

Passive Self-Interference Suppression for Full-Duplex Infrastructure Nodes

Evan Everett, Achaleshwar Sahai, and Ashutosh Sabharwal

Abstract— Recent research results have demonstrated the feasibility of full-duplex wireless communication for short-range links. Although the focus of the previous works has been active cancellation of the self-interference signal, a majority of the overall self-interference suppression is often due to passive suppression, i.e., isolation of the transmit and receive antennas. We present a measurement-based study of the capabilities and limitations of three key mechanisms for passive self-interference suppression: directional isolation, absorptive shielding, and cross-polarization. The study demonstrates that more than 70 dB of passive suppression can be achieved in certain environments, but also establishes two results on the limitations of passive suppression: (1) environmental reflections limit the amount of passive suppression that can be achieved, and (2) passive suppression, in general, increases the frequency selectivity of the residual self-interference signal. These results suggest two design implications: (1) deployments of full-duplex infrastructure nodes should minimize near-antenna reflectors, and (2) active cancellation in concatenation with passive suppression should employ higher-order filters or per-subcarrier cancellation.

I. INTRODUCTION

Conventional wireless devices operate in half-duplex mode to avoid the high-powered self-interference that is generated when transmission and reception coexist in time and frequency. Recent results [1]–[12], however, have demonstrated the feasibility of short-range full-duplex wireless communication by mitigating self-interference through a combination of passive suppression and active cancellation of the self-interference. Passive suppression is any technique to electromagnetically isolate the transmit and receive antennas, while active suppression is any technique to exploit a node’s knowledge of its own transmit signal to cancel the self-interference. Active cancellation, which has been studied extensively in the literature, functions by injecting a cancellation waveform into the receive signal path to null the self-interference. The cancellation waveform can be injected over a transmission line as in [4], [6], [8], [11], [12] or over the air using extra antennas as in [1], [5], [9], [10]. Passive suppression, however, has been studied only sparsely. With the aim to develop a better understanding of the capabilities and limitations of passive self-interference suppression, this paper presents a

measurement-based characterization of passive suppression mechanisms for full-duplex infrastructure nodes.

Our motivation for studying passive suppression stems from the observation that passive suppression accounts for a large portion of the total self-interference suppression in existing full-duplex designs. For example, of the 85 dB of self-interference suppression achieved in the full-duplex design in [11], 65 dB is attributed to passive suppression and only 20 dB to active cancellation. Moreover, recent characterizations of active cancellation [13]–[15] have shown that the amount of self-interference that can be actively cancelled is limited by fundamental radio impairments such as phase noise [13] and limited dynamic range [14], [15], but it is not known whether passive suppression will also encounter fundamental bottlenecks and what those bottlenecks may be. Thus, an opportunity exists to improve full-duplex performance through a richer understanding of capabilities and limitations of passive suppression mechanisms.

A. Contribution

Our first (and primary) contribution is an experimental characterization of passive self-interference suppression methods for full-duplex infrastructure nodes. The characterization focuses on three key passive suppression mechanisms: (1) directional isolation – the use of directional antennas such that the gain of the transmit antenna is low in the direction of the receive and visa versa, (2) absorptive shielding – the use of lossy materials to attenuate the self-interference, and (3) cross-polarization – the use of transmit and receive antennas in orthogonal polarization states. The characterization has two main goals, (a) to quantify the performance and limitations of the three mechanisms and combinations among the three mechanisms, and (b) to study the effect of passive suppression on the characteristics of the self-interference channel. To accomplish these goals, the three mechanisms were evaluated both in a best-case environment of an anechoic chamber, and a worst-case environment of a highly reflective room with metal walls. In each environment the self-interference channel was measured for several different antenna configurations both with and without the three passive suppression mechanisms in place. The characterization establishes two main results which are summarized below.

Result 1: Environmental reflections present a fundamental bottleneck in the amount of passive suppression that can be achieved. Comparing the measurements in the anechoic chamber to those in the reflective room, it is observed that although the passive mechanisms are quite effective in suppressing the direct path between antennas, none can predictably attenuate

The authors would like to thank the Antenna and Wireless Systems Branch of NASA’s Johnson Space Center for both technical advice and the use of the Antenna Test Facility. In particular, the authors would like to recognize Michael Khayat for his guidance in experiment design, and Wayne Cope for his EM measurement expertise. The authors also thank Jingwen Bai and Guarav Patel for their crucial assistance.

This work was partially supported by NSF Grants CNS 0923479, CNS 1012921 and CNS 1161596 and NSF Graduate Research Fellowship 0940902.

E. Everett, A. Sahai and A. Sabharwal are with the Department of Electrical and Computer Engineering, Rice University, Houston, TX, 77005 USA. e-mail: {evan.everett, as27, ashu}@rice.edu.

reflected self-interference paths. The design implication is that, in order to achieve high-levels of passive suppression, deployers of full-duplex infrastructure should seek to minimize nearby reflectors.

Result 2: Passive self-interference suppression, in general, causes a decrease in the coherence bandwidth of the residual self-interference channel. Since the passive suppression operates on the direct path between transmit and receive antennas, but not necessarily reflected paths, passive suppression can transition the self-interference channel from a predominantly line-of-sight channel to a multi-path dominated channel, where frequency selectivity occurs due to constructive and destructive combinations of reflected paths. The design implication is that when active cancellation is employed in concatenation with passive suppression, the active canceler should be designed to handle a frequency-selective self-interference signal.

As a second contribution, we present a prototype full-duplex infrastructure node, implemented in the WARPLab [16] framework, that combines the passive suppression mechanisms studied in the above characterization with the existing active cancellation mechanisms studied in [7], [11]. We evaluate the performance of this design in an outdoor full-duplex up-link, and see that total self-interference suppression exceeding 90 dB can be achieved, enabling the prototype to outperform a comparable half-duplex link even at ranges exceeding 100 m.

The above two results, when combined with previous works, paint a more complete picture of self-interference mitigation. Passive suppression is the first line-of-defense against self-interference, but is effective only against the direct path. The residual multi-path self-interference signal can be further suppressed by a per-subcarrier active cancellation strategy [7], [11] that is robust to frequency selectivity. Unfortunately, active cancellation also encounters bottlenecks due to RF impairments such as phase noise [13] and limited dynamic range [14], [15]. Because both passive and active suppression encounter limitations, it is important that full-duplex designs reach their full potential in both the active *and* passive suppression regimes.

B. Prior Art

Many of the full-duplex designs in the literature have employed one or more passive suppression method. In the full-duplex designs of [1], [4]–[7] the only passive suppression mechanism utilized is physical separation of transmit and receive antennas, i.e., free-space path loss. In [8], [11] heuristic strategies for laptop antenna placement are studied including “device-induced path loss,” in which transmit and receive antennas are shielded by laptop hardware such as the screen. Such techniques were used in [11] to achieve as much as 65 dB of passive suppression. Passive suppression for full-duplex relay systems in which the transmit and receive antennas are not collocated, and large amounts of path loss can be leveraged, are studied in [17]–[19]. In [17], [18] directional isolation and 4-5 m of antenna separation are leveraged to achieve 75-90 dB of passive suppression. In [19] an outdoor-to-indoor relay systems is studied, in which directional isolation and cross-polarization are leveraged. In a deployment

where the antennas are on opposite sides of a wall, 48 dB of passive suppression is observed, and when the indoor antenna is moved into the interior of the building, benefitting from large path loss, 60-80 dB of passive suppression is observed. In contrast to [17]–[19], our study focuses on more compact full-duplex deployments in which the antenna separation does not exceed 50 cm. Moreover, our study, which includes directional isolation, cross-polarization and absorptive shielding, is to our knowledge the first comprehensive study of passive self-interference suppression for wireless full-duplex systems. In particular, none of the previous designs have quantified the impact of environmental reflections on the performance of passive suppression.

C. Organization of Paper

The organization of the rest of the paper is as follows. In Section II we describe the experiment design in detail. In Section III we present the outcomes of the experiment, with special focus on the two main results listed above. In Section IV we analyze the impact of passive suppression on capacity. In Section V we present a prototype and evaluate the performance of combined passive and active suppression in a wideband OFDM full-duplex physical layer. In Section VI we have some system-level discussion on implementing passive suppression at infrastructure nodes. We conclude in Section VII.

II. EXPERIMENT DESIGN

The goal of the experiments presented in this paper is to characterize the performance of three key mechanisms for passive self-interference suppression in full-duplex infrastructure nodes, and to study how the three mechanisms affect the self-interference channel. To accomplish this goal, a network analyzer was used to characterize the self-interference channel between transmitting and receiving antennas as the three passive suppression mechanisms were applied to several different antenna configurations both in a reflective environment and in a nonreflective environment. The following is a detailed description of the experiment.

A. Experiment Setup

We begin by describing the directional antenna configurations studied, and then describe how absorptive shielding and cross-polarization are leveraged within in each configuration.

1) *Implementation of Directional Isolation:* The five infrastructure antenna configurations in Figure 1 were used to study passive self-interference suppression via directional isolation. The configurations differ in the directionality of the antennas and in their relative orientation. Configuration I, illustrated in Figure 1(a), uses 90° beamwidth¹ antennas which are pointed with 90° beam separation, so that there is little overlap in the coverage patterns. Configuration II, shown in

¹Throughout 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. The conical antennas patterns in Figure 1 are only figurative: actual patterns have a gradual rolloff in gain and also have side lobes.

