# Fine-grained Channel Access in Wireless LAN

Kun Tan
Microsoft Research Asia
Beijing, China
kuntan@microsoft.com

Ji Fang*
Microsoft Research Asia
and Beijing Jiaotong University
Beijing, China
v-fangji@microsoft.com

Yuanyang Zhang*
Microsoft Research Asia
and Beihang University
Beijing, China

Shouyuan Chen* Lixin Shi*
Microsoft Research Asia
and Tsinghua University
Beijing, China
v-lixshi@microsoft.com

Jiansong Zhang Yongguang Zhang
Microsoft Research Asia
Beijing, China
{jiazhang, ygz}@microsoft.com

## ABSTRACT

Modern communication technologies are steadily advancing the physical layer (PHY) data rate in wireless LANs, from hundreds of Mbps in current 802.11n to over Gbps in the near future. As PHY data rates increase, however, the overhead of media access control (MAC) progressively degrades data throughput efficiency. This trend reflects a fundamental aspect of the current MAC protocol, which allocates the channel as a single resource at a time.

This paper argues that, in a high data rate WLAN, the channel should be divided into separate subchannels whose width is commensurate with PHY data rate and typical frame size. Multiple stations can then contend for and use subchannels simultaneously according to their traffic demands, thereby increasing overall efficiency. We introduce FICA, a fine-grained channel access method that embodies this approach to media access using two novel techniques. First, it proposes a new PHY architecture based on OFDM that retains orthogonality among subchannels while relying solely on the coordination mechanisms in existing WLAN, carrier-sensing and broadcasting. Second, FICA employs a frequency-domain contention method that uses physical layer RTS/CTS signaling and frequency domain backoff to efficiently coordinate subchannel access. We have implemented FICA, both MAC and PHY layers, using a software radio platform, and our experiments demonstrate the feasibility of the FICA design. Further, our simulation results suggest FICA can improve the efficiency ratio of WLANs by up to 400% compared to existing 802.11.

## Categories and Subject Descriptors

C.2.1 [**COMPUTER-COMMUNICATION NETWORKS**]: Network Architecture and Design—*Wireless communication*

## General Terms

Algorithms, Design, Experimentation, Performance

*This work was performed while Ji Fang, Yuanyang Zhang, Shouyuan Chen and Lixin Shi were research interns at Microsoft Research Asia.


*SIGCOMM'10,* August 30–September 3, 2010, New Delhi, India.


## Keywords

Fine-grained channel access, OFDMA, Cross-layer, MAC

## 1. INTRODUCTION

Modern communication technologies are steadily advancing the physical layer (PHY) data rates in wireless local area networks (WLANs). For example, the latest ratified 802.11n standard [1] has boosted data rates to 600Mbps. This capacity growth is achieved primarily through wider channel bandwidths and advanced PHY techniques like MIMO (Multiple-Input Multiple-Output). Future standards like IEEE 802.11ac and 802.11ad are already poised to provide even faster PHY rates (>1Gbps) in the next few years.

However, the data throughput efficiency — the ratio between the network throughput and the PHY data rate — degrades rapidly as the PHY data rate increases due to the design of the current 802.11 medium access control (MAC) protocol. For example, given that most IP packets have a maximal transmit unit (MTU) size around 1500 bytes, the efficiency ratio in an 802.11n network at 300Mbps is only 20%. That is, the 300Mbps data rate can sustain an actual throughput of only 60Mbps [23].

The fundamental reason for this inefficiency is that the current MAC allocates the entire channel to one station as a single resource. This allocation strategy can become too coarse-grained when the channel width increases or PHY data rate increases. Even if a sender has a small amount of data to send, it still needs to contend for the entire channel. Such contention resolution time is therefore an overhead to the channel time used for data. Unfortunately, this overhead cannot easily be reduced due to constraints of current electronics and physical laws. As a result, the higher the PHY data rate, the lower the throughput efficiency will become.

One way to improve the MAC efficiency is to extend the useful channel time for data transmissions by sending larger frames. Indeed, IEEE 802.11n allows frame aggregation, i.e., sending multiple frames together in one contention period. However, when the PHY data rate increases, the aggregated frame size needs to increase as well: achieving an efficiency of 80% in a 300Mbps network would require frames to be as large as 23KB. This larger aggregated frame means longer delays as the sender must wait to collect enough frames before actual transmission, resulting in adverse effects to TCP, real-time applications like VoIP and video conferencing, and even Web browsing that involves chatty protocols or short-lived sessions.

We argue that a better way to improve WLAN efficiency is to effectively reduce the channel width and create more channels, where the channel width is commensurate with PHY data rate and typical frame size. Multiple stations can then contend for and use these smaller channels simultaneously according to their traffic demands, thereby amortizing MAC coordination and increasing overall efficiency. We call this method *fine-grained channel access* for high data rate WLANs.

It is, however, non-trivial to divide a wide channel band into multiple subchannels without losing useful channel bandwidth. One common practice is to allocate both edges of two adjacent subchannels as a "guard band" so that the useful transmissions are properly spaced to avoid interfering with each other. These guard bands can add up to significant overhead, though, especially if the number of subchannels is large. For example, 802.11a uses a 1.875MHz guard band at both edges of every channel. If a 20MHz channel is divided into four 5MHz subchannels, the overhead will amount to 75% of the total bandwidth. Further, the guard band width cannot be easily reduced due to power mask requirements and the difficulty of filter designs, independent of the width of a subchannel.

Orthogonal frequency division multiplexing (OFDM) is a well-understood PHY-layer technology that can eliminate the need to have guard bands, if the frequency and width of subchannels are strategically picked and transmission on each subchannel is synchronized in a way to become "orthogonal", and hence non-interfering, to one another. Although some cellular networks (*e.g.*, WiMAX [2] and 3GPP LTE [3]) have proposed using OFDM in channel multi-access (OFDMA), doing so requires tight synchronization among user handsets and they cannot support random access. It thus remains a new technical challenge for how to use OFDM-type channelization for fine-grained channel access among distributed and asynchronous stations in a random access WLAN, where it is impractical and unnecessary to achieve similar tight synchronization.

In this paper, we present the design and implementation of FICA, a novel cross-layer architecture based on OFDM that enables fine-grained subchannel random access in a high data rate WLAN. FICA introduces two key techniques to address the aforementioned challenges:

- FICA proposes a new PHY architecture based on OFDM. Solely relying on the coordination mechanisms provided by existing WLANs, carrier-sensing and broadcasting, FICA retains orthogonality among subchannels with low overhead.

- FICA employs a novel *frequency-domain contention method* that uses physical layer RTS/CTS signaling and *frequency-domain backoff* for contending subchannels. We show that frequency-domain contention is much more efficient than the conventional time-domain contention mechanism in a fine-grained channel access environment.

We have implemented a FICA prototype on the Sora software radio platform [24]. Our implementation demonstrates the feasibility of our key techniques for both PHY and MAC design. We further use detailed simulation to evaluate FICA in large-scale wireless environments under different traffic patterns. Our results show that FICA has up to a 4-fold gain in efficiency compared to existing 802.11n with all its optimizations.

In summary, this paper makes the following contributions. (1) We describe and examine the efficiency issue of current MAC protocols in the context of high-speed WLANs, and argue that this issue can be resolved by fine-grained channel access. (2) We design and implement FICA, a protocol that enables fine-grained subchannel random access in WLANs; (3) We demonstrate the feasibility of FICA with a prototype implementation on a software radio platform, and evaluate its performance using detailed simulation. To the best of our knowledge, FICA is the first system that enables fine-grained channel access in WLANs.

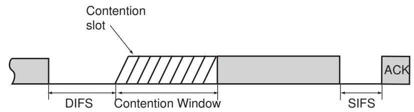


(a) Basic access

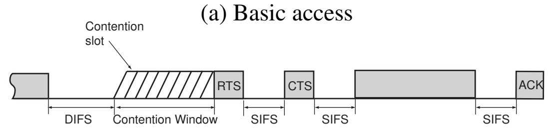


(b) With RTS/CTS handshake

**Figure 1: Illustration of CSMA/CA access method.**

The rest of paper is organized as follows. Section 2 provides a detailed analysis of the source of inefficiency in current MAC protocols. We then describe the design of FICA in Section 3 and evaluate its performance using simulation in Section 4. After describing the implementation of a FICA prototype using a software radio platform in Section 5, we evaluate its performance in Section 6. Finally, Section 7 discusses related work and Section 8 concludes.

# 2. BACKGROUND AND CHALLENGES

## 2.1 Inefficiency of Current WLANs

State-of-the-art MAC protocols in wireless LANs manage the whole channel (*e.g.*, 20/40MHz width) as a single resource. The MAC protocol arbitrates access among multiple potential senders and selects one as the winner, which then consumes the whole channel resource to transmit. If multiple senders transmit at the same time, collisions may happen and receivers will likely fail to decode the transmissions.

Current 802.11 WLANs use carrier sensing multiple access with collision avoidance (CSMA/CA) for their MAC protocol. When the channel is busy, all contending nodes wait until the channel becomes free. The MAC employs a random backoff scheme to avoid having multiple nodes transmitting simultaneously. Each node will randomly choose a number $b$ within a contention window $[0, CW)$, and wait for $b$ time slots before it starts transmitting. If a node detects a transmission detected its backoff period, it will freeze the backoff counter until the channel is free again. If two nodes randomly choose the same backoff time, their transmissions will eventually collide. A collision is usually detected by a missing acknowledgement (ACK) from the receiver. When a collision is detected, a sender will double its contention window $CW$ according to the binary exponential backoff (BEB) algorithm to further reduce the collision probability for the next transmission.

