

Auto-configuration of 802.11n WLANs

Mustafa Y. Arslan *
University of California, Riverside
marslan@cs.ucr.edu

Konstantinos Pelechrinis
University of Pittsburgh
kpele@pitt.edu

Ioannis Broustis †
University of California, Riverside
broustis@cs.ucr.edu

Srikanth V. Krishnamurthy
University of California, Riverside
krish@cs.ucr.edu

Sateesh Addepalli
Cisco Systems, Inc.
sateeshk@cisco.com

Konstantina Papagiannaki
Intel Labs, Pittsburgh
dina.papagiannaki@intel.com

ABSTRACT

Channel Bonding (CB) combines two adjacent frequency bands to form a new, wider band to facilitate high data rate transmissions in MIMO-based 802.11n networks. However, the use of a wider band with CB can exacerbate interference effects. Furthermore, CB does not always provide benefits in interference-free settings, and can even degrade performance in some cases. We conduct an in-depth, experimental study to understand the implications of CB. Based on this study we design an auto-configuration framework, ACORN, for enterprise 802.11n WLANs. ACORN integrates the functions of user association and channel allocation, since our study reveals that they are tightly coupled when CB is used. We show that the channel allocation problem with the constraints of CB is NP-complete. Thus, ACORN uses an algorithm that provides a worst case approximation ratio of $O(\frac{1}{\Delta+1})$, with Δ being the maximum node degree in the network. We implement ACORN on our 802.11n testbed. Our experiments show that ACORN (i) outperforms previous approaches that are agnostic to CB constraints; it provides per-AP throughput gains from 1.5x to 6x and (ii) in practice, its channel allocation module achieves an approximation ratio much better than $O(\frac{1}{\Delta+1})$.

1. INTRODUCTION

802.11n is based on the use of MIMO technology and promises drastically improved throughputs as compared to legacy 802.11 systems (a/b/g). In order to achieve high data rates (> 100 Mbps), 802.11n [1] supports a feature called channel bonding (CB). With CB, two adjacent channels can be combined to form a new, wider frequency band; this is expected to support transmissions at higher bit rates. 802.11n

*This work was supported in part by Cisco University grant, NSF NeTS grant 0626912 and Multi-University Research Initiative (MURI) grant W911NF-07-1-0318.

†Ioannis Broustis is now affiliated with Alcatel-Lucent.

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.

ACM CoNEXT 2010, November 30 – December 3 2010, Philadelphia, USA.

Copyright 2010 ACM 1-4503-0448-1/10/11 ...\$10.00.

uses 40 MHz channels with 108 OFDM subcarriers when CB is activated (as compared to 20 MHz bands with 52 subcarriers).

The use of CB however, has associated caveats. Communication links utilizing 40MHz channels project increased interference on other links and this can negatively impact the total network throughput [2], [3]. One might expect that CB would yield significantly higher throughputs in interference-free settings. However, the throughput performance with CB can be even worse than that with a single 20 MHz channel in many cases.

Channel bonding in interference-free settings: With a fixed transmission power (T_x), there is about a 3dB decrease in the signal power per subcarrier when CB is employed. Thus for a fixed T_x , the Bit Error Rate (BER) and the Packet Error Rate (PER) with CB is always greater than or equal to that without CB. As a consequence, the throughput observed with CB is almost always “less than double” of that without CB¹. Second, the use of a larger number of OFDM subcarriers increases the likelihood of experiencing errors. We find that links of poor quality are the ones that are most affected by these factors; the performance on such links with CB is worse than without CB. In fact, as discussed later, the existence of a single poor client in a cell that uses a wider band can degrade the long term throughput of the entire cell (due to the 802.11 performance anomaly [4]). Note that T_x cannot be increased beyond a specified maximum value; this value is the same for both 20 and 40 MHz channels (from the 802.11n specs [1]). In addition, increasing T_x may project additional interference on other links.

Channel bonding and interference: The use of CB projects interference over a larger spectral bandwidth (40 MHz as opposed to 20 MHz). Thus, the wider channels will have to be carefully assigned to cells in an enterprise WLAN. Moreover, note that due to the 3 dB reduction in the per-carrier signal power, transmissions with the wider bands are more susceptible to interference (i.e., the SINR is lower).

ACORN, our auto-configuration framework: The above observations suggest that CB must be administered with care. In this paper, we design ACORN, an Auto-ConfiguRation framework for enterprise 802.11n WLANs. ACORN jointly performs the functions of channel allocation and user asso-

¹In an ideal setting, one would expect that doubling the channel width would yield twice the throughput.

ciation. As discussed above, the existence of poor links impacts the performance of CB; therefore, intelligent user association can be significant in facilitating throughput gains from CB. In brief, clients make association decisions not in a selfish/greedy manner, but by considering the impact of their decision on neighboring cells. APs make decisions on whether or not to use CB by considering factors such as interference, towards achieving network-wide performance gains. To our best knowledge, ACORN is not only the first system to consider the application of CB in enterprise WLANs but also the first system to consider the use of two distinct bandwidths while performing joint channel allocation and user association.

The main contributions of our work are the following:

- We perform extensive lower layer experiments on WARP platform [5] and higher layer experiments on our 802.11n testbed towards validating that the performance can indeed degrade in interference-free settings if CB is used.
- We design ACORN, an auto-configuration system for 802.11n WLANs. ACORN jointly performs user association and channel allocation with CB.
- We show that the channel allocation module of ACORN achieves a theoretical worst case approximation ratio of $O(\frac{1}{\Delta+1})$ as compared to an optimal algorithm where Δ is the maximum node degree in the network (the problem is shown to be NP-complete).
- We implement ACORN on our 802.11n testbed and show that (i) it provides significant performance benefits over existing schemes designed for legacy 802.11 (which typically employ bands of a single width). Gains from 1.5x up to 6x are observed and (ii) in practice our channel allocation performs better than the theoretical worst case approximation.

The rest of the paper is organized as follows. In Section 2 we provide brief background on the 802.11n PHY and discuss related work. Section 3 describes our PHY and higher layer measurements that provide an understanding of the behaviors with CB. We present the design of, and analyze the algorithms included in ACORN in Section 4. Section 5 describes the system implementation and our experimental evaluations. Our conclusions form Section 6.

2. RELATED STUDIES

In this section, we provide a brief summary of the 802.11n PHY layer and describe related previous work on resource allocation and experimental studies on 802.11n.

802.11n Experimental Studies: 802.11n utilizes MIMO communications at the PHY with OFDM. Two modes of operations are feasible with 802.11n: (i) Spatial Division Multiplexing (SDM), which achieves higher data rates and (ii) Space Time Block Coding (STBC), which achieves higher reliability. Typically, vendors implement rate adaptation algorithms with 802.11n, which choose the mode of MIMO operations based on the link quality [6]. In [2], Shrivastava *et al.* identify that 802.11n networks are somewhat limited by the CSMA/CA access policy inherited from the 802.11

legacy protocol family. In addition, they provide insights on increased interference due to CB, using a previously proposed model [7]. However, their conclusions cannot explain why an isolated link does not always enjoy a higher throughput with CB, as compared to the single channel scenario, which is the case in practice. Visoottiviseth *et al.* [3] compare the performance between commodity 802.11g and 802.11n devices. However, this work does not examine CB in depth.

Common to the above studies is the conclusion that the negative impact of CB on 802.11n's performance is due to increased levels of interference. We find that there are additional effects which demonstrate that CB should be applied with care.

