Rice University

Full-duplex Infrastructure Nodes: Achieving Long Range with Half-duplex Mobiles

by

Evan Everett

A Thesis Submitted
in Partial Fulfillment of the
Requirements for the Degree

Master of Science

Approved, Thesis Committee:

_____

Dr. Ashutosh Sabharwal, *Chair*
Associate Professor of Electrical and Computer Engineering, Rice University

_____

Dr. Edward Knightly
Professor of Electrical and Computer Engineering, Rice University

_____

Dr. Behnaam Aazhang
J.S. Abercrombie Professor of Electrical and Computer Engineering, Rice University

Houston, Texas
April 2012

RICE UNIVERSITY

Full-duplex Infrastructure Nodes: Achieving Long Range with Half-duplex Mobiles
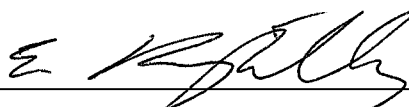
by

Evan Everett

A THESIS SUBMITTED
IN PARTIAL FULFILLMENT OF THE
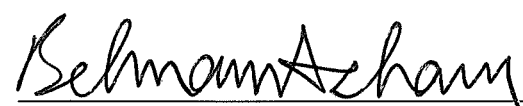REQUIREMENTS FOR THE DEGREE

Master of Science

APPROVED, THESIS COMMITTEE:

Dr. Ashutosh Sabharwal, *Chair*
Associate Professor of Electrical and Computer Engineering, Rice University

Dr. Edward Knightly
Professor of Electrical and Computer Engineering, Rice University

Dr. Behnaam Aazhang
J.S. Abercrombie Professor of Electrical and Computer Engineering, Rice University

HOUSTON, TEXAS
APRIL 2012

ABSTRACT

Full-duplex Infrastructure Nodes: Achieving Long Range with Half-duplex Mobiles

by

Evan Everett

One of the primary sources of inefficiency in today's wireless networks is the half-duplex constraint – the assumption that nodes cannot transmit and receive simultaneously in the same band. The reason for this constraint and the hurdle to full-duplex operation is self-interference: a node's transmit signal appears at its own receiver with very high power, desensitizing the receiver electronics and precluding the reception of a packet from a distant node. Recent research has demonstrated that full-duplex can indeed be feasible by employing a combination of analog and digital self-interference cancellation mechanisms. However, two glaring limitations remain. The first is that the full-duplex state-of-the-art requires at least two antennas and extra RF resources that space-constrained mobile devices may not be able to accommodate. The second limitation is range: current full-duplex demonstrations have been for ranges less than 10 m. At longer distances nodes must transmit with higher power to overcome path loss, and the power differential between the self-interference and the signal-of-interest becomes more that the current cancellation mechanisms can handle. We therefore present engineering solutions for answering the following driving questions: (a) can we leverage full-duplex in a network consisting mostly of half-duplex mobiles? and (b) can we extend the range

of full-duplex by achieving self-interference suppression sufficient for full-duplex to outperform half-duplex at ranges exceeding 100 m? In answer to the first question, we propose moving the burden of full-duplexing solely to access points (APs), enabling the AP to boost network throughput by receiving an uplink signal from one half-duplex mobile, while simultaneously transmitting a downlink signal to another half-duplex mobile in the same band. In answer to the second question we propose an AP antenna architecture that uses a careful combination of three mechanisms for passive suppression of self-interference: directional isolation, absorptive shielding, and cross-polarization. Results from a 20 MHz OFDM prototype demonstrate that the proposed AP architecture can achieve 90+ dB total self-interference suppression, enabling > 50% uplink rate gains over half-duplex for ranges up to 150 m.

Acknowledgements

First and foremost, I would like to acknowledge the collaborative contributions of Achal Sahai, Jingwen Bai, and Gaurav Patel to this work. The topology simulation results in Chapter 6 were primarily conducted by Mr. Sahai and Miss Bai and are included in this thesis with their permission. Miss Bai, Mr. Sahai, and Mr. Patel also spent long hours alongside the author writing code, preparing and running experiments, and analyzing data; their mission-critical contributions to this work are greatly valued by the author.

I would like also like to acknowledge Melissa Duarte, Chris Hunter, Michael Khayat, Tim Kennedy, Chris Dick, and Ashu Sabharwal. I have turned to each of these individuals for expert technical guidance related to the work in this thesis.

To Ashley.

Contents

List of Figures

List of Tables

Introduction

Current wireless devices operate in half-duplex mode – they do not transmit and receive simultaneously in the same band – which results in inefficient use of the resources available for communication. In cellular systems transmission and reception are orthogonalized in frequency (frequency division duplex), and in WiFi systems transmission and reception are orthogonalied in time (time division duplex). In both cases, the half-duplex constraint is a major source of wasted bandwidth. The hurdle to full-duplex operation, where devices transmit and receive simultaneously, is self-interference: the signal transmitted by a full-duplex node appears at its own receiver with very high power, overwhelming the packet-of-interest due to limited dynamic range of the receiver electronics (especially the analog-to-digital converter). Recent results [1, 2, 3] have demonstrated the feasibility of full-duplex wireless communication by suppressing self-interference via a combination of analog and digital cancellation, but two glaring limitations remain in full-duplex state of the art: resources and range. This thesis is dedicated to addressing and overcoming these limitations.

The first limitation of full-duplex is that of *resources*. The full-duplex state-of-the art requires two physically separated antennas and extra RF resources for analog cancellation at any full-duplex node. This is a problem for mobile devices, where every

millimeter of real-estate is at a premium. In this paper we identify an opportunity to leverage full-duplex even if the end-user devices remain half-duplex. We propose moving the burden of full-duplexing solely to access points, enabling an access point (AP) to boost network throughput by receiving an uplink signal from one half-duplex mobile node, while simultaneously transmitting (over the same frequency band) a downlink signal to another half-duplex mobile node.

The second limitation is *range*. Full-duplex is challenging because the self-interference is much more powerful than the received signal from a distant node. This large power differential is the *reason* that devices have remained half-duplex for so many years. As range between devices increases, the signal-of-interest is attenuated due to path loss, while the self-interference power remains the same or even gets worse if higher transmit power is needed to overcome the path loss. Thus the power differential that is the crux of the full-duplex challenge gets worse as range increases. To our knowledge, all published experimental results demonstrating full-duplex feasibility have been in line-of-sight conditions at less than 10 m range. In this paper, we propose three engineering solutions for passive suppression of self-interference: directional isolation, absorptive isolation, and polarization isolation, and introduce a novel access-point architecture for leveraging these mechanisms. Experimental evaluations shows that the architecture enables total self-interference suppression of 90+ dB, allowing full-duplex to outperform half-duplex even at ranges exceeding 100 m.

1.1 Full-duplex Infrastructure with Half-Duplex Mobiles

Our proposed scenario for leveraging full-duplex in a network with half-duplex mobiles is shown in Figure 1.1. A full-duplex access point, AP, can receive an uplink packet
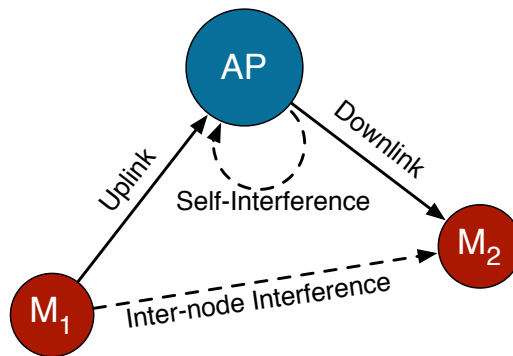
Figure 1.1: Full-duplex AP with half-duplex mobiles. Uplink suffers from self-interference while downlink suffers from inter-node interference

from one half-duplex mobile node, M_1, while simultaneously transmitting a downlink packet to another node, M_2. Such simultaneous uplink/downlink can boost network throughput without increasing the bandwidth. In current 802.11 networks, M_1 and AP would contend with each other for the channel. Assuming M_1 wins the contention, M_1 will transmit its uplink packet while AP defers, then at some later time AP will win the channel and transmit its downlink packet to M_2. However, if we can demonstrate that simultaneous uplink/downlink, as shown in Figure 1.1, is feasible at the physical layer, it will open doors to reduced contention and therefore higher overall MAC throughput. For example it could be possible to design a MAC protocol in which the AP does not have to contend with mobile users nodes for access to the medium: the mobile nodes compete among each other for deciding who gets to transmit on the uplink while the AP is transmitting on the downlink.

1.1.1 A tale of two interferences

Simultaneous uplink/downlink as shown in Figure 1.1 introduces two challenges: *self-interference* on the uplink and *inter-node interference* on the downlink. Strategies for suppressing self-interference have been proposed in [1, 2, 3, 4, 5, 6], but the suppression achieved in these works is only sufficient for establishing a short-range

($\sim$ 10 meters), line-of-sight full-duplex links. WiFi users expect to be able to wander 100+ meters from an outdoor AP, or several rooms away from an indoor AP, while maintaing high-rate service. Thus, although the problem of self-interference is not new, the problem of reaching practical ranges with full-duplex is unsolved.

In addition to the challenge of self-interference, leveraging full-duplex in the context of Figure 1.1 introduces the challenge of inter-node interference. If M_1 is to transmit to AP while AP transmits to M_2, then M_1's transmission can interfere with AP's transmission at M_2. In the case of bi-directional full-duplex, as studied in the previous works, the challenge at both receivers was self-interference, and each node can address this problem in the same way (analog and digital cancellation, etc.). But in the scenario of Figure 1.1 the uplink will suffer from self-interference while the downlink will suffer from inter-node interference, and these two interferences are very different in character. With self-interference the problem is how to communicate in the presence of a high-power interference that is partially known, whereas with inter-node interference the problem is how to communicate in the presence of a commensurate-power interference that is unknown, i.e. the traditional interference problem. We further elucidate how each of these challenge affect our goal of a deployable system to (a) achieve long-range full-duplex and (b) leverage full-duplex gains even with half-duplex mobile units.

1.1.2 Challenge 1: extending range in presence of self-interference

In practice, most manufacturers do not specify a single value for communication range, since it varies significantly from one radio environment to another. However, numerous field trials [7, 8, 9] have demonstrated that most APs can reach 100+ meters of range in line-of-sight, low-scattering environments, which reduces as the amount

of attenuation increases in different non-line-of-sight indoor environments.

Till date, all full-duplex experiments have been in line-of-sight environments with less than 10 meters of distance. The reason for such small distances can be traced to the low transmission power used in results till date; most experiments [2, 10] use a total 0–4 dBm (including transmit power and antenna gain) compared to 10–15 dBm used in WiFi equipment. One could argue that practicalities of experiments in a small lab space necessitate reducing transmit power to mimic a larger distance network. *However, size scaling by reducing powers does not apply to testing full-duplex systems as explained below.*

Consider the $M_1 \rightarrow AP$ link in Figure 1.1. Under a simple path-loss model for line-of-sight communication, where the power decay is proportional to the square of the distance, to increase the range from 10 meters to 100 meters, M_1 has to transmit with an additional 20 dB of power to achieve the same data rate. However, while the M_1 to AP distance is increasing, the distance between AP's transmit and receive antennas remains the *same*. This leads to very low signal-to-self-interference ratio (SSIR) for the $M_1 \rightarrow AP$ uplink as computed below. Assume that distances of links $M_1 \rightarrow AP$ and $AP \rightarrow M_2$ are 100 meters each. The free space path loss is 80 dB for a 100 meter link. To achieve a 20 dB SNR for the downlink, the AP transmits at 10 dBm assuming a -90 dBm noise floor; we note that these numbers are largely representative of typical WiFi hardware. Assume that the uplink mobile also transmits at 10 dBm.

The highest reported self-interference suppression is around 79 dB [5], which means that the residual self-interference power (while transmitting to M_2 in Figure 1.1) at the AP will be 10 dBm $-$ 79 dB $= -69$ dBm. The signal from M_1, after traveling 100 meters will have a signal power of 10 dBm $-$ 80 dBm $= -70$ dBm. This implies that the received uplink signal will have an SSIR of -70 dBm $+$ 69 dBm $= -1$ dB, which is too low to sustain any reasonable data rate on the $M_1 \rightarrow AP$ uplink.

Thus increasing transmit power decreases SSIR at the full-duplex AP, since increased transmit power means proportionally increased self-interference at the receiver. So while the AP $\rightarrow$ M$_2$ link can be made longer with higher power, the M$_1$ $\rightarrow$ AP link is severely impacted due to reduced SSIR. *Therefore, achieving long-range full-duplex will require significantly more self-interference suppression than has been achieved in the state-of-the-art.* In particular, in the above 100 m example, 90 -95 dB is required to produce a 10-15 dB SNR at the uplink. Thus our design goal will be achieving more than 90 dB of total self-interference suppression at the AP.

1.1.3   Challenge 2: inter-node interference

As mentioned earlier, full-duplex at the AP enables an uplink packet to be received from one mobile node simultaneously with a downlink packet being transmitted to another mobile, but as Figure 1.1 illustrates, this can leave the downlink mobile vulnerable to interference from the uplink mobile. One approach to avoid considering inter-node interference is to assume that all mobile nodes are also full-duplex (say in the future WiFi iterations) and bi-directional packet exchanges are the only supported full-duplex transmissions. That is, full-duplex is used only when both M$_1$ and AP or M$_2$ and AP simultaneously have a packet for one another.

Although this approach simplifies the design by getting rid of inter-node interference, it could severely limit the overall deployed utility of full-duplex APs. First, if AP has a packet for M$_1$ in its queue, but M$_1$ has no packet for AP (or visa versa), then full-duplex mode mode cannot be used. WiFi traffic can be highly asymmetric, and thus the opportunities for bi-directional full-duplex packet exchanges between two nodes may be limited. Second, till date (see Section 1.4 and also designs in this paper), full-duplex transceivers require physically separated transmit and receive antennas. However, most small form factor devices like smart-phones and tablets are

too space-constrained to support such antenna designs. For example, despite the success and maturity of MIMO technology, none of today's smart-phones have multiple antennas for MIMO support. Therefore it seems unreasonable to expect that full-duplex technology will change the game and motivate device designers to make room for more antennas. Future full-duplex designs may employ microwave circulators for single-antenna full-duplex, but circulators are heavy ferrite devices that will also be a burden for mobile devices. Thus, only leveraging full-duplex when *all* nodes involved are full-duplex equipped will likely limit full-duplex to backhaul links between infrastructure nodes. We therefore view the simultaneous uplink/downlink scenario of Figure 1.1 as an important opportunity for practical use of full-duplex, and thus make it the focus of this thesis.

1.1.4 A self-interference focus

The mechanisms and designs proposed in this thesis are primarily focused on the first challenge: combatting self-interference at the AP. It is important to point out that inter-node interference is dependent on the network topology, while self-interference is not. Topology analysis in Chapter 6 shows that there are indeed frequent opportunities in which inter-node interference is small enough for the benefit of simultaneous uplink/downlink to outweigh the cost of tolerating inter-node interference by treating it as noise. One could possibly design a MAC protocol to identify these opportunities.[1] Self-interference, however will always be present regardless of network topology. Thus it makes sense to focus the physical-layer effort on suppressing self-interference. MAC design for exploiting full-duplex opportunities is a future work.

---

[1]At worst, the MAC would allow simultaneous uplink/downlink for a given pair of nodes only if the inter-node interference is sufficiently small. At best, the MAC would opportunistically pair mobile nodes with weak inter-node channels for simultaneous uplink/downlink transmissions.

Coded Suppression	Bit-level	$C = \max_{p_X} I(X; Y)$
Active Suppression	Signal-level	$y[n] = h[n] * x[n]$
Passive Suppression	Physical-level	$\nabla \times \mathbf{E} = -j\omega\mathbf{B}$
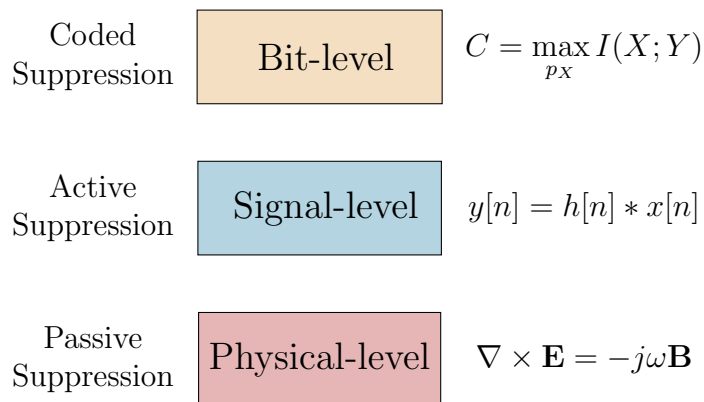
Figure 1.2: Three levels at which self-interference can be suppressed.

