

WiFi-XL: Extending WiFi to Wide Areas in White Spaces

Apurv Bhartia¹, Mahanth Gowda², Krishna Chintalapudi³, Bozidar Radunovic³, Ramachandran Ramjee³, Deeparnab Chakrabarty³, Lili Qiu¹, and Romit Roy Chowdhury²

¹Univ of Texas, Austin ²Univ of Illinois, Urbana-Champaign ³Microsoft Research

October 2014

Technical Report
MSR-TR-2014-132

Microsoft Research
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052

<http://www.research.microsoft.com>

1 Introduction

While WiFi has been an immense success as a local-area wireless technology, attempts to use it for wide-area wireless access have been failures [3, 5].¹ A key technical reason for this failure is the limited coverage of a WiFi access point (AP) [12].

The recent FCC white space ruling [10, 11] allows unlicensed transmissions over white spaces, i.e., TV spectrum that are allocated but remain unused. Better propagation of radio waves in white space frequencies coupled with FCC permitted maximum AP power of 4W ($40\times$ higher power than WiFi) suggests that simply adapting WiFi radios to white spaces as in IEEE 802.11af can liberate WiFi from its range limitations, enabling license-free wireless wide-area networks (WWAN).

In this paper we show that, due to two characteristics unique to WWANs, *using standard WiFi in white spaces will neither provide adequate capacity nor wide coverage*. Thus, we design, implement and evaluate WiFi-XL, a white space WWAN that extends WiFi-like spectrum sharing to wide area wireless networks.

Channel Trichotomy. We start with a detailed measurement study characterizing white space propagation in WWANs (Section 2). A unique characteristic of WWANs is that APs are deployed on towers at heights between 10-30m, while clients are mostly at ground level. As a consequence, we show that a typical WWAN deployment comprises of three different wireless propagation channels, namely, AP-AP, AP-Client and Client-Client, which we term channel trichotomy. In particular, the AP-AP channel in WWANs has far superior propagation characteristics compared to other channels.

While transmissions only occur between AP and its clients, channel trichotomy has important interference implications. For example, our measurements show that a WWAN AP has an *interference range 2.5 to 8 times its transmission range*. Thus, channel trichotomy amplifies white space signal propagation, creating a double-edged sword: the enhanced coverage of a WWAN white space AP comes at the cost of greatly expanded set of interfering APs, severely impacting overall network capacity.

Separating neighboring WWAN APs in frequency is

¹By wide-area coverage, we mean coverage of large areas similar to cellular base stations as opposed to long-distance WiFi *e.g.*, [20] where directional antennas can help WiFi traverse long distances.

therefore crucial. While automatic channel selection is a well-studied topic in WiFi [8], prior work relies on measurements for selecting an AP's channel. In WiFi, where the interference range is twice the transmission range, AP's use client-based scans to identify all interfering APs [18]. However, since interference range is greater than twice the transmission range in WWANs, even client-based scans can only identify a subset of interfering WWAN APs. Thus, measurement-based approaches are fundamentally unsuitable for WWANs.

Idle-quantum Hopping. WiFi-XL eschews measurement and, instead, uses a novel hopping scheme called idle-quantum hopping (Section 3). In idle-quantum hopping, each AP randomly hops to a channel configuration and is given a *quantum of idle time units*. The quantum determines the maximum amount of time the AP can idle/waste in a configuration (e.g., waiting time due to congestion or transmission time wasted due to loss/interference), after which the AP randomly hops to another configuration. This simple mechanism biases an AP to quickly get out of configurations which are congested (quantum depletes quickly) while spending longer times in configurations which have low congestion/loss (quantum depletes slowly). We prove that a network of APs, each locally running idle-quantum hopping, converges to a global configuration that is fair and maximizes overall network capacity, when APs are in a single interference domain. Further, our evaluations in general settings show that idle-quantum hopping achieves up to 34% higher aggregate throughput than WhiteFi [6], while simultaneously improving the throughput of the slowest flows by $4\times$.

Sub-noise Communication Support. A second characteristic of WWANs is that they provide wider coverage in sparsely populated areas by using higher AP transmit power. In white spaces, APs are allowed a maximum transmit power of 4W while mobile clients, which are battery constrained, are limited to 100mW [11]. When APs and clients use different transmit powers, range in WiFi is unfortunately determined by the *minimum* of AP's and client's power. For example, our measurements show that a 4W white space AP can reach a client 970m away at the 6dB minimum Signal-to-Noise ratio (SNR) supported by WiFi. However, 100mW client transmissions reach the AP at or above 6dB SNR only if the client is within 290m of the AP; from 970m away, clients reach the AP

at -10dB SNR! Thus, even though the 4W AP can potentially cover about 3 Km^2 , its effective coverage using WiFi is limited to only 0.26 Km^2 . In fact, clients roaming in about 80% of the 3 Km^2 area will reach the AP at sub-noise ($< 0\text{ dB}$) SNRs. Thus, supporting sub-noise SNRs is crucial for increasing coverage.

Cellular networks support sub-noise SNRs by using an efficient synchronous design that relies on the spectrum being licensed. A key challenge in supporting sub-noise SNRs in unlicensed spectrum is that the carrier sensing time increases dramatically, reducing efficiency to under 15% (Section 4).

Code Domain Contention (CDC). Instead, WiFi-XL uses a novel contention mechanism called CDC that is more efficient. CDC leverages a property unique to sub-noise SNR regimes that we term the *noise floor buffer*: multiple clients transmitting preambles at -10dB simultaneously only increases the AP noise floor marginally. Thus, CDC clients contend *concurrently* by transmitting a pseudo-noise code (preamble), randomly selected from a set of codes. The AP correlates to identify the winning client, that then transmits its uplink packet. Note that, in CDC, codes are used only during contention; packet transmissions may use the spectrally more efficient OFDM. Further, a unique feature of our CDC design is that the APs can detect and resolve almost all contention collisions.

To summarize, we make the following contributions:

- Highlight, through measurements and analysis, the challenges of channel trichotomy and sub-noise SNR support that impede WiFi effectiveness in WWANs.
- Propose, implement and evaluate idle-quantum hopping for combatting increased interference due to channel trichotomy.
- Propose, implement and evaluate CDC for efficiently supporting sub-noise SNR communication to enable range extension.

2 Measurement Study

Even though WiFi transmissions occur only between AP and its clients, successful operation of CSMA relies on clients and AP being able to sense each other’s transmis-

sions. Thus, the design choice of having both APs and clients use CSMA alike in WiFi relies on an implicit assumption that the three wireless channels, *AP-AP*, *AP-Client* and *Client-Client*, are the same. In fact, all existing WiFi literature is based on this assumption that these three channels are identical.

In this section, through a careful measurement study in wide-area settings, we show that these three wireless channels in WWANs (unlike in WLANs) are “significantly” different from each other in terms of propagation characteristics, a fact we refer to as the *channel trichotomy*. Specifically, we find that,

- *Client-Client channel is far inferior to the AP-Client channel.* This is because RF propagation between two clients at the street level faces far more obstructions compared to that between APs and clients with APs located on top of buildings or towers.
- *AP-AP channel is far superior to the AP-Client channel* as all APs are typically located much higher above the street level.

While there exist several measurement studies that characterize the signal propagation characteristics between APs and clients in wide area cellular settings, to the best of our knowledge, *our measurement study is the first attempt to characterize all three channels – Client-Client, AP-Client, AP-AP – and expose their impact on overall system design.*

2.1 RF Propagation models

RF propagation modelling is a well studied area and there exist a number of models for various settings *e.g.*, [14, 9]. A commonly used model is the Log Distance Path Loss (LDPL) model, given by

$$P_r(d) = P_0 - 10\gamma\log_{10}d + N \quad (1)$$

where d is the distance between transmitter and receiver, P_0 is the received power at 1m measured in dBm, γ is the path loss exponent, and N is Gaussian random variable that models the effects of shadowing.

Effect of AP placement height. There are several models *e.g.*, [9, 14], that have studied the affect of height on RF propagation. The most relevant model in our setting is the Egli Model [9] that applies to AP heights less than 30m. *The Egli model (eqn 2) predicts that doubling height of the*

transmitter/receiver results in a 6dB gain in the received power.

$$P_r(H_{AP}, H_C) = P_r(h_{AP}, h_C) + 20\log_{10}\left(\frac{H_{AP}}{h_{AP}}\right) + 20\log_{10}\left(\frac{H_C}{h_C}\right) \quad (2)$$

2.2 Measurement Methodology

We conducted signal strength measurements in two areas: *Urban*, with dense buildings three to four floors high, and *Suburban*, with open lawns and two to four floor buildings spaced wide apart. We were granted license to transmit in the 766-782 MHz frequency band and consequently all experiments were conducted at 774MHz centre frequency. The transmitter was a SFF-SDR Software Defined Radio platform [16]. We used two different devices simultaneously as receivers to collect the measurements – a spectrum analyser and another SFF-SDR device. The key advantage of using the SDR platform as a receiver was that by correlating against a long PN sequence, we could receive and measure RSS for signals below the noise floor. The spectrum analyser helped us calibrate our measurements in the SDR to actual dBm values. The measurement locations of the receivers were obtained using a GPS unit.

Client-Client Channel Measurements. For measuring the Client-Client channel, we used two cars. A transmitter and a receiver were placed on each car with antennas on top (about 1.5m high) and the cars were driven through the streets of *Suburban* and *Urban* to get measurements from several hundred pairs of locations.

AP-Client Channel Measurements. Our license only permitted us to transmit at 4W (36 dBm) power at two fixed locations – i) on a roof-top 15m high in an *Urban* area and ii) on a roof-top 10m high in a *Suburban* area. Measurements were collected at several hundred locations by placing the receiver antenna on top of a car and driving on the roads in these areas.

AP-AP Channel Measurements. Our transmitters were already on rooftops. For the receiver, we rented a van with a mast about 10m high (shown in Figure 7), allowing us to measure the received signal at a height of 10m. We were not allowed to drive in the streets with the mast up, so we were only able



Figure 7: Van with 10m mast

to conduct measurements at a locations where we could park the van. At each of these locations we measured RSS at the street level and at a height of 10m in order to quantify the gain (Δ) that an AP located at a height of 10m would see compared to the street level.