Orthogonal Frequency Division Multiplexing (OFDM): OFDM divides the allocated spectrum into smaller subchannels (subcarriers) which are orthogonal to each other. Each subcarrier carries data at lower rates; however together they maintain the total data rate. OFDM systems can cope better with narrowband interference and fading [8]. Lately, there are many studies that try to improve OFDM performance. For example, Rahul *et al.* [9] implement a downlink OFDM PHY for WLANs that performs rate adaptation per subcarrier. Subcarrier power allocation is examined in [10]. In this work, we are not interested in changing the PHY layer functions of 802.11n OFDM; we examine the effects of the PHY on the higher layers and design a solution towards improving performance without requiring any changes at the lower PHY / MAC layers.

Channel Width Adaptation: Commodity 802.11a/b/g hardware vendors have also implemented a function that uses channels of varying width. These systems can operate on 5, 10, 20 and 40 MHz bands [11]. In [12], the authors propose a channel width adaptation algorithm that can dynamically choose a channel width to satisfy an optimization criterion. Although the authors discuss the potential usage of the algorithm in WLANs, the algorithm is designed assuming two communicating nodes and its applicability in enterprise WLANs is not studied. We examine the impact of different channel widths on 802.11n WLAN performance by jointly considering user association and channel allocation.

Resource Allocation: Resource allocation in WLANs usually refers to power/rate control and channel selection. User association is very tightly related with these functions; therefore, in many cases they are studied together.

Static channel allocation is shown to have associated fairness issues and thus, Mishra *et al.* [13] proposed a dynamic algorithm based on frequency hopping. In [14], the same authors formulate the channel assignment problem as a graph coloring problem. Kumar *et al.* [15] formulate a utility optimization problem that accounts for user association. All of the above efforts however, study the problems of channel selection or user association independently. In [16] and [17] the problem of jointly performing frequency selection and user association is studied. Mishra *et al.* [16] provide a centralized approach, while Kauffman *et al.* [17] provide a distributed solution based on the Gibbs sampler. In [18], the authors study the interaction between channel allocation, power control and user association.

All the above (independent or joint) network optimizations are designed for legacy 802.11a/b/g networks operating with single channel widths. The complexities associated with CB make the above algorithms inadequate. To the best of our knowledge, we are the first to study the joint problem of channel selection and user association in 802.11n WLANs, considering two distinct channel widths.

3. CHANNEL BONDING IS NOT PANACEA

In this section, we first present and analyze our PHY layer experimental observations relating to CB. Later, we examine the impact of the PHY on the higher layers.

3.1 Channel bonding micro-effects

802.11n leverages the OFDM-based physical layer to implement CB. The OFDM sub-layer is inherited from legacy 802.11 systems (a/g). 802.11a/g use 64 subcarriers to form a single OFDM symbol of $4\mu\text{sec}$ duration. The OFDM subcarriers together form a 22MHz channel (or band). 16 of the subcarriers are used to form guard intervals and for carrying pilot tones; thus, 48 subcarriers carry data in an OFDM symbol. 802.11n increases the number of data subcarriers to 52 in a 20 MHz channel which results in a nominal bit rate of 65Mbps for a single data stream². CB utilizes two adjacent channels simultaneously and employs 108 data subcarriers over a total of 40MHz bandwidth. As one might expect, the nominal bit rates with 40MHz are slightly higher than double of their 20 MHz counterparts for the same modulation scheme. However, there is an important factor that can negatively impact the achievable throughput in a 40MHz channel; the SNR experienced is lower with CB than when a 20 MHz band is used. To understand why this is the case, let us look at the impact of increasing the number of OFDM subcarriers on (a) the thermal noise and (b) the energy per subcarrier.

Impact of CB on thermal noise: The total thermal noise in a 40 MHz channel is higher than that in a 20 MHz channel. The noise floor, N , in Wi-fi systems can be calculated as follows [19]:

$$N(\text{in dBm}) = -174 + 10 \cdot \log(B) \quad (1)$$

where B is the bandwidth of the channel in Hz. It is easy to see that the noise in a 40 MHz channel is about 3 dBm ($10\log 2$) higher as compared to a 20 MHz channel. If one assumes that the noise is uniformly distributed across the subcarriers, the noise per subcarrier can be expected to remain almost the same for both the 20 MHz and 40 MHz channels, and in theory there is just a 4% reduction.

Impact of CB on subcarrier energy: The 802.11n standard [1] mandates the use of the same maximum transmit power with and without CB. In an OFDM system, the transmit energy is evenly distributed across the subcarriers. Since CB uses 108 subcarriers and the total transmit power remains the same, the energy per subcarrier is theoretically reduced by 48% (approximately halved) as compared with

²This assumes a 800nsec guard interval (GI). The option of using a shorter GI (400nsec) is also available; this reduces the symbol duration to $3.6\mu\text{sec}$ and further increases the data rate. For more details on OFDM see [8].




Figure 1: PSD estimate with different channel widths.

that of 20MHz bands. Expressed in dB, this translates to about a 3 dB reduction in the power per sub-carrier. This in turn, can have an impact on the performance.

Considering the two factors together, we postulate that:

- For the same transmission power (T_x), the SNR of a signal is reduced by 3dB with CB. A reduction similar to this postulated value (52%), is seen in the energy per subcarrier in experiments (discussed below); the noise per subcarrier remains almost the same as predicted theoretically (4% reduction). Thus, **the SNR per sub-carrier is halved**; combined with the increased probability of error due to the larger number of subcarriers (each subcarrier experiences a different fade), a BER increase is expected.
- In order to achieve the same SNR on a link, with both 20 and 40MHz channels, we need to increase the transmission power in the latter case. However, this might not be possible given power budget constraints; in addition 802.11n dictates the use of the same maximum power with and without CB.
- One might expect these factors to primarily affect poor quality links since these links are the ones that are most susceptible to the lowered SNR.

With additive white Gaussian noise (AWGN), one can deduce the above observations from Shannon's theorem on the capacity of wireless channels [19]. The capacity C of an AWGN channel (in bits per second) with a bandwidth B (in Hz) is:

$$C = B \cdot \log(1 + SNR) \quad (2)$$

One can immediately see that for low SNR values the logarithmic term dominates. Thus, if increasing B decreases SNR (as is the case here when we transition from a 20 MHz channel to a 40 MHz channel), there may be regimes where the capacity decreases. The validity of these theoretical assessments in real settings has to be investigated, since the noise may not be AWGN in such settings. We next perform experiments for this purpose.

Experimental validation: We use the WarpLab framework with WARP [5] hardware to implement a basic OFDM system. We generate a random bitstream and modulate it using DQPSK. The Inverse Fast Fourier Transform (IFFT) is applied on the modulated I - Q (In-phase and Quadrature) samples. A cyclic prefix is then added. A Barker sequence is




Figure 2: Received constellations with different numbers of subcarriers; (a) 52 and (b) 108.

later prepended to facilitate symbol detection at the receiver. These samples are transmitted over the air using 2x2 STBC (Space Time Block Codes with two antennas - Alamouti) [20]; we use the STBC mode of transmission since on poor quality links, the auto-rate function of our 802.11n cards induces operations in this mode. At the receiver, the preamble sequence is detected and stripped. The cyclic prefix is removed and the remaining samples are fed into a Fast Fourier Transform (FFT) module. After demodulating the samples, the receiver obtains the bitstream. We implement the CB functionality by appropriately changing the subcarrier mappings, and using a 128-point FFT (as opposed to a 64-point FFT with a 20MHz channel).