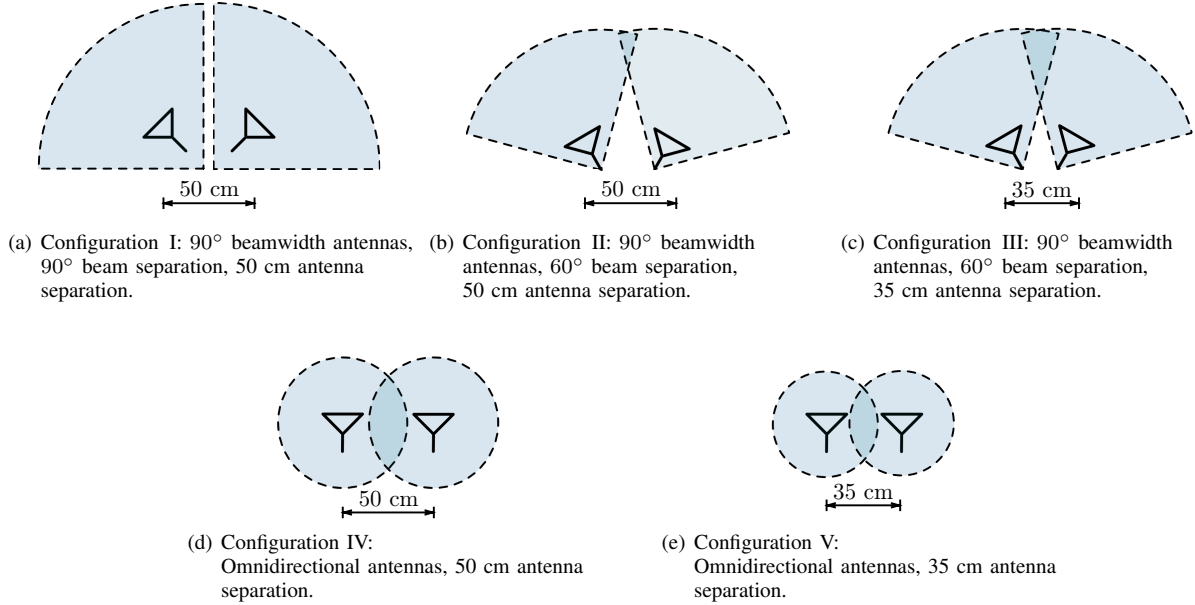


Fig. 1: Five architectures evaluated in the passive suppression characterization. We do not designate one antennas as the transmitter and one as receiver as these roles exchange.

Fig 1(b), also uses 90° beamwidth antennas, but with a beam separation of 60°, so that there is 30° overlap between the coverage zones. In both Configurations I and II the separation between antennas is 50 cm. Configuration III, shown in Figure 1(c), is identical to Configuration II, except that the separation between antennas is scaled down to 35 cm. The 90° beamwidth antennas used in Configurations I, II, and III are HG2414DP-090 panel antennas from L-com [20] designed for outdoor sectorized WiFi access points. For comparing to cases without directional isolation, Configurations IV and V, shown in Figures 1(d) and 1(e), were also studied. The 50 cm separation between antennas in Configuration IV is the same as in Configurations I and II, the 35cm separation in Configuration V is the same as in Configuration III. The omnidirectional antennas used in Configurations IV and V are HGV-2406 outdoor WiFi antennas from L-com [21].

2) *Implementation of Absorptive Shielding:* Absorptive shielding is realized by placing a slab of RF absorber material between the two antennas as illustrated in Figure 2. The material used in our experiments is Eccosorb AN-79 [22] 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 can provide up to 25 dB of absorption.

3) *Implementation of Cross-polarization:* The L-Com HGV-2406 antennas used in the directional-antenna configurations are dual-polarized as indicated in Figure 2, which means they have two ports: one excites a horizontally polarized mode, and the other a vertically polarized mode. By measuring the coupling between both co-polarized and cross-polarized ports, we can study the impact of cross-polarization on the self-interference channel. Most commercial omnidirectional anten-




Fig. 2: Shielding via RF absorber and cross-polarization via dual-polarized antennas.

nas support only a single polarization mode,² and such is the case for the L-Com HGV-2406 antennas used in the omnidirectional configurations of Figure 1, which only support vertical polarization. Thus passive suppression via cross-polarization was not studied in the omnidirectional configurations.

B. Environments

In order for the measurements to be as repeatable as possible, the first stage of the passive suppression characterization was performed in a shielded anechoic chamber to minimize reflections and interference from other sources. After the measurements were performed in the anechoic chamber, they were repeated in a highly reflective room to observe the effect of environmental reflections. The room used was roughly 30 ft. × 60 ft. with metal walls and a low metal ceiling, intended to represent a worst case reflection environment for practical deployments.

²Most commercial omnidirectional antennas are implemented as an electric dipole or an electric monopole, which only support a vertically polarized mode. Directional antennas, however, are commonly implemented as circular or rectangular patches, which can support two modes: one with vertical polarization and the other with horizontal polarization.

C. Measurements and Metrics

1) *Measurement Procedure*: For notational convenience let the variable,

$$\text{mech} \in \{\emptyset, D, A, C, DA, DC, AC, DCA\}, \quad (1)$$

denote a combination of passive suppression mechanisms, where \emptyset denotes no passive suppression, D denotes directional isolation, A denotes absorptive shielding, and C denotes cross-polarization. Note that the antenna configuration constrains which combinations are possible. Since directional antennas are used in Configurations I, II, and III, the possible values of mech are $\{D, DA, DC, DCA\}$ for these configurations. Similarly the possible values for mech in Configurations IV and V are $\{\emptyset, A\}$, since the antennas are omnidirectional and not dual-polarized. Therefore, for each of the directional-antenna configurations shown in Figure 1, the self-interference channel between antennas was measured for the following combinations of passive self-interference suppression mechanisms: directional isolation only (mech = D), directional isolation plus absorptive shielding (mech = DA), directional isolation plus cross-polarization (mech = DC), and directional isolation plus absorptive shielding plus cross-polarization (mech = DCA). For the omnidirectional configurations, the combination of mechanisms measured were no suppression (mech = \emptyset) and absorptive shielding (mech = A). For each combination of applied suppression mechanisms, mech, we directly measured the frequency response $H^{(\text{mech})}(f)$ of the self-interference channel between the two antennas using an Agilent E8364B [23] general-purpose network analyzer (PNA). The channel response, $H^{(\text{mech})}(f)$, was obtained by connecting the transmit and receive antenna ports to the network analyzer, configured to record the S-parameters of a two-port network at 20,000 uniformly spaced frequency components in the WiFi ISM band, $f \in [2.40, 2.48]$ GHz. Each frequency sample was averaged over 20 measurements to reduce noise.

2) *Metrics*: Let $\bar{\alpha}_P^{(\text{mech})}$ denote the amount of passive suppression achieved by the combination of mechanisms mech. We define $\bar{\alpha}_P^{(\text{mech})}$ to be ratio of the transmit power to the power of the self-interference incident on the receiver, which is the inverse of the power in the frequency response, $H^{(\text{mech})}(f)$. Thus we compute $\bar{\alpha}_P^{(\text{mech})}$ from the measured frequency response according to

$$\bar{\alpha}_P^{(\text{mech})} = \left(\frac{1}{N} \sum_{f=f_{\min}}^{f_{\max}} |H^{(\text{mech})}(f)|^2 \right)^{-1}. \quad (2)$$

In addition to studying the amount of suppression that different mechanisms attain, we also want to study how passive suppression affects the frequency-selectivity of the self-interference channel. The metric most often used to quantify the frequency-selectivity of a wireless channel is the coherence bandwidth, B_C . The definition of coherence bandwidth we adopt is the range of frequencies over which the channel response is at least 90% correlated [24]. The coherence bandwidth is difficult to measure directly, but can be computed from the power delay profile of the channel, $P(\tau)$, which can be easily measured. In our study, the power-delay profile is

obtained by taking the inverse discrete time Fourier transform (DTFT) of the measured frequency response. The RMS delay spread of the channel, σ_τ , can be computed from the power delay profile, which is related to the 90% coherence bandwidth via the approximation [25]

$$B_C \approx \frac{0.02}{\sigma_\tau}. \quad (3)$$

Let $P_0(\tau)$ denote the normalized power delay profile, that is $P_0(\tau) = \frac{P(\tau)}{\int_0^\infty P(\tau) d\tau}$. The mean delay, μ_τ is given by $\mu_\tau = \int_0^\infty \tau P_0(\tau) d\tau$, and the RMS delay spread, σ_τ , can then be computed as [24].

$$\sigma_\tau = \sqrt{\int_0^\infty (\tau - \mu_\tau)^2 P_0(\tau) d\tau}. \quad (4)$$

Taking σ_τ , computed as shown above from the measured power delay profile, and plugging into Equation (3) allows approximation of the coherence bandwidth from the empirical channel response. If we think of the normalized power delay, $P_0(\tau)$, as the ‘‘distribution’’ on the delay τ then the above equations are easy to interpret.

III. EXPERIMENT RESULTS

Table I shows the passive suppression, $\bar{\alpha}_P^{(\text{mech})}$, measured for all of the configurations, both in the anechoic chamber and in the reflective room. Table I(a) contains the suppression measurements for the directional-antenna configurations in the anechoic chamber, and Table I(b) compares these measurements to those made in the reflective room. Similarly, Table I(c) shows the passive suppression measurements for the omnidirectional configurations in the anechoic chamber, and Table I(d) shows the measurements for the omnidirectional configurations in the reflective room. In parenthesis under each entry is the uncertainty in each measurement due to network analyzer noise and calibration error³. See Appendix A for an explanation of how these uncertainty values are obtained. For easier visualization, we also plot the directional-configuration data shown in Tables I(a) and I(b) in Figure 3. To indicate how far the worst-case passive suppression at any particular frequency can differ from average passive suppression over the frequency band, the minimum suppression over the frequency band is indicated by a one-sided ‘‘error bar’’ (and does *not* correspond to the measurement uncertainty given in Table I).

First consider the passive suppression achieved for Configuration I in the low-reflection environment of the anechoic chamber, shown in the first cluster of Figure 3(a). Looking at the first bar in Figure 3(a) we see that when Configuration I is used without any absorptive shielding or cross-polarization, the passive suppression is 45 dB. Comparing to the omnidirectional Configuration IV in the first column of Table I(c), we see that the directional isolation (D) that comes from using the 90° beamwidth antennas with 90° beam separation provides

³The uncertainty values in Table I account for noise and calibration bias of the measurement equipment. The extent to which the fabrication of the antenna configuration is repeatable will also affect the uncertainty, but is very difficult to quantify. In the case of the anechoic chamber measurements, the reported uncertainties do not account for possible imperfections in the anechoic chamber, which is also hard to empirically quantify.

