Full Duplex Radios

Dinesh Bharadia
Stanford University
dineshb@stanford.edu

Emily McMilin
Stanford University
emcmilin@stanford.edu

Sachin Katti
Stanford University
skatti@stanford.edu

ABSTRACT

This paper presents the design and implementation of the first in-band full duplex WiFi radios that can simultaneously transmit and receive on the same channel using standard WiFi 802.11ac PHYs and achieves close to the theoretical doubling of throughput in all practical deployment scenarios. Our design uses a single antenna for simultaneous TX/RX (i.e., the same resources as a standard half duplex system). We also propose novel analog and digital cancellation techniques that cancel the self interference to the receiver noise floor, and therefore ensure that there is no degradation to the received signal. We prototype our design by building our own analog circuit boards and integrating them with a fully WiFi-PHY compatible 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.

Categories and Subject Descriptors

C.2.1 [**Computer Communication Networks**]: Network Architecture and Design—*Wireless communication*
General Terms: Algorithms, Design, Experimentation, Performance
Keywords: Full Duplex, Interference Cancellation, Non-linear Cancellation

1. INTRODUCTION

> "It is generally not possible for radios to receive and transmit on the same frequency band because of the interference that results."
> (Andrea Goldsmith, Wireless Communications [8])

The above quote captures a long-held assumption in wireless system design that radios have to operate in half duplex mode (i.e., either transmit or receive but not both simultaneously) on the same channel. Recent work has attempted to invalidate this assumption. Researchers at Stanford [11, 3], Rice [7, 6] and several other groups in industry and academia [14, 1] have proposed various designs to build in-band full-duplex radios. Full duplex, if possible, has tremendous implications for network design, not least of which is the fact that cellular networks could cut their spectrum needs by half. For example, LTE uses equal width separate uplink and downlink channels to enable radios to achieve full duplex. With an in-band full-duplex system we could use a single channel to get the same performance. Consequently, the problem has attracted significant attention, both from industry and academia and has spurred significant follow-up work.

To achieve full duplex, a radio has to completely cancel the significant self-interference that results from its own transmission to the received signal. Since WiFi signals are transmitted at 20dBm (100mW) average power, and the noise floor is around -90dBm, the transmit self-interference has to be canceled by $20\text{dBm}-(-90\text{dBm}) =$ 110dB to reduce it to the same level as the noise floor and render it negligible. If self-interference is not completely canceled, any residual self-interference acts as noise to the received signal and reduces SNR and consequently throughput. For example, if the received signal's SNR without full duplex is 25dB but is reduced to 5dB due to 20dB residual self-interference, then the throughput with full duplex is that achieved using two 5dB SNR links. This is significantly worse than using the original half duplex link with 25dB SNR and it is better to turn off full duplex in this case. To sum up, the amount of self-interference cancellation dictates overall throughput and is a figure of merit for any full-duplex design.

Prior designs have made significant progress on the self-interference cancellation problem [11, 5, 3]. However the best performing prior designs can at best provide 85dB of cancellation, which still leaves about 25dB of residual self-interference and therefore reduces the SNR of each direction of the full duplex link by 25dB. A calculation similar to the previous paragraph's shows that to see throughput benefits with these full-duplex designs, the half-duplex SNR of the link has to be extremely high (45dB or higher). In terms of range, the two nodes would have to be closer than 5m to see such high SNRs. Outside this range, it is better to turn off full duplex and use the traditional half duplex mode. To be fair however, these designs were intended for low-power, narrow-band, fixed rate protocols such as Zigbee where 85dB of self-interference cancellation is sufficient for full duplex. WiFi is far more demanding both in terms of bandwidth as well as cancellation.

Prior designs also need to have at least two antennas [11, 5] in place of the one used by half duplex systems (one each for transmit and receive and possibly more [3]). However, with two or more antennas, the argument for full duplex becomes weaker since the same doubling of capacity could be obtained by using the two antennas as MIMO antennas to spatially multiplex two independent packets in half duplex mode instead of using them for full duplex.

In this paper, we present the design and implementation of a full duplex WiFi radio that uses a *single antenna* [1] and delivers close

SIGCOMM'13, August 12–16, 2013, Hong Kong, China.
Copyright 2013 ACM 978-1-4503-2056-6/13/08 ...$15.00.

[1] Picasso [10] uses a single antenna, but it only allows the radio to simultaneously transmit and receive on *different* adjacent channels. Hence it fails to address the much harder problem of simultaneous TX/RX on the same channel. Our system does address this challenge, and offers novel and higher performance analog and digital cancellation techniques compared to Picasso.

to the theoretical doubling of throughput under all link SNR and distance ranges. Our key technical contributions are novel self-interference cancellation circuits and algorithms that provide the required 110*dB of self interference cancellation* for standard WiFi signals and thus eliminate all self interference to the noise floor. Our design is wideband: it works with the highest bandwidths (80MHz) and data rates used by the latest 802.11ac PHY in the 2.4GHz spectrum. We also experimentally demonstrate a complete full-duplex communication link which uses the full WiFi PHY (OFDM, constellations up to 256QAM and all the channel coding rates) and achieves close to the theoretically expected doubling of throughput. To the best of our knowledge, this is the first working implementation of a complete WiFi PHY single-antenna full-duplex link.

The reader might be wondering why full duplex is hard to realize. After all, as the sender knows the signal being transmitted, subtracting it should be relatively simple to implement. One of the key insight in this work is that in fact the *radio does not know what it is transmitting*. What it does know is the clean digital representation of the signal in baseband. However, once the signal is converted to analog and up-converted to the right carrier frequency and transmitted, the transmitted signal looks quite different from its baseband incarnation. The numerous analog components in the radio TX chain distort the signal in both linear and non-linear ways (analog circuits will create cubic and higher order components of the signal for example), add their own noise (e.g., power amplifiers add transmitter noise), are slightly inaccurate (e.g., your oscillator is tuned slightly off 2.45GHz), or delay it by different amounts at different frequencies and so on. In effect the transmitted signal is a complicated nonlinear function of the ideal transmitted signal along with unknown noise. Unsurprisingly, naively subtracting a “known” baseband version of the transmit signal without accounting for all these analog distortions does not work. As we will show in Sec 5 prior designs fail to account for these distortions and hence are limited to at best 85dB of cancellation.