Figure 1 illustrates the channel access timing diagram of the 802.11 MAC. Figure 1(a) is the basic access method, and Figure 1(b) shows channel access with the optional RTS/CTS handshake to handle hidden terminals. The Short Inter-frame Space (SIFS) is the shortest time interval required for a receiver to return a message to a sender. It is determined by Equation 1, where $t_{\text{rf\_delay}}$ is the delay incurred to transfer digital signals from the RF antenna to the processing unit, $t_{\text{proc}}$ is the time needed for the processing unit to operate on the incoming signals, and $t_{\text{TxRx}}$ is the time needed for the RF front-end to switch from receiving mode to transmitting. Normally, SIFS is about 10–16$\mu s$. The Distributed Inter-frame Space (DIFS) is determined based on SIFS and the backoff slot time, as shown in Equation 2. DIFS is defined to support priorities in CSMA/CA and should be larger than SIFS. The backoff slot time is critical. It is the minimal time needed for a node to sense the channel condition and acquire the channel. Slot time is determined by Equation 3, where $t_{\text{cca}}$ is the time for a node to measure

| Parameter | Value |
|---|---|
| $t_{\text{slot}}$ | $9\mu s$ |
| $t_{\text{sifs}}$ | 10–16$\mu s$ |
| $t_{\text{cca}}$ | $4\mu s$ |
| $t_{\text{TxRx}}$ | $\leq 5\mu s$ |
| $t_{\text{prop}}$ | $\leq 1\mu s$ |
| $t_{\text{preamble}}$ | 20–56$\mu s$ |

**Table 1: Timing parameters of 802.11.**

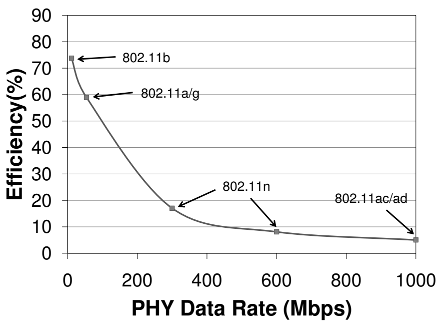


**Figure 2: Inefficiency of 802.11 MAC at high data rates with a typical Ethernet MTU (1500B).**

the channel energy to decide the channel status, and $t_{\text{prop}}$ is the time for the radio signal to reach the maximal distance of the network.

$$t_{\text{sifs}} = t_{\text{rf\_delay}} + t_{\text{proc}} + t_{\text{TxRx}}, \quad (1)$$

$$t_{\text{difs}} = t_{\text{sifs}} + 2 \cdot t_{\text{slot}}, \quad (2)$$

$$t_{\text{slot}} = t_{\text{cca}} + t_{\text{TxRx}} + t_{\text{prop}} + t_{\text{proc}}, \quad (3)$$

Using these values, we can build a simple analytical model to compute the efficiency ratio for CSMA/CA. Since a node chooses a random number uniformly from the contention window $[0, CW)$, the expected number of backoff slots is $W = \frac{CW}{2}$. Equation 4 gives the efficiency ratio for the basic access of CSMA/CA:

$$\eta = \frac{t_{\text{data}}}{t_{\text{slot}} \cdot W + t_{\text{difs}} + t_{\text{preamble}} + t_{\text{sifs}} + t_{\text{ack}} + t_{\text{data}}}, \quad (4)$$

where $t_{\text{data}}$ is the time used for data transmission, $t_{\text{preamble}}$ is the time used to transmit per-frame training symbols and $t_{\text{ack}}$ is the time used for the ACK frame.

Only $t_{\text{data}}$ is used for transmitting application data, while all other times are overheads. Some overheads are constrained by physical laws and current contraints in state-of-the-art radio electronics. For example, you cannot reduce $t_{\text{prop}}$ less than $1\mu s$ to cover a network with a radius of a few hundreds of meters. It is also difficult to reduce $t_{\text{TxRx}}$ since the RF circuit requires a few microsecond to settle down for sending or receiving. Others are needed for the correct operation of the protocol. For example, we need training symbols for reliable estimation of the wireless channel for each frame, thus $t_{\text{preamble}}$ is essential. The average backoff slots, denoted by $W$, reflects the ability of CSMA/CA to avoid collisions. Thus, to work well in normal network settings, we need a reasonably large $W$. ACKs are also needed to detect collisions and other losses, thus in general we do not want to remove $t_{\text{ack}}$. Table 1 outlines some timing parameters defined in 802.11. They remain similar across the different standards of 802.11a/g/n except for the preamble; since 802.11n uses MIMO, it requires more training symbols in its preamble.

Therefore, when the PHY data rate increases, only $t_{\text{data}}$ will be reduced proportionally, while the other parameters remain largely unchanged. As a consequence, the efficiency ratio $\eta$ decreases inversely proportionally. Figure 2 illustrates such a phenomenon: the efficiency quickly decreases from 60% at 54Mbps (802.11a/g) to less than 10% at 1Gbps (future 802.11ac/ad).

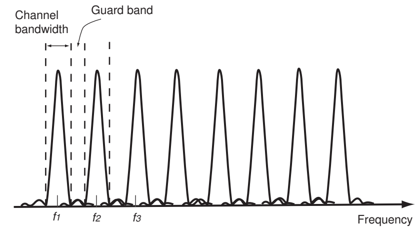


(a) Normal frequency division multiplexing

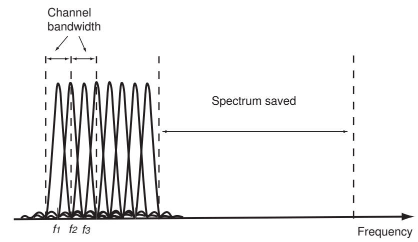


(b) OFDM

**Figure 3: OFDM achieves higher spectrum efficiency.**

As mentioned in Section 1, transmitting larger frames will improve the efficiency ratio, but such a frame-aggregation approach has practical limitations. Fine-grained channel access will be a better approach, if we can divide the whole channel into smaller sub-channels efficiently and allow different nodes to access different sub-channels simultaneously. Enabling concurrent transmission across sub-channels is in effect an aggregation and opportunity to amortize the MAC overhead across different nodes.

## 2.2 An OFDM Primer

Orthogonal Frequency Division Multiplexing (OFDM) has become increasingly popular in modern wireless communications [17]. It has been embraced by many existing wireless standards like IEEE 802.11a/g/n, WiMax [2], and by future standards like 3GPP LTE [3]. Cognitive radio technologies also mainly rely on OFDM to use non-contiguous spectrum bands for communication [19].

OFDM divides a spectrum band into many small and partially overlapping signal-carrying frequency bands called subcarriers. The subcarrier frequencies are chosen so that they are "orthogonal" to one another, meaning that cross-talk between subcarriers sums up to zero even though they are overlapping (Figure 3). OFDM can therefore pack subcarriers tightly together without inter-carrier interference, eliminating the need to have guard bands.

OFDM can be efficiently implemented using (inverse) Fast Fourier Transform (iFFT/FFT). In an OFDM system with FFT size $N$, each subcarrier has exactly the same width of $\frac{B}{N}$ and the subcarrier central points are located at frequencies of $f_c + \frac{2\pi n B}{N}, n = -\frac{N}{2}..(\frac{N}{2} - 1)$, where $f_c$ is the central frequency of the channel and $B$ is the channel width. Different modulations (*e.g.*, BPSK, QPSK, *etc.*) can be applied to each subcarrier independently. After modulating information onto each subcarrier, the sender performs an iFFT to convert the frequency domain presentation to $N$ time-domain samples which can be sent over the air. The time needed to transmit these $N$ samples is usually called the FFT period, which is equal to $\frac{N}{B}$ seconds. Thus, given a fixed channel width, a larger $N$ means a longer FFT period. Then, at the receiver side, the signal can be converted back to the frequency domain using the FFT, where each subcarrier can be demodulated independently.

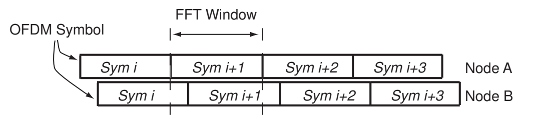


(a) Loss of orthogonality due to misaligned symbols

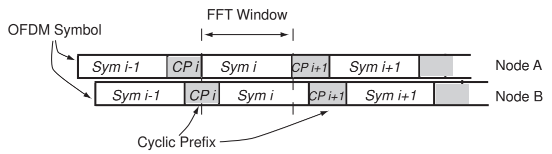


(b) Retaining orthogonality with a proper size CP

**Figure 4: Misaligned OFDM transmissions.**

When OFDM is used as a multi-access technology where multiple stations share the same channel, symbol timing alignment will be a critical issue. As shown in Figure 4(a), if OFDM symbols from two nodes misalign, the receiver may not be able to pick up an FFT window containing the same samples across all senders. Orthogonality will be lost and signals from both nodes will cause mutual interference. To ensure perfect symbol alignment, a multi-access technology called OFDMA has been proposed for OFDM cellular networks like WiMAX and LTE. OFDMA requires all mobile stations to maintain tight timing synchronization with the cellular base station (usually hundreds of nanoseconds). It requires a complex ranging scheme to measure the propagation delay and fine tune each mobile station's timing offset at the sample level granularity.

