

# Adaptive Subcarrier Nulling: Enabling Partial Spectrum Sharing in Wireless LANs

Xinyu Zhang and Kang G. Shin

Department of Electrical Engineering and Computer Science

The University of Michigan

Email: {xyzhang, kgshin}@eecs.umich.edu

**Abstract**—Emerging WLAN standards have been incorporating a variety of channel widths ranging from 5MHz to 160MHz, in order to match the diverse traffic demands on different networks. Unfortunately, the current 802.11 MAC/PHY is not designed for the coexistence of variable-width channels. Overlapping narrow-band channels may block an entire wide-band channel, resulting in severe spectrum underutilization and even starvation of WLANs on the wide-band. A similar peril exists when a WLAN partially overlaps its channel with multiple orthogonal WLANs.

In this paper, we propose to solve the problem of partial spectrum sharing using *Adaptive Subcarrier Nulling (ASN)*. ASN builds on the 802.11 OFDM PHY, but allows the radios to sense, transmit, detect, and decode packets through spectrum fragments, or subbands. An ASN transmitter can adapt its spectrum usage on a per-packet basis, by nulling the subbands used by neighboring WLANs, and sending packets through the remaining idle subbands. ASN preserves the 802.11 CSMA/CA primitives while allowing users to contend for access to each subband, and can opportunistically exploit the merits of wide-band channels via spectrum aggregation. We have implemented and evaluated ASN on the GNURadio/USRP platform. Our experimental results have shown ASN to achieve detection and decoding performance comparable to the legacy 802.11. Our detailed simulation in ns-2 further shows that ASN substantially improves the efficiency and fairness of spectrum sharing for multi-cell WLANs.

## I. INTRODUCTION

Most WiFi networks today operate with the default 20MHz bandwidth [1]. This bandwidth has been exhausted in the widely used 802.11g standard to provide up to 54Mbps data rate, but is becoming insufficient for throughput-demanding applications such as high-definition video streaming. The recently ratified 802.11n standard doubles the data rate using 40MHz channel width. The emerging 802.11ac [2] further enables Gbps wireless communications with 80MHz and 160MHz channels. On the other hand, narrow-band channels (5MHz and 10MHz) have also been incorporated in the recent 802.11 standard [1] to support WLANs with low throughput demands but high energy-efficiency requirements [3].

Although a variety of channel widths can be used, the spectrum is still a limited resource. For example, on the 2.4GHz ISM band used by 802.11b/g/n, the total spectrum width is only 83.5MHz. Hence, it is impractical to guarantee orthogonality between the channels used by every co-located WLAN, especially in the current high-density enterprise and public WiFi networks [4]. Thus, a WLAN often needs to share part or all of its spectrum with others. Most WiFi WLANs today reside on the three non-overlapping 20MHz channels 1, 6, and 11 specified by 802.11 [4], and thus, neighboring WLANs tend to be either orthogonal or sharing an entire channel. But as channel widths become more heterogeneous, partial spectrum sharing is unavoidable.

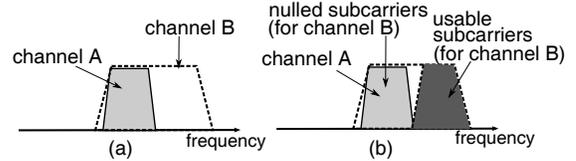


Fig. 1. (a) Partial-channel blocking problem in wireless LANs. (b) Adaptive subcarrier nulling (ASN) nulls the shared busy subband (containing a number of subcarriers) and leverages the non-overlapping subbands to send data.

The current 802.11 relies on CSMA/CA to coordinate transmitters on the same channel, but it is not inherently designed for partial sharing of the spectrum. An 802.11 transceiver treats an entire channel as a whole spectrum block to perform carrier sensing and packet transmission. It has to defer its transmission even if part of the spectrum is occupied (e.g., by a WLAN that has a narrower bandwidth as shown in Fig. 1(a), or resides on a partially overlapped channel). We refer to this problem as *partial-channel blocking*. Partial-channel blocking causes severe under-utilization of non-overlapped spectrum, which should otherwise be able to provide a higher throughput due to less contention. A more critical problem occurs when multiple narrowbands coexist and overlap with a wideband channel. With the 802.11 MAC, the wideband will be able to transmit only if all the narrowbands are idle, resulting in highly unfair channel access opportunities and even starvation of the wideband WLAN.

In this paper, we introduce a new mechanism called *Adaptive Subcarrier Nulling (ASN)*, to enable partial spectrum sharing between WLANs. ASN builds on the OFDM PHY used by 802.11g/n and other emerging standards [5], in which a channel comprises many small spectrum units called *subcarriers*. ASN groups the subcarriers into several *subbands*, and allows neighboring WLANs to share and contend for access to each subband. When a shared subband is occupied by one WLAN, another WLAN can opportunistically *null* the corresponding subcarriers in that subband, and use those non-overlapping subbands to send packets. ASN performs this adaptation on a *per-packet* basis, so as to fully utilize the available spectrum whenever possible, and to ensure fair access to shared spectrum. With ASN, the partial-channel blocking problem can be naturally solved (Fig. 1(b)).

Subcarrier nulling can be realized straightforwardly in the 802.11 OFDM PHY: instead of sending information bits (1 or -1), the transmitter can simply feed 0's to the subcarriers, resulting in zero power on the corresponding spectrum. However, it is nontrivial to ensure the receiver can correctly decode the remaining non-zero subcarriers. Since the transmitter decides on the set of subbands to be used for each packet, the receiver has no prior knowledge of the spectrum to be used by an incoming packet, yet it still needs to detect the packets, synchronize to

them, and then decode the information bits.

ASN meets these challenges by redesigning the preamble structure, packet detection and decoding algorithms in 802.11. It uses correlation-based algorithms to detect a packet and identify the subbands used by it. It further adapts the pilot-based approach in 802.11 to estimate the channel, and then decodes the bits carried by each subcarrier. In addition, ASN combines the time-domain energy sensing with frequency-domain spectrum sensing, so that a transmitter can identify the spectrum currently in use by neighboring WLANs. Although similar PHY layer problems have been addressed in non-contiguous OFDM (NC-OFDM) communications systems [6]–[9] (more details available in Sec. VI), ASN represents a complete 802.11 based NC-OFDM design that solves a network-level problem, *i.e.*, partial spectrum sharing for WLANs.

At the MAC layer, ASN retains the carrier sensing and backoff mechanism in 802.11, but makes the busy/idle decision based on the time/frequency domain spectrum sensing. ASN maintains a backoff counter for each subband, and allows decrementing the backoff counter if at least one subband is idle. This simple extension to 802.11 CSMA/CA (referred to as *ASN with direct access*, or ASN-DA) alleviates the partial channel blocking problem, but may cause certain transmitters to dominate a subband. Therefore, we propose an alternative protocol, *ASN with water filling access* (ASN-WF), which aligns the busy time of subbands by adapting the packet size, thereby balancing the access opportunities of different WLANs to shared subbands.

We have implemented an ASN prototype on the GNURadio/USRP platform. Our experimental results show that ASN can sense, synchronize, and decode partial spectrum, with a level of accuracy comparable to the legacy 802.11g that uses a full spectrum. We further use detailed simulation in ns-2 to evaluate ASN in multi-channel, multi-cell wireless LANs. Our experiments demonstrate that ASN significantly improves the throughput and fairness of spectrum sharing. In particular, when two WLANs of different widths coexist, it improves the total network throughput by up to 147.7%, by solving the partial-channel blocking problem. When multiple narrowband WLANs coexist with a wideband WLAN, ASN enables close-to-equal access to shared spectrum, providing an order of magnitude of throughput improvement for the wideband WLAN that tends to be starved by 802.11.

The remainder of this paper is organized as follows. In Sec. II, we experimentally study the problems caused by partial spectrum sharing and analyze the reasons behind them. In Sec. III, we introduce the detailed design of ASN’s channel sensing, detection, and decoding algorithms. Sec. IV describes the two medium access protocols for ASN. Sec. V presents the implementation and evaluation of ASN. Sec. VI discusses related work and finally, Sec. VII concludes the paper.