First, we investigate the spectral characteristics of the transmitted waveform. We obtain the power spectral density (PSD) of the transmitted signals. The same power T_x , is used for both 20 and 40MHz channels. PSD reflects the distribution of the energy with regards to the frequency content of the signal. Figure 1 depicts the PSD estimate of the spectrum with 52 and 108 subcarriers. It is evident, that there is an approximate 3dB reduction (-92 dB to -95 dB) in the energy per subcarrier when we increase the channel width. Note that in 802.11n systems the central frequency F_c is not the same for 20 and 40MHz channels (we depict the PSD using the same F_c for both in Figure 1). F_c is shifted and two adjacent, non-overlapping orthogonal bands are combined [21].

To elucidate the effects of CB on a received signal, we present a typical sample from the received QPSK constellations for both 20 and 40MHz channels in Figure 2. With 20MHz the received symbols are mostly clustered around the actual transmitted symbol on the I-Q plane. With CB, there is a higher uncertainty for the transmitted symbol due to the lowered energy per subcarrier. The signals are more vulnerable to fading and more likely to be erroneously decoded. The higher baud error rate (error rate of QPSK symbols) results in a higher BER.

To look at the implications of our last observation, we measure the Bit Error Rate (BER) with 20 and 40MHz channels. We use the OFDM reference design from [5]. On top of the OFDM PHY, we use the BERMAC implementation [5] to transmit packets and calculate the BER for various settings. We use a custom Java application that transmits back to back packets to an Ethernet switch using the *jpcap* API [22]. The BERMAC implementation loads specific buffers of the WARP boards (on both the transmitter and the receiver




Figure 3: BER for QPSK modulation with respect to (a) SNR and with respect to (b) T_x .




Figure 4: Uncoded PER for QPSK modulation.

boards) with the packet transmitted by the Java application; thus, the receiving node knows the *correct* payload contents for the BER calculation. In our experiments, we transmit a total of 9000 packets with a packet size of 1500 bytes and collect the BER statistics from the receiving board. We calculate the BERs using our WARP boards with 20 and 40 MHz channels, with QPSK. Note here that these are uncoded BERs (no forward error correction (FEC)). The BER results with respect to the measured SNR are presented in Figure 3 (a). As one can expect, for a fixed SNR, the BER does not depend on the channel width. We also plot on the same figure the theoretical bit error rates for the considered system from [19]; the theoretical BER formula depends only on the SNR per subcarrier and not on the bandwidth. We observe that the experimental curves fit well with the theoretical plots³. In particular, the coefficient of determination [23] is 0.8 and 0.89 for 20 and 40MHz, respectively. Figure 3 (b) presents the same set of BER measurements, but with respect to T_x . We notice that the wider channel exhibits a higher number of bits in error for a given T_x , thus corroborating our intuition (discussed earlier).

3.2 Effects of CB on higher layers

The PHY layer performance is not always directly exported to the higher layers (due to FEC, headers etc.). Therefore, the performance degradation in terms of BER with a 40MHz channel for a fixed transmission power, may or may not be reflected in the performance observed at the higher layers. To understand the effect at the packet level, we look at PER (Packet Error Rate). A small increase in the raw uncoded BER (when using a 40MHz channel) might result in no change in the PER on a commercial coded system like 802.11n. If so, the throughput enjoyed at the application

³We leverage this later in designing ACORN in Section 4.

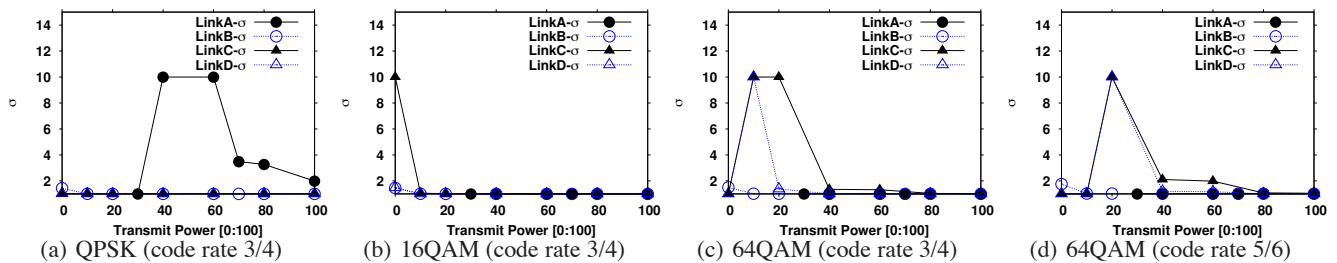


Figure 5: σ -values for different links, T_x , modulations (mod) and code rates (cod). For a given link, CB is beneficial ($\sigma < 2$) only beyond a certain power level. For lower power levels (lower SNR), CB hurts performance ($\sigma \geq 2$).

layer is practically doubled (ignoring MAC overhead). On the contrary, if the PER is also increased, this will result in less than a two-fold increase in the throughput or in some cases, could even result in a throughput reduction. Assuming the throughput T with transmission rate R is roughly $T = (1 - PER) \cdot R$, the throughput with a 40 MHz channel will be lower than that with 20 MHz if the following holds:

$$\sigma = \frac{1 - PER_{20}}{1 - PER_{40}} > \frac{R_{40}}{R_{20}} \approx 2, \quad (3)$$

where PER_x and R_x are the PER and the nominal transmission rate with the x MHz channels. Here, σ is simply the ratio of packet delivery probabilities achieved without and with CB.

PER performance: Using the previous experimental setup and results, we obtain the uncoded PER. Figure 4(a) presents the PER with respect to the SNR. As discussed, for a given SNR the BER does not depend on the channel width; thus, the uncoded PER is similar for the 20 and 40 MHz channels for the same SNR. However, for the same T_x , the PER with CB is much higher as compared to that without the feature (Figure 4(b)). Recall that with 40MHz, for the same T_x , the SNR per subcarrier is halved and this contributes to the performance degradation.

Experiments with commodity 802.11n cards: The performance of an 802.11n system might differ from what we have observed, given the implementation details (e.g. due to FEC). To examine the performance with commodity systems, we conduct experiments on our 802.11n testbed. Our testbed consists of 18, 2x3 802.11n nodes, each equipped with a Ralink mini-PCI card and three 5-dBi omnidirectional antennas. The testbed contains both indoor and outdoor links [24]. We operate on the 5GHz band, avoiding interference from other co-located WLANs in the 2.4GHz band.

With our initial 802.11n experiments, we examine the performance of the (coded) PER. We tune the transmission power on our links and we measure the PERs with both 20 and 40 MHz channels. In Figure 5, we plot the σ -values (Eq. 3) for various modulation schemes and code rates for 4 representative links (BPSK is omitted since it exhibits a similar performance as QPSK). Whenever σ is larger than 2, the throughput achieved with CB will be lower than that with a 20 MHz channel (inequality 3). Table 1 presents the observed SNR values (γ) when we have a crossover value of $\sigma = 2$. The common trend identified among all the cases where $\sigma \geq 2$

mod ^{cod}	QPSK ^{3/4}	16QAM ^{3/4}	64QAM ^{3/4}	64QAM ^{5/6}
$\sigma \geq 2$	-7dB	3dB	5dB	8dB
$\sigma < 2$	-4dB	5dB	7dB	11dB

Table 1: Experimental transition table for σ values.

is that this degradation in performance is observed for a certain range of transmission powers. This region maps to a 2-3 dB difference in SNR⁴. More specifically, for low T_x , the PER for both the 20 and 40 MHz channels is similar (and close to 1), resulting in $\sigma \approx 1$. As we increase the power, the PER with a 20 MHz channel drops faster (with respect to T_x), since the SNR at the receiver is 3dB higher as compared to the CB case. Thus $(1 - PER_{20})$ increases faster than $(1 - PER_{40})$ and the ratio σ can indeed assume values larger than 2; however, as we keep increasing T_x , the SNR with 40MHz also increases and therefore, the PER performance with the two cases become similar to each other again (almost no packet losses) and $\sigma \approx 1$.