Area	Path Loss (γ)		Δ (dB)
	Client-Client	AP-Client	AP-AP
<i>Suburban</i>	3.5	3.1	6
<i>Urban</i>	3.8	3.5	8

Table 1: Measurement-derived Channel Trichotomy Parameters

2.3 Observations

Figures 1- 3 depict the measured RSS as a function of distance for *Urban* and *Suburban* for each of the three communication channels. For AP-Client and Client-Client channels, the figures also depict an LDPL model (dashed line) fit to the measurement data. Table 1 summarizes the measurement derived path loss exponent (γ) values for the AP-Client and the Client-Client channels for *Urban* and *Suburban* areas. Further, it also depicts the average gain Δ in received power at a height of 10m compared to the street level (1.5m). *Our measurements indicate a much more conservative received power gain (6-8dB) for 10m AP height compared to the 16dB gain predicted by the Egli model ($H_C = 10, h_c = 1.5$ in eqn 2).*

A Model for WWANs. Based on our measurements, we arrive at the following model for WWANs for received power $P_{rec}(d)$ at a distance d , when the power 1m away from the transmitter is P_{xmit} ,

$$P_{rec}(d) = \begin{cases} P_{xmit} + \Delta - 10\gamma_{AP-C} \log_{10}(d) & AP - AP \\ P_{xmit} - 10\gamma_{AP-C} \log_{10}(d) & AP - Client \\ P_{xmit} - 10\gamma_{C-C} \log_{10}(d) & Client - Client \end{cases} \quad (3)$$

In Eqn 3, γ_{AP-C} , γ_{C-C} are the path loss exponents for AP-Client and Client-Client channels given in Table 2 (e.g., 3.1 and 3.5 respectively for *Suburban*), and Δ the gain AP-AP channel enjoys over the AP-Client channel (e.g., 6dB for *Suburban* and 8dB for *Urban*). The predicted value of the above model should be the expected average RSS value at a given distance.

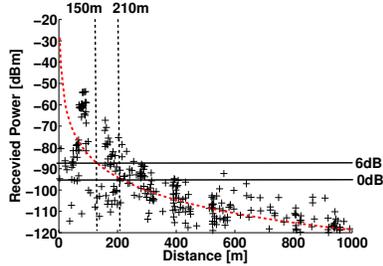


Figure 1: Propagation between clients in *Suburban*.

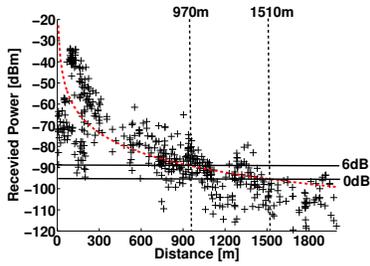


Figure 2: Propagation between APs and clients in *Suburban*.

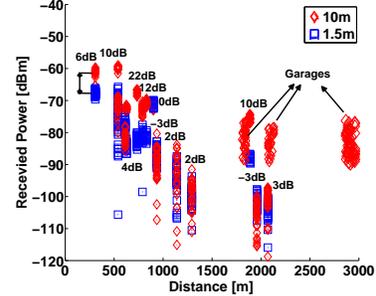


Figure 3: RSS at height of 10m and ground level *Suburban*.

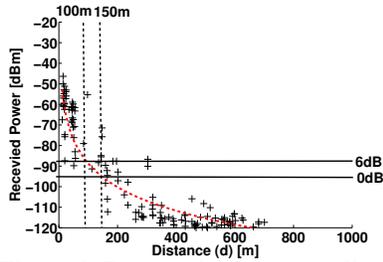


Figure 4: Propagation between clients in *Urban*.

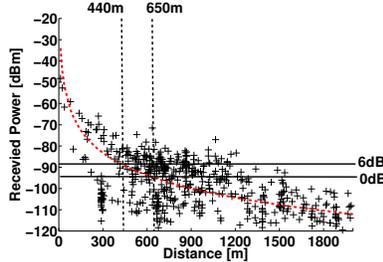


Figure 5: Propagation between APs and clients in *Urban*.

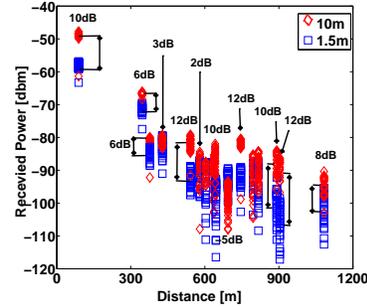


Figure 6: RSS at height of 10m and ground level in *Urban*.

Area	Transmission Range[m] (6dB SNR)					Interference Range[m] (0dB SNR)				
	Client-Client	AP-Client (100mW)	AP-AP (100mW)	AP-Client (4W)	AP-AP (4W)	Client-Client	AP-Client (100mW)	AP-AP (100mW)	AP-Client (4W)	AP-AP (4W)
<i>Suburban</i>	140	290	450	970	1510	210	440	680	1510	2350
<i>Urban</i>	100	150	250	440	740	150	230	390	650	1100

Table 2: Summary of Ranges

Transmission and Interference Ranges. Table 2 provides the transmission and interference ranges for APs transmitting at 100mW (default) and 4W (max FCC limit), calculated using the model in eqn 3 and the values in Table 1. The transmission range was chosen as the distance where received SNR drops to 6dB (with -95dbm as noise floor for 6MHz white space channels) and the interference range where the received SNR drops to 0dB. Based on Table 2, the key observations are:

- Transmitting at default power of 100mW, a WiFi AP in white spaces (at 774MHz) has a reach of 150m in *Urban* and 290m in *Suburban*. Transmitting at 4W power, a WiFi AP can potentially reach 440m in

Urban and 970m in *Suburban*. However, since the clients' transmit power cannot exceed 100mW, WiFi cell radius is limited to 150m in *Urban* and 290m in *Suburban* for both 100mW and 4W APs.

- Interference among APs is likely to be the dominant problem given the superior propagation characteristics of the AP-AP channel. For 100mW transmit powers, APs as far as $2.5\times$ the cell size (390m *Urban*, 680m *Suburban*) will cause interference while for 4W transmit power APs as far as $8\times$ the cell size (1.1Km, 2.35Km) will cause interference.
- As later discussed in Section 4, the hidden node problem in WiFi will be exacerbated for some uplink

(client to AP) transmissions since potentially interfering APs and clients will not be able to sense client transmissions. Thus, enabling RTS/CTS for uplink transmissions may be necessary.

2.4 Implications

These observations lead to two key implications:

- *Need for Judicious Frequency Usage.* Interference from large number of nearby APs (within 2.5 to 8 × cell radius) can bring down the network capacity dramatically. Thus, it is crucial for adjacent APs to operate on different channels.
- *Need to modify WiFi for Extending Range.* For sparsely inhabited regions, increased coverage may be desirable. However, since simply increasing the AP transmit power to 4W does not increase coverage, WiFi needs to be modified for range extension.

3 Idle-Quantum Hopping

Given the uncoordinated nature of deployments in unlicensed spectrum, a decentralized scheme that separates adjacent WWAN APs into different channels is desirable. Almost all existing schemes are based on selecting the “best” channel by trying to measure the number of competing APs and/or traffic load in each channel. In this section we argue that these schemes are fundamentally unsuitable for WWANs operating in white spaces due to *interferer blindness* and propose *IQ hopping*.

3.1 Interferer Blindness in WWANs

Consider the example in Figure 8 where R_{Tx} and R_{Int} are respectively the transmission and interference ranges of AP0. In WiFi, R_{Tx} is determined by the minimum SNR of 6dB required to decode packets and R_{Int} is chosen to be the distance from where transmissions are received at 0dB at AP0. It can be shown that $\frac{R_{Int}}{R_{Tx}} \leq 2$ and typically ranges between 1.4 to 2 (for $\gamma \in (2, 4)$).

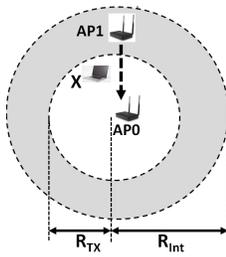


Figure 8: Interference in WLAN

Measurement of the number of interfering networks in a channel requires an AP0 to decode beacons from interfering APs. AP0 however, cannot decode packets transmitted from interfering devices such as AP1, that are further than R_{Tx} (area shaded in Figure 8) – about 60% of the total interfering devices. Consequently, in existing schemes [18], clients (*e.g.*, X in Figure 8) perform these measurements and report them to their AP. This allows the AP, to accurately detect all interfering APs/traffic load within a radius of $2R_{Tx}$ which includes all nodes in radius of R_{Int} in wireless LANs.

For WWANs, $\frac{R_{Int}^{AP-AP}}{R_{Tx}} > 2$ and can range between 2.5 at default 100mW transmit power to 8 at 4W (Section 2). Thus, even with client

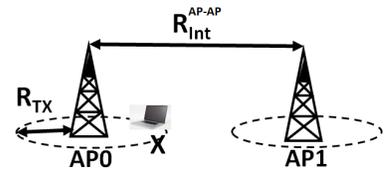


Figure 9: Interference in WWLANs

assistance, APs in WWANs will not detect interfering APs that are farther than $2R_{Tx}$ – almost 40% (at default power) to 95% (at 4W) of potential interferers. Thus, *in WWANs, both the AP and its clients will be blind to a large fraction of interferers and thus unable to measure the “goodness” of the channel correctly.*

3.2 IQ hopping

In this section, we start by describing an idealized version of IQ hopping that is simplified to its bare minimum in order to illustrate its simplicity and optimality. A practical version that takes variable channel width, implementation complexity, packet losses, etc. into account is described later in Section 3.6.