## II. MOTIVATION

The problem of partial spectrum sharing is akin to the well-explored effects of partially-overlapping channels in 802.11b WLANs [10]. In the 2.4GHz ISM band for 802.11b/g/n, 11 channels of 20MHz bandwidth each can be used, and adjacent channels’ center frequencies are separated by 5MHz. Hence,

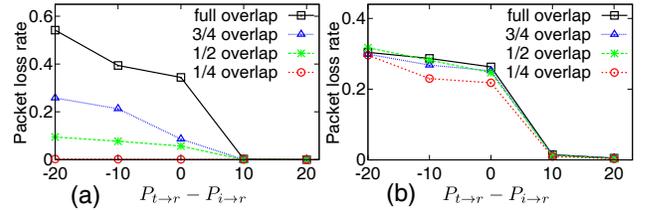


Fig. 2. Effects of partial-channel interference for 802.11b and 802.11g.

neighboring WLANs may have  $\frac{1}{4}$ ,  $\frac{1}{2}$ ,  $\frac{3}{4}$  or full overlap, if any. For 802.11b, interference from partially-overlapping channels is proportional to the amount of overlap, which may be much less than a full overlap, and thus, partially-overlapping channels can be simultaneously active in many cases [10]. However, does this apply to 802.11g, which builds on a distinct PHY layer? In this section, we answer this question with detailed experiments, and then discuss the advantages of ASN in OFDM WLANs.

### A. Partially-Overlapping Channels for 802.11b and 802.11g

We measure the interference caused by partially-overlapping channels using a small testbed that consists of a transmitter ( $N_t$ ), receiver ( $N_r$ ) and interferer ( $N_i$ ), which are laptops equipped with Atheros 5414 802.11b/g NIC, running on Mad-WiFi trunk-r4134.  $N_t$  and  $N_r$  use the same 20MHz channel, while  $N_i$  resides on a 20MHz channel that partially or fully overlaps with them, and its carrier sensing function is disabled.  $N_t$  continuously transmits ICMP Ping-broadcast packets to  $N_r$  at 100 pkts/second with packet size 1.4KB, while at the same time  $N_i$  emits Ping-broadcast packets with the same rate and size. We adjust the transmit power of  $N_t$  and  $N_i$ , thus varying the relative power received by  $N_r$  (denoted by  $P_{t \rightarrow r}$  and  $P_{i \rightarrow r}$ ) when  $N_t$  and  $N_i$  use the same channel.

Fig. 2(a) shows the packet loss rate of 802.11b (with 2Mbps data rate) subject to interference. When  $P_{t \rightarrow r} - P_{i \rightarrow r} < 10$ dB, interference may become detrimental to the data transmission. However, different fractions of channel overlap between  $N_i$  and  $N_t$  lead to disparate loss rates. When  $P_{t \rightarrow r}$  is 20dB lower than  $P_{i \rightarrow r}$ , a fully-overlapped channel suffers 54.7% loss, whereas a  $\frac{1}{4}$ -overlapped channel has nearly 0 loss. This result is consistent with existing measurements of 802.11b [10].

However, for 802.11g, packet loss rate is almost invariant to the channel overlap (Fig. 2(b)), *i.e.*, the effect of interference from a  $\frac{1}{4}$ -overlapped channel is comparable to that from a fully-overlapped channel. Therefore, existing approaches that exploit concurrent transmissions from partially-overlapping channels [10] are not applicable to 802.11g.

The distinct effects of partially-overlapping channels for 802.11b and 802.11g root in their PHY layers. The 802.11b PHY is based on DSSS (direct-sequence spread spectrum), which spreads one bit of information over an entire spectrum of 20MHz. Its SINR equals the total power of the non-interfered spectrum divided by that of the interfered part. For example, even when  $P_{t \rightarrow r} = P_{i \rightarrow r}$ , the resulting SINR is up to  $10 \log_{10}(4) = 6$ dB when  $\frac{1}{4}$  of  $N_i$ ’s spectrum overlaps with  $N_t$ . This SINR is enough to ensure close to 100% decoding probability at a low modulation level (*e.g.*, BPSK) [11].

In contrast, for the OFDM PHY used by 802.11g, a 20MHz channel is divided into 64 spectrum units (*i.e.*, *subcarriers*), each carrying one (or more) bits of information. An 802.11g packet comprises multiple *OFDM symbols* each occupying the

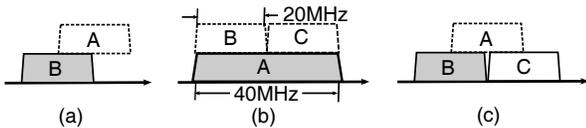


Fig. 3. Heterogeneous channel width or partially-shared channels cause inefficient or unfair spectrum usage in 802.11.

64 subcarriers and transmitted consecutively over time. When  $N_t$  and  $N_i$ 's spectrum overlap by  $\frac{1}{4}$  and  $P_{t \rightarrow r} = P_{i \rightarrow r}$ ,  $\frac{1}{4}$  of the subcarriers in each OFDM symbol will have an SINR of  $10 \log_{10}(1) = 0\text{dB}$ , which are unlikely to be correctly decoded. Equivalently,  $\frac{1}{4}$  of an interfered packet will be corrupted and is unlikely to be recovered. This is the reason why energy sensing is mandatory in 802.11g [1, Sec. 17.3.10.5] to sense and prevent interference from partially-overlapping channels.

### B. Why ASN?

Given that partially-overlapping channels in OFDM WLANs cannot transmit concurrently, coexistence of multiple channels faces several critical challenges, which can be solved by ASN.

1) *Partial-channel blocking*: As discussed in Sec. I, the partial-channel blocking problem occurs in the 802.11g WLAN when part of the channel is used by a co-located narrowband WLAN, and hence, the entire channel must suspend its transmission. In the example of Fig. 1, suppose WLAN A and B are 20MHz and 40MHz, respectively. Both transmit packets of the same size, but the transmission takes only 1 time slot for the 40MHz, and 2 for the 20MHz channel. Using 802.11, both WLANs have an equal chance to access the medium, resulting in mean spectrum utilization of  $(20 \times 2 + 40 \times 1) \frac{1}{3} \approx 26.7\text{MHz}$ . In contrast, ASN can activate the right half of channel B even when A is transmitting, thus maintaining 40MHz spectrum utilization at any time. With respect to individual spectrum usage, for legacy 802.11, the 20MHz channel would achieve  $\frac{20 \times 2}{3} \approx 13.3\text{MHz}$ , and the 40MHz channel achieves  $\frac{40}{3} \approx 13.3\text{MHz}$  — clearly, the 40MHz WLAN does not gain advantage when coexisting with a 20MHz one. With ASN, the 20MHz WLAN still has 13.3MHz channel utilization, but the 40MHz WLAN achieves  $\frac{20+20+40}{3} \approx 26.7\text{MHz}$ , thereby doubling its throughput.

A similar scenario occurs when two channels partially overlap, as shown in Fig. 3(a). Since the legacy 802.11g can only activate one channel at a time, its spectrum usage is only 20MHz. In contrast, by nulling the overlapping subcarriers and reusing the non-overlapping ones, ASN fully exploits the 30MHz spectrum, improving spectrum utilization by 50%.

2) *Channel starvation*: The CSMA mechanism in 802.11 may starve a wideband WLAN when it coexists with multiple narrowband WLANs. Fig. 3(b) illustrates the case when a 40MHz channel A partially overlaps with two orthogonal 20MHz channels B and C. With 802.11, A can transmit only if both B and C are idle, which occurs only when B and C finish their transmission approximately at the same time, and subsequently A wins the contention over both. Clearly, this is a rare case when B and C have backlogged traffic, so A will remain starved most of the time, although nominally it should have a higher throughput with larger bandwidth.