Furthermore, for a given link, the SNR value γ , at which we begin to see a PER decrease with 20MHz (which results in $\sigma \geq 2$), is higher as the modulation becomes more aggressive (Table 1). The reason for this is that with more aggressive modulations there is a higher SNR requirement to correctly receive packets.

Note that, for some robust links (e.g., link B in Fig. 5) the PER is extremely low for both the 20 and 40 MHz channels and here CB will provide huge benefits in terms of throughput. For poorer links, the difference in the PER performance can be significant and inequality 3 might be satisfied. In such cases, a 20 MHz channel is preferable.

Throughput performance: From the perspective of the end-user, the achievable throughput at the application layer is what is important. To examine the higher layer performance, we measure the achievable throughput with and without CB on our 802.11n testbed. The rate control algorithm of our cards is used. This proprietary algorithm not only adjusts the rates in response to packet successes/failures but also picks the best mode of operation (SDM or STBC) based on the channel quality. We consider both UDP and TCP. The maximum transmission power is used and we consider all of our links (24 in total) to capture a wide variety of link qualities. Figure 6 (a) depicts the results. We see that in

⁴When σ is > 10 , we cap its value at 10 for better visualization.




Figure 6: Throughput experiments (a) with rate control and (b) with fixed rates.

about 20% of our experiments the use of the conventional 20MHz channel yields higher throughput. The majority of these cases are clustered in the low throughput regimes. The SNR values observed during these experimental trials were smaller than 6dB, which conforms with our previous BER/PER observations in Table 1. In addition, approximately 30% of our TCP experiments yield better performance with 20MHz as compared to only 10% with UDP. TCP is more sensitive to packet losses and as a result even small PER increments can significantly degrade performance.

To understand the above observations, we experiment with fixed transmission rates. For every link on our testbed, we find through exhaustive search, the Modulation Coding Scheme (MCS [1]) which gives the highest (UDP) throughput with and without CB, considering both modes of 802.11n operations (SDM/STBC). The results from Figure 6 (b) imply that the optimal modulation scheme with 40MHz is almost always less ‘aggressive’ (smaller MCS) as compared with the one with 20MHz. This results in less than a two-fold throughput increase with CB as compared with 20MHz channels. Figure 6 (a) verifies this, since the vast majority of the points lie on the right side of the line $y = 2x$. This is again an artifact of the increased BER and PER with CB.

To summarize, our experimental study shows that CB does not always increase throughput. If one were to also consider the increased interference from wider channels, it becomes evident that channel assignment algorithms should be carefully designed for 802.11n WLANs. In addition, user association is even more critical in the case of 802.11n than in legacy systems; a poor client might hurt a cell utilizing CB. Next, we describe the design of our 802.11n auto-configuration framework ACORN, which accounts for all the above factors.

4. HARVESTING ACORN

ACORN is designed based on the understanding developed with the experiments described in the previous section. It consists of two modules: (a) a user association module and (b) a channel assignment module that accounts for CB. The operations of ACORN, are briefly depicted in Figure 7. In short, the functions of the two modules are the following:

a) The **user association module** tries to group users (or clients) of similar link qualities within the same cell. The basis for this assignment is that the presence of a few poor quality users in a cell can degrade the performance drastically with 40 MHz bands. Even the good clients suffer due to the performance anomaly with 802.11 [4]. In brief, the

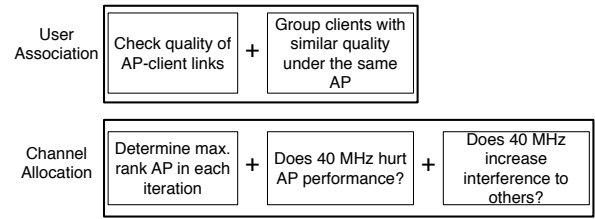


Figure 7: High level functions of ACORN.

distributed coordination function (DCF) used with 802.11 ensures equal long term medium access opportunities. Since poor clients occupy the channel for longer periods, the good clients are hurt. This effect would cause an AP with an associated poor client, to suffer from degraded throughput if it were to apply CB. If instead it does not apply CB, the potential throughput gains are lost. To address this, ACORN tries to ensure that each cell either has (preferably) all users with high quality links, or contains larger numbers of users with poor link qualities. In the former case, the cell can apply CB and in the latter it would simply use a 20 MHz channel.

b) The **CB module** exploits the application of the user association module. It assigns 40 MHz channels to those APs that achieve the highest improvements in throughput (clearly these are the APs that contain the most clients with good quality links). 20 MHz channels are assigned to APs that either suffer degradations in throughput with 40 MHz channels (due to the presence of poor clients) or do not achieve significant gains with CB.

Note that the objective of many of the previously proposed WLAN configuration schemes is to minimize the total transmission delay [17] [18]; this achieves fairness among the users. However, our objective is to *maximize the total network throughput*. In other words, we tradeoff some level of fairness for significant gains in the total network-wide throughput. This is the current trend in many commercial platforms for wireless communications (e.g., 3G/CDMA communications), which employ schedulers that give priority to high quality links (e.g., PF scheduler) [25] [26]; these schedulers target maximizing the total network throughput, allowing for some hit in terms of fairness across the users. In addition, many research efforts align with this direction (for example [27] and [28]). We assume saturated downlink traffic for analytical tractability of ACORN’s decisions. Most of the previous work also relies on this assumption [18] [29]. However, we show experimentally that ACORN helps even with unsaturated loads (as with TCP traffic).

4.1 User Association

A newly arriving client u usually has a set A_u , of serving APs to choose from. In order to pick the *best* AP for association, an objective must be in place. With ACORN, the decision is based on the pairing that achieves the maximum aggregate network throughput.

Gathering information for making user association decisions: In order to achieve our goal, u needs to know for every AP $i \in A_u$ the per client throughput of i , with and without u associated with it ($X_{w,u}^i$ and $X_{wo,u}^i$, respectively).

The client computes these values by obtaining a modified beacon. This beacon includes the number of clients associated with the AP (including u) K_i , the transmission delays per client d_{cl}^i , the aggregate transmission delay ATD_i of the AP and the channel access time M_i of the AP (if there is fully saturated traffic and no contention $M_i = 1$). Client u upon beacon reception calculates the above quantities as, (i) $X_{w,u}^i = \frac{M_i}{ATD_i}$ and (ii) $X_{wo,u}^i = \frac{M_i}{ATD_i - d_u^i}$ [17].

In order for the AP to be able to include these information in the beacon the user has first to associate with it. Other, more simplistic approaches for AP selection, do not require prior client association with the APs. For example, affiliation decisions that are based on the received signal strength (RSS) of the beacons, do not require each user to associate with the APs in range first. However, it is shown that cross layer information is important for user affiliation to deliver a good performance [29]; for instance simply looking at the RSS can lead to configurations with a few overloaded APs and other underloaded APs. In order to obtain the required information accurately, a user needs to associate with the AP and exchange traffic. Thus we implement a similar approach to the one proposed in [17] [18] to obtain the information required by our algorithm. Note however that with ACORN the decision on which AP to affiliate with, is different.

