# WiFi-NC : WiFi Over Narrow Channels

Krishna Chintalapudi*, Bozidar Radunovic†, Vlad Balan‡, Michael Buettner§,
Vishnu Navda*, Ramachandran Ramjee*, and Srinivas Yerramalli‡
* Microsoft Research India; † Microsoft Research UK; ‡ USC; §University of Washington

## ABSTRACT

The quest for higher data rates in WiFi is leading to the development of standards that make use of wide channels (e.g., 40MHz in 802.11n and 80MHz in 802.11ac). In this paper, we argue against this trend of using wider channels, and instead advocate that radios should communicate over multiple narrow channels for efficient and fair spectrum utilization. We propose WiFi-NC, a novel PHY-MAC design that allows radios to use WiFi over multiple narrow channels simultaneously. To enable WiFi-NC, we have developed the *compound radio*, a single wideband radio that exposes the abstraction of multiple narrow channel radios, each with independent transmission, reception and carrier sensing capabilities. The architecture of WiFi-NC makes it especially suitable for use in whitespaces where free spectrum may be fragmented. Thus, we also develop a frequency band selection algorithm for WiFi-NC making it suitable for use in whitespaces. WiFi-NC has been implemented on an FPGA-based software defined radio platform. Through real experiments and simulations, we demonstrate that WiFi-NC provides better efficiency and fairness in both common WiFi as well as future whitespace scenarios.

## 1. INTRODUCTION

Over the past decade, WiFi data rates have seen over a 100x increase. This was achieved through advances in physical layer wireless communication techniques (*e.g.*,OFDM, 64 QAM and MIMO) that provided increased spectral efficiency (bits/s/Hz). As further improvements in spectral efficiency become harder to achieve, using wider channels is being viewed as a solution to attain higher data rates. Today, 802.11n already allows for 40 MHz channels while the upcoming 802.11ac proposes 80 and 160 MHz channels.

*In this paper we argue against this obvious approach of merely increasing the channel width to increase wireless data rates*. Instead, we espouse the opposite – that the channels be no wider than existing 20 MHz WiFi channels and ideally be narrower, say 5 MHz or even 2 MHz. In order to achieve higher data rates then, unlike current day devices that operate over only one channel at a time, *we propose WiFi-NC a novel physical and MAC design that allows devices to run WiFi on several narrow channels, simultaneously and independently.* For example, a device must be able to use (transmit/receive on) eight 5 MHz channels instead of one 40 MHz channel. This diametrically opposite view is designed to address the following three key inefficiencies of current single (wide) channel systems (Section 3).

First, inefficiencies arise when heterogeneous radios coexist. While WiFi is designed to be fair to devices operating in the same channel, operation of 40 MHz devices nearby 20 MHz devices leads to starvation [1]. Consequently, 802.11n standard mandates devices to reduce their channel width to 20 MHz immediately upon detecting any coexisting 20 MHz device. As a result, in practical 802.11b/g/n deployments, 802.11n devices are often relegated to using only 20 MHz channels. To the best of our knowledge, no work has addressed this practical and common inefficiency in WiFi, which is bound to only get worse as 802.11ac devices with 80 MHz radios become available. In contrast, a 40/80 MHz WiFi-NC radio configured with two/four 20 MHz channels can make full use of its 40/80 MHz radio while still coexisting fairly with other 20/40 MHz networks.

Second, it is well-known that, due to MAC overheads such as backoffs, the increase in PHY data rates does not translate to commensurate gains in TCP/UDP throughput [24, 17]. To address this inefficiency, 802.11n standards support MAC-layer frame aggregation that allow frame sizes of up to 64KB, thereby reducing the relative impact of the MAC overhead. While frame aggregation works well for bulk data flows, other traffic such as TCP acks, VoIP packets and short HTTP flows are not amenable to such aggregation. The use of narrow channels in WiFi effectively also elongates packet transmission times relative to MAC overhead (for a given frame size, transmission time doubles when channel width is halved), thereby achieving higher throughput.

Third, fragmented spectrum can result in inefficient usage. For example, WhiteFi [4] uses variable width channels for operation in fragmented white spaces. However, the restriction of being able to use a single channel (wide or narrow) at a time, limits WhiteFi's ability to efficiently use free parts of the spectrum. For example, two 6 MHz narrow channels of a WiFi-NC radio can operate simultaneously on either side of a 6 MHz operating TV channel, while a single-channel system like WhiteFi will be restricted to choosing only one of

the two bands. The ability to use multiple narrow channels simultaneously, allows us to devise an optimal throughput maximizing spectrum selection scheme, $TMax$, that is not possible in single channel systems.

Related work that comes closest to WiFi-NC is FICA [24]. FICA splits a single OFDM physical channel into narrower *sub-channels* and allows different devices to access them. However, sub-channels differ from narrow channels in a very fundamental way – in FICA, a new transmission opportunity arises only when the entire wide channel is idle; then, transmissions by devices over different sub-channels are tightly time synchronized ($\approx 10\mu s$) Thus, FICA is essentially a wide single-channel system with sub-channels that are inter-dependent. While FICA addresses the MAC inefficiency issue, the lack of independence among sub-channels precludes FICA from solving the inefficiencies due to heterogeneous radio co-existence or fragmented spectrum.

The centerpiece of WiFi-NC is the *compound radio*, a novel design, that uses a *single physical wideband radio but provides the abstraction of several independent narrow band radios – radiolets*, to the MAC layer. Each radiolet allows for independent carrier sensing, transmission and reception of packets in its own narrow channel. Radiolets are entirely implemented as digital circuits and provide the low cost and form factor benefits of digital processing.

The fundamental challenge in designing a compound radio is enabling efficient *interference isolation* among the radiolets – a compound radio must be able to simultaneously carrier sense, receive, and transmit over different radio-lets without any inter-dependence. Note that, even if conventional radios supported full duplex communication [7, 12] over wide bands, we cannot simply divide the full duplex wide channel into multiple full duplex narrow channels and independently transmit/receive over these narrow channels. This is because while a full duplex radio will cancel out the spectral leakage of the narrow channel OFDM transmission at the *transmitter*, the spectral leakage can still cause severe degradation in adjacent narrow channels at the *receiver* (since the narrow channels are not synchronized). Thus, in order to achieve channel isolation, the compound radio uses *sharp elliptic filters* at both the transmitter and receiver (Section 5). These filters allows us to use very narrow guard bands between the radiolets (100KHz in our implementation), thereby paying an overhead of only 5% or 2% for a 2 MHz or 5 MHz narrow channel, respectively.

Another fundamental effect of using narrow channels is the need for *preamble dilation*. Since narrow channels transmit information at a slower rate, PHY layer preambles take longer to transmit. While the longer preamble results in only a small overhead for WiFi-NC (because data transmission times also dilate), a bigger issue is if this dilation results in increased carrier sensing time – in this case, WiFi-NC would need larger slots, which will severely affect channel utilization [17]. In order to avoid this problem, the compound radio uses energy detection to ensure carrier sensing time in WiFi-NC stays the same as in WiFi while employing cross-correlation over the dilated preamble in parallel for OFDM frame synchronization and frequency offset estimation.

We have prototyped the compound radio and WiFi-NC on a FPGA-based software defined radio platform. Through both real experiments on our testbed (Section 8) as well as extensive simulations (Section 9), we show that WiFi-NC is both more efficient and fair than WiFi. Further, while operating in whitespaces, we show that WiFi-NC is able to achieve up to 121% higher throughput than WhiteFi [5] in the presence of background transmitters.

While the use of narrow channels has significant efficiency benefits, the primary cost is increased logic/memory requirements both at the transmitter/receiver (e.g., transmit and receiver filters, decoding logic per narrow channel, etc.). However, as the FPGA/ASIC sizes grow benefiting from Moore's law, we believe that the additional logic/memory requirements will not pose a significant constraint.

In summary, our paper makes three key contributions:

- The simple insight that radios with multiple independent narrow channels instead of a single wide channel can improve the efficiency of WiFi in many practical settings such as heterogenous radio co-existence, at high PHY speeds and operation in fragmented white spaces.
- WiFi-NC, a novel PHY-MAC design that operates WiFi independently over multiple narrow channels, and its implementation in the form of a compound radio on a FPGA-based software defined radio platform.
- $TMax$ algorithm for maximizing throughput by optimal frequency selection for WiFi-NC radios operating in white spaces.

## 2. RELATED WORK

There has been tremendous amount of work targeted towards improving WiFi and wireless communication. We discuss a few papers that are most relevant to WiFi-NC here.

**Performance.** A number of papers [13, 15, 21, 23, 24] have proposed novel techniques to improve WiFi performance.

FICA [24] is closest to WiFi-NC in terms of advocating for fine-grain access. However, FICA proposes the use of *subchannels* for fine-grain access which is fundamentally different from the narrow channels of WiFi-NC. Subchannels in FICA require a *synchronous* system, where all nodes in carrier sense range must transmit within a few microseconds of each other. While it may be possible to time/frequency synchronize all APs under one management, FICA will not perform well in practical settings where WiFi APs from several autonomous systems (businesses/homes) co-exist and are not time/frequency synchronized. Furthermore, even with time/frequency sychronization, FICA does not tackle the inefficiencies that arise due to radios with different channel widths or operation in fragmented spectrum.