In general, such starvation effects occur whenever a WLAN partially shares spectrum with several other orthogonal WLANs (e.g., the scenario in Fig. 3(c)). Using ASN, the vulnerable

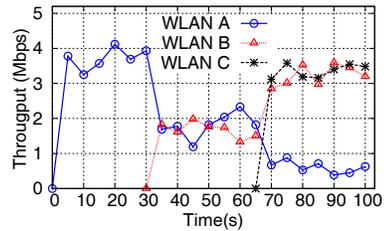


Fig. 4. Throughput evolution of 3 partially-overlapping WLANs.

WLAN can opportunistically null the busy part of the spectrum, and access the idle part, thus preventing starvation. It might seem feasible to achieve the same result by directly reducing the channel width of A and relocate it to the idle part of the spectrum. In practice, however, the channel switching time is in the order of several packets' duration [12], and the channel status may have already changed after relocating the channel. In ASN, a transmitter fixes its center frequency and maximum bandwidth, and performs subcarrier nulling on a per-packet basis, thus it needs not switch the channel and wait for the radio to stabilize.

3) *Experimental validation*: To validate the above motivating problems, we measure the throughput of three partially-overlapping WLANs running 802.11g (i.e., the scenario in Fig. 3(c)). Each WLAN consists of an AP and a client, with saturated downlink transmission, 6Mbps data rate and 1KB packet size. WLAN A, B and C are activated at 0s, 30s, and 65s, respectively. Fig. 4 shows the resulting throughput over time. When A and B are activated, only one of them can transmit at any time. Although they occupy 30MHz channels in total, the total throughput is similar to that of a single 20MHz WLAN (equivalent to the scenario in Fig. 3(a)). After all WLANs are activated, A's throughput is only around 17% of the two competitors' (the scenario in Fig. 3(c)). The same starvation effect would occur for the scenario in Fig. 3(b), where the access opportunity of A remains the same as in Fig. 3(c). Clearly, 802.11 results in inefficient and unfair spectrum usage in the presence of partially-shared channels.

## III. OFDM SUBCARRIER NULLING

A key challenge in realizing ASN is to ensure a node can sense partially-used channels, and can detect, synchronize, and decode a packet, without knowing in advance the spectrum used by the transmitter. In this section, we present the detailed design of ASN to address this challenge.

### A. ASN: An Overview

ASN allows a node to adaptively use a subset of subcarriers within its channel bandwidth. Observing that the channel bandwidth and overlap between channels in 802.11 is a multiple of 5MHz, ASN manages the spectrum in the unit of 5MHz subband, each comprising a group of 16 subcarriers. During carrier sensing, a transmitter senses the subbands within its bandwidth separately, and runs a CSMA/CA-like medium access protocol (Sec. IV) to schedule the transmission. The receiver uses a self-correlation algorithm to detect packets, and runs a cross-correlation with known sequence patterns to determine the subbands used by the transmitter and achieve synchronization. It then estimates the channel coefficients and decodes all subcarriers carrying information bits. In what follows, we detail

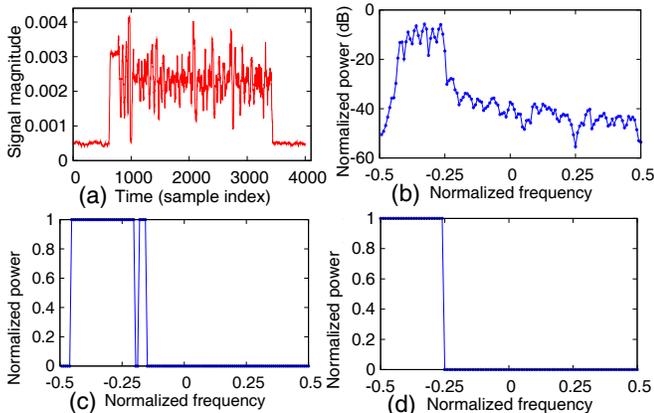


Fig. 5. Subband sensing in ASN (the transmitter’s channel has  $\frac{1}{4}$  overlap with the carrier sensing node): (a) receiving time domain samples and perform time domain energy detection (b) analyzing the PSD of samples (FFT size is 256) (c) regularizing the PSD (d) matching with an ideal overlapping pattern.

each step throughout this process. Without loss of generality, we assume the maximum bandwidth used by the transceivers is 20MHz.

### B. Sensing Subbands

An ASN-enabled transmitter needs to promptly identify the subbands currently in use. This is achieved by combining time and frequency domain energy sensing. Fig. 5 illustrates a typical procedure of subband sensing.

The time domain sensing is akin to the built-in carrier sensing primitive in 802.11g. It calculates the energy level via a moving average of the digital signals (*i.e.*, the sequence of discretized, complex samples provided by the radio’s analog-to-digital converter) within a short period, and declares a busy channel if the output exceeds the CCA (clear-channel assessment) threshold. The window size of the moving average is set to half of the length of an 802.11 preamble, to ensure a packet can be sensed promptly.

Time-domain sensing alone can sense a busy channel, but does not discriminate subbands. ASN needs to further analyze the frequency domain of the signals. Specifically, it calculates the power-spectrum density (PSD) of the recent  $N$  samples using FFT ( $N$  is called the FFT size). To ensure sufficient resolution,  $N$  needs to be larger than the number of subcarriers used by the entire channel ( $N = 64$  for a 20MHz channel).

Based on the PSD, ASN analyzes the power distribution and compares it with all possible channel-overlapping patterns. Intuitively, if the power is uniformly distributed over the entire spectrum, then the signals on the air come from a fully-overlapped channel; otherwise, only a fraction of the channel is occupied. The exact fraction of channel in use is hard to calculate, because different subcarriers may exhibit different power levels due to frequency-selective fading, and the imperfect hardware filter (used to confine the radio’s bandwidth) smears the boundary of the PSD curve. Fortunately, in 802.11g, the minimum separation between adjacent channels is 5MHz. Hence, for a 20MHz channel, for example, the overlapping pattern is one or a combination of only 4 possible overlapping cases:  $\frac{1}{4}$ ,  $\frac{1}{2}$ ,  $\frac{3}{4}$  and full overlap. Based on this observation, ASN first regularizes the PSD into a rectangular curve, compares it with all possible overlapping patterns, and then selects the one with maximum matching (Fig. 5(c) and (d)). The PSD

regularization is equivalent to thresholding points on the PSD curve with the frequency domain CCA threshold, which equals the time domain energy sensing threshold ( $-62\text{dBm}$  in 802.11g [1]) normalized by channel bandwidth.

Note that the complexity of time-domain sensing is the same as the RSSI calculation in typical communications systems, which is linear with respect to the number of incoming samples. Since frequency sensing is performed only after a sequence of signals pass the time-domain sensing, it takes constant time no matter how many samples come. The constant actually depends on the number of packets that cause the time-domain sensing to return “busy”.

### C. Packet Detection and Synchronization

In ASN, a receiver must be able to detect a packet and synchronize to it, without prior knowledge of the spectrum usage. Energy sensing alone is insufficient for packet synchronization. ASN meets this challenge by redesigning the preamble structure of 802.11g.

1) *Preamble structure in 802.11*: The original 802.11g preamble (also referred to as STF) lasts  $8\mu\text{s}$  and occupies all 64 subcarriers. From the frequency perspective, it comprises a random complex sequence spreading over every 4 subcarriers. Other subcarriers are set to 0. Owing to the duality between frequency-domain discretization and time-domain periodicity, the time domain of STF is a periodic signal that repeats every  $\frac{64}{4} = 16$  complex samples [13]. The receiver performs *self-correlation* between the latest 16 samples and previous 16 samples, which has an outstanding output only if two consecutive sequences of samples match (*i.e.*, an STF appears), and the corresponding output is comparable to the signal’s energy level [13]. After detecting the STF, the receiver further performs *cross-correlation* between the received STF samples and the original samples in the STF. An outstanding peak appears only when the received samples align with the known STF, and the peak position is used as a synchronization point marking the start of the packet.

2) *Preamble structure in ASN*: When subcarrier nulling is enabled, the random sequence in STF becomes shorter and vulnerable to noise. For example, when 48 subcarriers are nulled and the remaining 16 subcarriers are used for packet transmission, only 4 non-zero subcarriers remain in STF, which is insufficient for generating outstanding correlation output. Therefore, we modify the 802.11g preamble as follows.