Associating with an AP: Based on the information gathered from all the APs in range, u picks the AP $i^* \in A_u$, which maximizes the following utility function with respect to i :

$$U_{asoc}(u, i) = K_i \cdot X_{w,u}^i + \sum_{j \in A_u, j \neq i} (K_j - 1) \cdot X_{wo,u}^j \quad (4)$$

The first term on the right hand side of Equation 4 is the total throughput of the AP with which, u associates. The second term is equal to the total throughput achieved by the other APs in the range of u . Note here that K_j was defined as the number of clients associated with AP j , including client u . When a new client joins the cell, there could be a reduction in total throughput due to the increased transmission delay. The goal is to minimize this reduction and preferably maintain the throughput at the level prior to the client's arrival. Equation 4 minimizes the impact of a poor client v , in the network. v affiliates with an AP serving similar quality clients, since this association will minimize the total network throughput degradation (i.e., maximize the total network throughput) that could result from the 802.11 performance anomaly [4]. Clients with high quality links do not affect the throughputs due to the anomaly when they associate with their best APs. The pseudocode of the user association scheme is given in Algorithm 1.

4.2 Channel Bonding Decision/Allocation

Next, we describe the channel allocation module of ACORN. Table 2 enlists the notation used.

Problem Formulation: The channel allocation problem is cast as a graph coloring problem [14]. We apply the idea of the *interference graph (IG)* (also considered in [14] [13]). The set V of vertices of the interference graph $G(V, E)$ are the APs. An edge $e_{ij} \in E$, if APs i and j interfere with each other. In the classic graph coloring problem [30], the

Algorithm 1 User Association Algorithm

- 1: **Input:** $K_i, M_i, ATD_i, \forall i \in A_u$ and $d_u^i \forall i$'s clients
 - 2: **Output:** AP i that client u associates with
 - 3: **for** $i \in A_u$ **do**
 - 4: $X_{w,u}^i = \frac{M_i}{ATD_i}$
 - 5: $X_{wo,u}^i = \frac{M_i}{ATD_i - d_u^i}$
 - 6: $U_{asoc}(u, i) = K_i \cdot X_{w,u}^i + \sum_{j \in A_u, j \neq i} (K_j - 1) \cdot X_{wo,u}^j$
 - 7: **end for**
 - 8: **return** $i^* : U_{asoc}(u, i^*) \geq U_{asoc}(u, i), \forall i \in A_u$
-

V	Set of Access Points
Ch	Set of available 20/40MHz channels
$F : V \rightarrow Ch$	Channel assignment mapping
f_i	channel assigned at AP i
X_i	Throughput of AP i

Table 2: Notations for channel allocation algorithm.

objective is to assign colors from a given set of colors to the vertices of the IG, such that no adjacent vertices have the same color. Note here that, the notion of *color conflicts* differs in our case due to CB. For instance, let us assume colors c_i and c_j , corresponding to channels f_i and f_j (20MHz channels). Then the *composite color* $\{c_i, c_j\}$ corresponds to the 40 MHz channel derived from the combination of f_i and f_j . With this set up, the basic colors c_i and c_j do not conflict; however, each of them conflicts with the composite color $\{c_i, c_j\}$. In the rest of the paper we will refer to both basic and composite colors simply as colors. We relax the constraint of the above graph coloring problem, since our objective is to assign colors to the vertices (i.e., channels to the APs), so as to maximize the total network throughput. Formally, we seek to:

$$\max_F Y = \sum_{i \in V} X_i(F) \quad (5)$$

NP-completeness: The graph coloring problem with the above objective is NP-complete. The classic, NP-complete, **decision graph coloring problem** [30], tries to answer the following: "Given a graph $G(V, E)$ and k colors, can we color the vertices V such that $\forall e_{ij} \in E \rightarrow f_i \neq f_j$ ", where f_i is the channel assigned to AP i .

The total aggregate network throughput Y , is upper bounded by $Y^* = \sum_{i \in V} X_i^{isol}$, where $X_i^{isol} = \max\{X_i^{isol-20}, X_i^{isol-40}\}$

is the maximum possible throughput for AP i in an interference-free setting, achieved with either a 20 or a 40 MHz channel. Let F' be a solution to our problem, yielding a total aggregate network throughput of Y' . The solution is optimal **iff** $Y' = Y^* \Leftrightarrow G$ has a k -coloring. Thus, our problem is NP-complete.

Our approach: We design an algorithm that allows APs to decide upon the channel(s) to use. Later, we compute the approximation ratio of our algorithm relative to the optimal. Initially, all APs are assigned either a 20 MHz or a 40 MHz

channel at random. The algorithm is iterative and executed with a periodicity of T . In every iteration, the AP that can provide the maximum increase in the aggregate throughput by switching channels, is allowed to switch. The algorithm terminates, after a number of iterations K , either when no further improvement can be provided or when the improvements provided are very small. A pseudocode for our algorithm is given in Algorithm 2. Y_x^W represents the long term aggregate network throughput achieved at period W and after x iterations of the algorithm.

Algorithm 2 Channel Bonding Selection Algorithm

```

1: Input:  $Y_k^{T-1}, F^{T-1}$ 
2: Output:  $F^T$ 

3:  $Y_0^T = Y_k^{T-1}, k = 0$ 
4: repeat
5:   label (1)
6:    $AP = V, AP' = \emptyset$ 
7:   for  $i \in V$  do
8:      $k = k + 1$ 
9:     for  $c \in Ch$  do
10:       $Tmp_i(c) = \sum_{a \in V} X_a(F_{j \in AP'}^T, f_i = c, F_{j \in AP}^{T-1})$ 
11:    end for
12:    pick  $c_i^* : Tmp_i(c_i^*) \geq Tmp_i(c), \forall c \in Ch$ 
13:     $rank_i = Tmp_i(c_i^*) - Y_{k-1}$ 
14:    if  $\max_{i \in V} rank_i < 0$  then
15:      if  $|AP'| \leq 1$  then
16:        return  $F = (F_{j \in V}^{T-1}, F_{j \in V}^T)$ 
17:      else
18:        GOTO (1)
19:      end if
20:    else
21:      “winner” is AP  $m : rank_m \geq rank_n, \forall n \in V$ 
22:       $f_m^T = c_i^*, AP = AP / \{m\}, AP' = AP' \cup \{m\}$ 
23:       $Y_k^T = \sum_{a \in V} (F_{j \in V}^{T-1}, F_{j \in V}^T)$ 
24:    end if
25:  end for
26:  if  $AP \neq \emptyset$  then
27:    GOTO (1)
28:  end if
29: until  $Y_k^T < \epsilon \cdot Y_{k-1}^T$ 

```

In each iteration of the algorithm, the APs that have still not chosen new channel(s) (i.e., members of the set AP), estimate the aggregate throughput achieved with every possible channel, assuming that the other APs keep their current allocation (line 10).

Estimating throughput: In order to estimate the throughput on a new channel, an AP needs to take into account (i) the number of APs already residing on this new channel and (ii) the quality of the links to its clients on the channel. The first requirement is possible either with help from an administrative authority or the Inter Access Point Protocol (IAPP) [31]. For the second requirement, we assume that the link quality

on the different channels (of the same width) is not significantly different. Later in this section, we provide measurements that validate this assumption in indoor slowly varying settings, typical for enterprise 802.11n WLANs. However, this assumption does not hold when channels are of different width. In other words, the channel qualities to the clients may change if a channel of different width is used. To map the measured results from a 20 MHz channel on to a 40 MHz channel (or vice versa), we leverage the understanding obtained from our PHY layer measurements in Section 3.