**Coexistence.** SWIFT [22] tackles the problem of coexistence of wide band radios in the presence of narrow band de-

vices. The SWIFT radio detects narrow band transmissions and then weaves together the unused (non-contiguous) bands into one wireless link by transmitting only on the unoccupied frequencies. While both WiFi-NC and SWIFT support non-contiguous operation, SWIFT still uses the entire available, and potentially wide, band as a single channel, thus, suffering the same inefficiencies as WiFi.

**Variable Channel Width.** A few papers [19, 21] have proposed to eliminate the notion of fixed channels in WiFi. In [19], the idea is to adapt the channel-width of each AP based on its load. In FARA [21], each node contends for the entire 802.11 spectrum. However, these systems, by contending for large channel widths in an all-or-none fashion, suffer the same inefficiencies due to MAC time overheads as WiFi. FARA also supports frequency-aware rate adaptation varying the modulation/coding independently for each subchannel. WiFi-NC can benefit from FARA's independent modulation/coding for each of its narrow channels.

In [6], the authors argue that narrow channels can increase range and reduce power consumption. However, their channel width adaptation system only configures the radio to one of 5, 10, 20 or 40MHz channel for a *single* communicating pair of radios. Coexistence/fairness will be an issue if multiple networks are configured with different channel widths. In WiFi-NC, each radio can choose to use one or more independent narrow channels, thus, gaining the benefits of narrow channels without sacrificing coexistence.

**Guard bands.** A number of techniques to mitigate the problem of adjacent channel interference was studied in [11]. The authors found that the use of guard bands was the most efficient solution to the problem. The issue of appropriate size for the design of guard bands was considered in [27]. The authors show that the size of guard band needs to be adapted based on the location of the wireless nodes. However, the software digital filters used in [26, 27] were Hamming window filters that do not have the sharp cutoff properties of the elliptic filters used in WiFi-NC (Section 5). Thus, we are able to show that even a small fixed guard band is conservative enough for our needs. Moreover, a system like Ganache [27] can also help adapt the guard band size in WiFi-NC dynamically.

**Full-duplex.** Recently, full-duplex single channel wireless communication systems have been proposed [7, 12, 20]. The key challenge in these systems is eliminating the self-interference of the local transmitter. Note that, if these systems operate over the standard 20MHz WiFi channel, they would also suffer the same MAC overhead inefficiencies as WiFi. Full-duplex communication is an orthogonal feature to WiFi-NC and can be added to the narrow channels of WiFi-NC.

**Fairness.** 802.11-based wireless networks exhibit unfairness due to a number of reasons including hidden terminals [8], capture effect [16], exponential backoffs (short term unfairness) [9], etc. WiFi-NC is focused on the problem of unfairness that arises when two or more networks operate over frequency bands that overlap (Section 3).

**Overlapping Channels.** Authors in [3] show significant unfairness in chaotic WiFi deployments where WiFi channels of adjacent access points can overlap and argue that better channel allocation and power control can help improve efficiency and fairness. Similarly, authors in [18] propose a frequency hopping algorithm called MAXchop for avoiding unfairness in uncoordinated deployments. Compared to these approaches, the narrow channel model of WiFi-NC reduces the possibility of partial overlap in channels.

**White spaces.** Closest to WiFi-NC is WhiteFi [4], a WiFi-like system for TV white spaces. WhiteFi includes a spectrum assignment algorithm that maximizes a multichannel airtime metric called MCham and an algorithm called SIFT for detecting APs of varying channel widths. While WhiteFi supports variable channel width access, WhiteFi only supports contiguous operation over the channel. As we shall see in Section 7, the contiguous access restriction results in efficiency loss due to coexistence as well as due to the conservative behavior of the MCham metric. Since WiFi-NC supports independent narrow channels, non-contiguous operation through suppression of one or more narrow channels provides significant efficiency benefits when operating in fragmented white space spectrum. Authors in [26] propose Jello, a per-session FDMA system for latency sensitive applications such as streaming media. The focus is on utilizing (non-contiguous) white space spectrum over session durations rather than on a per-packet basis as in WiFi-NC. In addition, Jello does not consider fairness among distributed nodes, a key feature of WiFi-NC.

## 3. MOTIVATION FOR WIFI-NC

Existing wideband radios are monolithic, and access the channel in an "all-or-none" fashion. This design is inefficient in multiple settings and also leads to unfair channel access. We highlight these inefficiency and unfairness issues by using three examples.

**Example 1 - Heterogeneous Radios: Inefficiency in Frequency.** Consider the concurrent operation of two WLANs (Figure 1a) – an 802.11g WLAN1 (transmitter T1, receiver R1) operating on 20 MHz channel 3 and an 802.11n WLAN2 (T2, R2) operating on 40 MHz channel 3. As dictated by the 802.11n standard, the 802.11n radio detects the 20MHz transmitter and reconfigures itself to only operate on the 20MHz channel 3 (In fact, we were unable to get our 802.11n radio to operate in 40MHz mode in any of 2.4GHz channels in our lab due to this reason). Thus, transmitters T1 and T2 alternatively use the 20MHz band while 20MHz of frequency remains completely unused.

The above 802.11n mandate was a result of the observation that operation of 40 MHz 802.11n pre-standard devices (not subject to the above mandate) alongside 20 MHz 802.11g devices led to extreme unfairness and even starvation. To understand the reason for this unfairness consider the concurrent operation of three WLANs (Figure 1b) -

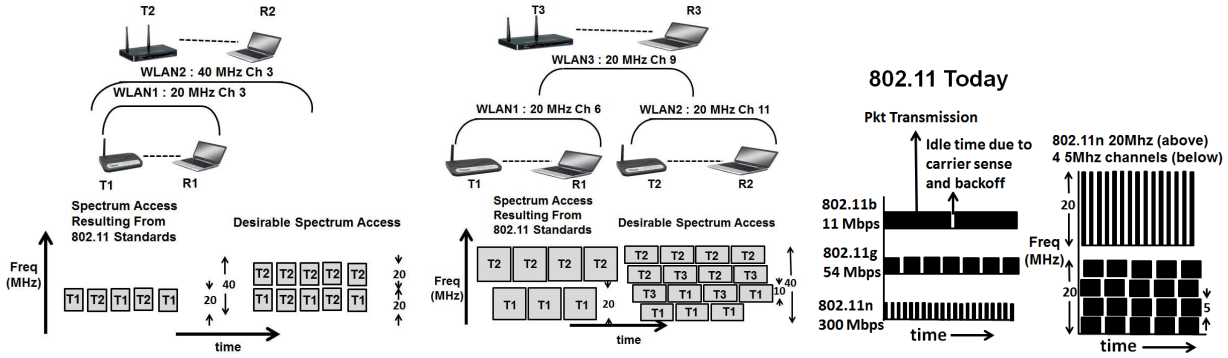


(a) Example I (b) Example I (partial overlapping) (c) Example II

1: Examples

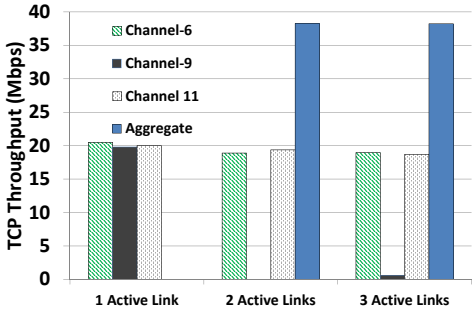


2: **Unfairness due to partial channel overlap**

WLAN1 (transmitter T1, receiver R1 ), operating on 20MHz Channel 6; WLAN2 (T2, R2) operating on 20 MHz channel 11; and WLAN3 (T3, R3) operating on 20 MHz channel 9. Since T3 is able to carrier sense both T1 and T2, 802.11 based CSMA dictates that it must wait until *both* T1 and T2 are not transmitting. However, T1 and T2 do not interfere with each other and may transmit whenever T3 is not transmitting. As depicted in Figure 1b, whenever T1 finishes transmitting a packet, T2 is still transmitting and vice-versa. Thus, T3 never finds its channel free for transmission resulting in its starvation.

To validate the above in a practical setting, we setup three identical 802.11b/g Netgear APs inside a lab area, and we had one client associated with each AP. The APs are configured to operate in channels 6, 9 and 11. The Figure 2 shows average TCP throughput at the clients for three different settings: (1) only one client is downloading at any time, (2) two clients R1 and R2 on non-overlapping channels, 6 and 11, downloading, and (3) all three clients are simultaneously downloading. From (1) and (2), it is clear that TCP flows on channel 6 and 11 are independent and do not interfere with each other when operating simultaneously. However, when all three flows are active, the client on channel 9 receives almost negligible throughput (570Kbps) compared to the other two clients (20 Mbps each) as depicted in case (3).

**Example 2 - MAC Overhead: Inefficiency in Time.** Another source of well-known inefficiency [24], illustrated in Figure 1c, arises from the fact that as the device bandwidth increases, while the time to transmit packets becomes smaller, the MAC overheads such as carrier sense and backoffs remain constant. 802.11n attempts to combat this unfairness by allowing for aggregated frames up to 64KB in size but this requires delaying frames at the interface in order to aggregate a large number of smaller packets and is not suitable for applications such as voip or short HTTP transactions. *An alternate approach that increases efficiency but does not require larger frame sizes is simply the use of narrow channels*. As seen in the Figure, reducing channel width from 20MHz to 5MHz simply results in quadrupling of the packet transmission time (the MAC overheads don't change). Thus, by elongating packet transmission time, narrow channels are able to better amortize the MAC time overheads.

