

SoftNull: Many-Antenna Full-Duplex Wireless via Digital Beamforming

Evan Everett, *Member, IEEE*, Clayton Shepard, Lin Zhong, *Senior Member, IEEE*,
and Ashutosh Sabharwal, *Fellow, IEEE*

Abstract—In this paper, we present and study a digital-controlled method, called SoftNull, to enable full-duplex in many-antenna systems. Unlike most designs that rely on analog cancelers to suppress self-interference, SoftNull relies on digital transmit beamforming to reduce self-interference. SoftNull does not attempt to perfectly null self-interference, but instead seeks to reduce self-interference sufficiently to prevent swamping the receiver’s dynamic range. Residual self-interference is then cancelled digitally by the receiver. We evaluate the performance of SoftNull using measurements from a 72-element antenna array in both indoor and outdoor environments. We find that SoftNull can significantly outperform half-duplex for small cells operating in the many-antenna regime, where the number of antennas is many more than the number of users served simultaneously.

Index Terms—Full-duplex, MIMO, massive MIMO, precoders, interference.

I. INTRODUCTION

FULL-DUPLEX wireless communication, in which transmission and reception occur at the same time and in the same frequency band, has the potential to as much as double the spectral efficiency of traditional half-duplex systems. The main challenge to full-duplex is self-interference: a node’s transmit signal generates high-powered interference to its own receiver. Research over the last ten years [1]–[10] has shown that full-duplex operation may be feasible for small cells, and the key enabler has been *analog cancellation* of the self-interference in addition to digital cancellation. Analog cancellation has been considered a necessary component of a full-duplex system, to avoid self-interference from overwhelming the dynamic range of the receiver electronics, and swamping the much weaker intended signal (see Appendix C for a detailed explanation of dynamic range and its impact full-duplex operation).

Many analog cancellation designs have been proposed for single-antenna [5], [7] and dual-antenna [2]–[4] full-duplex systems. However, current wireless base stations use many antennas (up to 8 in LTE Release 12 [11]), and next-generation

Manuscript received August 13, 2015; revised February 28, 2016 and August 26, 2016; accepted August 28, 2016. Date of publication September 23, 2016; date of current version December 8, 2016. This work was supported in part by NSF under Grant CNS 0923479, Grant CNS 1012921, and Grant CNS 1161596, and in part by the NSF Graduate Research Fellowship under Grant 0940902. The associate editor coordinating the review of this paper and approving it for publication was K. Haneda.

The authors are with the Department of Electrical and Computer Engineering, Rice University, Houston, TX 77005 USA (e-mail: evan.everett@rice.edu; cws@rice.edu; lzhong@rice.edu; ashu@rice.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TWC.2016.2612625

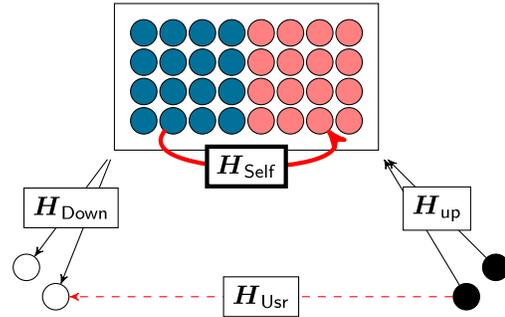


Fig. 1. Multi-user full-duplex system.

wireless systems will likely employ many more antennas at base stations. For example, discussions to include 64-antenna base stations have already been initiated in 3GPP standardization [12], and “massive” antenna arrays with hundreds to thousands of antennas have also been proposed [13]–[15]. Large arrays offer many benefits. For example, inter-beam interference can be managed with simple linear precoders/equalizers and inter-cell interference can be mitigated without requiring coordination between base stations [13]–[15]. Although adding more antennas incurs the cost of more RF chains, each RF chain can be built with lower-cost components. Therefore, each RF chain carries only a fraction of the power [14] as would be carried in a single-antenna base station.

As the number of base-station antennas increases, an important question is how to enable full-duplex with a large number of antennas. Full-duplex multi-user MIMO (MU-MIMO) communication would enable the base station to transmit to multiple downlink users and receive from multiple uplink users, all at the same time and in the same frequency band, as shown in Figure 1. Full-duplex with many antennas presents both challenges and opportunities. The complexity of analog self-interference cancellation circuitry grows in proportion to the number of antennas (which could potentially deter its adoption due to increased cost and complexity). At the same time, many-antenna full-duplex also presents an opportunity: having many more antennas than users served means that more spatial resources become available for transmit beamforming to reduce self-interference.

In this work, we investigate the possibility of many-antenna full-duplex operation with current radio hardware that can either send or receive on the same band but not both, i.e. *TDD¹ radios without analog cancellation*. We propose

¹We consider only TDD radios, because FDD radios, by design, do not transmit and receive in the same band and hence cannot be transformed into in-band full-duplex.

an all-digital² approach called *SoftNull*, to enable many-antenna full-duplex with only digital-domain modifications. In the *SoftNull* design, the array is partitioned into a set of transmit antennas and a set of receive antennas, and self-interference from the transmit antennas to the receive antennas is reduced by transmit beamforming. We envision that one method of using *SoftNull* will be a layer *below* physical layer, tasked to only reduce self-interference, and is agnostic to the upper layer processing. Thus *SoftNull* can operate on the output of algorithms for downlink MU-MIMO (such as zero-forcing beamforming) without modifying their operation.

Transmit beamforming to null self-interference has been considered previously [1], [8], [16]–[23], but to our knowledge, no prior work has included an experiment-based evaluation of many-antenna beamforming for full-duplex. The key departure in *SoftNull* design is that *our aim is not necessarily to null self-interference perfectly at each receive antenna*. Every null requires using one effective transmit antenna dimension. For a many-antenna system, self-interference is full rank and hence a nulling based self-interference scheme may end up using all available transmit degrees of freedom, leaving negligible degrees of freedom for actual downlink data transmission. Instead, our aim is to reduce self-interference to avoid saturating the analog-to-digital conversion in the receive radio chain. The *SoftNull* precoder minimizes the total self-interference power, given a constraint on how many effective antennas must be preserved, where “effective antennas” are the number of dimensions available to the physical layer for downlink communication. We find that the precoder to minimize total self-interference has a simple and intuitive form: the precoder is a projection onto the singular vectors of the self-interference channel corresponding to the D_{Tx} smallest singular values.

A. Contribution

Our contribution is an experiment-driven evaluation of the digital-controlled *SoftNull*-based full-duplex system using 3D self-interference channel measurements from a variety of propagation environments. The goal of the evaluation is to understand the conditions under which the *SoftNull* system outperforms a traditional half-duplex system, and quantify how close we can approach the performance of an ideal full duplex system. We collect channel measurements using a 72-element two-dimensional planar antenna array, with mobile nodes placed in many different locations, measuring self-interference channels and uplink/downlink channels both outdoors, indoors and in an anechoic chamber. The platform operates in the 2.4 GHz ISM band, with 20 MHz bandwidth. We use these real over-the-air channel measurements to simulate *SoftNull* and evaluate its performance extensively. The essence of the experimental results can be captured by the following two measurement-based conclusions.

²By “all-digital” we mean that the only modifications to standard half-duplex TDD systems are in the digital domain: no additional analog components are required beyond what is already present in the half-duplex radios.

B. Self-Interference Reduction

SoftNull enables a large reduction in self-interference while sacrificing relatively few effective antennas. However, the amount of reduction depends on the environment: more scattering results in less suppression. In an outdoor low-scattering environment *SoftNull* provides sufficient self-interference reduction while sacrificing only a few effective antennas. For example in the case of a 72-element array partitioned as 36 transmit antennas and 36 receive antennas, 50 dB of pre-analog self-interference reduction is achieved while sacrificing only 12 of the 36 available transit dimensions. Self-interference reduction via beamforming becomes more challenging in indoor environments due to backscattering. Since backscattering makes the self-interference channel less correlated, more effective antennas must be used to achieve the same reduction. With the same base station indoors, 20 of the 36 effective antennas need to be used to achieve 50 dB reduction.

C. Data Rate Gains Over Half Duplex

SoftNull can provide significant rate gains over half-duplex for small cells in the case when the number of transmit antennas is much larger than the number of users. The larger the path loss, the more challenging full-duplex operation becomes, because more self-interference reduction is required to suppress the self-interference to a power level commensurate to the power of the received uplink signal. For *SoftNull*, more path loss means more effective antennas must be used to suppress the self-interference to a level commensurate to the uplink signal power. Similarly, as the number of simultaneous users served increases, the cost of using effective antennas for self-interference reduction becomes more pronounced: not only is downlink power gain sacrificed, but downlink multiplexing gain is also sacrificed. For example, in the 72-antenna scenario mentioned above, with 12 users at 100 dB path loss, the data rate achieved by *SoftNull* is 12% *less* than half duplex, but for 4 users at 85 dB path loss the data rate improvement of *SoftNull* over half duplex more than 40%. For example, in the 72-antenna scenario mentioned above, with 4 users at 85 dB path loss the data rate improvement of *SoftNull* over half duplex more than 40%, but with 12 users at 100 dB path loss, the data rate achieved by *SoftNull* is 40% *less* than half duplex. We note however, that the trend in wireless deployments is moving towards smaller cells [24] (i.e. lower path loss) and towards operating in the regime where the number of antennas is much more than number of users served [13]–[15], therefore we foresee a large application space for *SoftNull*.

The rest of the paper is organized as follows. Section II describes the multi-user MIMO scenario under consideration and defines key variables and terms. Section III describes the design of the *SoftNull* system, in particular the self-interference suppression precoder, and gives a brief simulation example. Section IV describes the measurement setup. Section V presents the results of the measurement-driven performance evaluation. Concluding remarks are given in Section VI.

II. SYSTEM DEFINITION

We consider the multi-user system pictured in Figure 1. A base station is communicating with K_{Up} uplink users and K_{Down} downlink users. The base station is equipped with M antennas. We assume the base station uses traditional radios, that is each of the M antennas can both transmit and receive, but a given antenna cannot both transmit and receive at the same time. Therefore in full-duplex operation, M_{Tx} of the antennas transmit while M_{Rx} antennas receive, with the requirement that $M_{\text{Tx}} + M_{\text{Rx}} \leq M$. Note that choice of which antennas transmit and receive can be adaptively chosen by the scheduler, but study of such adaptation is left to future work. In half-duplex mode all antennas are used for either transmission or reception, that is $M_{\text{Tx}} = M_{\text{Rx}} = M$. The vector of symbols transmitted by the base station is $\mathbf{x}_{\text{Down}} \in \mathbb{C}^{M_{\text{Tx}}}$, and the vector of symbols transmitted by the users is $\mathbf{x}_{\text{Up}} \in \mathbb{C}^{K_{\text{Up}}}$.

The signal received at the base station is

$$\mathbf{y}_{\text{Up}} = \mathbf{H}_{\text{up}}\mathbf{x}_{\text{Up}} + \mathbf{H}_{\text{Self}}\mathbf{x}_{\text{Down}} + \mathbf{z}_{\text{Up}}, \quad (1)$$

where $\mathbf{H}_{\text{up}} \in \mathbb{C}^{M_{\text{Rx}} \times K_{\text{Up}}}$ is the uplink channel matrix, $\mathbf{H}_{\text{Self}} \in \mathbb{C}^{M_{\text{Rx}} \times M_{\text{Tx}}}$ is the self-interference channel matrix, and $\mathbf{z}_{\text{Up}} \in \mathbb{C}^{M_{\text{Rx}}}$ is the noise at the base station's receiver. The signal received by the K_{Down} downlink users is

$$\mathbf{y}_{\text{Down}} = \mathbf{H}_{\text{Down}}\mathbf{x}_{\text{Down}} + \mathbf{H}_{\text{Usr}}\mathbf{x}_{\text{Up}} + \mathbf{z}_{\text{Down}}, \quad (2)$$

where $\mathbf{H}_{\text{Down}} \in \mathbb{C}^{K_{\text{Down}} \times M_{\text{Tx}}}$, is the downlink channel matrix, $\mathbf{H}_{\text{Usr}} \in \mathbb{C}^{K_{\text{Down}} \times K_{\text{Up}}}$ is the matrix of channel coefficients from the uplink to the downlink users, and $\mathbf{z}_{\text{Down}} \in \mathbb{C}^{K_{\text{Down}}}$ is the noise at the receiver of each user.

Ongoing research is developing schedulers to select users such that the interference caused by uplink users on downlink reception is minimized [25]–[29] (and references within). Thus, we will make a simplifying assumption that $\mathbf{H}_{\text{Usr}} = 0$, allowing us to focus on self-interference; we acknowledge that future work should also characterize the role of inter-node interference on overall network rate. In half-duplex operation the above equations are simplified: the self-interference term is eliminated in (1), and \mathbf{H}_{up} is a $M \times K_{\text{Up}}$ matrix and K_{Down} is a $K_{\text{Down}} \times M$ matrix. The signaling challenge unique to full-duplex operations is how to design \mathbf{x}_{Down} such that the self-interference is small, while still providing a high signal-to-interference-plus-noise ratio (SINR) to the downlink users. One of the major challenges in many-antenna MIMO systems is the acquisition of downlink channel state information (CSI) at the base station, such that \mathbf{H}_{Down} can be used to compute the downlink precoder. The need to acquire CSI is universal to all many-antenna MIMO beamforming systems, and not specific to full-duplex implementation. A large body of work is devoted to studying methods for obtaining channel state information with low overhead, such as [13], [30]–[33], and references therein.

