MIT Libraries | DSpace@MIT

MIT Open Access Articles

Frequency-aware rate adaptation and MAC protocol

Massachusetts Institute of Technology

Frequency-Aware Rate Adaptation and MAC Protocols

Hariharan Rahul[†], Farinaz Edalat[ℵ], Dina Katabi[†], and Charles Sodini[†]
[†]Massachusetts Institute of Technology   [ℵ]RKF Engineering Solutions, LLC

ABSTRACT

There has been burgeoning interest in wireless technologies that can use wider frequency spectrum. Technology advances, such as 802.11n and ultra-wideband (UWB), are pushing toward wider frequency bands. The analog-to-digital TV transition has made 100-250 MHz of digital whitespace bandwidth available for unlicensed access. Also, recent work on WiFi networks has advocated discarding the notion of channelization and allowing all nodes to access the wide 802.11 spectrum in order to improve load balancing. This shift towards wider bands presents an opportunity to exploit frequency diversity. Specifically, frequencies that are far from each other in the spectrum have significantly different SNRs, and good frequencies differ across sender-receiver pairs.

This paper presents FARA, a combined frequency-aware rate adaptation and MAC protocol. FARA makes three departures from conventional wireless network design: First, it presents a scheme to robustly compute per-frequency SNRs using normal data transmissions. Second, instead of using one bit rate per link, it enables a sender to adapt the bitrate independently across frequencies based on these per-frequency SNRs. Third, in contrast to traditional frequency-oblivious MAC protocols, it introduces a MAC protocol that allocates to a sender-receiver pair the frequencies that work best for that pair. We have implemented FARA in FPGA on a wideband 802.11-compatible radio platform. Our experiments reveal that FARA provides a $3.1\times$ throughput improvement in comparison to frequency-oblivious systems that occupy the same spectrum.

Categories and Subject Descriptors C.2.2 [**Computer Systems Organization**]: Computer-Communications Networks

General Terms Algorithms, Design, Performance

Keywords Wireless, Cognitive Radios, Wideband, Rate Adaptation, Cross-layer

1   INTRODUCTION

Wireless technologies are pushing toward wider frequency bands than the 20 MHz channels employed by existing 802.11 networks. 802.11n already includes a 40 MHz mode that bonds together two 20 MHz bands [23]. Emerging ultra-wideband (UWB) technologies employ hundreds of MHz to support multimedia homes and offices [24, 50, 9, 40]. The FCC has recently permitted unlicensed

Figure 1: Frequency diversity across 100 MHz of 802.11a spectrum as observed by two receivers for transmissions from the same sender. The figure shows that the SNRs of different frequencies can differ by as much as 20 dB on a single link. Further, different receivers prefer different frequencies.

use of digital TV whitespaces that occupy 100-250 MHz of spectrum vacated by television bands in the analog-to-digital transition [12]. Furthermore, recent empirical studies show that the 802.11 channelization model which limits each node to a single 20 MHz channel can lead to severe load imbalance [19, 28, 37]. They advocate discarding channelization and allowing all nodes to access the entire 802.11 spectrum based on demand [19, 37]. This push towards wider bands is further enabled by the constantly lowering prices of high-speed ADC and DAC hardware [38, 31].[1] In particular, today, wireless cards that span over 100 MHz of spectrum can be built using off-the-shelf hardware components [35].

As wireless networks push towards wider bands, we can no longer afford to ignore frequency diversity. Specifically, multipath effects cause frequencies that are far away from each other in the spectrum to experience independent fading. Thus, different frequencies can exhibit very different SNRs for a single sender-receiver pair. Further, the frequencies that show good performance for one sender-receiver pair may be very different than the frequencies that show good performance for another pair. Fig. 1 shows empirical measurements of the SNRs across 100 MHz of the 802.11a spectrum, as observed by 2 clients for transmissions from the same AP (see §9 for experimental setup). The figure reveals that different frequencies show a difference in SNR of over 20 dB both for a single link and across links. Existing bitrate adaptation and MAC protocols however are frequency-oblivious. They assign the same bitrate to all frequencies and allocate the medium in a time-based manner, ignoring the fact that different frequencies work better for different sender-receiver pairs. Thus, current rate adaptation and MAC protocols can neither deal with the challenge nor exploit the opportunities introduced by the frequency diversity of wide bands or unchannelized 802.11.

---

[1]The wider the band, the faster the ADC and DAC have to sample the signal.

This paper presents the design and implementation of FARA, a frequency-aware wireless architecture. FARA is beneficial for both wireless LANs and mesh networks. Its design focuses on the 802.11a/b/g/n spectrum, but it can also be used in a cognitive mode over white spaces as discussed in §10.

FARA has four key components that together allow it to improve network throughput and balance load.

- *Per-frequency SNR estimation:* FARA leverages the existing OFDM system, which divides the entire frequency band into many narrow subbands. It devises a new approach to allow a receiver to use normal data packets, whether received correctly or incorrectly, to robustly estimate the SNR in each OFDM subband.
- *Frequency-aware rate adaptation protocol:* FARA uses its per-frequency SNR measurements to enable a transmitter to use different bitrates across different OFDM subbands. Specifically, a FARA receiver measures the SNR in each subband, maps it into an optimal bitrate using characterization tables for the receiver hardware,[2] and periodically reports this optimal bitrate for each subband to the transmitter.
- *Frequency-aware MAC:* A FARA transmitter acquires the medium using carrier-sense. However, once the medium is acquired, a transmitter that has traffic for multiple receivers can simultaneously transmit to all these receivers by preferentially allocating frequencies to receivers to maximize the overall throughput across these receivers.
- *Load-aware contention:* In contrast to existing channelized 802.11 networks, in FARA the entire frequency spectrum is available to all nodes without channelization. As a result, load balancing can be done using a small modification to CSMA where an AP or a router in a mesh network contends for the medium proportionally to its load.

We implemented FARA using the WiGLAN radio platform [35], and compared it to the current frequency-oblivious 802.11. Measurements from our indoor testbed reveal the following findings:

- Frequency diversity exists, and is stable over time. Specifically, our results show that different subbands in the 802.11a spectrum can show a difference in SNR of over 20 dB. Further, the SNR profile for individual subbands is relatively stable for periods upto 5 seconds, and hence can be communicated from receiver to transmitter with low overhead.
- FARA is effective at harnessing frequency diversity, delivering a median throughput gain of $3.1\times$ in our testbed.
- FARA's gains come both from exploiting frequency diversity within a single sender-receiver pair (frequency-aware rate adaptation), as well as across sender-receiver pairs (frequency-aware MAC). Typically, for our experimental scenarios, about 70% of the gains are due to frequency-aware rate adaptation, and 30% are due to the frequency-aware MAC.
- FARA's load-aware contention protocol is fair even when APs have a wide load disparity.

To the best of our knowledge, FARA is the first system to present frequency-aware rate adaptation and MAC protocols, and show through a prototype implementation and experimental evaluation, that frequency-awareness can improve the throughput of an 802.11 network.

---

[2]These characterization tables need to be calibrated only once for a particular receiver hardware.

2 RELATED WORK

Related work falls in the following areas.

(a) Measurement and Analysis of Frequency Diversity. The impact of multipath effects in creating varying signal strength across frequencies is well understood theoretically [47]. Also, multiple measurement studies [36, 43, 8, 29, 4] have demonstrated the existence of this frequency diversity in practice, showing that it occurs both at the low end of the RF spectrum, as in white spaces [3, 39], as well as the high end of the spectrum, as in 60 GHz technologies [36, 43]. Among these, the most relevant to our work are measurement studies in the 2.4 and 5.2 GHz spectrum [8, 29, 4] (corresponding to 802.11b/g and 802.11a, respectively), which report a difference in signal strength of as much as 20 dB between frequency bands both in line-of-sight (LOS) and non line-of-sight (NLOS) scenarios. Our work is motivated by these results, but differs from them significantly because it presents a rate-adaptation and MAC that can leverage frequency diversity to improve network throughput.