Note that there is a new inefficiency introduced due to narrow channels, namely, the guard band or the gap between two 5MHz channels. We show how this overhead can be kept very small (2% for 5MHz channels) in Section 5.

**Example 3 - Fragmented Spectrum:** The recent FCC ruling of T.V white spaces allows secondary devices to transmit in parts of the spectrum unoccupied by primary transmitters such as T.V broadcasts operating over 6 MHz channels. Such an opportunistic scenario often requires devices to operate in a fragmented spectrum. Consequently, a white space device with 40 MHz radio bandwidth may not find even a single continuous span of 40 MHz. A device that allows independent channel access and communication over several narrow channels (say, eight 5 MHz channels) will allow the use of fragmented spectrum more efficiently since the white space device can simply transmit around any occupied T.V. Channels. *This example shows that the future white space devices need to support non-contiguous operation which comes naturally to a device that has multiple independent narrow channels.*

## 4. WIFI-NC

In section 3 we saw that devices can achieve fairness, in-

creased efficiency and can potentially better use fragmented spectrum if they used multiple independent narrow channels instead of a single wide channel. Given that WiFi already provides fair access to devices operating in the same channel (narrow or wide) through CSMA and backoff, WiFi-NC simply allows devices to operate WiFi independently over multiple narrow channels.

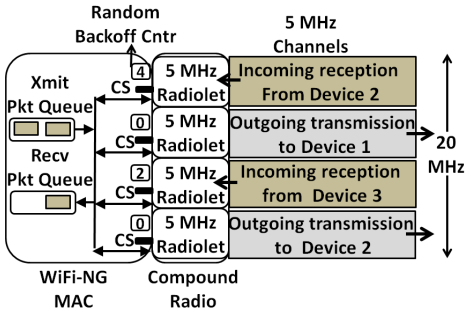


3: WiFi NC implements WiFi over several narrow channels

Figure 3 shows an illustration of WiFi-NC node configured with four 5MHz narrow channels using a 20MHz radio. The WiFi-NC MAC maintains independent random backoff counters and performs carrier sensing on each narrow channel. Whenever the backoff counter expires for a given narrow channel, a packet from the transmit queue is transmitted over the corresponding narrow channel. Similarly, packets can be received independently over narrow channels and placed in the receive queue. As can be seen from the Figure, the narrow channels are completely independent from each other. Thus, transmissions can be on-going simultaneously to different receivers (e.g., transmission to device 1 and 2), while other narrow channels can be in reception or carrier sensing mode. As we show in our evaluations (Section 9), this key property of independent narrow channels help WiFi-NC significantly outperform WiFi in terms of both efficiency and fairness in many common scenarios.

## 4.1 Exploring Design Choices

Off-the shelf radios allow operation on only one channel at time. In order to implement WiFi-NC there are several different alternatives. In this section, we consider these alternatives.

**Use multiple narrowband radios on the same device.** Several papers have advocated the use of multiple radios on a single node for better performance [2, 14]. Thus, one could consider implementing WiFi-NC using multiple narrow band radios. However, apart from several practical shortcomings such as cost, form factor, etc., there is also a fundamental drawback with such an approach – isolation requirement in the form of large guard bands between radios. For example, WiFi radios use a guardband of 3 MHz between two adjacent channels for interference free operation. This means that in order to create a compound radio of 80MHz with four 20 MHz radios, one would require guardbands worth 15 MHz (three 3 MHz in between the four radios and two on either side) - reducing spectral efficiency to 80%. This loss of spectral efficiency is further exacerbated as one uses narrower channels, say, 5MHz wide.

**Use the sub-carrier structure of OFDM itself to enable fine-grain access.** Prior work such as FICA [24] suggests that different nodes may use different sub-carriers within the same underlying physical channel to create sub-channels. However, in this approach actions such as Clear Channel Assessment (CCA), transmission and reception performed by different devices across all sub-channels must be tightly time synchronized. This is because in OFDM, sub-carriers overlap with each other and their accurate spacing and time synchronization is key to enable decoding at the receiver. Consequently, independent CCA, transmission and reception over sub-channels is not possible and leads to the same inefficiencies in co-existence between narrow and wideband devices described in Section 3.

**Our Approach - Compound Radio.** In order to enable multiple narrow channels, we propose a novel PHY-MAC design – the *compound radio*. The compound radio, *while using a single wideband physical radio device, performs digital processing to provide the abstraction of multiple independent radios to the MAC layer.* This is achieved by performing channelization digitally through digital filters and digital mixers as described in Section 5. Since digital filters allow for extremely cheap and high performance filters in comparison to analogue filters, digitally implemented adjacent channels require "very thin" guard bands (100 KHz in our implementation). Further, unlike overlapping sub-carriers in OFDM, these channels are completely separate from each other and have absolutely no cross-talk among them, allowing complete independent operation.

# 5. COMPOUND RADIO ARCHITECTURE

As discussed in Section 4, the compound radio provides an abstraction of multiple narrow-band radios while using only a single physical wideband radio. In this section, we start by describing the functioning of a conventional OFDM radio focusing only on the components that are necessary for providing the required background and then follow with our proposed architecture for the compound radio.

## 5.1 A Conventional Radio

As depicted in Figure 4 a typical radio transmitter or receiver consists of two key parts - an analogue front end and the digital baseband. Almost all the complex physical layer packet processing such as MIMO, OFDM, encoding and decoding etc. are implemented in the digital baseband since digital circuits provide the benefits of low cost, form factor and ease of implementation. However, as it is hard to design cheap digital circuits at clock rates of several GHz, the signal must first be down-converted from the carrier frequency (2.4/5 GHz) to the baseband frequency (20 MHz in case of 20 MHz channels) using the analogue frontend.

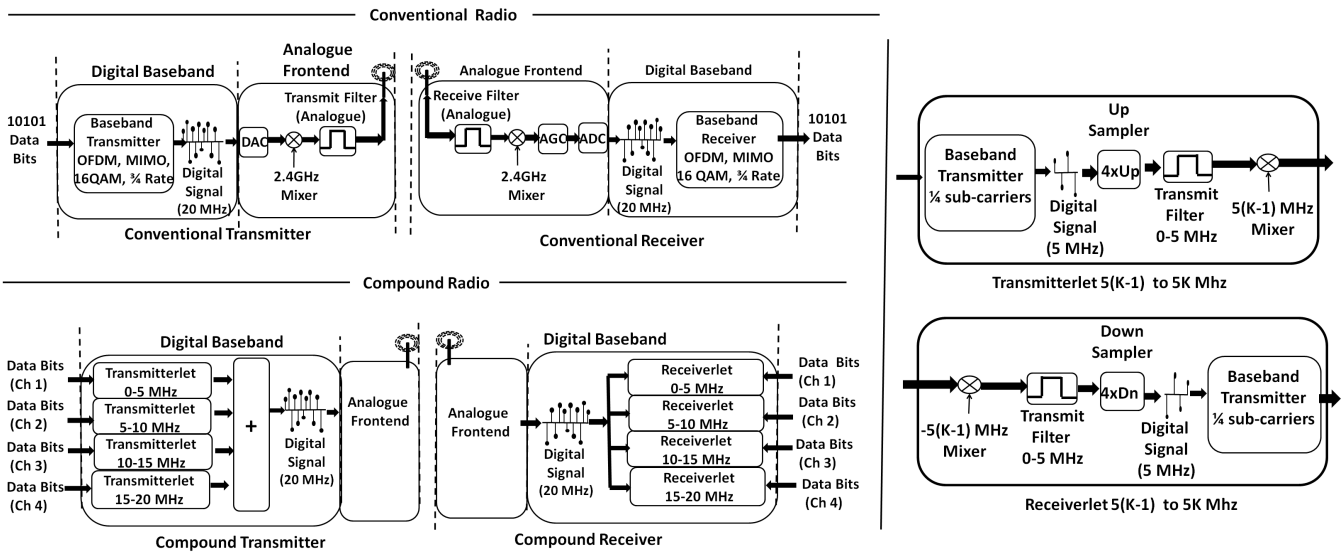


4: Conventional Radio and Compound Radio

**Analogue Transmit and Receive Filters.** In order to avoid interference from/to devices operating over adjacent channels, radios use transmit and receive filters in the analogue front end (Figure 4). These filters, only let frequencies within the bandwidth of the channel to pass through (say 2.4-2.42 GHz for 20 MHz channel 1).

**Mixer.** The mixer, at the receiver, is responsible for down-converting the received signal at carrier frequency (2.4 GHz) to baseband frequencies (0-20 MHz) to be presented to the digital baseband. At the transmitter, it up-converts the baseband signal to carrier frequency making it suitable for transmission.

**ADC and DAC.** These are used to convert between the analogue signal at baseband frequencies to digital signal at the receiver and vice-versa at the transmitter.