At a high level, idle-quantum (IQ) hopping takes a measurement-free approach by making a simple but remarkably effective modification to random hopping. In basic random hopping, upon changing a channel each AP generates a random quantum of time τ that dictates the duration for which the AP will stay in the current channel before it hops again. The simple change to obtain IQ hopping from random hopping is that instead of using the quantum τ as a measure of absolute time, IQ hopping uses the quantum τ as a measure of forced idle-time *i.e.*, time when the AP has packets but is unable to transmit them.

Figure 10 shows the Pseudo-code for an ideal version of IQ hopping and replacing the conditional in the IF statement in Line 6 with TRUE would give us the random hopping scheme. Note that, to avoid synchronization issues, the quantum τ is drawn from $Exp(\alpha)$ - an exponential distribution with mean $\frac{1}{\alpha}$ (Line 3, Figure 10).

```

1: repeat
2:   HopToRandom(Channels)
3:    $\tau = Exp(\alpha)$ 
4:   repeat
5:     for every clock tick of duration  $\delta$ 
6:       if Packets waiting but not being transmitted then
7:          $\tau = \tau - \delta$ 
8:       end if
9:     until  $\tau > 0$ 
10: until true

```

Figure 10: Pseudo-code for Ideal version of IQ hopping

The key intuition behind IQ hopping is that, in a congested channel, idle-time quantum remaining (τ) decreases quickly leading to the APs hopping out of the channel. If the AP has no competing traffic, it never switches its current channel.

3.3 Proof of Optimality of IQ hopping

For a contention domain with N nodes and K channels we prove that IQ-hopping converges to the optimal solution. More precisely, when $K \geq N$, IQ hopping converges to a situation where each channel has at most one node utilizing it. That is, each node obtains its own channel to transmit and stops hopping subsequently. Furthermore, the convergence is fast: in $O(N^2K)$ hops this solution is attained. When $K < N$, then the number of nodes in any channel converges to N/K . We leave the convergence rate as an open problem.

Theorem 1. *In a contention domain, IQ hopping converges to the following solution. (a) If $K \geq N$, then each channel has congestion at most 1, that is, each node obtain its own channel. (b) If $K < N$, then each channel has congestion N/K .*

Proof. We consider time indexed by hops, that is, each time a node hops from one channel to another, we note

that as a unit of time. If two nodes hop simultaneously, then we break ties arbitrarily.

Consider the *congestion* vector $\mathbf{x}^t := \langle x_1^t, \dots, x_K^t \rangle$ at time t , where x_i^t denotes the number of nodes in channel i just after the t th hop. Let v_t be the node which hops at time t . Let A_t be the channels with $x_i^t > 1$. Note v_t must have occupied a channel in A_t since otherwise it wouldn't have hopped. Let p_i^t be the probability that v_t was in channel i . Since the idle times are initialized to exponential random variables, we can precisely figure out p_i^t .

Claim 2. *If $A_t \neq \emptyset$,*

$$p_i^t = \frac{(x_i^t - 1)}{\sum_{i \in A_t} (x_i^t - 1)} > 0 \text{ for } i \in A_t \text{ and } p_i^t = 0 \text{ otherwise.} \quad (4)$$

Proof. This proof crucially uses the property of exponential random variables, In particular we use the following three properties.

- If $X \sim \exp(\lambda)$, then $c \cdot X \sim \exp(\lambda/c)$.
- If $X_i \sim \exp(\lambda_i)$ for $i = 1, \dots, k$, then $\min(X_1, \dots, X_k) \sim \exp(\lambda_1 + \dots + \lambda_k)$.
- If $X_i \sim \exp(\lambda_i)$ for $i = 1, \dots, k$, then $\Pr[X_i = \min(X_1, \dots, X_k)] = \lambda_i / \sum_{j=1}^k \lambda_j$.

Let us consider the situation after the $(t-1)$ th hop. Some vertex v has hopped to a channel and have set it's initial idle time to an exponential random variable with $\lambda = 1$. The idle times of all other nodes have been decremented by some amount. Since these are exponential random variables which are non-zero, the distribution of each random variable is still the same exponential distribution. Furthermore, these variables in a channel are independent, and variables across channels are also independent. Now the vertex v_t lies in channel i if its idle time at time $(t-1)$ divided by the rate of this channel is the smallest. The rate of decrement of vertices in channel i is $\frac{x_i^t}{x_i^t - 1}$. Therefore, idle time of a node in channel i divided by rate is also an exponential random variable with parameter $\lambda = x_i^t / (x_i^t - 1)$. The minimum idle time divided by rate among nodes in channel i is another exponential random variable with parameter $\lambda \cdot x_i^t \cdot \frac{x_i^t - 1}{x_i^t} = \lambda(x_i^t - 1)$.

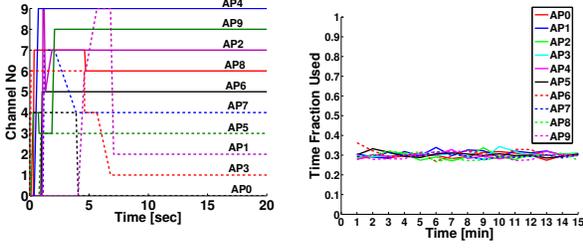


Figure 11: Self-Organization of IQ-Hopping in a single contention domain

Figure 12: Convergence of IQ-Hopping

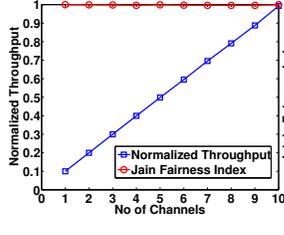


Figure 13: IQ-Hopping in a single contention domain with few channels

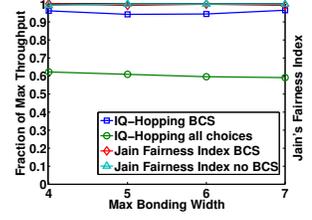


Figure 14: Performance of IQ-Hopping with channel bonding

Now the probability the minimum is from i , is precisely p_i^t . \square

Fix a particular channel i . Note that x_i^t follows a Markovian process whose transition probabilities depend on p_j^t for all the channels j . In particular, we get that x_i^{t+1} evolves as follows.

$$x_i^{t+1} = \begin{cases} x_i^t - 1 & \text{with probability } p_i^t. \\ x_i^t + 1 & \text{with probability } \frac{\sum_{j \neq i, j \in A_t} p_j^t}{K-1}. \\ x_i^t & \text{otherwise.} \end{cases}$$

This is because with probability p_i^t , v_t lies in channel i and the congestion on channel i decreases. Otherwise, with probability p_j^t , v_t belongs to some other channel j , and since v_t hops to a random channel other than j , x_i^t increments with probability $1/(K-1)$. With remaining probability nothing x_i^t remains the same.

Let (x_1^*, \dots, x_K^*) be the stationary distributions of these K (coupled) Markovian processes. It is not hard to see that the chains are irreducible and aperiodic, and therefore this is well defined. Let A^* be the channels i with $x_i^* > 1$. Let (p_1^*, \dots, p_K^*) be the resulting probabilities when x^* is substituted in (4). At stationary distribution, if $A^* \neq \emptyset$, we must have, for all i , $p_i^* = \frac{\sum_{j \in A^* \setminus i} p_j^*}{K-1}$.

We now claim that either $A^* = \emptyset$ or $A^* = \{1, 2, \dots, K\}$. That is, either all channels have stationary congestion ≤ 1 or all channels have stationary congestion > 1 . To see this, note that for $i \in A^*$, we have $Kp_i^* = \sum_{j \in A^*} p_j^*$. Adding over all $i \in A^*$ gives $(K - |A^*|) \sum_{i \in A^*} p_i^* = 0$. Since $p_i^* > 0$ for all $i \in A^*$, we get either $K = |A^*|$ or $A^* = \emptyset$. Note

that $N = \sum_{i=1}^K x_i^*$ since the sum of congestion vector is always the number of nodes. When $K \geq N$, we must have $A^* = \emptyset$ since $N \geq |A^*|$. That is, $x_i^* \leq 1$ for all i , that is, the congestion in each channel is at most 1 proving part (a). When $K > N$, we must have $A^* = \{1, \dots, K\}$. Otherwise $A^* = \emptyset$ implying $N \leq K$. Now, since $Kp_i^* = \sum_{j \in A^*} p_j^* = 1$, using (4), we get $K(x_i^* - 1) = \sum_{i=1}^K x_i^* - K = N - K$. This gives $x_i^* = N/K$ for all channels. \square

Theorem 3. *In a contention domain with $K \geq N$, IQ hopping converges in $2N^2K$ hops.*

Proof. When $K \geq N$, we can argue fast convergence by considering the “energy” function

$$\Phi(\mathbf{x}^t) := \sum_{i=1}^K (x_i^t)^2,$$

which is minimized at the optimum solution. The following claim shows that $\mathbf{E}[\Phi(\mathbf{x}^t)]$ monotonically decreases. Since $\Phi(x) \leq N^2$ for any x (this follows from the fact that $\sum_{i=1}^K x_i = N$ and the Cauchy-Schwartz inequality), the theorem follows.

Lemma 4. $\mathbf{E}[\Phi(\mathbf{x}^{t+1})] \leq \mathbf{E}[\Phi(\mathbf{x}^t)] - 2/(K-1)$.

Proof. For any fixed vector \mathbf{x}^t , we upper bound the conditional expectation $\mathbf{E}[\Phi(\mathbf{x}^{t+1})|\mathbf{x}^t]$ as

$$\mathbf{E}[\Phi(\mathbf{x}^{t+1}) - \Phi(\mathbf{x}^t) | \mathbf{x}^t] = \sum_{i \in A_t} \frac{p_i^t}{K-1} \sum_{j \neq i} 2(x_j^t - x_i^t + 1) \quad (5)$$

To see this, note that the summand in the second summation is the drop in the energy when a node moves from channel i to channel j . We now upper bound the RHS

of (5) by $-2/(K-1)$. This proves the lemma by taking expectations on the LHS and RHS.