1.2   The Case for Passive Suppression

As was emphasized in section 1.1.2, better self-interference suppression than is currently available in the state-of-art is needed to achieve long-range full-duplex links. The question then is "from where can this extra suppression come?". Figure 2.5 illustrates the three regimes in which self-interference can be suppressed. The top is the bit-level, where the tools of information theory can be used to develop transmission schemes that exploit the knowledge that the receiver has as to the structure of its self-interference; we call such approaches *coded suppression*. The middle is the signal-level, where the tools of analog and digital signal processing can be employed to cancel the self-interference; we call this *active suppression*. The final regime is the physical-level, where the tools of applied electromagnetics can be employed to mitigate self-interference before it impinges on the receiver electronics; we call this *passive suppression*.

Most of the prior work has been focused on active suppression (see Section 1.4). Despite this large body of work, the performance of analog cancellation has seemed to hit a ceiling at around 30-35 dB achieved suppression. Recent characterizations have demonstrated that this ~ 30 dB ceiling may be due to fundamental radio limitations such as phase noise [11]. Only limited work has been conducted on coded suppression

with some preliminary progress made in [12, 13, 14], but it is currently unclear whether any significant practical gains will emerge. Passive suppression, however, has received only a cursory treatment, but we contend that passive suppression has the potential for large improvements. In particular, considering that we are moving the burden of full-duplex from user devices to infrastructure nodes, the design space for passive suppression has been opened up for us.

It is not to say that performance at the other levels has saturated, but the physical-level certainly seems to be the regime most ripe for realizing big performance gains. Thus passive suppression will be this thesis' weapon for attacking self-interference and achieving long-range full-duplex. Moreover, we want to adopt the the approach of "prevention is better than cure" with regards to the disease of self-interference. Rather than focusing on a better cure (i.e. improving cancellation performance), why not do the best we can to *prevent* the self-interference from ever coupling onto the receive signal? Therefore, the primary thrust of this thesis is smart utilization of antennas at full-duplex APs so that a large amount of passive self-interference suppression is achieved.

1.3 Main Contributions

Achieving Long Ranges: We propose an antenna design for APs to achieve significantly more self-interference suppression compared to all prior reported methods, thus significantly boosting the SSINR of long uplinks. The antenna design uses a careful combination of three passive suppression techniques - directional isolation, absorptive shielding and cross-polarization, which achieves a *passive* self-interference suppression of 60 dB in indoor environments and up to 70 dB in outdoor environments. When combined with an active per-subcarrier analog and digital cancellation scheme proposed in [4, 5], we can achieve an average self-interference suppression of

94 dB, with peak suppression near 100 dB for a 20 MHz wideband OFDM system. These numbers are near 20 dB more than the best reported numbers in the literature and allow us to achieve 100 meters outdoor ranges with our design.

Gains with Half-duplex Mobiles: For long-range communication links, SNR at M_2 is already small. In the presence of interference by M_1's transmissions, the signal-to-interference-plus-noise ratio (SINR) at M_2 will be even smaller. The key finding is that even though interference from M_1 decreases the SINR of the $AP \to M_2$ downlink, the total network capacity (sum of both link rates) with our AP design can be significantly higher for many locations over the whole coverage region of a single AP. Peak gains in sum capacity can be as high as 60% for some opportune locations of nodes M_1 and M_2, but we focus on X%-percentile-area gain to measure what fraction of the coverage area of an AP can benefit at least X% over a half-duplex AP. Under realistic path-loss assumptions, our analysis shows that at least 80% of the area can gain X=30% over half-duplex and almost 50% of the area has X=50% rate gain over half-duplex.

1.4   Prior Art

1.4.1   Performance of previous designs

Over the last 15 years, feasibility of full-duplex for short range communication has been demonstrated in [1, 2, 3, 4, 5, 6, 15]. The longest-range feasibility demonstration was 8 meters, reported in [5]. The first demonstration in [1] was shown to achieve a range of 3 meters for line-of-sight communication with 0.1 MHz bandwidth. Recent narrowband implementations [2, 3] have been able to reach a range of up to 6.5 meters. More recently, a wideband 20 MHz full-duplex system with a line-of-sight range of 5 meters [6] and a 10 MHz system with range of 8 meters [5] have been demonstrated. In contrast, we propose a design which extends the range of full-duplex by one order

of magnitude to 100 meters for outdoor line-of-sight communication.

One reason for the limited range of the previous designs is insufficient self-interference suppression. In [1], the primary self-interference mechanism is MIMO null-steering. Multiple antennas are used, both at the transmitter and the receiver, to perform beamforming such that transmit and receive vectors are spatially orthogonal. When employed along with digital cancellation (digitally subtracting off the prediction of the self-interference from the received samples), the total self-interference cancellation is no greater than 60 dB. In [10], the primary self-interference suppression mechanism is "antenna cancellation": two transmit antennas are used, with one spaced a half-wavelength farther from the receive antenna than other so that a null is produced at the receive antenna. (this really just a low-complexity sub-case of the more general beamforming approach of [1]). The authors also propose using an off-the-shelf interference cancellation chip for analog self-interference cancellation. When combined with digital cancellation the proposed design provides up to 75 dB of suppression. However both the analog cancellation and antenna cancellation mechanisms proposed are inherently narrowband, and not scalable to wideband systems such WiFi. In [6, 16] a BALUN based inversion of transmit signal is used to perform broadband cancelation of self-interference in RF. A total suppression of 73 dB using a combination of automated RF and digital cancelation was reported in [16]. In [2, 5], an additional radio chain is used to feed the receiver with the broadband cancellation waveform that is the inverse of predicited self-interference signal. [5] takes advantage of higher passive suppression by placing the receive and transmit antennas on two ends of a laptop sized device. When employed along with digital cancellation, [2] achieves 72 dB while [5] achieves 79 dB suppression, the best reported suppression to date. In contrast, our proposed design achieves an average self-interference suppression of 94 dB which is pivotal in sustaining long range full-duplex communication.

1.4.2 Previous approaches to full-duplex infrastructure

Full-Duplex has been proposed and studied for infrastructure nodes in [17, 18, 19, 20, 21] (and references therein). In [17, 18], full-duplex communication has been studied in the context of repeaters/signal boosters where the self-interference suppression is driven by antenna isolation. [17] studies an architecture where the transmit and receive antenna are separated by 5 meters. In contrast, our design has an order of magnitude smaller antenna separation. In [19], boosters are proposed which rely only on adaptive signal processing for canceling the self-interference. More recently, [20, 21] study and analyze the design of full-duplex as a relay node where power control at the full-duplex relay node is used as an approach to mitigate self-interference and improve the end-to-end achievable rate. In contrast, our infrastructure node is designed to support independent uplink and downlink traffic which is a reasonable scenario in a WiFi-like network.

1.4.3 Limitations on the existing strategies for passive self-interference suppression

Another reason for the limited range of the previous designs is that the many of them the passive self-interference suppression mechanisms, although effective is suppressing self-interference, hurt the far field coverage. The lowered increased suppression comes at the cost of lowering SNR for the signal-of-interest.

In the "antenna cancellation" technique of [3], two transmit antennas are used, one spaced a distance d away from the receive antenna, and the other spaced $d + \lambda/2$ from the receive antenna, such that the superposition of the two patterns produces a null at the receive antenna. The problem is that since periodic nulls will occur not just at the receive antenna but throughout the entire coverage zone, isolation comes at the expense of coverage. In the "device-in-the-middle" approach of [5], the

line-of-site self-interference path is attenuated by placing the transmit and receive antennas on opposite sides of the device's shell. Here the coverage is degraded due to the fact that receive antenna is "blinded" to the region opposite the device, and the same for the transmit antenna. Furthermore, a device's conductive shell can detune the antenna and cause spurious reflections that distort radiation patterns.

In contrast, we seek achieve passive suppression is such a way that coverage is not degraded. The passive mechanisms of directional isolation, absorptive shielding, and cross polarization that we propose in Chapter 2 can all be employed to either enhance coverage or be coverage-neutral at the worst. In the AP architecture proposed in Chapter 3 theses mechanisms are carefully leveraged to ensure that far-field coverage is to affected.

1.5   Organization of Thesis

In Chapter 2 we introduce three electromagnetic mechanisms for passive self-interference suppression, and in Chapter 3 a general antenna and RF architecture for full-duplex access points is proposed for leveraging these three mechanisms. In Chapter 4, the performance of the proposed passive suppression mechanisms is evaluated by directly measuring the the amount of suppression achieved for different configurations. Chapter 5 presents results from a WARPLab prototype that quantifies the total suppression achieved with all mechanisms in place (both passive and active) and demonstrates gains in uplink rate (over half-duplex) at ranges exceeding 100 m. In chapter 6 the full-topology (both uplink and downlink) performance is evaluated via data-driven simulations aimed at capturing the impact of inter-node interference on the downlink as a function of the geographic location of the nodes. Chapter 7 summarizes and concludes the thesis.

Mechanisms for Passive Self-Interference Suppression

In this chapter, we introduce three mechanisms that can be leveraged for passive self-interference suppression at infrastructure nodes: directional isolation, absorptive shielding, and cross-polarization. This chapter provides a general background on each of the mechanisms: the physics for why each can provide passive self-interference suppression is described. The specific strategy for utilizing the three mechanisms at an access point is saved for Chapter 3.

2.1 Directional Isolation

Many commercial access points attain uniform coverage via an array of directional antennas, as depicted in Figure 2.1, as opposed to a single omnidirectional antenna [22]. A directional antenna architecture provides two advantages for traditional half-duplex access points. The first is that by selecting the antenna pointed in the direction of the client being served,[1] an AP can create a higher SNR link with that client, since

---

[1] Choosing the right antenna is usually achieved by listening for control packets (such as the request-to-send in the uplink case, and the clear-to-send in the downlink case) on all antennas, and comparing received power levels.

each individual directional antenna has a higher gain than an omnidirectional antenna. The second advantage is WiFi sectoring. By allocating each antenna to a separate WiFi channel, the AP creates several orthogonal sectors. Although sectoring does not increase the spectral efficiency at any one AP (added capacity is from added bandwidth), it can increase the net spectral efficiency of a network with many APs. Sectoring directionalizes interference such that another AP that uses the same set of channels can placed closer to the the first AP without causing interference than would be the case if sectoring were not employed [23]. Thus sectoring improves spectral efficiency my making the frequency-reuse pattern denser.
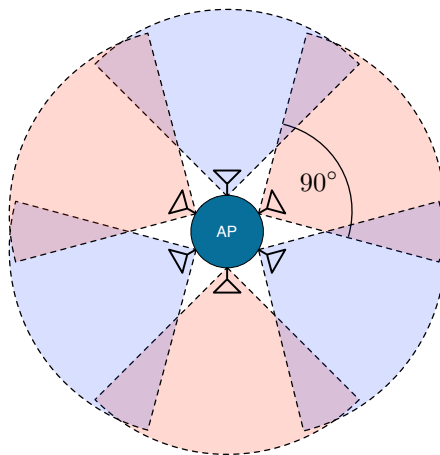
Figure 2.1: Access point with an array of directional antennas for 360° coverage.

Beyond the above two conventional benefits, Figure 2.2 shows how directionality can isolate a receive antenna from the transmit antenna's interfering signal. If, as in Figure 2.2(a), a single omni-directional transmit antenna is tasked with providing downlink coverage in all directions, and a single receive antenna is providing uplink coverage in all directions, then the transmit antenna will radiate directly onto the receive antenna, producing severe self-interference. In contrast, if an array of direc-

tional antennas is employed as in Figure 2.2(b), then the transmit energy is always directed *away* from the receiving antennas on the AP. Similarly the receive pattern is pointed *away* from the transmit antenna for uplink flows. Thus the coupling between transmit and receive antennas could be greatly reduced with directional antennas.

(a) Omni antennas   (b) Directional antennas
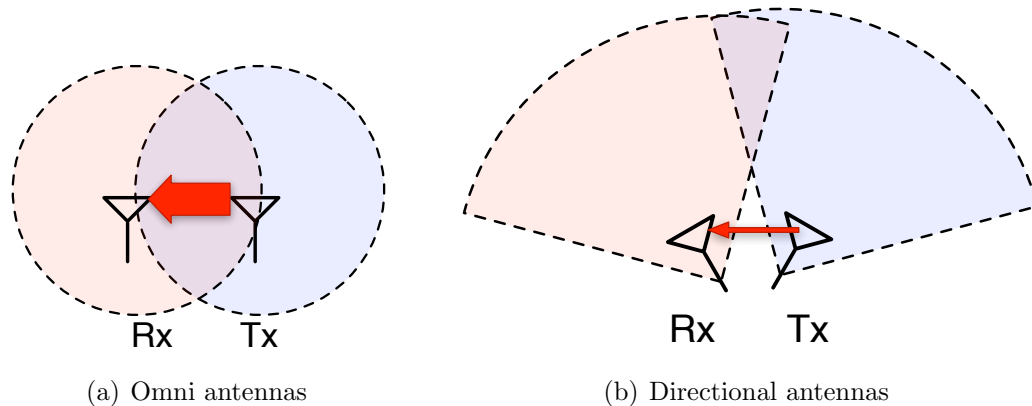
Figure 2.2: Directional Isolation. When omnidirectional antennas are used (*a*), the Tx antenna radiates directly onto the receive antenna, producing severe self-interference. In contrast, directionality (*b*) prevents the Tx antenna from radiating across the Rx antenna, suppressing self-interference.

However, just like practical filters are never perfectly rectangular, practical antennas are never perfectly directional. The cones depicted in figure 2.2(b) represent the 3dB antenna beamwidth: the span of angles for which the antenna gain is within 3dB of its max gain. Outside of the 3dB beamwidth, the gain is not zero, but decays gradually. This means the transmit antenna will still have non-zero gain in the direction of the receive antenna, and likewise the receive antenna will have non-zero gain in the direction of the transmit antenna. Thus perfect isolation is not achieved, and a residual self-interference signal will be indeed be present at the AP's receive antenna.

2.2   Absorptive Shielding

For client devices such as smartphones or laptops, user experience is paramount and antenna performance is secondary, but an access point can be designed from the ground up for efficient antenna utilization. With this in mind, it is perfectly reasonable to consider structurally shielding the access point antennas from one another for improved full-duplex performance.

Figure 2.3: When used in conjunction with directional antennas, absorptive shielding can provide self-interference isolation without obstructing the intended coverage zones.

The idea of "device-in-the-middle" structural shielding was proposed in [5, 24]. The idea is that when placing omnidirectional antennas at opposite sides of a device's electronics, passive isolation is achieved since the line of sight path from the Tx antenna to the Rx antenna is strongly attenuated. The disadvantage to this approach is that the isolation comes at the expense of coverage by potentially altering *far-field* antenna pattern. The transmit antenna cannot transmit efficiently to nodes located on the hemisphere opposite the device. In short, the design in [5, 24] alters both the near- and far-field patterns. However, when directional antennas are used at the access-point, shielding can indeed be placed between antennas to suppress self-interference *without* reducing coverage and altering far-field patterns. Figure 2.3 illustrates that when shielding can be placed between adjacent directional antennas

to reduce the side-lobes that cause self-interference, but without obstructing the antenna's intended radiation pattern.

This leads to the question of "what sort of shielding should be placed between the antennas to provide isolation?". Conductors such as aluminum and copper are commonly used to provide electromagnetic shielding. So one may be led to place metallic sheets between antennas, but such conductive materials provide isolation by *reflecting* the incident energy, as illustrated in Figure 2.4(a) which will alter the intended coverage pattern as follows. A conductor in the near-field of an antenna will couple with the antenna, detuning it (i.e. making it a less efficient radiator) and unpredictably distorting the coverage pattern. Moreover, even if the conductor is in the far-field of the antenna, it will reflect the energy back into the desired coverage area producing an interference pattern. What is needed are materials that provide isolation not by reflecting electromagnetic energy but by *absorbing* it, as illustrated in Figure 2.4(b) Thankfully, such materials are readily available.