**AGC.** Typical DAC circuits are designed to operate correctly for a specific input voltage range (say 0.5V to -0.5V). Thus, Automatic Gain Control (AGC) appropriately scales the analogue signal from the antenna to ensure that the signal from the antenna is within the desired voltage range.

**Baseband transmitter/recevier.** Generation and reception of packet including MIMO, OFDM, encoding, decoding, modulation and demodulation are handled by the baseband transmitter and receiver using digital circuits.

## 5.2 A Compound Radio

*The key idea behind the compound radio architecture is to use digital mixers and transmit/receive filters in the baseband to create narrow channels digitally.* Figure 4 depicts this idea for a compound radio that implements four 5 MHz narrow channels.

### *5.2.1* Compound Transmitter

The compound transmitter comprises $N$ *transmitterlets*, each responsible for transmitting data over one narrow channel of width $\frac{B}{N}$, where $B$ is the bandwidth of the analogue front end. Each transmitterlet consists of a baseband transmitter, an upsampler, a digital low pass filter and a digital mixer. The outputs of each of the transmitterlets are then added digitally and passed on to the analogue frontend which is identical to the analogue frontend of a conventional radio.

**Baseband Transmitter.** The baseband transmitter is identical to the baseband transmitter used in any conventional radio, except for two differences. First, since it operates over a channel that is $\frac{1}{N}$ the bandwidth, it uses $\frac{1}{N}$ number of subcarriers intended for the wide band channel. Second, it operates at $\frac{1}{N}$ the sampling frequency of that used for the wideband radio, since the required Nyquist sampling rate for the narrow channel is $\frac{1}{N}$ of that for the wide channel with bandwidth $B$. As discussed later in this section, this allows individual transmitterlets to operate at $\frac{1}{N}$ the clock rate and hence keep the total number of operations required per second across the $N$ transmitterlets the same as the wide band radio.

**Upsampler.** In order to match the sampling rate of the wide band radio, the digital samples from the baseband transmitter are upsampled by a factor of $N$. During upsampling, $N-1$ additional digital samples are inserted between two consecutive samples through interpolation. There are several ways to interpolate – in our implementation, we use a DFT based upsampler.

**Low Pass Filter.** A sharp low pass filter (described in detail later in the section), ensures that the signal is indeed limited to within 0 to $\frac{B}{N}MHz$.

**Mixer.** Prior to the mixer, the digital signals in all transmitterlets have frequencies between 0 - $\frac{B}{N}$ MHz. The dig-

ital mixer for the $k^{th}$ transmitterlet shifts these frequencies by $\frac{(k-1)B}{N}$ MHz to ensure that the digital signal emanating from it has frequencies in the range $(\frac{(k-1)B}{N}, \frac{kB}{N})$ MHz. The mixer essentially multiplies each digital sample by a complex sinusoid of frequency $\frac{(k-1)B}{N}$ MHz and can be cheaply implemented using a ROM and two digital multipliers.

#### *5.2.2* Compound Receiver

The compound receiver architecture is symmetric to that of a compound transmitter and consists of multiple *receiverlets* - each to receive packets over one narrow channel. Each receiverlet comprises, a mixer, a low pass filter, a down sampler and finally the baseband receiver. The mixer of the $k^{th}$ recevierlet downs shifts the frequency of the received signal by $\frac{(k-1)B}{N}$ MHz. This frequency downshifting ensures that frequencies corresponding to the $k^{th}$ recevierlet *i.e.,* in the range $(\frac{(k-1)B}{N}, \frac{kB}{N})$ MHz are mapped to the range $(0, \frac{B}{N})$ MHz. A low pass filter between $(0, B)$ MHz then extracts on the band corresponding to the receiverlet. A $\frac{1}{N}$ downsampler then reduces the sampling rate by a factor of $\frac{1}{N}$ by simply dropping $N-1$ consecutive samples after picking each sample. The baseband receiver is identical to the baseband receiver of a conventional radio except that it operates at $\frac{1}{N}$ the sampling rate and uses $\frac{1}{N}$ sub-carriers of that used for the wideband channel.

## 6. DESIGN CHALLENGES

We faced two fundamental challenges in the design of the compound radio.

**Interference Isolation.** In WiFi-NC, nodes must be able to carrier sense, receive and transmit simultaneously on adjacent narrow channels. Since we use OFDM in each narrow channel for efficiency, the leakage from OFDM transmissions into the adjacent narrow channels can be significant (Section 6.1). We need to be able to isolate this interference within each narrow channel.

**Preamble Dilation.** Since narrow channels inherently transmit information at a lower rate, it takes longer to transmit physical layer preambles. Given that WiFi uses preambles to perform clear channel assessment (CCA), longer preambles might potentially necessitate larger slots and hence adversely affect efficiency.

In the rest of this section, we describe in detail each of these challenges and the approach we use to address them.

### 6.1 Interference Isolation

Figure 5 depicts a possible scenario with three WiFi-NC nodes. Node A simultaneously transmits to nodes B and C over narrow channels 1 and 3. At the same time nodes B and C transmit to node A over narrow channels 2 and 4 respectively. The spectrum of each of these transmissions as seen at Node A is also depicted in Figure 5 – node A's receiver experiences very high interference from its own transmissions in narrow channels 1 and 3 (about -20 dBm at the receive antenna, assuming a transmit power of 20 dBm [12]). Node B and C are located far away and their signals at A are extremely weak, about -85 dBm and -80 dBm, respectively.

| Filter Type | Bandwidth 5 MHz | | Bandwidth 2 MHz | |
|---|---|---|---|---|
| | # Adds | # Mults | # Adds | # Mults |
| Chebyshev | 76 | 76 | 48 | 48 |
| Butterworth | 492 | 492 | 208 | 208 |
| Elliptic | 26 | 20 | 22 | 17 |

1: Filter Comparison - 60dB attenuation, 100KHz guardband

Let as assume that we limit guard bands between narrow channels to 100 KHz so that even for a 2 MHz narrow channel, spectral wastage is only 5%. Figure 5 also depicts the typical OFDM spectral leakage in the absence of filters. The power in the adjacent channels decays to approximately about -40 dBM in the adjacent channels. In order to provide perfect interference isolation, the transmit and receive filters must attenuate the OFDM spectral leakage to below noise levels (-90 dBm or lower). Thus, we require an attenuation of the transmit signal by least 50dB to provide interference isolation. Thus, in our implementation, we use *transmit and receive filters that provide an attenuation of about 60 dB within 100kHz.* Note that this represents an extreme scenario for WiFi-NC where the self-interference is maximum compared to the received signal.

Figure 6.1 shows the number of adders and mutlipliers required to achieve our design (100 KHz guardband, 60 dB attenuation) by different choices of filters. As indicated in Figure 6.1, Elliptic Filters [10] satisfy our requirements with the least number of elements. Consequently, we use Elliptic filters in our implementation.

We have implemented the compound radio on an FPGA based software defined radio platform (Section 8). Figure 6 depicts the transmitted spectrum measured at a distance of 1cm from the transmit antenna for a 16QAM, 3/4 coding (36 Mbps) OFDM transmission over a 5 MHz narrow channel. As seen from Figure 6, the spectral leakage due to OFDM is significantly high and decays to only about -60dBm even at a distance of 2MHz from the transmitted band. The figure also shows the spectrum when using our transmit filter. We can see that the spectrum decays to about -90dBm beyond the 100 KHz guard band.

#### *6.1.1* Effect of limited bits in ADC

While one can achieve self-interference isolation only by using sharp filters, this is possible only if the ADC of the radio is able to support a wide range of power levels. An ADC typically accepts as input an analogue signal that is within $\pm 0.5$ V (or a similar range). Consequently, received signal is typically scaled (by a gain controller) down (or up) to lie within this range. The range of an ADC is specified in bits. Each extra bit of the ADC allows for discerning signals with half the amplitude and hence one fourth the power – in other words, each bit provides 6 dB resolution.