III. SOFTNULL DESIGN

The physical layer design for SoftNull is depicted in Figure 2. We propose a two-stage approach. The first stage is standard MU-MIMO for which conventional precoding and equalization algorithms can be used. The second

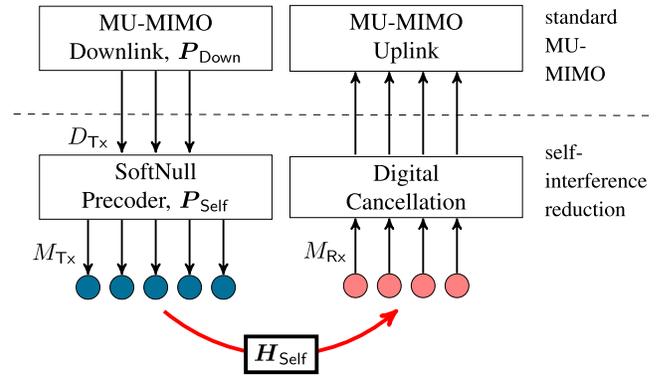


Fig. 2. SoftNull design. First stage is standard MU-MIMO. Second stage is self-interference reduction, with two components: SoftNull transmit precoder to reduce the self-interference, and receiver-side digital canceler to reduce residual self-interference.

stage is the self-interference reduction stage, which reduces self-interference via transmit beamforming and digital self-interference cancellation. The advantage of this two-stage approach is that SoftNull can be incorporated as a modular addition to existing MU-MIMO systems. The disadvantage is that the performance may be sub-optimal due to the two-stage constraint. Joint precoder design for MU-MIMO downlink and self-interference reduction is a topic for future work but is outside the scope of this paper.

The self-interference reduction stage of SoftNull has two components: a transmitter-side precoder to reduce self-interference and a receiver-side digital canceler to reduce residual self-interference. Digital cancellation is well understood, and we believe existing techniques are sufficient for practical use; see e.g. [6], [34]. Thus, in this section, we focus on the design of the SoftNull precoder. We assume that the decision on the partitioning of the transmit and receive antennas ($M_{\text{Tx}}, M_{\text{Rx}}$) is made by a higher layer operation, based on the network needs.

A. Precoder Design

As shown in Figure 2, the downlink precoder has two stages, a standard MU-MIMO downlink precoder, \mathbf{P}_{Down} , followed by the SoftNull precoder, \mathbf{P}_{Self} . The goal of the SoftNull precoder, \mathbf{P}_{Self} , is to suppress self-interference. The goal of the downlink precoder, \mathbf{P}_{Down} , is for the signal received by each user to contain mostly the signal intended for that user, and little signal intended for other users. The standard MU-MIMO downlink precoder, \mathbf{P}_{Down} , controls D_{Tx} effective antennas. The SoftNull precoder maps the signal on the D_{Tx} effective antennas to the signal transmitted on the M_{Tx} physical transmit antennas, as shown in Figure 2. Let $\mathbf{s}_{\text{Down}} \in \mathbb{C}^{K_{\text{Down}}}$ denote the vector of symbols that the base station wishes to communicate to each of the K_{Down} downlink users. We constrain both stages to be linear, such that \mathbf{P}_{Down} is a $D_{\text{Tx}} \times K_{\text{Down}}$ complex-valued matrix and \mathbf{P}_{Self} is a $M_{\text{Tx}} \times D_{\text{Tx}}$ matrix. The signal transmitted on the base station antennas is then $\mathbf{x}_{\text{Down}} = \mathbf{P}_{\text{Self}}\mathbf{P}_{\text{Down}}\mathbf{s}_{\text{Down}}$.

1) *Standard MU-MIMO Downlink Precoder*: The standard MU-MIMO downlink precoder, \mathbf{P}_{Down} , does not need to

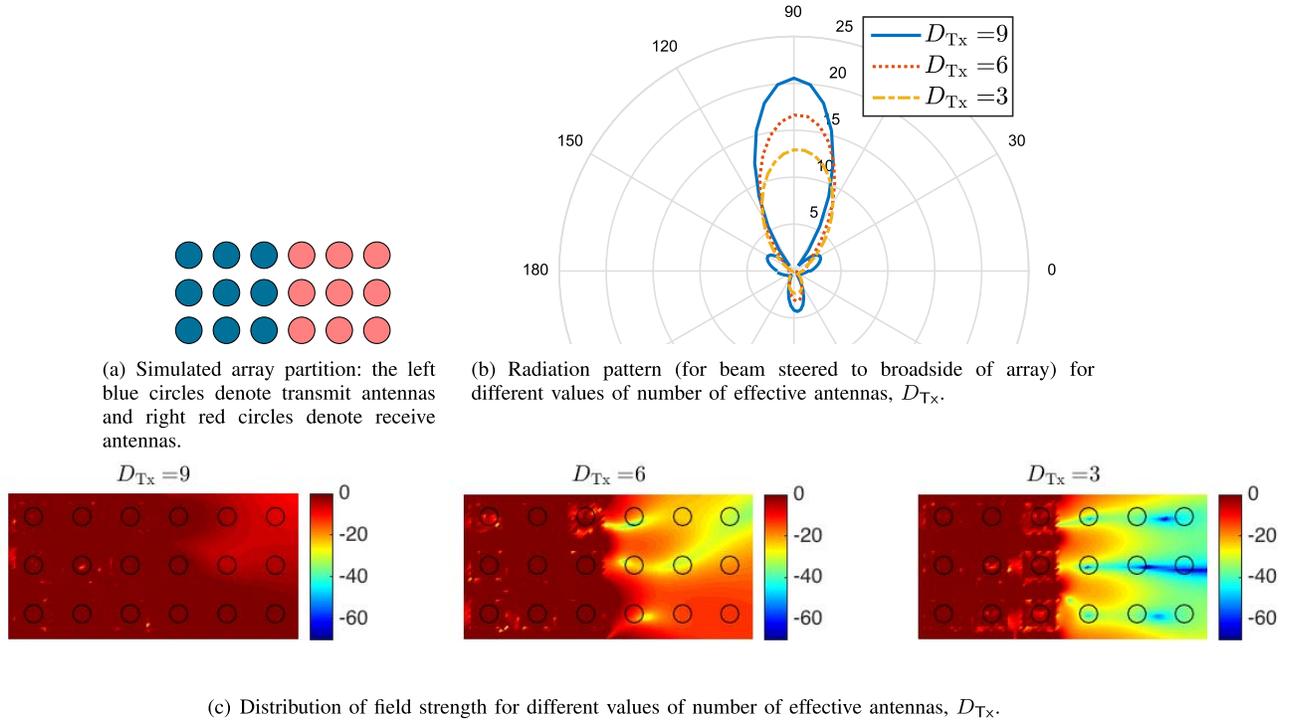


Fig. 3. Simulation example of SoftNull precoder operation. (a) The array simulated. (b) Gain of radiation pattern steered to broadside. (c) Distribution of field strength around the receive antennas.

have knowledge of both the self-interference channel and the downlink channel. Rather the downlink precoder, \mathbf{P}_{Down} , only needs to know the *effective* downlink channel, $\mathbf{H}_{\text{Eff}} = \mathbf{H}_{\text{Down}}\mathbf{P}_{\text{Self}}$, that is created by the SoftNull precoder operating on the physical downlink channel. Note that \mathbf{H}_{Eff} can be estimated directly by transmitting/receiving pilots along the D_{Tx} effective antennas. For the standard MU-MIMO downlink precoder, standard algorithms such as zero-forcing beamforming or matched filtering can be used. For example, in the case of zero-forcing beamforming, the MU-MIMO downlink precoder, \mathbf{P}_{Down} , is the Moore-Penrose (right) pseudoinverse of the effective downlink channel:

$$\mathbf{P}_{\text{Down}} = \mathbf{P}_{\text{Down}}^{(\text{ZFBF})} \equiv \alpha^{(\text{ZFBF})} \mathbf{H}_{\text{Eff}} (\mathbf{H}_{\text{Eff}}^H \mathbf{H}_{\text{Eff}})^{-1}, \quad (3)$$

where $\alpha^{(\text{ZFBF})}$ is a power constraint coefficient.

2) *SoftNull Precoder*: The goal of the SoftNull precoder is to reduce self-interference while preserving a required number of effective antennas, D_{Tx} , for the standard MU-MIMO downlink transmission. As shown in Figure 2 the SoftNull precoder has D_{Tx} inputs as effective antennas, and M_{Tx} outputs to the physical antennas. We assume that the SoftNull precoder has knowledge of the self-interference channel, \mathbf{H}_{Self} . Our goal is to minimize the *total* self-interference power while maintaining D_{Tx} effective antennas. Our choice to minimize total self-interference, rather than choosing a per-antenna metric is twofold: (i) Minimizing total self-interference gives the precoder more freedom in its design. Instead of creating nulls to reduce self-interference at specific antennas, it can optimize placement of nulls such that each null can reduce self-interference to multiple receive antennas. (ii) As is shown in the following, minimizing the total self-interference power leads to a closed-form solution. We therefore formulate the

precoder design problem as:

$$\begin{aligned} \mathbf{P}_{\text{Self}} &= \underset{\mathbf{P}}{\text{argmin}} \|\mathbf{H}_{\text{Self}}\mathbf{P}\|_F^2 \\ &\text{subject to } \mathbf{P}^H \mathbf{P} = \mathbf{I}_{D_{\text{Tx}} \times D_{\text{Tx}}}. \end{aligned} \quad (4)$$

The squared Frobenius norm, $\|\cdot\|_F^2$, measures the total self-interference power. The constraint, $\mathbf{P}^H \mathbf{P} = \mathbf{I}_{D_{\text{Tx}} \times D_{\text{Tx}}}$, forces the precoder to have D_{Tx} orthonormal columns, and hence ensures that D_{Tx} effective antennas are preserved for MU-MIMO downlink signaling.

It is shown in Appendix A that the above optimization problem (4) has the following closed-form intuitive solution. The optimal self-interference precoder is constructed by projecting onto the D_{Tx} left singular vectors of the self-interference channel corresponding to the smallest D_{Tx} singular values. Precisely,

$$\mathbf{P}_{\text{Self}} = \left[\mathbf{v}^{(M_{\text{Tx}} - D_{\text{Tx}} + 1)}, \mathbf{v}^{(M_{\text{Tx}} - D_{\text{Tx}} + 2)}, \dots, \mathbf{v}^{(M_{\text{Tx}})} \right], \quad (5)$$

where $\mathbf{H}_{\text{Self}} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^H$ is the singular value decomposition of the self-interference channel (\mathbf{U} and \mathbf{V} are unitary matrices and $\mathbf{\Sigma}$ is a nonnegative diagonal matrix whose diagonal elements are the ordered singular values) and $\mathbf{v}^{(i)}$ is the i th column of \mathbf{V} . Essentially, the SoftNull precoder is finding the D_{Tx} -dimensional subspace of the original transmit space, $\mathbb{C}^{M_{\text{Tx}}}$, which presents the least amount of self-interference to the receiver.

B. SoftNull Simulation Example

To help clarify the SoftNull design, we provide a simple simulation example that illustrates how SoftNull reduces self-interference by sacrificing effective antennas. Figure 3(a)

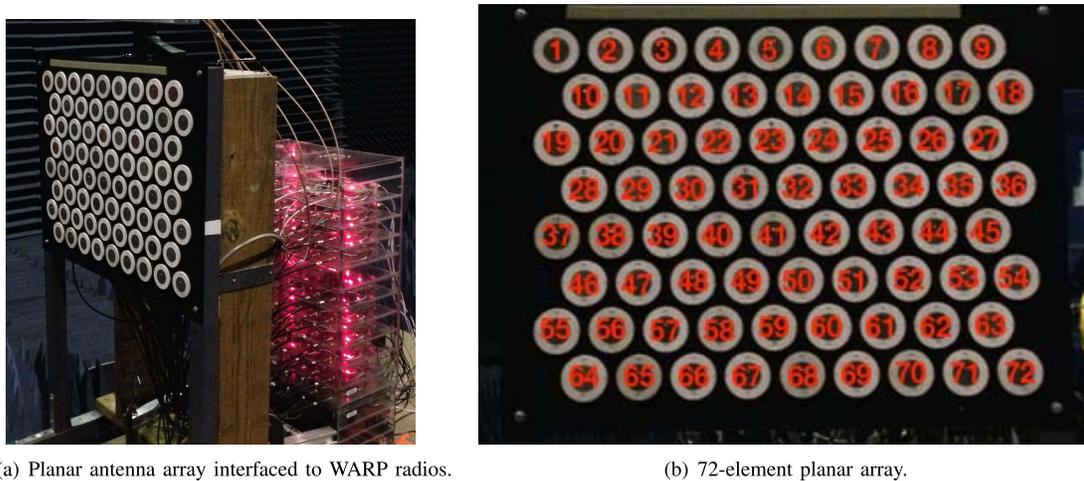


Fig. 4. Platform for channel measurements.

shows a 3×6 ($M = 18$) planar array that is the basis of the simulation. The space between adjacent antenna is half a wavelength. We consider an even $(M_{\text{Tx}}, M_{\text{Rx}}) = (9, 9)$, partition of transmit and receive antennas, as shown in Figure 3(a), where blue circles on the left correspond to the 3×3 transmit subarray, and the red circles on the right correspond to the 3×3 receive subarray.