(b) Subband Adaptive Modulation and Coding. Wired cable modem standards, such as Asymmetric Digital Subscriber Line (ADSL), High bit-rate DSL (HDSL), and Very high-speed DSL (VDSL) [44] adapt modulation and coding independently across OFDM subbands. Some prior theoretical or simulation studies have proposed a similar approach for wireless channels [27, 13, 7, 22, 30, 54, 32, 33, 34, 11], and a few papers [53, 48] have investigated the complexity of hardware implementations of such designs. Our work builds on this foundation but differs in two ways. First, FARA presents a rate adaptation algorithm that works in real-time and supports its design with empirical evaluation. Second, FARA augments its frequency-aware rate adaptation with a frequency-aware MAC, while all these prior studies focus only on a single link, and none of them exploits frequency diversity across multiple users.

(c) Opportunistic Communication Schemes. Prior work on opportunistic communication considers a scenario where two nodes that can hop between frequencies try and identify the best channel on which to communicate [42, 18]. This work proposes schemes to minimize exploration overhead while maximizing the probability of finding a high-performance frequency band. While FARA is similar in that it exploits frequency diversity, it differs from these schemes in both objective and mechanisms. First, these schemes focus on finding and using a small set of good frequencies assuming that they all have similar SNRs, while FARA allows a sender-receiver pair to operate over a wide set of frequencies that may differ drastically in SNR. To achieve this goal, FARA provides a rate adaptation scheme that uses different bitrates on frequencies with different SNRs. It also extends the 802.11 protocol to allow a transmitter to transmit simultaneously to multiple receivers, taking advantage of frequency diversity across them. Furthermore, FARA is implemented and evaluated in a wireless testbed, while prior work is simulation-based.

(d) Non-channelized 802.11 Protocols. Recent work has advocated using the 802.11 spectrum as a whole, and discarding the traditional fixed-width channel model [19, 37]. Specifically, ODS [19] allows all nodes to simultaneously access the entire 802.11 spectrum, and share it using code division multiple access (CDMA), while Moscibroda *et. al.* [37] dynamically assign non-overlapping frequencies to different APs proportional to their load. Similarly to these schemes, FARA allows all nodes access to the entire 802.11 spectrum based on their demands, and hence can provide load balancing, but, in contrast to the frequency-obliviousness of the prior work, FARA can exploit frequency diversity both for rate adaptation, and medium access, hence providing additional gains even when loads are balanced.

Figure 2: Multipath Effect Causes Frequency Diversity. Signals from different paths combine at the receiver constructively or destructively depending on their phases. Since the phase is a linear function of the frequency, the destructive and constructive patterns differ across frequency bands causing different frequencies to have different SNRs.

(e) FDMA Cellular Networks. Multiple cellular technologies such as Flash OFDM, GSM and WiMax [47, 2, 41] use frequency division multiplexing for medium access. Flash OFDM and GSM use pseudorandom frequency hopping, rather than assigning to each user those frequencies that work best for the user at that instant. WiMax on the other hand has two modes. For mobile and fast-changing channels, WiMax randomly assigns frequencies to users, rather than assigning the best instantaneous frequencies to each user. For static or moderately dynamic environments, WiMax allocates to each user a chunk of contiguous frequencies that work best for that user. FARA is designed for static or moderately dynamic networks and is similar to this latter mode of WiMax in that it assigns to each sender-receiver pair the frequencies that work best for that pair. FARA however differs from WiMax in two ways: First it does not limit a user to a contiguous chunk of frequencies, and allows a sender-receiver pair to use non-contiguous frequencies. Second, in contrast to WiMax, which uses the same bit rate across the whole chunk allocated to a user, FARA adapts the bit rate independently across different frequencies used by a sender-receiver pair.

3 FREQUENCY DIVERSITY

Frequency diversity is an intrinsic characteristic of RF propagation in multipath environments [47]. The wireless channel both attenuates the RF signal and changes its phase. Specifically, the channel shifts the signal's phase by $2\pi f \tau$, where f is the signal's frequency and τ is the path delay. In environments with multipath effects, the receiver ends up with multiple copies of the signal that traversed different paths with different delays, as shown in Fig. 2. These copies have different phases and hence may add up constructively or destructively. Since the phase of the received signal is a linear function of its frequency, different frequencies show different degrees of constructive and destructive signal patterns.

The effect of this frequency diversity is significant when examined across a wide spectrum, such as the entire 300 MHz of spectrum usable by 802.11a, or the 80 MHz usable by 802.11g. Past measurements show that different frequency bands within the wide 802.11a/b/g spectrum can differ by as much as 20 dB of SNR [4, 8, 29]. These results align with our own measurements shown in Fig. 1. The figure shows the SNRs of two 100-MHz channels in the range of 802.11a, for subbands of one MHz wide. The measurements are taken for two links in our testbed (from transmitter tx to receivers A2 and B3 in Fig. 6). The figure reveals that the SNR difference between frequencies is significant. Furthermore the SNR pattern is highly diverse both for a single link and across the two links.

Frequency diversity motivates bit rate adaptation schemes and MAC protocols that can leverage SNR differences across frequencies to increase network throughput.

4 FARA

FARA is a new architecture for static and moderately dynamic wireless networks, i.e., typical 802.11 environments. Similarly to recent proposals for channel bonding [23] and load balancing [19, 37], FARA advocates discarding the current channel notion and allowing all nodes to access a larger chunk of the 802.11 spectrum. FARA however recognizes that a wider spectrum increases frequency diversity. Its design harnesses frequency diversity via four components: per-frequency SNR estimation algorithm, frequency-aware rate adaptation, frequency-aware MAC protocol, and load-aware contention. Together these components significantly increase network throughput and balance the utilization of the 802.11 spectrum. In the following sections, we explain each of these components in detail.

5 PER-SUBBAND SNR ESTIMATION

FARA introduces a novel algorithm that allows a sender-receiver pair to estimate the performance of each frequency, i.e., its SNR, using normal data packets, whether received correctly or incorrectly. To do so, FARA leverages Orthogonal Frequency Division Multiplexing (OFDM), which is already implemented as part of the 802.11a/g/n physical layer [45, 46]. OFDM divides the used frequency spectrum into many narrow subbands. A subset of these subbands are called pilots and used to transmit a known bit pattern modulated at BPSK to allow the receiver to track the channel [20]. The other subbands are used for data transmission. A FARA receiver estimates the SNR for each OFDM data subband, for each sender, by estimating the signal power from all of the transmitted data, and leveraging the known bit pattern in the pilot bins to estimate noise.

In particular, the SNR in subband i, SNR_i is the ratio of the signal power in subband i, S_i, to the noise power, N_i. The receiver cannot directly measure the signal power; however, it can measure the received power in subband i, R_i, which is the sum of the signal power, S_i and the noise power, N_i. Thus,

$$SNR_i = \frac{S_i}{N_i} = \frac{R_i - N_i}{N_i}$$

$$SNR_i = \frac{R_i}{N_i} - 1 \qquad (1)$$

Note that the noise in a communication channel is typically the same for all subbands, i.e., white noise.[3] This is because noise comes from thermal noise in the receiver hardware, quantization, and digital computation errors, which are all independent of frequency. Thus, we can rewrite Eq. 1 as

$$SNR_i = \frac{R_i}{N_0} - 1, \text{where } N_0 = N_i, \forall i \qquad (2)$$

The received power, R_i, in a particular subband can be easily estimated by taking the square of the signal corresponding to that subband, and averaging this value across all data symbols in a packet.