Since our testbed platform uses 14 bit ADCs, it has a range of 84 dB which means the radio is sensitive to sig-

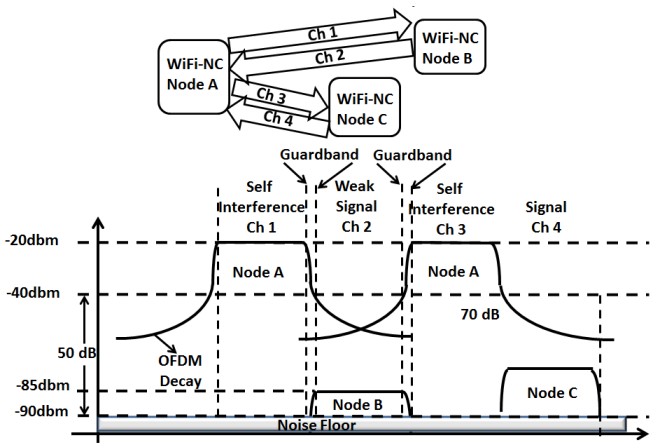


5: A possible scenario in WiFi-NC

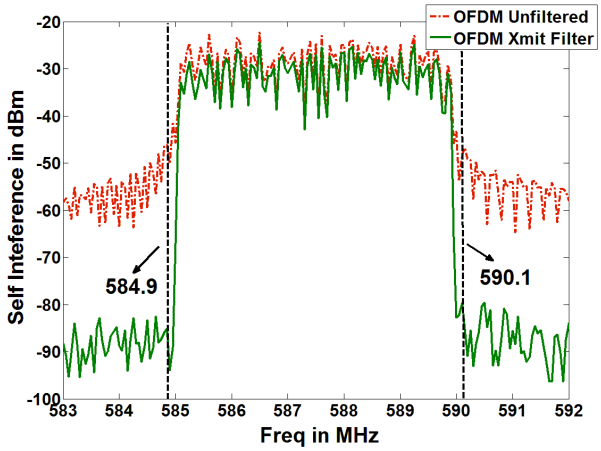


6: Spectrum of transmission with and without filters

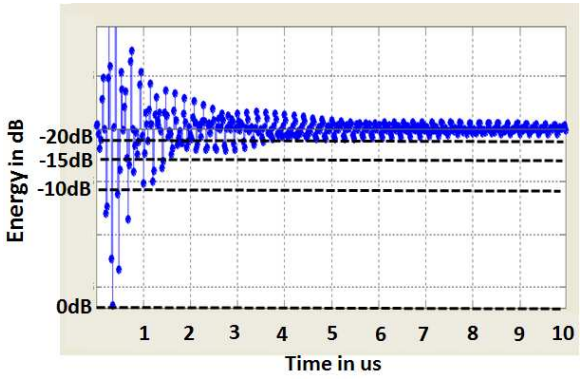


7: Filter induced multipath

nals that are 84 dB below the strongest received signal. In the face of -20dBm self interference then, a weak signal that is -85dBm effects only the last three to four bits of the ADC but is still discernible. However, many commercial systems use ADCs with 9 to 12 bits. Thus, for an ADC with 10 bits (60 dB range), this signal cannot be discerned at all since the last bit corresponds to -80dBm.

*In devices with fewer ADC bits, analogue self-interference cancelation [7] or signal-inversion using Balun transformer [12] can be used to reduce the strength of self-interference so that power levels are in the range of the ADC.* For example even a reduction of self-interference power by 25dB provided by Quellan QHx220 noise cancelers used in [7] or the 45dB over 40 MHz provided by Balun transformers [12] will permit devices with 9 bit ADCs to receive weak signals at -85 dBm while transmitting on adjacent channels.

Note that this cancellation is distinct from cancellation needed in full duplex systems [7, 20] – the cancellation here merely helps bring the power levels within the range of ADC so that transmit and receive CAN operate on *separate* channels while full duplex systems require cancellation of the full transmit signal so that one can receive on the *same* channel.

### *6.1.2* Filter Induced Interference

A filter restricts the spectrum of the signal by spreading (smoothing) it in time. The sharper the filter, the more the spreading. Figure 7 depicts the impulse response of our filter *i.e.,* the transmitted signal resulting from passing a single digital sample through the filter. As seen from the figure, the filter spreads the sample for several microseconds in time. This spreading in fact is the same effect as spreading due to indoor multipath environments. Such spreading results in self-interference between symbols termed Inter Symbol Interference (ISI).

**Need for longer Cyclic prefix (CP).** In order to combat ISI, OFDM uses the cyclic prefix, which pre-appends 25% of the OFDM symbol and extends the symbol. The spread version of the previous symbol, thus interferes with the cyclic prefix and does not adversely effect the original symbol. Typical spreading due to multipath in indoor environments is less than 800 ns, consequently, WiFi uses 800 ns cyclic prefix. However, use of sharp filters in the compound radio increases this spreading. As seen in Figure 7, the spreading decays by about 10dB within 800 ns and to about 15 dB within 1.6$\mu$s. While low data rate modulations such as BPSK require about 6 dB SNR, higher data rate modulations such as 16 QAM may require up to 14 dB SNR. In our implementation we found that a cyclic prefix about 1.6$\mu s$ long allowed for successful reception even at higher data rate modulation such as 16 QAM.

**Increasing number of Subcarriers.** Cyclic prefix is a wasteful part of the transmission and results in a decrease in spectral efficiency. In order to keep efficiency the same after extending the CP by a factor $\eta$, the symbol duration must also be stretched by the same factor $\eta$. In OFDM this is typically achieved by increasing the number of subcarriers by a factor of $\eta$. In our implementation we found that $\eta = 2$ was sufficient to combat the impluse response of the sharp 60 dB filters. Consequently, while WiFi uses 64 subcarriers in a 20 MHz band, WiFi-NC must use 128 subcarriers.

## 6.2 PHY Preamble Dilation

Physical layer preambles are crucial for packet reception and perform several key functions. Since narrow channels result in slower transmission of information, it would take longer to transmit WiFi's physical layer preambles in each narrow channel of WiFi-NC. In this section we describe the effects of this *preamble dilation* and describe our technique to address them.

### *6.2.1* Preamble Dilation in WiFi-NC

The WiFi preamble can be divided into two logical parts - the *pre-synchronization* and the *post-synchronization*. Let us look at the functions of these two parts:

**Pre-synchronization Preamble.** This part of the preamble is primarily responsible for three important functions. *C*lear channel assessment (CCA) to sense if carrier is idle, *O*FDM frame synchronization to detect OFDM symbol boundary for decoding and *f* requency offset estimation to correct for mismatches in carrier frequency between transmitter and receiver. WiFi uses pseudo-random noise (PN) sequences to perform these three functions. The performance of a PN sequence depends on the length (number of PN samples) of the sequence. A narrow channel that has $\frac{1}{N}$ the bandwidth will take $N$ times longer to transmit the same number of samples. Consequently, *this part of the preamble dilates by a factor of $N$, where $N$ is the number of radiolets.*

**Post-synchronization Preamble.** After the receiver is synchronized to the transmitter, it must estimate and compensate for the distortions caused by the wireless medium. To aid this, the transmitter sends training symbols, one (or more) for each OFDM subcarrier. The receiver then estimates the differences between the received and expected symbols and corrects for them. The key observation here is that *the number of training symbols is proportional to the number of subcarriers.* Thus, while a narrow channel with $\frac{1}{N}$ the bandwidth transmits $N$ times slower, the number of sub-carriers and hence the number of training symbols to be transmitted is also $\frac{1}{N}$ times lesser. Estimation of MIMO parameters, is based on a similar approach and is also proportional to the number of subcarriers. Consequently, since WiFi-NC uses 128 subcarriers instead of 64 used in WiFi (to counter the filter-induced interference), *this part of the preamble doubles in duration but is independent of channel width.*

**How much does the preamble dilate for WiFi-NC?**
The pre-synchronization preamble in WiFi is 4$\mu$s while the post synchronization preamble varies between 4 OFDM symbols (16$\mu$s in 802.11g) to 9 OFDM symbols (16$\mu$s in 802.11n). Thus, for a WiFi-NC transmitter with $N$ radiolets, the duration of the preamble transmission will be $4N + 32$ $\mu$s (802.11g) to $4N + 72$ $\mu$s (802.11n).

Premable dilation can potentially affect the performance of WiFi-NC in two ways. First, it can reduce the efficiency since dilated preambles take longer to transmit and second, it can mean requiring larger slot durations than 9 $\mu$s. We examine each of these next.

### *6.2.2* Effect of preamble dilation on efficiency

While the preamble transmission duration increases, so does the duration to transmit data. For example, for WiFi-NC with a 2 MHz narrow channel, the 802.11n 300 Mbps preamble will dilate from $40$ $\mu s$ to $112$ $\mu s$. At the same time, the time to transmit a 1500-byte packet elongates from $40$ $\mu s$ to $400$ $\mu s$. Thus, *the ratio of preamble transmission time to packet transmission time still reduces significantly from 100% to 28%*, resulting in significant overall gain in efficiency as the 20 MHz WiFi channel is reduced to a 2 MHz narrow channel in WiFi-NC.

### *6.2.3* Effect of preamble dilation on slot duration

The slot duration of WiFi is fixed to be 9$\mu$s, 4$\mu$s of which are allocated to perform CCA, 1$\mu$s allows for propagation delays and 5$\mu$s for switching from receive to transmit mode. Since WiFi nodes use the pre-synchronization part of the preamble to perform CCA, dilation of this part of the preamble implies that slots might also need to be dilated since CCA must be performed within one slot duration.

**Decoupling CCA from OFDM frame synchronization.**
While a sharp cross-correlation peak and hence a long PN sequence is crucial for frame synchronization, simply detecting an ongoing transmission reliably does not. A sudden and sustained increase in the energy distribution of the incoming signals can be used to reliably perform CCA. Thus, we decouple OFDM frame synchronization/frequency offset estimation from CCA. We use energy detection for CCA and in parallel perform cross-correlation over the dilated preamble. Thus, WiFi-NC nodes still perform CCA within 4$\mu$s like WiFi but the frame synchronization/frequency offset takes longer than WiFi (with no impact on slot duration). In our CCA implementation, we maintain a distribution of noise samples for the radio and infer packet transmission upon finding a large number of outliers within the 4$\mu$s window. Given that in 4 $\mu$s, about 160 consecutive samples can be processed, this approach of CCA is reliable even for weak signals as we evaluate in Section 8.

