Extended Abstract: Understanding the Effect of Access Point Density on Wireless LAN Performance

Mesut Ali Ergin, Kishore Ramachandran and Marco Gruteser
WINLAB, Electrical and Computer Engineering Department
Rutgers, The State University of New Jersey
North Brunswick, NJ, USA
{ergin, kishore, gruteser}@winlab.rutgers.edu

ABSTRACT

In this paper, we present a systematic experimental study of the effect of inter-cell interference on IEEE 802.11 performance. With increasing penetration of WiFi into residential areas and usage in ad hoc conference settings, chaotic unplanned deployments are becoming the norm rather than an exception. These networks often operate 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 802.11 network is determined by the number of active clients. Surprisingly, we find that with a typical TCP-dominant workload, cumulative system throughput is characterized by the number of interfering access points rather than the number of clients. We find that due to TCP flow control, the number of backlogged stations in such a network equals twice the number of access points. Thus, a single access point network proved very robust even with over one hundred clients. Multiple interfering access points, however, lead to an increase in collisions that reduces throughput and affects volume of traffic in the network.

Categories and Subject Descriptors: C.4 [Performance of Systems]: Measurement techniques, Performance attributes
General Terms: Experimentation, Measurement, Performance
Keywords: High-density WLANs, Intercell interference, Wireless testbeds, Real-world evaluation

1. INTRODUCTION

Wireless LAN's (WLAN) continuing market success and beginning integration into consumer devices such as printers, mobile phones, and music players, can be expected to lead to client node densities approaching or surpassing those of cellular telephony. This usage growth likely leads to increased co-channel interference. The conventional approach to mitigating inter-cell interference is careful frequency planning, to prevent neighboring cells from operating on the same channel, transmission power control, and carrier sense tuning to increase spatial reuse. Dynamic forms of these approaches such as distributed channel selection are also possible.

Even with such techniques, however, several factors render complete elimination of inter-cell interference difficult in WLANs. First, many WLANs, especially in residential or small business areas, are deployed in an unplanned, chaotic manner. Lack of coordination between neighbors and the restrictions of many ISP contracts 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 channel 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 [9]. Third, the number of available non-overlapping channels in the unlicensed spectrum for 802.11 a/b/g is still low, even with the 2006 US FCC allocation of 11 additional 802.11a channels. Therefore, 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 [9].

To date, little is known about the effects of such inter-cell interference on 802.11 system performance. Detailed analytical and simulation models exist for the medium access protocol's scalability [6, 11] and experimental studies have characterized scalability under TCP and UDP workloads [8] 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 [12] or conference settings [4, 14] and recent measurements in a dense deployment have detected performance anomalies [14, 13]. However, 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 high density WLANs. It complements earlier real-world measurements through experimentation with 100+ 802.11 nodes in a repeatable laboratory setting with controlled interference. Thus, it allows in-depth analysis through repeatable experiments, with precisely known configurations. In summary, our contributions with this paper include:

- Analysis of system performance using a realistic TCP dominated traffic mix in multi-cell WLANs. 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 [8] for 16 clients to a much larger network, with realistic client association patterns, and bursty traffic mixes. We also show that, in a multi-cell network, however, the collision rate increases significantly.
- Providing novel insights into the behavior of TCP in multi-

MobiCom'07, September 9–14, 2007, Montréal, Québec, Canada.
ACM 978-1-59593-681-3/07/0009.

cell WLANs. Due to TCP flow control the number of backlogged stations equals twice the number of access points, meaning that network efficiency is determined by the number of interfering access points, not the number of clients.

2. METHODOLOGY

In our evaluations, we use the 400-node ORBIT testbed [16] that allows flexible topologies. The main components of this integrated wired/wireless IP network are wired nodes hosting application servers, wireless access points (AP) and clients. 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.

2.1 Experimental Setup

Table 1 provides a list of the parameters we used in the experiments. To ensure that our results are representative of real-world behavior, we first carried out calibration tests comparing throughput of ORBIT machines configured as APs (that use the MadWifi driver[15]) with commercial Cisco (1200 series) and D-Link (DI-7XX series) APs. Since we did not observe a significant difference (less than 5%), we configured ORBIT nodes as APs with transparent bridging and used traffic sniffers (via tcpdump) on both the wired and wireless segments.

Metrics: For performance evaluation, we choose to use application-specific metrics and focus on system throughput and throughput fairness (i.e., Jain's fairness index [1]).

2.2 Traffic model

Workloads: Since inter-cell interference is affected by end-user workloads, we constructed a synthetic office workload in addition to bulk TCP-only workloads. The office workload is based on sniffer measurements obtained in our academic office environment for a single access point serving about 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. These measurements are reasonably consistent with (except for a 20% increase in web traffic) an earlier analysis of SIGCOMM 2001 conference traces covering 4 APs and 195 stations [4].