To estimate the link quality on a channel of different width, we adopt the following procedure. The input is the SNR at the current width. When we change the width (20/40MHz), there is a 3dB change in the SNR; this processing is performed by a *SNR calibration* module in our estimator. Using this calibrated SNR value, a *BER estimation* module calculates the theoretical coded BER (from [19]). Recall, from our PHY layer measurements that one can expect a reasonable match between the values computed with the theoretical formulas in [19] and the real experimentally observed BER.

Finally, using the BER knowledge we estimate PER. Here, we use the commonly used assumption (for example [32]), of independent, uniformly distributed bit errors within the packet and compute the PER as:

$$PER = 1 - (1 - BER)^L. \quad (6)$$

Note here that ACORN does not require the exact BER or PER values; it only needs a coarse estimate of the link quality i.e., a reasonable classification of good and poor links.

Once the possible performance on each of the available channels is estimated, the AP that can provide the maximum increase in the aggregate throughput by switching channels does so. This approach in essence, mimics the gradient descent algorithm; it greedily seeks to find the point that exhibits the maximum increase in the value of the objective function (the throughput).

The same procedure is repeated considering the APs that have not had an opportunity to switch. When no improvement is possible or the improvement is incremental (ϵ - line 29), the algorithm stops. In the current form of ACORN, $\epsilon = 1.05$ (i.e., if there is a 5% or less increase in the total aggregate throughput as compared to the previous iteration, the algorithm stops).

Approximation Ratio: Gradient-based optimization can be trapped in a local extremum. Given also that the channel allocation problem is NP-complete, we are interested in finding the worst case approximation ratio of our algorithm.

The maximum possible aggregate throughput is obtained when all APs operate in an interference free setting, and is equal to: $Y^* = \sum_{i \in V} X_i^{isol}$ as mentioned before. The worst

local extremum where our algorithm can be trapped, is the one in which every AP uses the same (20 or 40 MHz) channel. In other words, nodes are not just assigned conflicting colors, but are assigned exactly the same color. This is because, if they are assigned different conflicting colors (a composite color and a basic color) the achieved throughput will be higher (this is easy to verify). In this case, the throughput in isolation of every AP u will be reduced by




Figure 8: Link quality on different channels.

a factor of $\frac{1}{deg_u+1}$, assuming fully saturated traffic, where deg_u is the degree of node u ⁵. Consequently the long term total network throughput will be $Y_{worst} = \sum_{i \in V} \frac{1}{deg_i + 1}$.

$X_i^{isol} \leq \frac{1}{\Delta + 1} \cdot \sum_{i \in V} X_i^{isol}$, where Δ is the maximum node degree observed in the network. Thus, our algorithm exhibits a worst case approximation ratio of $O(\frac{1}{\Delta+1})$. As shown in Section 5, in practice, the channel bonding/allocation scheme performs much better.

Link quality on different channels: We assume that the quality of a link does not exhibit significant variations in terms of PER on different channels of the same width. In order to verify this, we conduct the following experiments on our testbed. For all the links, we measure back to back, the Packet Error Rate on the different channels, using the maximum transmission rate (MCS = 15 [1]). Figure 8 presents the results from 3 representative links. Our measurements demonstrate that, indeed, the variations across the different channels are negligible (for both 20 and 40MHz channels), making our assumption realistic. There are studies (e.g., [9]), that have reported variations in the link quality with different frequencies. However, these studies are on single antenna systems. In contrast, in our experiments, we utilize the MIMO PHY of 802.11n. The use of MIMO makes the performance stable and robust and decreases variations across the different frequencies (arising primarily due to fading in single antenna systems). ACORN can easily be modified, such that each AP scans (one at a time) all the available channels and gets more accurate information regarding the link quality to its clients. However, this would add more complexity and increase the convergence time of the system.

Periodicity of our algorithm: The periodicity T with which we apply channel allocation needs to be carefully chosen. If we apply it too often, then the hit in the throughput could be significant due to the overhead. If on the other hand, we activate channel allocation too infrequently, the topology might have significantly changed in the interim and the current allocation might provide poor performance.

In order to assess this tradeoff, we use data collected from 206 different (commercial) APs, in a time period spanning

⁵Note here that, the graph we consider is the interference graph of the network [14] with respect to the APs (vertices of the graph). Two APs interfere with each other either if they directly compete for the medium or if either competes with at least one of the other AP's clients.

more than 3 years from the CRAWDAD repository [33]. In particular, we extract the association duration of each user. Figure 9 depicts the CDF of the association duration. More than 90% of the associations last less than 40 minutes and the median is approximately 31 minutes. Based on these data, we run our channel allocation algorithm every 30 minutes.




Figure 9: CDF of user association durations.

5. EVALUATING ACORN

In this section, we provide implementation details of ACORN and its experimental evaluation on our 802.11n WLAN testbed.

5.1 Implementation Details

We implement our algorithms using the Click modular router (v1.6) [34]. We implement a user-level utility that runs both at the APs and the clients. We keep track of the SNR, the nominal rate and the association time per client by using dedicated functions implemented in our card's driver. The delay for each client is calculated and broadcast in a beacon as described in Section 4, along with the M_a values (defined in Section 4.1), the number of clients and the aggregate transmission delay of an AP. The client receives the beacons from every AP in its range and makes appropriate association decisions.

Calculating per client transmission delay and M_a : These metrics are not directly available since our hardware does not provide access to firmware. We implement a module where every AP calculates the delay of a client by utilizing our PER estimation procedure and the nominal rate for the client. We estimate M_a for an AP a by $\frac{1}{|con_a|+1}$ where con_a denotes the set of neighbor APs residing on the same channel as AP a . This estimation has very high accuracy when these APs can hear each other under saturated traffic. Accurate management and configuration of WLANs is of most interest in these regimes i.e., in dense deployments and with heavy loads.

5.2 Experimental Evaluations

Comparison with legacy 802.11 WLAN configuration systems: We start by randomly assigning initial channels to APs from the 5GHz band. Clients are then randomly activated one by one. Each client performs user association (i) as per **Algorithm 1** or, (ii) the algorithm described in [17]. The APs then perform channel selection either as per **Algorithm 2** or a modified version of [17]. We modify the frequency selection algorithm in [17] to implement a greedy strategy where APs aggressively use the (single width) 40 MHz chan-

	ACORN	Random Configurations (Descending order)
Network Tput (UDP)	259.2	201.63, 193.1, 188.56, 187.6, 184.62, 183.39, 169.62, 163.32, 160.47, 159.35
Network Tput (TCP)	178.93	161.7, 155.77, 134.78, 133.4, 130.64, 114.1, 109.4, 106.6, 103.41, 102.3

Table 3: ACORN achieves the highest network throughput (in Mbps) against 10 best (out of 50) manual configurations.

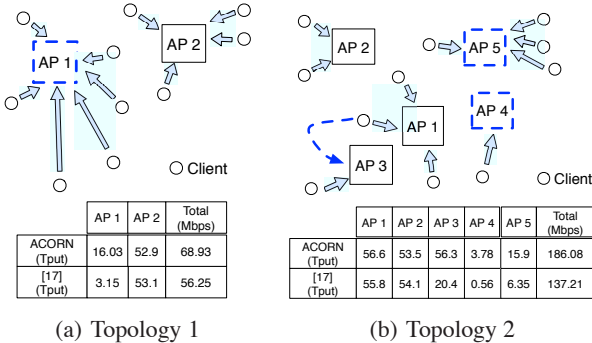


Figure 10: ACORN can provide throughput gains in interference-free deployments. Dashed APs use 20 MHz with ACORN and 40 MHz otherwise. Dashed arrows depict the user affiliation decisions taken by [17].