# 7. WIFI-NC IN WHITE SPACES

The key difference between operation in Whitespaces from that in the ISM band (2.4GHz) is that whitespace devices must avoid parts of spectrum occupied by primary users such as TV transmissions that use 6MHz wide channels. This leads to two key requirements for Whitespace devices. First, they must be able to operate on fragmented spectrum *i.e.*, no single continuous span of spectrum as wide as 40 MHz or even 20 MHz may be available. Second, devices need to judiciously pick which parts of the spectrum to transmit on, given that several other whitespace devices may be operating

– the *spectrum selection problem*.

The use of narrow channels allows WiFi-NC to efficiently use even narrow intermittent spaces between spectrum sections occupied by primary transmitters. Further, as we shall describe in this section, the ability to use multiple independent channels allows for a greedy distributed algorithm – $TMax$ that maximizes the total expected network throughput across all operating devices.

**Prior Approach - WhiteFi.** The problem of spectrum selection, has been examined in WhiteFi [4]. WhiteFi allows the flexibility to select among three possible analogue frontend bandwidths 5 MHz, 10 MHz and 20MHz. While, the ability to use narrower bandwidths allows WhiteFi to operate over 5 MHz or 10 MHz even when there is no span of continuous 20 MHz spectrum available, WhiteFi devices may use only one channel at time. The authors propose a metric called $MCham$ that each device maximizes greedily to determine the center frequency and bandwidth of operation. $MCham$ metric for a node $k$ with a certain center frequency $f$ and front-end bandwith $B$ is given by

$$MCham_k(f, B) = \frac{B}{5} \prod_{c \in (f,B)} \rho_k(c) \quad (1)$$

Here, $c$ corresponds to the 5MHz channels contained in the frequency span $(f - \frac{B}{2}, f + \frac{B}{2})$, and $\rho_k(c)$ corresponds to the expected share of node $k$ in a 5MHz channel $c$, given by,

$$\rho_k(c) = \max\left(R_k(c), \frac{1}{L_k^c}\right). \quad (2)$$

In equation 2, $R_k(c)$ refers to the fraction of residual airtime available in the channel $c$ and $L_k^c$ refers to the total number of contenders in the channel.

WhiteFi was faced with a key constraint, i.e., its radio only supported the notion of a *single channel* that operated in a contiguous manner over the full bandwidth. This creates *two key disadvantages:* 1) the need to choose an operating bandwidth (e.g., 5MHz) that may be lower than the full bandwidth of the radio (e.g., 20MHz); 2) the $MCham$ metric has to be *conservative* since a wideband radio cannot use the channel until all overlapping subchannels are free at the same time. Thus, the *product term* in Equation 1. Note that this coupling could also result in starvation, similar to the problem described in Section 3.

$TMax$ **Algorithm in WiFi-NC.** In the case of WiFi-NC, since the radio supports *independent narrow channels*, both disadvantages of WhiteFi disappear. The radio can always use its full available bandwidth since it can operate in a non-contiguous manner around any primary transmitters. Also, since the narrow channels are independent, the throughput available estimate need not be conservative and is simply the *summation* of throughput in each of its narrow channels. Thus, WiFi-NC uses a new metric called Throughput Maximal metric or $TMax$ for determining its frequency of operation, as

$$TMax_k(f) = \sum_{c \in (f)} \frac{B}{n} \rho_k(c) \quad (3)$$

where $n$ is number of narrow channels, $B$ is the analogue frontend's bandwidth. and $c$ is the set of all narrow channels in the range $(f - \frac{B}{2}, f + \frac{B}{2})$.

WiFi-NC nodes operating in whitespaces, periodically scan over the entire available parts of the spectrum computing the $TMax$ metric for part. They then greedily choose $f$ where the part of spectrum that maximizes $TMax$. When two or more regions of the spectrum have the same value for $TMax$, ties are broken by always choosing the lower frequency value for operation.

**Optimality of** $TMax$**.** It can be shown that, while each node greedily uses the $TMax$ algorithm, the scheme iteratively converges to maximizes the expected aggregate network throughput across all operating devices and hence the overall spectrum utilization. In the interest of space, we provide only a brief sketch of the *optimality proof.* Since nodes are not synchronized, without lack of generality, we assume only one node performs the frequency selection operation at any given time. Since the aggregate network throughput is simply a sum over all $TMax$ metrics for every node in the network, as at each step every node either greedily reduces it own metric or keeps it the same, the overall metric never increases and keeps converging towards a minima. Further, breaking ties in a manner so that nodes always pick the lowest available frequency region ensures that the parts of minimum spectrum are always collected in a single contiguous chunk. It can be then shown that when the choice to be picked always comes from a contiguous chunk of spectrum, the algorithm is also globally optimal and the algorithm converges in the face of dynamically leaving and entering devices.

## 8. RESULTS ON TEST BED

WiFi-NC has been implemented on a DSP/FPGA based software defined radio platform – the SFF SDR from Lyretech Inc. SFF SDR uses two Virtex-4 SX35 FPGAs and the DM6446 DSP processor from TI. The entire digital baseband of the compound radio and time-sensitive parts of MAC such as backoff counters and CSMA have been implemented on the FPGA. We used an off-the-shelf sub-gigahertz analogue radio front end provided by Lyretech that allows transmissions between 360 MHz to 960 MHz. The analogue radio front end supports two antennas – one for transmitting and one for receiving – each of which can be operated independently. While, the board itself supports two different bandwidths namely 10 MHz and 20 MHz, throughout our experiments we have used the 10 MHz option. Using our implementation of the compound radio we demonstrate that WiFi-NC allows devices to share the spectrum in a fair and efficient manner.

### 8.1 Self-Interference Isolation

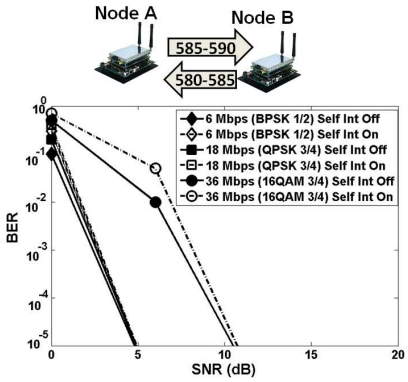


8: Self-Interference Isolation

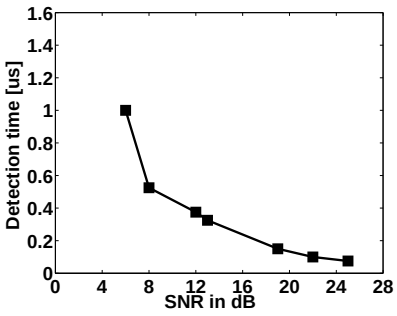


9: Performance of CCA

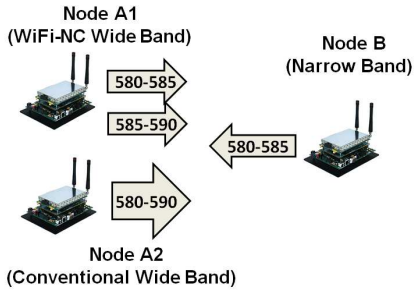


10: Experimental Setup Narrow Band Wide Band Fairness

In this experiment we ask the question "how well can a WiFi-NC device receive while transmitting simultaneously on an adjacent narrow channel?" Specifically, *we demonstrate that there is no difference in BER for WiFi-NC, whether or not it transmits over an adjacent channel over a wide range of SNRs and data rates* indicating perfect self-interference isolation. To answer this question we conducted an experiment with two WiFi-NC devices A and B as depicted in Figure 8. Node A transmits to Node B over the 5 MHz channel 585-590 MHz, while simultaneously Node B transmits to Node A over the channel 580-585 MHz. We measured the Bit Error Rates (BER) at Node A for various average SNR values generated by placing nodes at various distances. We then compared these values when only Node B transmits to Node A *i.e.,* in the absence of self-interference. We found that for data rates *up to 18 Mbps (QPSK with 3/4 coding rate) we were not able to see any bit error over* $10^6$ *bits with or without self-interference even at SNRs as low as 5dB and at narrow channel widths of 2 and 5 MHz.*

16 QAM and higher data rate modulations however, are more sensitive to SNR. Consequently, in order to investigate at these higher data rates we tried 36 Mbps (16 QAM, 3/4 coding rate). As seen from Figure 8, for 36 Mbps (16QAM 3/4 coding rate) require about 14dB for the same. *This performance is almost identical when there is not self-interference indicating that the channels are isolated from self-interference leakage from the adjacent channel.*

## 8.2 Efficacy of CCA Detection

As discussed in Section 6.2.3, using narrow channels leads to dilation of the pre-synchronization part preamble in time. Consequently, in order to ensure CCA in 4 $\mu$s, we use an energy detection based scheme described in Section 6.2.3 to perform CCA. In this section we evaluate the efficacy of our CCA detection scheme. There are two key measure for a successful CCA scheme – false detection rate and missed detection rate. When the CCA detector mistakes noise for an incoming packet we deem this a false detection. False detections are extremely expensive in receiver since this leads them falsely into the receive state in preparation of packet reception and wastes time. Missed detection on the other hand occurs when a genuine transmission is not detected, this leads to collisions. WiFi allows for a missed detection rate of 10% in 4$\mu$s. *In our design we ensured that the false detection rate was under 1% while missed detection rate was under 5%*. We transmitted the preamble from one device and at the receiver measured how much time would the receiver require before detecting the transmission correctly. Figure 9 depicts the performance of our CCA scheme. *As seen from Figure 9, even signals with SNR as low as 5 dB are detected in about 1$\mu$s, while at high SNR values CCA can be performed in a few hundred nano-seconds.* This is expected since, the higher the SNR, the more easily signal can be distinguished from noise.