This paper makes two key contributions over all prior work in this space. First, we design dynamic algorithms to estimate the distortions introduced by analog circuits and accurately model the actual self-interference being experienced by the received signal. Second, we design a novel programmable analog cancellation circuit using off-the-shelf components that allows us to implement the above algorithm in “analog” and dynamically cancel the self-interference. Such analog cancellation prevents receiver saturation from strong self-interference and allows us to use commodity radios. However, the analog cancellation stage does not completely cancel the self-interference. We complement it with a novel digital cancellation algorithm and implementation that cancels any remaining self-interference. Our digital cancellation algorithm differs from all prior work because it not only models the linear distortions, but also non-linear effects and other special effects such as oscillator noise. Thus, overall we use a hybrid analog-digital design that successfully models all linear and non-linear distortions as well as transmitter noise.

We implement our design via a combination of circuit designs and software implementations. Our analog cancellation is implemented on a PCB that we designed and populated using off-the-shelf components. We integrate our board with an off-the-shelf antenna and software radio transceiver [16, 15] based on test equipment from Rohde-Schwarz (RS) as well as on commodity WARP radios. We also implement our digital cancellation algorithms as well as a fully WiFi compliant PHY layer based on OFDM, supporting constellations up to the standard required 256QAM and all the channel coding rates. We deployed and evaluated our system in an indoor and noisy office environment in the 2.4GHz ISM band, operating the WiFi PHY over the 80MHz bandwidth on RS radios, and over the 20MHz bandwidth using WARP radios.

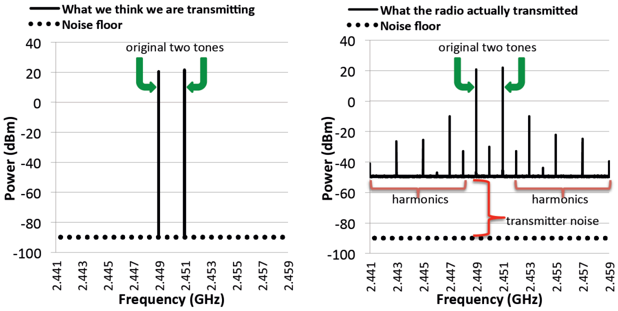

Figure 1: What we think we are transmitting in digital on the left side, and what the radio actually transmitted on the right side. The actual transmitted signal differs significantly from the two tones generated in digital baseband. Note transmitter noise and harmonics are generated in addition to the two main transmitter tones.

Our experiments demonstrate that our design delivers on the promise of full duplex. Under typical indoor deployment scenarios, our system delivers a median throughput gain of 87% in practice with WiFi radios which is close to the theoretically expected $2\times$. Looking into the cancellation itself, we show that our design consistently delivers the required 110dB of cancellation in a dense indoor office environment for both the RS 80MHz radios as well as the commodity 20MHz WARP radios. The system is robust to environmental changes, reflections, and can handle all the different constellations used in WiFi. We compare against the best known prior full duplex approaches [11, 7] and show experimentally that they can at best deliver 85dB of cancellation and therefore reduce the SNR of the received signal by at least 25dB.

2. THE PROBLEM

Full duplex, in theory, should be simple to accomplish. After all, we know the signal we are transmitting and we are only designing circuits and algorithms to subtract it from the received signal. The intuition follows from the conventional abstraction that the analog radio (also known as the RF front-end) is a black-box that takes the digital baseband signal, converts it to analog, up-converts it to the carrier frequency, scales it to the right power and sends it. In other words, the assumption has been that the radio preserves the original baseband signal except for power scaling and frequency shifting. In practice this abstraction turns out to be incorrect. Radios in fact significantly distort the signal being transmitted, relative to the digital baseband representation.

To demonstrate the distortions, we use the following experiment throughout this section. We take a software radio transceiver [16, 15] and send the following signal: two tones at 2.449GHz and 2.451GHz. In other words, we are sending an extremely simple signal, two sine waves with frequencies 1MHz away from the carrier frequency of 2.45GHz. We do this by creating a digital baseband signal with samples of the sine waves at -1MHz and 1MHz which the radio then up-converts to 2.45GHz and amplifies to 20dBm average transmit power (the power used by WiFi radios). We then compare the signal output of the antenna to what we would ideally expect if the radio did not introduce any distortions. This experiment serves as some sort of lower bound on the quality of radios. If radios cannot transmit even this simplest of signals without distortion then more complex signals such as WiFi are likely to be significantly distorted. Fig. 1 plots the ideal and actual transmitted signals’ spectra that resulted from our experimental set-up (we ensured that this was a clean environment with no other interference present in the environment at the time of the experiment).

Ideally, we expect to see only two tones at 2.451GHz and 2.449GHz as shown on the left side of Fig. 1. However in the transmitted signal, whose spectrum is plotted on the right side of Fig. 1, we can easily see that there are several other distortions present in addition to the two main tones that were transmitted. The main components in self-interference can be classified into three major categories:

1. **Linear Components**: This corresponds to the two main tones themselves which are attenuated and could consist of reflections from the environment. These are linear components because the received distortion can be written as a linear combination of different delayed copies of the original two tones.
2. **Non-Linear Components**: These components are created because radio circuits can take in an input signal x and create outputs that contain *non-linear cubic and higher order terms such as* x^3, x^5. These higher order signal terms have significant frequency content at frequencies close to the transmitted frequencies, which directly correspond to all the other harmonics we see on the right side of Fig. 1. Harmonics, as the name suggests, are signal distortions which occur at equally spaced frequency intervals from the transmitted frequencies. As the right side of Fig. 1 shows, we see spikes at frequencies 2.447GHz and 2.453GHz, that are spaced 2MHz apart from the two transmitted tones 2.451GHz and 2.449GHz, on either side.
3. **Transmitter Noise**: The general increase we see in the base signal level which we can clearly see on the sides of the two main tones is noise from the radio transmitter. A radio will of course always have noise, which works out to a noise power level of -90dBm [15]). But as we can see, the power at the side-bands is significantly higher, on the level of -50dBm, or 40dB higher than the receiver noise floor. This extra noise is being generated from high power components in the radio transmitter such as power amplifiers. In the radio literature this is referred to as broadband noise [12]. Further radios have phase noise generated by local oscillators (LO), which is typically of level of -40dBm, or 50dB above (not seen in the Fig. 1 because its hidden under the main signal component).