OFDM further has a built-in robustness mechanism called the cyclic-prefix (CP) [26] to guard against symbol misalignment due to multipath echoes. Each OFDM symbol is prefixed with a copy of the symbol's tailing samples so that the receiver can still find a proper FFT window as long as the misalignment is within a CP length (Figure 4(b)). CP is intrinsic to any OFDM system; in 802.11, the CP-to-symbol length ratio is 1:4 ($0.8\mu s$ to $3.2\mu s$).

## 2.3 Fine-grained Channel Access in WLAN

We propose to use fine-grained channel access to improve throughput efficiency in a high-data-rate WLAN. We divide the channel width into appropriately sized sub-channels commensurate with the PHY data rate and typical frame size, and further use OFDM on the whole channel to avoid wasting bandwidth on guard bands. The fundamental challenge with this approach is coordinating random access among multiple distributed and asynchronous nodes in a WLAN (potentially with multiple APs), without resorting to cellular-style tight timing synchronization.

Because coordination in a WLAN is distributed and decentralized in nature, it is impractical to have OFDMA-style global time synchronization. Not only would it introduce a great deal of system complexity, it would also likely require new hardware functionality beyond the current or emerging 802.11 standards. Further, OFDMA does not support random access and hence cannot be used directly in a WLAN.

Instead, we should use existing 802.11 coordination mechanisms, such as carrier-sensing and broadcast, to establish a rough symbol alignment among concurrent senders. We can leverage OFDM's intrinsic CP mechanism and lengthen it to suit the alignment scale, and further use a longer symbol length to maintain the same CP-to-symbol ratio. This approach calls for a new OFDM architecture specially designed for distributed coordination.

Having a longer symbol length, however, does have a negative impact that makes a conventional time-domain backoff scheme very inefficient. For example, if we can only guarantee a $10\mu s$ symbol alignment under current 802.11 coordination schemes, we will need a $40\mu s$ symbol length to keep the same guard-time overhead ratio. The reserved time slot for backoff, which has to be at least one OFDM symbol in length, will now increase proportionally. This raises another technical challenge: we need a new efficient MAC contention mechanism and a new backoff scheme. All these are necessary conditions for fine-grained channel access in high-data-rate WLANs.

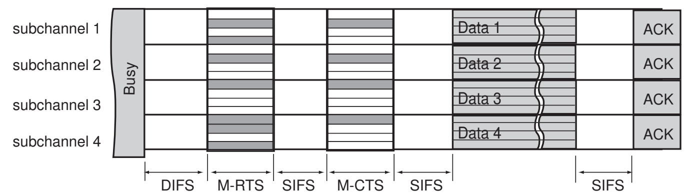


**Figure 5: FICA uplink media access with four subchannels per channel.**

# 3. FICA DESIGN

FICA is a cross-layer design that enables fine-grained channel access in high-rate wide-band WLANs. It is based on OFDM and divides a wide-band channel into a set of orthogonal subchannels, which different nodes can contend for individually. For the sake of clarity, we first assume a scenario consisting of a single access point and multiple mobile stations. We show later it is straightforward to extend FICA to the multi-AP case. Also, while in the following discussion we primarily use a 20-MHz channel as an example to explain the operation of FICA, we note that FICA is a scalable design that works for arbitrarily wide channel bands.

Figure 5 illustrates the basic uplink media access scheme for FICA. In this example, the channel is divided into just four subchannels and each subchannel contains a number of subcarriers. FICA follows the basic scheme of CSMA. A new transmission opportunity appears only when the whole channel is idle. Then, all stations try to contend for different subchannels after the channel is idle for a certain amount of time (DIFS). At this time, all nodes will transmit a special RTS signal simultaneously. This RTS signal is a specially-designed OFDM symbol, called *Multi-tone RTS* (see Section 3.3), in which each node embeds its contention information in a set of subcarriers for each subchannel it intends to access. All M-RTS signals are resolved at the AP, and the AP will broadcast the contention results in a corresponding *M-CTS* OFDM signaling symbol. Then, only the nodes assigned subchannels will use them for data transmissions; note that a node may contend for multiple subchannels based on its instantaneous traffic demands. The AP will then generate an acknowledgement on each subchannel where a data frame has been successfully decoded.

Downlink transmissions follow similar steps, but the AP will initiate an M-RTS signal and receiving stations may return an M-CTS. However, since FICA does not use random time backoff, it needs to separate uplink and downlink transmissions; otherwise, collisions would happen under bi-directional traffic. FICA does so by assigning different DIFS times to uplink and downlink transmissions, described further in Section 3.4.

We now present an analysis of the symbol timing misalignment problem in a CSMA-based WLAN. Based on this analysis, in Section 3.2 we describe the FICA PHY structure. We then describe how FICA achieves *frequency domain contention and backoff* in Section 3.3. We finish by discussing several related design issues in Section 3.5.

### 3.1 Symbol Timing Misalignment in WLANs

In a WLAN, transmissions from distributed nodes are coordinated based on carrier-sensing and overhearing broadcast frames (*e.g.*, RTS/CTS). Unlike conventional MACs that use these mechanisms to avoid simultaneous transmissions, FICA exploits simultaneous transmissions to enable concurrent access from different nodes but in orthogonal subchannels.

In FICA, carrier-sensing coordinates the transmissions of M-RTS. After they sense an idle channel for a fixed DIFS time, two nodes $A$ and $B$ may transmit their M-RTS symbols simultaneously. However, since there is always a delay to sense the channel and for a signal to propagate from one node to another, these two M-RTS symbols cannot be transmitted and received at exactly the same time. Figure 6(a) shows the worst case analysis of the symbol timing difference of two such M-RTS symbols received by the AP. Assume $A$ senses the channel idle for $t_{\text{difs}}$ seconds, and $A$ transmits M-RTS first. It will take at most $t_{\text{prop}}$ for the signal to arrive at node $B$. However, it may take $B$ at least $t_{\text{cca}}$ time to reliably sense the busy channel. If the signal arrives at node $B$ after $B$ starts a sensing slot, $B$ may not be able to assess a busy channel.[1] Thus, after the sensing slot, $B$ still declares an idle channel and sends its M-RTS as well. The radio signal of $B$'s M-RTS is actually emitted after $t_{\text{RxTx}}$. And it will take at most another $t_{\text{prop}}$ for the M-RTS to arrive at the AP. The total misaligned time is characterized by Eq. 5. The bound is tight in the worst case.

$$t_{\text{err\_rts}} \leq t_{\text{cca}} + t_{\text{RxTx}} + 2 \times t_{\text{prop}}. \quad (5)$$

If we use the values of the parameters listed in Table 1, we find that $t_{\text{err\_rts}}$ can be as large as $11\mu s$.

If coordination is performed by overhearing a broadcast frame, *e.g.*, M-CTS or DATA, the timing misalignment can be shorter because the two senders are synchronized better by a reference broadcast [6]. Figure 6(b) shows the worst case of symbol timing misalignment of data frames after receiving a M-CTS broadcast. Assuming the jitter for a local timer can be neglected since the waiting time is usually very small (a few microseconds), the timing misalignment is bounded by twice the propagation time, as shown in Eq. 6.

$$t_{\text{err}} \leq 2 \times t_{\text{prop}}. \quad (6)$$

Using the parameter values in Table 1, $t_{\text{err}}$ is about $2\mu s$ in a WLAN.

### 3.2 PHY Architecture

Based on the analysis in last subsection, FICA needs to provide a guard time sufficiently long to handle the symbol timing misalignment in a WLAN based on carrier-sensing ($11\mu s$) and broadcasting ($2\mu s$). We further need to include an additional $800ns$ guard time to account for the typical spreading delay in indoor environments [1]. We design two guard time sizes tailored to each coordination situation: a long cyclic prefix (CP) of $11.8\mu s$ and a short CP of $2.8\mu s$. The long CP is attached to M-RTS only, while a short CP is attached to every M-CTS, DATA, and ACK OFDM symbols, and therefore is the major overhead of concern.

To amortize the short CP overhead, we need a longer OFDM symbol for data. Recall from Section 2.2 that a longer OFDM symbol is achieved by applying a larger FFT size $N$. Although in theory one can choose any large FFT size $N$, there are a few practical considerations. First, a large $N$ requires more computational power to calculate the transform since the complexity of FFT is $O(Nlog(N))$. Second, with large $N$ the adjacent subcarriers are spaced very closely. Therefore, it is more sensitive to the frequency offsets of different nodes. Since multiple nodes may always have some small frequency offset (a few hundred Hz as we show later), the adjacent subcarriers should have enough separation to accommodate such an offset. In practice, subcarriers are usually separated by at least 15KHz [2, 3].

[1] This case depends on the way the node implements carrier sensing and threshold setting. Here we consider only the worst case.