Firstly, in the second summation we let j run over all channels and then *subtract* the case $j = i$. More precisely,

$$\sum_{i \in A_t} \frac{p_i^t}{K-1} \sum_{j \neq i} 2^{(x_j^t - x_i^t + 1)} = \sum_{i \in A_t} \frac{p_i^t}{K-1} \sum_{j=1}^K 2^{(x_j^t - x_i^t + 1)} - 2 \sum_{i \in A_t} \frac{p_i^t}{K-1} \quad (6)$$

The second summation in RHS of (6) is precisely $2/(K-1)$ since sum of probabilities is 1 (we are in the case when A_t is not empty; otherwise we are done). We now upper-bound the first summation as follows.

$$\begin{aligned} \sum_{i \in A_t} \frac{p_i^t}{K-1} \sum_{j=1}^K 2^{(x_j^t - x_i^t + 1)} &= \frac{2}{K-1} \sum_{j=1}^K \left(\frac{x_j^t \sum_{i \in A_t} p_i^t - \sum_{i \in A_t} (x_i^t - 1) p_i^t}{\sum_{i \in A_t} (x_i^t - 1) p_i^t} \right) \\ &= \frac{2}{K-1} (N - \sum_{j=1}^K \sum_{i \in A_t} (x_i^t - 1) p_i^t) \end{aligned}$$

We claim now that $\sum_{i \in A_t} (x_i^t - 1) p_i^t \geq 1$. Believing this for now, we get

$$\sum_{i \in A_t} \frac{p_i^t}{K-1} \sum_{j=1}^K 2^{(x_j^t - x_i^t + 1)} \leq \frac{2}{K-1} (N - K) \leq 0 \text{ since } N \leq K.$$

which when substituted in (6) implies the lemma. To prove the inequality, note that $x_i^t > 1$ for $i \in A_t$ which in turn implies $(x_i^t - 1) \geq 1$. In particular, $\sum_{i \in A_t} (x_i^t - 1)^2 \geq \sum_{i \in A_t} (x_i^t - 1)$, which implies $\sum_{i \in A_t} (x_i^t - 1) p_i^t \geq 1$ using (4). \square

General graphs. For any general graph, if the number of channels is greater than or equal to the number of nodes, each node will find and settle in a unique channel by running IQ hopping. Otherwise, finding the optimal channel assignment is harder than the NP-hard graph coloring problem. Nevertheless, in simulations with large number of randomly placed APs and limited channel availability as per the white space database, we find that IQ hopping performs better than existing schemes such as WhiteFi [6] in both aggregate capacity and fairness (Section 8).

3.4 Illustrative Examples

We now illustrate the functioning of IQ hopping through some simple examples to help provide the reader a feel for how IQ hopping works. We first show simulation results with 10 contending APs and 10 channels available (numbered 0 through 9). All APs start by selecting channel 0

and then run IQ hopping. In this simulation, packet losses only occur due to collisions, α is chosen to be 1 sec and there is only saturated downlink traffic. Figure 11 depicts the sequence of channel hops for each AP with time. The remarkable observation is that, at the end of 10 seconds, all the AP have self-organized and settled into a separate channel each!

When the number of available channels is less than number of APs. Figure 12 depicts the fraction of air-time obtained by each AP over a minute where there are only 3 available channels for 10 APs. As seen from Figure 12, all the APs keep hopping, however each AP obtains an average share of 0.3 (the Jain's Fairness Index was 0.99974). *This example demonstrates how IQ hopping converges to quickly give each AP its fair share.*

Figure 13 depicts the average fraction of time that each AP (total 10) had access to a channel as the number of channels is increased from 1-10. Figure 13 also depicts the Jain's Fairness Index for all the APs. As seen from Figure 13, the average airtime for each AP for k channels is almost equal to $\frac{1}{k}$ and the Jain's fairness index is close to 1 for all k . Thus, IQ hopping provides a fair share to every AP, while providing full aggregate utilization.

3.5 IQ hopping with Channel Bonding

\square Unlike WiFi, spectrum in white space is fragmented *i.e.*, available in non-contiguous pieces of different widths. For example, at the time of this writing, Kansas has a total of 21 whitespace channels (each 6MHz) available, in 7 non-contiguous pieces – 21, 23-29, 31-34, 39-42, 44, 47-49 and, 51. Consequently, in order to utilize the available spectrum efficiently, white space radios make use of channel bonding *i.e.*, combine one or more adjacent available channels in a single wider channel. Radio hardware however limits the maximum number of channels that may be bonded (typically 4-8 in present day hardware). Thus, for channel selection, APs have to choose a set of channels to bond. This dramatically explodes the number of choices to explore. In our Kansas example, there are 51 possible choices for maximum bonding width of 4. This explosion of choices significantly increases the overhead of measurement based approaches, making them further impractical in WWANs.

Bonded Channel Selection (BCS) In IQ hopping, instead of randomly picking from all possible choices,

we only pick from the set of non-overlapping bonded channels with maximum possible bonding width. In order to generate this set, we traverse all the channels from left to right combining as many contiguous channels as possible. Suppose that we can bond up to 4 channels, then we select the bonded-channel choices as $\{\{21\}, \{23,24,25,26\}, \{27,28,29\}, \{31,32,33,34\}, \{39,40,41,42\}, 44, \{47,48,49\}, \{51\}\}$. We refer to this scheme as *Bonded Channel Selection* (BCS) in the rest of the paper. We now argue that BCS is the best strategy if all devices use BCS.

Choosing Non-overlapping channels is the best strategy. Suppose for two APs, A and B, A chooses a bonded channel $\{1,2,3\}$ while B chooses $\{2,3,4\}$. Since they overlap they can only use their bonded channels half the time without causing collisions. Thus, both receive $\frac{3}{2}$ resulting an aggregate usage of 3. If however, A and B choose from the non-overlapping set $\{1,2,3\}$ and $\{4\}$, then they don't interfere and use an aggregate of 4 channels. Fairness can be achieved through alternating.

Choosing the maximum possible bonding option is better than smaller widths in practical settings. Suppose we have three APs and three channels $\{1,2,3\}$, suppose that APs can bond up to 3 channels each. If we assign a separate channel to each AP (e.g., A-1, B-2, C-3) then unless all three APs have traffic to send all the time some of the spectrum will be wasted since, no other AP will be able to use it. If however, all three APs choose to use $\{1,2,3\}$, then CSMA will let them get roughly a third share in highly loaded conditions and APs will also be able to use the spare spectrum if one or more APs are lightly loaded. Given load variations among APs is more likely in practice than all APs being fully saturated, choosing maximum bonded width is the better option.

Modifying the Quantum to Include Channel Width. In order to adapt the IQ hopping algorithm to use variable channel width, we need to capture the fact that the device obtains more throughput in the same time when using a wider channel. Consequently, if the AP is using a channel that is narrower than the widest possible option (*MaxWidth*), it can be considered a waste and the idle-time quantum needs to be decremented to reflect this. Thus, we modify IQ-Hopping as defined in Figure 10 so that τ is decremented by $\delta * (1 - \frac{CurrentChannelWidth}{MaxWidth})$ when the conditional in Line 6 of Figure 10 is FALSE.

Effectiveness of BCS – an example. To provide an in-

tuition to the reader as to how much BCS can help, we ran a simulation of IQ hopping for 10 APs using Kansas's list of channels with and without BCS. Figure 14 depicts the fraction of maximum possible throughput obtained for various values of maximum channel bonding width as well as the corresponding Jain's fairness index. As seen from Figure 14, IQ hopping with BCS performs significantly better than not using BCS and in fact achieves close to the maximum possible aggregate capacity for bonding widths of 4 and above.

3.6 IQ hopping in practice

For ease of exposition, we started with an idealized/simplified version of IQ hopping in Figure 10. While a few modifications are needed to the ideal version of IQ hopping due to practical considerations, these changes are not extensive. In fact, we have implemented IQ hopping on two platforms, madwifi and SORA. In this section, without going into the platform specific details (described in Section 7), we present a practically realizable pseudo-code for IQ hopping in Figure 15.

```

1: repeat
2:   HopToRandom(Channels)
3:    $\tau = \text{Exp}(\alpha)$ 
4:   repeat
5:     @ Periodic Timer Event with interval  $\Delta$ 
6:     if QSize > 0 then
7:        $\tau -= \Delta * \text{MaxWidth}$ 
8:     end if
9:     @ ACK Received Event
10:     $\tau += \text{CurrentChannelWidth} * (\text{Total Packet Time})$ 
11:    @ Pkt Received Event
12:     $\tau += \text{CurrentChannelWidth} * (\text{Total Packet Time})$ 
11: until  $\tau \geq 0$ 
12: until true

```

Figure 15: Pseudo-code of IQ hopping in practice

In order to determine whether there are packets to be transmitted we can use the queue size of packets waiting to be transmitted (Line 6). In a practical implementation, it is impossible to update τ on every clock tick. Rather, it can be done using a periodic timer (Line 5) whose granularity may typically be larger than the transmission time of a single packet. Consequently, in order to determine the time-frequency that was used productively we have to rely

on events such as ACKs and packets received. Each time the periodic timer goes off and queue is non-empty, by default we reduce τ by the timer interval Δ times the maximum channel width possible (Line 7). Further, each time an ACK is received, it means that some part of the spectrum was productively used, hence some of the τ must be replenished and this is given by the “Total Packet Time” (which includes packet transmission time, ACK transmission time) times the current channel width. Similarly, when the AP receives an uplink packet, even this represents useful time, consequently we once again replenish τ the same way as when ACK was received. Note that since we consider useful time only when ACKs are received, this version of IQ hopping has the side-benefit of implicitly also capturing the effects of losses due to bad channels (e.g., fading), cumulative interference from distant APs, temporary traffic/load variations in neighbouring APs, etc.

4 CDC: Extending WiFi’s Range

We now focus on the coverage of a white space WWAN AP, especially in rural or sparsely populated settings with few users who are spread across a wide-area. Covering a few square Kms (small village or region) with a single AP will be crucial for cost saving. For example, specifications of range networks’ GSM base station products show a simultaneous maximum voice call capacity as few as 7 calls, while service range shows coverage of large distances of 1-30Km [4].