First, we spread a non-zero random sequence over every 2 subcarriers in the STF, resulting in a time-domain sequence of period 32. Consequently, the cross-correlation peak results from correlation with a random sequence that has twice the length compared with the 802.11 sequence. For example, even when only a single subband (16 subcarriers) is used, 8 non-zero subcarriers are used to carry the random sequence, and therefore the STF becomes more resilient to noise.

Second, we assign different random sequences for different channel widths. For  $\frac{1}{4}$ ,  $\frac{1}{2}$ ,  $\frac{3}{4}$  and full channel width (corresponding to 1 to 4 subbands), each of them has a unique random sequence for STF. The receiver can easily identify the fraction of channel used by the transmitter by correlating the detected STF with all possible random sequences. The one that outputs peaks with the highest magnitude corresponds to the

sequence used by the transmitter, and the peaks are used as synchronization points.

Fig. 6 illustrates the packet detector’s output when a packet occupying one subband is received. The experiments runs on our prototype of ASN on the GNURadio/USRP platform (more details in Sec. V-A1). It can be seen that the self-correlation output is close to the energy level only at the preamble part; and is much smaller otherwise. Hence, it is used as a baseline for detecting the STF. In addition, when the length of the cross-correlation sequence mismatches the number of subbands used by the incoming packet, the output peaks have a much lower magnitude than those when the correct sequence is used.

#### D. Decoding Bits from Subbands

To decode a packet in 802.11g, the receiver first estimates the channel coefficients (including magnitude attenuation and phase distortion) of each subcarrier, and the frequency offset between transmitter and receiver, using an additional preamble following the STF, called *long-training field* (LTF). LTF comprises two duplicated versions of a random sequence (consisting of 1 and -1) of length 64 carried by the 64 subcarriers [1]. In ASN, when part of the channel is used, the random sequence is truncated accordingly (*i.e.*, the nulled subcarriers carry 0). To obtain the channel coefficients and frequency offset, the receiver performs self-correlation between the two truncated random sequences and normalizes it by the magnitude, similar to an 802.11 channel estimator [13]. To decode the bits, the receiver first performs IFFT over each 64 samples within an OFDM symbol, to obtain the complex samples corresponding to each subcarrier, and then normalizes the samples with the subcarrier’s channel coefficient. The normalized complex number is then mapped to the closest constellation point to obtain the digital information bits (for BPSK modulated bits, the constellation points lie at 1 and -1).

Due to temporal variation, the channel coefficients must be continuously updated when decoding the OFDM symbols. Moreover, the frequency offset estimation must be continuously refined, because even small errors in the initial LTF-based estimation may accumulate and result in decoding failure near the end of the packet. ASN updates the channel estimation using a *pilot* scheme similar to 802.11g. Specifically, among all the non-zero subcarriers in one OFDM symbol, several subcarriers (*i.e.* pilots) always send known bits. The phase drift between pilot subcarriers is used to update the frequency offset and channel coefficients [13]. When a partial-channel is used, ASN only uses 2 pilot subcarriers (due to reduced number of subcarriers available) instead of the 4-pilot scheme in 802.11 [1].

#### E. Managing Adjacent Channel Interference

Although different subbands are orthogonal, their PSD is imperfect and may leak power and cause interference to adjacent subbands used by other WLANs, referred to as adjacent channel interference (ACI). To alleviate ACI, the 802.11g OFDM PHY specified a guardband for each 20MHz channel. Among the 64 subcarriers (each is 312.5KHz), 6 are dedicated as guardband for the left border of the channel, and 5 for the right border [1], [14]. This guardband configuration is over-provisioned for most network topologies and under-utilizes spectrum [14].

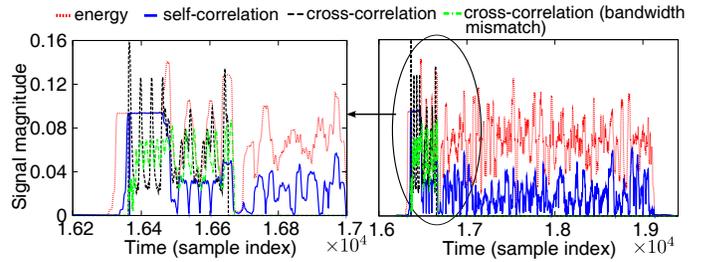


Fig. 6. Detector’s output when a packet arrives.

ASN employs fixed, but narrower guardbands. For a single subband, 1 subcarrier is used on the left boarder and 2 on the right boarder as guardband. Hence, two adjacent subbands are separated by 3 subcarriers, which is sufficient for most network topologies [14]. When all subbands are aggregated (*i.e.*, an entire channel is used), ASN restores the guardband size used by 802.11g.

In practice, harmful interference may still occur when links are closely located, even with a conservative guardband size as in 802.11. A larger guardband size may reduce such hazards, but at the cost of lower data rate. An optimal guardband configuration scheme would set the guardband according to the network topology and intensity of interference between links [14]. We leave such schemes as our future work.

### IV. ASN-AWARE MEDIUM ACCESS

When multiple WLANs partially overlap, a MAC protocol is necessary to arbitrate their contention for use of the shared subbands. We propose two MAC protocols, ASN-DA and ASN-WF, to achieve this objective.

#### A. ASN with Direct Access (ASN-DA)

The ASN-DA protocol adopts a CSMA/CA algorithm similar to the legacy 802.11, but manages sensing, backoff and transmission for each subband. When some subbands are busy, it opportunisticly nulls subcarriers in those subbands, aggregates the remaining subbands, performs backoff and sends packets through them. Fig. 7(a) illustrates a typical process of ASN-DA when two orthogonal 20MHz WLANs share the channel with a 40MHz WLAN (*i.e.*, the scenario in Fig. 3(b)).

Whenever a packet is queued, the transmitter first calls the PHY layer for time-frequency domain CCA. It freezes the backoff counter if the entire 20MHz channel is sensed busy. Otherwise, if at least one subband is idle, it generates a common backoff period for those idle subbands, using the binary exponential backoff algorithm in 802.11 [1]. Then, these subbands start decrementing the backoff counter for each idle time slot (specified to  $9\mu s$  in 802.11g).

When aggregating multiple idle subbands, ASN-DA must take into account the heterogeneity in their channel status, including the backoff-counter’s status and backoff window size. During the count-down process, a subband may be acquired by other WLANs, and its backoff counter must be frozen. Therefore, the initial idle subbands may end up with a different firing time. ASN-DA sends the queued packet through the set of subbands that first fire (*i.e.*, their backoff counters decrement to 0, and they remain idle for a DIFS period [1]). Meanwhile, other subbands will be frozen. Similar to 802.11, backoff windows of used subbands grows exponentially upon transmission failure.

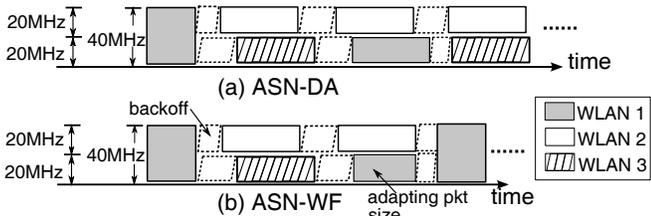


Fig. 7. ASN-Aware MAC protocols. WLAN A uses a 40MHz channel, and the other two are using 20MHz.

Therefore, not all subbands have the same backoff window size at any time. When aggregating subbands, ASN-DA generates the backoff counter based on the average of their backoff window size.

A drawback of ASN-DA is that it may lead to unfair access to shared subbands. For example, in Fig. 1, when the channel A acquires the shared subbands, it may constantly hold the subbands, while channel B can only access the remaining non-overlapping subbands. This problem may be alleviated by the post-backoff mechanism in 802.11 (*i.e.*, transmitters need to back off after successfully completing a transmission [1]), which may eventually grant the opportunity for channel B to acquire the subbands. But such opportunities are rare when both WLANs have backlogged traffic. Therefore, we design an alternative protocol, ASN-WF, to address this problem.