Bursty web traffic in our workload follows the self-similar ON-OFF traffic model in [5]. We emulate the HTML transfer, browser processing, and HTML object retrieval phases with *wget* page retrieval tool. User's thinking time X between page accesses is Pareto distributed, with parameters $k = 1.5$ and $x_m = 1$, as suggested in [5].

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 using the D-ITG traffic generator v.2.4.3 [3].

User Arrival Pattern: To measure performance in a more realistic manner, we extracted the user arrival patterns from WLAN traces of the 64^{th} IETF meeting [14]. In particular, we use the arrival pattern of the users returning from lunch between 12:30pm to 1:00pm. The IETF WLAN comprised over 150 APs and more than 700 users.

Table 1: Attributes Summary for Experiments

Attribute	Value
Radio Nodes	1GHz VIA C3 Processor, 512MB RAM, 20GB HDD
Wireless Interfaces	2 X Atheros AR5212 based mini-PCI 802.11a/b/g
PHY/LLC/MAC Used	IEEE 802.11a @ Channel 52
PHY Link Speed	up to 54 Mbits/s
Wireless output power	18 dBm
MAC retries	10
OS Used	Linux 2.6.18
Wireless Card Driver	MadWifi svn.21XX [15] and Atheros HAL v.0.9.4.5

3. NODE DENSITY AND PERFORMANCE

In this section, we systematically examine the effects of both client and AP density on system performance. We begin with the single-cell baseline and then analyze multi-cell AP configurations.

3.1 Effect of Client Density: Single AP case

This experiment measures time-to-association, per-station and cumulative throughput for a bulk TCP workload with up to 127 clients connected to one AP. The goal of this experiment is to (a) establish a cumulative throughput baseline with varying numbers of stations for comparison with the multi-AP case; and (b) determine the effect of dynamic client association patterns on system performance.

Of the two experiments we summarize here, "Test I" uses the previously described IETF-meeting client arrival pattern and "Test II" assumes arrival of all clients are completed before the experiment. Similar results were obtained for multiple-runs of the same experiment. Our main observations are:

- As the number of clients in the system increases, cumulative TCP throughput remains remarkably robust—even with 127 clients, throughput degradation does not occur as 802.11 MAC analysis under UDP saturation workload would normally suggest (note that downlink TCP traffic has ACK packets that travel in the upstream direction). The average steady-state system throughput values of 23.56 Mbits/s and 23.68 Mbits/s (for Test I and Test II respectively) compares favorably with 24.02 Mbits/s result obtained for a single flow TCP throughput experiment.
- When new clients arrive in a loaded WLAN, associations are frequently delayed, but the dynamic association patterns have no visible effect on system throughput.
- We observe high throughput fairness at steady-state (computed over the last three minutes of the experiment). Jain's fairness index values are 0.9620 for Test I, and 0.9434 for Test II.

Overall, a single-AP network performs efficiently with a downstream TCP workload irrespective of the number of clients it serves. Given that our workload is dominated by TCP, the observed system throughput robustness can be attributed to TCP flow-control [8]. The collision rate remains nearly invariant as the number of stations increases, confirming recent experimental and theoretical work [8, 7].

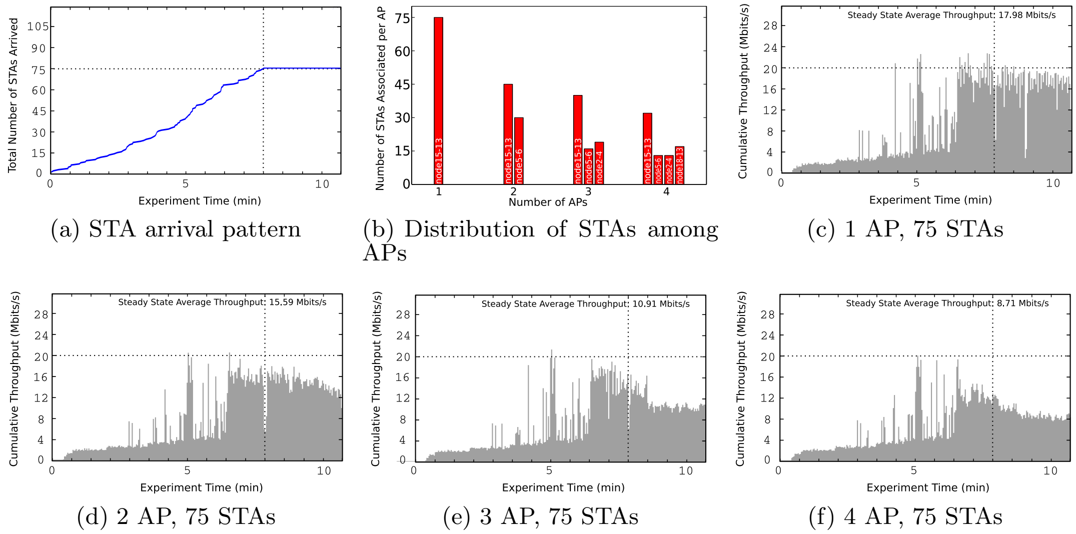

(a) STA arrival pattern
(b) Distribution of STAs among APs
(c) 1 AP, 75 STAs
(d) 2 AP, 75 STAs
(e) 3 AP, 75 STAs
(f) 4 AP, 75 STAs