We can get an accurate estimate of the noise power, N_0, by exploiting the fact that OFDM uses some subbands as pilots, which contain known data bits. Specifically, the received signal sample, $y_i[k]$, in subband i can be written as:

$$y_i[k] = H_i x_i[k] + n_i[k] \qquad (3)$$

---

[3]While channel noise is typically white, interference due to other technologies, say Zigbee, can differ across subbands. In this paper, we deal with interference the same way 802.11 does, i.e., via carrier sense. FARA however can leverage our previous work on SWIFT [40] to identify subbands occupied by other technologies and avoid them.

where H_i is the channel, $x_i[k]$ is the k^{th} transmitted signal sample in subband i, and $n_i[k]$ is the corresponding noise sample. The receiver knows H_i for all subbands because it is estimated using known OFDM symbols in the preamble [20]. In the case of a pilot subband, $x_i[k]$ is also known at the receiver since pilot subbands contain a known data sequence. As a result, the receiver can estimate the noise samples, $n_i[k]$, and the noise power, N_0, as:

$$n_i[k] = y_i[k] - H_i x_i[k] \qquad (4)$$
$$N_0 = E_{i,k}(n_i[k]^2) \qquad (5)$$

where the function $E(.)$ is the mean computed using all pilot bits across all symbols in the data packet.

Thus, every received packet allows the receiver to obtain a new SNR measurement for each OFDM subband. The receiver maintains a time weighted moving average of the SNR in each subband, which it updates on the reception of a data packet.

A few points are worth noting:

(a) What happens when the data packet is corrupted (i.e. does not pass the checksum test)? Even when the packet is corrupted, the receiver can still compute an accurate estimate of the per-subband SNRs. This is because the receiver can compute the average received power, regardless of whether the packet is corrupted or not. Furthermore, the receiver can still obtain an accurate estimate of the noise power since this only requires the pilots which are known, and sent at BPSK, which is the most robust modulation rate and hence allow synchronization and packet recovery even at low SNRs. Thus, FARA can get accurate estimates of the per-subband SNRs from every captured packet, including corrupted packets.

(b) How accurate are FARA's SNR estimates? We note that since FARA has access to the PHY layer, it can collect accurate SNR estimates. In particular, traditional estimates of the SNR use RSSI readings, which measure the received power of a few samples at the beginning of the packet (i.e., the AGC gain) [6], or infer the SNR using just the correlation of header symbols in the preamble of the packet [49]. In contrast, FARA exploits the known pilot bits to accurately estimate the noise power and utilize it in its SNR computation. Furthermore, FARA computes its signal and noise estimates over the whole packet and not just a few samples at the beginning of the packet, which allows it to obtain more stable estimates.

(c) Do different choices of bitrate affect the accuracy of FARA's SNR estimation? OFDM data subbands use a different modulation scheme depending on the choice of bitrate. The modulation scheme in a subband, however, does not affect our per-subband SNR estimate. The estimation of SNR involves only the measured power in each subband and hence can be performed on any packet independent of the modulation and coding schemes used by the transmitter.

6 FREQUENCY-AWARE RATE ADAPTATION

The goal of rate adaptation is to determine the highest bitrate that a channel can sustain at any point in time. Traditional 802.11 rate adaptation schemes are frequency-oblivious, and use the same modulation scheme and coding rate across all frequencies. Thus, they cannot exploit the frequency diversity present across the 802.11 spectrum. In contrast, FARA exploits this frequency diversity via a frequency-aware rate adaptation scheme that picks different bitrates for different frequencies depending on their SNRs.

6.1 PHY Architecture

In 802.11, a particular bit rate implies a single modulation scheme and code rate over all OFDM subbands in the entire packet. For
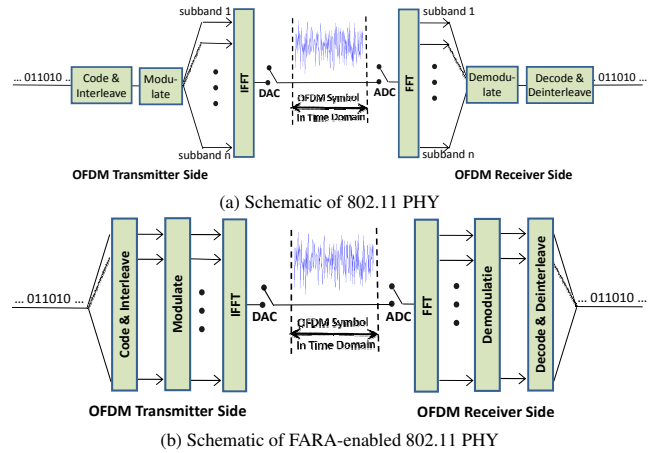

(a) Schematic of 802.11 PHY

(b) Schematic of FARA-enabled 802.11 PHY

Figure 3: OFDM PHY semantics with and without FARA. In FARA-enabled devices, the choice of modulation and FEC code rate is done independently for each OFDM subband.

Minimum Required SNR	Modulation	Coding
<3.5 dB	Suppress subband	
3.5 dB	BPSK	1/2
5.0 dB	BPSK	3/4
5.5 dB	4-QAM	1/2
8.5 dB	4-QAM	3/4
12.0 dB	16-QAM	1/2
15.5 dB	16-QAM	3/4
20.0 dB	64-QAM	2/3
21.0 dB	64-QAM	3/4

Table 1: Minimum required SNR for a particular modulation and code rate (i.e., bitrate). Table is generated offline using the WiGLAN radio platform by running all possible bit rates for the whole operational SNR range. The SNR field refers to the minimum SNR required to maintain the packet loss rate below 1% (see §9 for experimental setup).

example, a bitrate of 24 Mbps corresponds to 16-QAM modulation scheme and a half-rate code. 802.11 has 4 possible modulation schemes (BPSK, 4-QAM, 16-QAM, and 64-QAM), and 3 possible code rates (1/2, 2/3, and 3/4). In current 802.11, a transmitter implements a particular bitrate by first taking the input bit stream, passing it to the convolutional coder, and puncturing to achieve the desired coding rate. The bits are then interleaved, modulated and striped over the OFDM subbands, as shown in Fig. 3(a). The process is reversed on the receiver as shown in the figure.

FARA makes a few modifications to the existing 802.11 PHY layer, as shown in Fig. 3(b). Specifically, FARA employs the same set of modulation schemes and code rates supported by the existing 802.11. However, it allows each OFDM subband to pick a modulation scheme and a code rate that match its SNR, independently from the other subbands. Note that this design does not require additional modulation/demodulation or coding/decoding modules in the PHY layer. In particular, since we use standard 802.11 modulation and coding options, we only need to buffer the samples and process them through the same pipeline.

6.2 Mapping Subband SNRs to Optimal Bitrates

The receiver needs to map the average SNR in each subband to the optimal bitrate for that band. To do so, the receiver uses an SNR characterization table like the one in Table 1 that lists the minimum SNR required for a particular combination of modulation and cod-

ing rate, i.e., a particular bitrate. For each subband, the receiver picks the highest bitrate that can be sustained by the SNR of that subband. Subbands which have SNR too low to support even the lowest bitrate in Table 1 are not used *i.e.*, they do not have any power or data assigned to them, as such a decision will improve the overall throughput. Said differently, subband suppression is simply a special case of FARA's ability to use different modulation and coding rates for subbands based on their SNR.

Many hardware manufacturers already perform this calibration and can provide it as part of hardware specification sheets [10]. Even when the manufacturer does not provide the SNR characterization table, it can be computed using brute force by varying the transmission power and bitrates, and measuring the observed throughput and SNR [26]. We show in the results section that the table does not change with location or time, and thus the measurements can be done only once for each receiver.

6.3 Rate Adaptation Protocol

FARA's rate adaptation is receiver driven: a FARA receiver computes the optimal choice of bitrate on each subband, and feeds it back to the sender in ack packets. Specifically, FARA extends the 802.11 synchronous ack format with a field for bitrate feedback.