#### B. ASN with Water-Filling Access (ASN-WF)

The basic idea behind ASN-WF is to adapt the size of each packet, so that its duration (including the ACK) aligns with the earliest timestamp that another subband is expected to become idle (a typical procedure shown in Fig. 7(b)). From the time-domain perspective, ASN-WF attempts to “fill” the current idle subband, while maximizing the opportunity to aggregate with other subbands. To this end, ASN ensures multiple WLANs can have the opportunities to start from scratch and contend for the entire set of subbands within its channel bandwidth, thus preventing the case where a certain subband is exclusively occupied by one WLAN and achieving better fairness.

ASN-WF determines the busy duration of each subband based on the network allocation vector (NAV) provided by transmitters occupying the subband. The NAV is embedded in the header of a signaling packet preceding the actual data. It piggybacks the duration (number of time slots) that a packet plus ACK will occupy the subband. Since different subbands may be shared with different WLANs, the transmitter embeds the NAV into each subband that it uses to inform all those WLANs.

After completing each transmission, a transmitter usually has more idle subbands available to contend for. However, it still needs to start the normal CCA and backoff procedure for all idle subbands, in order to prevent unfair occupation. ASN-WF uses the same algorithm as in ASN-DA to increment/decrement the backoff window size.

## V. IMPLEMENTATION AND EVALUATION

In this section, we first validate the feasibility and performance of ASN on the GNURadio/USRP software platform. Since this platform does not yet support MAC-level functionalities due to its large response time, we use detailed simulation

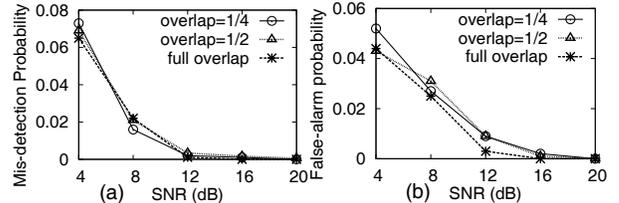


Fig. 8. Accuracy of carrier sensing for packets from partially-overlapping channels.

with ns-2 to evaluate the MAC-layer performance of ASN in multi-cell WLANs.

#### A. Performance of Subcarrier Nulling

1) *Implementation and experiment setup*: We implement ASN’s PHY-layer functionalities on GNURadio and test it on USRP. USRP is a software radio transceiver that converts digital symbols into analog waves carried by a center frequency within the ISM band. It can also receive analog signals via its RF front-end, and down-convert them into baseband signals, *i.e.*, discretized complex samples. The baseband signals are sent to a general-purpose computer running the ASN packet processing modules built atop the GNURadio library.

Our implementation is based on the 802.11g specification [1], but removes certain modules that are used to strengthen robustness to bit-errors, such as the interleaver and error-correction code. The transmitter module first maps digital information bits (0 and 1) to complex BPSK signals, and then modulates the BPSK signals into OFDM symbols. For each OFDM symbol, the pilot subcarriers are inserted according to the number of subbands to be used, following the description in Sec. III-D. The preambles (STF and LTF) are designed offline and prepended to each packet. At the receiver side, the time-frequency domain carrier sensing function and packet detector are running continuously. Once an STF preamble is detected, the receiver identifies the subbands in use and synchronizes to the packet. Then, the channel estimator and decoder follow immediately to decode all the OFDM symbols.

We run the carrier sensing, packet detection, and decoding algorithms of ASN on USRP2 radios equipped with the XCVR2450 daughterboard [15]. We set the maximum bandwidth of the USRP2 transceivers to 20MHz, and vary their center frequencies to create the overlapping patterns consistent with 802.11g channels. As a PHY layer prototype, we run a single pair of transmitter and receiver to demonstrate the feasibility of ASN.

2) *Carrier sensing of subbands*: To test the subband sensing capability of ASN, we adjust the transmit power and distance between the transceivers, thereby creating various levels of signal strength. Since the USRP2 radio does not have a direct mapping between the quantized signal magnitude and absolute power level (in dBm), we measure the relative signal strength (*i.e.*, SNR) instead. The SNR is estimated as  $SNR = \frac{E_s - E_N}{E_N}$ , where  $E_s$  is the average energy level of incoming samples when a packet is present, and  $E_N$  the noise floor, both smoothed using a moving average with the window size equal to half of the STF length. In 802.11 [1], packets must be accurately sensed by the energy detector when the signal strength is above -81dBm [1], while the noise floor (which is also the receiver sensitivity) of typical WiFi NIC is -96dBm. Thus, ASN must

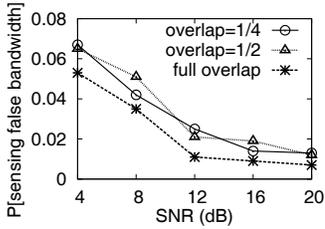


Fig. 9. Accuracy of sensing the fraction of overlapping spectrum.

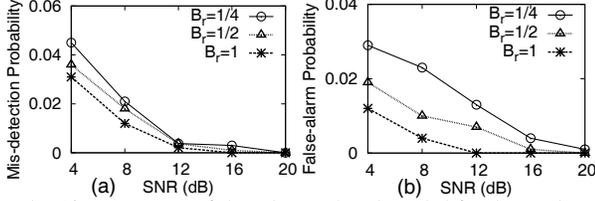


Fig. 10. Accuracy of detecting packets intended for the receiver.

be able to accurately sense a packet if its signal strength is 15dB above the noise floor.

In the experiments, the transmitter sends  $10^6$  packets with a constant inter-arrival time, bit-rate of 12Mbps and packet size 512 bytes. We use the mis-detection probability ( $P_m$ ) and false-alarm probability ( $P_f$ ) as the performance metrics.  $P_m$  is calculated by the fraction of timestamps where a packet is expected to arrive but fails to be sensed within the STF preamble duration; and vice versa for  $P_f$ .

Fig. 8 plots the resulting  $P_m$  and  $P_f$  under various levels of SNR and channel overlapping. When SNR is around 4dB, the CCA may miss packets or trigger false alarms with a relatively high probability (around 0.06). As SNR increases, both  $P_m$  and  $P_f$  decrease sharply. Above 12dB, both metrics approach 0. In addition, under the same level of SNR, the CCA performance remains almost the same for different levels of channel overlapping. It should be noted that the signal from a partially-overlapped channel is weaker than that from a fully-overlapped one. For example, for a  $\frac{1}{4}$ -overlapped channel (*i.e.*, overlapped by 1 subband), ASN must be able to detect its packets even though the SNR is  $10 \log_{10}(4) \approx 6$ dB lower than a packet from a fully-overlapped channel.

We further evaluate the accuracy of ASN's frequency domain CCA, *i.e.*, sensing the width of spectrum being used by an overlapped channel. The results in Fig. 9 show that the sensing error decreases with SNR, and approaches 1% when SNR is above 15dB. In addition, channels with a wider overlap are easier to be identified, since more matching points in the regularized PSD curve are available (Sec. III-B).

3) *Detecting packets*: To evaluate the accuracy of detecting a packet intended for the receiver, we configure the transmitter and the receiver to the same center frequency and maximum channel width. Under this setting, the transmitter may still send packets through a fraction of the channel. We denote  $B_r$  as the actual bandwidth that the transmitter uses relative to the channel bandwidth. Without loss of generality, we evaluate three cases:  $B_r = \frac{1}{4}, \frac{1}{2}$  and 1.

Fig. 10 shows the resulting  $P_m$  and  $P_f$ . We observe a similar trend as in the subband sensing experiments when SNR varies. However, both  $P_m$  and  $P_f$  are lower compared to pure energy sensing in Fig. 8, especially under low SNR. This is because the packet detector uses self-correlation and cross-correlation

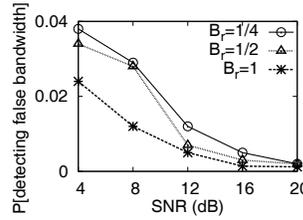


Fig. 11. Accuracy of detecting the bandwidth used by the transmitter.