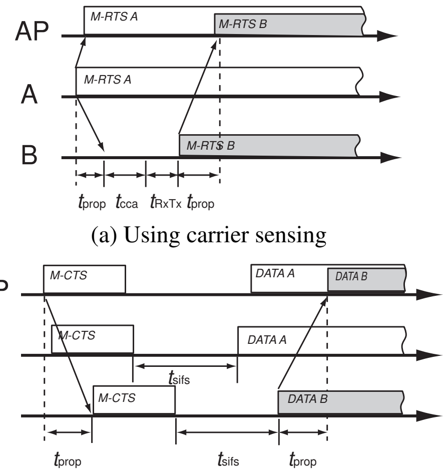


(a) Using carrier sensing

(b) Using a reference broadcast

**Figure 6: The worst-case symbol timing misalignment of two concurrent transmissions received at the AP in a CSMA WLAN: (a) symbol timing misalignment when coordination uses carrier sensing; (b) a reference broadcast synchronizes the senders better and timing misalignment is much tighter.**

In FICA, we choose the FFT size of the DATA OFDM symbol to be 256 points in a 20MHz channel (subcarrier width is 78.12KHz). Its FFT period is $12.8\mu s$, yielding a cyclic-prefix overhead of $18\%$ which is comparable to the CP overhead of 802.11 ($20\%$). ACK symbols have the same structure as DATA symbols. Since the M-RTS uses the long CP, it should contain as few symbols as possible. Thus, we choose a larger FFT size (512 points) for M-RTS to embed all necessary control information, which otherwise would require multiple OFDM symbols for a smaller FFT size. FICA applies the same FFT size to the M-CTS for format consistency. Table 2 summarizes the detailed time parameters of the OFDM symbol structure in FICA.

| Parameter | Value |
|---|---|
| $N_{\text{fft\_data}}$ | 256 points |
| $t_{\text{fft\_data}}$ | 12.8 $\mu s$ |
| $N_{\text{fft\_mrts}}, N_{\text{fft\_mcts}}$ | 512 points |
| $t_{\text{fft\_mrts}}, t_{\text{fft\_mcts}}$ | 25.6 $\mu s$ |
| $t_{\text{long\_cp}}$ | 11.8 $\mu s$ |
| $t_{\text{short\_cp}}$ | 2.8 $\mu s$ |
| $t_{\text{data\_sym}}$ | 15.6 $\mu s$ |
| $t_{\text{mrts\_sym}}$ | 37.4 $\mu s$ |
| $t_{\text{mcts\_sym}}$ | 28.4 $\mu s$ |

**Table 2: OFDM symbol timings in FICA.**

In principle, FICA can allocate each subcarrier independently to provide maximal flexibility. In practice, however, it is difficult for a node to use only a single subcarrier to transmit data for two reasons. First, using only one subcarrier suffers from frequency selective fading: if that subcarrier encounters deep fading, all data will be lost. Thus, it is essential for a wireless PHY to code across multiple subcarriers to achieve a spectrum diversity gain [17]. Second, although a preamble may be used for channel estimation and

| Configuration | FICA (Mbps) | 802.11n (Mbps) |
|---|---|---|
| 20MHz channel | 71.8 | 72.2 |
| 40MHz channel | 145 | 150 |
| 40MHz channel, 2xMIMO | 290 | 300 |
| 40MHz channel, 4xMIMO | 580 | 600 |

**Table 3: Example PHY data rates in FICA vs. 802.11n.**

compensation, the wireless channel may change during the period of data transmission. It is essential for the receiver to track the changes in the wireless channel. This tracking is typically done by adding an additional training subcarrier (pilot) along with other data subcarriers.

Therefore, following typical practice FICA groups a set of subcarriers into a subchannel and uses it as the basis for channel access.[2] Each subchannel contains sixteen data subcarriers and one pilot subcarrier. Thus, a single subchannel in FICA is 1.33MHz wide with a $6\%$ pilot overhead, which is comparable to 802.11 (7%). With different modulation modes, the PHY data rate of a single subchannel can range from 512Kbps (BPSK, 1/2 coding on each subcarrier) to 20.25Mbps (64QAM, 5/6 coding and four spatial multiplexing streams on four antennas).

With a 1.33MHz subchannel, a 20MHz 802.11 channel contains 14 orthogonal subchannels. FICA uses the remaining spectrum as guard bands separating adjacent wide-band channels. Note that it is also straightforward for FICA to support wider band channels, *e.g.*, 40–100MHz or wider. To support a 40MHz channel, for instance, we simply double the FFT size for all OFDM symbols; since our sampling rate is also doubled, the symbol period does not change. When using a 40MHz channel, we can save the guard bands that would be necessary to separate two 20MHz channels. Thus, we can have 29 orthogonal FICA subchannels with a 40MHz channel. Table 3 shows some example PHY data rates of FICA and compares them to 802.11n. As we will show later, although FICA has slightly lower data rates than 802.11n, the effective throughputs are actually dramatically higher due to diminished overheads.

## 3.3 Frequency Domain Contention

FICA uses the M-RTS/M-CTS signal exchange to avoid subchannel collisions. M-RTS/M-CTS use simple *binary amplitude modulation* (BAM) to modulate a single bit on each subchannel. Specifically, BAM uses On-Off Keying that maps a binary “0” to zero amplitude in a subcarrier, and it uses a random complex number on the unit circle ($e^{j\theta}$) for a binary “1”. Receivers can easily detect BAM symbols by comparing the energy on a subcarrier against a threshold, without the need to recover the symbol phase.

The basic idea of frequency domain contention works as follows. In M-RTS, a group of $K$ subcarriers are allocated for each subchannel, called a *contention band*. When a node contends for a subchannel, it will first randomly pick a subcarrier from the contention band and send a signal “1” using BAM. At the AP side, there may be multiple M-RTS signals superposed, and the AP may detect multiple ones on different subcarriers of the contention band. The AP then can arbitrate a winning node by selecting a subcarrier based on some predefined rules, *e.g.*, the one with highest frequency. Then the AP sends the identification of the winning subcarrier for each subchannel in an M-CTS signal. After receiving the M-CTS, each node compares the subchannel allocation broadcast by the AP to the subcarrier it has transmitted for the subchannel. If they match, the node will transmit data symbols on that corresponding subchannel after a SIFS delay.

There are a few issues that need to be considered.

**What if two nodes transmit on the same contention subcarrier?** It is quite possible for two nodes to choose the same random number and transmit on the same contention subcarrier. Thus, their transmitted energy is additive at the AP side. If the AP picks this subcarrier as the winner, a collision will happen since both nodes will consider themselves the winner and send data symbols on the same subchannel simultaneously. It is also possible (although the possibility is small) that two nodes’ signals are destructive and cancel each other, so the AP may not be able to detect the transmissions. In this case, though, the nodes involved will conclude that they were not assigned the subchannels they requested.

**How large should $K$ be?** Clearly, a large $K$ will have fewer collisions in a subchannel. A larger $K$, though, means more subcarriers are used and result in a larger signaling overhead: more subcarriers need a larger FFT size and therefore a longer FFT period. In FICA we set $K$=16, the initial contention window size in 802.11.

**Who is responsible for returning the M-CTS?** Any node that overhears M-RTS transmissions can arbitrate subchannel access. Simultaneous transmissions of M-CTS from different nodes are also allowed since, in a single broadcast domain, these M-CTS symbols are likely to have the same contention results and their energy is additive on each subcarrier. But it is still reasonable that only the potential receivers should return an M-CTS for an M-RTS based on the following considerations: (1) the irrelevant nodes can use power-save mode, since it is unfair for them to spend power responding to M-RTS signals not for them; (2) non-receiving nodes may unnecessarily create more exposed terminals if they reply with an M-CTS and disrupt transmissions that could otherwise happen without interference.

How we can specify receivers in an M-RTS? It is non-trivial since multiple nodes may transmit M-RTSs simultaneously to different receivers and the receiver information may be mixed. FICA resolves nnthis issue by using a membership vector of $m$ subcarriers to represent receiver information in the M-RTS. Before transmitting a M-RTS, a node will hash the receiver’s ID into a value between 0–$(m-1)$. Then, the corresponding subcarrier will carry a “1” bit. Any node receiving an M-RTS will check if the subcarrier corresponding to its ID has been set. If true, it should return an M-CTS. It is possible that multiple nodes have hash value collisions. But with a reasonably-sized $m$ — we use 40 bits in the current design — the number of station collisions should be small: 2–3 in a typical WLAN setting where one AP may be associated with dozens of stations [14].

### 3.3.1 M-RTS/M-CTS format

Table 4 summarizes the subcarrier allocation for the M-RTS and M-CTS symbols. We explain the M-RTS/M-CTS format using a 20MHz channel as a concrete example. M-RTS/M-CTS use a 512-point FFT in 20MHz channel (for wider channels, a proportionally longer FFT point size is used).

**Tag.** The tag band contains a predefined binary sequence to identify the type of this OFDM symbol. The tag band is interleaved on selected subcarriers whose indices are multiples of four. Thus, the tag band creates a repeating pattern in time-domain samples. The receiver can use this pattern to detect the symbol using correlation. In a 20MHz channel, the tag band contains 32 subcarriers. For M-RTS, the tag sequence starts with zero, $\{0, 1, 0, 1, ..., 0, 1\}$, while for M-CTS the sequence becomes $\{1, 0, 1, ..., 0, 1, 0\}$.

[2]The grouping of subcarriers may be arbitrary, and not necessarily contiguous. In our current implementation, however, we only group adjacent subcarriers into a subchannel.

| *M-RTS* | | *M-CTS* | |
|---|---|---|---|
| Tag | 32 | Tag | 32 |
| Contention band | $16 \times 14$ | Resolution band | $24 \times 14$ |
| NAV | 64 | NAV | 64 |
| Receiver band | 160 | Reserved | 48 |
| Guard band | 32 | Guard band | 32 |

**Table 4: Number of subcarriers allocated for the M-RTS and M-CTS symbols in a 20MHz channel.**

**Contention band and resolution band.** For M-RTS FICA allocates sixteen subcarriers to each subchannel for contention. So, a total of 224 subcarriers are allocated for a contention band for all 14 subchannels arbitrated using an M-RTS symbol.