Fields are computed using MATLAB’s full-wave method-of-moments electromagnetic solver to compute the electric field distribution produced by the transmit antennas in the vicinity of the receive antennas. The full-wave solver enables near-field effects as well as the impact of mutual coupling among the transmit antennas to be captured in the simulation. The antenna’s used are vertically polarized patch antennas, tuned to 2.4 GHz operation. The patch’s resonant dimension is $L = 48.7$ mm, and an air dielectric between patch and ground is assumed.

Figure 3(c) shows the radiated field distribution, in the vicinity of the received antennas, as a function of the number of effective antennas, D_{Tx} . First consider the case where $D_{\text{Tx}} = 9 = M_{\text{Tx}}$, in which no effective antennas are given up for the sake of self-interference reduction: all the receive antennas receive very high self-interference. Then, in the case where $D_{\text{Tx}} = 6$, and three effective antennas are given up for self-interference reduction, the SoftNull precoder essentially steers “soft” nulls towards the receive portion of the array, achieving a small amount of self-interference reduction. In the case of $D_{\text{Tx}} = 3$, the six effective antennas sacrificed allow the SoftNull precoder to significantly reduce self-interference at each receive antenna. The overall trend is that as more effective antennas are given up for the sake of self-interference reduction, the more freedom SoftNull has to create a radiated field pattern which reduces self-interference. Due to computation limits, we have considered a small array in this simulation, for which the self-interference reduction is modest. However, in the large-array experiments that follow, we will see that having a large array enables more design freedom, and thus better self-interference reduction. Figure 3(b) illustrates the downside to using effective antennas for self-interference suppression: reduced transmit gain. In Figure 3(b) we plot the

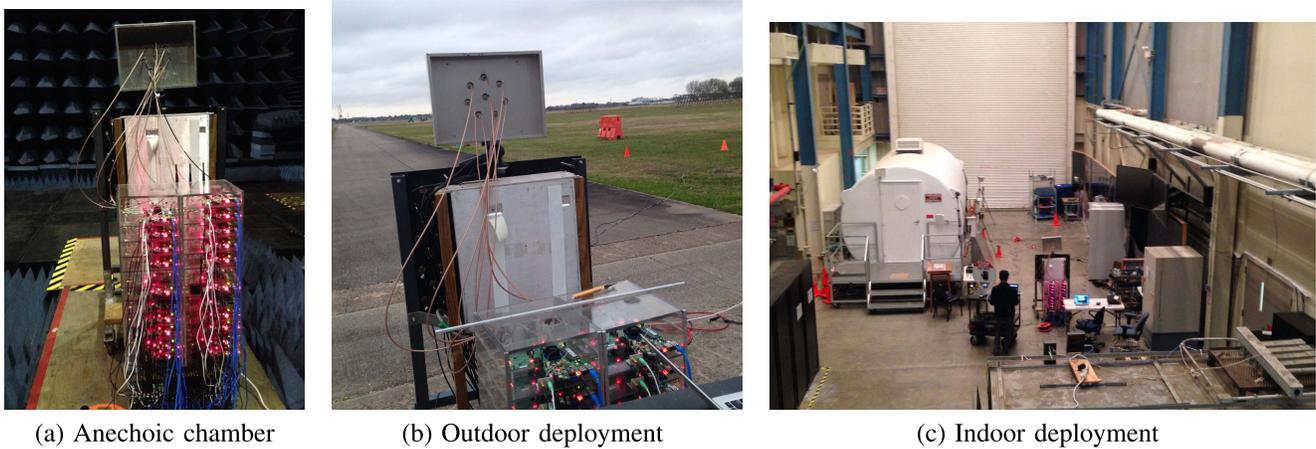
radiation pattern (i.e., far-field power gain) for a beam steered toward the array’s broadside. As we give up more effective antennas for the sake of self-interference reduction, the array gain is decreased. Therefore in the following sections, we will carefully evaluate *whether the benefit in self-interference reduction is worth the loss in beamforming gain*.

IV. CHANNEL MEASUREMENT SETUP

To evaluate the performance of SoftNull, measurements of real self-interference channels and array-to-client channels were collected using the ArgosV2 measurement platform, [35], shown in Figure 4(a). The platform consists of an array of 72 2.4 GHz patch antennas interfaced to 18 WARP v3 boards [36], each with 4 programmable radios. This platform enables 72 base station antennas (transmit or receive). Also four mobile clients are emulated using WARPv3 radios. The antenna array, shown in Figure 4(b) uses custom 2.4 GHz half-wave circular patch [37] antenna elements in a hexagonal grid spaced at 76 mm apart (0.6λ). The antenna elements have roughly 6 dBi gain at broadside. See [37], [38], and references therein for more details on patch antenna design and patch antenna radiation characteristics (input impedance, polarization, radiation patterns, etc.).

To collect traces of real self-interference channel for the 2D array, 20 MHz wideband channel measurements were performed in a diverse set of environments shown in Figure 5. The measured channel traces enable subsequent analysis to obtain an in-depth understanding of coupling between base station antennas, to explore Tx/Rx partitions, and to ultimately simulate real-world performance with clients. Measurements were taken in an anechoic chamber deployment, Figure 5(a), an outdoor deployment, Figure 5(b), and in a highly scattered indoor deployment, Figure 5(c). The outdoor deployment was in an open field, with very few obstructions to cause scattering. Finally the indoor deployment was in a very rich scattering environment, with metal walls and the array placed near a metallic structure as shown in Figure 5(c).

For each deployment, the four clients were placed at three different locations each and channel measurements were performed — both for the 72×72 self-coupling of the array,

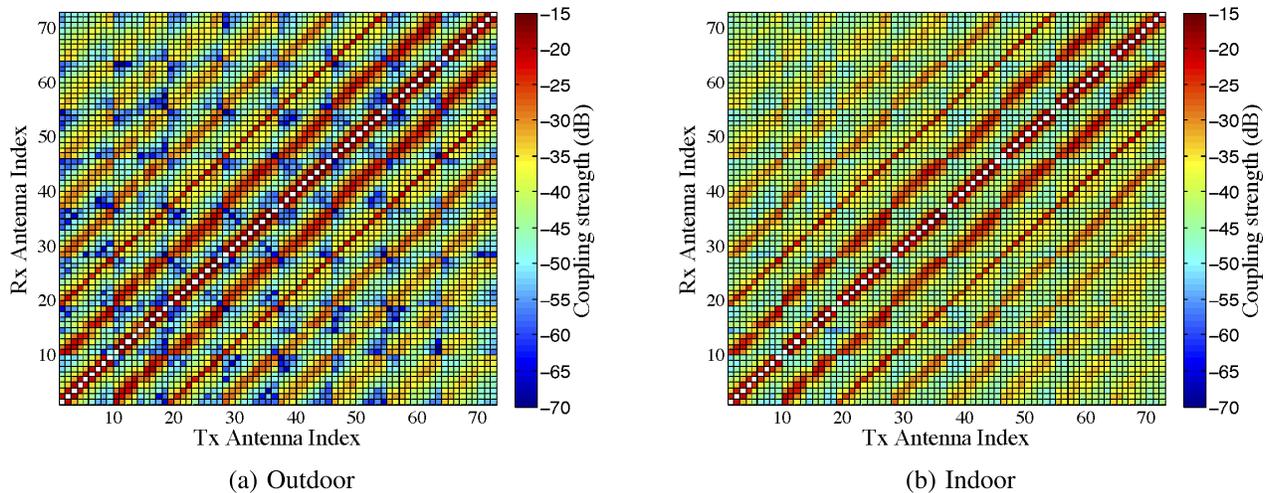


(a) Anechoic chamber

(b) Outdoor deployment

(c) Indoor deployment

Fig. 5. Experiment setup.



(a) Outdoor

(b) Indoor

Fig. 6. Array self-coupling.

and the 72×4 matrix of downlink/uplink channels for each placement. In all, more than 12 million wideband channels were measured, providing more than 40 GB of channel traces for the evaluation of SoftNull performance. Our channel traces are publicly available [39] for other researchers to leverage in their own simulations.

Here we present the results on of the amount of coupling (also called “crosstalk” and “mutual coupling” in the literature) among the elements of the antenna array. Figure 6 shows the strength of the coupling between all pairs of antennas on the array. The color of element (i, j) denotes the average strength of the channel from antenna i to antenna j , where the enumeration of the antenna elements is row-wise as shown in Figure 4(b). Figure 6(a) shows the average coupling among the outdoor measurements and Figure 6 (b) shows the average coupling indoors.³ The banded structure in the matrix is due to counting row-wise (every 9 elements the row ends and a new row begins). Only the coupling between distinct antennas

was measured (i.e. the hardware was not capable of measuring the reflection coefficient for individual antennas), hence the empty diagonal. As expected, there is significant variation in the amount of coupling over the array: antennas adjacent to one another can have a coupling as strong as -15 dB. In the outdoor environment, elements which are far apart can have coupling less than -60 dB, due to the separation and the directionality of the patch antenna elements. In the indoor environment, however, the the coupling between even far away antennas never falls much below -50 dB. The reason for increased coupling indoors is backscattering: for far apart antennas, the coupling due to backscattering from objects in the environment is actually stronger than the direct-path coupling. This phenomenon of backscattering becoming the self-interference bottleneck for well-isolated antennas was also reported in [40].

V. SOFTNULL EVALUATION

In this section, we utilize the collected channel traces to analyze the performance of SoftNull in several aspects. First, we consider how the partition of the whole array into transmit

³For the sake of space the results for the anechoic chamber have been left off, as they were very similar to the results for outdoors and did not yield additional insight. The full datasets can be downloaded at [39].

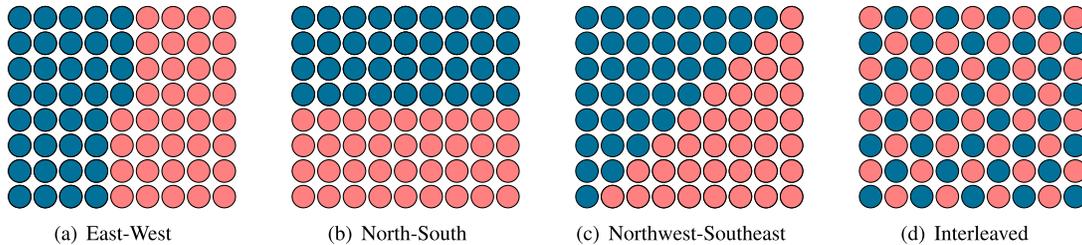


Fig. 7. Tx/Rx partitions heuristics: blue is transmit, red is receive.

and receive sub-arrays impact the ability of SoftNull to reduce self-interference. Next, we study the impact of the propagation environment on the total self-interference reduction by SoftNull. Finally, we consider the uplink and downlink data rate that SoftNull can deliver to clients, and compare with half-duplex and ideal full-duplex systems.

A. Antenna Array Partitioning

Previously, we presented the optimal precoder design of SoftNull, for a given $M_{\text{Rx}} \times M_{\text{Tx}}$ self-interference channel. We now consider how to best partition the M antennas into the set of M_{Tx} transmit antennas and M_{Rx} receive antennas. Due to the combinatorial nature of the problem, finding the optimal antenna sets is computationally difficult. For example, if $M = 72$ and $M_{\text{Tx}} = 36$, then there are $\binom{72}{36} \approx 4.4 \times 10^{20}$ possible combinations of transmit antenna sets. Therefore we focus on empirical insights using the traces collected via channel measurements to evaluate and compare several heuristic choices for partitioning the array.

1) *Heuristic Partitions Considered*: Intuitively, we recognize that SoftNull will perform the best when the power in the self-interference channel is concentrated within a fewer number of eigenchannels. It has been demonstrated both analytically and experimentally [41]–[44] that as the spread of the angles-of-departure from the transmitter to the receiver is decreased, the signal received at each receive antenna becomes more correlated. More correlated signals leads to the first few eigenvalues become more dominant, which is exactly the desirable situation for SoftNull. Contiguous linear partitions of the array (one side transmit, other receive) limit the angular spread of the angles-of-departure to/from the transmitter to the receiver, since all the interference is coming from only one “side” of the array. For example, in the North-South partition of Figure 7(b) the angular spread of angles-of-arrival is less than 180 degrees for all receive antennas, since all interference is coming from the “North.” Figure 7 shows three proposed partitions based on the above heuristic of linear contiguous partitioning to limit angular spread: East-West, North-South, and Northwest-Southwest partitions are shown in Figures 7(a), 7(b), and 7(c), respectively. In the figure, an even split between the number of transmit and receive antennas is assumed.