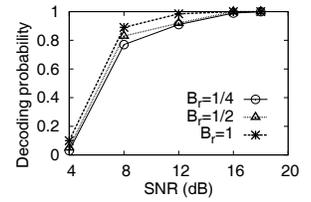


Fig. 12. Decoding probability of a packet.

to enhance resilience to noise, thus achieving higher accuracy.

We make an additional observation that a lower  $B_r$  may lead to lower detection performance, especially when SNR is low. A narrower bandwidth has fewer number of non-zero subcarriers in the STF preamble, corresponding to a shorter sequence for correlation-based detection, and are thus more susceptible to noise. Nevertheless, ASN can easily satisfy the requirement of accurate detection with above 15dB SNR, even if  $B_r = \frac{1}{4}$ .

Besides, ASN has to identify the packet's bandwidth (*i.e.*, the subbands in use). Recall the packet detector uses cross-correlation with known STF preambles to achieve this, and the accuracy is expected to be higher than a pure energy detector. This is justified in our experiment results in Fig. 11. Compared to the energy detector (Fig. 9), the detection error is typically reduced by 25% under low SNR, thus ensuring the correct channel width be fed to the decoder to recover the packet.

4) *Decoding packets*: After detecting a packet's preamble, the subsequent OFDM symbols can be decoded to recover the information bits. The decoding probability depends on the SNR level as well as on the accuracy of channel estimation. Our experimental results in Fig. 12(a) show that the decoding probability is close to 100% when SNR is above 12dB.<sup>1</sup> Below that SNR level, decoding rate drops sharply. Notably, the decoding performance of packets that partially use the channel is comparable to those on an entire channel, though slightly lower at low SNR due to fewer pilots used for channel estimation. It should be mentioned that error-correction codes may significantly boost the decoding performance, but are not used in our prototype implementation.

## B. Network Performance

The above experiments on USRP/GNURadio have shown the feasibility of ASN packet detector and decoder, and justified that it can achieve comparable performance with an 802.11 transceiver with full bandwidth under practical SNR settings. Due to the high latency of the user-space signal processing modules of USRP/GNURadio, we were unable to integrate the PHY directly with a MAC and evaluate it in a large network. Hence, we implement ASN in ns-2.34, and use the PHY-layer results to drive the network-level simulation. We modify the PHY parameters in ns-2 and ensure they are consistent with the default values in 802.11g. The original ns-2 uses a binary interference model that declares collision whenever two packets (partially) overlap. We implement an SINR based interference module that accumulates the power level of all interfering

<sup>1</sup>We measure the SNR of decoded packets via  $\text{SNR} = \frac{E_b}{N_0}$ , where  $E_b$  is the energy-per-bit, equivalently the average magnitude of decoded complex symbols.  $N_0$  is the noise energy-per-bit, equivalently the variance of the magnitude. This SNR metric accounts for the noise introduced by the decoder's channel estimator.

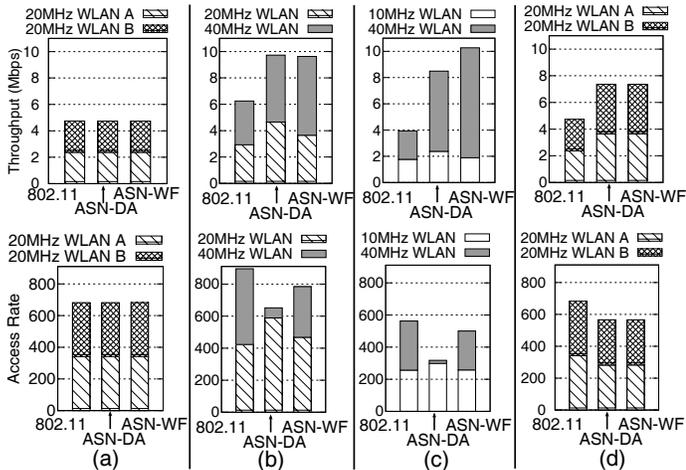


Fig. 13. Throughput and fairness when two WLANs share spectrum. (a) two 20MHz WLANs with full overlap. (b) a 20MHz WLAN overlap with a 40MHz WLAN (*i.e.*, the scenario in Fig. 1). (c) a 10MHz WLAN overlapping with a 40MHz WLAN. (d) two 20MHz WLANs overlapping by 10MHz (*i.e.*, the scenario in Fig. 3(a)).

packets, and declares a collision only if the SINR is below the decoding threshold (6dB for BPSK [11]). The collision model takes into account the possible partial overlap between packets from different channels. We have also incorporated the features of the ASN PHY: A transmitter can sense subbands separately and send packets even when certain subbands are occupied, and a receiver can detect and decode packets in each subband. The ASN-DA and ASN-WF MAC protocols are implemented on top of this modified PHY layer.

We compare three protocols: the legacy 802.11 MAC, ASN-DA and ASN-WF, using two performance metrics: throughput and *access rate*, *i.e.*, the number of transmission attempts (after CCA and backoff) per second on a shared subband. Access rate is used to study the fairness among different WLANs to access the shared spectrum. Ideally, all contenders should have the same rate of access to a shared subband, and thus, a WLAN with wider channel should achieve higher throughput.

1) *Two WLANs partially sharing spectrum*: We start with the case where two co-located WLANs are sharing part of the spectrum, each including one AP and two clients, both having saturated downlink traffic<sup>2</sup> running constant-bit-rate UDP file transfer. The packet size is fixed to 1KB. The data rate of 20MHz WLAN is fixed to 6Mbps, and that of 40MHz, 10MHz, 5MHz WLANs fixed to 12Mbps, 3Mbps, 1.5Mbps, respectively, the basic rate defined in 802.11g and 802.11n [1]. We assume the ACI is at least 6dB lower than the received signal strength for each receiver, such that ASN can be fully exploited without causing collision. This can be easily satisfied since the ACI is more than 20dB lower than the signal power even with a guardband size of 2 subcarriers [1], [14]. All our experiments run for 1000 seconds in simulation time, and the results are averaged over 10 repetitions with different random seeds. Fig. 13 shows the experimental results.

When the same channel width of 20MHz is used and the two WLANs' channels fully overlap (Fig. 13(a)), the legacy 802.11

<sup>2</sup>Although we only simulate downlink traffic, the direction of traffic does not affect the performance gain of ASN. As long as multiple links coexist and partially share spectrum, the unfairness and inefficiency of CSMA do occur and ASN becomes beneficial.

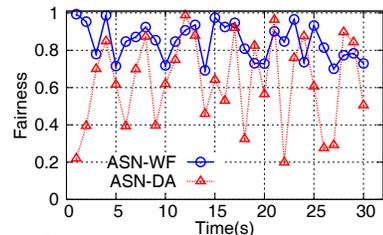


Fig. 14. Short-term fairness, with respect to access rate to the shared spectrum.

results in an equal share of throughput and access rate for both. ASN naturally downgrades to the legacy 802.11, and achieves the same level of performance.

The aggregate network throughput increases when a 20MHz WLAN shares its channel with a 40MHz channel (Fig. 13(b)), since the 40MHz WLAN takes much less time to send a packet. However, the 40MHz WLAN has almost the same throughput as the 20MHz one, although its data rate is twice as high. This is consistent with the motivating example in Fig. 1(a) — 802.11 results in an identical access rate of both WLANs to the shared 20MHz band, but at the same time, the 40MHz WLAN treats its entire channel as a single band, and accesses the non-overlapping 20MHz band at the same rate as the shared one, causing severe under-utilization of spectrum.

By contrast, with ASN-DA, both halves of the 40MHz channel can be opportunistically exploited at any time. Compared with the legacy 802.11, ASN-DA increases the throughput of the 20MHz(40MHz) WLAN by 58.7% (53.1%), and total network throughput by 55.7%. The downside of ASN-DA is that the 20MHz WLAN gains unfairly a high rate of access to the shared spectrum, leaving the 40MHz WLAN to only exploit the other 20MHz most of the time. This effect is mitigated by ASN-WF. Compared to 802.11, ASN-WF maintains similar throughput for the 20MHz WLAN, but increases throughput of the 40MHz WLAN by 80.1%. It achieves this by allowing fair access to the shared spectrum, while granting the non-overlapping spectrum exclusively to the 40MHz WLAN. In this sense, ASN-WF realizes the intuition that wider channels should gain higher throughput.