TABLE I: Passive Suppression Measurements

(a) <u>Anechoic Chamber</u>				(b) <u>Reflective Room</u>		
Applied Mechanism (mech)	$\bar{\alpha}_P^{(\text{mech})}$ Config I	$\bar{\alpha}_P^{(\text{mech})}$ Config II	$\bar{\alpha}_P^{(\text{mech})}$ Config III	$\bar{\alpha}_P^{(\text{mech})}$ Config I	$\bar{\alpha}_P^{(\text{mech})}$ Config II	$\bar{\alpha}_P^{(\text{mech})}$ Config III
Directional (D)	45.3 dB (0.12 dB)	39.7 dB (0.10 dB)	45.1 dB (0.12 dB)	37.7 dB (0.10 dB)	36.0 dB (0.09 dB)	37.1 dB (0.09 dB)
D + Absorber (DA)	55.4 dB (0.22 dB)	50.9 dB (0.16 dB)	55.9 dB (0.24 dB)	39.4 dB (0.10 dB)	38.1 dB (0.10 dB)	37.4 dB (0.10 dB)
D + Cross-polarization (DC)	63.9 dB (0.51 dB)	58.3 dB (0.28 dB)	56.9 dB (0.26 dB)	45.1 dB (0.12 dB)	44.9 dB (0.12 dB)	44.4 dB (0.11 dB)
D + Cross-pol + Absorber (DCA)	73.8 dB (1.5 dB)	72.5 dB (1.2 dB)	72.4 dB (1.2 dB)	45.9 dB (0.12 dB)	44.9 dB (0.12 dB)	45.7 dB (0.12 dB)

(c) <u>Anechoic Chamber</u>		(d) <u>Reflective Room</u>		(e) <u>Legend</u>	
Applied Mechanism (mech)	$\bar{\alpha}_P^{(\text{mech})}$ Config IV	$\bar{\alpha}_P^{(\text{mech})}$ Config V	$\bar{\alpha}_P^{(\text{mech})}$ Config IV	$\bar{\alpha}_P^{(\text{mech})}$ Config V	$\bar{\alpha}_P^{(\text{mech})}$ Configuration
None (\emptyset)	27.9 dB (0.08 dB)	24.5 dB (0.08 dB)	27.6 dB (0.08 dB)	25.1 dB (0.08 dB)	Value (Uncertainty)
Absorber (A)	48.0 dB (0.14 dB)	48.2 dB (0.14 dB)	40.1 dB (0.10 dB)	41.1 dB (0.10 dB)	

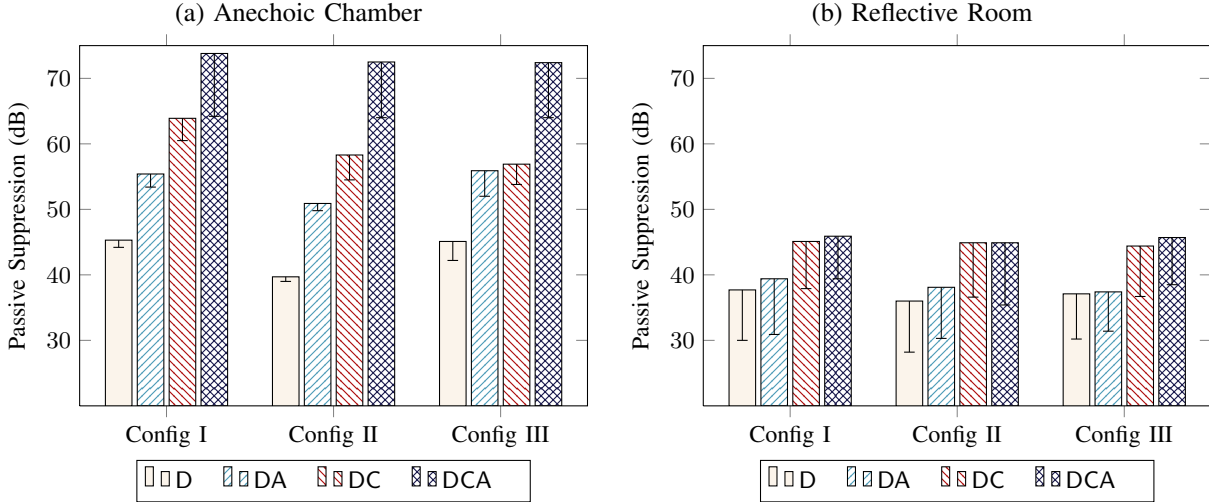


Fig. 3: Passive suppression measurements with directional isolation (D), directional isolation + absorptive shielding (DA), directional isolation + cross-polarization (DC), and directional isolation + cross-polarization shielding + cross-polarization (DCA). The worst-case (minimum) suppression over the frequency band is indicated by the error bar.

around 17 dB of passive suppression in addition to the 28 dB of path loss between antennas. Adding absorptive shielding (DA) gives an extra 10 dB of suppression for 55 dB total suppression. Similarly, adding cross-polarization (DC) gives an extra 18 dB of suppression for a total suppression of 66 dB. When all the mechanisms are applied in tandem, directional isolation absorptive shielding, and cross-polarization together (DCA) yield 74 dB total suppression. A similar trend is followed for the other directional-antenna configurations

shown in Figure 3(a).⁴ We see in Figure 3(a) that, when in the low-reflection environment of the anechoic chamber, leveraging directional isolation, cross-polarization, absorptive shielding, in tandem (DCA) provides more than 70 dB of passive self-interference suppression in all three directional-antenna configurations. In the reflective room, however, the total suppression is much less, leading to the following result.

⁴We see in the first row of Tables I(a) and Tables I(b) that the suppression with directional isolation (D) is greater for Configuration II than Configuration III, both in the anechoic chamber and in the reflective room. This is surprising because the antennas are closer in Configuration III than Configuration II. We believe this is an artifact of the near-field coupling between the specific antennas used, not evidence of a general trend.

A. Effect of Environmental Reflections

Result 1: Environmental reflections limit the amount of passive self-interference suppression that can be achieved.

Comparing Figures 3(a) and 3(b), we see that in every case the achieved suppression is less in the reflective room than in the anechoic chamber. Moreover, we observe that the incremental contribution of each applied mechanism is much less in the reflective room than in the anechoic chamber. In Configuration I, for example, adding absorber improves the achieved suppression by 10 dB in the anechoic chamber, but provides less than 3 dB added suppression in the reflective room.

Reason for Result 1: Environmental reflections cause a bottleneck in passive self-interference suppression, because although the proposed mechanisms are very effective in suppressing the *direct path* between the transmit and receive antennas, they are not necessarily effective in suppressing *reflected* self-interference paths. Once the direct-path self-interference is suppressed to the extent that it is no longer the dominant path, further efforts to passively suppress self-interference may prove ineffective. Directional isolation is effective in suppressing the direct path, but will not suppress a reflected path that exits within the beamwidth of the transmit antenna and enters within the beamwidth of the receive antenna. Similarly, absorptive shielding certainly suppresses the direct path, but will do nothing to suppress a self-interference path that includes a reflector outside the node. Finally, reflections do not, in general, preserve polarization. Thus cross-polarizing the transmit and receive antennas suppresses the direct path, but not necessarily the reflected paths.

We will now observe the time-domain response of the measured self-interference channels to verify the hypothesis that suppression of the direct path but not reflected paths is the reason for Result 1. We expect that for a given combination of passive suppression mechanisms the time response of the self-interference channel measured in the anechoic chamber will be nearly the same as that measured in the reflective room for an initial interval corresponding to the direct path delay, but afterward we expect the tail of the time response to be much more powerful in the reflective room than in the anechoic chamber. Moreover, we expect that when more passive suppression mechanisms are applied, the tail of the time response will become dominant over the direct path, explaining why adding more passive suppression mechanisms often does not improve the achieved suppression in the reflective room.

The hypothesized behavior is exactly what we observe in Figure 4, which compares the time domain channel responses measured in the chamber to those measured in the reflective room for four different combinations of passive suppression mechanisms. The time responses for omnidirectional antennas with no suppression mechanisms in place (Configuration IV, $\text{mech} = \emptyset$) are shown in Figure 4(a). In this case, the self-interference is dominated by the direct path in both environments, explaining why the value the average suppression, $\bar{\alpha}_P^{(\text{mech})}$, listed in Table I for this configuration is the same in both environments. Figure 4(b) shows that when we apply directional isolation (Configuration I, $\text{mech} = D$), the direct

path is significantly suppressed, but the reflection floor is not. We saw in Figure 3 that adding absorber between the directional antennas gives 10 dB improvement in the anechoic chamber, but less than 3 dB improvement in the reflective room; Figure 4(c) explains why. Comparing Figure 4(b) and Figure 4(c), we see that adding absorber suppresses the *direct path* by more than 10 dB in *both* environments, but in the reflective environment the direct path is no longer the dominant path after adding absorber, while in the anechoic chamber the direct path remains the dominant path even after adding the absorber. Thus the full benefit of adding the absorber is not realized in the reflective environment. Finally, we see in Figure 4(d) that adding cross-polarization in addition to absorptive shielding adds another 10 dB of suppression in the anechoic chamber but does little in the reflective room, because reflections have already become the bottleneck. ■

Design Implications of Result 1: All other factors being equal, a full-duplex infrastructure node should be deployed as far as possible from potential reflectors. In an environment that is unavoidably reflective, resources should not be expended attempting to passively suppress self-interference beyond the reflection strength.

B. Impact of Passive Suppression Frequency Selectivity of Self-interference Channel

Result 2: Passive self-interference suppression generally decreases the coherence bandwidth (i.e. increases the frequency selectivity) of the residual self-interference channel.