As a comparison, we also consider the interleaved partition shown in Figure 7(d). If our heuristic of minimizing angular spread is effective, then we would expect the interleaved partition to be a near worst-case partition. In the interleaved

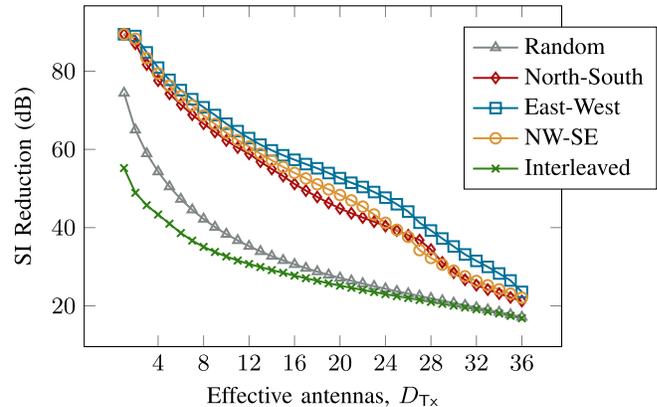


Fig. 8. Achieved tradeoff between self-interference reduction and number of effective antennas remaining for downlink signaling, D_{Tx} . Several different methods of partitioning 72-element array into an even $(M_{\text{Tx}}, M_{\text{Rx}}) = (36, 36)$ Tx/Rx split are compared.

partition receive antennas will experience interference arriving at every possible angle. In addition to the deterministic interleaved partition, we also compare against the average measured performance of 10,000 randomly chosen partitions.

2) *Evaluation of the Heuristic Partitions*: To assess the performance of these heuristics, we directly measured the self-interference channel response in an anechoic chamber using the 72-element rectangular array and a $(M_{\text{Tx}}, M_{\text{Rx}}) = (36, 36)$ partition of transmit and receive elements. We consider the self-interference channel measurements performed in the anechoic chamber, as this is the most repeatable scenario. Figure 8 shows the tradeoff between self-interference (SI) reduction and number of effective antennas, D_{Tx} . As D_{Tx} decreases from its maximum value of $D_{\text{Tx}} = M_{\text{Tx}} = 36$, the amount of self-interference reduction achieved by SoftNull improves. We remind the reader that D_{Tx} is the number of effective antennas preserved for downlink signalling, and thus $(M_{\text{Tx}} - D_{\text{Tx}})$ is the number of effective antennas leveraged for self-interference reduction. As D_{Tx} decreases we are “giving up” effective antennas for the sake of improved self-interference reduction. Therefore as D_{Tx} decreases, we expect to observe better self-interference reduction, as we see in the case of Figure 8. We see in Figure 8, that the tradeoff achieved for the contiguous partitions is much better than that achieved for the random and interleaved partitions.

Consider the tradeoff between self-interference reduction and effective antennas shown in Figure 8. Typical analog cancellation circuits provide 40–50 dB self-interference

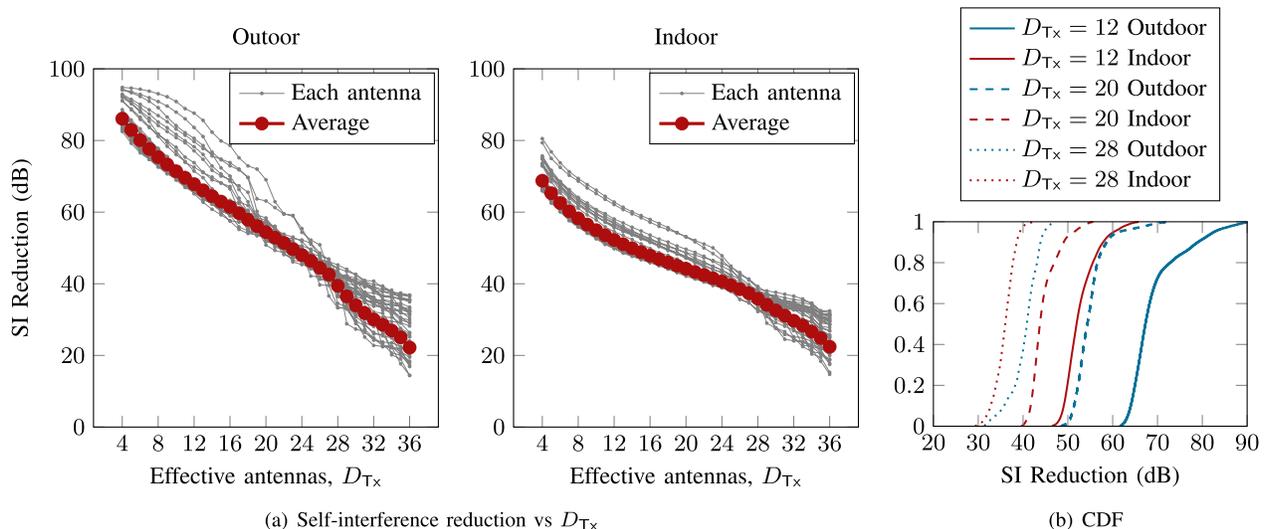


Fig. 9. Self-interference reduction achieved by SoftNull as a function of the number of effective antennas preserved, D_{Tx} . Better self-interference reduction is observed outdoors, due to less backscattering, which leads to a more correlated self-interference channel matrix.

reduction [18]. Therefore, an interesting point of observation is how many effective antennas can be preserved while achieving the ≈ 50 dB self-interference reduction similar to that of an analog canceler. For the random partition, we can only preserve 6 of the maximum 36 effective antennas while achieving > 50 dB self-interference reduction, but for all of the contiguous partitions, we can achieve > 50 dB reduction with at least 16 effective antennas preserved. The best performing partition is the East-West partition, which is in line with our heuristic: among the considered partitions, the East-West partition is the one with minimum angular spread between the transmit and receive partitions, since it splits the array along its smallest dimension (array is wider than tall). The interleaved partition performs even worse than the average of random partitions, emphasizing the importance of selecting contiguous partitions. Finally, note the large impact of the partition type on the tradeoff between self-interference reduction and number of effective antennas. For $D_{Tx} \in [3, 22]$, the East-West partition enables SoftNull to achieve more than 25 dB better self-interference reduction than the average of partitions chosen at random.

B. Impact of Scattering on Self-Interference Reduction

The scattering environment has a significant impact on the performance of SoftNull. We use the collected traces to study how the scattering environment impacts the tradeoff between self-interference reduction and effective antennas achieved by SoftNull. The 72-element array is used, with a $(M_{Tx}, M_{Rx}) = (36, 36)$, East-West partition of the transmit and receive elements, as shown in Figure 7(a). Figure 9(a) compares the tradeoff between self-interference reduction and preserved effective antennas in the outdoor deployment versus the indoor deployment. The thin gray lines correspond to the self-interference reduction achieved for each of the 36 antennas, while the thick red line corresponds to the self-interference reduction averaged over all 36 antennas. Figure 9(b) shows the empirical cumulative distribution function of the achieved self-interference reduction, both indoors

and outdoors, for a selection of values for the number of effective antennas preserved.

We see in Figure 9(a) that with all 36 effective antennas preserved the self-interference is only suppressed (passively) by 20 dB. But by giving up 16 effective antennas and preserving $D_{Tx} = 20$ effective antennas for the downlink, the self-interference is suppressed more than by 50 dB. We also see in Figure 9(a), however, that the self-interference reduction in the indoor deployment is not nearly as good as in the outdoor deployment. To achieve 50 dB self-interference reduction in the indoor deployment, 24 of the 36 effective antennas must be given up (as opposed to 16 outdoors), leaving $D_{Tx} = 12$ for downlink transmission. The same array was used in both environments, the only difference being the backscattering environment. The reason for better performance outdoors than indoors is that the backscattering reduces the correlation of the self-interference among antennas that is present in a low scattering environment. Less correlation makes it harder to suppress the self-interference at multiple antennas without giving up more effective antennas. More precisely, the SoftNull precoder projects the transmit signal onto the D_{Tx} singular vectors corresponding to the smallest D_{Tx} singular vectors. In other words, SoftNull reduces self-interference by avoiding the $(M_{Tx} - D_{Tx})$ dominant modes of the self-interference channel. Outdoors, the direct paths between antennas dominate any backscattered paths, leading to a more correlated self-interference matrix, and hence a large amount of the overall channel power resides in the dominant $(M_{Tx} - D_{Tx})$ modes (singular values). Therefore very good self-interference reduction is achieved by avoiding just these first few dominant modes. Indoors, however, multipath backscattering tends to decorrelate the self-interference channel and thus leads to a more uniform distribution of power over the modes. Therefore, in the indoor environment less self-interference is suppressed by avoiding the $(M_{Tx} - D_{Tx})$ most dominant modes.

Figure 9(b) shows the empirical cumulative distribution function of the achieved self-interference reduction, both

indoors and outdoors, for a selection of values for the number of effective antennas preserved, D_{Tx} . We see in Figure 9(b) that for small D_{Tx} there is much more variation in the achieved self-interference reduction outdoors than there is indoors. For the outdoor deployment with $D_{\text{Tx}} = 12$, the self-interference reduction for a given antenna can be as much 90 dB and as little as 62 dB, a 28 dB difference. Indoors, however there is much less variation. For the indoor deployment with $D_{\text{Tx}} = 12$, the difference between best and worst self-interference reduction is only 10 dB. More variation outdoors than indoors is also due to less backscattering outdoors than indoors. Outdoors, the backscattering is nearly nonexistent and direct paths between transmit and receive antenna dominate even for small D_{Tx} . The characteristics of the direct-path self-interference channel seen by each receive antenna vary greatly. For example, the receive antennas nearest the transmit antennas see less correlation among the transmit antennas (because of smaller angular spread) than the receive antennas farther away from the transmit antennas. Indoors, however, for smaller D_{Tx} the self-interference is dominated by backscattered paths. Unlike the direct paths, the characteristics of the direct-path self-interference channel seen by each receive antenna do not vary as much. Therefore for small D_{Tx} , we expect to see more variation in self-interference reduction over the array outdoors than we see indoors.

C. Achievable Rate Gains Over Half-Duplex

In the previous subsections, we observed that SoftNull enables the array to reduce self-interference by giving up a fraction of the available effective antennas for downlink, so that they may be used for self-interference reduction via beamforming. The question remains as to whether the gain in self-interference reduction is worth the cost of giving up the required effective antennas. In particular, we wish to understand the scenarios in which SoftNull can provide improved data rates over conventional half-duplex systems, and when SoftNull cannot. The self-interference channels used in simulation are the self-interference channel traces collected as described in Section IV. The uplink/downlink channel measurements are limited to 4 clients, therefore we use a mix of measured uplink/downlink channels and simulated uplink/downlink channels (especially in the case where more than four clients are considered).

Each trace consists of a 72×72 measurement of the self-interference channel, a 72×4 measurements of the downlink channel to the mobile clients, and a 4×72 measurement of the uplink channel. Each channel measurement is done over 64 OFDM subcarriers in channel 4 (20 MHz band) of the 2.4 GHz ISM band. We compute the uplink and downlink rates achievable for each trace by implementing the precoder and equalizer for the measured channels. We then compute achievable rates from the uplink and downlink signal-to-interference-plus-noise ratios achieved after equalization and precoding (assuming optimal channel codes). Since we are focusing in this study on small cells which use lower transmit power than macro cells, we assume a sum-power constraint at the array of 0 dBm and a client constraint of -10 dBm. Also,

we assume a receiver noise power of -95 dBm for both the array and client.

When simulating half-duplex MU-MIMO, we allow the array to use all $M = 72$ transmit antennas, and constrain the array to divide time equally between uplink and downlink operation. The half-duplex MU-MIMO downlink uses the standard zero-forcing precoder and the uplink uses the standard linear decorrelator equalizer. When simulating SoftNull we allow the array to operate the downlink and uplink simultaneously, but constrain the array to use $M_{\text{Tx}} = 36$ antennas that are for transmission only and $M_{\text{Rx}} = 36$ for reception. Based on the experiments and discussion in Section V-A, we choose the East-West partition of transmit and receive antennas shown in Figure 7(a). In simulation of SoftNull, we set the downlink precoder, \mathbf{P}_{Down} (applied before the SoftNull), to be the standard zero-forcing precoder, after which we apply the SoftNull precoder, \mathbf{P}_{Self} . Receive-side digital cancellation of the self-interference is also simulated for SoftNull. The effects of receiver limitations (in particularly AD quantization and LNA desensitization) are modeled via the dynamic range model of [16] and [17] (and references therein). In this dynamic range mode, Gaussian noise is added to the received signal in proportion to the power of the received signal+interference (see Appendix C for more explanation of this model). The constant of proportionality is the *dynamic noise figure*, D_0 , which we conservatively select to be 25 dB. Setting the dynamic noise figure to 25 dB effectively limits the amount of digital cancellation that can be achieved to no more than 25 dB. Therefore all other self-interference reduction must come via the beamforming performed by the SoftNull downlink precoder. We also compare against “ideal full-duplex,” where transmission and reception occurs at the same time but the self-interference is zero. Note that the ideal full-duplex rate will be less than twice the half-duplex rate, because we still assume that even for ideal full-duplex, $M_{\text{Tx}} = 36$ antennas that are for transmission and $M_{\text{Rx}} = 36$ for reception, as opposed to half-duplex which uses all antennas for both transmission and reception, but in separate time slots. The details on how achievable rates are computed from the simulation results are given in Appendix B.