2.1 Requirements for Full Duplex Designs

The above analysis suggests that any in-band full duplex system has to be able to cancel all the above distortions in addition to the main signal component itself, since all of these are within the frequency band we are transmitting and receiving on and act as strong self-interference to the received signal. In this section, we discuss how strong each of these components are for typical transceivers, and what are the requirements for full duplex. We will state all self-interference power levels relative to the receiver noise floor. The reason is that to implement full duplex, we need to cancel any self-interference enough so that its power is reduced to the same level as the receiver noise floor. There is no point in canceling beyond that since we won't see any benefits — the received signal's SNR will then be dictated anyway by the receiver noise floor which cannot be canceled or reduced, just as it is today in half duplex radios.

We use similar experiments for OFDM-wideband signals to quantify the power levels of the different distortions, shown in the left side of Fig. 2. In a typical WiFi radio using 80MHz bandwidth, the receiver has a noise floor of -90dBm (1 picowatt). First, since the main signal component is being transmitted at 20dBm (100mW), self-interference from the *linear main component is* $20 - (-90) = 110dB$ above the receiver noise floor. Second, we observed experimentally that *the non-linear harmonics are at* $-10dBm$*, or* $80dB$ *above the receiver noise floor*. Finally, the *transmitter noise is at* $-40dBm$*, or* $50dB$ *above the receiver noise floor*. Note that these numbers are consistent with other RF measurement studies reported in the literature [21] for standard WiFi radios.

There are four takeaways from the above analysis:

- Any full duplex system needs to provide 110dB of **linear self-interference cancellation** to reduce self-interference to the receiver noise floor. This will ensure that the strongest component (the main signal) which is 110dB above the noise floor will be eliminated.
- A full duplex system has to reduce non-linear harmonic components that are 80dB above the noise floor, so any full duplex technique has to provide at least 80dB of **non-linear self-interference cancellation**.
- Transmitter noise is by definition noise and is random. In other words, we cannot infer it by any algorithm. Hence the only way to cancel transmitter noise is to get a copy of it where it is generated, i.e. in the analog domain and cancel it there. This implies any full duplex system has to have an analog cancellation component that provides at least 50dB of **analog noise cancellation** so that transmitter noise is reduced to below the receiver noise floor.
- A final constraint is that RX chains in radios get saturated if the input signal is beyond a particular level that is determined by their ADC resolution. Assuming a 12 bit ADC resolution typically found in commodity WiFi radios, we have a theoretical 72dB of dynamic range, which implies that the strongest signal level that can be input to the radio relative to the receiver noise floor is -90dBm$+72 = -18$dBm. However, in practice it is necessary to leave 2 bits worth of margin, i.e a 12 bit ADC should be used as if it is a 10 bit ADC to reduce quantization noise. So the maximum input signal level can be -90dBm$+60 = -30$dBm. Since in WiFi, the transmitted self-interference can be as high as 20dBm, a full duplex system needs to have an analog cancellation stage that provides 60dB of self-interference reduction (we keep a further 10dB margin for OFDM PAPR where instantaneously an OFDM signal's power level can rise 10dB above the average power).

To sum up, any full duplex design needs to provide 110dB of linear cancellation, 80dB of non-linear cancellation, and 60dB of analog cancellation.

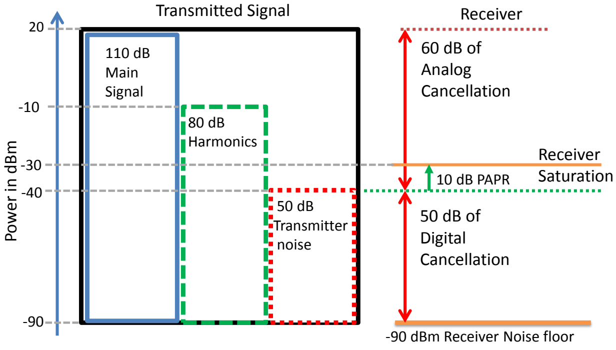

Figure 2: On the left hand side we see transmitted signal with sub-components. On the right hand side we see how this impacts the requirements of analog and digital cancellation.

2.2 Do Prior Full Duplex Techniques Satisfy these Requirements?

There are two state-of-the-art designs: ones which use an extra transmit chain to generate a cancellation signal in analog [6] and ones which tap the transmitted signal in analog for cancellation [11, 3]; both use a combination of analog and digital cancellation. Note that all these designs use at least two antennas for transmit and receive instead of the normal single antenna, and the antenna geometry ones use more than two.

Designs which use an extra transmitter chain report an overall total cancellation of 80dB (we have been able to reproduce their results experimentally). Of this, around 50dB is obtained in the analog domain by antenna separation and isolation between the TX and RX antennas of around 40cm (the designs also assume some form of metal shielding between the TX and RX antennas to achieve 50dB isolation). Note that this 50dB reduction applies to the entire signal, including linear and non-linear components as well as transmitter noise since it is pure analog signal attenuation. Next, these designs also use an extra transmit chain to inject an antidote signal [6, 9] that is supposed to cancel the self-interference in analog. However, the antidote signal only models linear self-interference components and does not model non-linear components. Further, it is incapable of modeling noise because by definition noise is random and cannot be modeled. Overall this extra cancellation stage provides another 30dB of linear self-interference cancellation in the best case. Thus, these designs provide 80dB of linear cancellation, 50dB of non-linear cancellation and 50dB of analog noise cancellation, falling short of the requirements by 30dB for the non-linear components. Hence if full duplex is enabled over links whose half duplex SNR is 30dB or lower, then no signal will be decoded. Further to see any throughput improvements with full duplex, the half duplex link SNR would have to be greater than 50dB.

