

MIDU: Enabling MIMO Full Duplex

Ehsan Aryafar¹, Mohammad (Amir) Khojastepour², Karthikeyan Sundaresan²,
Sampath Rangarajan², and Mung Chiang¹

¹Princeton University, Princeton, NJ, USA ²NEC Laboratories America, Princeton, NJ, USA
¹{earyafar, chiang}@princeton.edu, ²{amir, karthiks, sampath}@nec-labs.com

ABSTRACT

Given that full duplex (FD) and MIMO both employ multiple antenna resources, an important question that arises is *how to make the choice between MIMO and FD?* We show that optimal performance requires a combination of both to be used. Hence, we present the design and implementation of MIDU, the first MIMO full duplex system for wireless networks. MIDU employs antenna cancellation with symmetric placement of transmit and receive antennas as its primary RF cancellation technique. We show that MIDU's design provides large amounts of self-interference cancellation with several key advantages: (i) It allows for two stages of additive antenna cancellation in tandem, to yield as high as 45 dB self-interference suppression; (ii) It can potentially eliminate the need for other forms of analog cancellation, thereby avoiding the need for variable attenuator and delays; (iii) It easily scales to MIMO systems, therefore enabling the coexistence of MIMO and full duplex.

We implemented MIDU on the WARP FPGA platform, and evaluated its performance against half duplex (HD)-MIMO. Our results reveal that, with the same number of RF chains, MIDU can potentially double the throughput achieved by half duplex MIMO in a single link; and provide median gains of at least 20% even in single cell scenarios, where full duplex encounters inter-client interference. Based on key insights from our results, we also highlight how to efficiently enable scheduling for a MIDU node.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Wireless Communication*

General Terms

Design, Experimentation, Measurement, Performance

Keywords

Full Duplex, Multi-User MIMO

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiCom'12, August 22–26, 2012, Istanbul, Turkey.

Copyright 2012 ACM 978-1-4503-1159-5/12/08 ...\$15.00.

1. INTRODUCTION

A full duplex wireless device is one that can transmit and receive at the same time in the same frequency band, and typically it requires at least one Tx and one Rx antenna. The key challenge in realizing such a device lies in the self-interference (SI) generated by the Tx antenna at the Rx antenna. As an example, consider a Wi-Fi signal with a transmit power of 20 dBm. A Tx–Rx antenna separation of about 6–8 inches results in a path loss of about 40 dBm (depending on channel characteristics), resulting in a self-interference of at least -20 dBm. With a noise floor around -93 dBm, one would further require a self-interference cancellation of at least 73 dB to be able to decode the desired received signal. While one can solely employ digital interference cancellation techniques, current ADC's do not have a resolution to pass a received signal that is 73 dB less than the noise floor. Hence, several practical full duplex (FD) systems [7, 8, 13] have been proposed that couple RF cancellation along with digital cancellation to achieve the desired level of SI suppression.

RF cancellation may include a combination of antenna cancellation and analog cancellation. In [7], antenna cancellation was achieved by placing two Tx antennas asymmetrically at ℓ and $\ell + \frac{\lambda}{2}$ distance from the Rx antenna. This allows the transmit signals to add π out of phase and hence cancel each other. On the other hand, analog cancellation involves the generation of a π phase shift internally, coupled with the estimation and compensation of the SI channel [8, 13]. This allows for π phase shifters with a better frequency response over a wide-band channel (e.g., BALUN in [13]) to be employed, in contrast to the strong dependence on frequency (λ) posed by the antenna cancellation in [7]. While the existing schemes employ at least two antennas, one can also envision FD with a single antenna [14], where a circulator is used to isolate the Tx and Rx signals. However, owing to the lack of both path loss attenuation and RF cancellation, the required level of SI cancellation is significantly higher and hence hard to realize.

Given that at least 2 antennas are needed for a practical implementation of FD at Wi-Fi Tx power, an alternate approach would be to employ the multiple antennas as in half duplex (HD)-MIMO to increase the link capacity. With next generation wireless devices (access points, base stations, etc.) equipped with multiple antennas (often more than two), it is both timely and important to understand *how to best employ the available spatial degrees of freedom (antennas)*. When comparing FD with MIMO, we consider two models that are of practical interest – *antenna conserved*

(AC) and *RF chain conserved* (RC). While the AC model allows for FD to be realized in legacy MIMO nodes ($\# \text{ Tx/Rx chains} = \# \text{ antennas}$), the RC model allows for nodes that are designed with full duplex in mind ($\# \text{ Tx/Rx chains} \leq \# \text{ antennas}$). We show that the relative merits of FD and MIMO differ significantly depending on the model considered.

However, irrespective of the model considered, it turns out that there are several scenarios where the best strategy is not MIMO or FD in isolation but a combination of both. This in turn brings us to the next question as to *how to effectively realize a joint MIMO+FD system?*

We observe that the existing antenna cancellation [7] and analog cancellation [13] approaches cannot be readily extended to MIMO systems. Although one might envision an extension of [7] using two Tx and one Rx antenna for every transmitted/received MIMO stream, this would require antennas to be placed such that each of the Tx pairs (for each stream) lead to SI signals that are 180° out of phase at *every* Rx antenna, which is hard to realize. On the other hand, analog cancellation in [13], when extended to N stream MIMO, requires one to estimate the self-interference channel between every pair of N^2 Tx–Rx antennas. This in turn results in the use of N^2 variable delays and attenuators, each of which has to be auto-tuned to track the N^2 SI channels, which is hardly practical, specifically in frequency selective wide-band channels.

In addressing the above challenges, we design and prototype *MIDU*¹ – a node with joint MIMO and (full) DUPlexing capabilities. MIDU employs antenna cancellation with *symmetric* placement of antennas as its primary RF cancellation technique. Specifically, for a single stream transmission, it employs antenna cancellation with a symmetric placement of either two Rx antennas and one Tx antenna (which we refer to as Rx antenna cancellation), or in a dual manner, two Tx antennas and one Rx antenna (which we refer to as Tx antenna cancellation). In addition to avoiding the shortcomings of asymmetric antenna cancellation [7], we show that such a design provides large amounts of SI cancellation with several key advantages:

- It allows for a two-level design, whereby Tx antenna cancellation is followed by Rx antenna cancellation. Using antenna cancellation theory, we show that such a design has the potential to double the antenna cancellation gains because of its additive nature.
- It potentially eliminates the need for any other form of analog cancellation, which seems limited in practice due to the need for variable attenuators and delay elements and its subsequent lack of scalability to MIMO systems.
- More importantly, it scales very easily to MIMO systems, thereby enabling the co-existence of MIMO with FD.

We have prototyped MIDU on WARP radios [2] for a 3×3 narrow-band MIMO + FD system. Briefly, our experiments reveal that each level of antenna cancellation can contribute to 25–30 dB of SI suppression, while their combination can yield up to 45 dB of suppression, thereby eliminating the dependence on analog cancellation. Coupled with 25–30 dB of SI suppression from digital cancellation techniques as in [11, 13, 16], MIDU can enable FD for Wi-Fi transmit powers. We also evaluate the performance of MIDU as a relay between two hidden nodes and as an access point (AP) serving

¹pronounced as “MyDu”

multiple single antenna clients. For each of these schemes, we compare the performance of MIDU against half duplex MIMO (or MU-MIMO). Our results reveal that MIDU increases relay capacity up to 80%, and even in the presence of uplink-downlink interference in single cells, increases the median capacity by at least 20%.