Contention results are embedded in the *resolution band* in M-CTS. The resolution band occupies more subcarriers compared to the contention band in the M-RTS. There are 24 subcarriers for each subchannel. However, since BAM is less robust compared to phase-shift keying modulation (*e.g.*, BPSK), we improve BAM's reliability by replicating on multiple subcarriers and using soft-decoding. It can be shown that, by replicating twice, BAM can achieve the same performance of BPSK [25]. In FICA, we replicate a bit on four subcarriers when we require high reliability. Thus, with four-time replication, we can encode six bits $(24/4)$ to represent the resolution results for each channel. Currently, six bits can represent the winner number (1–16) and two other special states. If these 6 bits are all zero, it means the subchannel is not used (*e.g.*, no node has contended for this subchannel). Alternately, if the bits are all ones, the subchannel is reserved by the receiver; thus no transmission is allowed on this subchannel in this contention period.

**Receiver band.** The receiver band contains the hash vector for intended receivers across all subchannels. It contains 40 bits and each bit again is replicated on four subcarriers for reliability, for a total of 160 subcarriers.

**Network Allocation Vector (NAV).** The NAV band encodes the expected number of OFDM symbols in the following data frame. It contains 64 subcarriers and, also due to four times replication, can modulate 16 bits. There could be several ways to encode the data symbol number in the NAV band. One simple method is to use each single bit to present a certain number of OFDM symbols. Thus, sixteen bits can present numbers between 3–48, with a step of 3. As we discuss later, such a range is sufficient for FICA. In an M-CTS, only the highest bit is echoed back and all nodes that overhear the M-CTS should defer their contention for a corresponding period.

### *3.3.2 Frequency domain backoff*

In FICA, nodes will choose how many subchannels to request via contention based on their traffic requirements. If its traffic requirements are light, a node may request fewer subchannels, while a heavily loaded node may contend for as many as all subchannels for most efficient communication. However, if there are many nodes in a network contending for many subchannels, the collision avoidance mechanism provided by M-RTS/C-RTS may not be sufficient to represent all contenders. Using multiple M-RTS symbols can further scale collision avoidance, but at the cost of more signalling overhead as described above.

Instead, we use a novel *frequency domain backoff* scheme to scale FICA in a heavily-contended dense network. The basic idea is to control the maximum number of subchannels that one node can access based on the observed collision level. It is similar to existing widely-used congestion control mechanisms. Each node maintains a state variable, $C_{\max}$, which controls the maximum subchannels the node can access in the next transmission opportunity. Thus, when the channel is idle for DIFS, a node may pick up to $n$ subchannels to contend for, where $n = \min(C_{\max}, l_{\text{queue}})$ and $l_{\text{queue}}$ is the number of fragments in the node's local sending queue.

```
Initialize C_max:
C_max = C_total;

Update1: Reset to max
if collision detected in any subchannel then
    C_max = max(C_max/2, 1);
else
    C_max = C_total;
end if

Update2: AIMD
if p% subchannels have collisions and (p > 0) then
    C_max = max(C_max × (1 − p/100), 1);
else
    C_max = min(C_max + 1, C_total);
end if
```

**Figure 7: Pseudo-code of the frequency domain backoff algorithms. *Update1* emulates the behavior of 802.11 binary exponential backoff. *Update2* uses an AIMD strategy.**

The maximum subchannel count $C_{\max}$ is updated based on the contention situation on the channel. There can be multiple update strategies. For example, we can emulate the behavior of binary exponential backoff (BEB) used in 802.11. When a collision is detected on a subchannel that a node has transmitted on, it will reduce $C_{\max}$ by half. Once all transmissions are successful, the node resets $C_{\max}$ to the total number of subchannels.

Alternatively, we can use an additive increase/multiplicative decrease (AIMD) strategy. Assume after transmission, a node detects that $p\%$ of channels have collisions. The node then decreases $C_{\max}$ by $p\%$ (multiplicative decrease). This reaction is reasonable since $p$ reflects an estimation of the contention level in all subchannels. And when all accesses to subchannels succeed, the node increases $C_{\max}$ by one (additive increase).

It is straightforward to prove that both strategies converge if all contending nodes are within a single broadcasting domain. Figure 7 shows the pseudo-code of the adaptive $C_{\max}$ adjustment algorithms. Later, we evaluate these two strategies and find that AIMD is slightly better than the reset-to-max strategy.

## 3.4 Multiple access points and two-way traffic

Until now, we have only concentrated on uplink access. Downlink access follows the same process with roles reversed. In this case, the AP will send out an M-RTS. In FICA, it is possible for the AP to transmit simultaneously to multiple clients with a single transmission burst (but different frames to different subchannels). The receiver IDs are encoded in the *receiver band* of M-RTS. All receiving clients should return M-CTS to the AP. It is necessary for the downlink transmissions to go through such a contention process since in practice there can be multiple APs located nearby on the same wide-band channel. Thus, these APs should use M-RTS and frequency-domain backoff to contend for each sub-channel for transmission. The contention result is resolved by the receiving stations and fed back to each AP by M-CTS broadcasts.

However, there is another issue in presence of two-way traffic. Assume that an AP and one client have frames to exchange. Since FICA does not use time backoff by default, both nodes may send an M-RTS simultaneously. This may cause a failure if there are no other nodes that can send back an M-CTS. Further, even if there is a third node that receives these M-RTSs and sends an M-CTS back with the contention resolution, it is still possible that both the AP and the client have been granted some winning subchannels. So, they will both transmit data simultaneously but neither of them can receive its frame due to the simplex radio used.

To address this issue, FICA separates uplink and downlink traffic by assigning different DIFS times to the AP and stations. The one with a short DIFS has priority to access the channel by sending an M-RTS earlier. To ensure fairness between uplink and downlink traffic, we use a simple dynamic DIFS assignment strategy. We assign a fixed DIFS time to all mobile stations, and the AP has two different DIFS time settings. One is shorter than the DIFS of mobile stations, and the other one is longer. The AP chooses a DIFS time based on the following rules: (1) once an AP accesses the channel with a short DIFS, it will use the long DIFS for its next access; and (2) if an AP receives an M-RTS from stations, it will use a short DIFS for its next access. Note that this simple strategy ensures the fair interleaving of uplink and downlink traffic, but not the fairness among all nodes as 802.11 currently does. In effect, it gives the AP many more chances to transmit — which may be an appropriate strategy given the asymmetrical nature of WLAN application workloads. It is an interesting open question of what sort of fairness is best provided in FICA and remains future work.

## 3.5 Discussion

We end the description of FICA with a few additional points of consideration. For the issues raised, we broadly describe potential approaches for addressing them, but in general leave an exhaustive discussion of them as future work.

**Hidden terminals.** Until now, we have only described FICA within a single broadcast domain. How does FICA operate in a network with hidden terminals, and therefore when there are multiple overlapped contention domains? Hidden terminals cannot be coordinated by carrier sensing. As a result, there is a chance for M-RTSs to collide if the misalignment of symbols from hidden terminals exceeds the long cyclic prefix. To prevent persistent M-RTS collisions, once a node misses an M-CTS or it receives an invalid M-CTS after it sends an M-RTS, the node should wait for a random time before sending an M-RTS again. This random waiting time is necessary to de-synchronize the hidden terminals. However, since M-RTS/M-CTSs are short, we believe the chance of M-RTS/M-CTSs collisions will be small.

Another issue is that when a node contends for a subchannel in multiple domains, it may receive inconsistent resolution results from M-CTSs from different APs. This will effectively prevent the node from sending on that subchannel since the received information is a mixture from two M-CTSs. Only when the node is the winner in all contention domains it participates in should it be allowed to access the subchannels (implying all returned M-CTSs contain the same resolution results for intended subchannels). We believe such behavior is reasonable and follows the general principle of wireless system designs.

**Multi-user diversity.** FICA also enables an opportunity to exploit multi-user diversity in WLANs [26]. When a node chooses subchannels to access, it may also consider the quality of each subchannel. Moreover, heterogeneous modulation methods can be applied to different subchannels to match the conditions on that specific channel band. There is already much research on resource allocation for multi-user diversity in both single and multi-channel cases [7,12,15,21], and applying these ideas in a distributed system like FICA remains interesting open challenges.

**Backwards compatibility.** It is also possible for FICA to coexist with current 802.11. Since FICA is still based on CSMA, FICA nodes will defer if they sense a transmission of 802.11 nodes, and vice versa. It is also possible to retain fairness among these two types of nodes. For example, since FICA has a fixed access pattern with the M-RTS/M-CTS handshake, it is possible to choose an appropriate DIFS time (*e.g.*, equal to half of the $CW_{\text{min}}$ of 802.11) for FICA to be fair to 802.11 nodes with random time-domain backoff.

**Cellular vs. WLAN environments.** Finally, we note that the mechanisms designed in FICA may unlikely be applicable in cellular networks because coverage of a cellular base station is large (*e.g.*, a few kilometers). Thus, the propagation time is large as well, at least an order of magnitude larger than that in a WLAN. Consequently, even using broadcasting the synchronization accuracy is too coarse (*e.g.*, several tens of microseconds). Current OFDM-based WWANs already emply a relative long cyclic prefix ($4.69\mu s$) and FFT period ($66.67\mu s$) to handle a large delay spread due to multi-path fading in the wide area. A low-precision synchronization method in OFDMA will further enlarge the symbol FFT size, adding substantial engineering complexity to control frequency offsets and undermining the ability to handle Doppler effects in a mobile environment.