Figure 11: ACORN provides the highest throughput in settings with interference. X, Y, Z denote the channel widths (MHz) used by APs 1, 2 and 3.

nels⁶. Specifically, they scan 40MHz channels and select the one that minimizes the total noise and interference. We simply refer to this scheme as “[17]”.

ACORN significantly improves per-cell throughput in interference free settings: In these experiments, we activate saturated downlink UDP traffic from each AP to its clients. We evaluate ACORN on many different WLAN topologies. We employ all the twelve 20MHz channels available in the 5GHz band with both ACORN and [17]. Figure 10 quantifies our per-AP throughput observations with a few sample topologies (also depicted). Topology 1 consists of 2 APs. This is a sparse WLAN where clients are connected with poor quality links with AP1 and good clients are associated with AP2. We find that the user association with both ACORN and with [17] are identical. However, the use of the 20 MHz band provides a significant increase (4x throughput increase) for AP 1 because of its low-SNR client links.

⁶As one might expect, simply using the algorithm with 20 MHz channels results in lower rates and correspondingly lower throughput. For clarity and ease of discussion, we omit these results here.

In fact, with the 3 dB reduction in SNR, we observe that the poor clients are hardly able to communicate with the AP when it uses CB with [17].

Topology 2 includes 5 APs. We observe that with ACORN, when poor clients associate with an AP, the AP uses a 20 MHz channel. The same APs apply CB greedily with [17]. The presence of poor clients reduces the cell throughput of the corresponding AP. These effects are seen with AP 4 and AP 5 of the topology and for these, ACORN provides significant throughput improvements (6x for AP 4 and 1.5x for AP 5). We also observe that ACORN results in different user associations for APs 1 and 3; as discussed, ACORN tries to group clients with similar link qualities in the same cell. In contrast, [17] evenly divides the clients between these APs regardless of the specific client link qualities. Due to this, AP 3 achieves a higher throughput (1.8x) with ACORN since it serves only one good quality client. This results in more congestion at AP 1 with ACORN as compared to [17] since it has to serve more clients. However, interestingly, AP 1 can still achieve the *same total throughput with more clients*. We identify the reason behind this to be the following: since ACORN groups similar-quality clients in one cell, the aggregate throughput does not change despite the fact that *per-client* throughput is reduced; the performance anomaly of 802.11 does not take effect. We observe similar performance gains with ACORN in a variety of other deployments but do not present the results due to space constraints.

ACORN reduces interference: In dense deployments where channel availability is limited, ACORN reduces interference at the neighboring APs and provides even higher improvements in throughput. To showcase this, we experiment with a representative scenario in Figure 11 where the number of APs is not small relative to the number of available channels (unlike in the previous experiments). We have 3 APs that contend for channel access and there are four 20 MHz channels available. AP 1 serves a good quality client and APs 2 and 3 have poor clients associated with them. When all the APs aggressively use CB, they project interference on each other. In addition, APs 2 and 3 suffer because of the poor client links. Note that, with 4 channels, only one AP can use CB to achieve complete isolation. In such cases, it is essential for a channel allocation scheme to identify the best AP that should exclusively use CB. We observe that, ACORN identifies this AP and provides the highest throughput as compared to other possible allocations. It provides an almost 2x improvement over the scheme that aggressively allows CB operations at every AP; the aggressive allocation causes increased interference and thus, lower throughput.

Comparison with random manual configurations: We compare ACORN against a large set of manual configurations in terms of assigned channels and user associations. The purpose of this experiment is to see the effectiveness




Figure 12: Trajectory of our mobile client .




Figure 13: ACORN tracks the link quality and selects the channel width that gives the higher throughput over fixed channel widths.




Figure 14: ACORN can achieve approximation ratios better than $O(\frac{1}{\Delta+1})$ in practice.

of ACORN in yielding the highest network throughput. In this experiment (different from previous experiments), we measure the performance in terms of both UDP and TCP throughputs. The experiments evaluate ACORN’s performance with unsaturated TCP traffic; since our analysis assumes saturated traffic, we wish to examine if ACORN works well in these cases as well. Note that, we are mainly interested in scenarios that fully load the network (saturated conditions), since achieving balance and efficacy is most important in these regimes. However, these limited experiments demonstrate the applicability of ACORN in more generic settings. Tuning ACORN to provide benefits in all possible scenarios (if possible), is part of our future investigations.

For a randomly picked topology, we first run ACORN and obtain the total network throughput. Next, we configure APs with random channels (both 20 and 40MHz) and let each client associate with one of the APs in range with equal probability. We repeat this experiment for 50 different configurations and take the 10 best configurations. Table 3 tabulates our results. We observe that ACORN configures the network in a way that achieves the highest possible throughput as compared to what is achieved with these random configurations. We wish to point out that ACORN provides gains with TCP since congestion can still occur at shorter time scales; at such times, the use of ACORN provides benefits. The decisions relating to CB deliver higher performance since they are based on client link qualities and independent of the type of traffic. Although the set of random configurations is by no means exhaustive, the experiments do demonstrate the efficacy of ACORN in terms of yielding high throughputs.

How close to the optimal is ACORN channel allocation in practice?:

Next, we perform experiments to examine the approximation ratios achieved by the ACORN channel allocation module in practice. We choose 3 APs that *compete* for channel access in each experiment (i.e., $\Delta = 2$); 9 such sets of different APs are considered. We then associate clients with each of these APs. We then run saturated UDP downlink traffic from each AP to its clients in *isolation* for both 20 MHz and 40 MHz channels. We calculate Y^* , which is the best possible aggregate throughput, as

$$\sum_{i=1,2,3} \max(T_{20}^i, T_{40}^i),$$

where T_x^i is the throughput obtained by AP i using a channel width of x MHz. Note here that this maximum is achieved when we completely isolate the 3 competing APs i.e., they do not contend with each other. Subsequently, we run the ACORN channel allocation algorithm with 2, 4 and 6 orthogonal channels available. Note here that, 6 orthogonal channels are enough for all of the APs to simultaneously activate CB. Figure 14 depicts the total network throughput, T , obtained by ACORN in comparison with Y^* . Note that Y^* computed as above, is a *loose* upper bound, since complete isolation of the 3 APs is not always possible with less than 6 orthogonal channels. With 2 channels, ACORN does not perform worse than what is theoretically predicted; the aggregate network throughput is $\frac{Y^*}{3}$, since the medium access is shared among the contending APs. In the case of 6 channels, ACORN can achieve Y^* , since channel allocation isolates every AP and configures the best channel width for each AP. We observe some cases where ACORN performs very close to the optimal (what is possible with 6 channels) even with only 4 channels. Examining the cases with care, we find that there is at least one AP i such that $T_{20}^i > T_{40}^i$; ACORN accounts for this and configures the particular AP with a 20 MHz channel, leaving 3 channels for utilization to the other two APs.