When a sender first initiates communication with a receiver, it makes a conservative choice and uses the lowest bitrate on all subbands. The receiver uses this to obtain its first estimate of the SNRs, and hence, the bitrate, in each subband. In order to allow the sender to quickly jump to the correct bitrate, the receiver then sends the appropriate bitrate for each subband immediately in the ack response.

After this initialization, receiver feedback is sent in 802.11 synchronous acks, which we augment with a feedback field. FARA reduces the feedback overhead by exploiting the fact that bitrates typically do not change from one packet to the next, and even when they do, are likely only to change to neighboring bitrates on either side (*i.e.,* jump up or down by one bitrate) As a result, the subsequent ack packets only need to use a 2-bit field per subband to represent one of three choices: stay at the current bitrate, move up to the next highest bitrate, move down to the next lower bitrate. Further, since most of these field values are likely to represent staying at the current bitrate, the feedback information can be compressed drastically using run-length encoding, which is easy to implement in hardware.

We note that sending variable-length synchronous acks does not affect 802.11 behavior. An 802.11 network only requires acks to start within a fixed (SIFS) interval; variable-length acks will not affect network function as the next packet transmission will not start until the medium becomes idle.

What happens when packets or acks get lost? Since FARA's acks are incremental, loss of these acks could lead a sender and receiver to go out of sync. To address this problem, the receiver includes a sequence number with each ack, and stores the bitrate state that the sender would compute as a result of receiving that ack. The sender, in turn, includes the sequence number of the last ack that it has received in its data packets. Thus, when the receiver gets a data packet, it can look up the included ack sequence number in its stored state, and thereby infer the sender's bitrate state. It can then compute the incremental feedback to be included in the new ack with reference to that state. Note that, since FARA's acks are synchronous, the receiver only needs to store the state corresponding to the most recent ack sequence number received from the sender. It therefore needs only a small amount of state to track the sender's bitrate. The receiver also includes the original ack sequence number with respect to which the increment is computed, so that the sender can update its state correctly upon reception of the new ack.


Figure 4: Simplified FARA frame with three concurrent packets. A FARA sender transmits concurrently to multiple receivers, by allocating to each receiver a subset of the OFDM subbands. To ensure fairness, *i.e.*, to ensure that all clients obtain the same average rate, distant receivers that experience lower per-frequency SNRs are allocated more frequencies. Note that the frequencies used for one receiver are not necessarily consecutive. They are made consecutive in the figure to simplify the drawing. The frame header includes metadata about the intended nexthops as well as their allocated subbands.

7 FREQUENCY-AWARE MAC

Similar to 802.11, a FARA sender uses carrier-sense to access the medium. However, different from 802.11, when it accesses the medium, it transmits simultaneously to multiple nexthops, assigning each of them a non-overlapping set of OFDM subbands. The choice of concurrent nexthops, as well as the set of subbands assigned to them, aims to maximize throughput.

Determining the optimal assignment of subbands to concurrent receivers is a difficult problem. To see why, say that a FARA sender wants to deliver 3 packets to 3 nexthops in a single transmission. Fig. 4 shows the transmitted frame across time and frequency. As can be seen, each frame contains multiple concurrent packets intended for different nexthops. The rate for a particular nexthop is the sum of the rates of all the subbands assigned to that nexthop. The transmission time for a packet to that nexthop is therefore the packet size divided by the rate to that nexthop. Since all 3 packets in the figure are being transmitted concurrently, we would like to minimize wastage of medium time by equalizing the transmission time of the 3 packets. For equal packet sizes, this implies that the total rate assigned to the different nexthops are equal. If packet sizes are not equal, the rates need to be proportional to the packet size. This problem is NP-hard as can be demonstrated by a trivial reduction to the bin-packing problem [15]. Hence we seek a heuristic solution.

Assigning subbands to concurrent receivers: FARA's MAC protocol works as follows. A sender is configured with a maximum allowed number, N, of concurrent packets in a transmission. In practice, N is a small number between 2-5. FARA maintains pernexthop packet queues, as well as a global FIFO transmission queue which contains pointers to packets in the per-nexthop queues. For each transmitted frame, the sender picks upto N concurrent packets. It first picks the packet at the head of the global FIFO transmission queue, and determines the associated nexthop. It then randomly chooses upto $N-1$ other nexthops with non-empty queues, and picks the packets at the head of these queues. These packets will be transmitted concurrently in one frame. The random choice of nexthops ensures that FARA is fair to all nexthops while providing significant throughput gains, as we show in §9.4.

The FARA sender now needs to assign subbands to each packet as to equalize the transmission rate to all N receivers. The sender also wants to assign to each receiver its preferred frequencies, *i.e.*, the frequencies that achieve high SNRs for that receiver.

We use a randomized greedy approach for the subband assignment problem. The algorithm maintains two data structures:

- `SubbandAssignment`: Stores the current nexthop assignment for each subband.
- `RateCounter`: Stores the total rate currently assigned to each nexthop.

The algorithm first orders the N nexthops randomly. It initially assigns all subbands to nexthop 1. The `RateCounter` for that nexthop is assigned the sum of the rates that it would have obtained from all these subbands, and all other `RateCounter` values are set to 0.

At each step, we pick the nexthop $nmin$ with the smallest `RateCounter` value, breaking ties randomly. We now need to assign an additional subband to this nexthop so that it can achieve a higher rate. To do this, we pick the nexthop $nmax$ with the largest `RateCounter` value. For each subband i assigned to $nmax$, we compute $\Delta rate[i] = rate_{nmin}[i] - rate_{nmax}[i]$, where $rate_{nmin}[i]$ and $rate_{nmax}[i]$ are the rates that nexthops $nmin$ and $nmax$ would obtain from subband i respectively. We then change the `SubbandAssignment` for the subband with the largest $\Delta rate$ from $nmax$ to $nmin$, and update the corresponding `RateCounter` values accordingly.

We keep repeating this process and stop when we cannot increase the minimum rate, which means that the receivers have as close a rate to each other as possible.

The algorithm above aims to allocate to each nexthop the frequencies that work better for it than for other nexthops. This is achieved by assigning subbands to a receiver according to the decreasing order of $\Delta rate[i]$. It also aims for equal rates to all concurrent nexthops. This is achieved by moving subbands to the nexthop that has the minimum rate so far, and repeating until we can no longer increase the minimum rate.

The header of each transmitted frame includes the number of nexthops and their addresses, as well as a bitmap with the frequency assignment. This allows each nexthop to learn the frequency subbands used for its packet. In contrast to traditional 802.11 where each data packet is followed by one synchronous ACK, a data frame that encapsulates N packets is followed by N synchronous acks from the corresponding nexthops. The acking order is determined by the order of the nexthops in the header of the data frame, and the acks are separated by a SIFS.

7.1 Wireless LANs vs. Mesh Networks

FARA can be used both in wireless LANs and mesh networks. Further, our description of the protocol directly applies to both. We note, however, that the benefits of applying FARA differ between these two scenarios. Specifically, in a mesh network, any node typically has multiple neighbors which constitute its potential nexthops. Hence, a mesh sender can derive gains from both FARA's frequency-aware rate adaptation and MAC protocols.

In contrast, a wireless LAN has two types of nodes, APs and clients. Since an AP is associated with many clients, the downlink, which carries the bulk of the traffic, can benefit from both frequency-aware rate adaptation and MAC protocols. On the uplink, however, the client is associated with a single AP, and hence has only one potential nexthop. While a client does not benefit from a frequency-aware MAC, it can still benefit from a frequency-aware rate adaptation protocol.[4]

---

[4]One extension to FARA would be to allow concurrent senders, in addition to concurrent receivers. A fine-grained allocation of OFDM subbands to concurrent senders, however, would require the senders to be synchronized to within an OFDM symbol to avoid power leakage between subbands. We therefore leave this for future work.