1) *Achievable Rates for Measured Channels, $K = 4$ Clients:* We assume the number of uplink and downlink clients are the same and both equal to four, $K_{\text{Up}} = K_{\text{Down}} = K = 4$. In this case we can directly use the uplink and downlink channel traces that were measured. Figure 10 shows the achievable uplink, downlink, and sum rates achieved by SoftNull as a function of the chosen number of preserved effective antennas, D_{Tx} , and compares SoftNull’s performance to that of half-duplex as well as against ideal full-duplex.

First consider the results for the channels collected in the outdoor deployment, shown in Figure 10(a). The downlink rate achieved by SoftNull increases as D_{Tx} increases, since more effective antennas become available to beamform and thus create a better signal-to-interference-plus-noise ratio to the downlink clients. However, as D_{Tx} increases the uplink rate decreases because SoftNull can suppress less self-interference when more effective antennas are used for downlink beamforming. Note that once D_{Tx} is less than approximately

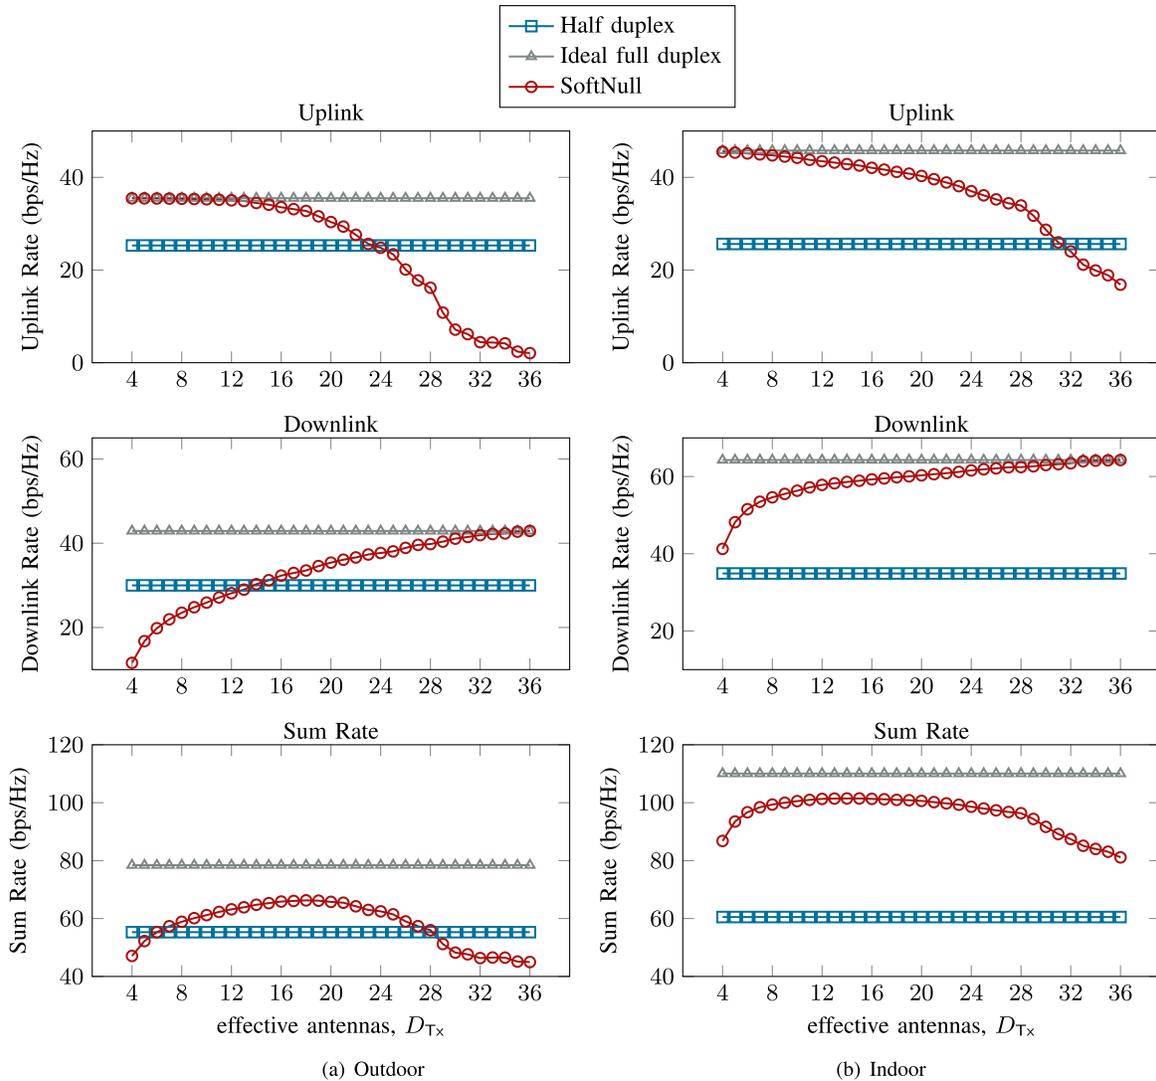


Fig. 10. Achievable rates of SoftNull versus half-duplex as a function of the number of effective antennas preserved, D_{Tx} . An east-west $(M_{Tx}, M_{Rx}) = (36, 36)$ partition is considered. SoftNull performs better indoors than outdoors because of less path loss.

12 the incremental gain in uplink rate from giving up each additional effective antenna is negligible. For example, at $D_{Tx} = 12$, the self-interference is sufficiently suppressed to no longer overwhelm the receiver, and digital cancellation removes remaining self-interference. Giving up more than 12 effective antennas improves the uplink only slightly but greatly decreases the downlink rate. There is a range of values for D_{Tx} for which SoftNull outperforms half-duplex both for the uplink and the downlink. The bottom plot of Figure 10(a) shows the sum rate (uplink rate plus downlink rate). We see that SoftNull outperforms half-duplex for $D_{Tx} \in [5, 28]$, achieving peak performance at $D_{Tx} = 18$. The achieved rate at $D_{Tx} = 18$ is 23% better than half-duplex, but still 15% less than the ideal full-duplex performance.

Now consider the performance of SoftNull for the indoor channel traces collected, shown in Figure 10(b). In the indoor environment SoftNull outperforms half-duplex for all values of D_{Tx} , with the best performance coming at $D_{Tx} = 14$, for which a 62% gain over half-duplex is achieved, but is still 12% less than ideal full duplex.

Figure 11 shows how the chosen antenna partition affects the achievable rate performance, in both the outdoor and indoor environment. We know from Section V-A (Figure 8) that contiguous partitions (East-West, North-South, and Northwest-Southeast) all provide good suppression, with the East-West partition providing slightly more suppression than the others. The interleaved partition provides very poor suppression. Based on these results from Figure 8, the achievable rate performance in the indoor environment, shown in Figure 11(b), is expected: the East-West partition provides the best data rates, with the other contiguous partitions close behind, and the interleaved partition leading to poor data rates. The achievable rate performance in the outdoor environment is surprising, at first glance: we see in Figure 11(a) that the North-South partition provides significantly better data rates than the other partitions, and the East-West partition (which has slightly better suppression) gives the lowest data rates among the three contiguous partitions. In the outdoor experiment, the users were in line-of-sight of the array and at the same ground level as the array. In other words, the

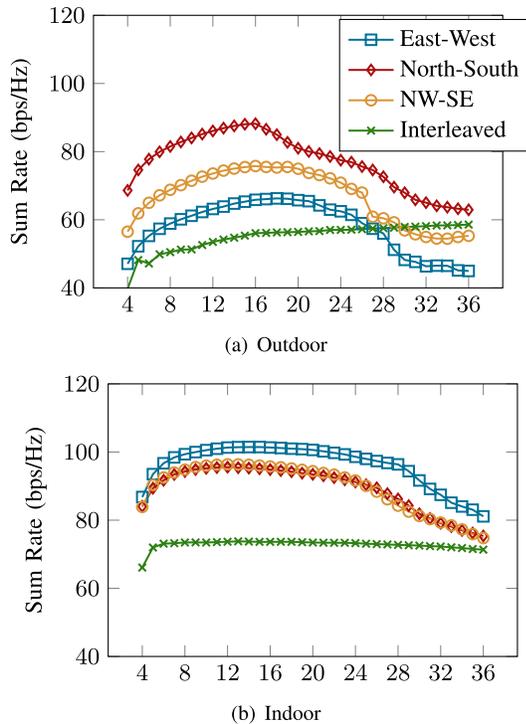


Fig. 11. Achievable rates for different antenna partitions.

paths between the users and the array vary in azimuth but are approximately the same in elevation. Of the three contiguous partitions, the North-South partition has the longest horizontal dimension, and shortest vertical dimension, meaning that it has the greatest azimuth resolution and the least elevation resolution. Because of the greater azimuth resolution, the vectors to each user are less aligned than they would be for other partitions, and thus the greatest SNR to each user can be achieved. The East-West partition has the least azimuth resolution, and hence has the worst performance of the contiguous partitions for the outdoor test. In the indoor environment, because of multipath scattering, the paths between the array and the users have significant variation in both azimuth and elevation. Therefore in the indoor environment the array's resolution has much less impact on beamforming performance. Because the East-West partition has slightly better self-interference suppression, the East-West partition yields the highest data rates. The take-away is that the choice of array partition depends not only on self-interference suppression performance, but also on matching the arrays' angular resolution to the angular distribution of the users that will be served.

At first, it seems surprising that the gains over half-duplex are *better* indoors than outdoors, when we saw in Figure 9(a) that the self-interference reduction achieved indoors is *worse* than that achieved outdoors. The difference is that the path loss for the channels measured indoors was much less than that measured outdoors. The clients indoors were necessarily placed closer to the array (4-9 m) because of limited space, but outdoors were placed much farther (9-15 m). Full-duplex always becomes more challenging as path loss increases.

Larger path loss means the uplink signal is weaker, and therefore more self-interference reduction is required to make

the self-interference commensurate in power to the uplink signal. Hence, full-duplex systems to date have only been demonstrated for small cells [18] (see Appendix C for more explanation and a numerical example on how path loss and dynamic range limits the amount of required self-interference suppression). For SoftNull in particular, larger path loss means more effective antennas must be given up to achieve better self-interference reduction. Larger path loss also means that more effective antennas are needed to achieve a sufficient signal strength on the downlink, therefore the cost of using effective antennas for the sake of reducing self-interference is greater. Because the path loss was greater in the outdoor deployment than the indoor, the gains of SoftNull are less for the outdoor deployment than for the indoor deployment. Even though the achieved self-interference is better outdoors than indoors, the benefit of better suppression does not compensate for the greater path loss.

It is also interesting to note in Figure 10, that although the ideal full-duplex rates are improved significantly indoors relative to outdoors, due to lower path loss, the half-duplex rates are only improved slightly indoors relative to outdoors. The reason is that half-duplex has a much larger beamforming gain than full-duplex, because it uses the full 72 antennas for transmission and reception, whereas ideal full-duplex only has 36 antennas (one partition for transmission, the other for reception). For half-duplex, with the large beamforming gain from 72 antennas, the received power is so large that the dominant noise source is not the thermal noise floor, which is independent of received power, but dynamic range limitations such as quantization noise, oscillator noise, etc, which are *proportional* to the received power (see Appendix C for details of the dynamic range model). Lower path loss increases the received signal power, but also increases the noise due to dynamic range limitations by an equal amount, hence the modest improvement in SNR for lower path loss. For ideal full-duplex, the beamforming gain is much less since there are only 36 antennas, and received power is low enough for the thermal noise floor to be the dominant noise source. Lowering the path loss increases the received signal power relative the thermal noise floor, therefore significantly increasing the SNR, and leading to higher rates.

We will next consider simulated uplink and downlink channels, so that path loss can be controlled for a fair comparison of outdoor versus indoor, and so that much larger path loss values can be considered.