Finally, we highlight the implications of a MIDU node on the design of the MAC itself. While the design of a scheduler (MAC) for single cell MU-MIMO (including client selection and precoding) is challenging in its own right, extending it to a MIDU system takes it to another level – spatial degrees of freedom now have to be carefully split between downlink and uplink for FD, with client selection and precoding jointly addressed for MU-MIMO in each direction. Since the search space for the problem is very large, we use key insights from analysis and experimental data to identify regions of pronounced FD gains, and hence provide guidelines for an efficient scheduling strategy with a tractable search space.

In summary, we are driven in this paper by three questions concerning full duplex MIMO: *Is it feasible? Is it scalable? What are the tradeoffs with half duplex MIMO?* In addressing these questions, we make the following contributions:

- Designed MIDU, a system that can enable both MIMO and FD in tandem.
- Enabled two stages of antenna cancellation with additive gains to yield as much as 45 dB suppression, thereby alleviating the dependence on analog cancellation and SI channel estimation.
- Built a prototype of MIDU and showcased the additive benefits of its two level cancellation as well as joint operation of 3×3 MIMO + FD in practice.
- Provided guidelines for the design of an efficient MAC for single cells employing MIDU nodes.

The next section describes the related work. The joint merits of FD and MIMO are discussed in the following section. Thereafter, we present the design of MIDU, performance evaluation, guidelines on scheduling, discussions, and finally, the concluding remarks.

2. RELATED WORK

Full Duplex. Recent work has proposed several techniques to enable FD at Wi-Fi transmit powers [7, 8, 13]. Fig. 1(b) shows the various components of a FD system that contribute to SI suppression, namely: antenna cancellation, analog cancellation, and digital cancellation.

Antenna cancellation arranges Tx and Rx antennas in such a way that SI is reduced at the Rx antenna. While [7] employed asymmetric placement of Tx antennas to generate a π phase shift between the transmitted signals at the RX, [9] employed directional Tx antennas that place a null at the Rx antenna. For a MIMO capable FD node, antenna placement must cancel SI at all of the receiving antennas, which is hard to realize in the proposed schemes.

Analog cancellation requires knowledge of the SI channel to create a copy of the SI signal in the RF domain and cancel it before the signal is digitized. While [16] uses a noise canceler chip for the purpose to achieve about 20–25 dB of SI suppression, [13] employs a combination of BALUN (balanced to unbalanced) transformer (to generate a negative copy of transmit signal) and variable delay and attenuators

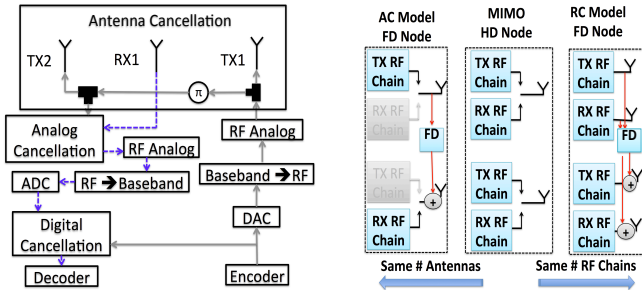


Figure 1: (a) FD Architecture, (b) Comparison Model.

(to track the SI channel), to yield a SI suppression of 25–30 dB in practice. In [8], a pseudo analog cancellation technique is introduced, in which an additional RF chain creates a canceling signal in RF from a digital estimate of SI in the base band, removing 35 dB of SI.

All these schemes require an estimation of SI channel between a Tx and Rx antenna, which becomes a scalability bottleneck for MIMO systems with FD.

Digital cancellation utilizes the digital samples of the transmitted signal in the digital domain and subtracts them from the received samples, removing up to 25 dB of SI [11, 12, 13]. However, the pseudo analog cancellation approach adopted in [8] limits the additional suppression from digital cancellation to only 4–5 dB. This restricts its total cancellation to less than 39 dB, limiting its applicability to small–medium range communications only.

While MIDU can also benefit from digital cancellation (as in [11, 12, 13]), in contrast to the above works, the large potential for SI suppression from MIDU’s antenna cancellation (albeit complementary to existing schemes), allows it to break away from the dependence on analog cancellation. This allows MIDU to easily scale to MIMO systems.

MIMO. In single user MIMO (SU-MIMO) systems (e.g., 802.11n [4] and BLAST [10]), the capacity of a point-to-point communication link is (theoretically) expected to scale with the number of antennas at the Tx and Rx (say N). However, in practice, the number of antennas at a base station or access point (N) is much more than those at the clients (n), limiting the performance of SU-MIMO to scale only with n . Multi-User MIMO (MU-MIMO) can be employed to overcome the limitation in such scenarios. Recent work [6, 17] has implemented MU-MIMO schemes, in which an AP can communicate with a number of clients simultaneously by utilizing the antennas that belong to a group of clients. As a baseline for comparison, we compare the performance of MIDU against such half duplex (HD)-MIMO schemes.

3. MIMO OR FD, OR BOTH?

To understand why we need a combination of MIMO and FD, we need to study the relative merits of FD and MIMO. While our primary focus is on performance, hardware complexity must also be taken into account for a fair comparison. We define two models that are of practical interest: *antenna conserved* (AC) and *RF chain conserved* (RC) models. The models are defined with respect to a legacy MIMO node as shown in Fig. 1 (b), where each antenna is associated with a pair of Tx and Rx RF chains, and are both important from different perspectives.

AC Model: Here, the node employing FD has the same

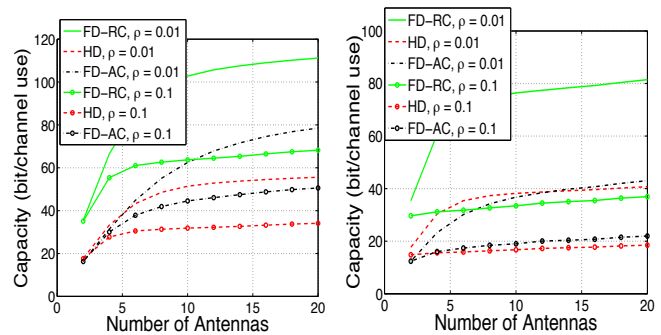


Figure 2: FD vs. MIMO Performance: (a) Perfect SI suppression, (b) SI suppression loss: 6 dB.

number of antennas (and RF chain pairs) as the legacy MIMO node and hence represents the case where FD is enabled directly on legacy MIMO nodes. Hence, when FD is enabled, depending on the split of antennas between downlink and uplink, a mix of Tx and Rx RF chains will be utilized in the FD mode. The total number of streams in FD (including both uplink and downlink) is the same as that in MIMO (either downlink or uplink). However, the full potential of FD is not leveraged in this case (half of the RF chains are not used as shown in Fig. 1(b)).

RC Model: Here, the node employing FD has more antennas than the number of RF chain pairs. Note that in HD-MIMO nodes, due to the half duplex nature, only half the number of RF chains (Tx or Rx) can be used in any transmission. However, this is not the case with FD, where all the RF chains can be effectively used as long as there are sufficient number of antennas available. For example, compared to a two antenna MIMO node supported by two pairs of Tx–Rx RF chains and capable of sending (or receiving) two streams, a corresponding FD node can use all the four RF chains through four antennas to send and receive two streams simultaneously. On the other hand, having the additional antennas on the legacy MIMO node would only contribute to diversity.

Note that the processing complexity of a transceiver lies predominantly in its RF chains and not its passive antennas. Hence, adding more antennas, while an issue for *form factor* and hence mobile devices, is not an obstacle for base stations and access points, which is where we expect FD to be predominantly employed. Therefore, when a node is designed taking FD into account, one can provision it with more passive antennas than the number of RF chains to leverage the potential of FD, which in turn is captured by the RC model.