Figure 5 shows measurements of the self-interference channel frequency response, $H^{(\text{mech})}(f)$, that illustrate the trend of increased frequency selectivity with increased passive suppression. Figure 5(a) plots the channel responses measured for Configuration I in the anechoic chamber. We see in Figure 5(a) that the variation of the self-interference channel gain with frequency increases as passive suppression mechanisms are applied.⁵ The same trend was observed in the other four configurations measured in the anechoic chamber. Figure 5(b) shows the channel responses measured for configuration I in the reflective room. We see in Figure 5(b) that with directional isolation alone, the self-interference channel is already quite frequency selective. Adding cross-polarization and absorptive shielding in Figure 5(b) does not significantly increase the frequency selectivity, but nor does it increase the average self-interference suppression, since the direct path has already been suppressed to the reflection floor as discussed in Result 1. Figure 5(b) hints that there is nothing special about directional isolation, cross-polarization, or absorptive shielding that causes the increase in frequency selectivity, but the fact that each mechanism serves to reduce the *direct path*.

Reason for Result 2: Result 2 is nearly a corollary of Result 1. Since the dominant direct path is suppressed, the

⁵The channel responses in Figure 5(a) also may appear to be growing more noisy as more suppression mechanisms are applied. The appearance of greater noise is an artifact of the dB scaling of the plot making the measurement noise of the much lower-power suppressed channel appear larger. The thermal noise floor in this measurement was -100 dB, as computed from the manufacturer data sheet [23] given the 20KHz IFBW and averaging factor of 20 used in the measurement.

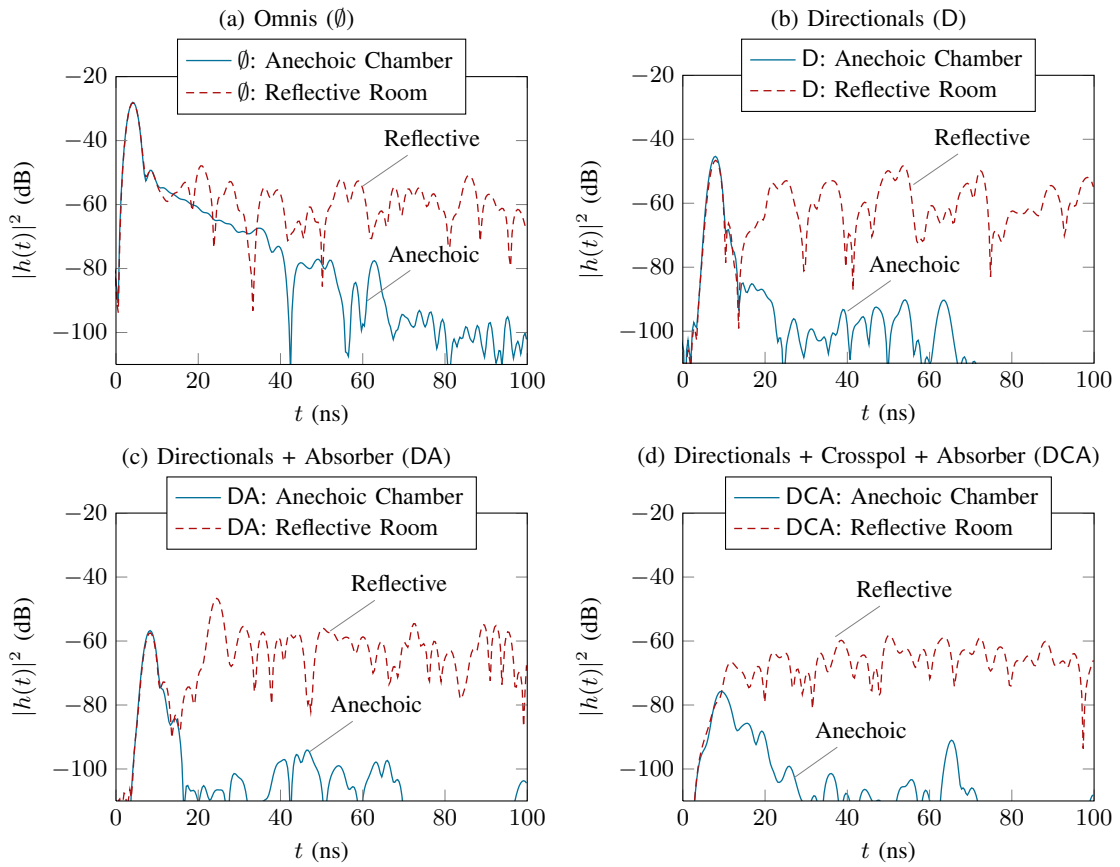


Fig. 4: Comparing time responses measured in an anechoic chamber against the time responses measured in a reflective room.




Fig. 5: Self-interference channel responses illustrating increased frequency selectivity with increased passive suppression.

residual channel after passive suppression is the superposition of many reflected paths. Passive suppression transforms the self-interference channel from a high-power line-of-sight-dominated channel to a low-power multi-path channel. It is well known that multi-path channels are frequency selective due to paths combining constructively at some frequencies and destructively at others. We will more precisely give the reason for Result 2 in two steps: first we will consider a simple analytical example that illustrates why the coherence bandwidth of the self-interference channel decreases as the direct path is suppressed. Second, we will present measurements of the coherence bandwidth for all five configurations in both environments to give further evidence of Result 2 and reinforce the intuitions gained from the analytical example.




Fig. 6: Example of a power-delay profile with both direct and reflected paths. In this example $P_D = 25$ dB, $P_R = 5$ dB \Rightarrow DRR = 20 dB, and $T_R = 20$ s.

1) *Analytical example:* Consider the “two-level” self-interference power delay profile shown in Figure 6, in which the first tap which contains the direct path and has power P_D , and the remaining $T_R - 1$ taps, all due to reflected paths, have power P_R , where T_R is the reflection duration of the environment. Denote the ratio of the direct-path power to the reflected-path power as the direct-to-reflected ratio, $\text{DRR} \equiv \frac{P_D}{P_R}$. As passive suppression mechanisms are employed, we know from Result 1 that DRR will decrease, since the direct path will be suppressed while the reflective paths maintain the same power.

Figure 7 plots the RMS delay spread, σ_τ , and the corresponding coherence bandwidth, $B_C = 0.02/\sigma_\tau$, vs. DRR for four different values of T_R . Suppressing the direct path (reducing DRR) will lead to a heavier tail in the normalized power delay profile, $P_0(\tau)$, and hence a larger RMS delay spread and smaller coherence bandwidth. When the passive suppression is such that the power in the direct path is the same as the power in the reflected paths, the RMS delay spread is maximum. Once $\text{DRR} < 0$ dB, the reflected portion of the power delay profile becomes dominant, so that further suppression of the direct path has little impact on the RMS delay spread and coherence bandwidth.

In summary, the analytical example shows the following trend. When no passive suppression is applied, the direct-path self-interference will be significantly stronger than reflected paths, and the coherence bandwidth will be relatively

high, such that the self-interference channel will be relatively frequency-flat. As passive suppression is applied, the reflected paths become a larger portion of the power delay profile, and the coherence bandwidth decreases such that the self-interference channel becomes more frequency-selective. In the limit of very strong passive suppression, the direct path is effectively eliminated, and the coherence bandwidth levels off, because the delay spread converges to the reverberation spread of the environment.

2) *Coherence Bandwidth Measurements:* Now let us check that the measurements in the passive suppression characterization follow the trend illustrated in the above analytical example. For each self-interference channel measurement, the the direct-path suppression and the coherence bandwidth were computed from the measured frequency response. The resulting tuples of direct-path suppression and coherence bandwidth are shown as a scatter plot in Figure 8(a). The dark red points in Figure 8(a) correspond to measurements made in the anechoic chamber, and the light blue points correspond to the reflective room. The combination of mechanisms used is labeled by the symbol as shown in the legend.⁶

The trend we observe in Figure 8(a) is the same as predicted from our analytical example. The anechoic chamber has a very low reflection floor, around -80 dB, so when no suppression is applied, the coherence bandwidth is large: between 10 and 25 MHz depending on the configuration. If absorptive shielding or cross-polarization is employed: the direct-path suppression is increased to 50-70 dB and the coherence bandwidth shrinks to 2-15 MHz. When both cross-polarization and absorptive shielding are applied together, the direct-path suppression is greater than 70 dB, near the reflection floor of the chamber, hence the coherence bandwidth begins to level off at around 3 MHz. In the reflective room, the trend is the same but much more abrupt, since the reflection floor is much higher: around -45 dB. In the reflective room, only when omnidirectional antennas are used with no suppression mechanisms ($\text{mech} = \emptyset$) is the direct path significantly stronger than the reflected paths (as was seen in Figure 4), in which case the coherence bandwidth is around 2 MHz. Once one or more suppression mechanisms are employed, the direct path is pushed below 40 dB, and the coherence bandwidth converges to the 0.6-0.9 MHz dictated by the long reverberations of the reflective room. Typical deployments will have a reflection floor between the two extremes presented here, but the trend will be the same. Passive suppression can transform a relatively frequency-flat self-interference channel to a highly frequency-selective channel, but the coherence bandwidth will be no worse than natural coherence bandwidth dictated by the reverberation depth of the environment. ■

Design Implications of Result 2: Most full-duplex systems will attempt to further suppress self-interference by employing active cancellation in addition to passive suppression. Self-interference is cancelled by leveraging foreknowledge of the

⁶Not evident in Figure 8 (to reduce clutter and focusing on the trend), is the configuration on which the measurement was made, which accounts for the sometimes wide variation in points corresponding to the same suppression mechanisms.