## 8.3 Narrow and Wide Band Device Coexistence

In this experiment we demonstrate that WiFi-NC allows narrow and wide band devices to coexist in a fair manner. The experimental setup is depicted in Figure 10. As shown in Figure 10, Nodes A1 and A2 are wideband devices operating over 10MHz while Node B is a narrow band device operating over a 5 MHz channel. Node A1 is a WiFi-NC device and uses two narrow 5 MHz channels while Node A2 is a conventional device and uses a single 10 MHz wide channel.

**Individual Links.** First, we measured the throughputs achieved by each device on each narrow channel individually without any other transmissions for measuring the base achieved throughput. All devices are operated at 54Mbps. As seen from Figure 11, the 10MHz wideband device achieves a throughput of about 16Mbps, while the achieved throughout over each 5 MHz narrow channel is about 10 Mbps.

**Conventional 10 MHz and Narrow Band.** Next, (A2 and B in Figure 11) device A2 and B are turned on to start transmitting. Since A2 uses a wide channel it shares the channel with B and vice-versa. Consequently, while A2 achieves roughly 9 Mbps of throughput, B achieves about 6 Mbps throughput.

**WiFi-NC 10 MHz and Narrow Band 5 MHz.** Finally, A1

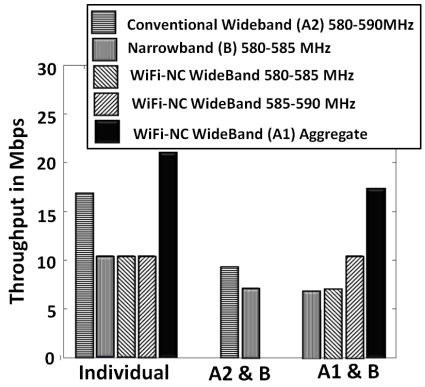


11: Narrow band wide band coexistence

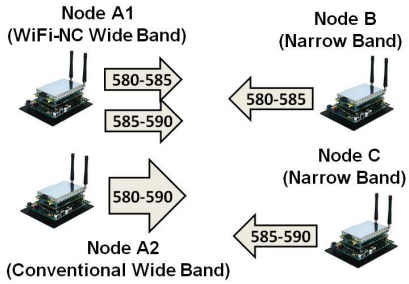


12: Experimental Setup for Starvation

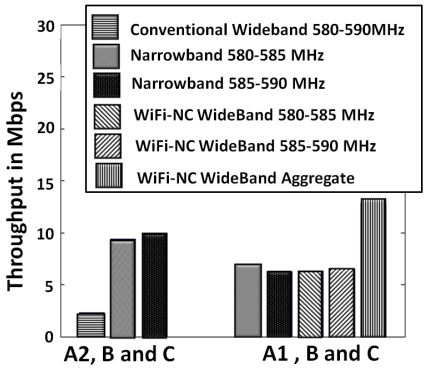


13: WiFi-NC solves Starvation

is turned on and operated alongside B. As seen from Figure 11 (A1 and B), A1 only shares the 5 MHz channel between 580-585 MHz with B and is able to completely use the channel 585-590 MHz without any contention. Consequently, A1 is able to achieve an aggregate throughput of about 15 Mbps while B achieves a throughput of about 6 MHz. *This demonstrates how WiFi-NC can help narrow and wideband devices gain fair access and thus keep the overall utilization high.*

## 8.4 WiFi-NC Avoids Possible Starvation

As discussed in Section 3, operation of devices with wide channels alongside those with narrow channels can result in starvation. In this experiment we demonstrate that WiFi-NC devices can avoid this stavation by using narrow channels. As depicted in Figure 12, there are four nodes used in this experiment. A 10 MHz wideband device, A2 operating over 580-590 MHz, a 10 MHz WiFi-NC device A1 using two narrow channels 5 MHz each and two narrowband devices B and C operating over non-overlapping 5 MHz bands 580-585 MHz and 585-590 MHz.

**Conventional 10 MHz and two 5 MHz devices.** We first, turn on devices A2, B and C which transmit packets while contending for channel access. As seen from Figure 13, node A2 achieves only about 2 Mbps out of a possible 16 Mbps (Figure 11) while devices B and C achieve most of the share in their respective bands. *This demonstrates how wideband devices can potentially suffer from extreme unfairness while operating alongside non-overlapping devices.*

**WiFi-NC 10 MHz and two 5 MHz devices.** Next, we turn off A2 and turn on A1 allowing the WiFi-NC wideband devices to transmit while contending for channel access. *As seen from Figure 13 the wiFi-NC device is able to share the narrow band channel fairly with each of the narrow band devices B and C and consequently avoid extreme unfairness.*

## 8.5 Efficincy in WiFi-NC

In order to measure efficiency gains in WiFi-NC with the use of narrower channels we implemented radiolets with 10, 5 and 2.5 MHz narrow channels on our platform. Since we

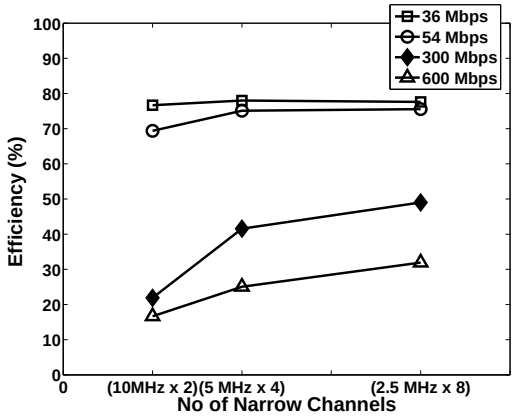


14: Increase in Efficiency with number of narrow channels

did not have a MIMO implementation, in order to create the effects of various data rates we used shorter packets with data in the packet scaled by the data rate. For example, to create the effect of 300 Mbps we used 36 Mbps with packets of 1500(36/300) bytes. Figure 14 shows the variation of efficiency with narrower channels. As depicted in Figure 14, efficiency increases with increase in increasing number of narrow channels and as expected, the increase is greater at higher data rates such as 300 and 600 Mbps.

# 9. SIMULATION STUDY

The testbed evaluation is restricted to small scale experiments involving a few devices and limited set of scenarios. Several questions regarding the performance of WiFi-NC require exploration. How do the choices of different channel width effect the efficiency of WiFi-NC? How does WiFi-NC perform with latency-sensitive media traffic such as VOIP compared to WiFi? How does WiFi-NC perform as a potential choice for white space usage? In order to answer these questions, we have implemented the compound radio PHY layer and WiFi-NC MAC layer as extensions to the Qualnet network simulator.

In our simulations, all nodes are within carrier sense range.

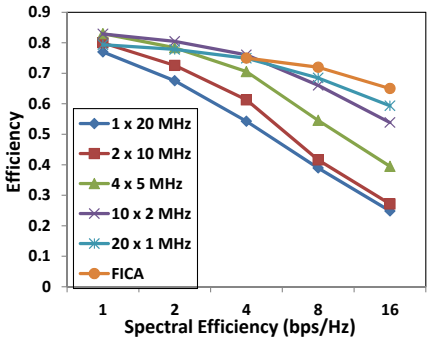


15: Efficiency for different channel sizes and spectral efficiencies

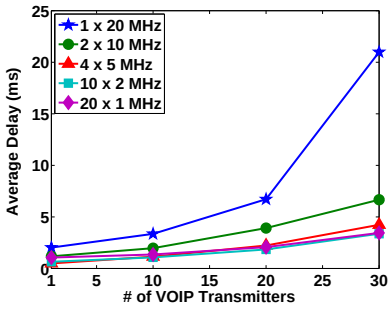


16: Average Latency for different channel sizes

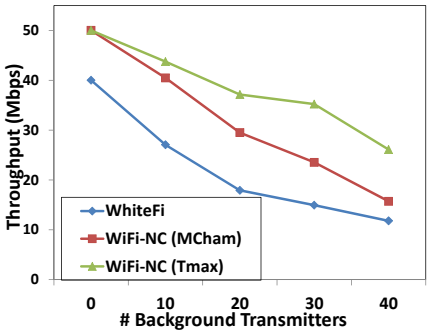


17: White space transmitter throughput

The narrow channel width and spectral transmission efficiency in bps/Hz (modulation rate etc.) is configured per node, along with the guard band size. Unless otherwise noted, we use a spectral efficiency value of 2.7 bps/Hz and a radio front-end bandwidth of 20MHz, which is equivalent to the 54 Mbps mode of 802.11a/g. Protocol overhead such as header size and ACK length is modeled on 802.11a. For all Wifi-NC configurations, we fixed the guard band size to 100 KHz, same as our prototype implementation.