The second design [11] gets a copy of the transmitted analog signal and uses a component called the balun (a transformer) in the analog domain to then create a perfectly inverted copy of the signal. The inverted signal is then connected to a circuit that adjusts the delay and attenuation of the inverted signal to match the self interference that is being received on the RX antenna from the TX antenna. We show experimentally in Sec. 5, that this achieves only 25dB of analog cancellation, consistent with the prior work's results. The cancellation is limited because this technique is very sensitive to and requires precise programmable delays with resolution as precise as 10picoseconds to exactly match the delay experienced by the self-interference from the TX to the RX antenna. Such programmable delays are extremely hard to build in practice, at best we could find programmable delays with resolution of $100-1000$picoseconds and these were in fact the ones used by the prior design [11]. Hence the cancellation circuit is never able to perfectly recreate the inverted self-interference signal and therefore cancellation is limited to 25dB in analog. However this design also uses two separate antennas separated by 20cm for TX and RX and achieves another 30dB in analog cancellation via antenna isolation. Hence a total of 55dB of self-interference reduction is obtained in analog, this cancellation applies to all the signal components (linear, non-linear and noise). The digital cancellation stage of this design also only models the linear main signal component, it does not model the non-linear harmonics that we discussed above. Thus we found that we obtain another 30dB of linear cancellation from digital in this design.

Overall, the second design provides 85dB of linear self-interference cancellation, 55dB of non-linear cancellation and 55dB of analog noise cancellation. Thus this design falls short of the requirements by 25dB (especially for the non-linear component). Hence if full duplex is enabled over links whose half duplex SNR is 25dB or lower, then no signal will be decoded. Further to see any throughput improvements with full duplex, the half duplex link SNR would have to be greater than 45dB.

3. OUR DESIGN

In this section we describe the design of our self-interference cancellation technique. Our design is a single antenna system (i.e. the same antenna is used to simultaneously transmit and receive), wideband (can handle the widest WiFi bandwidth of 80MHz as well as all the LTE bandwidths) and truly full duplex (cancels all self-interference to the receiver noise floor). The design is a hybrid, i.e., it has both analog and digital cancellation stages. Note that our hybrid cancellation architecture is not novel, similar architectures have been proposed in prior work [11, 20, 19]. The novelty of our work lies in the design of the cancellation circuits and algorithms, as well as their performance. To the best of our knowledge this is the first technique that achieves 110dB of cancellation and eliminates self-interference to the noise floor.

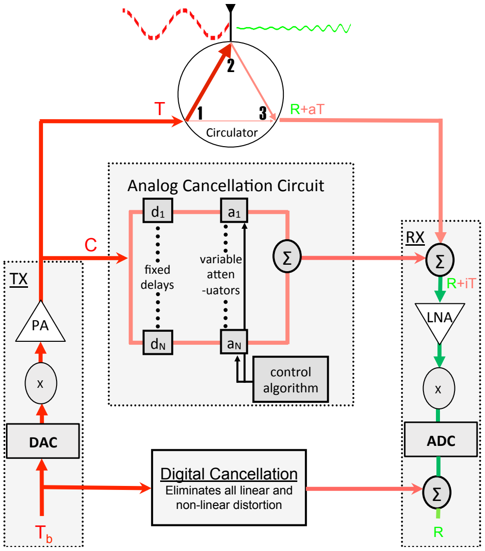

Figure 3: Full duplex radio block diagram. T_b is intended baseband signal we think we are transmitting, but in fact the transmit signal is T (red). The intended receive signal is R (green), however we see strong components of the red signal the RX side. Some of these red signals are undesirably leaked through the circulator. The analog cancellation circuit is trying to recreate a signal that matches the leaked interference signal for cancellation. The digital cancellation stage eliminates any residual self interference.

3.1 Analog Cancellation

We introduce a novel analog cancellation circuit and tuning algorithm that robustly provides at least 60dB of self-interference cancellation. Fig. 3 shows the high level design of the circuit and where it is placed in the radio architecture. A single antenna is connected to a circulator (at port 2), which is a 3 port device that provides limited isolation between port 1 and port 3 while letting signals pass through consecutive ports as seen in Fig. 3. The TX signal is fed through port 1, which routes it to the antenna connected to port 2, while the received signal from the antenna is passed from port 2 through to port 3. Circulator cannot completely isolate port 1 and port 3, so inevitably the TX signal leaks from port 1 to port 3 and causes interference to the received signal. From our experiments we find that the circulator only provides 15dB of isolation, i.e., the self-interference that is leaking to the RX circuit is reduced only by 15dB. To get to the noise floor, we still have to provide 95dB of cancellation, and at least 45 dB of that has to come in analog to ensure transmitter noise is sufficiently canceled and we do not saturate the receiver. We accomplish this using our novel analog cancellation circuit that we describe next. Note that when we report analog cancellation performance numbers, we include the 15dB of reduction we get from the circulator for simplicity of description.

Fig. 3 shows the design of our analog cancellation circuit. We tap the TX chain to obtain a small copy of the transmitted signal just before it goes to the circulator. This copy therefore includes the transmitter noise introduced by the TX chain. The copy of the signal is then passed through a circuit which consists of parallel fixed lines of varying delays (essentially wires of different lengths) and tunable attenuators. The lines are then collected back and added up, and this combined signal is then subtracted from the signal on the receive path. In effect, the circuit is providing us copies of the transmitted signal delayed by different fixed amounts and programmatically attenuated by different variable amounts. The key challenge is to pick the fixed delays, as well as to dynamically program the tunable attenuators appropriately so that the we maximize self-interference cancellation. Note that unlike prior work our design uses components that are all available off-the-shelf and is therefore easy to manufacture, we do not need sophisticated high resolution programmable delays that are hard to build like in prior work [11].

The design of our cancellation circuit is based on a novel insight: *we can view cancellation as a sampling and interpolation problem*. The actual self-interference signal has a particular delay and amplitude that depends on the delay d and attenuation a through the circulator. Our insight (the reason for which will become clear shortly) is that we should pick the fixed delays in our cancellation circuit such that they straddle the delay of the self-interference signal through the circulator. So if we have N fixed delay lines, $N/2$ of those lines should be placed at equidistant intervals all of which have delays that are less than the delay of the self-interference d, and we should do the same for the other half of the delays but greater than d. In practice it is hard to know the precise value of d since it is a function of how the circuit is put together, but we can always find the range over which it varies and place our fixed delays outside of that range on either side.

