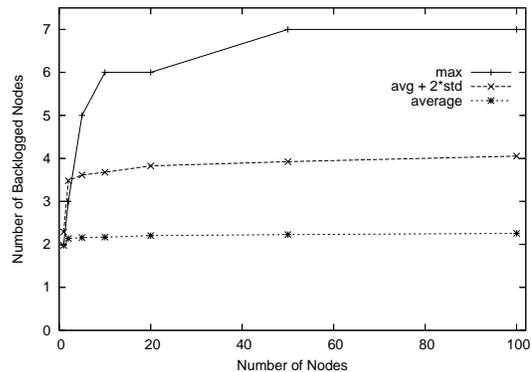


Figure 21: Average counting state in equilibrium for 2–100 wireless station simulation benchmark.

erage counting state, in equilibrium, of all backlog Markov chains for 2–100 stations in the simulation benchmarks of Figure 17. The average counting state is a little above 2, even when there are 100 contending wireless stations. This implies that the WLAN operates under an effective contention level of 2–3 wireless stations. Pertinent WLAN performance traits such as MAC throughput and collision rate also correspond to that of 2–3 wireless stations, albeit in saturation mode. The throughput and collision rate for 2 wireless stations may be discerned in Figure 3. Overall it is desirable to operate a WLAN at an effective contention level of 2 wireless stations, even in saturation mode.

4.3.2 Verification through Experiment

Simulation is suited for carrying out backlog Markov chain analysis because all relevant events can be logged with certainty and there is a global clock that makes event synchronization simple. In experimental benchmarking, there is system noise and synchronization of events—given only local clocks—becomes a technical problem. With the aid of AiroPeek NX [1], a state-of-the-art commercial WLAN sniffing and analysis tool, and tcpdump logging at the laptop and all pocket PCs, we are able to estimate the backlog Markov chain using sequential event ordering at the AiroPeek sniffer. This method works but with an inherent caveat: due to channel noise, collision, and system noise, not all frame transmission events are sensed by the sniffer. This leads to ambiguity and resultant frames that cannot be resolved. The fraction of such frames, however, is less than 2%, and we estimate the backlog chain after excluding these frames. Figure 22 shows the inferred Markov chain from a table top experiment with 15 pocket PCs. The average counting state is 2.59—a little higher than that of simulation-based benchmark results—which we attribute to the aforementioned in-

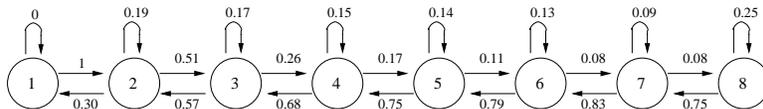


Figure 22: Backlog Markov chain inferred from table top client/server TCP benchmark experiment with 15 iPAQ pocket PCs.

accuracies. A strong negative drift is discernible consistent with the preceding backlog Markov chain analysis.

4.4 Dynamic Rate Shifting under TCP-over-WLAN

In Section 3.1.4 we showed that dynamic rate shifting implemented in WLAN cards easily confuse collision with channel noise resulting in significant throughput degradation even under moderate load when channel conditions are good. Figure 23 shows empirical TCP throughput in the indoor office environment for auto rate and fixed rates 11, 5.5, 2, and 1 Mbps as the number of contending iPAQ pocket PCs is increased from 1 to 14. Unlike in Figure 10, we observe a flat throughput curve under auto rate that is facilitated by the $n : 1$ traffic ratio of hot spot WLANs and the mitigating influence of TCP’s self-clocking mechanism. TCP-over-WLAN, by operating the shared baseband chan-

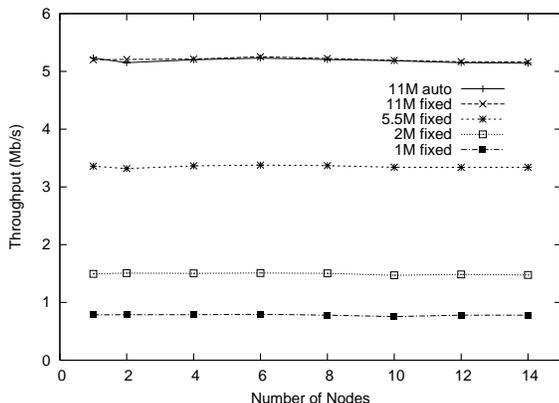


Figure 23: Empirical TCP throughput as a function of the number of pocket PCs for auto rate and fixed data rates 11 Mbps, 5.5 Mbps, 2 Mbps and 1 Mbps in indoor office environment.