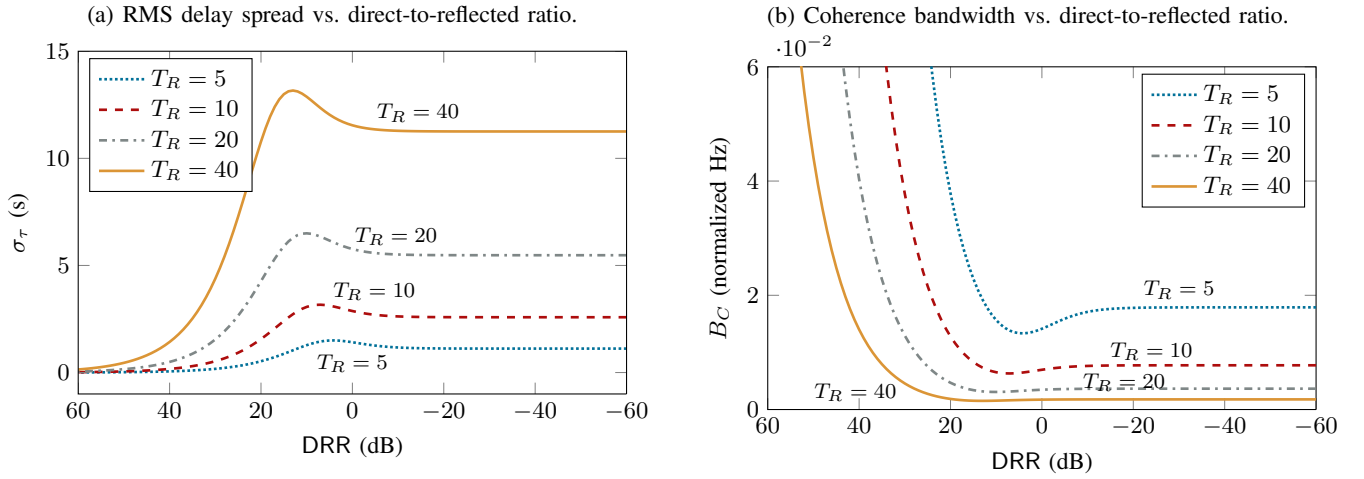


Fig. 7: Impact passive suppression on RMS delay spread and coherence bandwidth.

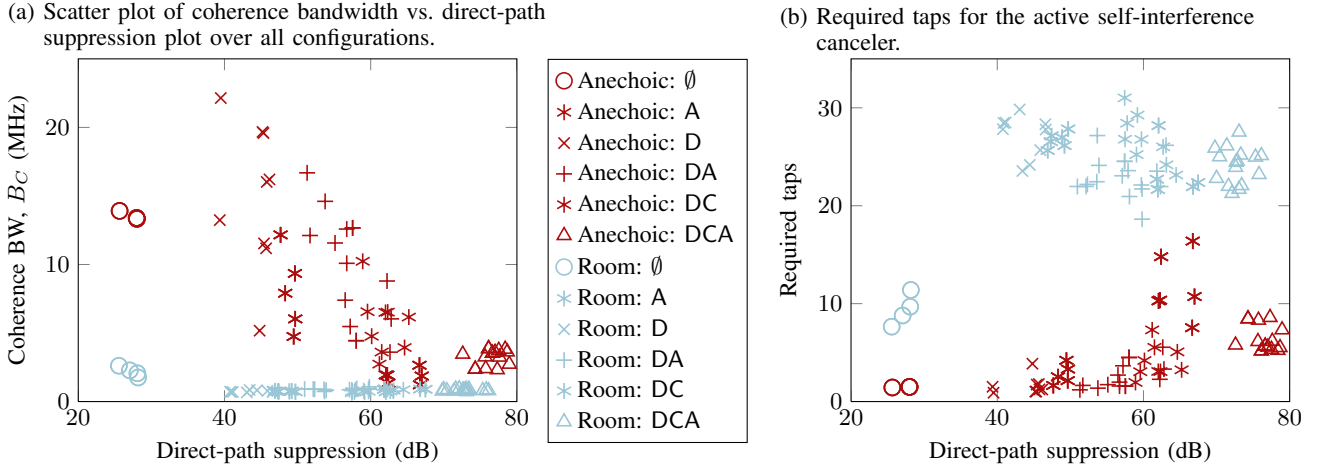


Fig. 8: Influence of direct-path suppression of frequency selectivity

signal to inject a cancellation waveform into the received signal path. To craft the cancellation signal, a copy of the transmit signal must pass through a filter whose response is the negative of the self-interference channel. Result 2 implies that the number of taps in this cancellation filter may increase when passive suppression is employed. Figure 8(b) takes the data shown in Figure 8(a), and computes the required number of taps in the cancellation filter, assuming the transmit signal has a bandwidth of 20 MHz. In the low-reflection environment only one or two taps are needed when the passive suppression is low (25-45 dB), but as the passive suppression gets strong (60-80 dB), 10-15 taps may be needed. In the high-reflection environment, 7-12 taps are needed when the passive suppression is low, but once the passive suppression is more than 40 dB, nearly 30 taps are needed. In an OFDM system, Result 2 implies that when active cancellation is employed in concatenation with strong passive suppression, it may be better to perform the active cancellation on a per-subcarrier basis as in [11], [12] rather than directly canceling the broadband self-interference as in [6].

C. Passive Suppression vs. Active Cancellation

The above discussion raises the following question: is there a tradeoff between passive suppression and active cancellation, and would one rather use one than the other, or always use both together? The results from this paper and those from related works are in agreement that *passive suppression and active cancellation should be applied together whenever possible*. In this paper, it is shown that passive suppression is limited by reflections, and active cancellation is needed to suppress the residual reflected paths. Conversely, in [13]–[15] it is shown that active cancellation is limited by transmitter noise, and thus passive suppression (which can suppress transmitter noise just as well as transmitter signal) is needed to improve performance. In [13] a preprint of this paper is cited and our Result 1 is used to show that although passive suppression can degrade active cancellation performance, it can only increase the *total* self-interference suppression. In the full-duplex system evaluation such as [2], [10], [11], the designs employing both passive suppression and active cancellation together always have the best performance. The summary of ours and the previous work is that although

the gains of passive suppression and active cancellation are not dB-additive, never is it the case that passive suppression alone or active cancellation alone will have better performance than both in concatenation. The question of which approach, passive suppression or active cancellation, is more *efficient* is still open, and is an interesting topic for future study.

IV. IMPACT OF PASSIVE SUPPRESSION ON SYSTEM METRICS

We will now analytically assess the impact that Results 1 and 2 have on the performance of a full-duplex infrastructure node. Consider a full-duplex infrastructure node that receives an uplink signal from one user while also transmitting to another user as shown in Figure 12(b). We will focus on uplink performance, since it is the uplink that suffers from self-interference.

A. Channel Model

The channel model must capture the limitations of both passive suppression and active cancellation. Prior work [13]–[15] has identified transmitter noise as the major factor limiting active cancellation, and this paper identifies environmental reflections as the major factor limiting passive suppression. Thus we model the signal received by the full-duplex infrastructure node at time n as,

$$y[n] = h_S x_S[n] * (x_S[n] + z_S[n]) + h_I[n] * (x_I[n] + z_I[n]) + z_R[n], \quad (5)$$

where x_I and x_S are the signals transmitted by the infrastructure node and the uplink user, respectively, h_I and h_S are the discrete baseband-equivalent impulse responses of the self-interference channel and the signal-of-interest channel, respectively, z_I and z_S are transmitter noise in the self-interference and signal-of-interest, respectively, z_R is the receiver noise, and “*” signifies convolution. For convenience, we assume the signal-of-interest channel, h_S , has only a single tap and the signal-of-interest is free of transmitter noise, as these will have little impact on self-interference suppression. Let the receiver noise and transmitter noise be drawn i.i.d. from circularly-symmetric gaussian distributions of variance N_R and N_T , respectively, and let the transmit signals, x_S and x_I , be subject to average power constraint of P_T . With the above assumptions, the received signal can be rewritten as

$$y[n] = h_S x_S[n] + h_I[n] * (x_I[n] + z_I[n]) + z_R[n], \quad z_I \sim \mathcal{CN}(0, N_T), \quad z_R \sim \mathcal{CN}(0, N_R). \quad (6)$$

Let us extend the “two-level” self-interference power delay profile discussed in section III-B to model the self-interference channel $h_I[n]$. We assume that there are many reflected paths of power P_R that extend over T_R taps and a single direct path of power P_D . Thus we model the first channel tap as Rician fading with κ -factor P_D/P_R and variance $P_D + P_R$, and the remaining $T_R - 1$ taps as Rayleigh fading with variance P_R . We assume the single-tap signal-of-interest channel, h_S , is Rayleigh fading with variance P_S . Per Result 1, passive suppression can reduce P_D , but not P_R . We assume each

channel tap is ergodic block fading with coherence time $T_C \gg 1$, and assume coding can be performed over a large number of coherence times so that the notion of ergodic capacity can be employed [26]. We assume both the user and the infrastructure node have perfect knowledge of the channel responses, and the infrastructure has knowledge of its own transmit signal, x_I , but has no knowledge of its transmitter noise, z_I .

B. Impact of Direct-path Suppression on Capacity

Theorem 1: The capacity of the full-duplex uplink channel described in eq. 6 is

$$C_{\text{Uplink}}^{\text{FD}} = \mathbb{E} \left[\int_{-\pi}^{\pi} \left(1 + \frac{(\nu - S_N(f))^+}{S_N(f)} \right) df \right], \quad (7)$$

where $S_N(f) = |\sqrt{N_T} H_I(f) + \sqrt{N_R}|^2$, $H_I(f)$ is the DTFT of the self-interference channel impulse response, $h_I[n]$, and ν is a water-filling parameter chosen such that $\int (\nu - S_N(f))^+ df = |h_S|^2 P_T$. The expectation is with respect to the channel responses, h_I and h_S .

Proof: Because the portion of the self-interference due to intended signal, $h_I * x_I$, is known, it can be cancelled by forming

$$y'[n] = y[n] - h_I[n] * x_I[n] = h_S x_S[n] + h_I[n] * z_I[n] + z_R[n], \quad (8)$$

so that the only self-interference that remains is the transmitter noise filtered by the self-interference channel response. Within a single coherence time, Equation 9 is equivalent to the well-known colored Gaussian noise channel with input power constraint $P = |h_S|^2 P_T$ and noise power spectral density $S_N(f) = |\sqrt{N_T} H_I(f) + \sqrt{N_R}|^2$, where $H_I(f)$ is the DTFT of $h_I[n]$. The capacity of the colored Gaussian channel is achieved via water-filling in the spectral domain and is given by [27]

$$C = \int_{-\pi}^{\pi} \left(1 + \frac{(\nu - S_N(f))^+}{S_N(f)} \right) df, \quad (10)$$

where ν is the “water level” chosen such that $\int (\nu - S_N(f))^+ df = P$. By averaging over many independent fades of the channel taps, the capacity given in Theorem 1 is obtained. ■

We want to compare full-duplex capacity to both half-duplex and “ideal” full-duplex capacity. Assuming that in half-duplex mode the uplink must evenly time-share with the downlink, the uplink capacity is easily shown to be

$$C_{\text{Uplink}}^{\text{HD}} = \mathbb{E} \left[\frac{1}{2} \log_2 \left(1 + \frac{2|h_S|^2 P_T}{N_R} \right) \right], \quad (11)$$

where the pre-log factor of $1/2$ accounts for the time sharing, and the factor of 2 in the numerator of the SNR term is due to the fact that the user can transmit with twice the power since it only transmits half the time. Similarly, the “ideal” full-duplex ergodic capacity (the capacity when self-interference channel is zero) is easily shown to be

$$C_{\text{Uplink}}^{\text{IdealFD}} = \mathbb{E} \left[\log_2 \left(1 + \frac{|h_S|^2 P_T}{N_R} \right) \right]. \quad (12)$$

(a) System parameters assumed.