At this stage we have leading and lagging copies of the transmitted self-interference signal, how might we use them to approximate the actual self-interference itself at some intermediate instant? If we take a step back, this is essentially an interpolation problem, similar to Nyquist digital sampling. In Nyquist digital sampling, we have discrete samples of the signal at a time period equal to the inverse of the sampling frequency. The Nyquist theorem [13] tells us that sampling (at the Nyquist rate) does not lose information, in other words we can always reconstruct the signal at any instant as a weighted linear combination of samples taken before and after the instant at which we want to recreate. The weights of the linear combination can be determined by using a standard algorithm called *sinc interpolation*. The basic idea is that you overlay sinc pulses at each sampling time instant and calculate the value of the sinc pulse at the time instant t where you wish to recreate the signal. This value gives the weight you should apply to this sample when you take the linear combination for reconstruction. We repeat this algorithm for every sample to determine the corresponding weight to apply to it. The value of the signal at time t is then given by the linear combination of all the samples with weights calculated by the sinc trick discussed above.

Our analog cancellation circuit is in effect implementing the same trick, at every instant we have copies of the signal at different equally spaced delays just like in digital sampling. The programmable attenuators essentially function as the weights we need to apply in the linear combination for reconstruction. Similar to digital sampling, we need to estimate the self interference at an instant d that lies somewhere in between these fixed delays $d_1, \ldots, d_N$ as shown in Fig. 4. To do so, the weights for each sample, i.e., the value of the attenuator that we need to set on each line i is equal to the value of the sinc pulse centered at the fixed delay d_i at instant d. If we adjust the attenuators for each delay line to those values, then we will be able to perfectly reconstruct the self interference and cancel it from the receive path. Fig. 4 shows this algorithm visually in action.

In practice however, there is an important difference with digital sampling. In digital, we can take linear combinations of a very large number of samples since memory is essentially free. To do that in analog we would need a correspondingly large number of delay lines. In practice, this is not possible due to a variety of reasons, ranging from space limitations to power consumption to electromagnetic radiations. Our key insight is that in interpolation, the samples that matter most are the ones that are closest to the instant t at which the signal is being reconstructed. Intuitively, the value of a signal at a much further/before time than t should not affect the value of the signal at t. This is reflected in the fact that the weights in the linear combination for these further out samples are nearly zero. This allows our analog circuit to therefore use a small number of delay lines and still approximate the self interference fairly well. We show in Sec. 5 that sixteen delay lines are sufficient to approximate the self interference signal leaking through the circulator. Further, we will show in Sec. 5 that our analog cancellation delivers at least 60dB cancellation comfortably exceeding the requirements we developed in Sec. 2.1.

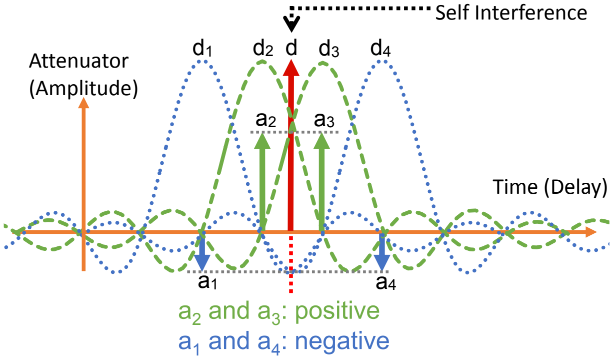

Figure 4: This figure shows how we can recreate the self interference signal which is located at instant d, positioned between the fixed delay lines d$_i$. The value of the attenuator a$_i$ for delay d$_i$ is given by the value taken by the sinc centered at d$_i$ at instant d.

3.2 Digital Cancellation

The goal of digital cancellation is to clean out any remaining residual self-interference. Assuming that analog cancellation provides 60dB, digital cancellation has to cancel the linear main signal component by another 50dB and non-linear components by another 20dB. We address each of these components separately.

3.2.1 Canceling Linear Components

The first part of digital cancellation eliminates the residual linear components of the self-interference. This consists of the main transmitted signal that is leaking over through the circulator after analog cancellation, as well as any delayed reflections of this signal from the environment. The reflections are also delayed and attenuated by different unknown amounts.

The basic idea is that this part of the self-interference can be modeled as a linear and *non-causal* function of the transmitted signal, as we know it in digital (recalling that we know the baseband IQ samples of the transmitted packet). The non-causal bit is important. Since we know the samples of the entire packet that was transmitted, we can use samples from the future to estimate the self-interference at the current instant. In other words, the received sample $y[n]$ at any instant can be modeled as a linear combination of up to k samples of the known transmitted signal $x[n]$ before and after the instant n. The parameter k is empirically chosen and is a function of the amount of memory in the channel. So we can write the equation as:

$$y[n] = x[n-k]h[k] + x[n-k+1]h[k-1] + \ldots + x[n+k-1]h[-k+1] + w[n]$$

where $h[k], \ldots, h[-k+1]$ represents the attenuations applied by the channel to the transmitted function, and $w[n]$ is the receiver noise floor.

How can we estimate the coefficients $h[n]$? We leverage the fact that most wireless transmissions have known packet preambles (e.g. WiFi uses a preamble of two known OFDM symbols at the start of the packet). Let the samples representing the preamble be $x_{pr}[n]$. Let the receive samples corresponding to the preamble be $y[0], \ldots, y[n]$.