## 9.1 Efficiency

To understand the trade-off of using smaller narrow channels, we experimented on a single 20 MHz wide-band link with different WiFi-NC configurations (number of channels x channel width). We measure the achieved bit rate when the link is saturated with 1500 byte back-to-back packets for different values of spectral efficiency (bps/Hz). Figure 15 shows the channel efficiency, which is computed as the ratio of the achieved bit-rate to the raw bit-rate, for different WiFi-NC configurations. For comparison, we also plot the channel efficiency numbers quoted by FICA [24] for similar data rates.

As spectral efficiency increases, the fixed protocol overheads become increasingly burdensome. As spectral efficiency increases, narrower channels provide much better channel efficiency than using a single wide channel. With a spectral efficiency of 16 bps/Hz, equivalent to a 320 Mbps bit rate across a 20 MHz band (similar to 300Mbps 802.11n), the 20 x 1 MHz configuration is 60% efficient compared to only 25% when using a single 20 MHz band. In comparison, FICA achieves an efficiency of around 65%. Thus, WiFi-NC is able to match the high efficiency of a synchronous system like FICA while still operating in a fully asynchronous manner.

## 9.2 Latency

Narrower channels increase throughput by elongating packet transmission times; this amortizes the cost of fixed overheads. However, longer transmission times increase latency which could be problematic for latency sensitive traffic such as VOIP. Surprisingly, we found that using *narrower channels actually reduce system latency* when there are multiple clients. In this experiment, along with the single bulk transmitter saturating the link, we add an increasing number of clients that each transmit 200B packets every 20 ms representing VOIP payload.

Figure 16 shows that, as the number of clients increase, narrow channels have lower latency compared to a single 20MHz channel. Additional clients means increased contention and likelihood of collisions. Since latency sensitive clients are not saturating the channel, using narrower channels reduces the number of contenders on any given band and thus reduces latency. In particular, using more channels reduces the incidence of packet collisions due to choosing the same slot for transmissions by up to 10% (not shown due to lack of space).

## 9.3 White Space Networking

WiFi-NC is also well suited for use in white space networks where fragmented spectrum is the norm. To demonstrate this, we simulated a white space network based on TV broadcasters in an urban area according to TV Fool [25]. This gave us 31 6 MHz TV channels with 12 incumbent transmitters. In order to study the impact of background traffic, among the open channels we randomly distribute narrowband background transmitters, each of which used a UDP stream to consume 1/3 the capacity of a single 6 MHz band, similar to the evaluation used in WhiteFi [4]. We then add a wide-band transmitter that can use up to 4 channels, and we measured its throughput for three different schemes, WhiteFi [4], WiFi-NC with the $MCham$ metric and WiFi-NC with the $TMax$ metric.

Figure 17 shows the throughput as the number of background transmitters increases. With 0 background transmitters, WhiteFi can use 4 channels concurrently and achieves a 40 Mbps throughput. However, as more background transmitters are added, $MCham$ is forced to select bands that use less than four channels to avoid background transmitters. This means that when there are 40 background transmitters, WhiteFi selects only a single channel, and achieves only 12 Mbps of throughput. In the case of WiFi-NC with $MCham$

metric, throughput is higher due to WiFi-NC's higher efficiency and also its ability to contend with each narrow band transmitter independently. With 20 background transmitters, this scheme increases throughput by more than 65% compared to WhiteFi. However, with 40 transmitters its gains reduce since the $MCham$ metric selects only one channel.

When WiFi-NC is used with the $TMax$ metric, $TMax$ always uses four adjacent channels and contends independently on each one (operating in a non-contiguous manner around incumbents). Thus, it is able to achieve much higher throughput than the other schemes, delivering a throughput gain of 121% over WhiteFi when there are 40 background transmitters.

## 10. CONCLUSION

In order to support gigabit wireless speeds, 802.11 standards are increasingly driven towards wide bandwidth, wide channel design. In this paper, we argue for support of multiple independent narrow channels inside the single wideband radio and propose WiFi-NG and compound radio, a novel design and implementation to achieve it. Through experiments and simulations, we show that WiFi-NG maintains high efficiency at high data rates, is able to fully and fairly utilize the wideband in the presence of coexisting networks, and works well in future white space scenarios where spectrum may be fragmented.

Given the prevalence of over billion WiFi devices, deployability of new approaches such as WiFi-NG can be a concern. However, since the same MAC is used in both WiFi and WiFi-NG, supporting *backward compatability* is easy. Upon detecting a 20 MHz transmission in its vicinity, the WiFi-NG radio simply reconfigures itself to use 20MHz wide channels. For a 40 MHz 802.11n radio, using two 20 MHz narrow channels can provide throughput and efficiency gains. As future radios sport 80 MHz or wider bands, WiFi-NG can provide significant value even when operating in backward compatability mode.

## 11. REFERENCES


[1] 802.11n Bonding Unfairness. http://www.zdnet.com/blog/ou/80211n-draft-110-a-kinder-gentler-neighbor/410.

[2] A. Adya, P. Bahl, J. Padhye, A. Wolman, and L. Zhou. A multi-radio unification protocol for IEEE 802.11 wireless networks. In *BroadNets*, 2003.

[3] A. Akella, G. Judd, S. Seshan, and P. Steenkiste. Self management in chaotic wireless deployments. In *ACM MobiCom*, 2005.

[4] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh. White space networking with wi-fi like connectivity. *ACM SIGCOMM*, 2009.

[5] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh. White space networking with wi-fi like connectivity. In *ACM SIGCOMM*, 2009.

[6] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, and P. Bahl. A case for adapting channel width in wireless networks. In *ACM SIGCOMM*, 2008.

[7] J. I. Choi, M. Jain, K. Srinivasan, P. Levis, and S. Katti. Achieving single channel, full duplex wireless communication. In *ACM MobiCom*, 2010.

[8] S. Gollakota and D. Katabi. Zigzag decoding: Combating hidden terminals in wireless networks. In *ACM SIGCOMM*, 2008.

[9] M. Heusse, F. Rousseau, R. Guillier, and A. Duda. Idle sense: An optimal access method for high throughput and fairness in rate diverse wireless lans. In *ACM SIGCOMM*, 2005.

[10] M. C. Horton and R. J. Wenzel. The Digital Elliptic Filter – A Compact Sharp-Cutoff Design for Wide Bandstop or Bandpass Requirements. *IEEE transactions on Microwave theory and techniques*, pages 307–314, May 1967.

[11] W. Hou, L. Yang, L. Zhang, X. Shan, and H. Zheng. Understanding cross-band interference in unsynchronized spectrum access. In *ACM CoRoNet*, 2009.

[12] M. Jain et. al. Practical, real-time, full duplex wireless. In *ACM MobiCom*, 2011.

[13] K. Jamieson and H. Balakrishnan. Ppr: Partial packet recovery for wireless networks. In *ACM SIGCOMM*, 2007.

[14] S. Kakumanu and R. Sivakumar. Glia: A practical solution for effective high datarate wifi-arrays. In *ACM MobiCom*, 2009.

[15] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Medard, and J. Crowcroft. Xors in the air: Practical wireless network coding. In *ACM SIGCOMM*, 2006.

[16] A. Kochut, A. Vasan, A. Shankar, and A. Agrawala. Sniffing out the correct physical layer capture model in 802.11b. In *ICNP*, 2004.

[17] E. Magistretti et. al. Wifi-nano: Reclaiming wifi efficiency through 800ns slots. In *ACM MobiCom*, 2011.

[18] A. Mishra, D. Agrawal, V. Shrivastava, S. Banerjee, and S. Ganguly. Distributed channel management in uncoordinated wireless environments. In *ACM MobiCom*, 2006.

[19] T. Moscibroda, R. Chandra, Y. Wu, S. Sengupta, P. Bahl, and Y. Yuan. Load-aware spectrum distribution in wireless lans. In *ICNP*, 2008.

[20] B. Radunovic, D. Gunawardena, P. Key, A. Proutiere, N. Singh, V. Balan, and G. Dejean. Rethinking indoor wireless mesh design: Low power, low frequency, full-duplex. In *WiMesh*, 2010.

[21] H. Rahul, F. Edalat, D. Katabi, and C. Sodini. Frequency-Aware Rate Adaptation and MAC Protocols. In *ACM MobiCom*, Beijing, China, September 2009.

[22] H. Rahul, N. Kushman, D. Katabi, C. Sodini, and F. Edalat. Learning to share: Narrowband-friendly wideband wireless networks. In *ACM SIGCOMM*, 2008.

[23] E. Rozner, Y. Mehta, A. Akella, and L. Qiu. Traffic-aware channel assignment in enterprise wireless lans. In *ICNP*, 2007.

[24] K. Tan, J. Fang, Y. Zhang, S. Chen, L. Shi, J. Zhang, and Y. Zhang. Fine Grained Channel Access in Wireless LAN. In *ACM SIGCOMM*, August 2010.

[25] TV FOOL. http://www.tvfool.com.

[26] L. Yang, W. Hou, L. Cao, B. Y. Zhao, and H. Zheng. Supporting demanding wireless applications with frequency-agile radios. In *USENIX NSDI*, 2010.

[27] L. Yang, B. Y. Zhao, and H. Zheng. The spaces between us: Setting and maintaining boundaries in wireless spectrum access. In *ACM MobiCom*, 2010.