3.1 FD vs MIMO Performance

Note that when a node has multiple antennas and employs more than two streams, FD by definition refers to a combination of FD and MIMO. This is because, while the available antennas can be split between downlink and uplink in FD, one would still need to employ MIMO within the available number of antennas in each direction to maximize the number of streams. Hence, the actual question that we are interested in understanding here is that, *given a set of antennas and RF chains, should one consider splitting the antennas between downlink and uplink through FD (with MIMO within each direction), or should one directly employ all the antennas solely towards MIMO in just one direction?*

To address this question, we contrast the performance of

FD-MIMO vs. HD-MIMO in a single link by comparing their respective capacities under both the AC and RC models. We also consider both perfect SI cancellation at the FD node, as well as when there is a remaining SI of 6 dB. We also incorporate the correlation between antenna elements on a node (ρ) that arises in practice and influences MIMO performance ($\rho = 1$ indicates perfect correlation). The results are presented in Fig. 2 as a function of the number of antennas in the HD-MIMO mode (corresponding to the same number of antennas in the AC model and half of the antennas in the RC model) at either ends of the link, as well as the antenna correlation factor and FD SI suppression. Two key observations can be made.

- In the AC model, even small antenna correlations ($\rho = 0.01$) saturates MIMO performance with increasing antennas. Hence, while FD does not have an advantage over MIMO theoretically in terms of number of stream transmissions, in practice with good SI suppression capability, it can deliver better performance even with a moderate number of antennas (about 4-6 in Fig. 2(a)). However, with imperfect SI suppression (loss of 6dB), such a transition happens at a higher number of antennas (12 in Fig. 2(b)) unless the antenna correlation is high.
- In the RC model, there is an advantage for FD as it has the potential to transmit twice as many streams as MIMO. This, coupled with MIMO's saturating performance, allows FD to deliver significant gains with a strong SI suppression capability (Fig. 2(a)). However, even for a moderate SI suppression capability (with a loss of 6dB), gains can be observed, even for a moderate number of antennas (two in Fig. 2(b)) at low antenna correlations.

In summary, we find that *irrespective of the model considered, the use of available antenna resources towards a combination of FD and MIMO is critical for optimal performance*. While the above comparison is with respect to a single point-to-point link, applicable to scenarios involving cellular backhaul links, mesh network links, relay links, etc., we will also discuss the relative impact in single cell point-to-multipoint scenarios in Section 5.4. We next detail how to realize such a joint MIMO and FD node.

4. DESIGN OF MIDU

MIDU employs antenna cancellation achieved through symmetric antenna placement (SAP) as its primary RF cancellation technique. We first outline the rationale behind our choice by identifying the advantages of SAP. Then, with the help of antenna cancellation theory, we show how to enable two levels of antenna cancellation (Tx and Rx antenna cancellation) with additive benefits. Finally, we show how it scales easily to incorporate MIMO, thereby completing the design of MIDU.

4.1 Symmetric Antenna Placement

Our antenna cancellation approach is based on a symmetric placement of antennas. Figure 3(a) illustrates our Rx antenna cancellation, where two Rx antennas are placed symmetrically at a distance ℓ from the Tx antenna. The signal received from one of the receive antennas is phase shifted internally using a fixed π phase shifter before being




Figure 3: (a) Receive Cancellation, (b) Transmit Cancellation.

combined with the other receive signal to help nullify the self-interference signal. Similar to Rx antenna cancellation, we can also have an analogous Tx antenna cancellation as shown in Fig. 3(b).

While the basic antenna configuration for cancellation is simple, we now highlight its significant potential to not only address the limitations of existing FD schemes, but also to allow for two levels of antenna cancellation and to leverage MIMO in tandem. Compared to the transmit antenna cancellation in [7], where the π phase shift was realized with asymmetric placement of Tx antennas (ℓ and $\ell + \frac{\lambda}{2}$), SAP has the following advantages:

- **Bandwidth Dependence:** Moving the π phase shift internally alleviates the bandwidth dependence (due to λ) of asymmetric antenna placement. Further, fixed π phase shifters have significantly better frequency responses over wide bandwidths compared to variable ones.
- **Tuning:** Since the received powers are similar, this avoids the need for tuning of the attenuation and phase of the self-interference signal, which are otherwise required to counteract the power difference due to asymmetric antenna placement.
- **Scalability:** Symmetric antenna placement allows for easy realization of several null points, therefore scaling to MIMO systems. Further, when compared with schemes such as [13] that do not consider antenna cancellation, SAP does not require an estimation of the SI channel between every pair of Tx and Rx antennas, which becomes a scalability issue for MIMO systems.

One limitation that was raised in [7] with respect to symmetric antenna placement, is its potential destructive impact on the far-field. However, the simulations used to highlight this observation used a free space path loss model. Note that while the SI channel can be modeled as free space, it is well known [18] that the far-field channels (indoors or outdoors) from the transmit antennas experience independent fading at any far-field receive point with sufficient separation (greater than λ which is 12.5 cm at 2.4 GHz) between the transmitting antennas (also validated in our experiments in Section 5.2). Hence, asymmetric antenna spacing does not provide an advantage over a symmetric placement with respect to impact on far-field. An analogous argument holds for signals received from far-field.

4.2 Understanding Antenna cancellation

To leverage antenna cancellation effectively, it is important to understand the notion of *signal nulling*. A signal is said to be “nulled” when two copies of the signal add π out of phase to cancel each other, thereby pushing the received signal strength to or below the noise floor. Let us consider transmit antenna cancellation for explaining the concepts.




Figure 4: Loci of Null Points: (a) Phase Offset + Antenna Spacing, (b) Only Antenna Spacing.

There are two parameters affecting the nulling process: relative phase and amplitude of the transmitted signals at the receiver. The relative phase between the two signals could be further controlled either by directly introducing a phase offset (ϕ) to one of the signals and/or by varying the relative distance between the transmit antennas with respect to the receive antenna.

Let d_t be the distance between the two Tx antennas, with d_1 and d_2 denoting the distance of the two transmit antennas with respect to a receive point, respectively. First, we consider the set of *potential* receive null points, when there is a phase offset $\phi = \pi$. Whether these null points can be realized in turn depends on the relative amplitude of the signals as well, which is discussed subsequently. Now the set of potential null points in a two-dimensional plane can be defined as the locus of the points satisfying $|d_1 - d_2| = k\lambda$, for some integer k , and includes the following (see Fig. 4(a)).

- The Perpendicular Bisector (PB) of the line joining the transmit antennas (i.e. $d_1 = d_2 = \frac{d_t}{2}$).
- A set of hyperbolas with the transmit antennas as the focal points. Each hyperbola intersects the line connecting the two transmit antennas at points that are $\frac{k\lambda}{2}$, $k \in \mathbb{Z}_+, \geq 1$ from the mid-point towards either one of the transmit antennas.
- If $d_t = m\lambda$, in addition to the above points, all points on the line passing through the two transmit antennas, besides those lying in between them, also contribute to the set of potential null points.

To understand scenarios where the relative phase is controlled only with just the help of antenna spacing (i.e., phase offset, $\phi = 0$), we note that the locus of the potential null points is now defined as those satisfying $|d_1 - d_2| = \frac{(2k+1)\lambda}{2}$ and consist of (see Fig. 4(b)),

- A set of hyperbolas with the transmit antennas as the focal points. Each hyperbola intersects the line connecting the two transmit antennas at points that are $\frac{(2k+1)\lambda}{4}$, $k \in \mathbb{Z}_+, \geq 1$, from the mid-point towards either one of the transmit antennas.
- If $d_t = \frac{(2m+1)\lambda}{2}$, in addition to the above points, all points on the line passing through the two transmit antennas, besides those lying in between them, also contribute to the set of potential null points.