Parameter	Value
Transmit power, P_T	0 dBm
Transmit noise power, N_T	-30 dBm
Receiver noise power, N_R	-90 dBm
Signal channel variance, P_S	-60 dB
Extent of reflections, T_R	32

(b) Capacity vs. direct-path strength.




Fig. 9: Numerical evaluation of ergodic uplink capacity of Eq. 7.

Note that neither half-duplex nor ideal full-duplex capacities depend on P_D or P_R .

Equations (7) (11) and (12) are numerically evaluated in Figure 9(b). The values assumed for the system parameters, shown in Figure 9(a), are typical of WiFi systems. We see in Figure 9(b) that once the direct path strength, P_D , is suppressed to the reflection level, P_R , the capacity no longer grows, which emphasizes the design implication of Result 1: that full-duplex infrastructure nodes should be deployed so as to neutralize nearby reflections.

C. Influence of Increased Frequency Selectivity on Active Cancellation


In the above analysis we have assumed no limitation on the number of taps in the self-interference canceler. We now relax this assumption to study how increased frequency selectivity influences the performance of an active canceler with $N_{\text{Tap}} \leq T_R$ taps. Denote the first N_{Tap} taps of h_I as $h'_I[n] = h_I[0, 1, \dots, N_{\text{Tap}} - 1]$. We assume that the infrastructure node has perfect knowledge of $h'_I[n]$ and subtracts the known transmit signal, filtered by $h'_I[n]$, from its received signal forming

$$\begin{aligned}
 y'[n] &= y[n] - h'_I[n] * x_I[n] \\
 &= h_S x_S[n] + \sum_{m=N_{\text{Tap}}}^{T_R-1} h[m] x_I[n-m] \\
 &\quad + h_I[n] * z_I[n] + z_R[n].
 \end{aligned} \tag{13}$$

We define the average passive suppression, α_P , to be the average ratio of the transmit power to the power of the self-interference incident on the receiver. Similarly, we define the average active cancellation, α_A , to be the average ratio of the power of the self-interference incident on the receiver to the power of the self-interference after cancellation. It is easy to show from Equation (13) that the average passive suppression is $\alpha_P = \mathbb{E} \left[\frac{1}{|h_I|^2} \right]$, and the average active cancellation is

$$\alpha_A = \mathbb{E} \left[\frac{|h_S|^2 (P_T + N_T)}{\sum_{m=N_{\text{Tap}}}^{T_R-1} |h_I[m]|^2 P_T + \sum_{m=0}^{T_R-1} |h_I[m]|^2 N_T} \right], \tag{14}$$

where the expectations are with respect to the channel gain distributions specified above.

Fig. 10: Passive suppression, α_P , and active cancellation, α_A , vs. direct-path strength with reflection strength of $P_R = -60$ dB and reflection extent of $T_R = 32$.

Equation (14) is evaluated numerically in Figure 10 for $N_{\text{Tap}} = 1, 24$ and 32 , where a reflection strength of $P_R = -60$ dB, a reflection extent of $T_R = 32$, and parameter values shown in Figure 9(a) are assumed. Figure 10 emphasizes the design implication of Result 2. We see that when the passive suppression is low, a single-tap canceler can have good performance, but as the passive suppression increases and the channel becomes frequency selective, the single tap canceler degrades. Only if the self-interference canceler employs the full-number of taps in the channel, $T_R = 32$, can it continue to cancel significant amounts of self-interference in the presence of strong passive suppression.

V. FULL-DUPLEX PROTOTYPE WITH COMBINED PASSIVE AND ACTIVE SUPPRESSION

We now investigate the performance of combined passive and active suppression in a prototype wideband OFDM full-duplex physical layer.

A. Prototype Full-Duplex Infrastructure Node Description

A full-duplex infrastructure node was prototyped using the WARPLab framework [16], in which baseband signals are processed offline in MATLAB, then up-converted and transmitted over-the-air at 2.4 GHz in real time. The prototype implements

a 20 MHz, 64-subcarrier OFDM PHY with a packet structure based on IEEE 802.11. The prototype infrastructure node was mounted outdoors on a 7 ft. pole just outside the Rice University Tudor Fieldhouse. The prototype was evaluated using antenna Configuration II (Figure 1(b)) with absorptive shielding, and with and without cross-polarization. Thus the mechanism combinations evaluated were $\text{mech} = \{\text{DA}, \text{DCA}\}$.

For active self-interference cancellation, we adopted the per-subcarrier analog and digital cancellation mechanisms demonstrated in [7], [11], [28]. Analog cancellation is implemented by estimating the self-interference channel via OFDM pilots, crafting a wideband cancellation waveform, and injecting the cancellation waveform into the receive signal path using an RF combiner, so that the self-interference is cancelled prior to the RF front end. Digital cancellation is implemented by estimating and subtracting, at baseband, the residual self-interference left after analog cancellation. Result 2 tells us that the coherence bandwidth of the self-interference channel will likely be high due to the passive suppression. Because the cancellation scheme is performed on a per-subcarrier basis, we expect that cancellation performance will be robust to such frequency-selectivity.

B. Prototype Evaluation Results

1) *Total suppression exceeding 90 dB*: One-hundred packets were transmitted by the infrastructure node at a power of 7 dBm, and the amount of self-interference suppression was measured by observing the received signal strength indication (RSSI) from the WARP radio. Figure 11(a) shows the empirical CDFs for the self-interference suppression. We see in Figure 11(a) that with directional isolation and absorptive shielding ($\text{mech} = \text{DA}$), 60 dB of passive suppression is achieved on the average. Adding cross-polarization ($\text{mech} = \text{DCA}$) increases the average passive suppression to 71 dB. These values are in line with the characterization of Section III: since the outdoor deployment has weak reflections, we expect the achieved suppression to be near that seen in the anechoic chamber, in which we saw 72 dB of passive suppression for $\text{mech} = \text{DCA}$. Furthermore, we see in Figure 11(a) that active cancellation adds around 25 dB of suppression, on the average, in both co-polarized and cross-polarized cases. In the cross-polarized case, the total self-interference suppression ranges from 87 dB to 100 dB, with 95 dB average suppression.⁷ Thus the analog and digital cancellation mechanisms of [11] are still able to provide 20-25 dB of self-interference cancellation, even after the self-interference is attenuated 70 dB by the passive mechanisms.

2) *Full-duplex uplink-rate improvements at 100+ m*: In a full-duplex infrastructure node with simultaneous uplink and downlink (such as a full-duplex relay or a full-duplex sectorized access point as shown in Figure 12(b)), the uplink to the infrastructure node is the link that will suffer from self-interference. Thus a good system-level metric for comparing

full-duplex performance against half-duplex is comparing the average achievable full-duplex uplink rate $R_{\text{Uplink}}^{\text{FD}}$ to the comparable average achievable half-duplex uplink $R_{\text{Uplink}}^{\text{HD}}$. To experimentally measure the average achievable rate, we employ the common practice of measuring the effective SNR (or SINR) for each packet, and computing the achievable rate for each packet from Shannon's formula [27] $R = \log(1 + \text{SNR})$. Averaging over all packets received, we obtain the average achievable rate. The effective SNR (SINR) of packet p is measured by computing the average error vector magnitude, $\overline{\text{EVM}}(p)$, which is the average distance between the received symbols (after equalization) in packet p and the original transmitted symbols. Average EVM is used to find the effective SNR using the common conversion $\text{SNR}(p) = 1/\overline{\text{EVM}}(p)^2$ [29].

We measure the signal-to-self-interference-plus-noise ratio (SSINR) for the full-duplex uplink by having the infrastructure node receive packets from the uplink mobile node while simultaneously transmitting. From the EVM statistics we measure the $\text{SSINR}_{\text{Uplink}}^{\text{FD}}$ for each packet. Similarly, to measure the SNR of the comparable half-duplex uplink, we compute $\text{SNR}_{\text{Uplink}}^{\text{HD}}$ from the EVM statistics of the packets received at the uplink, but without the infrastructure node transmitting simultaneously. The average achievable rates of the full-duplex/half-duplex links are thus computed from the measured SSINR/SNR values using the formulae,

$$R_{\text{Uplink}}^{\text{FD}} = \frac{1}{N} \sum_{p=1}^N \log_2[1 + \text{SSINR}_{\text{Uplink}}^{\text{FD}}(p)] \quad (15)$$

$$R_{\text{Uplink}}^{\text{HD}} = \frac{1}{N} \sum_{p=1}^N \frac{1}{2} \log_2[1 + \text{SNR}_{\text{Uplink}}^{\text{HD}}(p)], \quad (16)$$

where p is the packet index, and N is the total number of packets. Note that the pre-log factor of $\frac{1}{2}$ in Equation (16) is due to the assumption of even timesharing with downlink in half-duplex mode. To make the comparison fair in average power, the uplink mobile is given twice the transmission power (10 dBm rather than 7 dBm) in half-duplex mode, since we assume it defers to downlink half of the time. We note that the above method of characterizing a prototype's performance by computing achievable rates from measured SNR/SSINR is well-established in prior work such as [7], [30], [31]. For a very thorough explanation see Section III of [31] and references therein.