# 4. SIMULATION

We implemented an event-based simulator to study the performance of FICA in large-scale wireless networks and to compare its performance with 802.11n. The simulator can model both the CSMA MAC and an OFDM PHY that supports multiple subchannels. We study the performance primarily under a single AP network with varying number of stations. We assume only collisions will cause frame reception failures, and thus we focus on the performance of the MAC design. We also focus on only the uplink transmissions (the downlink behavior in this setting is analogous), and we apply various traffic patterns in a wide 40MHz channel with high data rates. For 802.11n, we also simulate MAC Service Data Units (MSDU) aggregation, the most efficient aggregation method defined in 802.11n [1].

For FICA, we use the same values of $t_{\text{sifs}}$ and $t_{\text{difs}}$ as in 802.11. The preamble in FICA requires three OFDM symbols for single and 2x MIMO and four for 4x MIMO. Using the three-symbol preamble as an example, $t_{\text{preamble}} = 46.8\mu s$, and counting another OFDM symbol for the ACK, *i.e.*, $t_{\text{ack}} = 15.6\mu s$, the subtotal per-access MAC overhead of FICA is $157.8\mu s$. Note that although FICA uses the M-RTS/M-CTS handshake, the overhead is comparable to that of 802.11 ($160\mu s$ with minimal contention window) due to the use of the PHY signaling mechanism. Eq. 7 gives a simple model for FICA's access efficiency,

$$\eta_{\text{fica}} = \frac{t_{\text{data}}}{t_{\text{difs}} + t_{\text{rts/cts/ack}} + t_{\text{preamble}} + 3 \times t_{\text{sifs}} + t_{\text{data}}}, \qquad (7)$$

where $t_{\text{rts/cts/ack}} = t_{\text{mrts}} + t_{\text{mcts}} + t_{\text{mack}}$.

Thus, to achieve an efficiency ratio of $80\%$, we need 40 DATA OFDM symbols. For different PHY data rates, frame sizes for the same efficiency correspond to a size of 400/800/1600 bytes at 145/290/580Mbps, respectively. We use these sizes as a rule of thumb for FICA nodes to fragment upper layer frames and send each fragment on one subchannel.

**No aggregation.** In this scenario, we first disable the frame aggregation of 802.11n as a lower bound. Figure 8 shows the throughput efficiency of 802.11n and FICA with two different frequency-backoff schemes: AIMD and Reset-to-Max (RMAX) (Section 3.3.2). The scenario simulates ten concurrent nodes where each node transmits UDP traffic corresponding to 1/10 of the PHY data rate with a frame size of 1500 bytes. As expected, with a 1500-byte frame

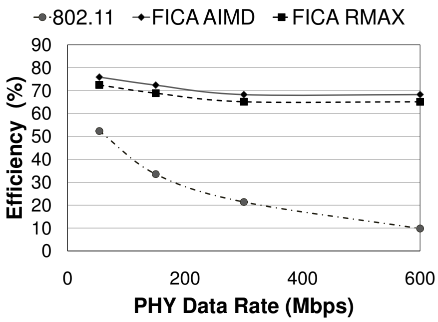


**Figure 8: Efficiency ratio of 802.11 and FICA with different PHY data rates. No frame aggregation is enabled.**

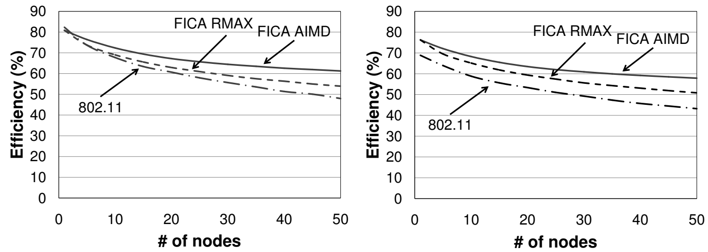


(a) (b)

**Figure 9: Full aggregation case. For 802.11, the maximal aggregated frame size is 28KB. All nodes are saturated. (a) 802.11 PHY 150Mbps; FICA 145Mbps. (b) 802.11 PHY 600Mbps; FICA 580Mbps.**

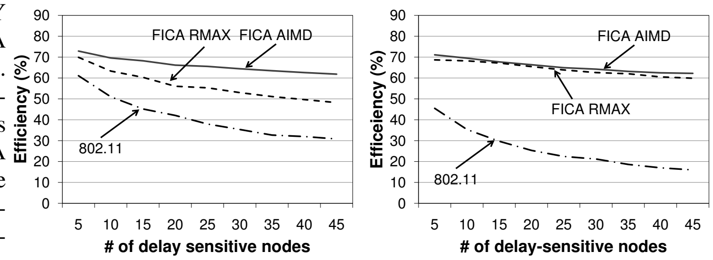


(a) (b)

**Figure 10: Mixed traffic. Five nodes are fully saturated. All other nodes have delay-sensitive traffic with a uniform distribution between 800Kbps to 5Mbps. (a) 802.11 PHY 150Mbps; FICA 145Mbps. (b) 802.11 PHY 600Mbps; FICA 580Mbps.**

current 802.11a/g rates only provide around a 50% efficiency ratio, and this ratio decreases rapidly with the increase of the PHY data rate. However, by enabling fine-grained channel access, FICA can achieve a much higher efficiency ratio in the same situation. This benefit is because different stations can access different subchannels simultaneously. Thus the per-access MAC overhead is amortized among all concurrent nodes. Also, we find that FICA AIMD has slightly better performance than FICA RMAX. As we will see in subsequent experiments, FICA AIMD consistently performs better. We hypothesize that this is because FICA AIMD adjusts $C_{max}$ much smoother compared to FICA RMAX. However, a deep analysis on the optimal frequency-domain backoff strategy remains future work.

This scenario is the worst case for 802.11n. We show this case to demonstrate how significant the MAC overhead can be at high PHY data rates, and that techniques like FICA or frame aggregation are indeed necessary for efficiency.

**Full aggregation.** Here, we show the best case of 802.11n with frame aggregation. In this experiment, all nodes are saturated so that the frame aggregation can work most efficiently. Figure 9 shows throughput efficiency with different numbers of contending nodes at two PHY data rates, 150Mbps and 600Mbps, respectively. In both cases, the efficiency of 802.11n has been significantly improved due to frame aggregation. Since all nodes are saturated, the aggregation level is very high: 12 frames (or 18KB) on average.

FICA still has slightly better performance than 802.11n even in this case, though, because FICA has slightly fewer collisions compared to 802.11n. To understand why, consider the operation of frequency domain contention. When there are many stations contending for a subchannel, if two stations happen to pick up the same subcarrier to send their signals, it does not necessarily result in a collision. A collision occurs only when the collided subcarrier is also chosen as the winner as nodes contend for subchannels. In the next contention period, all stations will pick a different random number again. This situation is unlike time-domain backoff used in 802.11: when two stations pick the same backoff slots they will eventually collide with each other.

**Mixed traffic.** Finally we evaluate a situation in between the two extremes. We have five saturated stations that always have full-sized frames to transmit. In addition, there are a variable number of nodes that have small but delay-sensitive traffic representing, for instance, video conferencing or Web browsing. We choose the load of this delay sensitive traffic uniformly from 800Kbps to 5Mbps, and the packet size from 800–1300 bytes. Figure 10 shows the efficiency results of this scenario as a function of the number of delay-sensitive nodes. With a few delay-sensitive nodes, the throughput efficiency of the network is significantly reduced for 802.11n. Since the delay sensitive flows cannot be aggregated, their access to the channel is much less efficient. Thus, the overall channel utilization is reduced. However, with FICA such nodes can request access to fewer subchannels, leaving the other subchannels for use by other nodes. Consequently, the overall network efficiency remains at a high level, improving upon 802.11 from 16% up to 4 times better at the high PHY data rates.

# 5. IMPLEMENTATION

We have also implemented the basic mechanisms of FICA using Sora, a fully programmable software radio platform based on commodity general-purpose PC architectures [24]. Our FICA implementation is based on SoftWiFi, a software implementation of the 802.11a/b/g PHY/MAC [24]. We make the following modifications: (1) we change the FFT size from 64-point to 256-point for DATA/ACK symbols and 512-point for M-RTS/M-CTS symbols; (2) we employ convolutional coding in each subchannel and decode data in each subchannel individually using the Viterbi algorithm; (3) we remove the random time-domain backoff in the CSMA MAC, and implement the M-RTS/M-CTS handshake after the channel is sensed idle.

FICA uses a PHY frame structure and synchronization algorithm similar to 802.11. A preamble precedes data symbols. The first symbol is used for symbol time synchronization (*i.e.*, finding the boundary of symbols). It employs a self-repeating pattern in the time domain so that the receiver can detect it using auto-correlation. The second symbol is used for channel estimation. To support 4x MIMO, another training symbol is needed. The last symbol encodes the Physical Layer Convergence Protocol (PLCP) header using BPSK and 1/2 convolutional coding. The PLCP header contains the modulation mode used in the following DATA symbols for the receiver to set the proper demodulating parameters.

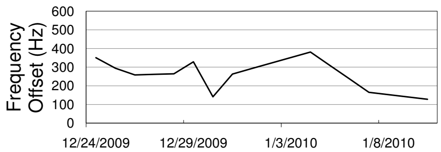


**Figure 11: Measured frequency offsets of two nodes over two weeks starting from a single calibration.**

There are two practical issues that our implementation addresses:

**Threshold setting.** With physical layer signal symbols (M-RTS/M-CTS), we need a proper energy threshold to reliably detect the bit modulated on each subcarrier. Since we target WLANs, where the signal-to-noise-ratio (SNR) is usually high, setting such a threshold is not a difficult task [5]. In FICA, we further use a self-calibration method to facilitate this threshold setting. We use blank subcarriers (carrying a bit "0") in the M-RTS/M-CTS tag band to calibrate the noise floor. Then, we choose a threshold that is a constant multiple of the noise floor. In our experiments, we find that a threshold 2–5dB higher than the noise floor works reliably in WLAN settings.