Evaluating ACORN with mobility: With ACORN, once an AP is assigned a 40 MHz channel, it can opportunistically use its allocated channels. In other words, the AP can opt out from using CB and only employ the 20 MHz channel (one of the two assigned). Since the other APs choose their frequencies based on the channels assigned to this particular AP, using either of the two 20 MHz channels will not change the interference on the neighboring APs. This mode of opera-

tion is particularly desirable in WLANs with mobile clients. Since an AP to client link quality can vary temporally, the AP can dynamically activate the desired width of operation (20 vs 40 MHz) with ACORN based on the measured link qualities of its clients. We experiment with a scenario that involves pedestrian mobility. We configure a laptop with the same card that we use in our testbed and use it as a mobile client. In this experiment, we have a single AP that has 2 static clients in addition to the laptop. First, we position the laptop close to the AP and start moving it away from the AP. Figure 12 depicts this trajectory with dark arrows. The AP transmits downlink UDP traffic to its clients and we record the aggregate throughput measured as a function of time. Figure 13(a) presents the time trace of the aggregate cell throughput during the mobile client's movement. We compare ACORN against a configuration that strictly uses a 40 MHz channel. We observe that ACORN uses the 40 MHz channel in the beginning and sustains this until the point where the link quality becomes poor for the mobile laptop (around 30 sec). From that point until the end of the experiment (the client stops at a location far from the AP), ACORN falls back to the 20 MHz mode and is able to sustain a cell throughput that is almost ten times that of a fixed 40 MHz channel. Note here that the poor quality link to the distant client affects the good clients as well due to the 802.11 performance anomaly [4]. In a similar experiment, we have our client moving towards the AP, the trajectory is depicted with striped arrows in Figure 12. We compare ACORN against a fixed 20 MHz configuration. We observe that ACORN uses the 20 MHz until it recognizes the improvement in link quality (SNR); at that time it switches to a 40 MHz channel (at around 10 sec) and is able to utilize the gains from CB.

6. CONCLUSIONS

Channel bonding (CB) allows 802.11n nodes to use a wider frequency band towards achieving high data rates. We show that applying CB has associated caveats and can degrade performance even in interference-free settings. We perform extensive experiments to understand why this is the case. Based on the understanding gained, we design and implement ACORN, an auto-configuration framework for WLANs. We show via extensive experimental evaluations that ACORN provides significant per-AP throughput gains (as high as 6x) over what is achieved when we apply prior proposed approaches designed for legacy 802.11 systems that are based on using single channel widths. To the best of our knowledge, ACORN is not only the first system to consider the application of CB in enterprise WLANs but also the first system to consider the use of two distinct bandwidths while performing joint channel allocation and user association.

7. REFERENCES

- [1] IEEE 802.11n Standard 2009.
- [2] V. Shrivastava, S. Rayanchu, J. Yoon, and S. Banerjee. 802.11n Under the Microscope. In *ACM IMC*, 2008.
- [3] V. Visoottiviset, T. Piroonsith, and S. Siwamogsatham. An Empirical Study on Achievable Throughputs of IEEE 802.11n Devices. In *ACM WinMee*, 2009.
- [4] M. Hesse, F. Rousseau, G. Berger-Sabbatel, and A. Duda. Performance Anomaly of 802.11b. In *IEEE INFOCOM*, 2003.
- [5] Wireless Open-Access Research Platform. <http://warp.rice.edu/>.
- [6] Ralink Technology. <http://www.ralinktech.com/>.
- [7] A. Mishra, V. Shrivastava, S. Banerjee, and W. Arbaugh. Partially overlapped channels not considered harmful. In *ACM SIGMETRICS*, 2006.
- [8] L. Hanzo, M. Munster, B.J. Choi, and T. Keller. *OFDM and MC-CDMA for Broadband Multi-User Communications, WLANs and Broadcasting*. IEEE WILEY, ISBN 0-470-85879-6.
- [9] H. Rahul, F. Edalat, D. Katabi, and C. Sodini. Frequency-aware Rate Adaptation and MAC Protocols. In *ACM MobiCom*, 2009.
- [10] C. Tang and V.J. Stulpman. A Gradient-Based Method for OFDM Sub-Carrier Power Allocation. In *IEEE VTC*, 2005.
- [11] Atheros AR5001J. <http://www.atheros.com/news/AR5001J.html>.
- [12] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, and P. Bahl. A Case for Adapting Channel Width in Wireless Networks. In *ACM SIGCOMM*, 2008.
- [13] A. Mishra, D. Agrawal, V. Shrivastava, S. Banerjee, and S. Ganguly. Distributed channel management in uncoordinated wireless environments. In *ACM MobiCom*, 2006.
- [14] A. Mishra, S. Banerjee, and W. Arbaugh. Weighted Coloring based Channel Assignment for WLANs. In *MC2R, Volume 9, Issue 3*, 2005.
- [15] A. Kumar and V. Kumar. Optimal association of stations and APs in an IEEE 802.11 WLAN. In *NCC*, 2005.
- [16] A. Mishra, V. Brik, S. Banerjee, A. Srinivasan, and W. Arbaugh. A client-driven approach for channel management in wireless LANs. In *INFOCOM*, 2006.
- [17] B. Kauffmann, F. Baccelli, A. Chaintreau, V. Mhatre, K. Papagiannaki, and C. Diot. Measurement-Based Self Organization of Interfering 802.11 Wireless Access Networks. In *INFOCOM*, 2007.
- [18] I. Broustis, K. Papagiannaki, S. V. Krishnamurthy, M. Faloutsos, and V. Mhatre. MDG: Measurement-Driven Guidelines for 802.11 WLAN Design. In *ACM MobiCom*, 2007.
- [19] T.S. Rappaport. *Wireless Communications*. Prentice-Hall. ISBN 0-13-042232-0.
- [20] S.M. Alamouti. A simple transmit diversity technique for wireless communications. In *IEEE Journal on Selected Areas in Communications*, October 1998.
- [21] Wi-Fi Wireless LAN Frequency, Bands and Channels. <http://www.networkdictionary.com/Wireless/Wi-Fi-Wireless-LAN-Frequency.php>.
- [22] JPCap Library. <http://netresearch.ics.uci.edu/kfujii/jpcap/doc/>.
- [23] R. Jain. *The Art of Computer Systems Performance Analysis*. John Wiley & Sons. ISBN 0-471-50336-3.
- [24] UCR Wireless Testbed. <http://networks.cs.ucr.edu/testbed>.
- [25] S. Z. Asif. Aligning business and technology strategies-an evolution of a third generation wireless technology. In *Engineering Management Conference*, 2002.
- [26] S. Bali, S. Machiraju, H. Zang, and V. Frost. On the performance implications of proportional fairness (PF) in 3G wireless networks. In *PAM*, 2007.
- [27] D.C. Tsilimantos, D.A. Zorbouti, G.V. Tsoulos, G.E. Athanasiadou, and D.I. Kalamani. Fairness and Throughput Trade-Off Analysis for UMTS WCDMA Network Planning. In *WPC*, 2009.
- [28] D. Panigrahi and F. Khaleghi. Enabling trade-offs between system throughput and fairness in wireless data scheduling techniques. In *WWC*, 2003.
- [29] K. Sundaresan and K. Papagiannaki. The Need for Cross-Layer Information in Access Point Selection Algorithms. In *ACM IMC*, 2006.
- [30] T. R. Jensen and B. Toft. *Graph Coloring Problems*. John Wiley & Sons. ISBN 0-471-02865-7.
- [31] IEEE. Recommended Practice for Multi-Vendor Access Point Interoperability via an Inter-Access Point Protocol Across Distribution Systems Supporting IEEE 802.11 Operation. In *IEEE Standard 802.1f*, July 2003.
- [32] M. Vutukuru, H. Balakrishnan, and K. Jamieson. Cross-Layer Wireless Bit Rate Adaptation. In *ACM SIGCOMM*, 2009.
- [33] CRAWDAD. <http://uk.crowdad.org/meta.php?name=ilesansfil/wifidog>.
- [34] R. Morris, E. Kohler, J. Janotti, and M. F. Kaashoek. The Click Modular Router. In *ACM Symposium on Operating Systems Principles*, 1999.