8 LOAD-AWARE CONTENTION

Since all senders in a FARA network have dynamic access to the entire frequency band, FARA naturally eliminates the problem of underutilizing the frequency spectrum due to inefficient 802.11 channel allocation [19, 37].

It is also straightforward to improve load balancing in a FARA network by exploiting prior work on load balancing for CSMA networks [21]. Specifically, FARA's contention-aware load balancing is based on two simple techniques. First, each AP or router contends for the medium by simulating contention from as many clients as have packets in its queue. Tracking the number of active clients is relatively simple. The AP or router keeps a hash table of counters. Whenever it receives a packet, it hashes the IP address of the nexthop, and increments the corresponding entry in the hash table. Whenever it transmits a packet, it hashes its nexthop IP and decrements the corresponding entry. Packets which arrive into a full queue are not counted. The number of active clients is equal to the number of non-zero entries in the hash table. This value needs to be updated only when an entry changes to or from zero. Say the number of active clients is N, the FARA AP picks N random contention slots and transmits in the smallest one as long as no other node transmits first.

The second technique scales the size of the contention window as a function of the number of contenders for the medium. Specifically, since FARA nodes contend for the entire medium without channelization, the average contention is higher. To deal with this issue, we leverage prior research on scaling the contention window with the level of contention. Specifically, IdleSense [21] updates the size of the minimum contention window depending on how long the medium is idle. FARA can use this result directly for its contention window scaling.

Combined, these two techniques allow a node to compete for the medium in proportion to its load, while ensuring that CSMA contention avoidance stays efficient.

8.1 Hidden Terminals

One concern with discarding channelization is that it might increase hidden terminal scenarios. FARA uses a simple solution that extends adaptive RTS-CTS activation [52], a commonly used mechanism to detect and address hidden terminals. Specifically, since FARA's SNR based rate adaptation allows the sender to converge to the correct rate within a few packets, a persistently high loss rate is a good indication of interference due to hidden terminals. Hence, a sender turns on RTS-CTS to a receiver whenever the loss rate to that receiver exceeds a configured threshold (20% in our case). FARA can also additionally leverage recent techniques to solve the hidden terminal problem such as [17].

9 PERFORMANCE

We have implemented a prototype of FARA in FPGA using the WiGLAN radio platform [35], and evaluated it in a wireless testbed.

(a) Hardware: We use the WiGLAN transceiver platform shown in Fig. 5. The radio board connects to the PC via the PCI bus, and acts like a regular network card. The radio spans 100 MHz of bandwidth around the 802.11a spectrum and its FPGA code implements standard 802.11 transmit and receive chains, including OFDM over BPSK, 4-QAM, 16-QAM, and 64-QAM modulations. It however differs from traditional 802.11 cards in that it does not use channelization, and hence allows a node to directly access the medium over a 100 MHz of spectrum.

Figure 5: The WiGLAN radio platform used in FARA's evaluation.

Figure 6: Testbed topology showing node locations. The node marked `tx` is used as a transmitter/AP. The other 17 locations are used for receivers.

(b) Implemented Infrastructure: Comparing FARA to a frequency-oblivious 802.11 system requires implementing an evaluation infrastructure that is suitable for running both standard 802.11 and FARA.

(a) Supporting 802.11: The WiGLAN board does not implement the 802.11 convolutional codes. Thus, we implement the 802.11 convolutional codes in software and apply them on the signal before passing it to the radio board to be modulated. Matlab has a reference implementation of 802.11 convolutional codes as part of its communication toolbox. It includes the scrambler, the convolutional coder, and the interleaver. We use this reference implementation to ensure that packets receive the same error protection that they would receive with a complete 802.11 implementation.

(b) Supporting FARA: We have implemented both FARA's rate adaptation algorithm and MAC protocol. Specifically, we augmented the FPGA code on the radio board to measure the SNR in each OFDM subband as explained in §5. The FPGA is also programmed to use SNR measurements to predict the optimal bit rate for each OFDM subband using the table in Fig. 1 and communicate it back to the sender. Finally, the frequency-aware MAC is implemented partially in software in the driver and partially in FPGA. The driver divides the subbands between potential nexthops, whereas the PHY code in the FPGA uses this subband assignment to transmit packets concurrently in one frame.

9.1 What is the Opportunity from Frequency Diversity?

Frequency diversity is a known property of wireless channels. However, if the performance of a frequency subband changes too quickly (say every millisecond), it will be hard to track it without excessive overhead. Exploiting frequency diversity in rate adaptation and MAC protocols requires the performance of the subbands to change slowly in comparison with the adaptation timescale.

Method. We use the topology in Fig. 6, where the node labeled `tx` transmits and the rest of the nodes receive. Since we have a total of 5 radio boards, we fix one of them as the transmitter and move the other boards to cover all the locations indicated in the figure. Each run lasts for 10 minutes, and is repeated 5 times. The receivers continuously measure the SNRs in all subbands and report the values as a function of time.

Results. Top graphs in Figs 7 (a) and Fig. 8 (a) show a plot of subband SNR for both non-line-of-sight (NLOS) and line-of-sight (LOS) channels. The transmission band is depicted as centered on 0
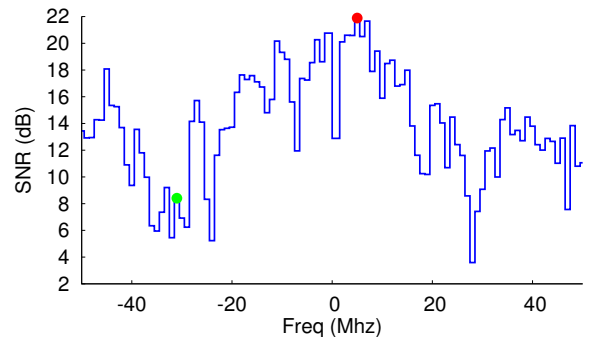

(a) Subband SNR for a Non-Line-of-Sight (NLOS) channel

(b) SNR over time for two typical NLOS subbands

Figure 7: Frequency Diversity in Non-Line-of-Sight (NLOS): The top graph presents the SNR in each 1-MHz wide subband across the 100 MHz band of our radio for a typical NLOS channel in our testbed. The graph shows that the subband SNRs can differ by more than 20 dB. The bottom graph shows the SNRs of two subbands in the top graph as a function of time. It reveals that the subband SNRs are stable over a multiple-second time period, thereby allowing an adaptive scheme to harness the frequency diversity.

and, the subbands are numbered from -50 to 50, as is conventional for baseband representation. The figures show that SNR differs significantly across subbands for both cases. Differences can be as high as 15–25 dB for the NLOS channel. The LOS channel is less diverse. Nonetheless its subband SNRs can vary by as much as 5–10 dB. Thus, a frequency-aware rate allocation scheme can derive benefits in both these channels.

(a) Subband SNR for a Line-of-Sight (LOS) channel

(b) SNR over time for two typical LOS subbands

Figure 8: Frequency Diversity in Line-of-Sight (LOS): The top graph presents the subband SNR for a typical LOS channel in our testbed. It shows that the subband SNRs can differ by more than 5-10 dB. While the variation is smaller than in NLOS channels, it is still significant. The bottom graph shows the SNRs of two subbands in the top graph as a function of time. It reveals that the subband SNRs are stable over a multiple-second time period, thereby allowing an adaptive scheme to harness the frequency diversity.

The bottom graphs in the same figures show how the SNR in a representative subband varies over time. As can be seen, the SNRs largely vary within only a narrow interval even over a period of several seconds, except for the rare deep fade. Hence, a rate adaptation scheme based on SNRs can successfully harness the frequency diversity.

9.2 Can We Robustly Map SNR to Best Bitrate?

Method. Harnessing frequency diversity in a rate adaptation scheme requires mapping an SNR value to the maximum sustainable bit rate, *i.e.*, to a combination of modulation and code rate. FARA uses a table look up for this mapping. Underlying our approach is an assumption that given an SNR value, one can determine the optimal combination that maximizes the throughput independent of location and time. Thus, in this experiment, we show that the SNR value robustly determines the best bit rate.