**Frequency offset calibration.** In OFDM-based multi-access networks, the frequency offset among simultaneous senders should be controlled within a certain range. Otherwise, it will undermine the orthogonality among subcarriers since the concurrent senders are not actually transmitting on orthogonal frequency due to the offset. Thus, in FICA, all nodes should have their frequency calibrated.

Fortunately, FICA can use the AP's frequency as a reference. Every station can hear the beacon from the AP and adjust its local numerically-controlled oscillator (NCO) to match the frequency of the AP. We note that frequency calibration is much simpler than time synchronization, since the frequency offset does not accumulate over time. Figure 11 shows the measured absolute value of the frequency offset of two Sora nodes over two weeks after a single calibration. The frequency offset is within a limited range of 500Hz. This range is quite small compared to the data subcarrier width ($0.63\%$ of $80KHz$), and its impact to orthogonality can be neglected in practice. Thus, we conclude that an infrequent frequency calibration process can support FICA very well.

Although Sora is the fastest programmable software radio platform that allows us to prototype FICA quickly, it still has a limitation that prevents our implementation from running in real-time. It takes too long to transfer the just-in-time modulated PHY frames from the CPU to the Radio Control Board (RCB) after receiving contention results in the M-CTS. As a result, it will miss the SIFS deadline by a small margin. To improve the situation, we will either need a faster version of Sora or take the step of moving part of the FICA implementation to hardware.

# 6. EVALUATION

In this section, we evaluate the feasibility of FICA using our prototype implementation on four Sora nodes: one serves as the AP and the others are stations. We first demonstrate how well concurrent transmissions can be coordinated in a WLAN. We evaluate the maximal symbol timing misalignment at the receiver for two concurrent transmissions coordinated by means of carrier sensing and broadcasting. We show that even with our existing software radio implementation, we can bound the maximal symbol timing misalignment within the range discussed in Section 3.1. Then, we evaluate the efficiency and reliability of detecting BAM-modulated PHY signals in M-RTS/M-CTS. Finally, we show the decoding performance of our FICA decoder for two concurrent FICA senders.

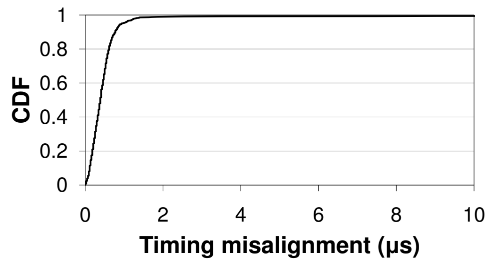


**Figure 12: CDF of symbol timing mismatch with broadcasting.**

We perform our experiments in the 2.4GHz band with 802.11b/g-compatible RF front-ends. Due to timing constraints, we have pre-stored all needed PHY frame samples on the RCB first. We also conduct the experiments late at night to minimize interference from other traffic in the same frequency band.

## 6.1 Symbol timing misalignment in a WLAN

As discussed in Section 3.1, there are two ways to coordinate concurrent transmissions in a WLAN: using a reference broadcast (*e.g.*, M-CTS) or carrier sensing on an idle channel. These two methods result in different performance requirements in aligning symbol times.

To measure the symbol timing misalignment with broadcasting, the AP broadcasts an M-CTS symbol and two stations, when they detect the M-CTS, wait for a fixed $200\mu s$ and return two different Pseudo-Noise (PN) signals. We use PN signals so that we can easily separate the two signals by cross-correlation, and precisely measure the time difference of the arrivals of the two signals at the AP by counting the number of samples between the two correlation peaks. Figure 12 shows the cumulative distribution function (CDF) of 1,000 tests. In over $95\%$ of the cases, the symbol timing difference of these two transmissions is less than $1\mu s$, and in $99\%$ cases the timing difference is less than $2\mu s$.

To characterize the maximal symbol time difference when coordinating with carrier sensing, we use the following method. After the AP sends out a broadcast symbol, one station waits for a fixed $200\mu s$ and transmit its PN signal. We adjust the waiting time of the second station incrementally from $200\mu s$ to $220\mu s$, performing around 100 tests for each value. Upon timeout, the second station performs carrier sensing first. If it senses a busy channel, it will cancel its transmission attempt and the AP will only receive one PN signal from the first station. Otherwise, if the second station senses an idle channel, it will send a different PN signal. The AP will then detect two PN signals and measure the time difference between their arrivals.

Figure 13 shows the results of this experiment. The $x$-axis shows the relative delay of the second station to the first station. The dark line shows the probability of the second station sensing the transmission of the first node. We see that if the relative delay is larger than $9\mu s$, the second node will always sense the first node's transmission and cancel its own sending attempt. The $9\mu s$ threshold reflects the turnaround time of Sora, which is about $5\mu s$, and the carrier sensing time of $4\mu s$ we used in this experiment. When the relative delay is less than $8\mu s$, the second node always senses an idle channel and concurrent transmissions occur.

The rectangular points mark the difference in the arrival time of the two signals, and the error bars indicate the minimal and maximal values measured. The timing misalignment linearly increases with the relative delay, as expected. Note that, when the relative delay is between 8–$9\mu s$, carrier sensing may not always work and results in larger variance. Overall, the symbol timing misalignment is within $10\mu s$, confirming our analysis in Section 3.1.

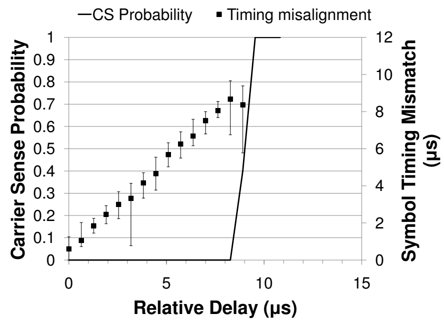


**Figure 13: CDF of timing misalignment with carrier sensing. The dark line shows the carrier-sensing probability and the rectangular points show the measured symbol timing difference.**

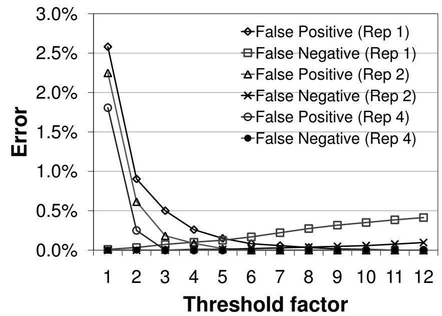


**Figure 14: Reliability of PHY M-RTS/M-CTS signaling.**

## 6.2 Reliability of PHY signaling

In this experiment, we evaluate the reliability of the PHY signaling implementation. Two nodes transmit the same M-RTS symbol simultaneously so that every subcarrier contains a superposition of the signals transmitted by the two senders. The AP detects the information on each subcarrier. Since we know what is exactly transmitted, we can detect false positives (*i.e.*, the AP detects one on a subcarrier that should be zero) and false negatives (*i.e.*, the AP detects zero where it should be one). We use the self-calibrated method to set detection threshold as described in Section 5: we measure the maximal energy $N_f$ on the blank subcarriers and set the threshold to be $\alpha N_f$, varying $\alpha$. We also experiment with different degrees of bit replication on subcarriers.

Figure 14 shows the error rates for 1,000 tests at various degrees of replication as we vary the threshold $\alpha$. With a large threshold, false positives decrease while false negatives increase slightly (although still less than $0.5\%$). Overall, though, there is quite a large space for threshold setting to provide good performance. In particular, when a bit is replicated on a few subcarriers (*e.g.*, 4), both false positive and negative rates are close to zero when $\alpha$ is in the range 3–5. We hypothesize that, with more simultaneous transmissions, the detection should be more reliable as BAM is essentially energy detection. With more transmissions superposed, the energy is additive at the receiver side. However, due to limited hardware availability we leave this for future study.

## 6.3 Demodulation performance

In this experiment, we compare the demodulation performance of FICA, where multiple nodes are allowed to simultaneously access different subchannels, to the conventional WLAN, where only a single node can access the whole channel. We have one FICA station access only the odd-numbered subchannels and another FICA station access only the even-numbered subchannels. Thus, these two stations maximally interleave their subchannels and should be more sensitive to mutual inter-subchannel interference if there is any. We fix the position of the AP and two stations and adjust the transmission power to get different signal-to-noise-ratios (SNR). For each SNR setting, we evaluate four different modulation schemes. We schedule the transmissions in the following way. For each transmission power setting, we let two stations access the channel simultaneously using FICA first, and then immediately one station transmits alone. Since these two transmissions are back-to-back, we assume their channel conditions should be similar. For each power setting, we send 1400 frames. Each frame is 400 bytes and uses one subchannel.

