Full Duplex MIMO Radios

Dinesh Bharadia Stanford University

Abstract

This paper presents the design and implementation of the first in-band full duplex WiFi-PHY based MIMO radios that practically achieve the theoretical doubling of throughput. Our design solves two fundamental challenges associated with MIMO full duplex: complexity and performance. Our design achieves full duplex with a cancellation design whose complexity scales linearly with the number of antennas, this complexity is close to the optimal possible. Further we also design novel digital estimation and cancellation algorithms that eliminate almost all interference and achieves the same performance as a single antenna full duplex SISO system, which is again the best possible performance. We prototype our design by building our own analog circuit boards and integrating them with a WiFi-PHY compatible standard WARP software radio implementation. We show experimentally that our design works robustly in noisy indoor environments, and provides close to the expected theoretical doubling of throughput in practice.

1 Introduction

Full duplex radios have garnered significant attention recently in academia and industry [14, 9, 19, 18, 13, 16, 15, 12, 22, 17, 21, 20]. Several efforts are now underway to include full duplex technology in future cellular 5G standards [1], as well as explore applications of the technology in current wireless infrastructure. However these efforts are hampered by the fact that there aren't viable and efficient full duplex designs that can work in conjunction with MIMO. Specifically, no current practical designs are known which can enable one to build a M antenna full duplex MIMO radio that can transmit and receive from all antennas at the same time and double the throughput. The best known prior MIMO full duplex system, MIDU [9] requires 4M antennas for building a full duplex M antenna MIMO radio, and even then fails to provide the needed self-interference cancellation for WiFi systems (20 MHz bandwidth) to achieve the expected doubling of throughput.

Recent work has however demonstrated that a single antenna (SISO) full duplex system is practically possible [11]. Specifically, it demonstrates the design and implementation of a cancellation system for a SISO system that completely cancels self-interference to the noise floor and consequently achieves the theoretical doubling Sachin Katti Stanford University

of throughput. A natural question therefore is why not just replicate the same design M times to build a MIMO M full duplex radio? After all, a MIMO radio can be conceptually and physically viewed as a collection of M single antenna SISO radios.

The challenge is cross-talk interference. When a full duplex MIMO radio transmits, the transmission from any one of the M antennas (interchangeably referred to as transceiver chains) propagates to the other antenna (chains) and causes a large amount of interference. For the sake of clarity, in this paper we will refer to the self-interference at a receive chain caused by a transmission from the TX-chain with which the receive chain shares an antenna as "self-talk", and the interference from a neighboring TX chain's transmission as "crosstalk". Since MIMO antennas are closely spaced due to size constraints, this cross-talk is extremely strong, almost 75-80dB stronger than the desired signal that is being received on that chain. Consequently, even if we have cancellation circuits and algorithms that cancel every chain's self-talk, there is an extremely strong crosstalk interference that can saturate the receive chain.

A naive solution is to introduce a separate copy of the cancellation circuit and DSP algorithm for each pair of chains that experiences cross-talk. If there are M antennas, then it would imply a total of M^2 circuits and DSP algorithms. In other words complexity grows quadratically with the number of antennas, which is untenable as MIMO systems go towards 4 to 8 antennas. Supporting 16 cancellation circuits and DSP implementations (for 4 antenna MIMO) on even a WiFi AP based form-factor is untenable (our analysis suggested that with the current SISO design we would need 400sq.cm of analog circuit area and a high-end Virtex FPGA that consumes 80W of power to accommodate the DSP computations). Complexity impacts more than space and power consumption, cancellation systems (both analog and digital) need to be tuned continuously to adapt to environmental changes. The time for tuning scales linearly with the complexity, hence it would take M^2 time longer to tune such a design's MIMO self-interference cancellation system. The best known prior algorithm for tuning [11] requires around a millisecond to tune, so we would need 16 ms to tune for a 4 antenna MIMO system which would be untenable even in a slowly changing environment like indoor WiFi (coherence times are on the order of tens of milliseconds), let alone mobile environments such as LTE.

A second problem is performance itself. The key metric is the residual interference left after cancellation at each receive chain, the residual directly translates to decrease in SNR for the desired received signal. As we will show in Sec. 3, even if one could accommodate a quadratic number of circuits and DSP cancellation implementations, the performance degrades linearly with the number of MIMO chains. In other words, the residual interference after cancellation at each receive chain increases linearly with *M*. This is due to the accumulation of the residual interference from all the cross-talk and self talk cancellation systems. Once again, as MIMO systems scale to support many antennas, this essentially limits the performance gains of full duplex.

This paper presents the design and implementation of a MIMO WiFi full duplex radio. Our M antenna full duplex MIMO radio uses each antenna for simultaneous transmit and receive, i.e., it uses the same number of antennas as a standard half duplex M-antenna MIMO radio unlike prior designs. The design uses slightly more than M cancellation circuits and DSP algorithms to cancel all the self and cross talks. In other words, complexity scales linearly with the number of chains, which is the best performance one could expect. Further, the performance does not degrade linearly with the number of MIMO chains, i.e., the residual interference is the same as the SISO design and does not increase linearly with the number of chains. We prototype our design and integrate it with the off-the-shelf WARP software radios [4] running a stock WiFi baseband and demonstrate experimentally that it achieves close to the theoretical doubling of throughput.

Our design solves the key challenge of efficiently and effectively achieving the MIMO full duplex using two major ideas as follows.

- First, a key insight is that MIMO chains are co-located, i.e., "they share a similar environment". Intuitively, the signals transmitted by two neighboring antennas (separated by a few cm) go through a similar set of reflectors and attenuations in the environment. Cancellation systems are essentially trying to model these distortions, so when we want to model cross-talk, we can reuse the work that has been done for modeling the chain's own self-talk interference. This results in a novel "cascaded" filter structure for cancellation that results in an overall design that has near-linear complexity scaling with the number of MIMO antennas.
- Second, the reason performance degrades linearly with the SISO replication based design is that each of the *M* independent cancellation algorithms for self-talk and cross-talk at a receive chain produce their own estimation error which add up to the linear degradation. Our

key insight here is to leverage the fact that we have *M* transmitters available that can concurrently send training symbols. Specifically, we design a training preamble for WiFi that allows each receive chain to estimate each of the self-talk and cross-talk channels at an error that is *M* times lower than the SISO design by combining information from all *M* training symbols. Consequently, in our design when the estimation errors add up for the self-talk and cross-talk cancellations, the overall error or residue is the same as a SISO system would have achieved, which is the best one can hope for. Further the algorithm is *modular and structured in a way that, if in the future the SISO full duplex design manages to improve its performance even further, the MIMO design in this paper immediately benefits.*

We prototype our design using our own custom designed analog cancellation circuits, implement the digital cancellation algorithms and integrate them with off-theshelf WARP radios. Our experiments demonstrate that in a 3×3 configuration, our system achieves a performance that leaves a negligible 1dB of self-interference after cancellation. We also show that our system achieves a 95% throughput gain over half duplex radios using a standard WiFi compliant OFDM PHY of 20MHz for 802.11n for all different modulations (BPSK, QPSK, 16QAM and 64 QAM) and coding rates of (1/2,2/3,3/4,5/6), supporting three streams for 3×3 MIMO.