4.1 WiFi range outdoors in white spaces

We used our WWAN Sora-based AP in the *Urban* setting for this measurement. We transmitted WiFi packets at the base rate from the AP at 4W while the client was in a car and we plot the packet reception success rate versus distance in Figure 16. From the figure, one can see the range of the 4W AP is around 450m, matching well with the ranges estimated from measurements in Section 2. The figure also shows the range at 100mW and one can see that it is significantly lower, between 100-150m.

4.2 CSMA with extended range

Sub-noise SNR support. As discussed in Section 1, the key requirement for extending range for WiFi in white spaces is enabling support for sub-noise communication. If we define range as distance where a 4W AP’s SNR drops to 6dB, we would need to support reception of a -10 dB SNR signal from a 100mW client. A key challenge in supporting sub-noise SNRs is enabling carrier sensing for sub-noise SNRs, which is our focus next.

Need for 320 μ s preamble. We performed wide-area measurements to determine the length of the PN-sequence required for detection of a -10 dB SNR signal. While the details are provided in Section 7, we find that 320 μ s preamble is necessary for reliable detection. The expected channel access overhead for this preamble can be computed to 3360 μ s which is an order of magnitude larger overhead in 802.11af as discussed in in Section 5.

Need for RTS/CTS. In addition to preamble dilation, RTS/CTS becomes essential when CSMA is used with extended range. We analyze the amount of hidden terminals when CSMA is used in WWAN settings in Section 5 and find that while CSMA works well for downlink transmissions with negligible hidden nodes, up to 90% of APs in the region may act as hidden nodes when CSMA is used for uplink transmissions. Thus, RTS/CTS is mandatory for uplink when CSMA is used for extending range.

Efficiency under 15%. When the extended channel access overhead is combined with RTS/CTS, efficiency (defined as fraction of airtime with useful data) of the extended range CSMA drops to under 15% (as described in Section 5) even for a 6MHz narrow channel. In comparison, 802.11af efficiency for 6MHz channels is over 80%. In congested conditions, RTS/CTS collisions will increase, further undermining efficiency. Thus, while CSMA can be made to work with extended range, it will come at the cost of severely limited cell capacity.

4.3 Other Contention Mechanisms

Frequency Domain Contention. A recent approach for achieving efficient channel access has been to enable concurrent frequency domain contention using tones [24, 22]. However, FCC white space regulation on clients’ transmission also dictates that the client power should be evenly spread over the entire channel [11]. Thus, concen-

trating client energy into tones violates FCC regulation for operation in white spaces.

Cellular support for Sub-noise. Cellular network protocols such as LTE support sub-noise SNRs efficiently [15]. However, efficiency in cellular comes at the cost of higher initial packet latency (e.g., initial packet access latency can be as high as 1sec after which a data channel is dedicated to this client for 10+sec [15]). Thus, collisions in LTE random access channel is not a major concern since the latency cost of collision is amortized over the duration of the reservation. However, in unlicensed spectrum, we cannot have dedicated channels and, thus, collisions are a key concern.

4.4 Code Domain Contention (CDC)

In this section, we propose a novel contention mechanism for clients’ uplink traffic – code domain contention (CDC) – providing *high efficiency at sub-noise SNR regimes*. Further, CDC is *almost collision-free* as it can detect and resolve potential collisions very efficiently as a part of the contention process.

Figure 17 depicts the operation of CDC pictorially. CDC begins when the AP initiates it by transmitting a Solicit Contention Message (SCM). Upon receiving SCM, all clients that intend to transmit, pick a random PN-sequence from a common bank of PN-sequences (16 in our current implementation) and transmit them simultaneously to the AP. For example, in Figure 17, the three clients C1, C2 and C3 that intend to transmit pick PN3, PN5 and PN3 respectively. Each client waits a small amount of random time (between 0 to $26\mu\text{s}$ in our current implementation) to artificially introduce a random jitter in the transmissions.

Collision detection. The AP maintains a bank of correlators, one for each of the PN sequences, that it uses to correlate against the received signal. Thus, the AP detects all the transmitted PN-sequences. The jitter introduced by the clients allows the AP to detect potential collisions. As seen in Figure 17, the AP detects two clear peaks for PN3 indicating a collision while PN5 is received in the clear. Now the AP can pick PN5 as the winner, for example, using the rule “Pick the PN sequence with the lowest collision-free index”, thereby avoiding collisions. A signal snapshot of a collision of three nodes based on our experiments is illustrated in Figure 18. This ability

to detect collisions, which is not possible in frequency domain contention techniques, dramatically improves efficiency when several clients contend at the same time. Further since PN-sequences have identical properties to noise, they have an even energy spread across the entire channel and hence do not violate FCC regulations.

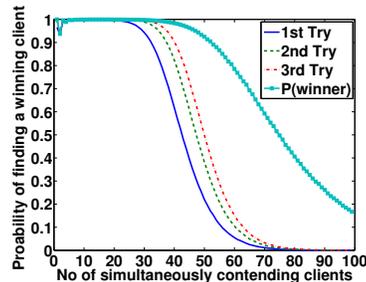


Figure 19: Graceful Degradation in CDC with increasing contenders

How AP picks the winning client A significant advantage of detecting collisions in the operation of CDC is that it can ignore colliding PN sequences and pick only among the non-colliding ones. Thus, in CDC the AP simply picks the non-colliding PN sequence with the smallest id. Thus, there needs to be only one client that picks a unique PN sequence for CDC contention to succeed. Figure 19 ($P(\text{winner})$) depicts the probability that at least one of the contending clients will pick a unique PN sequence. As seen from Figure 19 indicates that with 16 PN sequences, even up to 60 clients there will be a potential winner with a probability greater than 90%. Thus, CDC can potentially accommodate large occasional traffic bursts very well.

Why CDC is efficient – the Noise Floor Buffer. In order to be detected, a client C1 located at the cell edge must transmit a long PN-sequence ($320\mu\text{s}$ as seen in Section 7) against which the AP must correlate to detect the transmission. Now suppose another client C2 transmits a completely different PN-sequence to the AP, this PN-sequence will interfere as noise to C1’s PN-sequence. However, since received power of C2’s signal is -10dB , $10\times$ below the noise floor, it does not significantly add to the noise floor (the noise floor will rise by 0.4dB). Thus, the AP can correlate C2’s preamble as well in the same $320\mu\text{s}$, thereby allowing two nodes to contend without incurring

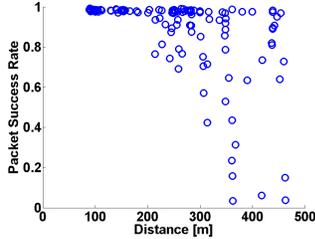


Figure 16: WiFi wide-area reception

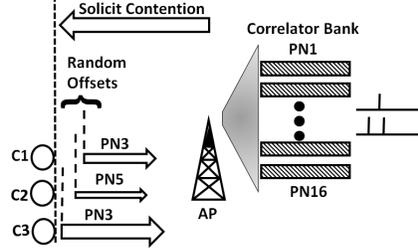


Figure 17: CDC Operation

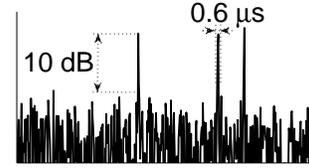


Figure 18: Collision example with 3 peaks

the overhead of two separate slots.

As more and more clients contend at the same time, the noise floor keeps increasing gradually until it becomes 3dB when 11 clients contend. Since 3dB SNR corresponds to doubling of the noise floor, by increasing the PN-sequence by a factor of two ($640\mu\text{s}$), we can tolerate up to 11 contending clients simultaneously. Compare this with CSMA which would require a minimum of 11 collision-free slots (*i.e.*, several tens of actual slots) of $320\mu\text{s}$ each to accommodate 11 clients. The reason we get these gains using PN-sequences is because the noise floor of the receiver effectively acts as buffer to shield clients from each others interference.

Power control. A key requirement for successful functioning of CDC is that no client’s PN-sequence should overwhelm another’s due to high received power. Thus, all clients, irrespective of how close or far they are from the AP, should make sure that their received powers at the AP are at -10dB SNR. Since, CDC is initiated by the AP, the clients can accurately estimate its received power and thus adjust their transmit powers accordingly as we show in our evaluation in Section 7.

Graceful degradation. Note that CDC is designed for wide-coverage in sparsely populated regions. However, one key advantage of using PN-sequences is graceful degradation of detection with increase in clients. In other words, if more than 11 contenders contend simultaneously, then the SNR will decrease, resulting in only a few (randomly chosen) PN-sequences being detected. The AP can choose a winner among these and still remain fair.

Increasing the number of number of clients increases the noise floor and leads to increased false alarms, *i.e.*, detection of a PN sequence when none was transmitted. Thus, the AP will attempt to solicit a packet from a client

with no data. In such a case the client need not respond and after a timeout of $320\mu\text{s}$ the AP concludes that this was a false alarm and then tries soliciting the next winner in the contention. Figure ?? depicts the probability of finding a winner correctly as the number of contending clients increases for the 1st, end and 3rd attempts. As depicted in Figure ??, CDC can find a correct winning client with a probability greater than 90% up to 30+ clients within the first try and up to 40+ clients for the next try. There is no significant advantage in trying any further attempts, thus we limit trials in CDC to 2 after which AP simple resends the SCM.

Accommodating larger bursts than 40 clients. When the number of contending clients is much larger than 40, the noise floor will increase by a factor of 5 or 7dB (as noise floor now becomes $1 + \frac{40}{10}$ times greater). The AP has $320\mu\text{s}$ to estimate this increase in noise floor and hence can detect this increase extremely reliably. Upon detecting this, the AP simply requests the clients to not-contend with a 50% probability (backoff probability), thus reducing the number of contenders by half for the next round of CDC. In general the AP and fine-tune the backoff probability finely after measuring the net received power during CDC.