The achievable rates of the full-duplex uplink and comparable half-duplex uplink were measured as the range from of the uplink mobile to the infrastructure node was varied from 50 to 200 meters in order to encounter a variety of path losses. Figure 11(b) plots the percent improvement of the full-duplex uplink rate, $R_{\text{Uplink}}^{\text{FD}}$, versus the half-duplex uplink rate, $R_{\text{Uplink}}^{\text{HD}}$, as a function of the encountered path loss.⁸ We see that at 86 dB path loss, the improvement over

⁷In the empirical CDF for "w/ Cross-pol: passive + active", there is a slight flat spot at 95 dB. We believe this is an artifact of the finite-sampled (i.e. histogram) approximation of the distribution, and not evidence of a multimodal distribution of self-interference suppression.

⁸As 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.

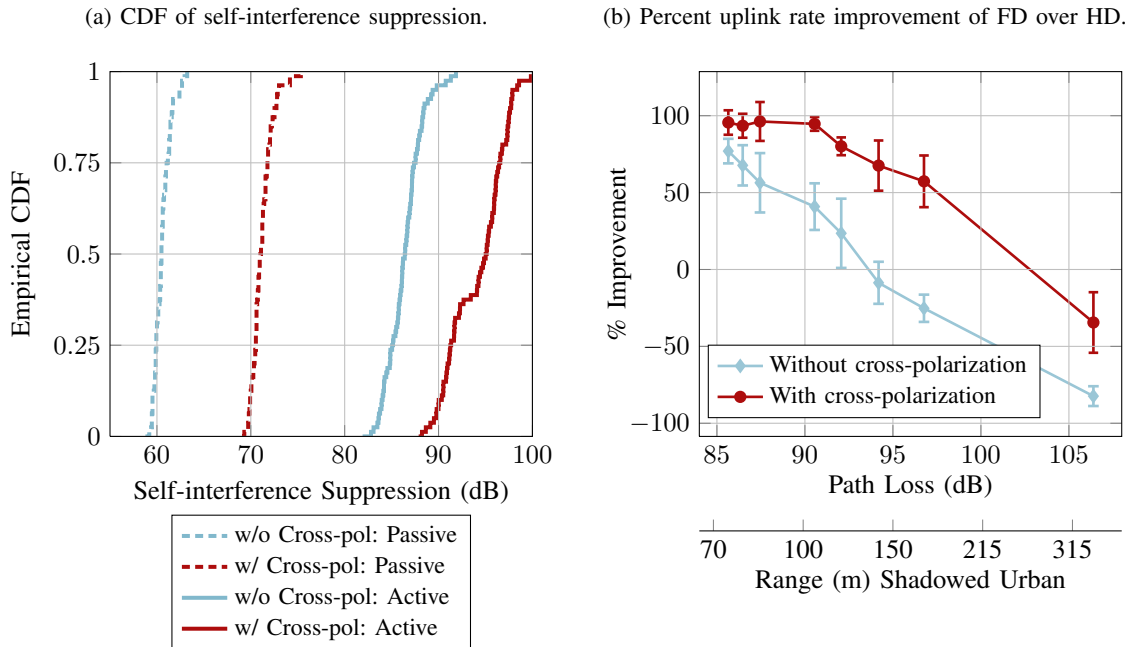



Fig. 11: Full-duplex prototype evaluation results.


half-duplex is 86% when directional isolation and absorptive shielding are applied together with active cancellation. When cross-polarization is also employed, the gain over half-duplex at 86 dB path loss is 96%, meaning that the effect of self-interference is nearly eliminated, since the rate is nearly doubled. In Figure 11(b), a second x-axis is also shown, which maps path loss to effective range in a typical urban channel [32]. With the best design, including cross-polarization, we expect to see full-duplex outperform half-duplex for up to 200 m range. Even without cross-polarization, the prototype full-duplex design outperforms half-duplex at effective ranges exceeding 100 m. Of particular note is the significant impact of cross-polarization. Figure 11(a) 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 11(b): the same gains over half-duplex can be attained at 10 dB more path loss with cross-polarization than without.

VI. DISCUSSION: IMPACT ON THE USER'S ACCESS TO THE INFRASTRUCTURE NODE

There are two main scenarios in which an infrastructure node can leverage full-duplex capability to enhance its service to users: two-node full-duplex and three-node full-duplex, illustrated in Figure 12. Note that two-node full-duplex requires the user device to be full-duplex capable, while three-node does not. Due to the extra hardware burden that full-duplex would place on mobile devices, we expect the first applications of full-duplex will be three-node scenarios, where the users are half-duplex, hence the emphasis on infrastructure nodes rather than user nodes in this paper. The interference from the uplink user to the downlink user in three-node full-duplex can be managed by either scheduling non-interfering users as in [33], or canceling the inter-user interference using advanced physical layers as in [34].



(a) Two-node full-duplex: infrastructure node simultaneously transmits and receives to a single full-duplex user.



(b) Three-node full-duplex: infrastructure node receives from one user while transmitting to another user.

Fig. 12: Scenarios for full-duplex transmissions.

A. Impact of Directional Isolation

1) *Supporting Full-duplex and Half-duplex Users:* We make the distinction between two-node and three-node full-duplex, because directional isolation is often better suited for three-node full-duplex applications. For example, in Configuration I (Figure 1(a)), the two antennas cover disjoint areas, and cannot transmit and receive in the same direction, hence Configuration I is poorly suited to support two-node full-duplex transmissions. However, Configuration I is indeed well-suited to support three-node full-duplex transmissions, when the uplink user is covered by one antenna the downlink user is covered by the other antenna as in Figure 12(b). Configurations II and III (Figures 1(b) and 1(c)), both of which use 90° beamwidth antennas with 60° separation, also support three-node full-duplex transmissions but can only support a two-node full-duplex transmission if the full-duplex-equipped user is in the 30° intersection of the two coverage patterns. Configurations IV and V (Figures 1(d) and 1(e)) which use omnidirectional antennas support both two-node and three-node full-duplex scenarios regardless of location of the user.

2) *Managing directional access:* In a directional antenna system, control must be employed to ensure that any user is serviced by the antenna(s) with the strongest gain in the user's direction. Many algorithms have been proposed for adaptive directional antenna selection [35]–[37], and commercial access points with directional antennas that implement dynamic user selection are in the market [38], [39]. It is outside the scope of this paper, but modest extension of the control algorithms used in these existing systems can be used to schedule concurrent three-node full-duplex transmissions with directional antennas.

B. Impact of Cross-polarization

Cross-polarization at the infrastructure node can also affect users. Assume the infrastructure node transmits with horizontal polarization and receives with vertical polarization. The impact of cross-polarization on the user nodes depends on the environment as described below.

1) *Non-line-of-sight environments:* In a non-line-of-sight environment, studies have shown that scattering leads to signals becoming *de-polarized* [40]: the received wave is the superposition of many randomly polarized components. Thus the user is agnostic to the polarization state at the infrastructure node. A single-antenna user with any arbitrarily polarized antenna will be able to receive from the infrastructure node's horizontally polarized transmit antenna and transmit to the infrastructure nodes' vertically polarized receive antenna equally well.⁹

2) *Line-of-sight environments:* In a line-of-sight environment, however, cross-polarization at the infrastructure node influences the polarization requirements of the user. If the user has only a single vertically polarized antenna, for instance, then it will have poor SNR when receiving the horizontally-polarized downlink. The ideal solution would be for users to

⁹The optimal design at the user device would have three orthogonally polarized antennas, combining power from all three so that all incident energy is captured. But this is not specific to full-duplex, but the case also for half-duplex in any non-line-of-sight environment.

match the infrastructure polarization: a half-duplex user with a dual-polarized antenna could transmit on the vertical port and receive on the horizontal port. Dual-polarized antennas are available in form factors suitable for mobile phones [41]. The downside to polarization matching is that it requires the user device to either be in a stable orientation or to track its orientation. Smartphone sensors make tracking orientation feasible, and as orientation information was leveraged for antenna-pattern switching in [42], so it could be leveraged for polarization matching. A simpler solution is for each user device to use a circularly polarized antenna, which couples with both vertical and horizontal polarization, but with a 3 dB loss due to polarization mismatch [43].

VII. CONCLUSION AND FUTURE WORK

We have seen that a careful utilization of the passive self-interference suppression mechanisms of directional isolation, absorptive shielding, and cross-polarization can significantly improve the performance of full-duplex links, but some fundamental limitations are encountered. The amount of passive suppression is limited by reflected paths which are difficult to suppress passively. Moreover, because passive suppression eliminates the once-dominant direct self-interference path, the residual self-interference channel can be much more frequency selective. We have shown that a design that takes these capabilities and limitations into account can achieve a full-duplex link that outperforms a comparable half-duplex link even a ranges exceed 100 m.

A future step in this research direction is to develop a general theory for passive suppression. For example, given a size constraint, a coverage requirement, and a description of the scattering/reflection environment, one would like to solve for the antenna design and placement of absorbers which optimizes system capacity. Such a theory will require not only new analysis but new models, and we hope the results in this paper are a first step towards developing the models required for a general theory of passive suppression.

APPENDIX A

MEASUREMENT DETAILS AND UNCERTAINTY ANALYSIS

As mentioned in Section II the frequency response of the self-interference channel was measured using a network analyzer. We now discuss details how the measurement setup, and discuss the associated uncertainty.

A. Measurement Details

A two-port network analyzer, such as the the Agilent E8364B [23] used in our study, measures the S-parameters of an arbitrary two-port RF network by sweeping a continuous wave at the input(s) and measuring the amplitude and phase of the output(s). In our case, the two ports of interest are the ports of the infrastructure node antennas, and S_{21} , the transmission coefficient, gives the amplitude and phase of the channel. A full two-port calibration was performed with the Agilent N4691B [44] calibration kit connected *after* the cables and connectors, thus correcting for any cable loss or coupling, and placing the phase and amplitude reference directly at

the intended antenna ports. An intermediate frequency (IF) bandwidth of 20 kHz was used, as this provided a good balance between noise floor and frequency sweep time. To lessen the effect of noise, each frequency point was averaged over 20 repeated measurements.