Now for a potential null point to be realized, the two transmit signals must arrive at the receive point with *equal* amplitude but π out of phase. Due to symmetry, this can be easily achieved on the PB with an equal transmit power from the two transmit antennas. Hence, all null points on

the PB are realizable. However, for a null point on a hyperbola, it is easy to see that different transmit powers will be required from the two transmit antennas. Further, this will vary from one point to another on the same hyperbola as well as across hyperbolas. Hence, for fixed (and potentially different) transmit powers from the two transmit antennas, at most two null points on each hyperbola may be realizable. Note that we do not have null points on the PB when $\phi \neq \pi$. Given that the null points on the hyperbolas are hard to realize, this limits the applicability of asymmetric antenna spacing based approaches (e.g., [7]) to two level antenna cancellation. This limitation is compounded in the case of MIMO. This important property of realizing null points on the PB, when transmit signals are phase shifted by π , is leveraged for two purposes: (1) extend the transmit antenna cancellation to a two-level transmit and receive antenna cancellation scheme, and (2) to realize FD communication together with MIMO.

4.3 Two Level Antenna cancellation

Given that the above properties of transmit antenna cancellation analogously apply to receive cancellation as well, we can easily extend our proposed scheme to employ two stages (transmit and receive) of antenna cancellation in tandem. In the first stage two transmit antennas transmit at equal power and π out of phase signals that destructively interfere at any point on the PB of the transmit antennas. Now, place two Rx antennas symmetrically on the PB of the transmit antennas as shown in Fig. 5(a), such that the Tx and Rx sets of antennas are on each other's PB. While the transmit signals add destructively at each Rx antenna, the signals received from the two Rx antennas are further combined 180 degrees out of phase to provide the second level of antenna cancellation. Although four antennas are employed to achieve two levels of antenna cancellation, the number of RF chains used is still only two (one for forward and another for reverse streams).

The isolation (in dB) achieved by these two stages of cancellation are additive in theory, although in practice the cancellations might not be perfectly additive. In fact under ideal conditions even a one stage cancellation should provide a perfect null. However, gain imbalance or a slight phase offset between the signals may prevent one from achieving a perfect null, wherein a residue of the self-interference signal remains. We can now establish the following property.

PROPERTY 1. *Under small gain imbalance and/or phase offset (from imprecise antenna placement or imperfect RF devices) between the transmit and receive cancellation paths, the self-interference cancellation provided by two levels of antenna cancellation are additive (in dB scale).*

For details refer to [3].

4.4 Scaling to MIMO systems

Realizing null points on a straight line is critical because it facilitates the design of MIMO transmit and receive antenna arrays. Hence, our proposed two-level antenna cancellation solution based on phase offset can be readily extended to MIMO systems by using standard linear antenna array configurations. In particular, to generate an $N \times M$ MIMO + FD system, we start by placing two sets of antennas (N transmit and M receive) on two perpendicular axis to allow for $N \times M$ MIMO (in each direction of FD) as shown in Fig.




Figure 5: (a) 2-Level cancellation, (b) MIMO.

5(b) where $(N, M) = (3, 3)$. Then, to enable this $N \times M$ system with full duplex, we use an equal number of transmit (N) and receive (M) canceling antennas and place them in a symmetric position on the opposite side of their respective axis. The MIMO transmit streams from the N transmit and their respective canceling antennas will add out of phase at each of the receive antennas in the first stage of cancellation. The composite received signals at each of the M receive antennas are then further combined out of phase with their respective canceling antennas to provide the second level of cancellation. It is worth pointing out that only such symmetric antenna configurations can be extended to generic MIMO systems without the need for variable attenuators and delay elements. Again note that, while $2(N + M)$ antennas are employed for achieving two levels of antenna cancellation with FD, the total number of RF chains required is only $N + M$, which is the minimum required to enable $N \times M$ HD-MIMO communication in either direction. Fig. 5(b) shows the antenna structure for a 3×3 MIMO + FD node.

To summarize, the key aspect in MIDU's design is its symmetric antenna placement for antenna cancellation that not only provides two levels of additive RF cancellation, thereby alleviating the dependence on analog cancellation, but also realizes it through a structure that scales seamlessly to MIMO systems.

5. IMPLEMENTATION AND EXPERIMENTAL EVALUATION

In this section we provide experimental results on the performance of MIDU. We first perform extensive channel measurements to verify some of the assumptions in the design of MIDU. Next, we evaluate the cancellation design by measuring the amount of self-interference cancellation for each of the two levels of antenna cancellation and for the case when the two levels are combined. Finally, we compare the performance of MIDU to an equivalent half duplex MIMO system.

5.1 Measurement Setup

We performed channel measurement experiments using the WARP Vertex-4 FPGA boards. Our implementation is based on the WARPLab framework [2]. In this framework, all WARP boards are connected to a host PC through an Ethernet switch. The host PC is responsible for baseband PHY signal processing, while WARP boards act as RF front-ends to send/receive packets over the air. Since the baseband processing is performed in MATLAB, there is a 50 ms delay between consecutive transmissions.

We construct an access point (AP) in an indoor open space (low multi-path) environment. There is a distance of at least 20 m between our setup and the nearest wall. We mount our antennas on a wooden board, which itself is at the height of approximately 1.8 m from the earth surface. The carrier frequency was centered at 2.484 GHz (2.4 GHz, channel 14), and we implemented a single subcarrier narrowband system with a bandwidth of 625 KHz. We used 2.4 GHz 3dBi dipole antennas. We discuss the implications of rich scattering environment, wider bandwidths and antenna type in Section 7. In our experiments we construct 100-bit packets, and use the BPSK modulation scheme. The beginning of each packet is augmented with a preamble header. Each preamble is composed of training samples and an appropriate pilot tone for channel estimation between a transmitter and the receiver.

5.2 MIDU Feasibility

Three assumptions are critical in order for our full duplex system to work effectively: (i) Channel between a Tx and Rx antenna is only dependent on the distance between the two in the near field; (ii) Channel symmetry is stable over time; (iii) Symmetric antenna placement does not create nulls in the far-field. We now proceed to verify each of these assumptions.

The channel between a Tx and Rx antenna is only dependent on the distance between the two in the near field. The first assumption implies that an Rx antenna with similar distances from two Tx antennas would observe channels with equal amplitude and phase. This allows us to realize several null points in the BP of the two Tx antennas through the use of a π phase shifter, and thus enable MIMO.

We mount two antennas on our AP. The two antennas are attached to a rod with adjustable distance. We fix the two Tx and Rx antennas at a distance of 40 cm and send 40 back-to-back packets and measure the average channel (amplitude and phase) across all the packets.

Once the channel measurements are taken, we rotate the rod (with Tx and Rx antennas still attached to it) on a circle around the Rx antenna, and perform the same experiment at 7 other locations. These locations are shown as 8 dark squares on the circle around the Rx as shown in the picture of Fig. 6(a). Once these measurements are taken, we perform the same experiment on 3 more circles around the Rx by increasing the distance between the Tx and Rx antennas.

For each circle, we calculate the average and standard deviation of the *average channels* measured over the 8 selected points. Fig. 6(b) depicts the corresponding results. For all the circles, standard deviation of gain is less than 1%. Similarly, the standard deviation of phase is around 2° out of 360° . The results in Table 6(b) confirm that in an open-space environment, the near field channel is dependent only on the distance between the transmitter and receiver.

Symmetry stability over time. We now measure the amount of variation in the self-interference channel between two Tx and two Rx antennas over time. While in the design of MIDU two TX(RX) antenna pairs are connected to only one TX(RF) chain, such a setup would not allow us to measure all of the four self-interference channels simultaneously. Therefore, for benchmarking, we perform an experiment in which we connect each antenna to a separate RF chain. We next construct a preamble header in which the training samples are followed by two non-interfering pilot tones for cor-




Figure 6: (a) h -dist measurement setup, (b) h -dist variation results, (c) stability setup, (d) stability results.

rect channel estimation between all of the Tx and Rx pairs. Fig. 6(c) shows our experiment’s setup. We next store the packet with the modified preamble in the two Tx RF chains, and send the same packet continuously for a duration of ten minutes. For each packet transmission, we measure the four self-interference channels ($h_i, i = 1...4$).