Figure 1: Cumulative TCP throughput with inter-cell interference. Steady state throughput when all clients are associated is reduced by about 50% with four cells compared to single cell (same number of clients).

As a result, secondary effects that have been observed in practice [14] such as inefficient bit-rate adaptation do not manifest themselves.

3.2 Effect of AP Density

We now consider deployments with multiple interfering APs in close proximity, as described in recent war-driving studies [2] and short-term conference deployments [14, 13]. While these studies provide performance characterizations of such deployments, we concentrate here on isolating the effects of client interference from AP interference in controlled experiments.

The following results were obtained from experiments using the workload and the dynamic client arrival pattern described in Section 2.2. Figures 1(c)-1(e) shows TCP system throughput in an experiment for different numbers of interfering access points in communication range. Note that the final number of clients across experiments remains constant at 75—what differs is the number of access points used to serve the same number of clients. All APs use the same SSID, thus clients select the AP with the highest Received Signal Strength Indication (RSSI) at their position. Our observations are summarized in the following paragraph.

Increasing the number of APs from 1 to 4 while maintaining the same number of clients degrades system throughput by about 50%. Comparing steady-state, in the last 3 min of the experiments, across Figures 1(c) to 1(e), average cumulative throughput degrades from 17.98 Mbits/s for one AP to 8.71 Mbits/s for four APs. Spikes in the first part of the traces are due to bursty web traffic and unsaturated channel conditions. Since the workload and number of clients remain unchanged across the four experiments, we conclude that cumulative TCP throughput is governed by the number of active (i.e. interfering) APs, rather than the number of clients.

To further determine whether the throughput loss in multi-AP installations can be attributed to increased MAC contention, we design another experiment in which we do not consider dynamic client arrivals, bitrate adaptation, or bursty web traffic to focus on MAC contention in a baseline scenario. To better understand degree of station activity in the network, we look at collision-based PER, instead of TCP-throughput. Figure 2(a) shows the mean collision rate as the number of interfering clients is increased for one and two AP networks.[1] In this experiment, 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 one to two APs more than doubles the average PER from 11% to nearly 28%. In calibration experiments, we have confirmed that <1% packet loss in our setup occurs on each individual link, suggesting that nearly all observed packet errors in this scenario are due to collisions.

Models from neither Bianchi [6] nor Kumar et.al. [11] explain the reduction in cumulative throughput and increase in collisions for the multi-cell AP case. In fact, the empirical collision-PER with a single AP (for 4, 8, and 16 STAs) matches the value predicted by Bianchi's Model [6] with two backlogged stations. Similarly, the empirical collision-PER with two APs (for 4, 8, and 16 STAs) matches the value predicted by Bianchi's Model with four backlogged stations. Collision PER from the MATLAB simulations of Bianchi's Model using the same experiment parameters are provided in 2(b). These results lead us to believe that this effect is due to the interaction of TCP with the 802.11 MAC protocol. We analyze this further in the following section.

4. TCP IN MULTI-CELL NETWORKS

The reduction in cumulative throughput for multi-cell networks raises the question: *why is TCP Reno over WLAN robust against intra-cell congestion but not against inter-cell congestion and interference?* Gong and Marbach's model of a TCP over 802.11 system [10] predicts that in a single BSS case, on average, two stations will be backlogged, irre-

[1]Note that we reduced the bit-rate to 36 Mbits/s so that only negligible frame errors occur due to channel noise and path loss, all remaining frame errors are due to interference.

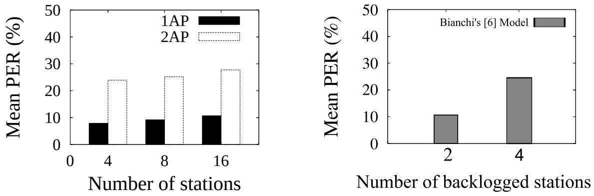

(a) Empirical Results (b) Bianchi's Model

Figure 2: (a) Empirical collision-caused packet error rate (PER) with up to 16 clients and 2 APs on ORBIT testbed. (b) Collision-PER predicted by Bianchi's Model [6] (via MATLAB simulations).

spective of the number of clients. It also predicts that if n additional flows (in the form of IBSS networks) are added as interferers, the expected number of active (i.e., backlogged) nodes at any *operating point* would be $2(1+n)$. The authors of [10] prove the perpetual existence of (at least one) operating point.