As in the previous experiment, the `tx` node in Fig. 6 transmits and the rest of the nodes measure the received SNRs in each subband. For this experiment, we treat each subband completely independently, *i.e.*, we assign it its own modulation, convolutional FEC code, and checksum. The separate checksum allows us to decide whether the bits in a particular subband are decoded correctly, independent from the bits in other subbands. The sender's transmissions use all 802.11 bit rates in a round robin manner, assigning the same rate to all subbands. For each received packet, the receiver reports


Figure 9: Mapping SNRs to a bitrate (i.e., a modulation and code rate): Plots the throughput per 1-MHz subbband as a function of SNR for each choice of modulation and code rate. It shows that for any SNR, the optimal choice of modulation and code rate is fairly clear.
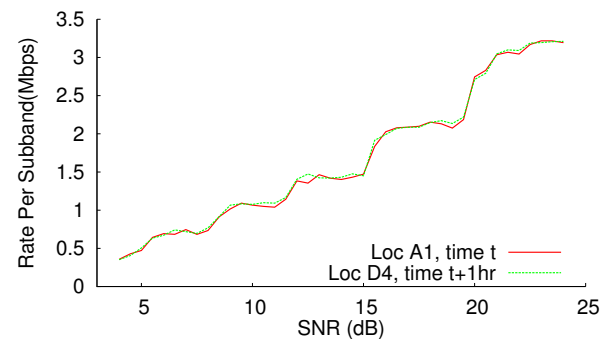

Figure 10: Stability of the relation between SNR and the optimal modulation and code rate across time and space. The figure shows the envelope of the functions in Fig. 9, for measurement collected at receivers `A1` and `D4` at different times. The two envelopes match closely showing that the SNR dictates the best modulation and code rate.

the SNR in each subband and whether the bits in that subband have passed the checksum test. We aggregate this information across all subbands and all receivers. We plot in Fig. 9 the bits per second per 1-MHz subband, *i.e.*, the throughput of a single 1-MHz subband as a function of its SNR, for all 802.11 modulation and code combinations.

Results. Fig. 9 shows that the per-subband SNR clearly determines the optimal modulation and code rate. For example, when the subband SNR is 17 dB, the optimal choice is the third from the top, i.e., 16-QAM and code rate of 3/4. Using any higher modulation or code rate reduces the probability of decoding the bits in that subband and brings the per-subband throughput close to zero. Using any lower modulation and code rate reduces the subband throughput. On the other hand, when the subband SNR is below 3.5 dB, no combination of modulation or code rate works. In this case, it is better not to transmit in that subband, i.e., to suppress that subband.

Fig. 10 plots the envelope of the curve in Fig. 9 for two different locations and times. It shows that the mapping of SNR to a modulation and code rate is stable across time and space. Thus, mapping subband SNRs to bit rates requires only a table lookup which reports the SNR values that cause a transition from one set of modulation and code rate to the next. In fact, Table 1, which we presented in §6.2, summarizes the information in the previous figures and is all that a bitrate adaptation protocol needs to map SNRs to bitrates.

(a) Frequency-aware rate adaptation for a 100 MHz channel

(b) Frequency-aware rate adaptation for a 20 MHz channel

Figure 11: FARA Rate Adaptation: FARA's frequency-aware rate adaptation achieves higher throughput than SampleRate's frequency-oblivious rate adaptation at all locations, with gains varying from $1.4\times$ to $3.6\times$ for a 100 MHz wide channel, and $1.1\times$ to $1.5\times$ for the 20 MHz channel.

9.3 Gains of Frequency-Aware Rate Adaptation

Now that we have established the existence of frequency diversity, its stability which makes it amenable to be harnessed by a rate adaptation protocol, and the robustness of the mapping from SNR to optimal bitrate, we measure the experimental gains from a frequency-aware rate adaptation protocol.

Method. Again we use the topology in Fig. 6. We fix the sender in location tx and randomly pick a receiver location. We repeat the experiment for all receiver locations shown in Fig. 6. For each location, we compare two schemes. The first is FARA's frequency-aware rate adaptation as described in §6. The second uses SampleRate [5], a well known rate adaptation scheme that assigns the same bitrate to all subbands. Each run lasts for ten minutes, and is repeated five times. We look at the benefit of frequency-aware adaptation for two scenarios: a standard 20 MHz 802.11 channel, and a wide 100 MHz channel.

Results. Fig. 11 shows that FARA's frequency-aware rate adaptation achieves significantly higher throughput than a frequency-oblivious algorithm such as SampleRate. Specifically, for a standard 20 MHz channel, a frequency-aware rate adaptation scheme increases the throughput by $1.24\times$. These gains become even higher as we move to wide and bonded channels, where FARA's rate adaptation improves the average throughput by $2.1\times$ over SampleRate.

The throughput gain is larger for receivers with worse channels. For example, some of the worse receivers experience a throughput gain that is as high as $3.5\times$. This is due to FARA's ability to avoid bad frequency bands. Specifically, SampleRate's frequency-oblivious rate adaptation experiences significant errors from sub-


Figure 12: Gains from a Frequency-aware Architecture: The figure plots two CDFs. The dashed line is the CDF of the ratio of client throughput under FARA to its throughput in traditional 802.11 networks which use SampleRate and CSMA MAC. The solid line is the CDF of the ratio of client throughput under FARA with a CSMA MAC and traditional 802.11 with SampleRate and CSMA. The CDFs show that FARA provides on average $3\times$ throughput gain. 70% of the gain comes from FARA's frequency-aware rate adaptation, and 30% is due to its frequency-aware MAC protocol.

bands that have very low SNRs and hence cannot support even the lowest transmission rate. To compensate for such bad subbands, SampleRate has to drastically lower its average transmission rate and increase coding across all subbands. In contrast, FARA suppresses subbands with less than 3.5 dB SNR and does not need to reduce the rate of every subband to compensate for the extra errors from such bad subbands.

Also, the throughput gain for NLOS channels is typically higher than the gain for LOS channels, because these channels see higher frequency diversity due to the greater prevalence of multiple paths with similar attenuation. Interestingly, location A2 shows significant throughput gain even though it has a LOS channel to tx, because it is within a passage that provides multiple opportunities for reflected waves that together create significant frequency diversity.

9.4 Gains of Frequency-Aware MAC

We now examine the throughput improvement provided by a frequency-aware MAC over a frequency-oblivious MAC.

Method. We again use the topology in Fig. 6. We collect measurements by transmitting from node tx to four random receiver nodes. We consider only four concurrent receivers because we have a total of five radio boards (including the transmitter). However, we can experiment with various scenarios by choosing different receiver sets. We run the experiment 10 times for each set of receivers, and repeat for a variety of receiver sets. We compare two MAC protocols: first, a frequency-oblivious CSMA MAC, where a sender checks whether the medium is available and transmits the packet at the head of its queue, and second, FARA's frequency-aware MAC as described in §7. Note that FARA transmits four packets in every frame and hence has less medium sensing overhead. Thus, to ensure that the differences between the two MACs are due only to frequency diversity, and not medium access overhead, we allow the sender to transmit its packets without waiting for an idle medium. This optimization favors the baseline MAC, and is possible because we have only a single sender in each experiment. Note that both FARA and the CSMA MAC use the same spectrum of 100 MHz.