When the width of the narrowband WLAN reduces from 20MHz to 10MHz (Fig. 13(c)), the total network throughput decreases when running 802.11, although more non-overlapping spectrum is available. ASN-DA shows similar trends of throughput and access rate as in the case of 20/40MHz spectrum sharing. Compared to 802.11, it improves throughput by 34.4% and 181.7% for the 10MHz and 40MHz WLAN, respectively, and the total throughput by 115.5%. ASN-WF achieves the same throughput as 802.11 for the 10MHz WLAN, but improves that of the 40MHz WLAN by 286.5%, and the total throughput by 147.7%. In summary, the spectrum underutilization of 802.11 gets severer as the ratio of the shared spectrum to the channel bandwidth decreases, and ASN becomes more important in such cases.

Fig. 13(d) plots the experimental results for the case where two 20MHz channels overlap with each other by 10MHz (*i.e.*, the scenario in Fig. 3(a)). For the legacy 802.11, the same throughput is achieved for the case where two 20MHz channels fully overlap with each other (Fig. 13(a)). ASN-DA and ASN-WF can fully exploit non-overlapping spectrum, increasing the throughput of both WLANs by 54.4%. Both protocols lead to

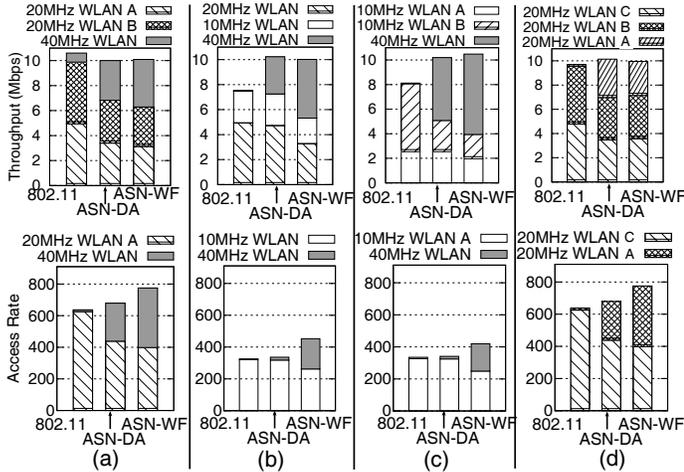


Fig. 15. Throughput and fairness when multiple WLANs of different channel width coexist.

an equal access rate to the shared spectrum.

The above evaluation focuses on the long-term access rate to shared spectrum. In the short-term, however, ASN-DA may result in dominant access to certain subbands. We investigate this effect for the case of 20/40MHz WLAN coexistence (*i.e.*, the scenario in Fig. 1). We define short-term fairness by the ratio between the minimum and maximum access rate (to the shared 20MHz spectrum) of the two WLANs averaged over a short period (*e.g.*, 1 second). Fig. 14 shows the temporal variation of short-term fairness. The fairness level of ASN-DA ranges from 0.2 to 0.98 and exhibits a significant variation. In contrast, ASN-WF maintains much more stable fairness, ranging from 0.67 to 0.99. This justifies the effectiveness of the water-filling approach in ASN-WF for arbitrating fair access to shared spectrum.

2) *Multiple WLANs sharing subbands*: When multiple WLANs of different channel widths coexist, the one partially overlapping with multiple other WLANs may be starved (Sec. II). In this section, we justify the effectiveness of ASN as a countermeasure. Without loss of generality, we first explore the case where a single 40MHz WLAN overlap with several other orthogonal narrowband channels (*i.e.*, the scenario in Fig. 3(b)). The network traffic settings are the same as above.

Fig. 15(a) shows that, when the 40MHz WLAN coexists with two orthogonal 20MHz WLANs running 802.11, its throughput approaches 0, while the 20MHz WLANs have similar throughput to the case without any contenders. This is because the 40MHz WLAN can hardly find any slots when both 20MHz contenders are idle. Accordingly, its access rate to a shared 20MHz band is close to 0, far below the equal sharing objective.

Using ASN-DA, the 40MHz WLAN can opportunistically transmit over one of the 20MHz subbands, thus achieving a similar level of throughput as the two 20MHz WLANs. However, its access rate to each 20MHz band is only around  $\frac{1}{2}$  compared to that of the 20MHz narrowband WLAN, resulting in low fairness. Again, ASN-WF alleviates this problem and enables close-to-equal access to the shared spectrum.

When the two narrowband WLANs reduce their channel width (*i.e.*, the case for 20/10/40MHz and 10/10/40MHz coexistence in Fig. 15(b) and Fig. 15(c)), the 40MHz channel remains starved when running 802.11. In contrast, by enabling access

to partially-shared spectrum, ASN improves the throughput by an order of magnitude. Owing to efficient usage of the non-overlapping spectrum, the total network throughput is also increased by around 29%.

Fig. 15(d) shows a case where three 20MHz channels partially overlap with each other (*i.e.*, the scenario in Fig. 3(c)). Consistent with our measurement (Sec. II), for legacy 802.11, the WLAN that shares spectrum with the other two orthogonal channels is starved. With ASN, all the WLANs achieve a similar level of throughput and access rate to the shared spectrum. Therefore, ASN is both necessary and effective whenever partial spectrum sharing occurs.

## VI. RELATED WORK

Besides the problem of partially-overlapping channels [10] that we discussed in Sec. II, researchers have explored other related problems and proposed their solutions.

*Fine-grained channel access*. FICA [16] reduces the MAC-layer overhead of high-rate WLANs by splitting a channel into multiple subchannels, and allowing contention for subchannels. It uses a frequency-domain backoff algorithm distinct from the traditional CSMA, and thus cannot coexist directly with current 802.11 WLANs. Moreover, FICA requires tight synchronization (with accuracy below  $0.8\mu\text{s}$ ) between all nodes that contend for spectrum. Similar approaches have been proposed to extend the OFDM-based multiple access scheme in WiMax to WiFi WLANs [17]. In contrast, ASN retains the distributed, asynchronous CSMA/CA mechanism, and can be deployed directly and coexist with current 802.11 WLANs. It targets the spectrum under-utilization problem that occurs when 802.11 WLANs partially share spectrum with each other.

*Channel width adaptation*. The proposal of variable-width channels in recent IEEE standards, such as 802.11-2007 [1], 802.11n [18], and 802.11ac [5], has generated interests in adaptively changing channel width. Chandra *et al.* explored the benefits of adapting channel width to balance the trade-off between throughput and energy-efficiency [3]. Subsequent efforts [19] proposed to assign spectrum of different widths to WLANs according to their traffic load, similarly to the notion of traffic-aware channel assignment [12]. With diverse channel widths, the partial channel sharing problem becomes inevitable, and hence, ASN can be used to further enhance such protocols.

*Narrowband-wideband coexistence*. In [20], a MAC/PHY mechanism called SWIFT is proposed to enable the coexistence between OFDM-based ultra-wideband (UWB) system and the WiFi WLANs that have a relatively narrower bandwidth. SWIFT allows UWB radios to identify the busy channels and then null them to prevent interference to WiFi. However, it identifies busy spectrum by poking the WiFi devices with a jamming tone and observing their backoff reaction. It enables UWB to achieve long-term coexistence with WiFi by evacuating the spectrum where WiFi resides on. ASN adopts OFDM subcarrier nulling similar to SWIFT, but is able to perform such adaptation on a per-packet basis, via a non-intrusive way of spectrum sensing. Using ASN, a WiFi WLAN can access spectrum with short-term fairness even if it is shared with other WLANs.

An alternative approach, Remap [21], is proposed to facilitate the coexistence between partially-overlapping 802.11 channels. Remap resolves the collision due to concurrent access to shared spectrum, by shuffling the OFDM subcarriers and harvesting diversity from repeated collisions. It can be combined with ASN as a means of collision resolution, since the CSMA/CA in ASN alone does not guarantee collision-free transmissions.