(a) Reflective shielding   (b) Absorptive Shielding

Figure 2.4: Self-interference shielding. A metallic shield (*a*), provides effective shielding, but distorts the radiation pattern of the transmit antenna. An absorptive shield (*b*) suppresses self-interference without disturbing antenna patterns.

RF absorber material is commonly used to line anechoic chambers, reduce radar visibility in defense equipment, and control electromagnetic interference (EMI) in high-frequency electronics. RF absorbers work by using lossy materials to dissipate electromagnetic energy as heat. According to [25], there are two fundamental types of RF absorber: *standing-wave absorber* and *free-space absorber*. Standing-wave ab-

sorber is for suppression of unwanted radiation that is coupled to a structure, such as the walls of a waveguide or a coaxial transmission line. Free-space absorber is for suppression of waves propagating in the "open-air". Free-space absorber is the obvious choice for shielding an Rx antenna from a Tx antenna at a full-duplex AP, hence in this thesis the usage of free-space absorber alone is studied. However in practical devices, self-interference may couple from the Tx antenna to the Rx antennas not only via free-space radiation, but also via standing waves excited along the structure of the device. For example, in a laptop the transmit antenna may excite currents in the laptop's conductive shell that couple to the receive antenna. Characterizing which coupling mechanism (free-space vs. standing-wave) dominates self-interference in common form-factors, and designing absorber solutions based on the results may be a fruitful area of future study.

Free-space absorbers may be further classified as either (a) resonant absorber or (b) broadband absorber. A resonant absorber can provide huge amounts of absorption over a narrow band, but in order to suppress of a 20 MHz wideband OFDM signal, broadband absorber is the only choice. Broadband RF absorber functions somewhat like an "air resistor." It is made from a polyurethane foam embedded with carbon particles that make the material conductive but lossy (just like a they do in a common carbon-composite resistor). The incident electric field excites currents in the conductive carbon foam, but encounters $I^2 R$ losses as the incident energy is dissipated as heat by the carbon [25]. But making the material lossy is not sufficient for creating an effective absorber. The design of a broadband absorber is a tradeoff between *loss* and *impedance match*. When an electromagnetic wave encounter a step in characteristic impedance, reflection occurs, and a lossy, carbon-loaded foam will have much different characteristic impedance than the 377 Ω of free space. Thus a gradient is needed to transition from free space to lossy material so that the wave can
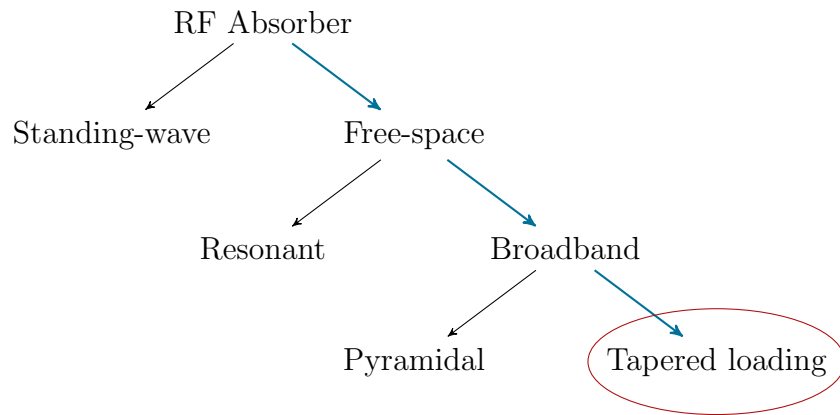
Figure 2.5: A Taxonomy of RF absorber materials. The blue path signifies design choices appropriate for self-interference in full-duplex WiFi systems.

"make it" to the lossy material and be dissipated as heat.

There are two ways that this impedance gradient is realized (a) pyramidal structure and (b) tapered loading [25]. Pyramidal absorber is the kind one is accustomed to seeing in anechoic chambers. The pyramidal cones provide a gradual transition from air to absorber, thus producing a nice impedance gradient. Pyramidal absorber provides higher absorption levels than its tapered loading counterpart, but the pyramidal structure must have thickness on the order of a wavelengths [25]. At 2.4 GHz pyramidal absorber may need to be as thick as 1 ft., too big for practical access-points. A tapered loading absorber is a simple slab of material, but where the lossy carbon particles are applied in a gradient with low concentration at the interface and high concentration at the back. The absorption performance improves with increasing slab thickness, but the slab does not have to be order-of-wavelengths thick to be effective. Thus free-space, broadband, taper loading absorber is the most reasonable design choice for absorptive shielding for passive self-interference suppression.

2.3 Cross-polarization

The *polarization* of an electromagnetic wave can be loosely defined as the direction in which the electric field vector is oscillating [26]. In a vertical dipole antenna (assume the z axis is vertical), electric current oscillates back and forth along the length of the dipole producing an electric field that oscillates in the z direction as it propagates. Hence vertical dipoles produce vertically polarized radiation. If a horizontal dipole antenna (say in the y direction) attempts to receive this vertically polarized wave, it will be unsuccessful. The electrons in the horizontal dipole are spatially constrained to move only in the y direction, but the vertically varying electric field can only exert a force on the electrons in the z direction. Hence the vertically polarized wave does not excite any current in the horizontal dipole. One would say that the vertical and horizontal dipoles are *cross-polarized*: they are orthogonal in "polarization space."

Figure 2.6: Four basis polarization states: vertical, linear, right hand circular, and left hand circular.

In addition to linearly polarized radiation (such as vertical, horizontal, or any angle in between) antennas can also produce circularly polarized radiation.[1] In a circular polarized wave, the tip of the electric field vector rotates around the axis of propagation, as opposed to just oscillating up and down in a plane parallel to to the axis of propagation as in the linear case. Circular polarization comes in two flavors depending on the direction of the electric field's rotation: right hand circular polarization (RHCP) and left hand circular polarization (LHCP). Circular polarization is

---

[1]In general the polarization of a time-harmonic wave is elliptical, of which linear and circular are special cases.

sometimes used in mobile wireless communication because coupling between circularly polarized antennas is orientation independent. Two RHCP antennas can receive each other's signal with no polarization loss, regardless of the rotational orientation between the two antennas. Linearly polarized antennas must be aligned in order to avoid polarization loss.

(a) Orthogonal polarization pairings

(b) Half-power polarization pairings

(c) Full-power polarization pairings
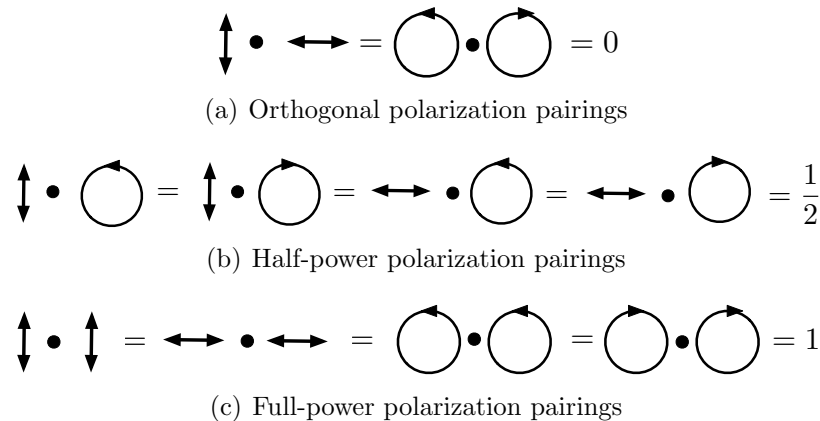
Figure 2.7: Polarization efficiency all possible combinations of basis states.

Figure 2.6 shows the four basic antenna polarizations. Most commercial antennas are designed to produce one of these four polarizations. *Polarization match factor* is the fraction of the total power of an electromagnetic wave that can be received by an antenna of a given polarization [26]. If the transmitted wave has the same polarization as the receive antenna, then the polarization match factor is one, and no power is lost, as is shown in Figure 2.7(c). Conversely if the transmitted wave and received antenna are cross-polarized, (eg. VP→HP or LHCP→RHCP), then the polarization match factor is zero as shown in Figure 2.7(a), and zero power is transferred between the two antennas. Figure 2.7(b) shows that any linearly polarized antenna can receive half the power of a circularly polarized wave, and vice versa for a circular antenna receiving a linear polarized wave. This can be explained by the fact that a circularly polarized wave can be described mathematically as the superposition of a horizontally polarized wave with a vertically polarized wave out of phase by 90°. A vertically

polarized antenna "catches" the vertical part, but not the horizontal part, and thus receives half the power of the incident wave.

In typical wireless applications, the goal is to have a high polarization match factor, so that the SNR for a link is a high as possible. However, mitigating self-interference, we can turn cross-polarization around to our advantage and aim for a *low* polarization match factor, so that the self-interference is passively suppressed.

(a) Co-polarized Tx and Rx

(b) Cross-polarized Tx and Rx

Figure 2.8: If the uplink and downlink antennas have the same polarization (a), self-interference is strong, but if uplink and downlink are cross-polarized, passive self-interference suppression is achieved.

Figure 2.8 shows how cross-polarization can be utilized for full-duplex self-interference suppression at an infrastructure node such as an AP. By transmitting on one polarization and receiving on an orthogonal polarization, the self-interference channel is cross-polarized, and hence self-interference is passively suppressed. For instance, an AP could transmit with horizontal polarization, and receive with vertical polarization as depicted in Figure 2.8(a). In section 3.3 we address the issue of how mobile devices communicate that has uplink and downlink on orthogonal polarizations. In short, in NLOS environments, scattering leads to the uplink signal becoming depolarized by the time it reaches the AP and the downlink signal becoming depolarized by the time it reaches the AP, hence mobile nodes are agnostic to the polarization configuration at the AP. In line-of-sight environments the mobiles will need to match the AP's polarizations, either by using a single circularly polarized antenna (which can couple with both H-pol uplink an V-pol downlink with 3 dB loss) or a dual-polarized antenna.

2.4   Summary

We have introduced three mechanism for isolating the transmit and receive antennas from each other for passive self-interference suppression. Directional isolation directs transmit energy away from receive antennas, absorptive shielding converts a large portion of the transmitted self-interference power into heat before it impinges on the receive antenna, and cross-polarization puts the self-interference in a polarization state that is orthogonal to the receive antenna. Each of these mechanisms are imperfect: directional antennas have side-lobes, RF absorber has leakage, and practical antenna never produce perfectly linear polarizations, hence perfect orthogonality from cross-polarization is impossible. Moreover, all these mechanisms are vulnerable to reflected self-interference. Obviously, both directional isolation and absorptive shielding are ineffective in suppressing a wave transmitted in the intended direction, but reflected back to a receive antenna. Reflections also change a wave's polarization, so that even if the transmit and receive antennas are perfectly cross-polarized, the reflected self-interference may not be cross-polarized to the receive antenna. Nonetheless, it is our hope that the combination of these imperfect mechanisms will lead to a much lower-power self-interference, such that long-range (100+ m) full-duplex can become practical. If the reflected components of the the self-interference become dominant, the the passive mechanisms will have done their job.

In the following chapter we discuss how the above three mechanism can be utilized in practical access point. In Chapter 4, measure the suppression achieved by each of the mechansims and different combinations of the mechanisms.

Architecture for Full-duplex Access Points

We now present a general design strategy for single-channel full-duplex APs. While our overall design includes both passive and active suppression mechanisms, according to our "prevention is better than cure" approach, the architecture especially targets passive suppression of self-interference before it hits the receive antenna.

3.1   Antenna Architecture

We propose a novel antenna architecture for an AP, shown in Figure 3.1 for passive suppression of self-interference while providing complete 360° coverage to the mobile devices. The general architecture, denoted by (N, θ_B), is a circular array of N directional antennas of beamwidth of θ_B, each of which can be used for Tx or Rx. 360° coverage is ensured as long as $N\theta_B > 360°$.

The goal of the proposed antenna architecture is to systematically leverage the three passive suppression mechanism introduced in Chapter 2: (i) directional isolation, (ii) absorptive shielding, and (iii) cross-polarization, as explained below.

Figure 3.1: (N, θ_B) full-duplex AP antenna architecture, in this figure N = 6, $\theta_B = 90°$

3.1.1 Leveraging directional isolation

If two omnidirectional antennas are tasked with providing coverage in all directions (one for uplink and one for downlink), then the transmit antenna will radiate directly onto the receive antenna, producing severe self-interference. The topology of Figure 1.1, however, emphasizes that when the AP participates in a full-duplex transmission with a pair of half-duplex mobiles, the uplink and downlink will in general be in different directions. If this *directional diversity* is exploited by an array of directional antennas as in Figure 3.1, then the transmit energy is always directed *away* from the receiving antenna on the AP, and similarly, the receive antennas are pointed *away* from any possible transmit antenna. Thus, the self-interference is passively suppressed from this directional isolation.

To leverage directional isolation, the access point will be equipped with N directional antennas of beamwidth[1] θ_B. These antennas will be circularly arrayed around the access point as shown in Figure 3.1. Increasing the directionality of the an-

---

[1]Throughout the paper when we use the term "beamwidth" we mean the 3 dB beamwidth - the span of angles for which the gain of the antenna is within 3 dB of its maximum gain.

tennas (i.e. decreasing θ_B), will improve directional isolation and lead to better self-interference suppression, but will require more antennas at the AP to ensure complete coverage. In general, $N\theta_B > 360°$ is required to ensure that any half-duplex mobile in range of the access pointed has at least one AP antenna pointed in its direction. One can achieve increasing self-interference suppression while maintaining coverage by letting θ_B get smaller, and increasing the number of antennas such that $N\theta_B > 360°$ is still satisfied. However, the size constraint of the access point puts a lower limit on θ_B, and an upper limit on N. The size of the access point will not grow linearly with decreasing θ_B, but quadratically, for decreasing θ_B not only increases the required number of antennas, but also increases the required size of each individual antenna, since the directivity of an antenna is proportional to its size. We want to limit our discussion to practical AP sizes, less that 50 cm for an indoor AP and 70 cm for outdoor AP. Therefore in our implementation experiments, we use a conservative configuration of $N = 6$, $\theta_B = 90°$. In which case there is 30° of overlap between each of the antenna's pattern, so that uniform coverage is conservatively ensured.

Although $N\theta_B > 360°$ is a sufficient condition for ensuring coverage to any one mobile node, simultaneous uplink/downlink transmissions can only be performed if the uplink mobile and downlink mobile are in different "sectors," since the full-duplex state-of-the art requires separate antennas for transmit and receive. Similarly, $N\theta_B > 360°$ is not sufficient for supporting a bi-directional transmissions with a full-duplex mobile device,[1] since this would require two AP's antennas (one for Tx and one for RX) to be pointed at the mobile. Thus $N\theta_B > 720°$ is required for the AP to support (a) simultaneous uplink/downlink to HD nodes located in the same sector, and (b) bidirectional full-duplex with full-duplex equipped devices.

This architecture differs from cellular sectoring in that the N antennas are meant

[1]this is not the focus of this paper, since this would require, but we include the discussion for completeness.

to be interfaced to a single RF front end via an RF switch, as opposed to cellular sectoring in which N antenna is connected N independent RF front ends, each allocated to orthogonal frequency bands.

3.1.2   Leveraging absorptive shielding

Absorptive shielding is achieved by placing RF absorber material in between each of the antennas as shown in Figure 3.1. The absorptive material is positioned such that the absorber obstructs the direct path between adjacent antennas without obstructing the coverage zone. Free-space RF absorber designed for antenna isolation is commercially available from many vendors. The performance of the absorber affected by its thickness in wavelengths: lower frequency requires thicker absorber to achieve the same isolation. One absorber supplier [27] offers sheets of RF absorber ranging in thickness from 0.25 inches to 4 inches. Only thicknesses exceeding 1 inch provide appreciable isolation. For the 1-inch absorber, the suppression is 15 dB at 2.4GHz, and the highest suppression is 23 dB, for the 4-inch absorber. These thinknesses are quite reasonable even for compact access-points. If multiple antenna can be fit onto the AP, it should not be a problem fitting absorber between them.