Results. Fig. 12 plots the CDFs of the ratio of the throughput in FARA to the throughput in traditional 802.11 which uses Sam-
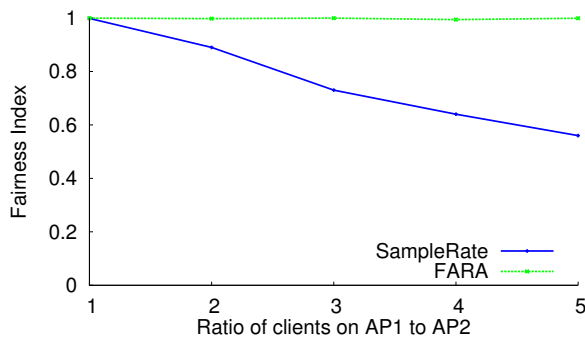
Figure 13: Load balancing with FARA: The figure plots the Jain's Fairness Index as a function of the ratio of the number of clients on AP1 to those on AP2. Note that fairness is optimal when the index is 1, and is worst when the index is $1/n$, where n is the total number of clients.

pleRate and a CSMA MAC. The CDF is computed across all receivers in our testbed and all runs. The graph contains two CDFs, one for a full-fledged FARA, and one for FARA after replacing its frequency-aware MAC with a frequency-oblivious CSMA MAC. The figure shows that a full-fledged FARA improves the median throughput by $3.1\times$ over a traditional SampleRate based CSMA MAC. The figure also demonstrates that about 30% of this gain is due to FARA's frequency-aware MAC while 70% is due to its rate adaptation scheme, showing that both mechanisms contribute significantly to the throughput improvements. Finally, it shows that all clients achieve significantly higher throughput with FARA than with traditional 802.11, which shows that FARA is beneficial to all nodes.

9.5 Load Balancing

We now demonstrate that, with FARA, multiple APs sharing the same frequency spectrum can achieve load balancing by using the load aware contention scheme described in §8.

Method. We use a modified version of the topology in Fig. 6. Specifically, we put two transmitters (i.e., two APs) around the location `tx`, and place their corresponding receiver nodes at `C3` and `D4`, respectively. Since we have a small number of radio boards, we make each board simulate a number of clients. AP1 simulates a varying number of backlogged clients, ranging from 1 to 5, on the link to `C3`, while AP2 always has only one backlogged client at `D4`. This setup allows us to experiment with scenarios with imbalanced loads, where AP1 has up to 5 times the number of clients of AP2. We perform an infinite download to each client and compute the per client throughput averaged over the first 10 minutes.

Using the per client throughput, we compute the Jain Fairness Index for the network, as follows [25]:

$$Fairness\ Index = \frac{(\sum_1^n x_i)^2}{n \sum_1^n x_i^2},$$

where x_i is the throughput of client i. The network shows optimal fairness when the index is 1, and is completely unfair when the index is $\frac{1}{n}$, in which case only one client has traffic.

We compare channelized 802.11 to a single channel FARA-equipped 802.11. Since we have only two APs, we limit both FARA and the channelized 802.11 to a total bandwidth of 40 MHz. In the case of channelized 802.11, the 40 MHz band is divided into two channels of 20 MHz and each AP is assigned a different channel. In the case of FARA, only one channel of 40 MHz is used for both APs.

Results. Fig. 13 plots the Jain fairness index as a function of the ratio of the number of clients at AP1 to those at AP2. The figure shows that the Jain fairness index remains close to 1 for FARA, whereas, for traditional 802.11, it drops linearly as the difference in load between the APs increases. This is because, in traditional 802.11, different APs operate on different channels and hence the single client on AP2 enjoys a throughput that is about the sum of the throughputs of all clients on AP1. In contrast, FARA discards channelization and further allows each AP to contend for the medium in proportion to the number of active clients. Hence, it allows the nodes to achieve a fairer throughput distribution.

10 Discussion

While the results in this paper have been presented in the context of 802.11, FARA applies to a wider variety of scenarios. Specifically, measurement studies show the existence of frequency diversity in the WiMax, UWB, and the 60 GHz range [2, 14, 36, 43]. All of these technologies use OFDM and have static or moderately dynamic applications, where the per-subband SNRs change relatively slowly [51, 24, 16]. FARA naturally extends to these scenarios.

FARA can also be extended for cognitive operation, and applied to the newly introduced whitespaces. The FCC has recently opened up for unlicensed access 100-250 MHz of digital whitespaces vacated by television bands as part of the analog-to-digital transition [12]. These whitespaces demonstrate significant frequency diversity [3, 39]. Further, they are expected to be used for several static and low mobility scenarios such as fixed wireless broadband access in rural and urban areas, as well as data connectivity inside the home, where FARA could provide significant throughput benefits. To do so, FARA needs to detect which subbands are occupied by the primary owner of the whitespace, and avoid these occupied subbands. FARA can leverage much prior work on detecting and agreeing upon occupied subbands [1, 40], and avoiding them by suppressing these occupied subbands, as in the cognitive PHY of [40], and as discussed in §6.2.

11 Concluding Remarks

This paper addresses the challenge and the opportunity of frequency diversity presented by the growing trend of wireless systems to use wider frequency bands. It demonstrates that a frequency-aware design of the physical, link and MAC layers offers significant throughput improvements both for a single client and for a network of clients, as compared to current frequency-oblivious rate adaptation and medium access schemes.

References

[1] I. F. Akyildiz, W.-Y. Lee, M. C. Vuran, and S. Mohanty. NeXt Generation/Dynamic Spectrum Access/Cognitive Radio Wireless Networks: A survey. In *Computer Networks Journal*. Elsevier, September 2006.

[2] J. G. Andrews, A. Ghosh, and R. Muhamed. *Fundamentals of WiMAX: Understanding Broadband Wireless Networking*. PrenticeHall, 2007.

[3] S. R. Banerjee, R. Jesme, and R. A. Sainati. Investigation of spatial and frequency diversity for long range UHF RFID. In *IEEE Antennas and Propagation Society International Symposium*, San Diego, CA, 2008.

[4] H. L. Bertoni. Propagation effects observed indoors. `http://eeweb.poly.edu/faculty/bertoni/docs/06IndoorObserve.pdf`, 2005.

[5] J. Bicket. Bit-rate selection in wireless networks. Master's thesis, Massachusetts Institute of Technology, 2005.

[6] J. Camp and E. Knightly. Modulation Rate Adaptation in Urban and Vehicular Environments: Cross-layer Implementation and Experimental Evaluation. In *ACM Mobicom 2008*, San Francisco, CA, September 2008.

[7] P. Chow, J. Cioffi, and J. Bingham. A practical discrete multitone transceiver loading algorithm for data transmission over spectrally shaped channels. *IEEE Transactions on Communications*, 48, 1995.

[8] H. K. Chung and H. L. Bertoni. Indoor propagation characteristics at 5.2 GHz in home and office environments. *Journal of Communication and Networks*, 4:176–188, 2002.

[9] CNN. Cutting the cord to flat-screen TVs. `http://www.cnn.com/2008/TECH/01/03/wirelesshigh.def.ap/index.html`, January 2008.

[10] F. Edalat, J. K. Tan, K. M. Nguyen, N. Matalon, and C. G. Sodini. Measured Data Rate from Adaptive Modulation in Wideband OFDM Systems. In *IEEE International Conference on UWB*, Waltham, MA, Sept. 2006.

[11] E. Eleftheriou and S. Olcer. Low-density parity-check codes for digital subscriber lines. In *Proceedings of IEEE International Conference on Communications (ICC)*, volume 3, pages 1752–1757, 2002.

[12] Second Rep. and Order and Memorandum Opinion and Order, November 2008. FCC 08-260.

[13] R. F. Fischer and J. B. Huber. A new loading algorithm for discrete multitone transmission. In *Proceedings of IEEE Global Telecommunications Conference (GlobeCom)*, volume 1, pages 724–728, 18-22 November 1996.

[14] J. Foerster. Channel modeling sub-committee report final, February 2003. IEEE P802.15 Wireless Personal Area Networks.

[15] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.

[16] New Wireless Standard Promises Ultra-Fast Media Applications. `http://www.sciencedaily.com/releases/2009/01/090122161953.htm`, January 2009. Science Daily.