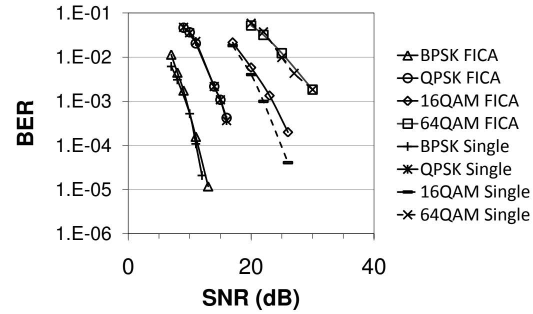


**Figure 15: Demodulation performance of FICA compared to conventional WLAN where only a single node can access channel at a time.**

We use the classic bit-error-rate (BER) to SNR plot to illustrate the demodulation performance. Figure 15 shows the results for FICA as well as the case where only one node accesses the channel. The BER value shown is measured before the Viterbi decoder; after Viterbi, most of the errors are corrected. Clearly, all curves are very close to each other, including the high rate modulations like 64QAM which are very sensitive to interference. Thus, we conclude that with FICA different nodes can transmit on different subchannels simultaneously without interfering with each other.

# 7. RELATED WORK

The application of OFDM for multiple access in WLANs is limited. In [18], Rahua, *et al.*, developed FARA that implements downlink OFDMA in a WLAN and per-subcarrier rate adaptation. But, since there is only one transmitter (the AP), symbol alignment is not an issue. In [5] and similarly in [11], OFDM has been used as a simple form of concurrent channel access. Nodes may modulate one bit of ACK information on different subcarriers after receiving a broadcast frame. However, FICA is a new PHY/MAC framework for WLANs that enables data communication over fine-grained subchannels to improve overall network efficiency.

Physical layer signaling, usually with Binary Amplitude Modulation, has been used previously to assist MAC protocols. In [8,13], busy tones are used to indicate channel occupancy to mitigate the hidden terminal problem. Recently, SMACK [5] uses a physical layer ACK, and MCBC [20] uses a PHY RTS/CTS handshake to facilitate a multi-round leader election protocol in a heavy-contention environment; after the protocol completes, the winner obtains access to the medium. FICA similarly shares the idea with SMACK and MCBC to apply PHY signaling based on simple BAM modulation, but FICA has the broader goal of enabling fine-grained channel access in high data rate WLANs.

Coordination using broadcast in local area networks has been previously exploited for time synchronization in reference-broadcasting synchronization, which provides microsecond-level synchronization precision [6]. Our results further confirm that microsecond-level coordination accuracy is practical in WLANs.

There is extensive work in the literature to improve 802.11 MAC performance by fine-tuning the backoff scheme [9, 10]. But these approaches still consider the channel as one resource unit where only one radio can work on one channel at a time. Multi-channel MAC protocols [16] have been studied to improve wireless network performance by using more orthogonal channels that are separated by guard bands. In contrast, FICA improves the performance of WLANs by increasing channel access efficiency. Using OFDM, FICA creates a fine-grained structure (*i.e.*, subchannels) for multi-access inside a wide-band channel without guard bands. Thus, a FICA node can adjust the portion of the spectrum it accesses based on its traffic demands, while other nodes can use the remaining spectrum simultaneously. This property shares some similarity to the adaptive channel width demonstrated in [4]. FICA is complementary to that work by providing a concrete means for adaptive fine-grained subchannel access in WLANs.

The inefficiency of the 802.11 MAC has also been discussed before for supporting VoIP traffic [22,27]. In [27], a TDMA approach is used to reduce the contention overhead for CSMA in 802.11. In this paper, we argue that the inefficiency of 802.11 MAC is a fundamental bottleneck as the PHY data rate increases for all traffic, not just VoIP traffic. We further argue that this inefficiency issue should be resolved by enabling fine-grained channel access.

# 8. CONCLUSION

This paper addresses the inefficiency issue of MAC protocols in current WLANs as the PHY data rate increases. The fundamental reason of this inefficiency lies in the fact that the current MAC protocol allocates the entire wide-band channel as a single resource. Such allocation become too coarse-grained for general traffic demands as the channel width or the PHY data rate increases.

We argue that this inefficiency issue should be resolved using fine-grained channel access in high data rate WLANs. We present the design of FICA, a new cross-layer design that enables fine-grained subchannel random access based on OFDM. FICA addresses challenges in both PHY and MAC design due to the asynchronous and distributed nature of WLANs. First, FICA proposes a new PHY architecture based on OFDM that retains orthogonality among subchannels solely relying on the coordination mechanisms in existing WLANs. Second, FICA employs frequency-domain contention that uses physical layer RTS/CTS signaling, and frequency domain backoff to efficiently coordinate subchannel access.

We have implemented FICA on the Sora software radio platform. With our prototype, we validate the feasibility of the FICA cross-layer design. Our detailed simulation results further indicate that FICA can outperform 802.11 with frame aggregation with up to 400% under different traffic settings.

# 9. ACKNOWLEDGEMENTS

The authors would like to thank Geoff Voelker (UCSD), Hui Liu (University of Washington), our shepherd Kyle Jamieson (UCL), and the anonymous reviewers for their valuable comments and suggestions.

# 10. REFERENCES

[1] IEEE standard for local and metropolitan area networks part 11; amendment 5: Enhancements for higher throughput. *IEEE Std 802.11n-2009*.

[2] IEEE standard for local and metropolitan area networks part 16: Air interface for fixed broadband wireless access systems. *ANSI/IEEE Std 802.16-2004*.

[3] 3GPP TS 36.201-820:. Evolved universal terrestrial radio access (E-UTRA); long term evolution (LTE) physical layer; general description.

[4] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, and P. Bahl. A case for adapting channel width in wireless networks. *SIGCOMM Comput. Commun. Rev.*, 2008.

[5] A. Dutta, D. Saha, D. Grunwald, and D. Sicker. SMACK: a smart acknowledgment scheme for broadcast messages in wireless networks. In *SIGCOMM '09.*, New York, NY, USA, 2009. ACM.

[6] J. Elson, L. Girod, and D. Estrin. Fine-grained netwrok time synchronization using reference broadcasts. In *5th Symposium on Operating Systems Design and Implementation (OSDI)*, 2002.

[7] A. Goldsmith. *Wireless Communication*. Cambridge University Press, 2005.

[8] Z. J. Haas and J. Deng. Dual busy tone multiple access (dbtma)-a multiple access control scheme for ad hoc networks. *Communications, IEEE Transactions on*, 50(6):975–985, 2002.

[9] Y. He, R. Yuan, J. Sun, and W. Gong. Semi-random backoff: Towards resource reservation for channel access in wireless lans. In *ICNP*, 2009.

[10] 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. *SITCOM Compute. Common. Rev.*, 35(4):121–132, 2005.

[11] B.-S. Kim, S. W. Kum, and R. L. Ekl. Ofdma-based reliable multicasting mac protocol for wlans. *IEEE Transactions on Vehicular Technology*, 57(5), 2008.

[12] D. Kivanc, G. Li, and H. Liu. Computational efficient bandwidth allocation and power control for ofdma. *IEEE Transactions on Wireless Communications*, 2003.

[13] L. Kleinrock and F. Tobagi. Packet switching in radio channels: Part i–carrier sense multiple-access modes and their throughput-delay characteristics. In *IEEE Transactions on Communications*, 1975.

[14] D. Kotz and K. Essien. Analysis of a campus-wide wireless network. *Wirel. Netw.*, 11(1-2):115–133, 2005.

[15] X. Liu, E. K. P. Chong, and N. B. Shroff. Opportunistic transmission scheduling with resource-sharing constraints in wireless networks. *IEEE Journal on Selected Areas in Communications*, 19, october 2001.

[16] J. Mo, H.-S. So, and J. Walrand. Comparison of multichannel MAC protocols. *Mobile Computing, IEEE Transactions on*, 7(1):50–65, Jan. 2008.

[17] J. G. Proakis and M. Salehi. *Digital Communications*. McGraw Hill, 2008.

[18] H. Rahul, F. Edalat, D. Katabi, and C. Sodini. Frequency-aware rate adaptation and mac protocols. In *Mobicom*, 2009.

[19] H. Rahul, N. Kushman, D. Katabi, C. Sodini, and F. Edalat. Learning to share: narrowband-friendly wideband networks. *SIGCOMM Comput. Commun. Rev.*, 38(4):147–158, 2008.

[20] B. Roman, F. Stajano, I. Wassell, and D. Cottingham. Multi-carrier burst contention (MCBC): Scalable medium access control for wireless networks. In *WCNC 08*, 2008.

[21] Z. Shen, J. G. Andrews, and B. L. Evans. Adaptive resource allocation for multiuser ofdm with constrained fairness. *IEEE Transactions on Wireless Communications*, 2005.

[22] S. Shin and H. Schulzrinne. Experimental measurement of the capacity for voip traffic in IEEE 802.11 wlan. In *Infocom 2007*, 2007.

[23] V. Shrivastava, S. Rayanchu, J. Yoon, and S. Banerjee. 802.11n under the microscope. In *IMC '08.*, 2008.

[24] K. Tan, J. Zhang, J. Fang, H. Liu, Y. Ye, S. Wang, Y. Zhang, H. Wu, W. Wang, and G. M. Voelker. Sora: High performance software radio using general purpose multi-core processors. In *NSDI 2009*.

[25] Q. Tang, S. Gupta, and L. Schwiebert. BER performance analysis of an on-off keying based minimum energy coding for energy constrained wireless sensor applications. In *ICC 2005*, 2005.

[26] D. Tse and P. Vishwanath. *Fundamentals of Wireless Communications*. Plenum Press New York and London, 2005.

[27] P. Verkaik, Y. Agarwal, R. Gupta, and A. C. Snoeren. Softspeak: Making VoIP play well in existing 802.11 deployments. In *NSDI*, 2009.