3.1.3   Leveraging cross-polarization

Now we we turn to the question of how cross-polarization can be leveraged within the (N, θ_B) architecture. One option would be to have alternating polarizations among the N antennas: every other antenna would be vertically polarized and the rest horizontally polarized, such that any adjacent pair is cross-polarized. But what if the antenna in whose sector the downlink mobile resides and the antenna in whose sector the uplink mobile resides have the same polarization? Such an alternating-polarization configuration would only *partially* leverage cross-polarization. What we

would like is for *each* of the N antennas to be able to use either polarization, so that *any* Tx/Rx pair can be cross-polarized. This can be accomplished by using dual-polarized antennas at the AP.

We proposed that each of the N antennas at the AP will be dual-polarized so that any antenna can use vertical *or* horizontal polarization. Each antenna will use horizontal polarization when it is transmitting and vertical polarization when it is receiving. In this case any pair of antennas involved in a full-duplex transmission will be cross-polarized. A dual-polarized antenna is a single antenna with two-ports: one port for each of two possible polarizations. The most common dual-polarized antennas have one port for horizontal polarization and the other for vertical. Horizontal/vertical dual-polarization is often realized by having one port feed the antenna at a location that produces horizontal polarization and the other port feed the antenna at a location that produces vertical polarization.

(a) Uplink from M1 and downlink to M2 (b) Uplink from M1 and downlink to M3
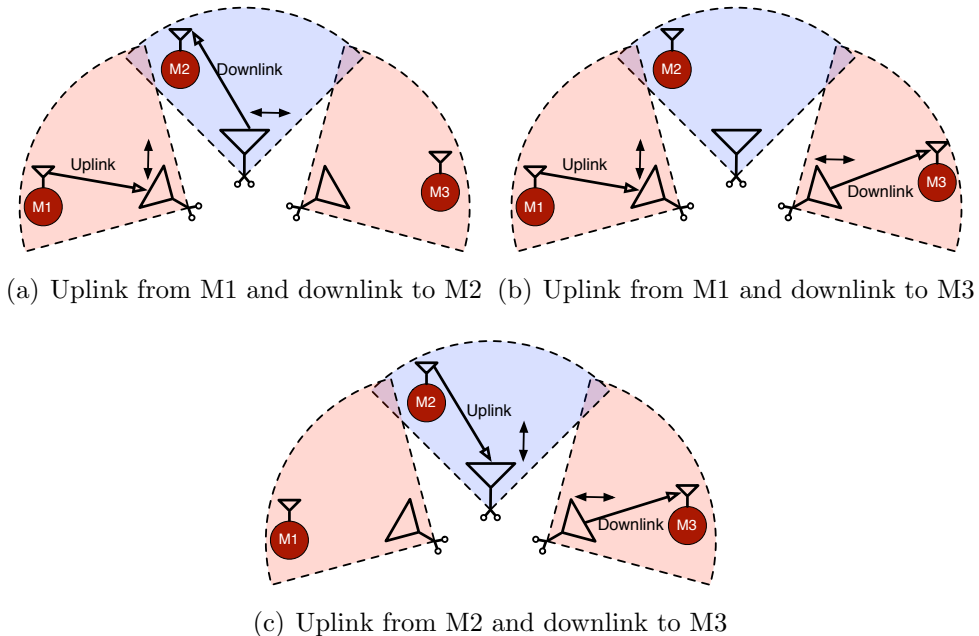
(c) Uplink from M2 and downlink to M3

Figure 3.2: If dual-polarized antennas are used, then any antenna can be either a transmitter or receiver, but always with uplink/downlink cross-polarized.

Figure 3.2 illustrates the advantage of dual-polarized antennas. Consider three

mobile nodes, each of which resides in the coverage zone of a different AP antenna. Figure 3.2 illustrates three possible full-duplex uplink/donwlink configurations. Note that if the antennas have only a fixed polarization, then there is no way that the transmit and receive antennas can be cross-polarized in all three cases. For example, if the left and right antennas are horizontally polarized and the middle is vertically polarized, then cross-polarization can be achieved in (a) and (b) but not in (c). However, if the antennas are dual-polarized, then the middle antenna can transmit with horizontal polarization in (a), but receive with vertical polarization in (c). Thus by using dual-polarized antennas, and requiting every transmission to be horizontally polarized and every reception to be vertically polarized, Tx/Rx cross-polarization is always ensured. The RF architecture for interfacing the transceiver electronics to each of the N dual-polarized antennas is discussed in the following section.

The design choice of having the AP always transmit with horizontal polarization, but receive with vertical polarization has the potential to impact the design of the antennas at the mobile devices. We discuss the mobile design implications of cross-polarization in Section 3.3

3.2   RF Architecture

The RF architecture for the full-duplex AP is shown in Figure 3.4. The primary features are (i) smart antenna switching for utilizing the direction and polarization degrees of freedom offered by the (N, θ_B) antenna architecture, and (ii) analog and digital self-interference cancellation for "cleaning up" self-interference that cannot be suppressed passively.
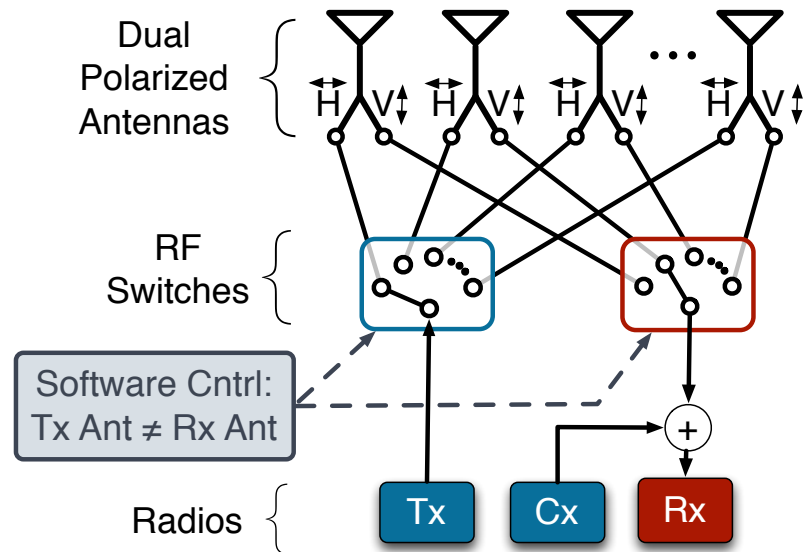
Figure 3.3: RF architecture for full-duplex AP

3.2.1   Antenna switching

Since we are constraining our focus to SISO APs, we consider the case of a single transmit path and single receive path. Dual-polarized antennas such as [28] have two ports: one for vertical polarization (V-pol) and the other for horizontal polarization (H-pol). An RF switch in the transmit path allows transmission on any one of the H-pol ports of the N antennas, as seen in Figure 3.4. Similarly, an RF switch in the receive path allows reception on any one of the V-pol ports of the N antennas. Thus Tx and Rx will always be on orthogonal polarizations.

The switches are under software control, and higher-layer information is used to associate each mobile in the network with one of the AP antennas, such that the optimal antenna for each transmission is selected. Such software-controlled switching among directional antennas at the AP has been extensively studied in works such as [29], and is therefore not addressed in this work. We instead study the worst case of simultaneous transmission and reception on adjacent antennas. The only difference between our approach and that of previous work on directional switching is we allow

simultaneous transmission and reception on separate antennas. Note, however, that we do *not* allow the switches to connect both the receive chain and the transmit chain to the same antenna. We will study the impact of this constraint in Section 5.

3.2.2   Active cancellation

For active self-interference suppression, we adopt the wideband analog and digital cancellation mechanisms demonstrated in [4, 5]. In analog cancellation, the over-the-air self-interference channel is estimated at the beginning of each packet from OFDM pilots. These per-subcarrier channel estimates allow the AP to craft a wideband cancellation waveform that is the inverse of the predicted self-interference. The cancellation waveform is transmitted over a wire using a dedicated cancellation (Cx) radio and is combined with the over-the-air received signal to cancel self-interference prior to it impinging upon the analog-to-digital converter, thus avoiding A/D quantization issues. Digital cancellation is implemented by estimating and subtracting, at baseband, the residual self-interference left after analog cancellation.

3.3   Impact on Mobile Nodes

Leveraging directional isolation and absorptive shielding at the AP, comes for "free" at the mobile nodes: the RF and antenna architecture at the mobiles are agnostic to these mechanisms. Leveraging cross-polarization, however, does require us to consider the impact at mobile devices. In the cross-polarization strategy we have proposed, uplink to the AP is vertically polarized and downlink is horizontally polarized.

3.3.1 Impact of cross-polarization in NLOS

Reflection of a electromagnetic wave changes its polarization. In and indoor non-line-of-sight (NLOS) environments, multiple reflections lead to signals between nodes becoming *de-polarized*: the received wave is the supper position of many randomly polarized components. Thus for indoor applications, a single-antenna mobile with any arbitrarily polarized antenna will be able to both receive from a horizontal polarized antenna at the AP and transmit to a vertically polarized antenna equally well. Thus Tx/Rx cross-polarization at the AP does not affect antenna design at the mobile, nor does it hurt the mobile's SNR. In indoor, NLOS conditions, leveraging cross-polarization at the AP also comes for "free" at the mobile devices. Line-of-sight conditions, however, offer both a challenge and an opportunity at the mobile devices.

3.3.2 Impact of cross-polarization in LOS

In line-of-sight conditions, however, Tx/Rx cross-polarization at the AP does impact the antenna design at the mobile node. If a mobile node has only a single vertically polarized antenna, then it will have poor SNR when receiving the AP's downlink; likewise a node with a single horizontally polarized antenna will have poor SNR on the uplink. The ideal solution would be for the mobile nodes to also use dual-polarized antennas as shown in Figure 3.4(a). Dual-polarized antennas are readily available in form factors suitable for mobile phones [30]. In this configuration, all uplink signals in the network would be vertically polarized, and all downlink signals horizontally polarized. The downside to this strategy is that it requires the mobile node to be in a stable orientation, or for the mobile device to track it's orientation such that effective polarization switching can be performed. For larger WiFi devices such as a laptop, for the orientation is relatively static and predictable, and the configuration of Figure 3.4(a) is quite feasible. For smart phones, a wealth of orientation sensors such

as compasses and accelerometers make tracking orientation quite feasible. Just like this sensor information was leveraged for mobile phone antenna-pattern switching in [31], so it could be leveraged for polarization switching.

(a)   Dual-polarized half-duplex mobile
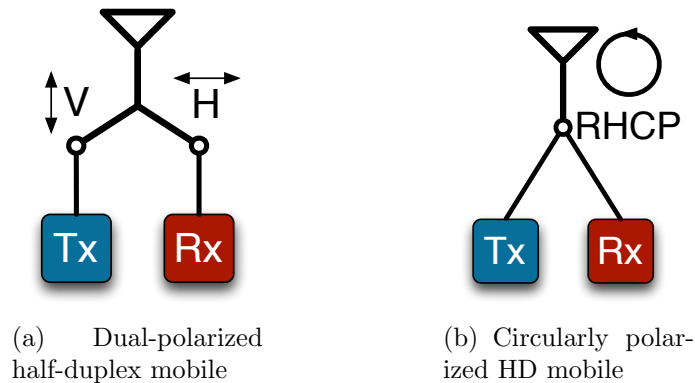
(b) Circularly polarized HD mobile

Figure 3.4: RF Architecture at AP and mobile for leveraging polarization diversity, and for switching among directional antennas. Corresponding single-antenna architecture at the half-duplex mobile.

The dual-polarization strategy of Figure 3.4(b) has an advantage in addition to providing high SNR for both uplink and downlink: suppressing *inter-node* interference. Since the mobile nodes transmit and receive on orthogonal polarizations, the interference from mobile A at mobile B introduced by simultaneous transmissions is now greatly suppressed. Thus, cross-polarization of the uplink from the downlink affords two simultaneous benefits: suppression of self-interference at the AP, and suppression of inter-node interference at the the downlink mobile.

However, in low-complexity devices in which the orientation problem cannot be solved, a simpler solution is for each of the mobile nodes to use a single circularly polarized antenna as shown in Figure 3.4(b). In this case the single antenna at the mobile node can both transmit to the AP's horizontally polarized uplink antenna and receive from the AP's vertically polarized downlink antenna, but with a 3 dB power loss in each case since the polarization efficiency is now 1/2 (see Figure 2.7(b)). In this configuration, the self-interference suppression benefits of cross-polarization at

the access-point remain, but interference from node A at node B can now be very strong, a problem we will address in the sequel.

Evaluation of Passive Suppression Mechanisms

In this chapter and in Chapter 5, we seek to evaluate the ability of the proposed full-duplex AP architecture to overcome the problem of self-interference. That is, how well can the AP receive an uplink signal from a distant node while simultaneously transmitting? We will perform this evaluation in two steps. In the first step, discussed in this chapter, we will evaluate the performance of only passive suppression using a network analyzer. The passive evaluation will also identify an (N, θ_B) antenna architecture that balances the tradeoff between number of directional elements and amount of suppression. Then in Chapter 5 will use this antenna architecture and evaluate the complete RF architecture with analog and digital cancellation for a 20 MHz OFDM modulation.

The passive suppression measurements were performed using an Agilent N5224A general-purpose network analyzer (PNA), in a $55 \times 75 \times 65$ ft. open room (Martel hall) to characterize multiple realizations of of the (N, θ_B) antenna architecture proposed in Section 3.1.

Figure 4.1: Passive suppression vs. separation angle

4.1 Suppression from Directional Isolation

To characterize the impact of directional isolation on passive self-interference suppression, two directional elements were used, a 90° element L-Com HG2414DP [28] and a 60° element L-Com HGV-2406U [32]. For comparison against an omnidirectional antenna the L-Com 6409 [33] was used. In each case, the antennas were mounted on poles and separated by distances of 30cm to characterize an indoor AP and 50cm to characterize an outdoor AP.

Figure 4.1 shows the amount of average passive separation in dB as the angle θ shown in Figure 3.1 is varied for each directional antenna. No absorber was used and transmit and receive antennas were both vertically polarized. The "average" is over frequency in the 2.40 - 2.48 GHz ISM band, not over repeated trials. First observe that the isolation between omni antennas is 23 dB for 30 cm separation and 27 dB for 50 cm separation. This was expected since the dominant path is the direct path between two antennas, and the observed suppression is nearly equal to free-space path-loss.

Second, as expected, the passive suppression of directional antennas improves as the angle θ increases from 0° to 180°. The HGV-2406U [32] 60° beamwidth antenna has a small side-lobe from 90° to 180°, hence the suppression flattens after 90° separation and even degrades at 180°. The 60° beamwidth antennas do not have significant side-lobes[1] and suppression is nearly monotonic with θ.

Third, note that the directional antennas provide increased isolation, compared to omni, even for 0° angular separation, since they are focusing energy outwards perpendicular to the antennas and hence less energy is radiated laterally.

Next we zoom into one interesting angle on the plot – the 60° angular separation. At 50 cm antenna separation, the 60° beamwidth antenna achieves a suppression of 53 dB (26 dB more than omni) and the 90° beamwidth antenna achieves a suppression of 40 dB (13 dB more than omni). So from the self-interference suppression point of view, 60° beamwidth is enticing. However, an angular separation of 60° for 60° beamwidth means that there will be a decreased SNR, by approximately 3 dB, in many areas of AP coverage. The loss of 3 dB is because the 60°/90° beamwidth is 3 dB beamwidth, so at edge of each sector, the antenna gain is 3dB less. Hence, we will choose to use a $(N = 6, \theta_B = 90°)$ antenna architecture, with 50 cm separation, in several subsequent experiments.

4.2   Suppression from Absorptive Shielding

Suppression via absorptive shielding was evaluated using Eccosorb AN-79 [27] free-space RF absorber. AN-79 is a broadband, tapered loading absorber made from polyurethane foam impregnated with a carbon gradient. It is a 4.25 inch slab that can be cut to fit the application. The manufacturer's data sheet indicates that AN-79

---

[1]Higher directionality generally comes at the expense of side-lobes, hence the 90° antenna can get by wit small side-lobes but the 60° antenna cannot.

can provide up to 25 dB of absorption.