[17] S. Gollakota and D. Katabi. ZigZag Decoding: Combating Hidden Terminals in Wireless Networks. In *ACM SIGCOMM 2008*, Seattle, WA, August 2008.

[18] S. Guha, K. Munagala, and S. Sarkar. Jointly optimal transmission and probing strategies for multichannel wireless systems. *Information Sciences and Systems, 2006 40th Annual Conference on*, pages 955–960, March 2006.

[19] R. Gummadi and H. Balakrishnan. Wireless networks should spread spectrum based on demands. In *Proc. ACM Hotnets, Calgary*, Oct 2008.

[20] J. Heiskala and J. Terry. *OFDM Wireless LANs: A Theoretical and Practical Guide*. Sams Publishing, 2001.

[21] 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. *SIGCOMM Comput. Commun. Rev.*, 35(4):121–132, 2005.

[22] D. Hughes-Hartogs. *Ensemble Modem Structure for Imperfect Transmission Media*. U.S. Patents no. 4,679,227 (July 1987), 4,731,816 (March 1988), and 4,833,706 (May 1989).

[23] Wireless channel bonding. `http://www.intel.com/support/wireless/sb/CS-025343.htm`. Intel Technologies.

[24] Enabling high-speed wireless personal area networks. `http://www.usb.org/wusb/docs/Ultra-Wideband.pdf`, 2005. Intel White Paper.

[25] R. Jain, W. Hawe, and D. Chiu. A quantitative measure of fairness and discrimination for resource allocation in shared computer systems. Technical Report DEC-TR-301, DEC, September 1984.

[26] D. Jiang, Q. Chen, and L. Delgrossi. Optimal data rate selection for vehicle safety communications. In *Vehicular Ad Hoc Networks*, 2008.

[27] I. Kalet. The multitone channel. *IEEE Transactions on Communications*, 37:119–124, February 1989.

[28] S. Kandula, K. C.-J. Lin, T. Badirkhanli, and D. Katabi. FatVAP: Aggregating AP Backhaul Capacity to Maximize Throughput. In *5th USENIX Symposium on Networked Systems Design and Implementation*, San Francisco, CA, April 2008.

[29] S.-C. Kim, H. Bertoni, and M. Stern. Pulse propagation characteristics at 2.4 GHz inside buildings. *Vehicular Technology, IEEE Transactions on*, 45(3):579–592, Aug 1996.

[30] S. K. Lai, R. S. Chen, K. B. Lataief, and R. D. Murch. Adaptive trellis coded MQAM and power optimization for OFDM transmission. In *Proceedings of IEEE Vehicular Technology Conference*, volume 1, pages 290–294, 16-20 May 1999.

[31] Lattice announces low-cost FPGA-based ADC interface reference design solution. `http://www.latticesemi.com/corporate/newscenter/productnews/2008/r080114announceslowcostfp.cfm`, 2008. Lattice Semiconductors.

[32] V. K. N. Lau and M. D. Macleod. Variable-rate adaptive trellis coded QAM for flat-fading channels. *IEEE Transactions on Communications*, 49:1550–1560, September 2001.

[33] Y. Li and J. Moon. Increasing data rates through iterative coding and antenna diversity in OFDM-based wireless communication. In *Proceedings of IEEE Conference on Global Telecommunications (GlobeCom)*, volume 5, pages 3130–3134, November 2001.

[34] V. Mannoni, D. Declercq, and G. Gelle. Optimized irregular low-density parity-check codes for multicarrier modulations over frequency-selective channels. *EURASIP Journal on Applied Signal Processing*, 204(10):1546–1556, 2004.

[35] N. Matalon. An Implementation of a 5.25 GHz Transceiver for High Data Rate Wireless Applications. Master's thesis, MIT, EECS, July 2005.

[36] D. M. Matic, H. Harada, and R. Prasad. Indoor and outdoor frequency measurements for mm-waves in the range of 60 ghz. In *Proceedings IEEE 48th Vehicular Technology Conference*, pages 567–571, 1998.

[37] T. Moscibroda, R. Chandra, Y. Wu, S. Sengupta, P. Bahl, and Y. Yuan. Load-aware spectrum distribution in wireless LANs. In *International Conference on Network Protocols*, Oct 2008.

[38] Y. N. Papantonopoulos. High-speed ADC technology paves the way for software defined radio. http://www.rfdesignline.com/showArticle.jhtml?articleID=201202962, 2007.

[39] S. Plass, A. Dammann, S. Kaiser, and K. Fazel. Space-time-frequency diversity in the next generation of terrestrial digital video broadcasting. In *Multi-Carrier Systems and Solutions*, pages 101–110. Springer Verlag, Netherlands, 2009.

[40] H. Rahul, N. Kushman, D. Katabi, C. Sodini, and F. Edalat. Learning to Share: Narrowband-Friendly Wideband Networks. In *ACM SIGCOMM 2008*, Seattle, WA, August 2008.

[41] S. Redl, M. Weber, and M. W. Oliphant. *GSM And Personal Communications Handbook*. Artech House, 1998.

[42] A. Sabharwal, A. Khoshnevis, and E. Knightly. Opportunistic spectral usage: bounds and a multi-band CSMA/CA protocol. *IEEE/ACM Trans. Netw.*, 15(3):533–545, 2007.

[43] A. G. Siamarou and M. O. Al-Nuaimi. Wideband propagation measurements for indoor Rician fading radio channels at 62.4 GHz. In *Proceedings IEEE 53rd Vehicular Technology Conference*, pages 449–453, 2001.

[44] T. Starr, J. Cioffi, and P. Silverman. *Understanding Digital Subscriber Line Technology*. Prentice Hall PTR, 1999.

[45] Local and metropolitan area networks requirements part 11: Wireless LAN medium access control (MAC) and physical layer (PHY). http://standards.ieee.org/getieee802/download/802.11a-1999.pdf, 1999. IEEE Standards Association.

[46] Local and metropolitan area networks requirements part 11: Wireless LAN medium access control (MAC) and physical layer (PHY). http://standards.ieee.org/getieee802/download/802.11g-2003.pdf, 2003. IEEE Standards Association.

[47] D. Tse and P. Vishwanath. *Fundamentals of Wireless Communications*. Cambridge University Press, 2005.

[48] J. Veillcux, P. Fortier, and S. Roy. An FPGA implementation of an OFDM adaptive modulation system. *3rd Intl IEEE-NEWCAS Conference*, 2005.

[49] M. Vutukuru, H. Balakrishnan, and K. Jamieson. Cross-Layer Wireless Bit Rate Adaptation. In *ACM SIGCOMM 2009*, Barcelona, Spain, August 2009.

[50] C. Wilmot. Intel demonstrates fast new UWB WPAN at IDF Taiwan. http://www.tweaktown.com/articles/968/intel_demonstrates_fast_new_ultrawideband_wpan_at_idf_taiwan/index.html, 2006.

[51] Applications for WiMAX. http://www.wimax.com/education/wimax/ims, 2009. WiMax.com.

[52] S. H. Y. Wong, H. Yang, S. Lu, and V. Bharghavan. Robust rate adaptation for 802.11 wireless networks. In *MobiCom '06: Proceedings of the 12th annual international conference on Mobile computing and networking*, 2006.

[53] M. Wouters, G. Vanwijnsberghe, P. V. Wesemael, T. Huybrechts, and S. Thoen. Real time implementation on FPGA of an OFDM based wireless LAN modem extended with adaptive loading. *Proceedings of the 28th European Solid-State Circuits Conference*, pages 531 – 534, September 2002.

[54] Y. J. Zhang and K. B. Lataeif. Single- and multi-user adaptive pragmatic trellis coded modulation for OFDM system. In *Proceedings of IEEE Wireless Communications and Networking Conference*, volume 1, pages 9–14, 16-20 March 2003.