2 The Problem

In this section, we describe the nature of interference in a MIMO full duplex radio and then discuss the architectural challenges in designing a cancellation system.

Self-talk or cross talk (or for that matter any transmitted signal) is made up of three major components [5, 10, 2]:

- Linear Signal: This is the signal that the baseband modem wanted to transmit and is then distorted by channel reflections. Its linear because it can be represented as a linear combination of delayed and summed copies of the same signal that arise from environmental multipath reflections.
- Non-linear Signal: This is the signal that is generated due to non-linear transformations that the linear signal goes through when it is passed through analog radio components such as mixers, power amplifiers in the transmit chain.
- Transmit Noise: This is the noise that is generated by active TX chain components such as power amplifiers and local oscillators (we club things such as broadband noise and phase noise into this term for the sake of brevity).

The relative strengths of these components depends on the quality of the radio. Fig. 1a tabulates the strengths of the different components we empirically measured for a

Power and Interference relative to noise floor of -85 dBm

	Power level in dBm	Cancellation needed in dB
Total TX signal	20	105
Linear component	20	105
Non-linear component	-10	75
Transmitter Noise	-20	65

(a) The different components of the transmitted signal (self-talk) for a typical WiFi radio. The second column tabulates the amount of self-talk cancellation needs to eliminate the corresponding self-talk component to the noise floor.

Power and Interference relative to noise floor of -85 dBm

MIMO FD, Receiver 1	Power in dBm		Cancellation needed (dB)			
	Self- talk	Cross talk 1	Cross talk 2	Self talk	Cross talk 1	Cross talk 2
Overall signal at antenna 1	15	-9	-15	100	76	70
Linear component	15	-9	-15	100	76	70
Non-linear component	-15	-39	-45	70	46	40
Transmitter noise	-25	-49	-55	60	36	30
Cancellation Requirement						
MIMO FD, Receiver 1 S	elf-talk		Cross	Talk 1	Cross	Talk 2

	een van	0.000	0.000 runt 1
Analog cancellation	65 dB	41 dB	35 dB
Digital cancellation	35 dB	35 dB	35 dB

(b) Interference components and cancellation requirements for 3 antenna MIMO full duplex. The first table describes the levels of different interference components (linear, non-linear and transmit noise) that make up self-talk and cross-talks at one receiver in a 3 antenna MIMO radio. Cross-talk 1 is from the neighboring antenna and cross-talk 2 is from the farther neighboring antenna. The second table lists the overall cancellation needed, here the values are bumped up by 5dB relative to the first table to ensure that even when the residues left from the selftalk and the two cross-talk cancellations are added up, the overall noise floor does not go up (else it would go up by 5dB if the cancellation requirement for each component did not have a 5dB margin).

Figure 1: Requirement tables

commodity 20dBm WiFi SISO radio, and the amount of cancellation needed to eliminate them in a full duplex system. Note that this is a cheap radio widely used in many commercial WiFi devices [2, 4], so we believe this is representative of the WiFi radios in general.

The above analysis is of course true even for a single antenna radio without MIMO, and recent work [11] describes how to design a cancellation techniques that eliminates all these sources of self-talk. However, what is unique with MIMO is cross-talk. In other words, the interference that results at a receive chain due to a transmission from a neighboring co-located MIMO antenna/chain. In a 3 antenna full duplex MIMO radio, each receiver chain would see two cross-talk signals from the other two antennas as seen in Fig. 3.

Cross-talk is slightly weaker than the self-talk generated by the chain's own transmission, but is still quite strong and has all the above three enumerated components. Like the earlier SISO design [11], the transmit noise component of the cross-talk signal has to be canceled in the analog domain, whereas the non-linear and linear components could be canceled in both analog and



Figure 2: Prior best performing SISO full duplex design. The figure on the right shows an equivalent conceptual filter based view of self-talk cancellation. The filter is parameterized by its complexity, the number of taps. The filter subsumes both analog and digital cancellation.

digital domains. Fig. 1b tabulates the strengths of the various components that make up a cross-talk and self-talk signal in a typical 3-antenna MIMO WiFi radio with 20dBm transmit power (note that the power is divided equally among all three transmitters, so the power out of each antenna is 15dBm).

2.1 Why cant we reuse the SISO full duplex design by replicating it?

At first glance, the MIMO interference cancellation problem looks quite similar to a SISO full duplex problem, only replicated a few times. After all the cross-talk signal that needs to be canceled looks like an attenuated version of a chain's own self-talk signal that the SISO design manages to cancel completely. So why couldn't we replicate the SISO design $M^2 - M$ times for each of the cross-talk signals in a *M* antenna MIMO radio and be done with it (shown in Fig. 3)?

To understand the reason this might not work, it will help to have a conceptual understanding of what a SISO self-talk cancellation system accomplishes. At its core, the self-talk cancellation technique can be thought of as shown in Fig. 2. The input is the baseband signal that is being transmitted, to which transmit noise is added and the combined signal is passed through a linear and nonlinear unknown transfer function that captures the distortions introduced by the analog components and the wireless channel and is denoted by H. What the cancellation circuits and algorithms are trying to accomplish is estimate and recreate this unknown transfer function H as accurately as possible (to the tune of 105dB resolution), and then pass a copy of the input baseband transmitted signal and noise through this estimated transfer function to recreate the self-talk and cancel it. The estimated transfer functions are modeled using tunable analog and digital FIR filters, for example the prior SISO design's analog cancellation circuit requires 12 delay-attenuation



Figure 3: **SISO Replication Based Design**: Shows a 3 antenna full duplex MIMO radio, using nine SISO cancellation circuits (SISO replication design). This design uses in total 9N taps for M=3 assuming each circuit requires N filter taps. In the general case this design would require $M^2 \dot{N}$ for a M antenna full duplex MIMO system.

taps that each represent a single analog FIR filter tap, and what is being controlled is the weight on each tap (practically this translates to controlling the attenuator on that delay-attenuation analog line). A similar FIR filter structure is used for digital cancellation and the challenge is calculating the weights to use on each of the taps. So the key challenge the SISO self-talk cancellation system is solving is calculating a set of FIR filter weights that can accurately model this unknown and time-varying transfer function.

Consequently, there are two metrics that characterize this estimation circuits and algorithms.