Figure 12 compares the outdoor and indoor performance of SoftNull for path loss values of 70, 85 and 100 dB. At 2 GHz these values correspond roughly to distances of 50, 300, and 1000 m in outdoor line-of-sight conditions (i.e. assuming path loss exponent of 2). Indoors, these path loss values correspond to distances of 3, 10, and 30 m for indoor non line-of-sight conditions (i.e. assuming a path loss exponent of 3) [45]. The self-interference channels are the channels measured indoors and outdoors, but the uplink and downlink channels are i.i.d. Rayleigh distributed channels, with a controlled path loss. Both outdoor and indoor, as the path loss increases the gain of SoftNull over half-duplex decreases. More path

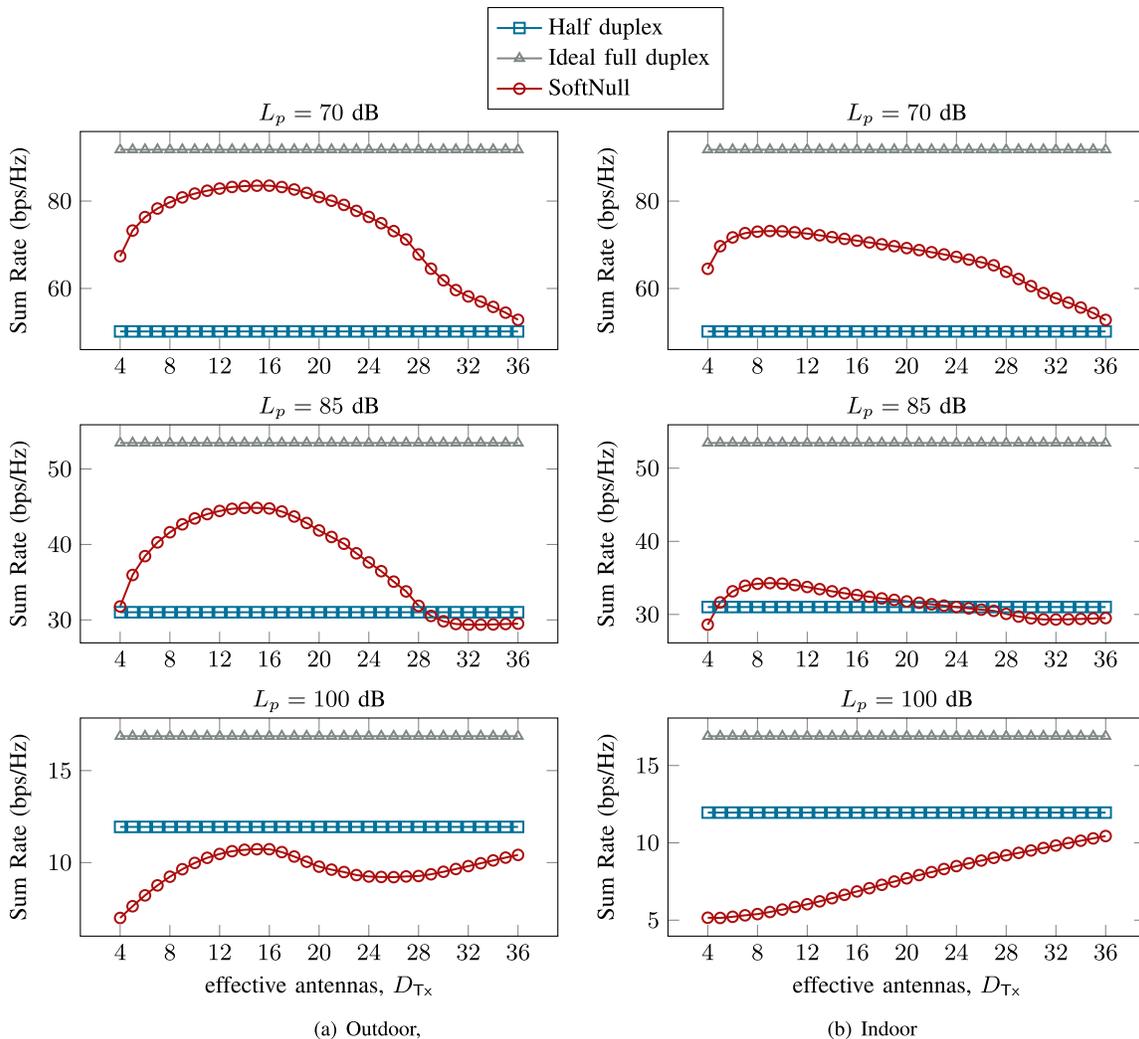


Fig. 12. Achievable sum rates of SoftNull versus half-duplex as a function of the number of effective antennas preserved, D_{T_X} , for several different values of path loss, L_p . An east-west $(M_{T_X}, M_{R_X}) = (36, 36)$ partition is considered. The larger the path loss, the more effective antennas must be given up to achieve sufficient self-interference reduction. SoftNull performance is better outdoors, because of better self-interference reduction due to more correlation in the self-interference channel.

loss means that more self-interference must be suppressed to make self-interference commensurate to the uplink received signal strength, therefore we see in Figure 12 that as the path loss increases, the optimal value of D_{T_X} decreases in order to provide the needed self-interference reduction. Also, as expected, when the path loss is the same in both environments, the performance is better outdoors than indoors, because (as discussed above) SoftNull can better reduce self-interference outdoors than indoors. For the outdoor deployment, SoftNull can significantly outperform half-duplex at 70 and 85 dB paths loss, but at 100 dB path loss, SoftNull cannot outperform half-duplex for any value of D_{T_X} . At 100 dB path loss indoors, too many effective antennas must be given up to achieve the required self-interference reduction, and the downlink rate suffers accordingly. Similarly in the indoor environment, where full-duplex is more challenging because of less self-interference suppression achieved, SoftNull can significantly outperform half-duplex at 70 dB path loss, just barely outperforms half-duplex at 85 dB paths loss, and underperforms half-duplex at 100 dB path loss.

D. Varying Number of Clients

Figure 13 shows how the performance of SoftNull compares to that of half-duplex as a function of the number of uplink clients and downlink clients (which are assumed to be equal). For half-duplex, adding clients is an opportunity to provide higher sum rates via spatial multiplexing. Likewise for full-duplex with SoftNull, adding more clients provides a multiplexing opportunity, but adding more clients also means that fewer effective antennas can be given up for the sake of self-interference reduction. For each of the data points plotted for SoftNull, we assume SoftNull selects the D_{T_X} value that maximizes the sum rate. The simulations were carried out for path loss values of 70, 85, and 100 dB. We see in Figure 13 that SoftNull outperforms half-duplex for small numbers of clients (except when the path loss is large). However, as more clients are added, eventually a point is reached at which SoftNull underperforms half-duplex. Half-duplex benefits from the fact all 72 antennas can be used for uplink or downlink transmission and that no effective antennas

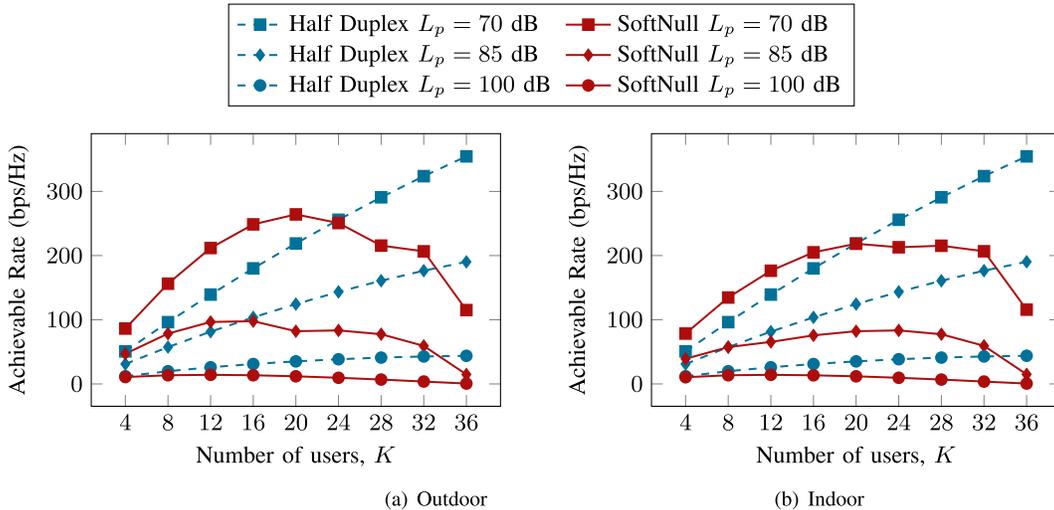


Fig. 13. Achievable rate of SoftNull and half-duplex as function of the number of clients being served, K . Results are for array with $M = 72$, and $(M_{\text{TX}}, M_{\text{RX}}) = (36, 36)$, both outdoor (a) and indoor (b). For low number of clients SoftNull outperforms half-duplex, and for larger number of clients SoftNull underperforms half-duplex operation.

must be sacrificed for self-interference suppression. Therefore, in systems where the number of users served is similar to the number of antennas (i.e. when $K \approx M$), the spatial-multiplexing benefit from having access to more antennas in half-duplex outweighs the time-multiplexing benefit from transmitting and receiving simultaneously in full-duplex via SoftNull. However, in systems where the number of users served in a given time slot is small relative to the number of antennas (i.e. when $K \ll M$) – the regime-of-interest for many-antenna wireless (also called “massive MIMO”) [13]–[15] – the benefit of SoftNull is more pronounced. We see in Figure 13 that the threshold on K below which SoftNull outperforms full-duplex depends heavily on the path loss. Consider the outdoor case, shown in Figure 13. When the path loss is only 70 dB, serving more clients increases SoftNull performance up to $K \leq 20$, after which SoftNull cannot serve more clients without allowing prohibitive self-interference. Therefore the rate decreases for $K > 20$, and SoftNull underperforms half-duplex for $K > 20$ clients. As the path loss increases, more effective antennas must be given up to sufficiently suppress self-interference and avoid swamping the uplink signal, and hence fewer clients can be served. For 85 dB path loss, SoftNull outperforms half-duplex when $K \leq 12$. And in the 100 dB path loss case SoftNull underperforms half-duplex for each value of K simulated. In the indoor environment, more effective antennas must be used to achieve the same self-interference reduction, therefore the value for K at which SoftNull outperforms half-duplex is smaller. For 70 dB path loss SoftNull outperforms half-duplex for $K \leq 16$, and for 85 dB path loss SoftNull outperforms half-duplex for $K \leq 4$. At 100 dB path loss, SoftNull strictly underperforms half-duplex in the indoor deployment.

VI. CONCLUSION

SoftNull provides an opportunity to enable full-duplex operation with current base station radios without requiring additional circuitry for analog cancellation. The primary intuition

behind the SoftNull precoder is that the self-interference need not be perfectly nulled; we only need to sacrifice the minimum number of effective antennas required to sufficiently suppress self-interference. Our analysis based on the channels measured using a 72-element array shows that when the path loss is not too large, sufficient self-interference reduction can be achieved while only using a portion of the effective antennas for self-interference suppression. However, due to additional backscattering indoors, more effective antennas must be used for self-interference reduction indoors to achieve the same level of self-interference reduction outdoors. SoftNull also performs quite well in the regime where $M_{\text{TX}} \gg K$. We note that the trend in wireless deployments is towards larger arrays, and SoftNull can provide a new full-duplex transmission mode using only half-duplex hardware. We would like to point the readers to a extension of this work in the related thesis [46], in which a software-defined radio prototype of SoftNull is presented.