*Subband nulling for OFDM networks.* Subband nulling has also been used for different purposes. For example, MPAP [22] enables WiFi and ZigBee APs to operate on the same radio platform by nulling certain WiFi subcarriers and allocating them to ZigBee. In the context of OFDM cellular networks, there have also been proposals to null subbands that experience deep fading [23], or cause severe interference to adjacent cells [24], and reallocate the power to usable subbands. In contrast, ASN nulls subbands that are already occupied by existing WLAN cells, and uses the remaining subbands to transmit data to improve the fairness and efficiency of multi-cell WLANs.

*Non-contiguous OFDM (NC-OFDM) for cognitive radio networks (CRN).* The PHY-layer challenges of ASN resemble those in non-contiguous OFDM (NC-OFDM), a key enabling technology for CRN where available spectrum tends to be scattered over a wide range. Poston *et al.* [6] demonstrated the feasibility of NC-OFDM using a software radio based prototype, which was implemented by directly nulling the subcarriers of an OFDM communications system. Qu *et al.* [7] proposed two decision-theoretic algorithms for detecting active OFDM subbands occupied by primary users in CRN. The detection algorithms rely heavily on *a posteriori* probability of each subband's being active, which must be obtained via extensive training and is more suitable for static networks. In [8], a packet synchronization mechanism for NC-OFDM is proposed, which leverages a cyclic pattern of OFDM symbols, and is suitable for CRNs with unknown preambles. In [9], another PHY-layer challenge, *i.e.*, detecting which subcarriers are occupied, is addressed by modifying the random sequence in the 802.11g preamble. ASN's subband detection algorithm is based on a similar rationale, but becomes much simpler by leveraging the specific channel overlapping patterns in 802.11.

In summary, algorithms have been proposed to solve various PHY-layer problems in NC-OFDM communications. ASN's PHY layer can be considered as a specific NC-OFDM, but it represents a complete 802.11-based system design that includes subband sensing, detection, synchronization, and decoding. More importantly, it uses such a PHY layer to solve network-level problems, *i.e.*, partial channel blocking and wideband starvation which, to the best of our knowledge, have not been discussed elsewhere.

## VII. CONCLUSION

In this paper, we investigated the inefficiency and unfairness of 802.11 in coordinating partial spectrum sharing between WLANs, which occurs due to partially-overlapping channels or coexistence of heterogeneous channel widths. We proposed an innovative solution, ASN, that opportunistically splits the channel into subbands, nulls busy subbands, aggregates idle subbands, and transmits packets through them. We designed a set of OFDM packet processing algorithms that enable

an ASN receiver to sense, detect, and decode the packets without prior knowledge of the subbands to be used by the transmitter. We also proposed two ASN-aware MAC protocols that are 802.11-compatible, but enable efficient and fair access to partially-shared spectrum in wireless LANs. Our design was validated with implementation and experimentation on the GNURadio/USRP platform and the ns-2 simulator. As future work, we plan to extend ASN to facilitate the spectrum sharing in the whitespace networks where spectrum tends to be fragmented and partial spectrum sharing becomes unavoidable.

## REFERENCES

- [1] IEEE Standard, "802.11<sup>TM</sup>: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," 2007.
- [2] G. Hiertz, D. Denteneer, L. Stibor, Y. Zang, X. Costa, and B. Walke, "The IEEE 802.11 Universe," *Communications Magazine, IEEE*, vol. 48, no. 1, 2010.
- [3] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, and P. Bahl, "A Case for Adapting Channel Width in Wireless Networks," in *Proc. of ACM SIGCOMM*, 2008.
- [4] A. Akella, G. Judd, S. Seshan, and P. Steenkiste, "Self-Management in Chaotic Wireless Deployments," in *Proc. of ACM MobiCom*, 2005.
- [5] IEEE 802.11ac TG, [http://www.ieee802.org/11/Reports/tgac\\_update.htm](http://www.ieee802.org/11/Reports/tgac_update.htm).
- [6] J. Poston and W. Horne, "Discontiguous OFDM Considerations for Dynamic Spectrum Access in Idle TV Channels," in *Proc. of IEEE DySPAN*, 2005.
- [7] D. Qu, J. Ding, T. Jiang, and X. Sun, "Detection of Non-Contiguous OFDM Symbols for Cognitive Radio Systems without Out-of-Band Spectrum Synchronization," *IEEE Transactions on Wireless Communications*, vol. 10, no. 2, 2011.
- [8] J. Acharya, H. Viswanathan, and S. Venkatesan, "Timing Acquisition for Non Contiguous OFDM Based Dynamic Spectrum Access," in *Proc. of IEEE DySPAN*, 2008.
- [9] S. Feng, H. Zheng, H. Wang, J. Liu, and P. Zhang, "Preamble Design for Non-Contiguous Spectrum Usage in Cognitive Radio Networks," in *Proc. of IEEE WCNC*, 2009.
- [10] A. Mishra, V. Shrivastava, S. Banerjee, and W. Arbaugh, "Partially Overlapped Channels Not Considered Harmful," in *SIGMETRICS*, 2006.
- [11] M. Y. Arslan, K. Pelechrinis, I. Broustis, S. V. Krishnamurthy, S. Addepalli, and K. Papagiannaki, "Auto-Configuration of 802.11n WLANs," in *Proc. of ACM CoNext*, 2010.
- [12] E. Rozner, Y. Mehta, A. Akella, and L. Qiu, "Traffic-Aware Channel Assignment in Enterprise Wireless LANs," in *Proc. of IEEE ICNP*, 2007.
- [13] K. Jamieson, "The SoftPHY Abstraction: from Packets to Symbols in Wireless Network Design," Ph.D. Thesis, MIT, 2008.
- [14] L. Yang, B. Y. Zhao, and H. Zheng, "The Spaces Between Us: Setting and Maintaining Boundaries in Wireless Spectrum Access," in *Proc. of ACM MobiCom*, 2010.
- [15] Ettus Research LLC, "Universal Software Radio Peripheral (USRP)," <http://www.ettus.com/>.
- [16] K. Tan, J. Fang, Y. Zhang, S. Chen, L. Shi, J. Zhang, and Y. Zhang, "Fine-Grained Channel Access in Wireless LAN," in *SIGCOMM*, 2010.
- [17] J. Gross, O. Puñal, and M. Emmelmann, "Multi-user OFDMA Frame Aggregation for Future Wireless Local Area Networking," in *Proceedings of IFIP-TC 6 Networking Conference*, 2009.
- [18] IEEE Standard, "802.11n: Enhancements for Higher Throughput," 2009.
- [19] T. Moscibroda, R. Chandra, Y. Wu, S. Sengupta, P. Bahl, and Y. Yuan, "Load-Aware Spectrum Distribution in Wireless LANs," in *Proc. of IEEE ICNP*, 2008.
- [20] H. Rahul, N. Kushman, D. Katabi, C. Sodin, and F. Edalat, "Learning to Share: Narrowband-Friendly Wideband Networks," in *Proc. of ACM SIGCOMM*, 2008.
- [21] L. E. Li, K. Tan, H. Viswanathan, Y. Xu, and Y. R. Yang, "Retransmission  $\neq$  Repeat: Simple Retransmission Permutation Can Resolve Overlapping Channel Collisions," in *Proc. of ACM MobiCom*, 2010.
- [22] Y. He, J. Fang, J. Zhang, H. Shen, K. Tan, and Y. Zhang, "MPAP: Virtualization Architecture for Heterogeneous Wireless APs," *SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 1, 2010.
- [23] B. C. Jung, Y.-J. Hong, D. K. Sung, and S.-Y. Chung, "Adaptive Sub-Band Nulling for OFDM-Based Wireless Communication Systems," in *Proc. of IEEE WCNC*, 2007.
- [24] M. S. Kang and B. C. Jung, "Decentralized Intercell Interference Coordination in Uplink Cellular Networks using Adaptive Sub-Band Exclusion," in *Proc. of IEEE WCNC*, 2009.