- Complexity: can be quantified by the number of filter taps that are used in the implementations that represent the estimated \hat{H} . The more taps we need, the more analog circuitry is needed as well as DSP resources in FPGA to implement them. Keeping the number of taps low is important so as to reduce the space and power consumed by analog circuits [8] and DSP logic for FIR implementations (the baseline is the SISO design that requires 12 analog taps and 132 digital FIR taps). To get a sense of the impact, 12 analog taps consume roughly 24sq.cm of board area. A second consequence of complexity is the amount of time it takes us to re-tune the cancellation when the environment changes (including things such as temperature). The larger the number of taps, the longer it will take to tune since there are more variables to estimated. When cancellation is being tuned, the radio cannot be operated in full duplex mode. Hence tuning time is pure overhead, and needs to be minimized.
- Estimation error: A second key metric is estimation error which manifests as residual interference left after cancellation and directly reduces the SNR of the desired received signal. A perfectly accurate cancellation system would leave no residue. The baseline for



Figure 4: **Cascaded Cancellation Design**: Shows a 3 antenna full duplex MIMO radio design with cascaded filter structure for cancellation. The structure is shown for receiver chain 1 only, but the same structure is repeated for the other chains. For, self-talk cancellation we have N filter taps on every chain. Further we have C and D taps feeding in a cascading fashion at the input of the N tap self-talk cancellation circuit. Notice cross talk 1 is stronger so we need more taps (C > D) as compared to cross talk 2. However both *C* and *D* are significantly smaller than *N*.

this metric is the best performing prior SISO self-talk cancellation design that leaves 1dB of residue over the noise floor. In other words, the receiver noise floor is increased by 1dB and therefore the SNR of the received signal is also decreased by 1dB. To put this number in context, this is extremely accurate since at most normal receive link SNRs, a 1dB decrease will have negligible impact. The reason for this residue is estimation and quantization error in the algorithms that calculate the weights for the filter taps used in analog and digital cancellation. Estimation error is inevitable and cannot be avoided, but its important to keep it as small as possible.

How well would the SISO replication based design for MIMO perform on these two metrics? The optimal scenario given the SISO design is that the overall complexity of a M antenna full duplex MIMO radio would be $M \times$ the complexity of the SISO design, and it would have the same estimation error as the SISO design. We cannot do better than a linear increase in complexity and no increase in estimation error.

However, the SISO replication based design doesn't come close to achieving these optimal targets. Its complexity is $M^2 \times$ the complexity of the SISO design. This is because it requires us to replicate the SISO design for each cross-talk factor, and therefore we need a total of M^2 versions of the SISO design. In terms of taps this implies $12 * M^2$ taps in analog circuits alone, along with the

corresponding increase in digital cancellation FIR taps.

Second, this design's estimation error turns out to be even more worse. At each receiver chain, we show in Sec. 3.2 that the residual interference scales linearly with the number of MIMO chains M. Intuitively the reason is that each replica of the SISO design is running an independent estimation algorithm for determining the values of the filter taps to use for cancellation. Since at each receiver chain we have M versions of the SISO design running, we will have a $M \times$ increase in estimation error and consequently the interference residue.

3 Design

We present a new cross talk cancellation technique for full duplex MIMO which is scalable and efficient. The key technique behind our MIMO cancellation design is a cascaded filter structure. Specifically, we exploit the fact that in MIMO, cross-talk and self-talk are correlated, since they share a similar environment (or similar set of multi-path reflection and attenuation profiles in the channel). Further, cross-talk across chains is naturally reduced compared to the chain's own self-talk because of physical antenna separation. We exploit these insights to design a low complexity and low error cross-talk cancellation system. For canceling the chain's own self-talk we use the design from prior work [11].

3.1 Reducing Complexity: The Cascade

Our design builds on a key insight: co-located MIMO antennas share a similar environment. In other words the transfer function that transforms the cross-talk signal at a receive chain has a close relationship with the transfer function that the chain's own self-talk undergoes. Intuitively this is because the near-field environment around a radio looks essentially the same to neighboring antennas since they share the same reflectors in the environment, and the distance to these reflectors is almost the same from the closely-spaced antennas. The only difference is the additional delay experienced by the cross-talk signal compared to the chain's own self-talk. Note that this does not mean that MIMO self-talk and cross-talk channels are completely correlated, they will still have independent phase that arises from the slightly different delays experienced by self-talk and cross-talk signals at RF frequencies. But the amplitudes of the multi-path reflections for self and cross-talks from the same reflector tend to be nearly the same since the distance to the reflector is almost equal from the closely spaced MIMO antennas.

The above insight can be mathematically modeled as a cascade of transfer functions. Let $H_i(f)$ and $H_{ct}(f)$ be the transfer functions of the chain's own self-talk and cross-talk respectively. The overall relationship between these functions can be modeled as follows:

$$H_{ct}(f) = H_c(f)H_i(f) \tag{1}$$



(a) Cancellation performance in the frequency domain for the cascaded design and the replication based design with the same complexity for a 3 antenna MIMO full duplex radio operating a WiFi PHY in a 20MHz band at 0dBm TX power.

Resource Comparison between SISO replication and Our design					
	SISO replication design	Our design			
Analog Cancellation taps (3X3)	108 (12*9)	56 (reduced by 1.92x)			
Digital Cancellation taps (3X3)	1188 (132*9)	485 (reduced by 2.45x)			
Tuning time (3X3)	9 ms (1ms*9)	.024 ms (reduced by 375x)			
Analog Cancellation taps (mXm)	O(M ² N)	O(MN)			
Digital Cancellation taps (mXm)	O(M ² R)	O(MR)			
Tuning time (mXm)	O(M ²)	O(M)			

(b) Table showing the reduction in complexity and tuning time with the cascaded design compared to the replication based design for both a 3 antenna full duplex MIMO radio as well as the general case of a *M* antenna full duplex MIMO radio.

Figure 5: Cascade Design Evaluation.

where $H_c(f)$ is the cascade transfer function. The key observation is that $H_c(f)$ which cascaded with $H_i(f)$ results in the cross-talk transfer function, is an extremely simple transfer function. Typically $H_c(f)$ is a simple delay that corresponds to the fact that the two antennas are separated and the cross-talk signal experiences slightly higher delay compared to the self-talk.

How might we exploit this insight? The idea is to mimic the cancellation design in a cascade similar to the equation above as seen in Fig.4. Specifically, we could design simple low-complexity analog cancellation circuits and digital cancellation filters that model the cascade function $H_c(f)$. These circuits and filters would then feed into the cancellation circuits and digital cancellation filters for the chain's own self-talk cancellation and thus reuse all that circuitry to model the cross-talk channel. Remember that the circuits and digital filters for the chain's own self-talk are modeling $H_i(f)$, hence the cascaded structure is essentially recreating the above equation. So the only additional complexity compared to the optimal MIMO design would be from the circuits and filters that model the cascade transfer function $H_c(f)$.