Table 4.1: Comparison of different combinations of passive isolation mechanisms for 50 cm antenna separation. Directional values are for 90° beamwidth antennas at 60° separation.

Configuration	Avg. Suppression	Minimum	Maximum
Omnidirectional	27.2 dB	26. dB	28.0 dB
Directional	39.9 dB	36.5 dB	50.3 dB
Directional + Absorber	45.5 dB	41.4 dB	65.5 dB
Directional + Cross-polarization	54.5 dB	50.3 dB	64.9 dB
Directional + Absorb + Cross-pol	62.2 dB	55.8 dB	83.4 dB

Table 4.1 compares the isolation achieved for several different combinations of isolation mechanisms. The first two rows of table Table 4.1 show that when placing absorber between omnidirectional antennas spaced 50 cm, the isolation increases from 27 dB to 47 dB: a 20 dB improvement in self-interference suppression. However, when the absorber was placed between directional antennas, for which the isolation is already 40 dB without absorber, the isolation only improves by 6 dB.

This result was initially surprising: RF absorber is a passive, linear suppression mechanism, thus the amount of absorption should not depend on the incident power. The best explanation is that for directional antennas, the self-interference is not near as dominated, like in omni-antennas, by the direct path between antennas. So while the absorber can attenuate the direct path passing through the absorber, it cannot impact the reflected path.We will show in the following section that absorptive isolation will be most effective in outdoor applications where energy from reflective paths between the AP antennas may be much lower.

4.3 Suppression from Cross-polarization

In Row 3 (Directional configuration) of Table 4.1, vertical polarization mode is used at both the receive and transmit antennas, so this serves as a baseline to measure the impact of cross-polarization. Row 5, shows the measured suppression when the transmit antenna instead uses the horizontal polarization mode. Cross-polarization of the Tx and Rx antennas suppresses by an additional 15 dB. The last row shows that when all three mechanisms are applied in tandem (for 50 cm separation, 60 degree separation angle), the average suppression is 62 dB. So in general, the three mechanisms applied in tandem provide a total of an additional 35 dB of passive suppression beyond the omni antennas.

4.4 Major Insights

The following two major insights can be derived from the above results.

4.4.1 Not distance, but *angle* and *polarization* have high suppression impact

As seen in Figure 4.1, increasing the antenna separation from 30 cm to 50 cm reduces the self-interference power by no more than 4dB. Directional separation of the antennas can improve the isolation by 10-20 dB, and cross-polarization can add 10-15 dB more isolation on top of that. Thus the solution to designing access points that are robust to full-duplex self-interference is not to make them big, but to make them "smart." This is encouraging that we do not need large amounts of "empty" space for separating the antennas.

4.4.2 Passive suppression introduces frequency selectivity
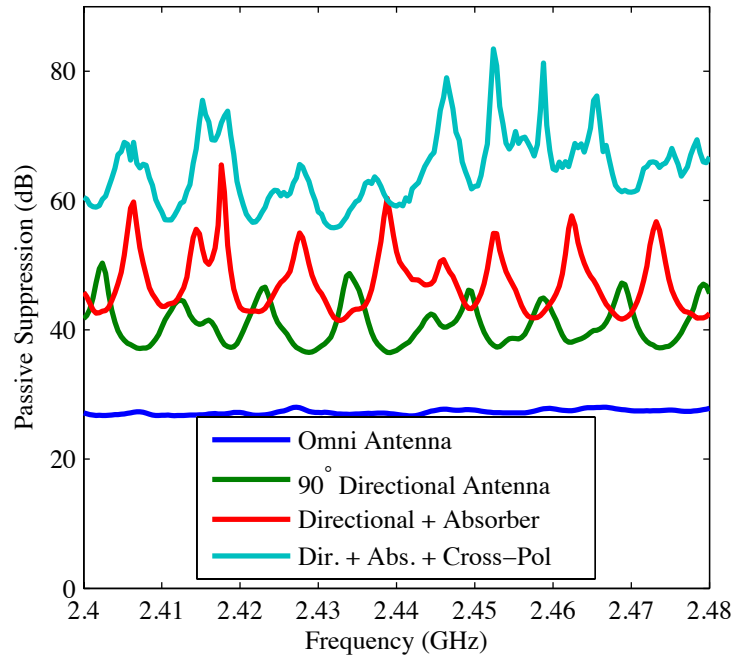
Figure 4.2: Passive suppression vs. frequency for combinations of the three mechanisms

Figure 4.2 shows us that when the self-interference is a simple line-of-sight signal between two omni antennas, the self-interference channel is nearly frequency flat. This is because the line-of-sight path dominates the received power, and the line of sight channel is governed by free-space path loss, which changes only slightly with frequency. However, we also see in Figure 4.2 that when we add passive isolation mechanisms, the self-interference channel becomes frequency selective. Notice the channel response for "Dir.+ Abs.+ Cross-pol" in Figure 4.2 from 2.45 GHz to 2.46 GHz. We see that when directionality, absorber, and cross-polarization are applied in tandem, the amplitude of the self-interference channel can change by more than 20 dB within a 10 MHz band.

This is to be expected. All of the passive isolation mechanisms are designed to

knock out the line-of-sight path between Tx antenna and the Rx antenna. The residual signal that does arrive at the receive antenna will likely be the superposition of several several waves that have "gotten around" the passive suppression mechanism. Since these waves will have random phases, they will add constructively at some frequencies and destructively at others. In other words, by passively suppressing the direct self-interference, we've essentially changed the self-interference from a frequency-flat, line-of-sight signal to a frequency-selective multi-path signal (thankfully with much lower average amplitude): the residual signal that does arrive at the receive antenna is due to multi-path reflections from surroundings. However, the degree of frequency selectivity may depend on the environment. In a smaller room than the one in which these measurements were performed, the residual self-interference may be *more* frequency selective due to closer scatters, but in an outdoor environment, the amount of frequency selectivity introduced by the passive suppression mechanisms may not be a great.

This frequency selectivity of the residual self-interference has an impact on the choice of active self-interference cancellation mechanisms at the AP receiver. A WiFi OFDM signal spans either 20 MHz or 40 MHz, and Figure 4.2 shows that the power of the residual self-interfernce can vary by more that 20 dB over such bandwidths. In the balun cancellation method of [6], a single amplitude and phase is applied to an inverted version of the transmit waveform before it is added to the received signal to cancel the self-interference. Thus the self-interference channel is being tracked with a single amplitude and phase. Tracking a self-interference channel with the characteristics observed in Figure 4.2 with a single amplitude and phase will obviously result in poor self-interference cancellation. In the the approach of [2], the cancellation waveform is crafted at baseband from the broadband per-subcarrier estimates of the self-interference channel. Hence the broadband ODFM RF cancellation waveform

produced from these per-subcarrier channel estimates can cancel self-interference signal that is very frequency selective. For this reason, in the full-system experiments described in Chapter 5, we use the analog cancellation method of [2] to cancel the the residual self-interference remaining after the passive isolation mechanisms.

Physical-Layer Evaluation

The passive suppression measurements in the previous sections show that the three proposed passive mechanisms provide excellent self-interference suppression for full-duplex APs, but they alone are not sufficient to achieve desired long ranges. In this section, we discuss the results of a software-defined radio implementation to measure (a) the total suppression with all mechanisms in place, passive and active and (b) achievable rate gains for only the uplink. We will evaluate the sum-rate gains, of both uplink and downlink, in Chapter 6.

5.1 Evaluation Methodology

5.1.1 Prototype Wideband OFDM Full-Duplex PHY

The prototype full-duplex PHY used at the AP is a WARPLab [34] implementation of the real-time wideband OFDM full-duplex physical layer presented in [5]. In the WARPLab framework, waveforms are transmitted over the air in real-time using the WARP platform [35], but are crafted and processed off-line in MATLAB. The prototype communicates using 20 MHz, 64-subcarrier OFDM waveforms within a

packet structure that mimics 802.11a. All experiments were performed at a center frequency of 2.484 GHz (channel 14 of the 2.4 GHz ISM band)

5.1.2 Metrics

As discussed earlier, the purpose of the prototype evaluation is to measure the amount of self-interference suppression achieved and measure the rate gain that full-duplex provides on the uplink. To measure the achieved self-interference suppression, the strength of the residual self-interference is measured using received signal strength indication (RSSI) pilots. And to measure the rate gains, we compute achievable rates from the error vector magnitude (EVM) of the received frames. Each of these metrics is discussed in greater detail below.

5.1.2.1 Self-interference suppression from RSSI measurements

The method for measuring the achieved suppression is the same as described in [4], but we mention it here for completeness. Like most commercial radios, the Maxim MAX2829 transceiver IC [36] used on WARP provides a received signal strength indication (RSSI) to the baseband processor. This RSSI signal allows us to measure the power of the RF waveform incident on the receiver, since the gain of the stages proceeding the transceiver chip are known deterministically and can be backed out of the RSSI reading. The header of the the mobile's packet has a small "silent period" during which the AP transmits "suppression pilots". During the suppression pilots, the only signal received at the AP is its own self-interference, hence the RSSI observed during this portion of the header allows the AP to measure the power of its self-interference signal. To measure the achieved passive suppression, we just observe the RSSI of the suppression pilots without employing any active suppression. To measure the achieved suppression when analog cancellation is added, we simply observe the

RSSI of the self-interference when the AP is employing analog cancellation. If P_{Tx} is the known transmit power, P_{RSSI} is the RSSI of the self-interference (measured when the signal-of-interest is turned off), and G_{RX} is the lumped gain of the stages preceding the transciever in the receive chain, then the measured RF self-interference suppression, α_{RF}, is

$$\alpha_{\text{RF}} = P_{\text{Tx}} - (P_{\text{RSSI}} - G_{\text{RX}}). \tag{5.1}$$

The total self-interference suppression includes digital cancellation which cannot be inferred from the RSSI, since digital cancellation is performed post-RSSI-measurement at baseband. To measure the suppression contributed by digital cancellation, we compute the difference of the squared magnitude of the digital samples corresponding to the suppression pilots before digital cancellation to the squared magnitude after cancellation to get α_{Dig}, the amount of self-interference suppression due to digital cancellation. Now the total acheived self-interference can be computed as

$$\alpha_{\text{Tot}} = \alpha_{\text{RF}} + \alpha_{\text{Dig}}. \tag{5.2}$$

5.1.2.2 Achievable Rate from EVM Statistics

We compare system-level performance of the full-duplex uplink to that of a half-duplex uplink, by transmitting frames in both cases (half-duplex and full-duplex) and measuring effective SNR from the error vector magnitude (EVM) statistics of the received frames. EVM is the distance (in the complex plane) of the received symbol after channel equalization and demodulation from the symbol actually transmitted. The average EVM for over a frame i of symbols, $\overline{\text{EVM}}(i)$, can be used to measure the effective SNR from frame i using the common conversion [37]

$$\text{SNR}(i) = \frac{1}{\overline{\text{EVM}}(i)^2}. \tag{5.3}$$

(a) Full-duplex uplink: AP transmits frames while receiving, effective effective $\mathsf{SSINR}^{\mathrm{FD}}_{\mathrm{Uplink}}$ is measured from EVM statistics

(b) Half-duplex uplink: AP only receives, effective effective $\mathsf{SNR}^{\mathrm{HD}}_{\mathrm{Uplink}}$ is measured from EVM statistics
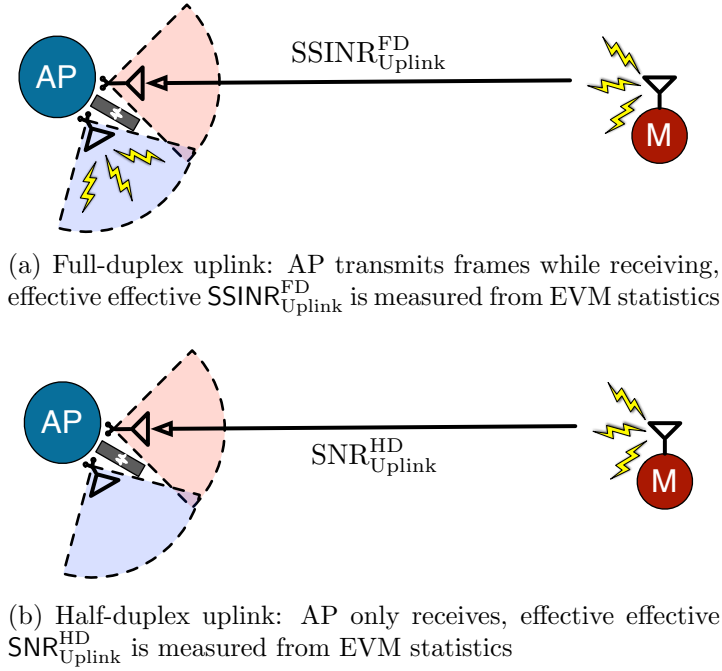
Figure 5.1: Experiment setup for evaluating performance of full-duplex uplink

As shown in Figure 5.1(a), we measure the SSINR for the full-duplex uplink at the AP by having the AP receive frames from the uplink mobile while simultaneously transmitting to a "dummy" downlink mobile on its nearest antenna. From the EVM statistics for each frame, we measure the $\mathsf{SSINR}^{\mathrm{FD}}_{\mathrm{Uplink}}$ for each frame. Similarly, to measure the SNR for half-duplex, we compute $\mathsf{SNR}^{\mathrm{HD}}_{\mathrm{Uplink}}$ from the EVM statistics of frames received from the uplink mobiles, but without transmitting simultaneously, as shown in Figure 5.1(b).

From these EVM-based SNR measurements, we can compute the ergodic achievable rate for the mobile-to-AP uplink from the SNR measurements using Shannon's formula. Let i be the index for the frames transmitted, and the N be the total number of frames transmitted. The half-duplex ergodic achievable rate, $R^{\mathrm{FD}}_{\mathrm{Uplink}}$, and the full-duplex ergodic achievable rate, $R^{\mathrm{HD}}_{\mathrm{Uplink}}$, are then computed by

$$R_{\text{Uplink}}^{\text{FD}} \;\;=\;\; \frac{1}{N}\sum_{i=1}^{N} \log_2[1 + \mathsf{SSINR}_{\text{Uplink}}^{\text{FD}}(i)] \qquad\qquad (5.4)$$

$$R_{\text{Uplink}}^{\text{HD}} \;\;=\;\; \frac{1}{N}\sum_{i=1}^{N} \frac{1}{2}\log_2[1 + \mathsf{SNR}_{\text{Uplink}}^{\text{HD}}(i)]. \qquad\qquad (5.5)$$

Note that in (5.5) a $\frac{1}{2}$ pre-log factor is added due to the half-duplex constraint. This $\frac{1}{2}$ factor assumes that in half-duplex mode the AP performs a 50/50 time-split between the uplink and downlink. The optimal time-split is in general not 50/50, but since in this experiment we do not incorporate a third node to receive the AP's downlink transmission, we assume that the SNR at the downlink mobile is the same as the SNR for the uplink mobile, in which case the optimal time split would indeed be 50/50. Under this 50/50 time split assumption, the uplink mobile will transmit twice as often when the AP operates in full-duplex mode as it would when the AP operates in half-duplex. Hence, in order for the comparison to be fair in terms of expended transmit power, the uplink mobile will transmit with twice the power (3 dB) in the half-duplex case as in the full-duplex case.

Equations 5.4 and 5.5 remind us that the advantage of the full-duplex uplink is that even though it will have a lower SSINR due to residual self-interference, it does not have to share time with the downlink. Thus operating the uplink in full-duplex mode allows the uplink mobile to transmit for twice the time duration as in the half-duplex uplink case. Full-duplex will "win" over half-duplex if the pre-log benefit of no longer time sharing outweighs the in-log cost of the residual self-interference.

5.2 Indoor Full-Duplex Uplink with Directional Isolation Only

Here we present performance evaluation results from a prototype indoor AP that only employs one of the mechanisms of Chapter 2: directional isolation [38]. Characterization of the full (N, θ_B) architecture that leverages all three mechanisms will come in the next section. The goal of this study is to better understand the impact of directional isolation in a full system. In particular, we wish to characterize how uplink performance is impacted by the angle between transmitting and receiving antennas. The characterization will supplement the suppression results in Section 4.1, in enabling an intelligent choice for N and θ_B, the number of antennas and beamwidth of each antenna, respectively, in the (N, θ_B) architecture.