Then the above channel equations can be written specifically for the preamble as:

$$y = Ah + w$$

where A is Toeplitz matrix of $x_{pr}[n]$.

$$A = \left(\begin{array}{ccccc} x_{pr}(-k) & \dots & x_{pr}(0) & \dots & x_{pr}(k-1) \\ \dots & \dots & \dots & \dots & \dots \\ x_{pr}(n-k) & \dots & x_{pr}(n) & \dots & x_{pr}(n+k-1) \end{array} \right).$$

Our goal is to find a maximum likelihood estimate of the vector h, i.e.,

$$minimize \ \ ||y - Ah||_2^2$$

Note that the matrix A is known in advance since we know the values of the preamble samples. Hence it can be *pre-computed.* Additionally, we know from prior work [2] that the coefficients for the above problem can be computed by multiplying by the ith received sample of the preamble, as the samples arrive serially as follows:

$$h = \sum (y_i a_i^{\dagger})$$

where $a_i^{\dagger}$, is the ith column of pseudo inverse of A matrix. Thus our estimation algorithm computes the linear distortions that the transmitted main signal has gone through for every packet, and is capable of dynamically adapting to the environment.

3.2.2 Canceling Non-Linear Components

The second task for digital cancellation is to eliminate the residual non-linear components whose power is around 20dB after being reduced by 60dB due to analog cancellation. However, it is quite hard to guess the exact non-linear function that a radio might be applying to the baseband transmitted signal. Instead, we use a general model to approximate the non-linear function using Taylor series expansion (as this is a standard way to model non-linear functions)[4]. So the signal that is being transmitted can be written as:

$$y(t) = \sum_m a_m x_p(t)^m$$

where $x_p(t)$ is the ideal passband analog signal for the digital representation of $x(n)$ that we know.

The above general model contains a lot of terms, but the only ones that matter for full duplex are terms which have non-zero frequency content in the band of interest. A little bit of analysis for passband signals (taking the Fourier transform) of the equation above reveals that the only terms with non-zero energy in the frequency band of interest are the odd order terms (i.e., the terms containing $x_p(t)$, $x_p(t)^3$, $x_p(t)^5$ and so on), so we can safely ignore the even order terms. The first term that is the linear component, i.e., the terms for $x_p(t)$ is of course the one corresponding to the main signal and is estimated and canceled using the algorithm discussed in the previous section. In this section, we focus only on the higher-order odd power terms. We can therefore reduce the above model and write it in the digital baseband domain as:

$$y(n) = \sum_{m \in \ odd \ terms, n=-k,\dots,k} x(n)(|x(n)|)^{m-1} * h_m(n)$$

where $h_m[n]$ is the weight for the term which raises the signal to order m and is the variable that needs to be estimated for cancellation, and k is the number of samples in the past and future which significantly influence the value of the signal at instant n.

To estimate these coefficients, we can use the same WiFi preamble. The WiFi preamble is two OFDM symbols long of length 8μs, and assuming a sampling rate of 160MHz, it consists of a total of 1280 digital samples at the Nyquist sampling rate. However, if we look at the above equation, the number of variables $h_m(n)$ that we need to compute is a function of $2k$ (i.e., how far in the past and future is the current self-interference signal influenced by) and the highest value of m that exhibit strength greater than the receiver noise floor. A naive model assuming that just the $1, 3, 5, 7, 9, 11$th order terms matter, and that upto 128 samples from both the future and the past influence the self-interference signal at any instant [2] would require us to estimate $128 * 2 * 6 = 1536$ variables using 1280 equations. Clearly, this is under-determined system, would increase the noise floor significantly.

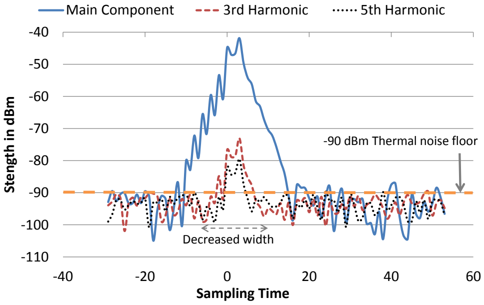

Figure 5: Signal strength of various harmonics that make up the transmitted signal. Note that higher order harmonics are much weaker relative to main component and therefore any reflections of these harmonics have to be quite closely spaced in time for them to be stronger than the receiver noise floor.

In practice we found empirically that many of these variables do not matter, that is their value is zero typically. The reason is that higher order terms have correspondingly lower power since they are created by the mixing of multiple lower order terms and each mixing reduces power. So the 7^{th} order term has lower power than the 5^{th} order term which has lower power than the 3^{rd} order term. Fig. 5 shows a plot of the strength of of the main signal and higher order non-linear terms relative to the receiver noise floor. As we can see higher order terms have weaker strength relative to the main signal, and consequently their multipath components also decay quickly below the receiver noise floor. In other words, far fewer than 128 samples from the past and future impact the value of the self interference harmonic component at this instant. We find empirically that for indoor WiFi systems, across all the non-linear higher orders, a total of only 224 such variables are all that we need to estimate which we can easily accomplish using the WiFi preamble (over-determined system of linear equation). Hence our digital cancellation algorithm calculates all these coefficients using the WiFi preamble and applies them to recreate the harmonics and cancel them. The method for estimating the coefficients is the same as the one used in the linear digital cancellation step described by Eq. 3.2.1, but the matrix A is formed using the higher order odd powers of the preamble samples.

3.2.3 Complexity

The complexity of digital cancellation is the same as solving 1280 (say W, width of preamble in general) linear equations with 224 unknowns. Further the matrix that forms the linear equations is known in advance (this is the known preamble trick as discussed above). Hence the pseudo-inverse of this matrix can be pre-computed and stored. Thus the complexity of digital cancellation reduces to $O(W)$ multiplications. The design is therefore relatively simple to implement and can be efficiently realized in hardware.

3.3 Dynamic Adaptation of Analog Cancellation

To provide a robust full duplex link, we need to ensure that sufficient cancellation is maintained to reduce self interference to the

[2] The number of samples required is a function of the amount of multipath, the higher the mutlipath, the higher the number of samples in the past and future that matter but 128 is the number suggested by the WiFi standard and is equal to the length of the WiFi OFDM Cyclic Prefix

noise floor, even as things such as environment, transmit power, temperature and other such parameters change. These changes would clearly reduce the cancellation achieved by any static configuration, since they change the distortions that are imposed by the self interference. Digital cancellation can cope since it essentially estimates these distortions on a per-packet basis, however analog cancellation might be degraded and hence performance might be worsened. In this section, we describe how we can quickly tune the analog circuit to provide the required amount of cancellation (60dB at least).