The natural question is of course, what is the complexity of the circuits and filters that model the cascade function? We empirically measure it by implementing different complexity cascade cancellation circuits and filters and finding out what complexity suffices to meet the cancellation requirements described in the previous section. We then provide the SISO replication based design the same number of taps and compare the cancellation performance. We conduct the experiment by transmitting a 20MHz WiFi OFDM 3×3 MIMO signal, and attempt to cancel it using the two approaches.

Fig. 5a plots the results for the impact of the cascaded structure on the efficacy of analog cancellation. As we can see the cascaded structure consistently provides 18dB better performance than the SISO replication based design with the same complexity and meets the required 65dB of analog cancellation. The number of analog taps required to realize this performance using the cascaded design is tabulated in Fig. 5b. For a typical 3 antenna MIMO WiFi radio with 12cm separation between antennas (typical of APs), the antenna separation itself provides about 24dB of isolation, so we need another 41dB of cross-talk cancellation in analog (see Table. 1b for requirements). As we can see we need only four analog taps with the cascaded structure compared to the 12 taps required by the naive design for canceling cross-talk at an adjacent antenna and only two taps, when canceling to the farther out antenna as shown in Fig. 4. The cascaded design therefore requires $1.92 \times$ lower number of taps compared to the SISO replication design for a 3 antenna full duplex MIMO radio as seen in Fig. 5b. The reduction factor approaches the optimal $3 \times$ number as the number of antennas increases.

To verify the improvement for digital cancellation, we conduct a similar experiment with the same setup (20 dBm TX power). However, we provide the SISO replication design the required number of taps to meet the requirement on analog cancellation so we can specifically evaluate the benefits for digital cancellation with cascading. As seen in Fig. 5b, with a cascaded structure we need a total of 485 taps to cancel self-talk and cross-talk to the noise floor for a 3 antenna MIMO radio. Further, for the SISO replication based design using the same number of taps (485), the residual interference is still an additional 7dB. To achieve the same performance as our cascaded design with the SISO replication based design, we would need 1188 or $2.45 \times$ more taps as tabulated in Fig. 5b. Once again the reduction factor approaches the optimal number M and the number of antennas (M) grows. Finally in terms of cancellation performance, a 7dB increase in noise floor or reduction in desired signal's SNR is quite high by itself, and when we take into account the reduction in cancellation for analog of 18dB, we are looking at a 25dB reduction in overall cancellation for the SISO replication based design with the same complexity as our cascaded structure.

There are two main benefits to reducing complexity: **Size, Cost and Tuning time reduction**: Each additional filter tap increases the size of cancellation boards in analog and FPGA resource consumption in digital cancellation. For analog cancellation, our circuits consumed 110sq.cm of board area compared to nearly 216sq.cm for the SISO replication based design for a 3-antenna MIMO full duplex system. For example, we found experimentally that reducing the number of digital filter taps from 1185 to 485 for a 3 antenna MIMO radio means that a lower class Xilinx Kintex series FPGA has sufficient DSP resources to implement the cancellation, whereas the SISO replication based design would require the higher end Virtex FPGA [6]. This translates to enormous power savings, a Virtex FPGA consumes nearly 80W of power whereas a Kintex consumes only 40W on twice as less [7]. Power reduction translates to less heat and consequently simpler AP designs. Finally to ultimately realize the design in compact boards or in ICs, reducing the number of taps as much as possible is a must. A final consequence is the tuning time to compute the weights for each of these taps also reduces linearly with lesser number of taps (tuning time is pure overhead since during tuning the radio cannot be used for communication).

Transmit power waste: The amount of power that needs to be couple off from the transmit paths to powering cancellation circuits depends linearly on the number of taps in the cancellation circuits. This is because each tap is of course only useful if some copy of the transmitted signal is passed through it, and in addition each tap has loss associated with it that adds up. Thus reducing number of taps helps reduce TX power waste, which in turn translates to lesser battery consumption especially for mobile devices.

3.2 Reducing Interference Residue: Joint Training

The goal of digital cancellation is to clean out any remaining residual self-interference. Once again, a natural question is why not reuse the digital cancellation algorithms designed for SISO? In other words, for each receive chain in a M antenna full duplex MIMO radio, run M separate digital cancellation algorithms that estimate the chain's own self-talk and the other M - 1 crosstalk interference components. These algorithms work by estimating the distortion experienced by each of the interference (both for linear and non-linear components). They then apply the estimated distortion functions to the known baseband copy of the transmitted signal and subtract it from the received signal.

The above approach doesn't work because every additional and independent digital cancellation algorithm we use in the receive chain linearly increases the residual interference after cancellation. In other words, performance worsens linearly with the number of MIMO chains. To see why, we start with describing why even a simplified SISO digital self-interference cancellation algorithm will have some residual interference that cannot



c) Our Design based 2X2 MIMO Full Duplex

Figure 6: This figure shows the transmitted and received packets for a SISO full duplex, 2 antenna MIMO full duplex with the traditional training technique, and our design with the novel training technique. Notice the training symbol structure in the last figure, this allows us to reduce the estimation error by half for the self-talk and cross-talk components for a 2 antenna MIMO radio.

be canceled.

Digital cancellation works in two stages, first there is a training phase and then cancellation phase. The training phase uses training symbols (e.g. the WiFi preamble), and the assumption is that there is no desired received signal from the other full duplex node. The training symbols are used to estimate the self-interference. Lets say the training self-interference symbol is s as seen in Fig. 6.a. The self- interference symbol is being received after transmission from the same radio (for simplicity assume there is no distortion from the channel), and the receiver adds its own noise n_1 (variance σ^2) to the received signal (this noise comes from effects such as quantization in the ADC). Hence the received signal y_1 can be written as,

$$y_1 = s + n_1$$

The best estimate of the self-interference s in this case is simply y_1 However this estimate \hat{s} has some estimation error, which in this case is simply the power of the receiver noise as show below:

$$\hat{s} = y_1, \ E((s-\hat{s})^2) = E(n_1^2) = \sigma^2$$

How can we use this estimate to cancel subsequent selfinterference? For simplifying the description, lets assume the packet that is being transmitted and is acting as self-interference is simply the same training symbol repeated throughout the packet (real world packets are of course not trivial like this, but this assumption does not change the basic insight below). To cancel this selfinterference throughput the packet, the algorithm will simply subtract the above estimate from the overall received signal. Lets say x is the actual desired received signal, the overall signal received is y, and the signal after cancellation, are given by:

$$y = x + s + n_3$$

$$\underbrace{y - \hat{s}}_{\text{cancellation}} = x + \underbrace{s - \hat{s}}_{estimation \ error = \sigma^2} + \underbrace{n_3}_{\text{RX noise}}$$

As we can see, the estimation error shows up as residual interference with variance of σ^2 . As the best known prior design has shown this is on the order of 1dB over the half-duplex noise floor.