5.2.1 Experiment Setup

The antennas used in the experiment were standard 2.4 GHz rectangular patch antennas [39]. These vertically polarized antennas have 5 dBi gain and 85° beamwidth. The configuration of the antennas for the experiment is shown in Figure 5.2(a). One antenna was used for transmission and the other for reception. The antennas were mounted such that they pivot around a common axis. The distance from the axis to the antennas was 18 cm. This mounting apparatus allowed control of the angle between transmit and receive directional antennas, so that performance as a function of the angle between the antennas could be measured. In Figure 5.2(a) the antennas are at 30° separation. Figure 5.2(b) shows a front view of the prototype full-duplex AP with the directional antennas at 45° separation.

The experiment was carried out in an open hallway in Duncan Hall at Rice University, where long-range line-of-site channels could be obtained. Figure 5.3 visualizes

(a) 5dBi directional patch antennas at the prototype full-duplex base station

(b) Front view prototype full-duplex base station with directional antennas mounted

Figure 5.2: Prototype full-duplex AP with directional antennas.

the experiment setup. At distances of 10 and 15 meters, we varied the angle between the transmit and receive antennas from 30 to 180 degrees. From each of the dots shown in Figure 5.3 the mobile transmitted 150 frames to the AP. The first 50 frames were half-duplex transmissions: the AP was not transmitting while receiving. From these first 50 frames we compute the half-duplex uplink rate $R_{\text{Uplink}}^{\text{HD}}$. During the second 50 frames the AP transmitted to a dummy downlink mobile while it received the frames form the uplink mobile, and the AP employed both analog cancellation and digital cancellation to suppress the self-interference. From these full-duplex frames we compute $R_{\text{Uplink}}^{\text{FD}}$, the rate achieved when both analog and digital cancellation are employed. Finally, in the last 50 frames the AP did not employ RF cancellation, and only canceled the self-interference digitally at baseband. From these frames we compute the achievable rate for digital cancellation alone, $R_{\text{Uplink}}^{\text{FD (Dig. Cx)}}$.

To quantify the benefits of using directional antennas at the AP, we also collected data for frames transmitted when using omnidirectional antennas. Changing from directional antennas to omnidirectional antennas has two results: (1) the self-interference at the AP will likely be stronger, since the transmit antenna will be radiating directly onto the receive antenna, (2) the power of the received signal from

Figure 5.3: Procedure for measuring full-duplex performance as a function of range and angle between Tx and Rx antenna at the AP.

the mobile will be weaker, since the receive omnidirectional antenna has a smaller gain than the directional antenna. The goal of this experiment was to characterize the benefits of directional antennas in mitigating self-interference, not in improving link quality. We therefore wanted to study the first effect in isolation from the second. For this reason, we empirically determined a mobile-to-AP distance for which the received signal strength (RSSI) at the AP with omnidirectional antennas was nearly equal to the RSSI when directional antennas were used. We then transmitted 50 frames at each of these effective distances to measure compare full-duplex to half-duplex in the case of omni antennas at the AP. Hence instead of taking measurements at 10 m and 15 m, as in the directional antenna case, measurements were taken at the "effective distances" of 7.0 m and 13.3 m.[1]

---

[1]Because the self-interference power will not change significantly with the angle between omni-directional antennas, only one measurement at each of the two distances is performed, and in the plots that follow we assume that the same values would have been measured at all angles.

5.2.2 Results

Figure 5.4 visualizes the measured percentage uplink rate improvement full-duplex achieves over half-duplex, assuming a 50/50 time-sharing with downlink in the half-duplex case. The regions are colored according to percent improvement over the half-duplex achievable rate attained at each angle and distance. The dots represent the coordinates of the actual measurements, and the rest of the region's coloring is obtained via interpolation. Green indicates an improved rate over half-duplex: brightest green corresponding to the ideal 100% gain (i.e. doubling of half-duplex rate). Black corresponds to full-duplex being on-par with half-duplex and red indicates full-duplex underperforming half-duplex. The top plot is the case of analog and digital cancellation applied in tandem, the second is digital cancellation alone, and the last is when omnidirectional antennas rather than directional antennas are used (in this case both RF and digital cancellation are employed).

5.2.2.1 Analog + Digital Cancellation

We see in Figure 5.4(a) that full-duplex performs quite well when directionality is exploited and both analog and digital cancellation are employed. At 10 m range full-duplex outperforms half-duplex by more than 60% as long as the antennas are separated by at least 45°, and at 15 m range full-duplex outperforms half-duplex by at least 50% for angles ranging from 90° to 150°. The best performance is achieved at (10 m, 120°), where a near 95% improvement over half-duplex is achieved; this means we are approaching the ideal doubling of rate that full-duplex promises. However, as the angle between antennas gets small, performance degrades. When 45° separation is approached, the gains over half-duplex are small, and in the region around (15 m, 30°) we actually see the color fade from black to dark red: at (15 m, 30°), full-duplex is underperforming half-duplex. We now turn to the received signal strength values

(a) Analog + Digital Cancellation with Directional Antennas

(b) Digital Cancellation Alone with Directional Antennas

(c) Analog + Digital Cancellation with Omni Antennas

Figure 5.4: Percent improvement over the half-duplex achievable rate as a function of mobile-to-AP distance and angle between antennas (*a*) when directional antennas are employed and both analog and digital cancellation are performed (*b*) with directional antennas and digital cancellation only, and (*c*) with omnidirectional antennas and both analog and digital cancellation.

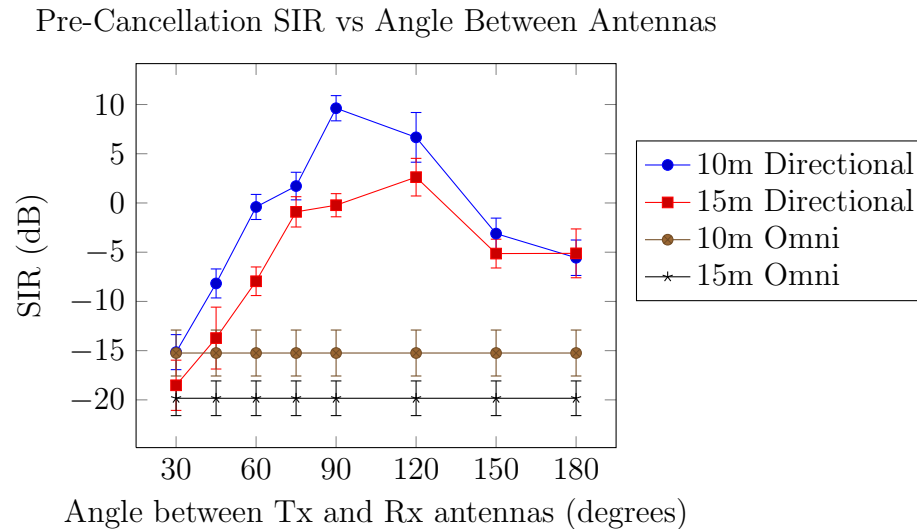Pre-Cancellation SIR vs Angle Between Antennas

Figure 5.5: Pre-cancellation signal to self-interference ratio (SIR) as computed from RSSI readings

to understand why full-duplex fails in this region.

Figure 5.5 plots the pre-cancellation signal-of-interest to self-interference ratio (SIR) as a function of the angle between antennas for each of the distances evaluated in the experiment. The SIR values are obtained from the radios' average RSSI readings over the frames transmitted. Figure 5.5 helps us understand why full-duplex is underperforming at small angles. At around 75° the pre-cancellation SIR begins to fall off rapidly with decreasing angle due to the coupling between the Tx and Rx antennas becoming stronger as the angle between them gets smaller. At (15 m, 30°) the self-interference is nearly 20 dB more powerful than the signal-of-interest, and in this regime the cancellation mechanisms do not suppress the self-interference sufficiently for full-duplex to outperform half-duplex.

5.2.2.2 Digital Cancellation Alone

Figure 5.4(b) shows that when directionality is exploited, full-duplex can achieve significant rate improvements over half-duplex even without employing extra hardware

for RF cancellation. At 120° the full-duplex rate is around 60% higher than the half-duplex rate, and full-duplex continues to out-performs half-duplex for angles from 60° to 150° at 10 m and from 90° to 130° at 15 m. However, at 60° the gains over half-duplex are marginal and as the angle get smaller the self-interference becomes too powerful to be suppressed via digital cancellation alone, and the rate falls below the half-duplex rate. The fact that performance degrades for smaller angles is expected, but it is surprising that performance also degrades for large angles. At 180°, when the antennas are pointed in opposite directions, full-duplex actually underperforms half-duplex when RF cancellation is not employed. Let us look to the pre-cancellation SIR values for an explanation of this decreased performance at large angles.

We see in Figure 5.5 that SIR starts off small when the angle between the antennas is small, and the direct coupling is strong. As the angle increases the SIR increases, since the self-interference is becoming weaker as the antennas become more isolated. The SIR reaches a maximum somewhere around $90-120°$, and then begins to decline as the angle increases further. The self-interference is actually much stronger at 180° than 90°. There are three possible causes for this surprising increased coupling when the antennas are facing opposite directions. One possibility is an antenna back-lobe. The patterns included in the antenna data sheet [39] indicate a small back-lobe, but the back-lobe does not seem strong enough to produce the observed 10 dB swing in SIR. Another possibility is that the increased coupling is an artifact of room-specific reflections – that it is not direct coupling between antennas causing the lower SIR, but a reflected component. This could partially be the case, but when performing a pilot study in a different room, similar effects were observed. The final possibility is that the increased self-interference at large angles is due to a near-field coupling effect that would not be captured in the far-field antenna patterns. A future work is to perform full-wave electromagnetic simulations to determine the exact mechanism

causing this observed increase in self-interference when antennas are pointed away from each other.

5.2.2.3 Omnidirectional comparison

Figure 5.4(c) shows the performance when omnidirectional antennas are employed rather than directional antennas. With omni antennas there is obviously no angular variation in performance. We see that the self-interference is too strong to be suppressed enough for full-duplex to be preferable to half-duplex at the distances evaluated. The distances evaluated here are longer than those evaluated in [2, 5], where full-duplex was shown to effective with omnidirectional antennas. This result shows that as the distance between devices increases, and the signal-of-interest attenuates, *passive suppression* is needed to attenuate the self-interference in order for full-duplex to be effective. Comparing omni vs. directional performance at (15 m, 90°), we see that with directional antennas and RF + digital cancellation full-duplex outperforms half-duplex by ∼75%, but when the directional antennas are interchanged with omni antennas the pre-cancellation SIR shifts from a benign ∼0 dB SIR to a challenging ∼−20 dB SIR, and the full duplex achieved rate is ∼75% *less* that what is achieved with half-duplex. Hence achieving passive suppression by exploiting directionality makes a huge impact on system performance.

5.2.3 AP Architecture Implications

One of the conclusions of the passive suppression measurements of Chapter 4 was that a 6-antenna AP, with each antenna separated by 60°, is reasonable design choice for leveraging directional isolation without resorting to excessive antennas. Figure 5.4 shows that using directional antennas separated 60° instead of omni antennas pushes the full-duplex improvement from 0% to +70% at 10 m range, and from −40% to

+50% at 15 m.

Therefore for ranges of 10-20 m, directional isolation alone is sufficient for enabling full-duplex, and this is about largest line-of-sight range we could study indoors. But even with the 60° directional isolation, gains start falling off after 15 m. In order to achieve our 100 m goal in an outdoor setting, we will need to employ all of the passive suppression mechanisms introduced in Chapter 2. This is the topic of the following section.

5.3 Outdoor Full-Duplex Uplink at 100+ Meters

The overarching goal of the three passive suppression mechanisms introduced in Chapter 2, and the (N, θ_B) architecture for leveraging them (Chapter 3) is that they would enable long-range full-duplex links, when employed along with the pre-existing active suppression mechanisms. In this section we put the mechanisms to the test, to see whether or not the architecture achieves this goal. Here we present a performance evaluation of a prototype outdoor full-duplex AP, receiving uplink packets from a mobile node at ranges of 50-150 m, while transmitting a full power.


Figure 5.6: Uplink test setup

5.3.1 Experiment setup

Per our conclusions in Section 4.1, we prototype a $(N = 6, \theta_B = 90°)$ antenna architecture using the HG2414DP dual-polarized 90° antennas with 50 cm separation.

Since 6 antennas are used, the worst case self-interference occurs between two antennas separated 60°.

Figure 5.7: Photo of outdoor experiment setup: AP directional panel antennas with absorber mounted between.

To evaluate the worst-case self-interference in the 6-antenna AP, the experiment setup depicted in Figure 5.6 was used in which the receive antenna is pointed 60° away from the receive antenna's direction. This setup is shown in Figure 5.7. An average transmit power was used at the mobile and at the AP. The AP's directional antenna has a gain of 14 dBi, giving an effective isotropic radiated power (EIRP) at the AP of 21 dBm. This high EIRP allows the received signal strength at a downlink mobile to be in acceptable WiFi ranges (-80 to -60 dBm) for typical 100 to 200m outdoor path losses of 80 to 100 dB. The uplink mobile uses 6 dBi gain omnidirectional antenna.

When comparing to a half-duplex uplink, the mobile node is allowed to transmit with twice the transmit power (3 dB more) than in the full-duplex case, so that the comparison to half-duplex is fair in terms of average power (assuming an equal uplink/downlink time split).

Since the goal of the experiment was to study self-interference suppression at the AP and its effect on the *uplink* rate at the AP, the AP will be transmitting at full power on the other directional antenna in full-duplex mode and silent in half-duplex mode. The AP was placed in a fix location, and the distance between the uplink

mobile and the AP was varied from 50 to 150 m. At each location the mobile node transmitted frames to the AP, and statistics were recored.

5.3.2 Total average suppression exceeding 90 dB

First we characterize total amount of self-interference suppression achieved. Figure 5.8 plots the empirical CDFs for the amount of cancellation achieved both with and without cross-polarization at the AP. The solid blue CDF for the amount of passive suppression achieved by directional isolation and absorptive shielding (no active cancellation) with the AP transmitting and receiving on the same polarization. Note that this is nearly 10 dB more suppression than was observed in the indoor measurements, this corroborates the hypothesis that in the indoor measurements, the absorber was not adding much contribution because reflected paths had become dominant. Here is the outdoor setting, reflections are minimal, and absorptive shielding has a big impact: enabling 60 dB passive suppression without any cross-polarization. The dotted blue CDF is the total (passive + active) suppression achieved after analog and digital cancellation are performed. Analog and digital cancellation add another 25 dB of active suppression for a total suppression of 86 dB on the average.