Fig. 6(d) presents the stability results among the SI channels by measuring the variation of each SI channel (h_i) and the ratio between them. Fig. 6(d) shows that each SI channel observes small variations (less than 1% for both phase and gain). This is intuitive as the two antennas are fixed at a close distance. Fig. 6(d) further reveals that the variation in phase difference between two SI channels is less than 0.3%, implying that when the two transmissions are synchronized, the channel observed is even more stable.

Impact of symmetric antenna placement on the far-field clients. We have performed experiments to verify that symmetric antenna placement at the Rx or Tx would not cause severe reduction in SNR at far-field locations. We observed similar results with Rx cancellation, thus here we only present the results with Tx cancellation. We perform an experiment in which an Rx client moves in two different circles around the AP with a distance of 10m and 3m from the AP respectively. In each circle, the client performs SNR measurements at 10 different equally spaced locations. In each location, the AP first transmits ten packets using a single antenna. Next, we change the antenna configuration by pairing the original antenna with a canceling transmit antenna (consisting of a transmit antenna and a 180° phase shifter). In this setup, the canceling antenna is at a distance of 40 cm from the first antenna, and the two Tx pairs are able to cancel 25 dB of SI for a carefully placed Rx antenna on the AP (discussed more in the next section). We repeat the experiment by sending ten back-to-back packets and measuring the SNR at the client’s location. The experiment is repeated for all the 20 client locations.

Fig. 7(d) compares the received SNR for the two antenna configuration schemes. It shows that the achieved SNR can be up to 4 dB lower than the single antenna scheme. However, Fig. 7(d) also shows that Tx cancellation’s achieved SNR can be up to 4 dB higher than the single antenna scheme. Note that Tx cancellation uses two antennas for transmission. If the second antenna used by Tx cancellation has a higher gain than the first antenna, the combined effect can even increase the SNR. Fig. 7(d) confirms that there is no particular correspondence to increase/decrease in SNR, as one would expect from far-field fading [18].

5.3 Self-Interference Cancellation Evaluation

We now investigate MIDU’s SI cancellation performance for Tx cancellation and Rx cancellation separately, as well

as when the two levels of cancellation are combined. Finally, we measure the amount of cancellation when more than one Tx pair is active.

Cancellation over wire. Prior to our over-the-air measurements, we first verify if a commercial phase shifter can cancel two equal signals over wires. For the transmit signal, we used the Agilent MXG vector signal generator that creates a sinusoid signal at the 2.4 GHz center frequency. The signal was passed through a splitter; one of the resulting copies was next passed through our phase shifter; and then the two paths were combined using a combiner. The output of the combiner is next connected to an Agilent CSA spectrum analyzer to observe the received signal power. In this setup, we observed that for a 0 dBm transmitted signal power, one level of cancellation was able to achieve 35 dB of cancellation. We next employed two phase shifters in our system (one on each path), with one set to 0° and the other set to 180° of phase shift. With this setup we were able to cancel 90 dB of self-interference. *This implies that using a phase shifter on each path is necessary to maintain the correct symmetry and account for phase shifter induced insertion loss and delay.*

Wireless experiment setup. We use the antenna configuration setup depicted in Fig. 5(b). There are 6 transmit antennas symmetrically placed on a straight line, and 6 Rx antennas on the bisector perpendicular. The distance between $T_1T'_1$ is set to 40 cm, and the same distance is set between $R_1R'_1$. The rest of the antennas are placed at 20 cm distance from the adjacent antennas. Each antenna pair is further connected to two phase shifters, with phase shifter values set to 0° and 180°, respectively. The 6 Tx antennas and the 6 Rx antennas are connected to 3 Tx RF chains and 3 Rx RF chains, respectively. We use the WARPLab experimental setup of Section 5. However, now we set the transmitting RF chains in a continuous transmission mode for a duration of 10 sec, in order to correctly measure the received signal power on the spectrum analyzer. In continuous transmission mode, a transmitting WARP board continuously transmits the same packet without any delay between back to back packets.

Transmit Cancellation. In this experiment, we set the transmit antenna pair fixed as $T_1T'_1$, and measure the amount of transmit self-interference cancellation at each of the 6 receive antennas as a function of transmit power. Specifically, for each Rx antenna we first measure the Rx signal strength when only T_1 is active. Next, we measure the received signal strength on the spectrum analyzer when both T_1 and T'_1 are active simultaneously. We repeat this set of measurements by varying the transmit power from -5 to 15 dBm, and for each of the 6 Rx antennas.

Fig. 7(a) shows the amount of Tx cancellation for each of




Figure 7: (a) Tx cancellation, (b) Rx cancellation, (c) 2-Level and MIMO cancellation; (d) Impact in far-field.

the Rx antennas as a function of transmit power. We observe that a one level Tx cancellation results in 22–30 dB of self-interference cancellation for each of the Rx antennas. Similar Tx cancellation results were also observed in [7]. We further observe that the cancellation results remain relatively flat across different transmit powers, suggesting that the amount of cancellation does not depend on the Tx power but on the precise placement of the Tx and Rx antennas.

Receive Cancellation. In this experiment, we fix the receiving antenna pair as $R_1R'_1$, and measure the amount of receive SI cancellation for each of the transmitting antennas. Since we observed in Fig. 7(a) that the amount of cancellation does not depend on the transmit power, we use a fixed transmit power of 15 dBm. We first measure the received signal power when only R_1 is connected to the spectrum analyzer. Next we measure the received signal power when R_1 and R'_1 are combined. We calculate the amount of Rx cancellation by subtracting the two values. The experiment is repeated for each of the Tx antennas.

Fig. 7(b) shows the measured Rx cancellation values. Similar to Tx cancellation results of Fig. 7(a), we observe that one level Rx cancellation can provide cancellation values between 20 to 30 dB. We also observe that although cancellation measurements are performed at different time instants, symmetric transmitting antennas observe almost similar cancellation values.

2-Level and MIMO cancellation. We now perform experiments to validate if the two levels of cancellations are additive, and further measure the cancellation results when more than one Tx antenna pair is transmitting (i.e., MIMO). We measure the 2-Level cancellation results, by measuring the resulting amount of cancellation for each $T_iT'_iR_iR'_i$ combination. We measure MIMO cancellation by activating all transmit antennas while using $R_iR'_i$ for $i = 1 \dots 3$ (i.e., 3x1 MIMO). MIMO cancellation results in Fig. 7(c) are shown as $T_{all}R_iR'_i$.

Fig. 7(c) shows the 2-Level and MIMO cancellation results. The x-axis in Fig. 7(c) is the i variable, denoting the receiving antenna pair for both 2-Level and MIMO cancellation. Fig. 7(c) reveals that the 2-Level combination can provide up to 45 dB of cancellation. This in turn verifies that *Rx cancellation after Tx cancellation provides additive cancellation gains*. Fig. 7(c) also shows that when all three antenna pairs are active, the 2-Level cancellation gains reduce by up to 5 dB. Note that if the second active antenna pair makes interference levels equal to the noise power at $R_1R'_1$, the resulting observed cancellation at $R_1R'_1$ would decrease by 3 dB. In order for $T_2T'_2$'s addition to have no impact on $R_1R'_1$, the resulting interference should be far lower than the noise power. Further, note that while one additional trans-

mit pair can cause 3–4 dB of additional interference, with 2 additional transmit pairs the observed reduction is 5 dB. This implies that each additional transmit pair would only slightly reduce the amount of self-interference cancellation.