SISO Replication based MIMO design: Its now easy to see why a design for MIMO that simply uses M replicas of the digital cancellation algorithm at each receive chain for the self-talk and the M-1 cross-talk interference signals increases the estimation error roughly by a factor of *M*. The training symbol structure for a 2×2 MIMO transmission is shown in the Fig. 6.b. above, essentially there are two training symbols s_1 and s_2 sent over two slots from the two different transmit chains. The algorithms at a particular receive chain use these symbols like in the SISO case to estimate the self-talk and the crosstalk, and each of them will have their own estimation error. When these estimates are used for cancellation, the estimation errors add up, and the overall estimation error (or residual self-interference) at each receive chain is theoretically two times the SISO case. The math below shows the above intuition formally. First, the estimates for the self-talk and cross-talk symbols are given by:

$$\hat{s}_1 = y_1, \ E((s_1 - \hat{s}_1)^2) = \sigma^2$$

 $\hat{s}_2 = y_2, \ E((s_2 - \hat{s}_2)^2) = \sigma^2$

When canceling to attempt to recover the desired received signal x, we can calculate the estimation error as follows:

$$y = X + s_1 + s_2 + n_3$$

$$\underbrace{y - \hat{s_1} - \hat{s_2}}_{\text{cancellation}} = X + \underbrace{s_1 - \hat{s_1}}_{\sigma^2} + \underbrace{s_2 - \hat{s_2}}_{\sigma^2} + \underbrace{n_3}_{\text{RX noise}}$$

As we can see, the estimation error shows up as residual interference with variance of $2 * \sigma^2$, both self-talk and cross-talk estimation introduce σ^2 error. We can recursively show that for a general M antenna full duplex MIMO radio, the estimation error and consequently residual interference on each receive chain goes to $M\sigma^2$. Our Design: Our key contribution is a novel training symbol structure and estimation algorithm that reduces the estimation error for each interference component at each receiver chain (self-talk or cross-talk) to σ^2/M for a full duplex $M \times M$ MIMO radio. The key insight is to re-design the training symbols to reduce the estimation error. Specifically instead of sending training symbols

from each of the transmit chains separately in consecutive time slots, we send a combination of all of them from each transmitter in parallel. The idea is to actually leverage the fact that there are two transmitters that could be leveraged to transmit training information jointly and thereby improve accuracy, there is no need to treat each of them separately. Doing so requires an intelligent joint training symbol design so that each symbol can be estimated as a linear combination of the received transmissions. Fig. 6.c. shows the main idea.

We use a similar set of equations as before to show formally why this works. As seen in Fig. 6.c., the training symbols are transmitted by chain 1 and chain 2 simultaneously. In time slot 1, transmitter 1 and 2 transmit s_1 and s_2 , respectively. And in time slot 2, transmitter 1 and 2 transmit s_1 and $-s_2$ respectively. Receiver 1, receives the combined symbols in time-slot 1 and time-slot 2, y_1 and y_2 . Thus: $y_1 = s_1 + s_2 + n_1$, $y_2 = s_1 - s_2 + n_2$

Lets assume the rest of the transmissions from the two chains are just repetitions of the same symbols s_1 and s_2 respectively (again this is for description simplicity and suffices to explain the insight). We need to get estimates for the data symbols s_1 and s_2 using the two training symbols y_1 and y_2 . The best estimates are given by:

$$\hat{s}_1 = \frac{y_1 + y_2}{2}, \ E(s_1 - \hat{s}_1)^2 = E((\frac{n_1 + n_2}{2})^2) = \frac{\sigma^2}{2}$$

 $\hat{s}_2 = \frac{y_1 - y_2}{2}, \ E(s_2 - \hat{s}_2)^2 = E((\frac{n_1 - n_2}{2})^2) = \frac{\sigma^2}{2}$

As we can see, the error in each of these estimates (self-talk and cross-talk) is $\sigma^2/2$. Now when these estimates are used for cancellation, the following equation results:

$$y = x + s_1 + s_2 + n_3$$

$$\underbrace{y - \hat{s_1} - \hat{s_2}}_{\text{cancellation}} = x + \underbrace{s_1 - \hat{s_1}}_{\frac{\sigma^2}{2}} + \underbrace{s_2 - \hat{s_2}}_{\frac{\sigma^2}{2}} + \underbrace{n_3}_{\text{RX noise}}$$

As we can see the residual interference is only σ^2 , rather than the $2\sigma^2$ that would have resulted from the SISO replication based design. Further, we can show by recursion that this residual is the same as the SISO design, i.e. there is no linear increase with the number of MIMO chains as the number of antennas increases. The training symbols for the general *M* antenna case are designed such that appropriate linear combinations of the received symbols at any chain can be used to estimate the self-talk and cross-talk symbols individually, the only requirement for designing these training symbols is that the linear combinations in multiple time slots at a receive chain are not degenerate.

4 Robust MIMO Interference Cancellation

Interference cancellation needs to be robust to enable consistent full duplex operation in the face of frequent channel changes. To accomplish this, both analog and digital cancellation need to continuously tune their filter taps to maintain cancellation. The main bottleneck is tuning analog cancellation, since digital cancellation can be tuned on a per-packet basis in software as prior work has shown [23, 19, 11]. Tuning analog circuits requires measuring the residue in digital and then sending control signals to analog components, which is relatively slow. Minimizing the amount of time required to tune here is therefore critical, since during the time spent tuning packets likely cannot be received. We focus on this problem in this paper and re-use the algorithms from prior work for tuning digital cancellation.

The prior SISO full duplex design demonstrated a technique to tune a single analog cancellation in around a millisecond. However, as before if we were to naively replicate the same algorithm for all the self-interference components, we would need M^2 ms for a M antenna full duplex MIMO radio (e.g. 9ms for a 3 antenna full duplex). Such a high overhead is untenable for moderately mobile environments where the channel changes on average every 60ms (e.g. WiFi hotspots).

In this paper we propose a novel technique that reduces tuning time by three orders of magnitude, i.e. an algorithm that tunes the circuit in 8μ s. Note that this algorithm also applies to the SISO case, and therefore improves on the best known prior SISO design too. Our insight is to model the cancellation circuit as a filter whose response we are tuning to match as closely as possible the frequency response of the self-interference channel. Like prior work, we estimate the frequency response of the cancellation circuit for different combinations of filter tap values. The pre-calculated response is represented in a matrix A, whose each column is the frequency response of the analog cancellation circuit for a particular value of the filter tap at K different frequencies in the band of interest (e.g. K=128 for a 20MHz bandwidth in our current prototype for WiFi). Now assuming H(f)is the frequency response of the self-talk channel in the frequency domain (i.e. the channel introduced by the antenna, circulator and any strong environmental reflections), the analog cancellation tuning problem reduces to:

$$\min_{x} ||H - Ax||^2$$

Where, *H* is the column consisting of H(f) at different frequencies, and *x*, represents a binary indicator vector for selecting the corresponding filter tap values as in [11].