The red CDFs are the passive and total (passive + active) suppression achieved with cross-polarization of Tx and Rx antennas at the AP. We see that cross-polarization adds another 10 dB of passive suppression over the suppression provided by directional isolation and absorptive shielding, for 72 dB of total passive suppression from the three mechanism. Once again, this is $\sim$ 10 dB better than the suppression measured indoors, due to the absorptive shielding and directional isolation being more effective in a lower-reflection environment. When active suppression is added, the total self-interference suppression ranges from 87 dB to 100 dB, averaging around 94 dB total suppression. Analog and digital cancellation is still able to provide 20-25
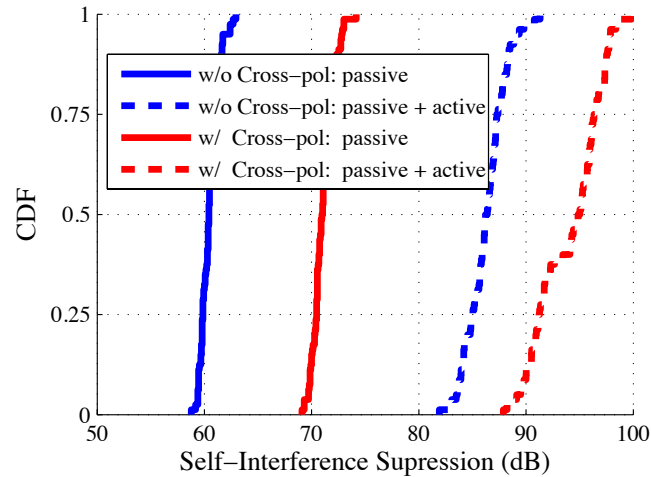
Figure 5.8: Total self-interference suppression of 20MHz OFDM signal

dB of self-interference suppression, even after the self-interference is knocked down by 70 dB by the passive mechanisms. Note that as the amount of suppression increases, the variability of the suppression also increases. Nonetheless, for more 90% of the frames transmitted, the total self-interference suppression was at least 90 dB when the three passive mechanisms: directional isolation, absorptive shielding, and cross-polarization are applied in tandem with analog and digital cancellation. *To our knowledge this is the best reported self-interference suppression to date.*

5.3.3 Full-duplex rate improvements at 100+ m

In an outdoor area on a university campus, the achievable rate for the full-duplex and half-duplex uplinks were measured as the range from of the uplink mobile to the AP was varied from 50 to 150 meters. As expected, the encountered path loss was not monotonic with range due to shadowing effects, hence rate-vs.-range curves are noisy and difficult to interpret. Instead, we measured the encountered path loss for each mobile node location so that the performance could be indexed by this meaningful, repeatable parameter. Figure 5.9 plots the percent improvement in achievable rate

of full-duplex uplink over the half-duplex uplink (assuming half-duplex uplink must share time equally with the downlink).
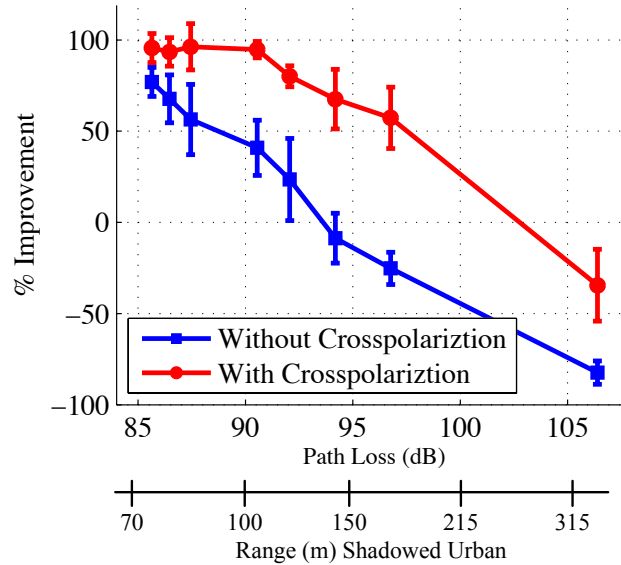

Figure 5.9: Percent rate improvement of FD uplink over HD uplink

We see that at 86 dB path-loss, the improvement over half-duplex is 86% when directional isolation and absorptive shielding are applied together with analog and digital cancellation. When cross-polarization of Tx and Rx is also employed, the gain over half-duplex at 86 dB path loss is 96%. In Figure 5.9, a second x-axis is also shown, which maps pathloss to effective range in a typical urban channel. With the best design, including cross-polarization, we expect to have full-duplex gains (albeit small) for up to 200 m. As a result, we will use 200 m range for further sum rate analysis in Chapter ch:internode. Even without cross-polarization, our design achieves 100+ m range, thereby meeting our goal of long-range full-duplex.

Of particular note is the significant impact of cross-polarization. Figure 5.8 showed that cross-polarization provides an extra 10 dB of total suppression, and we see the system-level impact of this extra 10 dB in Figure 5.9. Notice that with cross-polarization, the same gains over half-duplex can be attained at 10 dB more path loss

with cross-polarization that without. In other words, cross-polarization allows the full-duplex AP to handle 10 dB more path loss than it could otherwise, and 10 dB more translates to several tens of meters (nearly 100) of added range.

5.4   PHY Evaluation Summary

We have seen that with omnidirectional antennas, i.e. no passive self-interference suppression, analog and digital cancellation can only enable a full-duplex uplink to outperform a comparable half-duplex uplink when the range is less than 10 m (see Figure 5.4). When directional isolation and and absorptive shielding are employed in an $(6, 90°)$ architecture, with 50 cm antenna separation, total self-interference suppression of 85 dB is achieved, enabling the full-duplex uplinks to have > 50 % improvement over half-duplex for path loss up to 87 dB (90 m in shadowed urban environment). Adding cross-polarization gives total suppression of 94 dB, enabling a full-duplex uplink to outperform half-duplex by > 50 % for range up to 150 m. Thus the large passive suppression achieved by utilizing directional isolation, absorptive shielding, and cross-polarization enables our goal of extending the range full-duplex to typical Wi-Fi distances.

Full Topologoy Evaluation: Impact of Inter-node Interference

In Chapter 1, we introduced a scenario in which a full-duplex AP can network efficiency by receiving an uplink signal from one node while simultaneously transmitting to another node. For convenience, we refresh Chapter 1's illustration of this topology in Figure 6.1. Recall that simultaneous uplink and downlink introduces two challenges: self-interference at the AP, and inter-node interference from the uplink mobile at the downlink mobile. The main thrust of this thesis is addressing the first problem: employing passive suppression mechanisms so that the uplink signal can overcome self-interference. All previous chapters have focused on addressing self-interference on the uplink, since the mechanisms we have proposed only assist the uplink. In this chapter we evaluate the full network performance when inter-node interference at the downlink is also taken into account.

Figure 6.1: Full-duplex AP with half-duplex mobiles. Uplink suffers from self-interference while downlink suffers from inter-node interference

6.1 Evaluation Methodology

Modeling of self-interference at the AP is driven by experiment results in Chapter 5. To analytically capture the impact of self-interference at the AP, we assume that the self-interference is suppressed by a constant amount. In our field trials, 90 dB or more suppression was observed for 90% of the frames transmitted (see discussion about Figure 5.8 in Section 5.3.2). Thus we assume that the AP faces a self-interference floor that is 90 dB less than its transmit power.

We model attenuation of the signal-of-interest and inter-node interference using a log-distance path loss model with path-loss exponent η [40]. To model a realistic typical deployment, we use a reference distance of 10 m, thus the path loss model consists of the following two parts: free space path loss for the first 10 meters and path loss with a loss exponent η beyond 10 meters. The attenuation (in dB) of the signal at a distance d from the transmit antenna is given by:

$$
L_p(d) = \begin{cases} 32 + 20\log_{10}(f_c d) & \text{if } d \leq 10 \\ 60 + 10\eta\log_{10}(d/10) & \text{d} > 10 \end{cases}
$$

where f_c is the carrier frequency in GHz, which we set to 2.4 GHz.

We consider an outdoor urban environment where line-of-sight to the AP is avail-

able[1], and assume that mobiles are equipped with dual-polarized antennas to match the AP's cross-polarization of uplink and downlink. Measurement campaigns in such environments have shown that co-polarized signals (AP↔mobiles) encounter a path-loss exponent of 2, while cross-polarized signals (mobile↔mobile) see a path-loss exponent of 4, due the cross-polarization eliminating the line-of-sight [41]. Thus we assume $\eta = 2$ for uplink and downlink and $\eta = 4$ for inter-node interference.[2]

To assess net performance of the topology of Figure 6.1, we compare the achievable sum-rates of full-duplex to that of half-duplex. The sum-rate of full-duplex is given by

$$
\begin{aligned}
R_{\text{sum}}^{\text{FD}} &= R_{\text{Downlink}}^{\text{FD}} + R_{\text{Uplink}}^{\text{FD}} \\
&= \log_2(1 + \text{SINR}_{\text{Downlink}}) \\
&\quad + \log_2(1 + \text{SSINR}_{\text{Uplink}}).
\end{aligned}
\tag{6.1}
$$

For half-duplex, the achievable sum-rate is

$$
\begin{aligned}
R_{\text{sum}}^{\text{HD}} &= \frac{1}{2}\log_2(1 + \text{SNR}_{\text{Uplink}}) \\
&\quad + \frac{1}{2}\log_2(1 + \text{SNR}_{\text{Downlink}}).
\end{aligned}
\tag{6.2}
$$

To compare full-duplex with half-duplex, we assume that the uplink mobile is at a fixed distance from the AP, while the downlink mobile is located within a circular region of radius 200 meters from AP. The AP architecture we evaluate is the same $(6, 90°)$ architecture evaluated in Section 5.3: the AP has 6 antennas each covering a 60° region, and 90° beamwidth for each antennas ensures that the gain is uniform

---

[1]Such was the case in our experiments.

[2]Even without cross-polarization, the mobile-to-mobile channel will encounter higher path loss than the mobile-to-AP channel simply because the infrastructure AP will be strategically deployed (tower or top of a building) for unobstructed coverage, while the path between mobiles may be quite obstructed.

over each region. To match the parameters of the implementation Section 5.3, in the received power calculations we assume both the AP and the mobile node transmits with 7 dBm power, and we assume 14 dBi gain antenna gain at the AP, and 6 dBi gain at the mobile nodes. We assume that the AP chooses to communicate on the particular directional antenna which has maximum gain in the direction of the mobile node. The AP architecture, as described in Section 3, does not allow the same antenna at AP to act as transmitter and receiver simultaneously. Therefore we exclude computing the gain of full-duplex over half-duplex when both uplink and downlink mobiles are in the same "sector". As shown in Figure 6.3, we compute the percentage gain of full-duplex over half-duplex only in an angular region between 30° and 330° since the uplink mobile is fixed on the positive side of the horizontal axis.
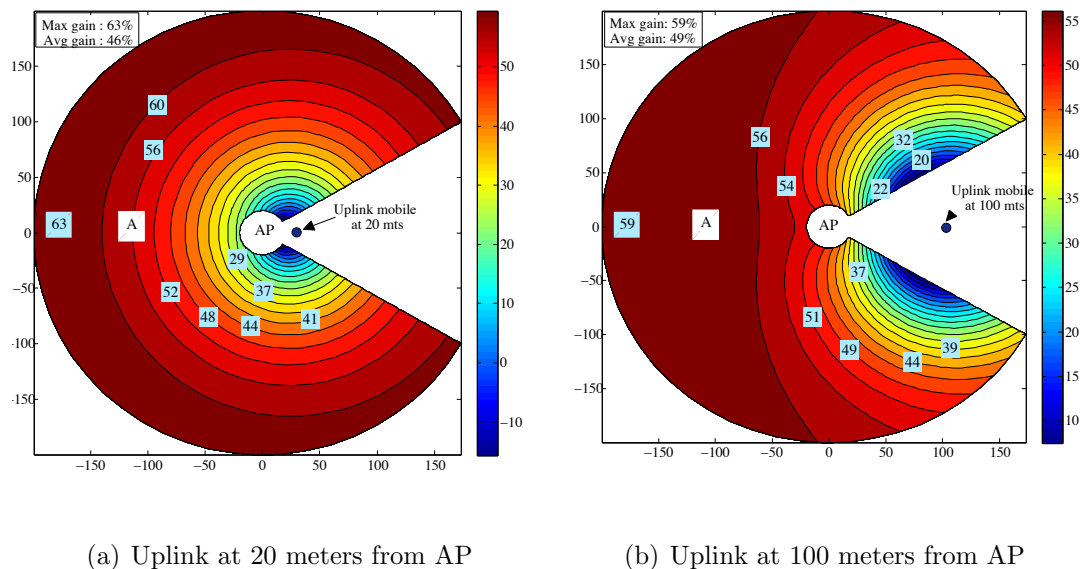

(a) Uplink at 20 meters from AP   (b) Uplink at 100 meters from AP

Figure 6.2: The overall gain (in %) in sum-rate due to full-duplex system over the half-duplex counterpart.

6.2 Evaluation Results

Figure 6.3 plots the sum-rate percentage gain of full-duplex over half-duplex as a function of the geographical location of the downlink mobile for a fixed uplink mobile location. In Figure 6.2(a) the uplink mobile is placed 20 m from the AP at location $(x, y) = (20\text{m}, 0\text{m})$, and in Figure 6.2(b) the uplink mobile is 100 m from the AP at $(100\text{m}, 0\text{m})$. As expected, performance of full-duplex is poor when the downlink mobile is close to uplink mobile, and interference from the uplink at the downlink is strong. For example, when the uplink mobile is at $(20\text{m}, 0\text{m})$ and the downlink is at $(25\text{m}, 25\text{m})$ full-duplex has worse performance than half-duplex: the cost of orthogonalization is less than the cost of tolerating interference from the uplink mobile However, we see that when the downlink mobile is separated from the uplink mobile by a reasonable distance, full-duplex provides significant sum-rate gains over half-duplex.

6.2.1 Peak gains

From Figure 6.3, we observe that the maximum gain achieved by our design is 63% and 59% over half-duplex when uplink mobile is located at 20 meters and 100 meters, respectively.

Figure 6.2(a) shows that when the uplink mobile is near to the AP, full-duplex performance is best when the downlink mobile is far from the AP. Full-duplex is therefore advantageous in *asymmetric range* scenarios in which the uplink is near to the AP, but the downlink is far. This is because when the uplink is near the AP, the uplink signal encounters little path loss, and the residual self-interference is far below the signal-of-interest. Similarly, since the downlink mobile is far form the AP, it is also far from the uplink mobile, and inter-node interference is tolerable.

Figure 6.2(a) shows that when the uplink mobile is far from the AP, full-duplex

Table 6.1: Sum-rate Gain of Full-duplex over Half-duplex

Uplink mobile distance	Gain of Ideal FD	Max. Gain of our design
20	86%	63%
40	85%	61%
60	85%	60%
80	84%	59%
100	84%	59%

performance is best when the downlink mobile is in the hemisphere opposite the uplink mobile. Note that in the opposite asymmetric case to the one discussed earlier, when the uplink is far and the downlink is close, the gains over half-duplex are high only if the downlink and uplink mobiles are on opposite sides of the AP.

No matter the location of the uplink mobile, the peak gains occur when the downlink mobile is separated as far as possible from the uplink, and inter-node interference has the least impact. Table 6.1 shows the comparison of maximum gain achieved by our design with that of ideal full-duplex for different uplink distances. We note that achieving 100% gain over half-duplex is not possible even if self-interference and inter-node interference were made zero (ideal full-duplex). This is because half-duplex is allowed to use 3 dB more power than full-duplex so as to ensure equal average power consumption, since node will be transmitting for twice the time in full-duplex mode.We notice that the percentage gain gap between our design and the ideal gains is only 20%, indicating that for an outdoor setting, our design is operating close to the ideal full-duplex.

6.2.1.1 Percentile area gains

The second key observation from Figure 6.3 is that the average/maximum gains for both uplink mobile distance of 20 and 100 meters are approximately the same. This

(a) Interference and noise floors

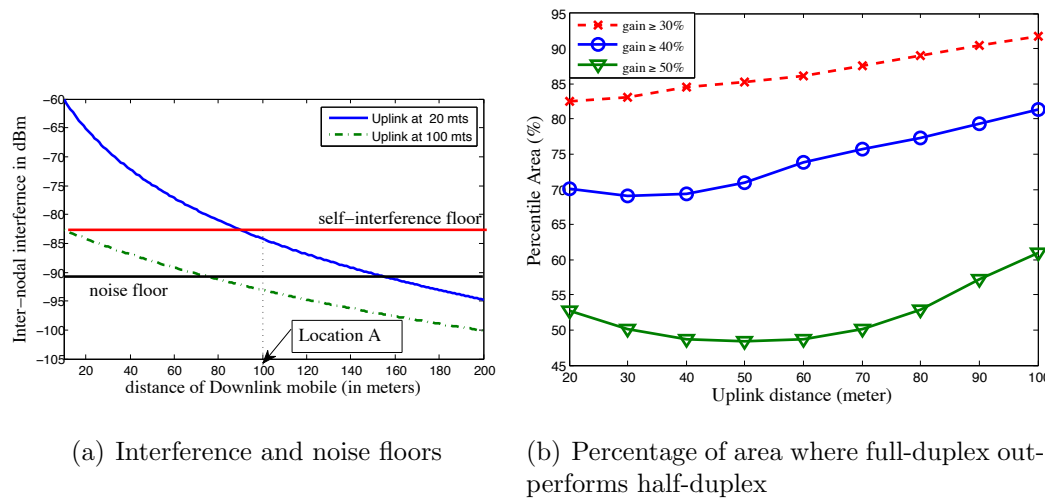(b) Percentage of area where full-duplex out-performs half-duplex

Figure 6.3: (a) Self-interference power with distance of the downlink mobile is shown for fixed uplink mobile distance. (b) Percentage of area where full-duplex out-performs half-duplex by at least 30%, 40% and 50% for varying distance of uplink mobile is shown.

implies that to a good extent the gain of full-duplex over half-duplex is insensitive to the distance of the uplink mobile.