Fig. 13: Uncertainty curve for the network analyzer measurement setup presented in Section II

B. Uncertainty Discussion

Assessing uncertainty by merely computing empirical standard deviation over the 20 points measured would only quantify noise and not quantify the uncertainty due to imperfect calibration. Thankfully, Agilent provides a software tool called *Uncertainty Calculator* [45] that computes the uncertainty associated with a measurement setup from the following key parameters: the IF bandwidth (20 kHz in our case) which affects the thermal noise, the averaging factor (20 in our case), and the calibration instrument used. *Uncertainty Calculator* uses proprietary statistical characterizations of Agilent hardware to perform a root sum square uncertainty analysis and compute uncertainty as a function of the power of the quantity measured. Figure 13 plots the uncertainty curve generated for our setup, from which the uncertainty entries in Table I are taken.

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.
- [2] A. K. Khandani, "Methods for spatial multiplexing of wireless two-way channels," US Patent US 7817 641, October 19, 2010.
- [3] 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.
- [4] 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.
- [5] J. I. Choi, M. Jain, K. Srinivasan, P. Levis, and S. Katti, "Achieving single channel, full duplex wireless communication," in *MobiCom 2010*.
- [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>
- [7] M. Duarte, C. Dick, and A. Sabharwal, "Experiment-driven characterization of full-duplex wireless systems," May 2011, to appear in *Wireless Communications, IEEE Transactions on*. [Online]. Available: http://warp.rice.edu/trac/wiki/TransWireless2011_FullDuplex
- [8] A. Sahai, G. Patel, and A. Sabharwal, "Pushing the limits of full-duplex: Design and real-time implementation," Rice University, Tech. Rep. TREE1104, 2011.
- [9] M. A. Khojastepour, K. Sundaresan, S. Rangarajan, X. Zhang, and S. Barghi, "The case for antenna cancellation for scalable full-duplex wireless communications," in *Proceedings of the 10th ACM Workshop on Hot Topics in Networks*, ser. HotNets-X. New York, NY, USA: ACM, 2011, pp. 17:1–17:6. [Online]. Available: <http://doi.acm.org/10.1145/2070562.2070579>
- [10] E. Aryafar, M. A. Khojastepour, K. Sundaresan, S. Rangarajan, and M. Chiang, "Midu: enabling mimo full duplex," in *Proceedings of the 18th annual international conference on Mobile computing and networking*, ser. Mobicom '12. New York, NY, USA: ACM, 2012, pp. 257–268. [Online]. Available: <http://doi.acm.org/10.1145/2348543.2348576>
- [11] M. Duarte, A. Sabharwal, V. Aggarwal, R. Jana, K. Ramakrishnan, C. Rice, and N. Shankaranayanan, "Design and characterization of a full-duplex multi-antenna system for wifi networks," *submitted to IEEE Transactions on Vehicular Technology*, October 2012. [Online]. Available: <http://warp.rice.edu/trac/wiki/DesignFDWiFi>
- [12] M. Duarte, "Full-duplex wireless: Design, implementation and characterization," Ph.D. dissertation, Rice University, April 2012. [Online]. Available: <http://warp.rice.edu/trac/wiki/DuartePhDThesis>
- [13] A. Sahai, G. Patel, C. Dick, and A. Sabharwal, "On the impact of phase noise on active cancelation in wireless full-duplex," *Vehicular Technology, IEEE Transactions on*, vol. 62, no. 9, pp. 4494–4510, 2013.
- [14] B. Day, A. Margetts, D. Bliss, and P. Schniter, "Full-duplex MIMO relaying: Achievable rates under limited dynamic range," *Selected Areas in Communications, IEEE Journal on*, vol. 30, no. 8, pp. 1541–1553, september 2012.
- [15] —, "Full-duplex bidirectional MIMO: Achievable rates under limited dynamic range," *Signal Processing, IEEE Transactions on*, vol. 60, no. 7, pp. 3702–3713, july 2012.
- [16] WARPLab framework. [Online]. Available: <http://warp.rice.edu/trac/wiki/WARPLab>
- [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.
- [18] C. Anderson, S. Krishnamoorthy, C. Ranson, T. Lemon, W. Newhall, T. Kummert, and J. Reed, "Antenna isolation, wideband multipath propagation measurements, and interference mitigation for on-frequency repeaters," in *SoutheastCon, 2004. Proceedings. IEEE*, mar 2004, pp. 110–114.
- [19] K. Haneda, E. Kahra, S. Wyne, C. Icheln, and P. Vainikainen, "Measurement of loop-back interference channels for outdoor-to-indoor full-duplex radio relays," in *Antennas and Propagation (EuCAP), 2010 Proceedings of the Fourth European Conference on*, April 2010, pp. 1–5.
- [20] L-Com. 2.4 GHz 90 degree dual-polarized antenna. [Online]. Available: <http://www.l-com.com/item.aspx?id=22078>
- [21] —. HGV-2406 2.4 GHz 6 dBi omnidirectional antenna. [Online]. Available: <http://www.l-com.com/wireless-antenna-24-ghz-6-dbi-omnidirectional-antenna-n-female-connector>
- [22] Emmerson and Cuming Microwave. Eccosorb free-space absorber. [Online]. Available: <http://www.eccosorb.com/products-eccosorb-an.htm>
- [23] A. Technologies. (June, 2013) E8364B network analyzer data sheet. [Online]. Available: <http://www.home.agilent.com/en/pd-72337-pn-E8364B>
- [24] A. Goldsmith, *Wireless Communications*. Cambridge University Press, 2005.
- [25] W. Lee, *Mobile Cellular Telecommunication*. McGraw Hill, 1989.
- [26] D. N. C. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.
- [27] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. John Wiley and Sons Inc., 2006.
- [28] A. Sahai, G. Patel, and A. Sabharwal, "Pushing the limits of full-duplex: Design and real-time implementation," *CoRR*, vol. abs/1107.0607, 2011.

- [29] 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.
- [30] E. Aryafar, N. Anand, T. Salonidis, and E. W. Knightly, "Design and experimental evaluation of multi-user beamforming in wireless lans," in *Proceedings of the sixteenth annual international conference on Mobile computing and networking*, ser. MobiCom '10. New York, NY, USA: ACM, 2010, pp. 197–208. [Online]. Available: <http://doi.acm.org/10.1145/1859995.1860019>
- [31] M. Duarte, A. Sabharwal, C. Dick, and R. Rao, "Beamforming in MISO systems: Empirical results and EVM-based analysis," *Wireless Communications, IEEE Transactions on*, vol. 9, no. 10, pp. 3214–3225, october 2010.
- [32] T. S. Rappaport, *Wireless Communications: principles and practice*. Prentice Hall, 1996.
- [33] N. Singh, D. Gunawardena, A. Proutiere, B. Radunovic, H. Balan, and P. Key, "Efficient and fair mac for wireless networks with self-interference cancellation," in *Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt), 2011 International Symposium on*, 2011, pp. 94–101.
- [34] J. Bai and A. Sabharwal, "Decode-and-cancel for interference for cancellation in full-duplex networks," in *Asilomar Conference on Signals, Systems and Computers*, 2012.
- [35] 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.
- [36] Y.-B. Ko, V. Shankarkumar, and N. Vaidya, "Medium access control protocols using directional antennas in ad hoc networks," in *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 1, 2000, pp. 13–21 vol.1.
- [37] R. R. Choudhury, X. Yang, R. Ramanathan, and N. H. Vaidya, "Using directional antennas for medium access control in ad hoc networks," in *Proceedings of the 8th annual international conference on Mobile computing and networking*, ser. MobiCom '02. New York, NY, USA: ACM, 2002, pp. 59–70. [Online]. Available: <http://doi.acm.org/10.1145/570645.570653>
- [38] Xirrus Inc., "Sectorized Wi-Fi benefits," Xirrus Inc., Tech. Rep., November 2007.
- [39] —. XN and XR series wireless arrays. [Online]. Available: <http://www.xirrus.com/Products/Wireless-Arrays.aspx>
- [40] 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.
- [41] 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.
- [42] A. Amiri Sani, L. Zhong, and A. Sabharwal, "Directional antenna diversity for mobile devices: characterizations and solutions," in *MobiCom 2010*. New York, NY, USA: ACM, 2010, pp. 221–232.
- [43] D. M. Pozar, *Microwave and RF Design of Wireless System*. John Wiley and Sons Inc., 2000.
- [44] Agilent N4691B calibration kit. [Online]. Available: <http://www.home.agilent.com/en/pd-491095-pn-N4691B>
- [45] Agilent. (June, 2013) Network analyzer uncertainty calculator. [Online]. Available: <http://www.home.agilent.com/agilent/software.jsp?cc=US&lc=eng&ckey=1000000418:epsg:sud&nid=-536902639.350794.02&id=1000000418:epsg:sud&cmpid=93648>



Achaleshwar Sahai received a Dual Degree (B.Tech and M.Tech) in Electrical Engineering from the Indian Institute of Technology Madras, Chennai in 2008. Currently he is pursuing a Ph.D. in Electrical and Computer Engineering at Rice University, Houston, TX. His research interests include information theory and design of wireless communication systems.



Ashutosh Sabharwal (S'91-M'99-SM'04) received the B. Tech. degree from the Indian Institute of Technology, New Delhi, India, in 1993 and the M.S. and Ph.D. degrees from The Ohio State University, Columbus, in 1995 and 1999, respectively. He is currently a Professor with the Department of Electrical and Computer Engineering, Rice University, Houston, TX. His research interests include information theory, communication algorithms and experiment-driven design of wireless networks. He received the 1998 Presidential Dissertation Fellowship Award.



Evan Everett (S'06) received the B.Eng degree in Wireless Engineering and the B.S degree in Physics from Auburn University in 2010. He received the M.S. degree in Electrical Engineering from Rice University in 2012, where he is currently pursuing a Ph.D. His research focus is theory and design for full-duplex wireless communication. He received the 2010 Comer Award for Excellence in Physical Science from Auburn University, and he is a National Science Foundation Graduate Research Fellow.