The efficacy of the tuning that results from the above problem depends on the accuracy in the measurement of H(f). We can measure H(f) using the preamble of the received interference signal y(t) (e.g. the first two OFDM symbols of a transmitted WiFi packet which are known preamble symbols). The challenge is measuring the frequency response of the interference channel accurately. The accuracy is limited by the linearity of the transmit-receive chain, which is 30dB, By this we mean that any initial measurement can only have an accuracy of 30dB. The main reason is that the transceiver produces non-linearities which act as noise to the channel estimation algorithm. In other words the received interference signal y(t) has non-linearities that are only 30dB below the main linear signal component. Our key contribution in this paper is a technique to accurately measure this channel quickly in the presence of non-linearities and tune analog cancellation.

Source of error and its magnitude: The transmitter produces non-linearities 30 dB lower than the transmitted signal. To show mathematically, say x(t) is the baseband signal that is being transmitted after up-conversion and amplification, we can write

$$x_{tx}(t) = x(t) + a_3 x(t)^3 + a_5 x(t)^5 + a_7 x(t)^7 + \dots + w(t)$$

This transmitted signal $x_{tx}(t)$ is somewhat known to us because we know x(t), however its non-linear components and the transmit noise w(t) are unknown. This signal further undergoes the circulator and antenna channel H(f) (which we wish to estimate), so when its received at the receiver the frequency domain representation of the received signal is given by:

$$Y(f) = H(f) * \mathscr{F}(x(t) + a_3 x(t)^3 + \ldots) + \text{ transmit noise}$$

Here, a_3 is around $10^{(-30/20)}$, i.e., its 30 dB lower. Further transmit noise distortion is 40 dB lower than the signal level of x(t). The challenge is that our channel estimation algorithm is only going to use its knowledge of x(t) to estimate the channel H(f), and the other terms in the received interference signal limit the accuracy of the estimation to 30dB (the estimation noise is 30dB lower). Accurate, Iterative method: The key idea is to run the estimation algorithm in an iterative fashion. Remember that the WiFi preamble has two OFDM symbols, each of length $4\mu s$. After the first OFDM symbol, we solve the above equation to produce an inaccurate estimate of the interference channel H_a and tune the cancellation circuit to achieve (at best) 30dB of cancellation (we cannot cancel more than our estimation accuracy). Now when we obtain the second preamble symbol, we know that the non-linearities and the transmit noise components that were producing the error are reduced by 30dB. We can exploit this fact by the following trick:

We transmit one OFDM symbol to estimate the inaccurate H_a , which can be written as a function of accurate H as, $H_a = H + e_1$. Note e_1 is 30 dB lower then H. We use the same algorithm as [11] to optimize the following,

$$\min_{x} ||H_a - Ax||^2$$

which produces the solution as \hat{x} , which gives us the values to use in the filter taps. We program the cancellation

circuit using these values and achieve a 30 dB cancellation. Next, when we transmit second OFDM symbol and measure the channel response we get:

$$H_b = (H - A\hat{x}) + e_2$$

Notice that e_2 is 30 dB lower then $H - A\hat{x}$ and $H - A\hat{x}$ is 30 dB lower than H. So in essence e_2 is 60 dB lower than H. Define,

$$\tilde{H} = H_b + A\hat{x}$$
$$\tilde{H} = H + e_2$$

Thus, we can this new estimate \tilde{H} with an error that is 60 dB lower. We use this estimate to re-tune the optimization algorithm and find a solution \tilde{x} that tells us what values to use for the analog filter taps. This new solution provides 60 dB cancellation. Further, we only needed two OFDM symbols of $4\mu s$ each to get to this cancellation.

Extension to Cascaded Filter Structure: The above description is for a single cancellation circuit, but our MIMO design has a cascaded structure of multiple circuits. This leads to a combinatorial explosion in the parameter space that makes the problem NP hard to solve if we use the above approach. In this subsection we present a trick to approximate the overall combinatorial problem via two reduced complexity problems which can be solved using the same technique as the SISO one presented above.

We describe the algorithm in the context of tuning the cancellation circuits at receiver 1 for self and cross-talk in a 2 antenna MIMO radio. Lets say H_{11} is the self-talk channel response and H_{12} is the cross-talk channel response. The general tuning problem can be stated as:

$$\underset{x_{1},x_{2}}{\text{ninimize}} t \tag{2}$$

subject to
$$norm(H_{11} - A_1 x_1) \le t$$
 (3)

$$norm(H_{12} - (A_1x_1) \odot (A_2x_2)) \le t$$
 (4)

Where, \odot represents the element wise multiplication of the column, and *t* represents the analog cancellation achieved, and A_1 is the response of the self-talk cancellation board with N taps in Fig.4 and A_2 is the response of the cascade cancellation board with C taps. The second constraint Eq. 4 renders the problem irreducible to a convex solvable form, and in fact the columnwise multiplication of the indicator variable vectors explodes the problem space and makes it a NP hard combinatorial problem.

We use a novel trick to approximate and help solve this problem practically. Since the first constraint in Eq. 3 is trying to find $A_1x_1 = H_{11}$, we can approximate A_1x_1 in the next constraint, Eq.4 with H_{11} which is known (since we measured H_{11}). This is of course an approximation, but it suffices to solve for x_2 using this substitution since we are after all trying to emulate the same cascaded channel response structure using our circuits as described in Sec. 3. Thus instead of a cascade of unknown variables, the new problem to solve is

$$\begin{array}{ll} \underset{x_1, x_2}{\text{minimize}} & t \\ \end{array} \tag{5}$$

subject to $norm(H_{11} - A_1x_1) \le t$ (6)

$$norm(H_{12} - H_{11} \odot (A_2 x_2)) < t$$
 (7)

This new problem is no longer a combinatorial problem. This can be reduced to an integer program, which can be solved using randomized rounding in fraction of micro seconds practically [11]. Thus in effect the substitution trick reduces the non-tractable combinatorial problem into a tractable problem, whose solution can be found using the techniques described above. The tuning time for each MIMO chain is still two OFDM symbols, and the overall tuning time for the MIMO radio therefore scales linearly with M, the number of chains.

5 Evaluation

In this section, we experimentally demonstrate that our MIMO full duplex design almost completely cancels all self-talk and cross-talk interference to the noise floor with a low-complexity design. We also show that this translates to a doubling of throughput for the link performance.

We implement our design using four WARP v2 boards for building a 3×3 MIMO full duplex link. We design our own boards for analog cancellation and integrate them with the WARP boards. At each receive chain, we have analog circuits with 12 taps for the self-talk cancellation, 4 taps for the first cross talk and 2 taps for the farthest transceiver. In total we have 56 taps in the analog cancellation circuits for a 3 antenna full duplex MIMO radio, and total of 485 filter taps in digital cancellation. Since the WARP cannot generate 20dBm transmit power, we use an external off-the-shelf power amplifier [3].