Our empirical results in Figure 2(a), when compared to Bianchi's predictions in Figure 2(b), reveal the applicability of the same model to multi-cell BSS networks. We show that multi-cell interference can be treated as additional flow interference, extending the multiple IBSS subnets case studied in [10], thus yielding the following generalization:

Consider multiple 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 have always data to send. At the operating point for TCP Reno with IEEE 802.11, for a network topology consisting of i BSSs, the expected number of backlogged nodes is* $2i$.

An intuitive explanation for this result follows from the fact that TCP requires an equal packet flow rate in both up and downlink directions (i.e., DATA and ACK). Regardless of the direction of traffic workload, the volume of flow per unit time is severely limited by AP's channel access probability. Adding more APs increases cumulative channel access probability in favor of APs, hence exciting more clients (proportional to the number of APs added) to become active.

Note that this result is not a function of downlink dominance of the traffic [7]. We have also carried out simulations to compare a TCP-uplink dominated scenario and observed that direction of application traffic did not change the throughput or collision rates significantly ($<1\%$).

This analysis together with the empirical results in Section 3.2 show that the network efficiency in a TCP-dominated network is primarily a function of the number of interfering APs. This insight also suggests revisiting some of the existing high-density WLAN solutions, such as collision window (CW) adaptations based on the number of clients, to make use of the number of actively interfering cells within an area.

5. CONCLUSIONS

In this work, we have investigated the effect of inter-cell interference on 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 in a testbed with more than one hundred nodes. We found that

- Cumulative throughput degrades significantly, by about 50% with four access points, while it remains remarkably robust with over one hundred clients in the single cell case.
- TCP's flow control leads to an average of $2i$ stations concurrently backlogged in the network, where i is the number of active access points. Thus, the collision rate increases with the number of access points.

These findings suggest practical algorithms and recommendations that can largely recover cumulative throughput loss from inter-cell interference (e.g., CW optimizations based on the number of APs).

6. ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grants No. CNS 0335244 and CNS 0627032.

7. REFERENCES

[1] R. Jain and D. Chiu. A Quantitative Measure Of Fairness And Discrimination For Resource Allocation In Shared Computer Systems. DEC TR-301, September, 1984.

[2] A. Akella, G. Judd, et al. Self-management in chaotic wireless deployments. In *Proc. ACM MOBICOM*, pp. 185–199. 2005.

[3] S. Avallone, S. Guadagno, et al. D-ITG distributed internet traffic generator. In *Proc. QEST 2004*, pp. 316–317. September 2004.

[4] A. Balachandran, G. M. Voelker, et al. Characterizing user behavior and network performance in a public wireless LAN. In *Proc. SIGMETRICS*, pp. 195–205. 2002.

[5] P. Barford and M. Crovella. Generating representative web workloads for network and server performance evaluation. In *Proc. ACM SIGMETRICS*, pp. 151–160. July 1998.

[6] G. Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE JSAC*, 18(3):535–547, May 2000.

[7] R. Bruno, M. Conti, et al. Performance modelling and measurements of TCP transfer throughput in 802.11-based WLAN. In *Proc. ACM MSWiM*, pp. 4–11. 2006.

[8] S. Choi, K. Park, et al. On the performance characteristics of WLANs: revisited. *SIGMETRICS Perform. Eval. Rev.*, 33(1):97–108, 2005.

[9] Extricom, Inc. `http://www.extricom.com`.

[10] Y. Gong and P. Marbach. Interaction of rate and medium access control in wireless networks:: the single cell case. In *Proc. ACM MOBIHOC*, pp. 178–189. 2006.

[11] P. Gupta and P. Kumar. The capacity of wireless networks. *IEEE Trans. on Info. Theory*, 46:388–404, 2000.

[12] T. Henderson, D. Kotz, et al. The changing usage of a mature campus-wide wireless network. In *Proc of ACM MOBICOM*, pp. 187–201. September 2004.

[13] A. P. Jardosh, K. Mittal, et al. IQU: Practical queue-based user association management for WLANs. In *Proc. ACM MOBICOM*, pp. 158–169. September 2006.

[14] A. P. Jardosh, K. N. Ramachandran, et al. Understanding congestion in IEEE 802.11b wireless networks. In *Proc. USENIX IMC*. October 2005.

[15] MADWiFi. *Multiband Atheros Driver for WiFi.* http://madwifi.org, 2007.

[16] D. Raychaudhuri, I. Seskar, et al. Overview of the orbit radio grid testbed for evaluation of next-generation wireless network protocols. In *Proc. IEEE WCNC*, vol. 3, pp. 1664–1669. March 2005.