4.5 Multi-network scenario

CDC as described in Section 4 was for a single network, *i.e.*, a single AP and its clients. How does one then use CDC in a multi-AP scenario? During the CDC phase, when client transmissions to the AP are carefully power controlled, if neighbouring networks interfere, they can severely disrupt its operation. The key observation from Section 5 is that the superior range of AP-AP link not

only protects the AP from interference but can also be used to protect its clients (similar to the protection provided by AP's CTS message in the RTS-CTS example in Section 5). This suggests that AP's can gain access to the channel using standard WiFi (802.11af) based CSMA, thereby shutting out interference from neighboring networks, and then initiate a CDC. We discuss this in more detail below.

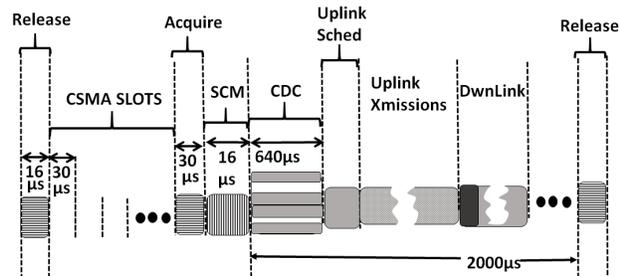


Figure 20: Operation of CDC in multi-network setting

CDC in multi-network setting. The typical operation of CDC in a multi-network scenario is depicted in Figure 20. APs contend in a WiFi (802.11af) like manner using a $16\mu\text{s}$ PN-sequence – the *Acquire* PN-sequence. Note that a $16\mu\text{s}$ PN-sequence is sufficient for sensing between APs since APs are transmitting at 4W. Note that the slot duration is $30\mu\text{s}$, accounting for the preamble, $10\mu\text{s}$ of propagation delay (3.3Km) and $4\mu\text{s}$ for tx-rx switch time. Upon gaining access to the channel, the AP can either transmit a downlink packet, initiate start of code-domain contention for enabling client uplink transmissions using an SCM as discussed in the previous section, or request certain clients to transmit their uplink packets based on a previous CDC.

In order for APs to resume CSMA, the AP that acquired the channel transmits a separate *Release* PN-sequence notifying other APs that they are free to contend once again. The Release PN-sequence is necessary as strategies such as having a time duration in the Acquire message will not function correctly since distant, possibly interfering, APs may not be able to decode messages from the AP that has acquired the channel. Further, if there is no uplink or downlink traffic to transmit, the AP can simply release the channel immediately by transmitting a Release PN-sequence.

Aggregation for Efficiency. In WiFi, channel contention is performed on a per-packet basis. CDC could also use

the same approach and schedule exactly one downlink or uplink packet for every channel access. However, this implies that the channel access overhead applies to every packet, reducing efficiency of the system. Since CDC has explicit acquire and release primitives, and has a complete picture of the requests from all uplink contenders, CDC can do better by *scheduling multiple packets* (uplink and/or downlink) when it gains access to the channel subject to a maximum *aggregation time*, $\tau = 2\text{ms}$ in our implementation). This is analogous to frame-aggregation in 802.11 standards where multiple packets are sent back-to-back after the AP gains access to the channel.

5 Inefficiency of using CSMA for Extending Range

Since CSMA is extremely effective for sharing in WiFi, we start by analyzing the feasibility of CSMA for WWANs, using the WWAN model.

5.1 CSMA in WiFi

In order to facilitate understanding of our analysis, we first provide brief background using WiFi CSMA.

Carrier Sensing (CS) in WiFi. WiFi devices (APs and clients alike) sense the channel for ongoing transmissions before transmitting to avoid collisions. To facilitate this detection, the first $4\mu\text{s}$ ($16\mu\text{s}$ for 5 MHz narrow channels) of every transmission comprises a well known Psuedo-Noise (PN) sequence [17]. WiFi devices, detect the onset of a packet transmission by continuously listening to the channel and correlating against this PN sequence. This scheme allows WiFi devices to detect transmissions received at 0dB SNR with 90% probability.²

Transmission range in WiFi. Figure 21 depicts a WiFi transmission between two devices A and B, located x distance apart. First A transmits a data packet and then B transmits an acknowledgement (ACK) indicating successful receipt. Distance x can be at most the maximum transmission range, $R_T = R_{6dB}$, since WiFi supports receiving packets at a minimum of 6dB SNR [2].

²The WiFi standard dictates that transmissions with a received power of 87.7 dBm must be detected with a 90% or higher probability. The noise floor in typical WiFi cards is between 88-90 dBm

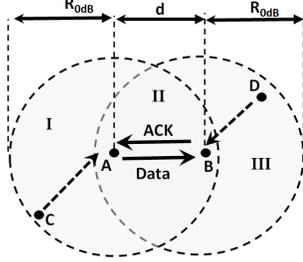


Figure 21: CSMA in WiFi

γ	2	3	4	5
$\eta_{WiFi}(\gamma)[\%]$	17	21	22	24

Table 3: Percentage of hidden nodes in CSMA (WiFi)

Interference range in WiFi. In order for the transmission in Figure 21 to succeed, both the data as well as ACK should succeed. Since any transmission that reaches A or B at 0dB or greater SNR is interference, devices such as C or D in Figure 21 within range R_{0dB} of A and B, respectively, must be silenced. Thus, the total interference region $I(x)$ is simply the area covered by the two intersecting circles of radius R_{0dB} centred around A and B.

Hidden nodes in WiFi. In Figure 21, when A initiates its data transmission, all nodes in the area covered by the circle of radius R_{0dB} centred around A (regions I and II) will detect the transmission and be silenced. However, devices in region III (e.g., D), spanning an area $H(x)$, deemed *hidden nodes* will not be able to detect A's transmission and can cause interference at B if they initiate transmission.

Expected fraction of hidden nodes. The ratio $\eta(x) = \frac{H(x)}{I(x)}$ is then the fraction of hidden nodes. Assuming uniform node distribution, the expected fraction of hidden nodes can be obtained by averaging over all possible locations of B, within a circle of transmission radius R_T around A, and is given by,

$$\eta_{WiFi} = \frac{1}{\pi R_T^2} \int_0^{R_T} \frac{H(x)}{I(x)} 2\pi x dx = \int_0^{R_T} \frac{H(x)}{I(x)} \frac{2x}{R_T^2} dx \quad (7)$$

Table 3 lists η_{WiFi} for various path loss exponents γ . As seen from Table 3, CSMA in WiFi may suffer from

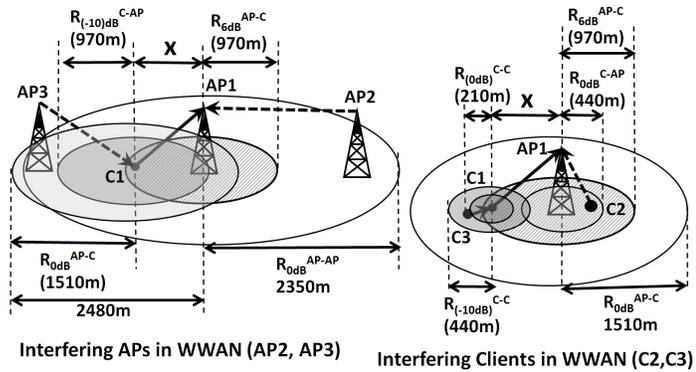


Figure 22: Interference in WWANs

hidden nodes in up to 17-24% of cases.

RTS-CTS While WiFi provides RTS-CTS to deal with hidden nodes, RTS-CTS is seldom used in practice as it results in high overheads for 80+% of the cases when there are no hidden nodes.

5.2 Interference in a WWAN using CSMA

As discussed in Section 2, unlike a WLAN, a WWAN comprises three different channels (AP-AP, AP-Client and Client-Client) and two different kinds of devices (APs and Clients). Consequently, unlike WiFi, there are four different interference ranges in a WWAN: R_{0dB}^{AP-AP} , R_{0dB}^{AP-C} , R_{0dB}^{C-AP} and R_{0dB}^{C-C} . Here $R_{\alpha dB}^{X-Y}$ is the distance at which a transmission from device type X is received at device type Y at α dB SNR; device types being Clients (C) and APs (AP).

Performing CSMA for sub-noise transmissions As discussed earlier, transmissions from clients may need to be detected at -10dB SNR at the AP. In order to enable this, clients must use longer PN-Sequences for doing CSMA instead of the standard WiFi PN-sequences [21]. Specifically, a PN-sequence of length K times longer can be used to detect an ongoing transmission at $-10 \log_{10} K$ dB SNR. Thus, a PN-sequence at least $10 \times$ longer ($K = 10$) will be required by the clients.

Interfering APs in WWANs : All APs located within a radius of R_{0dB}^{AP-AP} (2350m) centred around AP1 (e.g., AP2), can interfere at AP1 (while it is receiving an ACK or data from C1). Similarly, all APs located within a radius of R_{0dB}^{AP-C} (1510m) centred around C1 (e.g., AP3,

channel (Section 7) and the slot must accommodate detection, delay due to propagation and tx-rx switch time [17]. Assuming an average of 8 slots (for a single client), the total air-time efficiency *i.e.*, fraction of time useful data is transmitted is under 15%. While the use of frame aggregation can potentially improve efficiency in some limited scenarios (e.g., bulk download) it does not help for short http sessions or voip traffic. Finally, as the number of clients increase, increased RTS-CTS and packet collisions can further reduce efficiency. Thus, while RTS-CTS can be made to work in a WWAN, it will come at the cost of severely limited cell capacity. *As we discuss in Section 4.4 CDC offers the benefit of higher efficiency and almost collision-free contention.*

6 Implementation Details

To test various aspects of IQ-Hopping and CDC we had to use three different platforms – madwifi, SORA, and USRP since each had its limitations and advantages.

Implementation of IQ hopping on madwifi.

Implementing IQ hopping on madwifi shows that the scheme can be implemented on any commodity WiFi AP with relatively minor modifications. Also this allowed to test IQ hopping in the presence of commodity WiFi APs running various kinds of traffic such as UDP, VoIP, ftp, etc. In all these experiments we used dell machines equipped with netgear cards running MadWiFi0.9.4.