5.4 Comparing MIDU with MIMO

In this section we compare the performance of HD-MIMO and MIDU systems in a relay architecture, and a point-to-multipoint (PtMP) single-cell architecture. In the relay setup, all nodes are MIMO capable. In the point-to-multipoint architecture, the AP has multiple antennas, whereas the clients each have a single antenna. We first consider the RC model in which both MIDU and HD-MIMO have the same number of TX/Rx RF chains, before turning to the AC model in Section 5.4.2.

5.4.1 Relay Architecture

We first consider a setup in which a HD-MIMO node (N_2) acts as a relay between two hidden HD-MIMO nodes (N_1 and N_3). Fig. 8(a) depicts our experiment setup. We next consider the case in which the relay node (AP) is a MIDU node. Note that when a MIDU node is considered as a relay, N_1 and N_3 can be active simultaneously (with one transmitting and the other node receiving). Thus, the architecture would be similar to a Point-to-Point (PtP) MIDU link with zero self-interference at the second node.

Performance Metric. We use the signal to noise plus interference ratio (SINR) observed by each receiving antenna, or the corresponding Shannon capacity (C) as our performance metric. For a PtP MIMO link with N Tx and N Rx antennas, we measure the SINR for each receive antenna and calculate the corresponding capacity. We calculate the link capacity as the summation of per-antenna capacities.

For the HD-MIMO relay architecture, we assume a TDMA scheme with an equal amount of time for each of the two links of Fig. 8(a). Thus, the aggregate capacity of the HD-MIMO system is equal to $\frac{C_{N_1N_2}^{HD-MIMO}}{2} + \frac{C_{N_2N_3}^{HD-MIMO}}{2}$. When a MIDU node is deployed as a relay, the two links can be active at the same time. Thus, the aggregate capacity is equal to $C_{N_1N_2}^{MIDU} + C_{N_2N_3}^{MIDU}$.

Note that the overall end to end throughput of a system is dependent on the specific MAC protocol implementation (including rate adaptation, etc.) and is an active research area. Shannon capacity is a measure of physical layer capacity and is used here to provide an upper bound on the throughput that would be achieved by any MAC protocol.

HD-MIMO Implementation In our implementation, *all the physical layer processing is done at the relay node* (N_2). Nodes N_1 and N_3 are located such that they have a strong SNR at the relay station (AP), however, they are




Figure 8: (a) Relay architecture, (b) Capacity.

hidden from each other. Each of the nodes in the relay architecture of Fig. 8(a) has 3 Tx RF chains and 3 Rx RF chains.

During $N_1 \rightarrow N_2$ transmission, each independent data stream is transmitted from a separate antenna at N_1 . The relay node first obtains the channel information between its three Rx antennas and the transmitting antennas at N_1 . It then employs a standard technique termed *zero force (ZF)* [18] filtering to separate the received streams. During $N_2 \rightarrow N_3$ transmission, N_3 receives interference free streams on separate receive antennas. Similarly, the relay node first obtains the channel information between its three antennas and the receive antennas at N_3 . It then uses ZF beamforming to transmit independent streams to each of N_3 's receiving antennas. In our ZF implementation, equal power is assigned to each ZF weight vector (for details of ZF filter/beamformer refer to [18]).

$N_2 \rightarrow N_3$ SINR Measurement. We use the received signal strength (RSS, in dBm) value reported by the radio boards for our SINR measurements. We take the following approach to measure the SINR for a receiving antenna k at N_3 . The MIMO transmitter (N_2) first performs ZF beamforming by maintaining the power associated to k , and setting the rest of the powers to zero. The measured RSS value would then correspond to the signal power. We next re-run the experiment by setting the power associated to k as zero and maintaining the powers associated to the other receive antennas. The measured RSS value would then correspond to the noise and inter-stream interference. By subtracting the two RSS values, we measure the SINR. We take 10 such SINR measurements and report the average value for each data point.

$N_1 \rightarrow N_2$ SINR Measurement. During N_1 to N_2 transmission, the AP is responsible for separating the data streams in the base band. In order to measure the SINR for a stream k , the MIMO transmitter first only transmits stream k , and the receiver applies the ZF weights to obtain the received signal samples (\mathbf{y}_k). The base-band signal power (P_s) is then calculated as $E\{\mathbf{y}_k \mathbf{y}_k^* \}$ (* denotes the conjugate transpose). Next, the MIMO transmitter (N_1) sets the power allocated to stream k as zero while maintaining the power associated to the other streams, and the receiver uses the previous weights to obtain the interference samples (\mathbf{z}_k) caused by other streams. The base-band interference power (P_i) is then calculated as $E\{\mathbf{z}_k \mathbf{z}_k^* \}$. We report the $10 \times \log_{10} \frac{P_s}{P_i}$ as the SINR value.

MIDU Implementation. We change the antenna configuration at the AP to a 2-Level antenna cancellation scheme as described in Section 4. We keep nodes N_1 and N_2 at the same locations. Similar to HD-MIMO implementation, all of the physical layer processing is performed at the relay node.

However, unlike the HD-MIMO mode, the transmission and reception is performed simultaneously.

SINR Measurement. $N_2 \rightarrow N_3$ SINR Measurement in MIDU is similar to HD-MIMO. However, the imperfect self-interference cancellation can reduce the performance of $N_1 \rightarrow N_2$ transmission. In our implementation of MIDU on the WARP boards, we observed only 4-5 dB of SI remaining above the noise floor as reported by the WARP boards.

This remaining SI is due to the presence of multi-path components in our environment, which can be suppressed by employing conventional digital cancellation techniques [11, 12, 13]. However, given the remaining small margin for SI suppression with WARP, we do not consider it in our implementation. Thus our results would be a lower bound on MIDU's performance.

We measure the $N_1 \rightarrow N_2$ SINR similar to HD-MIMO SINR measurement. However, during noise measurement for stream k , we also activate the transmitting N_2 streams to measure the combined impact of SI and received inter-stream interference.

Evaluation. Fig. 8(b) depicts the aggregate capacity of HD-MIMO and MIDU as a function of the number of independent data streams. We first observe that MIMO capacity does not scale linearly as the number of data streams is increased. In fact, while 2 streams increase the aggregate capacity by 80% compared to having only 1 stream, 3 streams decrease the aggregate capacity by 7% compared to having only 2 streams. This is due to MIMO capacity saturation when all degrees of freedom are used as discussed in Section 3. When all degrees of freedom are used to enable spatial multiplexing, per-link SINR can potentially decrease, reducing the benefits of spatial multiplexing.

Fig. 8(b) also compares MIDU's performance to the HD-MIMO scheme. We observe similar capacity saturation trend due to the incorporation of MIMO. However, we observe that MIDU on average increases HD-MIMO's capacity by over 80%.

5.4.2 Single Cell Architecture

We now compare the relative gains of HD-MIMO and MIDU in a single cell scenario.

Setup. We consider a setup in which we deployed 6 client nodes scattered around the AP. Each of the clients has only a single antenna for transmission or reception. The clients are placed in a manner that they all have a strong link to the AP, while some of them can be hidden from each other.