To understand this observation, we consider the interplay between self-interference and inter-node interference. Consider that the downlink mobile is situated at Location A which is at a distance of 100 meters from AP on the negative side of the horizontal axis, as shown in Figure 6.3. The path-loss model tells us that the inter-node interference faces a path-loss of 103 dB and 112 dB for uplink mobile distance of 20 and 100 meters, which translates to an absolute interference power of −84 and −93 dBm, respectively. From Figure 6.3(a), we note that the self-interference floor for the same transmit power at AP is -83 dBm (assuming 90 dB cancelation). The self-interference floor is higher than the inter-node interference floor (i.e. self-interference, not inter-node interference is the aum-rate bottleneck) when the downlink mobile is at Location A or anywhere further than that as seen in Figure 6.3(a). Also, at Location

A or anywhere further (a reasonably large region), the inter-node interference approaches the thermal noise floor which implies that the performance of full-duplex is non-ideal only because of the self-interference floor. Since this phenomenon is uplink mobile distance independent (almost), the average/maximum gains too are almost independent of uplink mobile distance.

Figure 6.3(b) shows the plot of the percentage of area in the circular region where the gain is at least X% as the uplink-to-AP distance varys. The plot is shown for gains of at least 30%, 40% and 50%. Fixing the percentage gain desired, the percentile area does not change by more than 10% even when the uplink mobile is moved from 20 to 100 meters.

6.2.2 Comparisons to prior full-duplex designs

The main feature of our design surpassing prior full-duplex designs is the high overall suppression achieved by techniques described in Section 3.
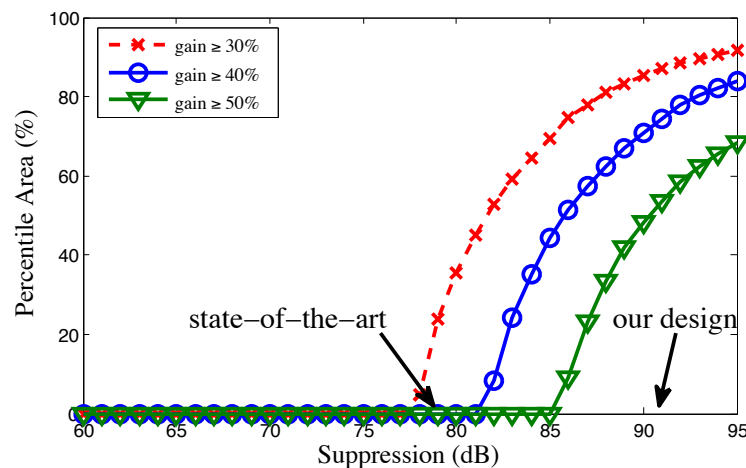

Figure 6.4: Percentage of area where full-duplex out-performs half-duplex by 30%, 40% and 50% as a function of achieved self-interference suppression. Uplink mobile is fixed at 50 meters from AP.

Figure 6.4 shows the progression in the percentage of area in the circular region within 200 meters where the gain is at least X% more than half-duplex. We can see the effect of higher suppression on the percentile gains. For having any percentage gains above 30% of full-duplex over half-duplex, an overall suppression of more than 75 dB is needed. With the current state-of-the-art suppression techniques which offers less than 79 dB suppression [5, 6, 16], *not* more than 24% area has 30% gain. On the other hand, our proposed design which has reliably more than 90 dB overall suppression has at least 50% gain in as much as 50% of the whole area, which is appreciably better than the state-of-the-art.

6.2.3   Take-away Message: Opportunistic MAC Needed

The main take-away from the above results is that there *are indeed* opportunities for simultaneous uplink/downlink to an AP in which inter-node interference is tolerable, and given the high levels of self-interference suppression demonstrated in Chapters 4 and 5 as high as 60% gains over half-duplex can be achieved. However, performance is highly dependent on the topology, as Figure 6.3 shows. Figure 6.3 the locations where full-duplex under-performs half-duplex are small, but we have taken the optimistic assumption of path-loss-exponent 2 for the signal-of-interest and 4 for inter-node interference. In a case where the mobiles are not in line-of-sight of the AP, such as an indoor office deployment, both inter-node interference and signal-of-interest may encounter a path-loss-exponent of 4. In this case, topologies where internode interference prohibits simultaneous uplink/downlink will be much more frequent.

This high variability of full-duple vs. half-duplex performance suggests that need for a new MAC. Design of a MAC protocol decides when to have full-duplex simultaneous uplink/downlink and when to orthogonalize is of utmost importance in the future work. The best solution would be an opportunistic full-duplex MAC proto-

col. The MAC must be "opportunistic" in the sense that if the AP is preparing to transmit receive a signal from some uplink mobile, it will "find" a downlink mobile, who is out-of-range of the uplink mobile, to whom it can transmit while receiving the uplink.

Conclusion

There is a new opportunity to reap the benefits of full-duplex without asking mobile devices to become "full-duplex" themselves: access points supporting simultaneous uplink/downlink with half-duplex devices. We have presented the design of a novel RF/antenna architecture for full-duplex access-points that achieves 90+ dB of total self-interference suppression by leveraging three mechanisms: directional isolation, absorptive shielding, and cross-polarization. Field tests using the new design demonstrated high-rate full-duplex links not for ranges not on the order to 10 meters as in the previous work, but on the order of 100 meters. Although simultaneous uplink/downlink introduces interference between mobile devices, path loss analyses revealed large geographic regions for which the benefits of simultaneous transmission outweigh the cost of tolerating interference: the overall network capacity gain can be as high as 60%. The obvious next step is the development of MAC protocols to identify such "full-duplex wins" scenarios and initiate simultaneous uplink and downlink transmissions at the AP.

References

[1] D. W. Bliss, P. A. Parker, and A. R. Margetts, "Simultaneous transmission and reception for improved wireless network performance," in *Proceedings of the 2007 IEEE/SP 14th Workshop on Statistical Signal Processing*. Washington, DC, USA: IEEE Computer Society, 2007, pp. 478–482. [Online]. Available: http://portal.acm.org/citation.cfm?id=1524876.1525079 1, 1.1.1, 1.4.1

[2] M. Duarte and A. Sabharwal, "Full-duplex wireless communications using off-the-shelf radios: Feasibility and first results," in *Proc. 2010 Asilomar Conference on Signals and Systems*, 2010. 1, 1.1.1, 1.1.2, 1.4.1, 4.4.2, 5.2.2.3

[3] J. I. Choi, M. Jain, K. Srinivasan, P. Levis, and S. Katti, "Achieving single channel, full duplex wireless communication," in *MobiCom 2010*. 1, 1.1.1, 1.4.1, 1.4.3

[4] M. Duarte, C. Dick, and A. Sabharwal, "Experiment-driven characterization of full-duplex wireless systems," May 2011, submitted to *Wireless Communications, IEEE Transactions on*. [Online]. Available: http://warp.rice.edu/trac/wiki/TransWireless2011_FullDuplex 1.1.1, 1.3, 1.4.1, 3.2.2, 5.1.2.1

[5] A. Sahai, G. Patel, and A. Sabharwal, "Pushing the limits of full-duplex: Design and real-time implementation," *CoRR*, vol. abs/1107.0607, 2011. 1.1.1, 1.1.2, 1.3, 1.4.1, 1.4.3, 2.2, 3.2.2, 5.1.1, 5.2.2.3, 6.2.2

[6] M. Jain, J. I. Choi, T. Kim, D. Bharadia, S. Seth, K. Srinivasan, P. Levis, S. Katti, and P. Sinha, "Practical, real-time, full duplex wireless," in *Proceedings of the 17th annual international conference on Mobile computing and networking*, ser. MobiCom '11. New York, NY, USA: ACM, 2011, pp. 301–312. [Online]. Available: http://doi.acm.org/10.1145/2030613.2030647 1.1.1, 1.4.1, 4.4.2, 6.2.2

[7] K. Chebrolu, B. Raman, and S. Sen, "Long-distance 802.11b links: performance measurements and experience," in *Proceedings of the 12th annual international conference on Mobile computing and networking*, ser. MobiCom

'06. New York, NY, USA: ACM, 2006, pp. 74–85. [Online]. Available: http://doi.acm.org/10.1145/1161089.1161099 1.1.2

[8] A. Sheth, S. Nedevschi, R. Patra, S. Surana, E. Brewer, and L. Subramanian, "Packet loss characterization in wifi-based long distance networks," in *INFO-COM 2007. 26th IEEE International Conference on Computer Communications. IEEE*, may 2007, pp. 312 –320. 1.1.2

[9] J. Ott and D. Kutscher, "Drive-thru internet: Ieee 802.11b for "automobile" users," in *23rd Annual Joint Conference of the IEEE Computer and Communications Societies*, ser. INFOCOM 2004, 2004. 1.1.2

[10] J. Choi and M. Jain and K. Srinivasan and P. Levis and S. Katti, "Achieving single channel, full duplex wireless communication," in *Proc. of the ACM Mobicom*. Illinois, USA: ACM New York, NY, USA, 2010. 1.1.2, 1.4.1

[11] A. Sabharwal, A. Sahai, E. Everett, G. Patel, A. Amiri, and L. Zhong, "Revisiting basic assumptions in wireless: full-duplex and directional communication," in *Information Theory and Applications Workshop*, 2012. 1.2

[12] E. Everett, D. Dash, C. Dick, and A. Sabharwal, "Self-interference cancellation in multi-hop full-duplex networks via structured signaling," September 2011, to Appear: *49th Annual Allerton Conference on Communication, Control, and Computing*. 1.2

[13] B. P. Day, D. W. Bliss, A. R. Margetts, and P. Schniter, "Full-duplex bidirectional mimo: Achievable rates under limited dynamic range," in *Proc. Asilomar Conference on Signals and Systems*, November 2011. 1.2

[14] B. P. Day, A. R. Margetts, D. W. Bliss, and P. Schniter, "Full-duplex mimo relaying: Achievable rates under limited dynamic range," 2011. [Online]. Available: arXiv:1111.2618 1.2

[15] B. Radunovic, D. Gunawardena, P. Key, A. P. N. Singh, V. Balan, and G. Dejean, "Rethinking indoor wireless: Low power, low frequency, full duplex," Microsoft Technical Report, Tech. Rep., 2009. 1.4.1

[16] M. Jain, J. I. Choi, T. Kim, D. Bharadia, S. Seth, K. Srinivasan, P. Levis, S. Katti, and P. Sinha, "Practical, real-time, full duplex wireless," 'http://www.sigmobile.org/mobicom/2011/slides/124-practical-slides.pdf', New York, NY, USA, pp. 301–312, 2011. [Online]. Available: http://doi.acm.org/10.1145/2030613.2030647 1.4.1, 6.2.2

[17] W. Slingsby and J. McGeehan, "Antenna isolation measurements for on-frequency radio repeaters," in *Antennas and Propagation, 1995., Ninth International Conference on (Conf. Publ. No. 407)*, apr 1995, pp. 239 –243 vol.1. 1.4.2

[18] C. Anderson, S. Krishnamoorthy, C. Ranson, T. Lemon, W. Newhall, T. Kummetz, and J. Reed, "Antenna isolation, wideband multipath propagation measurements, and interference mitigation for on-frequency repeaters," in *Southeast-Con, 2004. Proceedings. IEEE*, mar 2004, pp. 110 – 114. 1.4.2

[19] H. Suzuki, K. Itoh, Y. Ebine, and M. Sato, "A booster configuration with adaptive reduction of transmitter-receiver antenna coupling for pager systems," in *Vehicular Technology Conference, 1999. VTC 1999 - Fall. IEEE VTS 50th*, vol. 3, 1999, pp. 1516 –1520 vol.3. 1.4.2

[20] T. Riihonen, S. Werner, R. Wichman, and Z. Eduardo, "On the feasibility of full-duplex relaying in the presence of loop interference," in *Signal Processing Advances in Wireless Communications, 2009. SPAWC '09. IEEE 10th Workshop on*, june 2009, pp. 275 –279. 1.4.2

[21] T. Riihonen, S. Werner, and R. Wichman, "Hybrid full-duplex/half-duplex relaying with transmit power adaptation," *Wireless Communications, IEEE Transactions on*, vol. 10, no. 9, pp. 3074 –3085, september 2011. 1.4.2

[22] Ruckus, "802.11n sector access point," Online. "http://www.ruckuswireless.com/products/zoneflex-outdoor/7762-s", 03 2012. 2.1

[23] V. K. Garg, *Wireless Communications and Networking*. Morgan Kaufmann, 2007. 2.1

[24] A. Sahai and G. Patel and A. Sabharwal, "Asynchronous Full-duplex Wireless," in *Proc. of IEEE COMSNETS*, Bangalore, India, 2012. 2.2

[25] C. Tong, *Advanced Materials and Design for Electromagnetic Interference Shielding*. CRC Press, 2009. 2.2, 2.2

[26] W. L. Stutzman and G. A. Thiele, *Antenna Theory and Design*. John Wiley and Sons Inc., 1998. 2.3, 2.3

[27] Emmerson and Cuming Microwave. Eccosorb free-space absorber. [Online]. Available: http://www.eccosorb.com/products-eccosorb-an.htm 3.1.2, 4.2

[28] L-Com. 2.4 GHz 90 degree dual-polarized antenna. [Online]. Available: http://www.l-com.com/item.aspx?id=22078 3.2.1, 4.1

[29] R. Ramanathan, J. Redi, C. Santivanez, D. Wiggins, and S. Polit, "Ad hoc networking with directional antennas: a complete system solution," *Selected Areas in Communications, IEEE Journal on*, vol. 23, no. 3, pp. 496 – 506, march 2005. 3.2.1

[30] Y. Liu, Y. Yuan, and K. Contractor, "A method to achieve a dual-polarized planar inverted f antenna (pifa)," in *Mobile Technology, Applications and Systems, 2005 2nd International Conference on*, nov. 2005, pp. 1 –5. 3.3.2

[31] A. Amiri Sani, L. Zhong, and A. Sabharwal, "Directional antenna diversity for mobile devices: characterizations and solutions," in *MobiCom '10: Proceedings of the sixteenth annual international conference on Mobile computing and networking*. New York, NY, USA: ACM, 2010, pp. 221–232. 3.3.2

[32] L-Com. 2.4 GHz 9 dBi yagi antenna. [Online]. Available: http://www.l-com.com/item.aspx?id=21852 4.1

[33] ——. 2.4 GHz 6 dBi omnidirectional antenna. [Online]. Available: http://www.l-com.com/productfamily.aspx?id=6409 4.1

[34] WARPLab framework. [Online]. Available: http://warp.rice.edu/trac/wiki/WARPLab 5.1.1

[35] Wireless open-access research platform (WARP). [Online]. Available: http://warp.rice.edu/ 5.1.1

[36] M. Inc. MAX2829 single-/dual-band 802.11a/b/g world-band transceiver ICs. [Online]. Available: http://www.maxim-ic.com/datasheet/index.mvp/id/4532 5.1.2.1

[37] H. Arslan and H. Mahmoud, "Error vector magnitude to SNR conversion for nondata-aided receivers," *Wireless Communications, IEEE Transactions on*, vol. 8, no. 5, pp. 2694 –2704, May 2009. 5.1.2.2

[38] E. Everett, M. Duarte, C. Dick, and A. Sabharwal, "Empowering full-duplex wireless communication by exploiting directional diversity," in *Asilomar Conference on Signals, Systems and Computers*, October 2011. 5.2

[39] Telex. 2405aa 5 dBi patch antenna. [Online]. Available: http://www.wlanantennas.com/datasheets/wlan_antenna_2405.pdf 5.2.1, 5.2.2.2

[40] T. S. Rappaport, *Wireless Communications: principles and practice*. Prentice Hall, 1996. 6.1

[41] C. Perez-Vega and J. L. G. Garcia, "A simple approach to a statistical path loss model for indoor communications," in *27th European Microwave Conference*, vol. 1, sept. 1997, pp. 617 –623. 6.1