APPENDIX A OPTIMAL PRECODER SOLUTION

The optimal precoder is the solution to the optimization problem

$$\mathbf{P}_{\text{Self}} = \underset{\mathbf{P}}{\text{argmin}} \|\mathbf{H}_{\text{Self}} \mathbf{P}\|_F^2, \quad (6)$$

$$\text{subject to } \mathbf{P}^H \mathbf{P} = \mathbf{I}_{D_{\text{TX}} \times D_{\text{TX}}}. \quad (7)$$

Via simple manipulations, the optimization function in can be rewritten as

$$\mathbf{P}_{\text{Self}} = \underset{\mathbf{P}}{\text{argmin}} \frac{1}{2} \text{Tr} \left(\mathbf{P}^H \mathbf{H}_{\text{Self}}^H \mathbf{H}_{\text{Self}} \mathbf{P} \right), \quad (8)$$

$$\text{subject to } \mathbf{P}^H \mathbf{P} = \mathbf{I}_{D_{\text{TX}} \times D_{\text{TX}}}. \quad (9)$$

Problem (8) is a convex optimization problem with equality constraints, and can thus be solved by the method of Lagrange

multipliers. The Lagrangian for Problem (8) is

$$L(\mathbf{P}, \mathbf{\Lambda}) = \frac{1}{2} \text{Tr} \left(\mathbf{P}^H \mathbf{H}_{\text{Self}}^H \mathbf{H}_{\text{Self}} \mathbf{P} \right) - \frac{1}{2} \text{Tr} \left(\mathbf{\Lambda} \mathbf{P}^H \mathbf{P} \right),$$

$$\mathbf{\Lambda}^H = \mathbf{\Lambda} \quad (10)$$

The gradient of the Lagrangian with respect to \mathbf{P} is $\nabla_{\mathbf{P}} L(\mathbf{P}, \mathbf{\Lambda}) = \mathbf{H}_{\text{Self}}^H \mathbf{H}_{\text{Self}} \mathbf{P} - \mathbf{P} \mathbf{\Lambda}$. Therefore the stationary points of the Lagrangian must satisfy

$$\mathbf{H}_{\text{Self}}^H \mathbf{H}_{\text{Self}} \mathbf{P} = \mathbf{P} \mathbf{\Lambda} \quad (11)$$

Let $\{\lambda_i\}_{i=1}^{M_{\text{Tx}}}$ and $\{\mathbf{v}_i\}_{i=1}^{M_{\text{Tx}}}$ denote the eigenvalues and eigenvectors of $\mathbf{H}_{\text{Self}}^H \mathbf{H}_{\text{Self}}$, and we assume that the set of eigenvalues is ordered from largest to smallest λ_i . It is easy to check that (11) is satisfied if $\mathbf{\Lambda}$ is diagonal, and the j th column of \mathbf{P} is an eigenvector of $\mathbf{H}_{\text{Self}}^H \mathbf{H}_{\text{Self}}$, with the j th diagonal element of $\mathbf{\Lambda}$ set to the corresponding eigenvalue. Note that the columns of \mathbf{P} need not be distinct to satisfy (11). For example, the columns of \mathbf{P} could each be the same eigenvector. However, \mathbf{P} must also satisfy the constraint $\mathbf{P}^H \mathbf{P} = \mathbf{I}_{D_{\text{Tx}} \times D_{\text{Tx}}}$. Note that since $\mathbf{H}_{\text{Self}}^H \mathbf{H}_{\text{Self}}$ is a hermitian $M_{\text{Tx}} \times M_{\text{Tx}}$ matrix, there exists a set of M_{Tx} orthonormal eigenvectors of $\mathbf{H}_{\text{Self}}^H \mathbf{H}_{\text{Self}}$. Hence, $\mathbf{P}^H \mathbf{P} = \mathbf{I}_{D_{\text{Tx}} \times D_{\text{Tx}}}$ is satisfied when the D_{Tx} columns of \mathbf{P} consist of D_{Tx} distinct eigenvectors of $\mathbf{H}_{\text{Self}}^H \mathbf{H}_{\text{Self}}$. Therefore, any matrix \mathbf{P} whose columns are distinct eigenvectors of $\mathbf{H}_{\text{Self}}^H \mathbf{H}_{\text{Self}}$ is a both a feasible point and a stationary point of the Lagrangian and thus a candidate for an optimal solution. It remains to determine which choice of eigenvectors for the columns of \mathbf{P} leads to the smallest value of the objective function $\frac{1}{2} \text{Tr}(\mathbf{P}^H \mathbf{H}_{\text{Self}}^H \mathbf{H}_{\text{Self}} \mathbf{P})$. It is fairly obvious that the eigenvectors corresponding to the smallest eigenvalues will lead to the optimal value, but we include the details below. Let $\mathcal{P} \subset \{1, 2, \dots, M_{\text{Tx}}\}$ denote the set of D_{Tx} indices corresponding to which eigenvectors of $\mathbf{H}_{\text{Self}}^H \mathbf{H}_{\text{Self}}$ constitute the columns of \mathbf{P} . One can check that

$$\frac{1}{2} \text{Tr} \left(\mathbf{P}^H \mathbf{H}_{\text{Self}}^H \mathbf{H}_{\text{Self}} \mathbf{P} \right) = \frac{1}{2} \sum_{i \in \mathcal{P}} \lambda_i, \quad (12)$$

which is minimized when $\mathcal{P} = \{M_{\text{Tx}} - D_{\text{Tx}} + 1, M_{\text{Tx}} - D_{\text{Tx}} + 2, \dots, M_{\text{Tx}}\}$, the indices corresponding to the smallest eigenvalues of $\mathbf{H}_{\text{Self}}^H \mathbf{H}_{\text{Self}}$. Therefore, the solution to Problem (8) is

$$\mathbf{P}_{\text{Self}} = \left[\mathbf{v}^{(M_{\text{Tx}} - D_{\text{Tx}} + 1)}, \mathbf{v}^{(M_{\text{Tx}} - D_{\text{Tx}} + 2)}, \dots, \mathbf{v}^{(M_{\text{Tx}})} \right]. \quad (13)$$

Note that the eigenvectors of $\mathbf{H}_{\text{Self}}^H \mathbf{H}_{\text{Self}}$ are equal to the right singular vectors of \mathbf{H}_{Self} , where $\mathbf{H}_{\text{Self}} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^H$, is the singular value decomposition of \mathbf{H}_{Self} . Therefore, an equivalent characterization of the optimal solution is that the columns of \mathbf{H}_{Self} are drawn from the right singular vectors of \mathbf{H}_{Self} corresponding to the D_{Tx} smallest singular values.

APPENDIX B RATE COMPUTATION DETAILS

Uplink and downlink achievable rates are computed as follows. We use the common method of computing ergodic achievable rates from the per-packet signal-to-interference-plus-noise ratios [47]–[49]. Let the variable

$\mathbf{S} \in \{\text{HD}, \text{SoftNull}, \text{IdealFD}\}$ be an index into the set of schemes we are comparing: half-duplex, SoftNull, and ideal full-duplex. Let $\alpha_{\text{Down}}^{(\mathbf{S})} \in [0, 1]$ denote the fraction of time scheme \mathbf{S} allocates for downlink communication, and likewise $\alpha_{\text{Up}}^{(\mathbf{S})} \in [0, 1]$ denote the fraction of time allocated for uplink. Obviously, since the schemes ideal-duplex and SoftNull are full-duplex schemes in which uplink and downlink are always simultaneously active, $\alpha_{\text{Down}}^{(\text{IdealFD})} = \alpha_{\text{Up}}^{(\text{IdealFD})} = 1$ and $\alpha_{\text{Down}}^{(\text{SoftNull})} = \alpha_{\text{Up}}^{(\text{SoftNull})} = 1$. However, for half-duplex, uplink and downlink must occur on separate time slots (or frequency bands). We assume an even allocation between uplink time and downlink time, that is $\alpha_{\text{Down}}^{(\text{HD})} = \alpha_{\text{Up}}^{(\text{HD})} = \frac{1}{2}$. When simulating scheme \mathbf{S} , for each channel realization p , we measure the signal-to-interference-plus-noise ratio at each downlink user j , $\text{SINR}_{\text{Down}}^{(\mathbf{S})}(p, j)$. The average ergodic achievable downlink rate is computed by summing over the downlink users and averaging over channel realizations

$$R_{\text{Down}}^{(\mathbf{S})} = \alpha_{\text{Down}}^{(\mathbf{S})} \frac{1}{N_p} \sum_{p=1}^{N_p} \sum_{j=1}^{K_{\text{Down}}} \log_2[1 + \text{SINR}_{\text{Down}}^{(\mathbf{S})}(p, j)], \quad (14)$$

where N_p is the number of channel realizations. Likewise, the average ergodic achievable uplink rate is computed by summing the rates over the uplink users and averaging over channel realizations

$$R_{\text{Up}}^{(\mathbf{S})} = \alpha_{\text{Up}}^{(\mathbf{S})} \frac{1}{N_p} \sum_{p=1}^{N_p} \sum_{i=1}^{K_{\text{Up}}} \log_2[1 + \text{SINR}_{\text{Up}}^{(\mathbf{S})}(p, j)]. \quad (15)$$

The uplink SINR for SoftNull, $\text{SINR}_{\text{Up}}^{(\text{SoftNull})}$ will be much smaller than $\text{SINR}_{\text{Up}}^{(\text{HD})}$ because of residual self-interference. Similarly $\text{SINR}_{\text{Down}}^{(\text{SoftNull})}$ will be smaller than $\text{SINR}_{\text{Down}}^{(\text{HD})}$ because of effective transmit antennas sacrificed by SoftNull for self-interference suppression. However, SoftNull can still outperform half-duplex since SoftNull due to the multiplexing gain of operating uplink and downlink concurrently, captured by setting $\alpha_{\text{Down}}^{(\text{SoftNull})} = \alpha_{\text{Up}}^{(\text{SoftNull})} = 1$ whereas $\alpha_{\text{Down}}^{(\text{HD})} = \alpha_{\text{Up}}^{(\text{HD})} = \frac{1}{2}$.

We note that Equations 14 and 15 are based on two assumptions: (1) that optimal channel codes are used, such that the ideal Shannon rate is achieved, and (2) that the channel codes can extend over multiple channel realizations. Although these assumptions may not hold in practice, they provide a fair way of comparing system performance among different schemes, without the results being particular to a certain coding scheme and modulation rate. The sum rates are simply the sum of the uplink and downlink rates explained above.

APPENDIX C EXPLANATION AND EXAMPLE OF DYNAMIC RANGE LIMITATION

In wireless communication theory, one often assumes a fixed noise floor, based on thermal noise. For example, one could assume a thermal noise floor of -90 dBm, and since

the noise floor is fixed, if the received signal power is -80 dBm, we would say the SNR is 10 dB, and if the received signal power is -70 dBm we would say the SNR is 20 dB. However, the model of a fixed noise floor, is only true for low SNR. In practice there are noise sources whose power *scales proportionally to the received power level*, such as A/D quantization noise, oscillator noise, amplifier non-linearities, etc [17]. Collectively, such impairments are called *dynamic range limitation*. One way to effectively model dynamic range limitation is by adding a gaussian noise term whose power level is proportional to the received signal power [16], [50], [51], i.e., having a “dynamic range noise floor” that is fixed *relative* to the received power. For example, a common assumption is that the dynamic range floor is 30 dB below the received power level. Consider our example of a -90 dBm thermal noise floor, and a dynamic range noise floor which is 40 dB below the received power level. This means that for received power levels less than -50 dBm, dynamic range limitations have little effect, because thermal noise is the dominant noise source. But if the received signal power level is more than -50 dBm, the dynamic range floor becomes the dominant noise source, limiting the SNR to be no more than 40 dB, even for very high received signal power.

Now let us consider the impact of limited dynamic range on full-duplex in a small cell, for the same example given above where the dynamic range noise floor which is 40 dB below the received power level. Let us assume the base station transmits at 0 dBm and the uplink users also transmit with power 0 dBm, and let us further assume that there is 80 dB path loss between the base station and the user(s). In this example the power of the desired uplink signal received at the base station will be -80 dBm. Now let’s say that the self-interference reduction provided by the base station is only 20 dB (this is about what could be achieved by merely separating the receive and transmit antennas by a foot). Then the self-interference received by the base station is at -20 dBm. The self-interference dominates the received power level, and raises the dynamic range noise floor to -20 dBm $- 40$ dB = -60 dBm. Even if the residual self-interference is digitally cancelled perfectly, the SNR for the desired uplink signal is

$$-80 \text{ dBm} - (-60 \text{ dBm}) = -20 \text{ dB},$$

which is much too low of an SNR to support a useful communication link. However if the amount of self-interference reduction is improved from 20 dB to 50 dB, then the received power level is reduced to -50 dBm, and thus the dynamic range noise floor is reduced to -50 dBm $- 40$ dB = -90 dBm. Now, after the the residual self-interference is cancelled digitally, the the SNR of the uplink signal is

$$-80 \text{ dBm} - (-90 \text{ dBm}) = 10 \text{ dB},$$

a positive SNR, which means that a useful communication link can be operated on the uplink in conjunction with downlink transmission.

We remark that the amount of self-interference suppression required to achieve a positive uplink SNR is dependent on the *path loss* between the base station and the uplink user(s). If the path loss had been 100 dB instead of 80 dB, then 70 dB

of suppression, rather than 50 dB, would be required to achieve an uplink SNR of 10 dB. That is, the amount of required self-interference suppression is proportional to the amount of path loss between base station and users.