The AP is equipped with multiple antennas and has 3 Tx and 3 Rx RF chains. In HD-MIMO mode, the AP uses three fixed antennas. In addition, the AP uses all of its antennas to transmit (receive) to (from) up to 3 clients. We denote an $M \times N$ transmission strategy by the AP, as a scheme in which the HD-MIMO AP transmits to M clients during downlink, and receives from N clients during uplink. For a given M and N , we select M out of the six clients as our downlink clients, and N among the remaining clients as our uplink clients, and measure the resulting capacity. We denote the resulting capacity by $C_{M \times N}^{HD-MIMO}$. For a given M and N , there are $\binom{6}{M} \times \binom{6-M}{N}$ different selections. We repeat the same experiment for all possible client selections, and denote the average capacity as $\bar{C}_{M \times N}^{HD-MIMO}$. Finally, we repeat the experiment for all selections of $M = 1, 2, 3$ and $N = 1, 2, 3$ leading to a total of 500 sub-topologies. Our capacity measurements are based on the measurement setup in




Figure 9: (a) CDF Overall, (b) $C_{M \times N}$ Comparison, (c) CDF with 3 downlink clients; (d) Capacity with maximum allowed M and N .

the previous subsection. However, instead of obtaining the channel matrix from multiple antennas at a single node, we obtain the channel information from multiple single antenna clients and use the ZF technique accordingly.

We next change the antenna configuration to the setup described in Fig. 5(b), and perform the same set of experiments with a MIDU AP. Note that in an $M \times N$ transmission scheme by MIDU, the M downlink clients and N uplink clients are active at the same time. We measure the resulting capacities based on our measurement setup in the previous section. However, since not all of the clients are hidden from each other, uplink transmission in MIDU can cause interference on the downlink clients. Thus, when measuring inter-stream interference for each of the downlink clients, we also activate all of the selected uplink transmissions to accurately measure the resulting interference at the downlink nodes. We denote MIDU's average $M \times N$ capacity as $\bar{C}_{M \times N}^{MIDU}$.

Evaluation. Fig. 9(a) plots the cdf of all the $\frac{\bar{C}_{M \times N}^{MIDU}}{\bar{C}_{M \times N}^{HD-MIMO}}$ values across all possible 500 sub-topologies. According to the results shown in Fig. 9(a), MIDU achieves a better performance compared to HD-MIMO in over 80% of all the sub-topologies. In addition, for almost 50% of all sub-topologies, MIDU gains over HD-MIMO by at least 20%.

Fig. 9(b) shows the average $\bar{C}_{M \times N}$ results for different M and N values. We first observe that MIDU provides larger gains compared to HD-MIMO when total number of uplink (UL) and downlink (DL) streams is small (less than 4). Note that with a small number of UL and DL clients, the total interference present in the system is lower and thus full duplex (FD) can provide significant gains.

In order to better understand the impact of UL \rightarrow DL interference problem in FD we also plot: (1) cdf of the $\frac{\bar{C}_{M \times N}^{MIDU}}{\bar{C}_{M \times N}^{HD-MIMO}}$ for $M = 3$ and varying number of uplink clients in Fig. 9(c), and (2) the maximum HD-MIMO and MIDU capacities over all $M \times N$ sub-topologies as a function of the maximum allowed M and N . Two key inferences can be made:

From Fig. 9(c) we observe that *in order to leverage the maximum gains from FD, the number of streams on DL and UL must be kept dis-proportionate (asymmetric)*. The reason is as follows. If the number of UL (DL) streams from (to) HD clients is increased for a fixed number of DL (UL) streams, then it increases the interference faced by the clients receiving DL streams, thereby limiting FD performance significantly. While this limitation is alleviated with FD clients, it cannot be avoided.

From Fig. 9(d) we observe that *FD yields diminishing gains as the number of streams is scaled in either direction*. Note that, when the number of streams is small, it is pos-

sible to pick DL and UL clients that do not interfere with each other, allowing the increased number of streams from FD to yield maximal gains. However, with more streams, the impact of interference dominates and the few additional streams in FD (due to asymmetry) tend to contribute less gain compared to the baseline (larger) number of streams in MIMO.

FD gains in perspective. While these inferences indicate the restricted nature of FD gains in single cell MU-MIMO (due to UL \rightarrow DL interference), in practice the total number of multi-user streams, in either DL or UL direction, is small compared to the number of antennas at an AP. This is due to many practical reasons such as feedback overhead, computational complexity, minimum SNR supported, etc. For example, Arraycomm [1] APs (base stations) only support up to 4 streams simultaneously, with 12 antennas at the AP. This implies that FD can provide significant gains under the RC (RF Chain Conserved) model in realistic PtMP architectures. The impact of interference between UL and DL on FD gains is more pronounced in AC (Antenna Conserved) model, making it less favorable for adoption in such single cell MU-MIMO scenarios.

To contrast these results with the single-link results in Section 5.4.1, FD has significant potential under both AC and RC models in single-link scenarios. These scenarios include back-haul, relay, mesh network links (owing to the lack of DL-UL interference), and single-cell SU-MIMO (which is still predominantly considered in cellular networks due to the complexity of scheduling and realization of MU-MIMO gains).

6. IMPLICATIONS FOR MAC DESIGN

Challenges: Our results in Fig. 9 clearly reveal the need for scheduling on a MIDU node to identify opportunities for FD gain. While the design of a scheduler (MAC) for single cell MU-MIMO (client selection and precoding) is challenging in its own right, extending it to a MIDU system takes it to another level - spatial degrees of freedom now have to be carefully split between downlink and uplink for FD, with client selection and precoding being jointly addressed for MU-MIMO in each direction. While the search space for the problem is very large, we can use the two insights from our evaluation in Section 5.4.2 to make the search space tractable.

Insights: To ensure that our inferences are not specific to the FD implementation in MIDU, we have also evaluated the relative performance of ideal FD-MIMO and HD-MIMO through extensive simulations in larger topologies and hundreds of channel realizations.

We perform simulations with 8 uplink and 8 downlink




Figure 10: Large-scale simulations of FD gains in a single-cell architecture: (a) FD gains with forced uplink users, (b) Maximum FD gains, (c) Number of uplink users that result in maximum FD gains.

clients, and with varying number of antennas at the AP. Since we assume an equal number of users in the uplink and downlink, the average throughput in the uplink and downlink are similar due to channel reciprocity. Therefore, we define the *ergodic DL capacity* [15] for a given number of antennas as the HD-MIMO capacity. Next, for a given selection of UL users, we calculate the FD capacity that is the sum of the ergodic downlink and uplink capacities while considering the impact of UL clients interference on the downlink clients. FD capacity is then averaged over all UL user combinations for a given total number of UL users and across all different channel realizations.

Fig. 10(a) depicts the corresponding throughput gain of FD to HD-MIMO. From Fig. 10(a), we observe that as the number of antennas and therefore the number of DL streams increases, the FD gains decrease. We further observe that, as the number of the uplink clients increases, the FD gains start to decrease.

We next characterize the number of uplink users that result in maximum FD gains and the corresponding gains across all different channel realizations. For a given number of antennas and for each channel realization, we calculate the maximum FD gain, as well as the number of uplink clients that result in the maximum FD gains. Fig. 10(b) depicts the average maximum FD gain as a function of number of uplink clients, while Fig. 10(c) depicts the number of uplink users that result in maximum FD gains. From these results, we observe that as the number of streams in DL direction is increased, only a small number of clients in the reverse direction can be scheduled simultaneously. In fact, the results in Fig. 10(c) show that with six downlink streams, in more than 90% of all channel realizations only one uplink client provides additional FD gains, and none of the realizations achieved FD gains with three or more uplink clients.

These simulation results re-emphasize our earlier inferences in Section 5.4.2. Hence, for maximal benefits in single-cell networks, where MU-MIMO operation is feasible, it is better to *operate FD in small, asymmetric configurations – where the maximum number of streams is kept small in either direction, while the active number of streams between DL and UL is disproportionate*.