The goal of tuning is to pick the attenuation values $a_1, \ldots, a_N$ such that self-interference is minimized. More formally,

$$\min_{a_1,\ldots,a_N} \quad (y(t) - \sum_{i=1}^{N} a_i c(t-d_i))^2$$

where $c(t)$ is the reference signal that is tapped from the transmit path, $y(t)$ is the self interference, $d_1, \ldots, d_N$ are delays associated with the taps as shown in Fig. 3.

A simple and obvious technique to solve the above problem in practice is a iterative gradient descent algorithm, which other prior works in full duplex have also used to tune their own analog cancellation [11]. However, we found that this algorithm is extremely slow (requires nearly 40ms) because of the larger number of variables (16) that need to be estimated in our design unlike prior work. That's an unacceptable overhead, since we found empirically that we need to re-tune analog cancellation once every 100ms on average in our setup. So taking 40ms to tune implies a 40% overhead.

Our key contribution here is an approach that solves the tuning problem in the frequency domain. The idea is that the self interference $y(t)$ can be modeled in the frequency domain as a function of the tapped signal $c(t)$ as

$$\mathbf{Y}(f) = \mathbf{H}(f)\mathbf{C}(f)$$

where $\mathbf{H}(f)$ is the frequency domain representation of the distortion introduced by the circulator, antenna and the environment and $C(f)$ is the frequency domain representation of the tapped signal. Recall that the tapped signal is essentially a scaled replica of the transmitted signal input to the circulator, hence the above equation can be written in terms of the tapped signal. This frequency response $\mathbf{H}(f)$ is easier to measure, it is essentially an FFT of the self interference channel which can be measured using the WiFi preamble. In fact, standard OFDM is doing exactly this, it is estimating the frequency domain channel using the preamble and pilot symbols.

The goal of the optimization problem then is to pick the attenuator values such that the overall frequency domain response of the cancellation circuit approximates $\mathbf{H}(f)$ as closely as possible. So the above optimization problem can be restated as

$$\min_{a_1,\ldots,a_N} \quad (\mathbf{H}(f) - \sum_{i=1}^{N} \mathbf{H_i^{a_i}}(f))^2 \qquad (1)$$

where, $\mathbf{H_i^{a_i}}(f)$ is the frequency response for delay line i for attenuation setting of a_i.

How might we solve this problem? The problem is two fold. First, we have to find the frequency response of each delay line of the cancellation circuit for every attenuation value, i.e., $\mathbf{H_i^{a_i}}(f)$. Second, once we have the frequency response of the self-interference channel $\mathbf{H}(f)$, we need to search on the space of possible attenuation values for every delay line(attenuator), to come up with best possible solution to the optimization problem. Each delay line can take 128 different attenuation values, and there are 16 delay lines, so in total we have $128^{16} = 2^{112}$ values, a computationally expensive search.

Modeling the frequency response of delay lines $\mathbf{H_i^{a_i}}(f)$: Measuring the frequency response of individual delay line is impossible — The entire circuit is well connected, thus isolating individual delay line is impossible. Our key observation, is if we can measure the frequency response of a delay line at one attenuation value, then the datasheet of the attenuator provide measurements called S parameters (specifically frequency response measurements between different ports of a device) that can be used to extrapolate the frequency response of the delay line for all attenuation values. The S parameter data provides the relative change in frequency response with changing attenuator value. To calculate the frequency response at this initial point, we use the following trick. We set the attenuators for all the lines to their highest attenuation setting, except the one being measured. The idea is to essentially emulate a board where none of the delay lines, except the one being measured, let any signal through. The highest attenuation value approximates that setting but doesn't fully accomplish that, hence we apply a second least squares fit to find a more accurate response (collecting more data for different attenuation's for this delay line, keeping the rest all others at highest attenuation setting). Then, the frequency response of this delay line for all 128 attenuator values can be calculated. We repeat this process for all the delay lines in the circuit. Note that all of this has to be done once and can be stored, since this frequency response of the delay line and attenuation is independent of the environment or other such changing parameters.

Optimization Algorithm : Now to actually find the attenuation settings in real time to optimize the cancellation, we use the following algorithm.

1. Measure the frequency response of the self interference $\mathbf{H}(f)$ using the WiFi preamble. This is relatively simple since we have two OFDM symbols and as part of the baseband decoding we can perform an FFT to measure the frequency response.
2. Solve the frequency domain integer linear optimization problem posed in Eq. 1 by relaxing it to a linear program and then use random rounding to find a solution for attenuator settings, which achieves required cancellation of 60dB. The intuition behind the algorithm is that it reduces the search space of attenuator values to a polynomial set compared to the exponential search space. This is due to the fact that we are looking for a point which provides required cancellation, instead of the optimal point (achieving optimal point is a NP hard problem). Note all the aforementioned calculations are offline and are implemented using the frequency response model. Essentially the model is used for looking up the frequency response of the circuit, for any combination of attenuator values. This offline algorithm implementation is therefore extremely fast — a non-optimized C++ implementation takes less than 1μ sec to converge.

In practice, we find that offline solution calculated above might yield a point that provides an analog cancellation of $45-50$dB due to manufacturing variation of attenuator (the S parameter data provided is accurate to 2%, thus every attenuator has its own response different from the provided standard data). To further improve the cancellation, we use an additional gradient descent step. Typically, gradient descent takes several hundreds of iterations, however here since we are starting the descent from a much more accurate starting point, the gradient descent converges to the required point in 10-12 iterations. So in the worst case, we show experimentally that analog cancellation tuning can take around 900-1000μs. Assuming we have to do such tuning once every $100ms$ (which is what we needed in our testbed), that represents less than 1% overhead for tuning.

4. IMPLEMENTATION

Fig. 6 shows the prototype of a single full duplex radio. To implement it we designed our own analog circuit boards for cancellation and integrated them with existing software radios. We also implemented the digital cancellation algorithms in the software radio. Below we discuss the different pieces.