First, we tapped the kernel interrupt generated on packet transmissions. This interrupt returns a descriptor to the driver which contains several parameters, *e.g.*, if an ACK was received, number of retransmissions, data rate, etc. These parameters were used to compute the amount of useful time spent on a channel. We used kernel timers to decrement idle-time and to decide when to hop. After taking the decision to hop, the AP also broadcasts 5 CSA (Channel Switch Announcement) packets (34 bytes each) indicating the newly selected channel before actually changing its channel. When the associated clients hear these CSAs, they hop to the indicated channel. To increase robustness, we modified the driver to ensure that CSAs are sent out on the high priority transmission queue. In the rare case that clients lose all channel switch messages, the client simply initiates re-association.

width of 6MHz to match our testbed hardware constraints

Implementation on SORA. On SORA we were able to try WiFi and IQ hopping together in white space frequencies both at 4W using a power amplifier in our outdoor experiments. With our experimental license expiring at one of the locations since the time of our measurements, we were limited to transmitting at 4W from only one of the two locations – the urban area from a height of about 15m. This SORA node acted as the AP. Two other SORA nodes which acted as clients and transmitted at 100mW. We were able to drive in the city in a car with one SORA client as mobile and one SORA client as fixed for the evaluations. The mobile SORA client helped us test CDC in a real wide area environment subject to vagaries of wireless propagation in WWANs.

Implementation on USRP. We had 8 USRP nodes at our disposal, but only 4 SORA devices. The USRP testbed provided us with an opportunity to scale evaluation of CDC up to 7 devices contending simultaneously with the 8th node as the AP. We used USRP to test how CDC's collision detection scales for a larger number of nodes. Having no suitable power amplifier for USRP, however, we were limited in our transmissions only up to 10mW. Due to real-time constraints for power control in CDC, we implemented CDC entirely on the FPGA on the USRP boards using VHDL.

7 Testbed Results

In this section we present a testbed evaluation of the two components of WiFi-XL namely, IQ-Hopping and CDC.

7.1 IQ-Hopping Evaluation on MadWiFi

In this section we ask the question, *How well does IQ-Hopping coexist with commodity WiFi devices and adapt to changing traffic situations?* In order to answer this question, we conducted two experiments.

MadWiFi Experiment 1. In this experiment we used six nodes. Four of these were used to create two AP-Client pairs running standard WiFi – AP1-C1 and AP2-C2. While the other two formed an AP-Client pair, $AP_{IQ} - C_{IQ}$ that ran IQ hopping. AP1-C1 and AP2-C2 ran saturated UDP traffic. Through the course of the experiment AP1 and AP2 changed their channels. We first

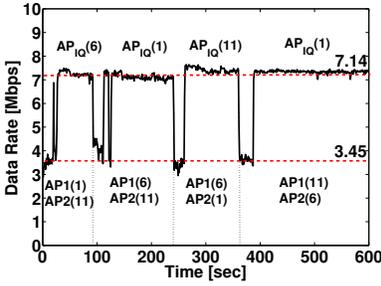


Figure 24: How IQ-Hopping helps APs avoid saturated channels.

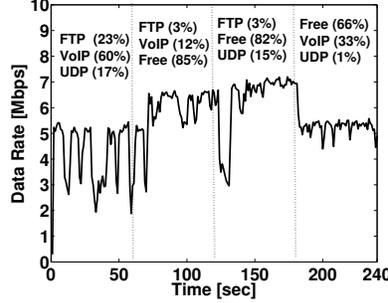


Figure 25: How IQ-Hopping functions across real traffic

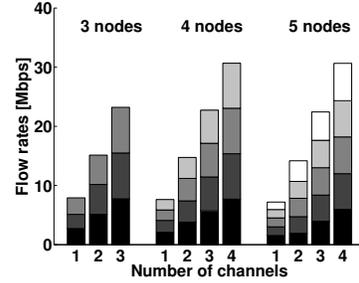


Figure 26: How IQ-Hopping functions across real traffic

ran iperf on regular WiFi in a free channel and an occupied channel to find the baseline achieved throughputs to be 7.14 and 3.45 Mbps respectively (indicated as dotted horizontal lines in Figure 26). The channels they were on are indicated in Figure 26, for example, in the interval 90-240 sec, AP1(6)AP2(11) indicates that AP1 and AP2 were on channels 6 and 11 respectively. As seen from Figure 26 AP_{IQ} correctly moved to the empty channel after each channel change, usually within about 20s. The average throughput of AP_{IQ} in the entire experiment was 6.6Mbps – a loss of about 7% from the maximum possible.

MadWiFi Experiment 2. In this experiment we used eight nodes with three AP-Client pairs serving as background traffic and the last pair running IQ hopping. On channels 1, 6 and 11 we had FTP traffic, VoIP traffic and saturated UDP traffic respectively. Initially all three traffic sources were on but once every 60s one of them was switched off as depicted in Figure 25. In Figure 25 we also provide the fraction of time spent in each of the channels. In 0-60 seconds, IQ hopping spent most time (60%) in VoIP channel and the least in the UDP channel (17%) – this is as desired since VoIP is the least congested background traffic while UDP is the most. When UDP was turned off between 60-120s, IQ hopping spent 85% on the free channel and only 3% on the FTP channel. Similarly, in other sections as well the fractions of time IQ hopping spent in the channels is decreasing order of the amount of congestion they cause.

7.2 IQ-Hopping Evaluation on SORA

In this experiment we placed SORA WiFi APs in a single contention domain, so that they all interfere with each other and varied the number of devices to be 3, 4 and 5 while also varying the available number of channels from 1 to 4, each 5 MHz wide. We choose packet duration to correspond to 1024B packets transmitted with 16-QAM and 3/4 coding, corresponding to 9 Mbps (equivalent to 36 Mbps rate of WiFi). We run each experiment for 5 minutes. The average flow data-rates are shown in Figure 30 (e). Each bar corresponds to one experiment, with a given number of nodes and channels. Each rectangle in a bar denotes a data rate of a single flow, and the height of a bar is the sum of all rates. As we can see, IQ hopping fully utilizes the available bandwidth and achieves almost perfect fairness in all cases.

7.3 Evaluation of CDC

We implemented CDC in two platforms – the SORA software defined radio and USRP. For both these devices we used 5 MHz channels (instead of 6MHz whitespace channels) due to hardware limitations of the platforms.

What are the minimum lengths of the preambles to provide the desired outdoor coverage? We first verify the range of the AP downlink preamble (transmitted at 4W). We set the size of the preamble to the size of the standard WiFi STS preamble, which is 32 μ s for a 5MHz channel (160 samples). While the AP transmits the preamble, we place the other Sora node in a car and we

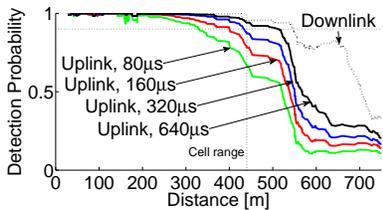


Figure 27: Collision Detection in CDC

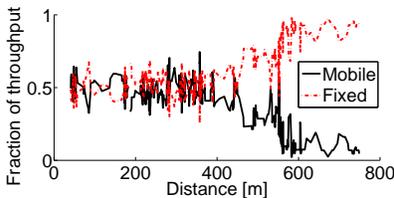


Figure 28: Uplink throughputs in a two-client network.

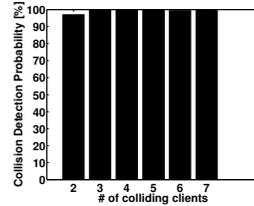


Figure 29: Collision Detection in CDC

drive it around, logging the preamble detection rate and the GPS position. The detection rate as a function of the distance to the AP is shown in Figure 28. We see that the detection range matches well with the expected cell range in Section 2.

Next, we evaluate the required size of the uplink preamble. We consider preambles of $80 \mu s$, $160 \mu s$, $320 \mu s$, $640 \mu s$ (taking 400, 800, 1600 and 3200 samples, respectively). We transmit the uplink preambles from the node in the car at 100 mW and we receive them at the AP. We log the reception rate at the AP and the position of the car. We plot the reception rates as functions of the distance for various preamble sizes. We see that the uplink preamble size of $320 \mu s$ suffices to achieve 90% detection rate throughout the cell range.

Note that due to the complexity of the outdoor experimental setup, this experiment has only two contending nodes. From the previous experiment with multiple contending nodes we conclude that we need to double the preamble size to be able to support multiple contending nodes. Hence we choose the preambles of size $640 \mu s$ for our CDC design.

How efficient is CDC in practice? In this experiment we evaluate the performance of CDC in a wide area setting. We place one client inside the building, 10m from the AP. We place the other client in the car. Both client transmit at a maximum of 100 mW but are power controlled so that they reach the AP at -10dB SNR. We set the length of the uplink preamble to $320 \mu s$ (to make the evaluation fair, since there are only 2 clients in the experiment). We drive the car and record the GPS position of the mobile node. At the AP we record the fraction of uplink contention resolution cycles acquired by the mobile and the fixed client. The results are depicted in Figure 28.

As expected, there is some variation in the actual frac-

tions due to varying channel conditions of the mobile client. When the mobile client is within the cell range, both mobile and fixed client get approximately equal share of the medium. Note that the mobile client occasionally gets a larger share than the fixed client. This is due to the imperfection of our CDC power control implementation. However, the oscillations are small, showing that the power control is reasonably accurate even when the mobile node is a few hundreds of meters away from the AP.

How efficient is CDC collision detection scheme? Next we evaluate the capability of CDC to detect code contention collisions due to multiple nodes choosing the same codeword to contend with. In this experiment we simply force multiple clients to transmit the same PN-sequence instead of choosing randomly and count the number of collisions instead. In order for the AP to infer that there has been no collision, it must observe only a single correlation peak even though multiple clients used the same code. Figure 29 depicts the collision detection probability as more and more clients choose the same PN-sequence. As seen from Figure 29 collision detection probability is close to 100% and improves as more clients choose the same PN-sequence.