Scheduling Strategy: Guided by these conclusions, one can devise the scheduler to look at asymmetric configurations for FD only when the number of streams is small. This significantly reduces the search space for client selection on DL and UL for FD, making it comparable to MU-MIMO scheduling. Thereafter, conventional MU-MIMO scheduling

Freq (GHz)	2.484	2.474	2.464	2.454
Cancellation (dB)	45	38	17	9

Table 1: Phase shifter’s frequency response impact on cancellation.

and precoding algorithms [5] can be directly employed in either direction to evaluate the relative benefits of FD-MIMO over HD-MIMO and select the better strategy. On the other hand, when the number of streams is large, the scheduler can directly employ MU-MIMO.

7. DISCUSSION

Wide Bandwidth. Our implementation of MIDU cancels 45 dB of SI over a 625 KHz bandwidth signal. However, most of the current and future technology standards use bandwidths much wider than that (e.g., up to 20 MHz LTE and 40 MHz 802.11n). Therefore, it is important for SI cancellation techniques to cancel SI over wide-band.

We perform an experiment in which we change the center frequency of the WARP boards and measure the resulting SI cancellation at different frequencies. The phase shifters are tuned such that they provide 45 dB of SI cancellation at 2.484 GHz. Table. 1 shows MIDU’s SI cancellation results. We observe that MIDU provides 37 dB of SI cancellation at 2.474 GHz, implying that MIDU can *still provide over 35 dB of SI cancellation for a 20 MHz bandwidth signal*. However, at 2.464 GHz, SI cancellation is only 17 dB, limiting the SI cancellation to less than 20 dB for a 40 MHz bandwidth signal. While channel symmetry holds across different frequency bands, a phase shifter’s ability to create a signal inverse is precise only for a small bandwidth signal. Hence, the performance of MIDU is limited by the manner in which a signal inverse is generated. Recent research [13] suggested that a transformer (termed BALUN) can provide π phase shift over large bandwidths. Therefore, one can potentially employ BALUNs in lieu of phase shifters in MIDU’s implementation to provide SI cancellation over wider bandwidths.

Impact of multi-path interference. MIDU is able to provide 45 dB of SI cancellation in an open-space indoor (less multi-path) environment and can potentially yield higher cancellation outdoors. However when measured in an indoor multi-path rich environment, we observed only 15 dB of SI cancellation. While MIDU can help with SI cancellation of the LoS component, channel symmetry may not hold for strong NLoS components. In such scenarios, if the additional level of suppression required is less than 25 dB, digital cancellation would suffice. However, if more cancella-




Figure 11: Example Tx and Rx antenna placement with orthogonal polarization.

tion is required, then estimation and compensation of the SI channel becomes necessary. One can employ a pseudo analog cancellation approach [8] for this purpose. While this approach by itself is not sufficient to provide the desired level of cancellation (as discussed in Section 2), it can be combined with MIDU to address multi-path without comprising on its scalability. However, other analog cancellation approaches (e.g., [13]) cannot be employed as they incur the limitation of scalability, and sensitivity of SI compensation for wide-band frequency-selective channels.

Antenna Type and Polarization. In our implementation of MIDU, all Tx and Rx antennas were placed such that they have similar polarization. Due to the symmetric structure of MIDU, one can envision a placement of antennas in which Tx and Rx antennas have orthogonal polarization, potentially increasing the amount of SI cancellation. Fig. 11(b) depicts an example of such antenna placement.

Further, as MIDU requires only similar radiation patterns by Tx or Rx antennas, it can easily work with an array of directional antennas (e.g., directional/sectorized antenna placement in 3G). As part of our future work, we are investigating these schemes in order to provide additional 20–30 dB SI cancellation, thereby potentially enabling FD-MIMO in pico and macro cellular towers.

8. CONCLUSION

We presented the design and implementation of MIDU, the first MIMO full duplex wireless system. We showed that MIDU’s design allows for 2 levels of antenna cancellation, eliminates the need for variable attenuators and delays, and, most importantly, easily scales to MIMO. We implemented a prototype of MIDU, and compared its performance to HD-MIMO. Our results demonstrated that MIDU has significant potential in both point-to-point and point-to-multipoint schemes. We also provided guidelines for scheduling in practical MIMO full duplex implementations.

9. ACKNOWLEDGEMENTS

We would like to thank our shepherd, Yunhao Liu, and the anonymous reviewers for their valuable feedback, which helped in improving the presentation of the paper. This work was in part supported by NSF grant CNS-1011962 and AFOSR grant FA9550-09-1-0643.

10. REFERENCES

- [1] ArrayComm. Available at: <http://www.arraycomm.com>.
- [2] Rice University WARP project. Available at: <http://warp.rice.edu>.
- [3] Technical Report. Available at: <http://www.nec-labs.com/~karthiks/midu.pdf>.

- [4] Antenna selection and RF processing for MIMO systems. *IEEE 802.11-04/0713r0*, 2004.
- [5] W. Ajib and D. Haccoun. An overview of scheduling algorithms in mimo-based fourth-generation wireless systems. *IEEE Network*, 19(5):43 – 48, Oct 2005.
- [6] E. Aryafar, N. Anand, T. Salonidis, and E. Knightly. Design and experimental evaluation of multi-user beamforming in wireless LANs. In *Proceedings of ACM MobiCom*, Sep 2010.
- [7] J. Choi, M. Jain, K. Srinivasan, P. Levis, and S. Katti. Achieving single channel, full duplex wireless communication. In *Proceedings of ACM MobiCom*, Sep 2010.
- [8] M. Duarte, C. Dick, and A. Sabharwal. Experiment-driven characterization of full-duplex wireless systems. 2011. Available at: http://warp.rice.edu/trac/wiki/TransWireless2011_FullDuplex.
- [9] E. Everett, M. Duarte, C. Dick, and A. Sabharwal. Empowering full-duplex wireless communication by exploiting directional diversity. In *Forty-Fifth Asilomar Conference on Signals, Systems, and Computers*, Nov 2011.
- [10] G.D. Golden, R.A Valenzuela, P.W. Wolniansky, and G.J. Foschini. V-BLAST: an architecture for realizing very high data rates over the rich-scattering wireless channel. In *Proceedings of URSI International Symposium on Signals, Systems, and Electronics*, Pisa, Italy, Sep 1998.
- [11] S. Gollakota and D. Katabi. ZigZag decoding: combatting hidden terminals in wireless networks. In *Proceedings of ACM SIGCOMM*, Aug 2008.
- [12] D. Halperin, T. Anderson, and D. Wetherall. Taking the sting out of carrier sense: interference cancellation for wireless LANs. In *Proceedings of ACM MobiCom*, Sep 2008.
- [13] M. Jain, T.M. Kim, D. Bharadia, S. Seth, K. Srinivasan, P. Levis, S. Katti, and P. Sinha. Practical, real-time, full duplex wireless. In *Proceedings of ACM MobiCom*, Sep 2011.
- [14] M. Knox. *Self-jamming Cancellation Networks for Full Duplex Communication*. PHD Thesis, Polytechnic University, 2008.
- [15] L. Li and A. J. Goldsmith. Capacity and optimal resource allocation for fading broadcast channels - part I: ergodic capacity. *IEEE Transactions on Information Theory*, 47(3):1083–1102, Mar 2001.
- [16] B. Radunovic, D. Gunawardena, P. Key, A. Proutiere, N. Singh, V. Balan, and G. Dejean. Rethinking indoor wireless mesh design: Low power low frequency, full-duplex. Technical Report, MSR-TR-2009-27, Microsoft Research.
- [17] K. Tan, H. Liu, J. Fang, W. Wang, J. Zhang, M. Chen, and G. Voelker. SAM: enabling practical spatial multiple access in wireless LAN. In *Proceedings of ACM MobiCom*, Beijing, China, Sep 2009.
- [18] D. Tse and P. Viswanath. *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.