REFERENCES

- [1] D. W. Bliss, P. A. Parker, and A. R. Margetts, “Simultaneous transmission and reception for improved wireless network performance,” in *Proc. IEEE Workshop Statist. Signal Process.*, Aug. 2007, pp. 478–482.
- [2] J. I. Choi, M. Jain, K. Srinivasan, P. Levis, and S. Katti, “Achieving single channel, full duplex wireless communication,” in *Proc. MobiCom*, New York, NY, USA, 2010, pp. 1–12.
- [3] M. Duarte and A. Sabharwal, “Full-duplex wireless communications using off-the-shelf radios: Feasibility and first results,” in *Proc. Asilomar Conf. Signals Syst.*, Nov. 2010, pp. 1558–1562.
- [4] M. Duarte, C. Dick, and A. Sabharwal, “Experiment-driven characterization of full-duplex wireless systems,” *IEEE Trans. Wireless Commun.*, vol. 11, no. 12, pp. 4296–4307, Dec. 2012.
- [5] M. Jain *et al.*, “Practical, real-time, full duplex wireless,” in *Proc. 17th Annu. Int. Conf. Mobile Comput. Netw.*, New York, NY, USA, 2011, pp. 301–312. [Online]. Available: <http://doi.acm.org/10.1145/2030613.2030647>
- [6] M. Duarte *et al.*, “Design and characterization of a full-duplex multi-antenna system for WiFi networks,” *IEEE Trans. Veh. Technol.*, vol. 63, no. 3, pp. 1160–1177, Mar. 2014.
- [7] D. Bharadia, E. McMillin, and S. Katti, “Full duplex radios,” in *Proc. ACM SIGCOMM Conf.*, Oct. 2013, pp. 375–386.
- [8] T. Riihonen, S. Werner, and R. Wichman, “Mitigation of loopback self-interference in full-duplex MIMO relays,” *IEEE Trans. Signal Process.*, vol. 59, no. 12, pp. 5983–5993, Dec. 2011.
- [9] A. Sahai, G. Patel, C. Dick, and A. Sabharwal, “On the impact of phase noise on active cancellation in wireless full-duplex,” *IEEE Trans. Veh. Technol.*, vol. 62, no. 9, pp. 4494–4510, Nov. 2013.
- [10] E. Aryafar, M. A. Khojastepour, K. Sundaresan, S. Rangarajan, and M. Chiang, “MIDU: Enabling MIMO full duplex,” in *Proc. 18th Annu. Int. Conf. Mobile Comput. Netw.*, New York, NY, USA, 2012, pp. 257–268. [Online]. Available: <http://doi.acm.org/10.1145/2348543>
- [11] D. Astely, E. Dahlman, G. Fodor, S. Parkvall, and J. Sachs, “LTE release 12 and beyond [accepted from open call],” *IEEE Commun. Mag.*, vol. 51, no. 7, pp. 154–160, Jul. 2013.
- [12] *Study on Elevation Beamforming/Full-Dimension MIMO for LTE*, document RP-141831, TSG RAN Meeting 66, 3GPP, Dec. 2014.
- [13] T. L. Marzetta, “Noncooperative cellular wireless with unlimited numbers of base station antennas,” *IEEE Trans. Wireless Commun.*, vol. 9, no. 11, pp. 3590–3600, Nov. 2010.
- [14] E. G. Larsson, O. Edfors, F. Tufvesson, and T. L. Marzetta, “Massive MIMO for next generation wireless systems,” *IEEE Commun. Mag.*, vol. 52, no. 2, pp. 186–195, Feb. 2014.
- [15] F. Rusek *et al.*, “Scaling up MIMO: Opportunities and challenges with very large arrays,” *IEEE Signal Process. Mag.*, vol. 30, no. 1, pp. 40–60, Jan. 2013.
- [16] B. Day, A. Margetts, D. Bliss, and P. Schniter, “Full-duplex bidirectional MIMO: Achievable rates under limited dynamic range,” *IEEE Trans. Signal Process.*, vol. 60, no. 7, pp. 3702–3713, Jul. 2012.
- [17] B. P. Day, A. R. Margetts, D. W. Bliss, and P. Schniter, “Full-duplex MIMO relaying: Achievable rates under limited dynamic range,” *IEEE J. Sel. Areas Commun.*, vol. 30, no. 8, pp. 1541–1553, Sep. 2012.
- [18] A. Sabharwal, P. Schniter, D. Guo, D. W. Bliss, S. Rangarajan, and R. Wichman, “In-band full-duplex wireless: Challenges and opportunities,” *IEEE J. Sel. Areas Commun.*, vol. 32, no. 9, pp. 1637–1652, Sep. 2014.
- [19] T. Riihonen, S. Werner, R. Wichman, and Z. Eduardo, “On the feasibility of full-duplex relaying in the presence of loop interference,” in *Proc. IEEE 10th Workshop Signal Process. Adv. Wireless Commun.*, Jun. 2009, pp. 275–279.
- [20] B. Yin, M. Wu, C. Studer, J. Cavallaro, and J. Lilleberg, “Full-duplex in large-scale wireless systems,” in *Proc. Asilomar Conf. Signals Syst. Comput.*, Nov. 2013, pp. 1623–1627.
- [21] H. Q. Ngo, H. Suraweera, M. Matthaiou, and E. Larsson, “Multipair massive MIMO full-duplex relaying with MRC/MRT processing,” in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2014, pp. 4807–4813.
- [22] B. Li, D. Zhu, and P. Liang, “Small cell in-band wireless backhaul in massive multiple-input multiple-output systems,” in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2015, pp. 1838–1844.

- [23] Z. Zhang, Z. Chen, M. Shen, and B. Xia, "On capacity of two-way massive MIMO full-duplex relay systems," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2015, pp. 4327–4332.
- [24] V. Chandrasekhar, J. G. Andrews, and A. Gatherer, "Femtocell networks: A survey," *IEEE Commun. Mag.*, vol. 46, no. 9, pp. 59–67, Sep. 2008.
- [25] A. Tang and X. Wang, "A-duplex: Medium access control for efficient coexistence between full duplex and half duplex communications," *IEEE Trans. Wireless Commun.*, vol. 14, no. 10, pp. 5871–5885, Oct. 2015. doi: 10.1109/TWC.2015.2443792.
- [26] J. Y. Kim, O. Mashayekhi, H. Qu, M. Kazadiieva, and P. Levis, "JANUS: A novel mac protocol for full duplex radio," Stanford Univ., Stanford, CA, USA, Tech. Rep., 2013.
- [27] 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 *Proc. Modeling Optim. Mobile Ad Hoc Wireless Netw. (WiOpt)*, May 2011, pp. 94–101.
- [28] Q. Gao, G. Chen, L. Liao, and Y. Hua, "Full-duplex cooperative transmission scheduling in fast-fading MIMO relaying wireless networks," in *Proc. Int. Conf. Comput., Netw. Commun. (ICNC)*, Feb. 2014, pp. 771–775.
- [29] C. Karakus and S. N. Diggavi, "Opportunistic scheduling for full-duplex uplink-downlink networks," *CoRR*, Jun. 2015. [Online]. Available: <http://arxiv.org/abs/1504.05898>
- [30] C. Shepard *et al.*, "Argos: Practical many-antenna base stations," in *Proc. 18th Annu. Int. Conf. Mobile Comput. Netw.*, New York, NY, USA, Oct. 2012, pp. 53–64. [Online]. Available: <http://doi.acm.org/10.1145/2348543.2348553>
- [31] H. Q. Ngo, E. G. Larsson, and T. L. Marzetta, "Massive MU-MIMO downlink TDD systems with linear precoding and downlink pilots," in *Proc. 51st Annu. Allerton Conf. Commun., Control, Comput. (Allerton)*, Oct. 2013, pp. 293–298.
- [32] K. T. Truong and R. W. Heath, "Effects of channel aging in massive MIMO systems," *J. Commun. Netw.*, vol. 15, no. 4, pp. 338–351, 2013.
- [33] J. Jose, A. Ashikhmin, T. L. Marzetta, and S. Vishwanath, "Pilot contamination and precoding in multi-cell TDD systems," *IEEE Trans. Wireless Commun.*, vol. 10, no. 8, pp. 2640–2651, Aug. 2011.
- [34] D. Bharadia and S. Katti, "Full duplex MIMO radios," in *Proc. 11th USENIX Conf. Netw. Syst. Design Implement.*, Berkeley, CA, USA, 2014, pp. 359–372. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2616448.2616482>
- [35] C. Shepard, H. Yu, and L. Zhong, "ArgosV2: A flexible many-antenna research platform," in *Proc. 19th Annu. Int. Conf. Mobile Comput. Netw.*, 2013, pp. 163–166.
- [36] *Wireless Open-Access Research Platform (WARP)*, accessed on May 7, 2015. [Online]. Available: <http://warp.rice.edu/>
- [37] C. A. Balanis, *Antenna Theory: Analysis and Design*, 3rd ed. Hoboken, NJ, USA: Wiley, 2005.
- [38] I. Bahl and P. Bhartia, *Microwave Solid State Circuit Design*. Hoboken, NJ, USA: Wiley, 1980.
- [39] (Jan. 2015). *Softnull Data Traces*. [Online]. Available: <https://www.dropbox.com/sh/qdzf6q5qofyp5vu/AAAr1yihm3yAyo>
- [40] E. Everett, A. Sahai, and A. Sabharwal, "Passive self-interference suppression for full-duplex infrastructure nodes," *IEEE Trans. Wireless Commun.*, vol. 13, no. 2, pp. 680–694, Jan. 2014.
- [41] A. S. Y. Poon, R. W. Brodersen, and D. N. C. Tse, "Degrees of freedom in multiple-antenna channels: A signal space approach," *IEEE Trans. Inf. Theory*, vol. 51, no. 2, pp. 523–536, Feb. 2005.
- [42] A. S. Y. Poon, D. N. C. Tse, and R. W. Brodersen, "Impact of scattering on the capacity, diversity, and propagation range of multiple-antenna channels," *IEEE Trans. Inf. Theory*, vol. 52, no. 3, pp. 1087–1100, Mar. 2006.
- [43] H. Bolcskei, M. Borgmann, and A. J. Paulraj, "Impact of the propagation environment on the performance of space-frequency coded MIMO-OFDM," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 3, pp. 427–439, Apr. 2003.
- [44] Q. H. Spencer, B. D. Jeffs, M. A. Jensen, and A. L. Swindlehurst, "Modeling the statistical time and angle of arrival characteristics of an indoor multipath channel," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 3, pp. 347–360, Mar. 2000.
- [45] T. S. Rappaport, *Wireless Communications: Principles and Practice*, Englewood Cliffs, NJ, USA: Prentice-Hall, 1996.
- [46] E. Everett, "Full-duplex wireless with large antenna arrays," Ph.D. dissertation, Dept. Elect. Comput. Eng., Rice Univ., Houston, TX, USA, 2015. [Online]. Available: <https://scholarship.rice.edu/handle/1911/87867>
- [47] D. N. C. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge, U.K., Cambridge Univ. Press, 2005.
- [48] M. Duarte, A. Sabharwal, C. Dick, and R. Rao, "Beamforming in MISO systems: Empirical results and EVM-based analysis," *IEEE Trans. Wireless Commun.*, vol. 9, no. 10, pp. 3214–3225, Oct. 2010.
- [49] E. Aryafar, N. Anand, T. Salonidis, and E. W. Knightly, "Design and experimental evaluation of multi-user beamforming in wireless LANs," in *Proc. 16th Annu. Int. Conf. Mobile Comput. Netw.*, New York, NY, USA, Feb. 2010, pp. 197–208. [Online]. Available: <http://doi.acm.org/10.1145/1859995.1860019>
- [50] R. M. Gray and T. G. Stockham, Jr., "Dithered quantizers," *IEEE Trans. Inf. Theory*, vol. 39, no. 3, pp. 805–812, May 1993.
- [51] W. Namgoong, "Modeling and analysis of nonlinearities and mismatches in AC-coupled direct-conversion receiver," *IEEE Trans. Wireless Commun.*, vol. 4, no. 1, pp. 163–173, Jan. 2005.



Evan Everett (S'06–M'13) received the B.Eng. degree in wireless engineering and the B.S. degree in physics from Auburn University in 2010, and the Ph.D. degree in electrical engineering from Rice University in 2016. He is currently a Research Scientist with Numerica Corporation, Fort Collins, CO, USA. 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. He is a Graduate Research Fellow of the National Science Foundation.



Clayton Shepard is currently finishing his Ph.D. degree in electrical and computer engineering at Rice University, where he also earned the B.S. and M.S. degrees. He is a member of the Rice Efficient Computing Group and a founding member of the Argos project that developed the first massive MIMO base station as part of his M.S. thesis. His research interests include mobile and wireless systems. He was an Intern and Visiting Researcher with Motorola's Advanced Research and Technology Lab in 2007 and 2008, respectively. He was an Intern with Bell Labs, Alcatel-Lucent, in 2011 and 2012. He also interned with Microsoft Research where he worked on MU-MIMO in TV white spaces in 2013 and 2014. He won the NDSEG fellowship award in 2011.



Lin Zhong (S'02–M'05–SM'12) received the B.S. and M.S. degrees from Tsinghua University in 1998 and 2000, respectively, and the Ph.D. degree from Princeton University in 2005. He is currently a Professor with the Department of Electrical and Computer Engineering, Rice University. He co-authored a paper, which was recognized as one of the 30 most influential papers in the first ten years of the Design, Automation, and Test in Europe Conference. His research interests include mobile and wireless systems. He received the Best Paper Awards from the ACM MobileHCI in 2007, the IEEE PerCom in 2009, the ACM MobiSys in 2011, 2013, and 2014, and the ACM ASPLOS in 2014. He is a recipient of the US National Science Foundation CAREER Award, the Duncan Award from Rice University, and the Rockstar Award from the ACM SIGMOBILE.



Ashutosh Sabharwal (S'91–M'99–SM'06–F'14) received the B.Tech. degree from IIT Delhi, New Delhi, India, in 1993, and the M.S. and Ph.D. degrees from The Ohio State University, Columbus, OH, USA, in 1995 and 1999, respectively. He is currently a Professor with the Department of Electrical and Computer Engineering, Rice University, Houston, TX, USA. His research interests include information theory, communication algorithms, and experiment-driven design of wireless networks. He received the 1998 Presidential Dissertation Fellowship Award.