8 Simulation Study

In this section, we evaluate the performance of IQ hopping on networks larger than our test-beds. In particular, we want to understand how IQ hopping compares to other variable channel-width channel-selection schemes, MCham (the greedy measurement-based scheme in WhiteFi [6]) and Random hopping by implementing them on Qualnet. For our evaluation, we used the channel availability information [1] from two locations, Austin (with 8

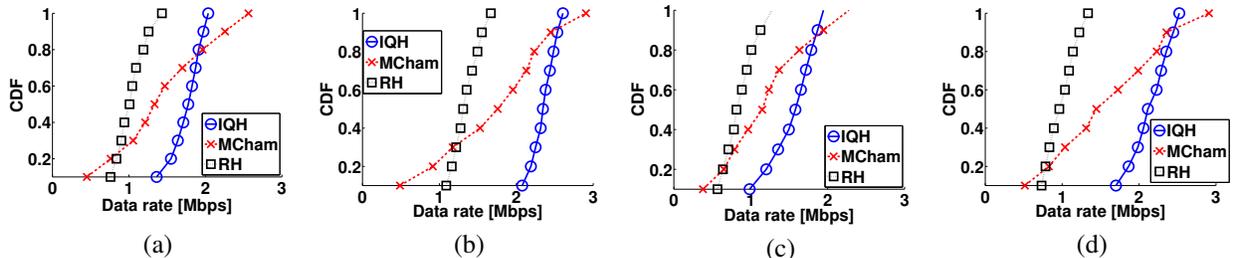


Figure 30: (a) IQ Austin, degree 3, (b) IQ Ithaca, degree 3, (c) IQ Austin, degree 6, (d) IQ Ithaca, degree 6.

available TV channels) and Ithaca (with 16 available TV channels). We then evaluate the performance of a large network with 20 APs and 20 clients. We parameterize network topologies by placing AP at random locations such that the average AP interference degree is d . We place one client per AP at the edge of the coverage range for fair comparison across schemes, ensuring that the down-link data rate is always 1.5 Mbps (corresponds to 6 Mbps WiFi rate). Each AP has saturated traffic for its client and we use the propagation model parameters of the urban scenario (Section 2).

We simulate 10 random topologies for each combination of channel availability (Austin and Ithaca) and evaluate over interference degree of 3 and 6 respectively. For each topology we find the 10%, 20%, \dots , 100% percentile of the data rates distribution. We average values for each percentile over 10 topologies, and we plot the average of the percentile values in Figure 30 (a)-(d).

We now make the following observations from the figures. First, we see that IQ hopping significantly outperforms both MCham and random hopping. Specifically, on an average it provides 20%, 33%, 27% and 34% gains in total throughput over MCham and 67%, 73%, 77% and 69% over random hopping in scenarios depicted in Figure 30 (a)-(d), respectively. Second, compared with MCham, IQ hopping increases the throughputs of the lowest 10% of the flows by upto 4 times. Third, both IQ and random hopping achieve Jain’s fairness index above 0.97 in all cases, whereas the fairness index of MCham varies from 0.82 - 0.87. To summarize, IQ hopping achieves significant gains both in terms of achievable throughput and fairness over MCham.

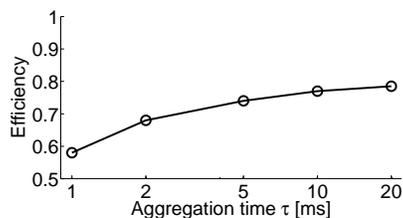


Figure 31: WiFi-XL efficiency.

8.1 Efficiency of WiFi-XL

Finally, we want to evaluate the efficiency of WiFi-XL. To analyse this, we simulate an AP with 30 clients. Each client is distributed randomly across the coverage region of the AP and we use the data rate that can be supported for the client’s SNR. We use an HTTP traffic generator, which generates realistic web traffic between clients and its AP, based on historical data. For this experiment, we implement the entire WiFi-XL protocol. We measure the efficiency as the fraction of total time that was used to send actual data packets. We vary the length of the aggregation time τ from 1 ms to 20 ms and we plot the efficiency for each τ . The results are shown in Figure 31 (f). We see that the efficiency of WiFi-XL is comparable to WiFi, achieving close to 70% for $\tau = 2$ ms and approaching 80% for $\tau = 20$ ms,

9 Related Work

Mesaurement study. There have been several measurement studies of white space availability in different parts of the world [25, 7] and they find between 50-60% of TV channels unused on average. But given restrictions on transmission and channel fragmentation, usable channels

vary geographically. For example, estimates for southern Rhineland in Germany reveals availability of *only four 24 MHz channels in urban or four 16 MHz channels in a rural settings* [23]. Similarly, the spectrum database for U.S. [1] show wide range of channel availability (0-28), with fewer available channels in urban centers. Thus, a WiFi-XL like mechanism for sharing channels in both time and frequency is critical.

Standards. IEEE 802.11af is currently a draft standard for white space networks [2]. It takes the new 802.11ac gigabit WiFi specification and simply down-clocks it to support white space TV channel bandwidths of 6/8 MHz.

Systems. WhiteFi [6] is a WiFi-like system for TV white spaces, similar to 802.11af. Thus, WhiteFi will face the same issues as 802.11af, namely, lack of wide-coverage. SenseLess [19] presents the design of a white space database that uses an accurate propagation model and terrain data to identify white space availability at various geographic coordinates. Weeble [21] is a system that allows networks with two different power levels to coexist, e.g., bluetooth and WiFi. WiFi-XL can use Weeble to enable a mixed network deployment with 100mW APs and 4W APs.

Contention. Zigzag decoding [13] is a technique to combat hidden terminals by recovering packet collisions through interference cancellation. However, given sub-noise communication and the high fraction of hidden nodes, applying zigzag to WWAN settings will be very challenging.

Frequency Selection. A large amount of prior work [8] has focused on solving the automatic frequency selection problem based on measurements and most assume fixed channel widths. WhiteFi [6] uses measurement and the MCham metric for channel selection with variable channel widths. However, as we show in this paper, this can lead to inefficiency and unfairness.

10 Conclusion

We show that two key characteristics of a WWAN network in white spaces, namely, channel trichotomy and sub-noise communication support, make the network design challenging, and simple adaptation of existing protocols such as WiFi to unlicensed WWANs does not work well. Therefore, we design and build WiFi-XL, based on

two novel techniques, idle-quantum hopping and code domain contention, and show using testbed experiments and simulations that WiFi-XL is able to fairly and efficiently run a WWAN in unlicensed spectrum.

References

- [1] Google White Space Spectrum Database. <https://www.google.com/get/spectrumdatabase/channel/>.
- [2] Ieee 802.11af, project status. http://www.ieee802.org/11/Reports/tgaf_update.htm.
- [3] Municipal Wi-Fi: A proven failure. <http://www.theamericanconsumer.org/2011/05/20/municipal-wi-fi-a-proven-failure/>.
- [4] Range Networks OpenCell Series. <http://www.rangenetworks.com/products/open-cell-series>.
- [5] Where's All the Free Wi-Fi We Were Promised?, 2010. <http://www.technologyreview.com/view/421084/wheres-all-the-free-wi-fi-we-were-promised/>.
- [6] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh. White Space Networking with Wi-Fi like Connectivity. In *ACM SIGCOMM*, August 2009.
- [7] D. Chen, S. Yin, Q. Zhang, M. Liu, and S. Li. Mining Spectrum Usage Data: a Large-scale Spectrum Measurement Study. In *ACM MOBICOM*, September 2009.
- [8] S. Chiochan, E. Hossain, and J. Diamond. Channel Assignment Schemes for Infrastructure-Based 802.11 WLANs: A Survey. *IEEE Communications Surveys and Tutorials*, 12(1):124–136, 2010.
- [9] J. Egli. Radio propagation above 40mc over irregular terrain. *Proceedings of the IRE (IEEE)*, 45(10):1383–1391, 1957.
- [10] FCC. Fcc adopts rules for unlicensed use of television white spaces, November 2008.
- [11] FCC. Unlicensed operation in the tv broadcast bands: Second memorandum opinion and order, September 2010.
- [12] E. Fraser. A postmortem look at Citywide WiFi. *Journal of Internet Law*, 14(2):6–11, August 2010.
- [13] S. Gollakota and D. Katabi. Zigzag decoding: combating hidden terminals in wireless networks. In *ACM SIGCOMM*, 2008.
- [14] M. Hata. Empirical formula for propagation loss in land mobile radio services. *IEEE Transactions on Vehicular Technology*, August 1980.

- [15] H. Holma and A. Toskala. *LTE for UMTS: Evolution to LTE-Advanced, Second Edition*. 2011.
- [16] LYRETECH. <http://www.lyretech.com/>.
- [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] R. Murty, R. Chandra, T. Moscibroda, and V. Bahl. Senseless: A database-driven white spaces network. *IEEE Transactions on Mobile Computing*, 2012.
- [20] R. Patra, S. Nedeveschi, S. Surana, A. Sheth, L. Subramanian, and E. Brewer. WiLDNet: Design and Implementation of High-Performance Wifi-based Long Distance Networks. In *NSDI*, 2007.
- [21] B. Radunovic, R. Chandra, and D. Gunawardena. Weeble: Enabling Low Power Nodes to Coexist with High Power Nodes in White Space Networks. In *ACM CONEXT*, 2012.
- [22] S. Sen, R. R. Choudhury, and S. Nelakuditi. No time to countdown: migrating backoff to the frequency domain. In *ACM MobiCom*, 2011.
- [23] L. Simic, M. Petrova, and P. Mahonen. Wi-Fi, but not on Steroids: Performance Analysis of a Wi-Fi-like Network Operating in TVWS under Realistic Conditions. In *ICC*, 2012.
- [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] J. van de Beek, J. Riihijarvi, A. Achtzehn, and P. Mahonen. UHF white space in Europe - a quantitative study into the potential of the 470-790MHz band. In *DySPAN*, 2011.