We compare against the SISO replication based design primarily. This is the straightforward replication of the recently published SISO full duplex design as discussed at the start of Sec. 3. We compare against two variants of this design. One is a design that fully replicates the analog and digital cancellation implementations for all selftalk and cross-talk cancellations. As discussed before the complexity of this design is a factor of two higher for analog and $2.5 \times$ higher for digital compared to our design. We call this design SISO Replication. However to make an apples to apples comparison with our design we also implement a SISO replication design with the same complexity as our design. The difference compared to our design is that, it neither use the cascaded structure nor the novel estimation algorithm, but simply replicates the SISO design with lower number of taps. We experiment with the tap distribution between self-talk and



Figure 7: Spectrum plot after cancellation of various selftalk and cross-talk components for RX1 of a 3×3 full duplex system using our design.

cross-talk to obtain the best overall cancellation. We call this compared approach **SISO Low Complexity Replication**.

Unless stated otherwise, all experiments are conducted by placing the two full duplex nodes at various locations in our department building. At each location, we repeat the experiment ten times and calculate the average performance.

5.1 Can we cancel all the interference for 3 antenna full duplex MIMO ?

The first claim made in this paper is capability of canceling all of the interference for the 3×3 MIMO. To prove this, we experimentally test if we can fully cancel a WiFi 802.11n 20MHz signal upto a max transmit power of 20dBm for a 3×3 MIMO. To demonstrate we first pick one instance of this experiment, and show the spectrum plot of the received self-interference after various stages of cancellation in Fig. 7. Remember, that in analog we first cancel the chain's own self-talk leaking through the circulator, and then the cross-talk from the other two antennas. Finally, we apply our digital cancellation step to clean up the residual. We see that overall in analog we achieve 68-70dB of self-interference cancellation after all three stages. This satisfies the requirements outlined in Sec. 2.

We now place the node at several different locations in the testbed. At each location we vary the overall TX power from 16dBm to 20dBm and plot the average cancellation for each power across all locations. At each location and for each power, we conduct 40 runs. The goal is to show that we can consistently cancel to the noise floor for a variety of transmit powers up to and including the max average TX power of 20dBm. In each instance of the above experiment, we also measure the increase in noise floor due to any residual self-interference that is not canceled. Note that the increase in noise floor represents the SNR loss the received signal will experience when the node is used in full duplex mode. Fig. 8 plots



Figure 8: Increase in noise floor vs TX power on the left side and Cancellation vs TX power on the right side. For different MIMO cancellation designs, we present the performance of a full duplex 3 antenna full duplex MIMO system.

the average cancellation and the increase in noise floor as a function of TX power.

Fig. 8 shows that our 3-antenna MIMO full duplex design cancels the entire self interference almost to the noise floor. In case of max average transmit power of 20dBm [11], the noise floor is increased by 1.6dB over each receive chain's noise floor. The SISO replication design increases the noise floor by 4dB per receive chain, while the SISO low complexity replication approach increases the noise floor by 25dB. Finally, the performance of our design and the SISO replication design scales with increasing TX power, while the other replication based design is limited due to its inability to cancel the increasing transmit noise and non-linearities due to the reduced number of taps available to it.

5.2 Scaling with the number of MIMO antennas

A question with MIMO is how does full duplex performance scale with increasing number of transmit chains. The ideal case would be to maintain the same level of cancellation at each RX chain as the number of transmit antennas increase, starting from one antenna. In other words, even with increasing number of transmit antennas and cross-talk components that need to be canceled, we retain the same performance as if there was a single transmit antenna and a single self-interference signal to deal with. Fig. 9 plots the increase in the noise floor at one receive chain as we go from one transmit chain to three transmit chains for a MIMO radio for both our design as well as the SISO replication technique. The overall TX power is fixed to be 20dBm (additional 10 dB of PAPR for WiFi [11], i.e., total 30 dBm) to adhere to ISM band EIRP requirements. Hence if we use a single transmit chain, then all the 20dBm is used for a single antenna. If we use two chains, then each antenna produces a 17dBm signal and so on.

As we can see from the figure, our design maintains a near-constant performance even as we go from one to three transmit chains. In other words, the performance is



Figure 9: Increase in noise floor at a RX chain as the number of MIMO chains and consequently the number of cross-talk components increase from 1 to 3. With our design we observe a 2.5 dB improvement for 3×3 MIMO per RX chain compared to the SISO replication design.

roughly the same regardless of the number of cross-talk components (We do wish to note that we could not go beyond three transmit chains due to hardware limitations, verifying the above claim for higher number of transmit chains is future work). On the other hand, the SISO replication design shows the noise floor increasing linearly with increasing number of transmit chains, a fact we provided theoretical intuition for in Sec. 3.2. Thus this design will look worse as we scale to higher MIMO configurations. We omit the SISO low complexity replication approach because its results are significantly worse.

5.3 Dynamic Adaptation

An important metric for analog cancellation is how quickly can it be tuned, and how often do we need to tune? The best know prior technique [11] required around 1 millisecond to tune a single SISO analog cancellation circuit. So for a 3×3 MIMO, applying the same algorithm will take at least 9ms for the SISO replication based design. In this section we show the efficacy of our new tuning algorithm which cuts the tuning time to $8\mu s$ per receive chain. Fig. 10 shows the tuning time as a function of the amount of analog cancellation. To achieve the 70dB analog cancellation, our algorithm takes $8\mu s$ per chain, for a total of $24\mu s$ for the full radio. The prior work as we can see take a millisecond per chain. The interesting takeaway is that both schemes achieve 40dB of analog cancellation fairly quickly (with one preamble symbol, i.e. $4\mu s$), but our scheme covers the final 30dB in one more step of $4\mu s$, while the prior scheme takes an exponential number of symbols to achieve that. The reason for this improvement is precisely our ability to get a precise measurement of the selfinterference channel using the trick described in Sec. 4.

A second question is how often one needs to tune? This depends on the environment and the amount of analog cancellation that needs to be maintained. In this paper, we tune for challenging indoor environments which



Figure 10: Tuning time for analog cancellation. The first figure shows the three orders of magnitude improvement in tuning time with our algorithm compared to the best known prior approach. The second figure shows how often this tuning algorithm needs to be run for an indoor environment.



Figure 11: CDF of throughput gain relative to half duplex 3×3 WiFi MIMO. Our 3×3 MIMO system provides a median gain of 95% relative to half duplex, whereas the SISO replication design only provides a $1.36 \times$ relative gain.

have strong multi-path (this is the main source of analog cancellation degradation). We define a near-field coherence time which depends on the amount of analog cancellation and is essentially the time for which that analog cancellation can be maintained on average before the circuits need to be retuned. Fig. 10 plots the near-field coherence time for three different analog cancellation targets. As we can see, to maintain an analog cancellation of 70dB, we need to retune roughly every 60ms. Given that our tuning overhead is $24\mu s$, our overhead is negligible.