nel at an effective contention level of 2–3 wireless stations, is able to keep the collision rate in check preventing frequent occurrence of two consecutive missing 802.11 ACK frames that trigger downshifting to fallback rates. With 14 pocket PCs, less than 0.5% of IEEE 802.11 data frames carrying TCP payload are transmitted at fallback rate 5.5 Mbps (or below) under dynamic rate shifting—data frames with 2 and 1 Mbps data rates are essentially nonexistent—whose low occurrence stems from the fact that less than 0.1% of 802.11 ACK frames are missing back-to-back. The higher fraction (0.5%) of data frames transmitted at fallback rates is due to an asymmetry in the downshift/upshift procedure: upshift in data rate is instituted more conservatively than downshift. Nonetheless, their low frequency renders their performance impact negligible.

5. CONCLUSION

We have studied WLAN and TCP-over-WLAN performance by incorporating the influence of inter-layer dependencies, including physical layer diversity, that significantly impact performance. We have shown that DCF throughput degrades gracefully under increasing offered load and multiple access contention, but fairness and jitter undergo a sudden “phase transition” at a critical offered load. In contrast to throughput, MAC layer fairness and jitter which can benefit from traffic controls aimed at operating the system outside the saturation region. We have shown how dynamic rate shifting implemented in vendor cards can degrade WLAN throughput under moderate contention, which may underlie part of the negative perception of IEEE 802.11 DCF performance. We have shown that TCP-over-WLAN achieves scalable throughput facilitated by the self-regulating actions of DCF and TCP congestion control. As a consequence, we have shown that the TCP data/ack collision problem does not manifest as a significant performance concern. TCP operating over WLAN also exerts a mitigating influence on dynamic rate shifting—an instance of inter-layer protocol dependence—allowing rate shifting to focus on frame errors stemming from channel noise by reducing the occurrence of bursty frame errors from collision. Additional details and performance results, including the influence of BER and variable workload (e.g., Poisson packet arrivals and heavy-tailed file transport), the mitigating effect of TCP-over-WLAN on fairness, and transparent receiver-side control are discussed in an extended paper under preparation.

6. ACKNOWLEDGMENTS

This research was supported, in part, by NSF grants ANI-9875789, EIA-9972883 and ANI-0082861, and KOSEF grant R01-2004-000-10372-0 and BK-21 grant. Part of the research was carried out while Sunwoong Choi was visiting the Network Systems Lab at Purdue University. We would like to thank Namgi Kim for conducting corridor experiments that benchmark WLAN throughput under fixed data rates and dynamic rate shifting. We would also like to thank WildPackets for free use of their AiroPeek NX WLAN sniffing software.

7. REFERENCES

- [1] AiroPeek NX, WildPackets, Inc. <http://www.wildpackets.com/>.
- [2] The network simulator – ns-2. <http://www.isi.edu/nsnam/ns/>.
- [3] IEEE 802.11 part 11: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications, Aug. 1999.
- [4] N. Abramson. The Aloha system—another alternative for computer communications. In *Proc. Fall Joint Comput. Conf. AFIPS Conf.*, pages 281–285, 1970.