5.4 Does Full Duplex Double Throughput?

A final question is whether all this cancellation performance translates to a the desired doubling of overall throughput. We show experimentally the throughput gains of our 3×3 MIMO full duplex design compared to the SISO replication based design. Two full duplex 3-antenna MIMO nodes are placed at different locations and we send 1000 packets in full duplex mode between them, and then send 1000 packets for each direction of the half duplex mode. We repeat this experiment for each bitrate that is available in WiFi. We pick the bitrate which maximizes the overall throughput for all of the compared full duplex designs and half duplex respectively. We repeat this experiment for 50 different locations. We found the received power of the links varied uniformly between -45 to - 80dBm, across locations as found in typical indoor deployments. To put these numbers in perspective, this implies that the SNR of the links in half duplex mode ranges from 5dB to 40dB.

We plot the throughput for half duplex and full duplex designs in Fig. 11. Note that all of these throughput numbers account for the overhead introduced by the periodic analog cancellation tuning. As we can see, our full duplex system achieves a median throughput gain of $1.95 \times$ over the half duplex mode, but the SISO replication based design with full complexity only achieves a $1.36 \times$ gain. The reason is the higher increase in noise floor from the SISO replication based design. For example, if the link SNR in half duplex mode is 10dB, a 4dB increase in noise floor will result in worse overall throughput for full duplex compared to running the link in half duplex mode. Our ability to keep the noise floor constant results in a performance close to the theoretical optimum.

The SISO replication based design with lower complexity is quite poor, in fact in 70% of the scenarios, the throughput was zero. This is because it increases the noise floor by at least 25dB which acts as noise and if the SNR is below 30dB no signal is decoded (WiFi requires a minimum of 4 - 5dB SNR to decode the lowest rate packet). As the half-duplex link SNR increases, the performance improves but is still not sufficient to beat the system throughput achieved by half duplex. The reason is that even if the link half-duplex SNR is 35dB, it implies that we only have two 10dB links for full duplex. The throughput achieved with a single 35dB half duplex link is still higher than two 10dB links. Consequently the only region where we could find improvements for full duplex over half duplex with this design was when the link SNR was greater than 38dB.

6 Conclusion

This paper brings towards completion a line of work on PHY layer of full duplex radios, and shows that practical full duplex is achievable for the most common wireless protocols and for MIMO while using commodity radios. The cancellation techniques developed in this paper are fundamental and apply to a wide variety of problems [17] where self-interference cancellation is needed. While this work wraps up work on board level realizations of full duplex, much work remains in realizing these designs in a chip. Tackling these problems is future work.

References

- Huawei Sets Out Its 5G Stall. http://www.lightreading.com/ document.asp?doc_id=703466&init_gateway=true.
- [2] Power Amplifier Data-sheet. http://datasheets.maximintegrated. com/en/ds/MAX2828-MAX2829.pdf.
- [3] Power Amplifier Data-sheet. http://www.minicircuits.com/pages/ npa/PGA-105+_NPA.pdf.
- [4] WARP Project. http://warpproject.org.
- [5] White paper by NI on Understanding Dynamic Hardware Specifications. http://www.ni.com/white-paper/5529/en.
- [6] Xilinx 7 Series FPGA Overview. http://www.xilinx.com/support/ documentation/data_sheets/ds180_7Series_Overview.pdf.
- [7] Xilinx Power Estimator Tool. http://www.xilinx.com/products/ design_tools/logic_design/xpe.htm.
- [8] E. B. And. An analog cmos high-speed continuous-time fir filter, 2000.
- [9] E. Aryafar, M. A. Khojastepour, K. Sundaresan, S. Rangarajan, and M. Chiang. Midu: enabling mimo full duplex. In *Proceedings of the 18th annual international conference on Mobile computing and networking*, Mobicom '12, pages 257–268, New York, NY, USA, 2012. ACM.
- [10] J. Bardwell. Tech Report. http://www.connect802.com/download/ techpubs/2005/commercial_radios_E0523-15.pdf.
- [11] D. Bharadia, E. McMilin, and S. Katti. Full duplex radios. In Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM, SIGCOMM '13, pages 375–386, New York, NY, USA, 2013. ACM.
- [12] J. I. Choi, M. Jain, K. Srinivasan, P. Levis, and S. Katti. Achieving single channel, full duplex wireless communication. In *Proceedings of the sixteenth annual international conference on Mobile computing and networking*, MobiCom '10, pages 1–12, New York, NY, USA, 2010. ACM.
- [13] M. Duarte, C. Dick, and A. Sabharwal. Experiment-driven characterization of full-duplex wireless systems. *CoRR*, abs/1107.1276, 2011.
- [14] E. Everett, A. Sahai, and A. Sabharwal. Passive self-interference suppression for full-duplex infrastructure nodes. *CoRR*, abs/1302.2185, 2013.
- [15] E. Fear, S. Hagness, P. Meaney, M. Okoniewski, and M. Stuchly. Enhancing breast tumor detection with near-field imaging. *Microwave Magazine*, *IEEE*, 3(1):48–56, 2002.
- [16] S. Gollakota, H. Hassanieh, B. Ransford, D. Katabi, and K. Fu. They can hear your heartbeats: non-invasive security for implantable medical devices. *SIGCOMM Comput. Commun. Rev.*, 41(4), Aug. 2011.
- [17] S. S. Hong, J. Mehlman, and S. Katti. Picasso: flexible rf and spectrum slicing. In Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication, SIGCOMM '12, pages 37–48, New York, NY, USA, 2012. ACM.
- [18] Y. Hua, P. Liang, Y. Ma, A. Cirik, and Q. Gao. A method for broadband full-duplex mimo radio. *Signal Processing Letters*, *IEEE*, 19(12):793–796, dec. 2012.
- [19] M. Jain, J. I. Choi, T. Kim, D. Bharadia, S. Seth, K. Srinivasan, P. Levis, S. Katti, and P. Sinha. Practical, real-time, full duplex wireless. MobiCom '11, pages 301–312, New York, NY, USA, 2011. ACM.
- [20] S.-C. Jung, M.-S. Kim, and Y. Yang. A reconfigurable carrier leakage canceler for uhf rfid reader front-ends. *Circuits and Systems I: Regular Papers*, *IEEE Transactions on*, 58(1):70–76, jan. 2011.
- [21] M. Knox. Single antenna full duplex communications using a common carrier. In Wireless and Microwave Technology Conference (WAMICON), 2012 IEEE 13th Annual, pages 1–6, 2012.
- [22] 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. In *Wireless Mesh Networks (WIMESH 2010), 2010 Fifth IEEE Workshop on*, pages 1–6, 2010.
- [23] B. Radunovic, D. Gunawardena, A. Proutiere, N. Singh, V. Balan, and P. Key. Efficiency and fairness in distributed wireless networks through self-interference cancellation and scheduling. Technical Report MSR-TR-2009-27, Microsoft Research, 2009.