- [5] A. Balachandran, G. Voelker, P. Bahl, and P. Rangan. Characterizing user behavior and network performance in a public wireless LAN. In *Proc. ACM SIGMETRICS '02*, pages 195–205, 2002.
- [6] H. Balakrishnan, V. Padmanabhan, S. Seshan, and R. Katz. A comparison of mechanisms for improving TCP performance over wireless links. *IEEE/ACM Trans. Networking*, 5(6):756–769, 1997.
- [7] G. Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE Journal on Selected Areas in Communications*, 18(3):535–547, 2000.
- [8] D. Boggs, J. Mogul, and C. Kent. Measured capacity of an Ethernet: myths and reality. In *Proc. ACM SIGCOMM '88*, pages 222–234, 1988.
- [9] F. Cali, M. Conti, and E. Gregori. IEEE 802.11 wireless LAN: Capacity analysis and protocol enhancement. In *Proc. IEEE INFOCOM '98*, pages 142–149, 1998.
- [10] F. Cali, M. Conti, and E. Gregori. IEEE 802.11 protocol: Design and performance evaluation of an adaptive backoff mechanism. *IEEE Journal on Selected Areas in Communications*, 18(9):1774–1786, 2000.
- [11] I. Cuinas and M. Sanchez. Measuring, modeling, and characterizing of indoor radio channel at 5.8GHz. *IEEE Transactions on Vehicular Technology*, 50(2):526–535, 2001.
- [12] M. Grossglauser and D. Tse. Mobility increases the capacity of ad-hoc wireless networks. In *Proc. IEEE INFOCOM '01*, pages 1360–1369, 2001.
- [13] P. Gupta and P. Kumar. The capacity of wireless networks. *IEEE Trans. Inf. Theory*, 46(2):388–404, 2000.
- [14] W. Haitao, P. Yong, L. Keping, C. Shiduan, and M. Jian. Performance of reliable transport protocol over IEEE 802.11 wireless LAN: Analysis and enhancement. In *Proc. IEEE INFOCOM '02*, pages 599–607, 2002.
- [15] V. Joel. Exploding the myth of WLAN performance. *Telephony Online*, October 2004.
- [16] H. Kim and J. Hou. Improving protocol capacity with model-based frame scheduling in IEEE 802.11-operated WLANs. In *Proc. ACM MobiCom '03*, pages 190–204, 2003.
- [17] S. Kjesbu and T. Brunsvik. Radiowave propagation in industrial environments. In *26th Annual Conference of the IEEE Industrial Electronics Society*, volume 4, pages 2425–2430, 2000.
- [18] L. Kleinrock and F. Tobagi. Packet switching in radio channels: Part I - carrier sense multiple-access modes and their throughput-delay characteristics. *IEEE Trans. Commun.*, 23(12):1400–1416, 1975.
- [19] R. Knopp and P. Humblet. Information capacity and power control in single-cell multiuser communications. In *Proc. ICC '95*, 1995.
- [20] D. Kotz and K. Essien. Analysis of a campus-wide wireless network. In *Proc. ACM MobiCom '02*, pages 107–118, 2002.
- [21] U. Kozat and L. Tassiulas. Throughput capacity of random ad hoc networks with infrastructure support. In *Proc. ACM MobiCom '03*, pages 55–65, 2003.
- [22] B. Liu, Z. Liu, and D. Towsley. On the capacity of hybrid wireless networks. In *Proc. IEEE INFOCOM '03*, pages 1543–1552, 2003.
- [23] S. Lu, T. Nandagopal, and V. Bharghavan. A wireless fair service algorithm for packet cellular networks. In *Proc. ACM MobiCom '98*, pages 10–20, 1998.
- [24] R. Metcalfe and D. Boggs. Ethernet: distributed packet switching for local computer networks. *Communications of the ACM*, 19(7):395–404, 1976.
- [25] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose. Modeling TCP throughput: A simple model and its empirical validation. In *Proc. ACM SIGCOMM '98*, pages 303–314, 1998.
- [26] K. Pentikousis. TCP in wired-cum-wireless environments. *IEEE Communications Surveys*, Fourth Quarter(7):2–14, 2000.
- [27] R. Rom and M. Sidi. *Multiple Access Protocols: Performance and Analysis*. Springer-Verlag New York, 1990.
- [28] N. Spring, M. Chesire, M. Berryman, V. Sahasranaman, T. Anderson, and B. Bershad. Receiver based management of low bandwidth access links. In *Proc. IEEE INFOCOM '00*, pages 245–254, 2000.
- [29] S. Sridhar, D. Hogenboom, and B. Willemsen. Microwave experiments on chaotic billiards. *J. Stat. Phys.*, 68(239), 1992.
- [30] D. Tang and M. Baker. Analysis of a local-area wireless network. In *Proc. ACM MobiCom '00*, pages 1–10, 2000.
- [31] R. Tingley and K. Pahlavan. Space-time measurement of indoor radio propagation. *IEEE Trans. Instrum. Meas.*, 50(1):22–31, 2001.
- [32] D. Tse. Optimal power allocation over parallel Gaussian channels. In *Proc. International Symposium on Information Theory*, 1997.
- [33] N. Vaidya, P. Bahl, and S. Gupta. Distributed fair scheduling in a wireless LAN. In *Proc. ACM MobiCom '00*, pages 167–178, 2000.
- [34] A. Vasani and U. Shankar. An empirical characterization of instantaneous throughput in 802.11b WLANs. Technical Report CS-TR-4389, UMIACS-TR-2002-69, Department of Computer Science and UMIACS, University of Maryland College Park, 2002.
- [35] G. Xylomenos and G. Polyzos. TCP and UDP performance over a wireless LAN. In *Proc. IEEE INFOCOM '99*, pages 439–446, 1999.
- [36] G. Xylomenos, G. Polyzos, P. Mähönen, and M. Saaranen. TCP performance issues over wireless links. *IEEE Communications Magazine*, 39(4):52–58, 2001.
- [37] C. Yuen and P. Marbach. Service differentiation in random access networks. Technical Report CSRG-472, Computer System Research Group, University of Toronto, Jul. 2003.