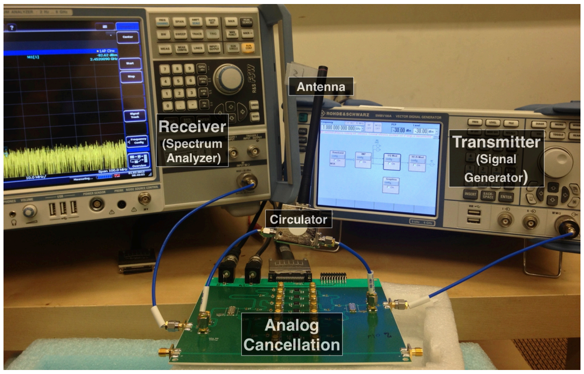

Figure 6: Experimental set-up of our full duplex transceiver

Analog Cancellation Board: The analog cancellation board is a 10×10 cm PCB board designed and built using Rogers 4350 material. The fixed delay lines are implemented using micro-strip trace lines of different fixed lengths. The attenuators are programmable step Peregrine PE43703 [17] attenuators which can be programmed in steps of 0.25dB from 0 to 31.5dB for a total of 128 different values.

Radio Transceiver and Baseband: Our goal was to design and implement a full duplex system that was capable of supporting the latest WiFi protocol 802.11ac with least 80MHz of bandwidth in the 2.4GHz range and 20dBm average TX power. Unfortunately none of the widely used software radios, such as USRPs or WARPs, support such high performance; at best they are capable of supporting 20MHz bandwidths. For that reason, we prototyped our design using radio test equipment from Rohde and Schwarz. For our transmitter, we used a SMBV 100A vector signal generator [16] to send our desired WiFi signals. Since the SMBV is not capable of generating 20dBm power, we use an external power amplifier [18]. For the receiver, we use the RS spectrum analyzer [15].

A practical concern is how to kick-start re-tuning of analog cancellation. Specifically if analog cancellation drops below a threshold, then the receiver might get saturated and the feedback needed to tune is distorted. To tackle this we implemented an AGC via a digital tunable attenuator in front of the LNA. The idea is that if the baseband detects that the receiver is getting saturated, then it programs the attenuator to a large value which brigs the whole signal down to within the dynamic range. After cancellation is tuned, this attenuation is turned off. The FSW is capable of receiving 100MHz signals at 2.45GHz, down-converting and digitizing it to baseband, and then giving us access to the raw IQ samples, which we can then freely process using our own baseband algorithms. The noise floor of this receiver is -90dBm at 100MHz bandwidth. It has a 16 bit ADC capable of sampling a 100MHz signal, however to ensure that we are only using resources found in commodity WiFi cards we configure the ADC to only use 12 bits of resolution.

The IQ samples are transported via ethernet to a host PC, on which we implement our cancellation and baseband software. We implemented a full WiFi-OFDM PHY that can be configured to operate over all the standard WiFi bandwidths (20MHz, 40MHz, and 80MHz). We support all the WiFi constellations from BPSK to 64-QAM for 40MHz, and 256 QAM for 80MHz. We also support all the channel codes with coding rates (1/2, 2/3, 3/4 and 5/6 convolutional coding). Finally we also implement our digital cancellation algorithm in software on the same host PC.

However to show that our design is general and does not benefit from using expensive test equipment, we also develop an implementation using standard WARP radios. Due to their radio limitations, these results will be for 20MHz signals which is the widest that the WARP supports.

5. EVALUATION

In this section we show experimentally that our design delivers a complete full duplex WiFi PHY link. We prove the claim in two stages. First, we show that our design provides the 110dB of self interference cancellation required to reduce interference to the noise floor. We also show experimentally that the received signal is received with almost no distortion in full duplex mode (the SNR of the received signal is reduced by less than 1dB on average), and that these results are consistent across a wide variety of bandwidths, constellations, transmit powers and so on. Second, we take the next step and design a working full duplex communication WiFi link. We show experimentally that it delivers close to the theoretical doubling of throughput expected from full duplex.

We start with an experimental evaluation of the cancellation system. We define two metrics we use throughout this section:

- *Increase in noise floor*: This is the residual interference present after the cancellation of self interference which manifests itself as an increase in the noise floor for the received signal. This number is calculated relative to the receiver noise floor of the radio of -90dBm. For example, if after cancellation we see a signal energy of -88dBm, it would imply that we increased the noise floor by 2dB.
- *SNR loss*: This is the decrease in SNR experienced by the received signal when the radio is in full duplex mode due to any residual self interference left after cancellation. To compute this we first measure the SNR of the received signal when the radio is in half duplex mode and there is no self interference, and then with full duplex mode. The difference between these two measured SNRs is the SNR loss.

We compare our design against two state-of-the-art full duplex systems presented in prior work.

- *Balun Cancellation*: This design [11] uses a balun transformer to invert a copy of the transmitted signal, adjust its delay and attenuation using programmable attenuators and delay lines and cancel it. The design also uses two antennas separated by 20cm one each for TX and RX which automatically provides 30dB of self interference reduction. We implement this design and optimize it to produce the best performance.
- *Rice Design*: This design uses an extra transmit chain in addition to the main transmit chain. The extra chain generates a cancellation signal that is combined with the signal on the receive chain to cancel self interference. This design also uses two antennas and to make a fair comparison we use a 20cm separation as the balun based design. However we also provide results with 40cm separation since that was the value used in the prior work. We implement this design by using an extra signal generator as an extra transmit chain for cancellation.

Note that our design uses a single antenna and therefore does not have the benefit of the 30dB of self interference reduction that prior schemes enjoy from using two physically separate antennas.

5.1 Can we cancel all of the self interference?

The first claim we made in this paper is that our design is capable of canceling all of the self interference for the latest operational WiFi protocols. To investigate this assumption, we experimentally test if we can fully cancel a 80MHz WiFi 802.11ac signal upto a max transmit power of 20dBm (all of which are the standard parameters used by WiFi APs), as well as the smaller bandwidths of 40MHz and 20MHz. We conduct the experiment by placing our full duplex radio in different locations in our building. Further we increase the transmit power from 4dBm to 20dBm (typical transmit power range). For

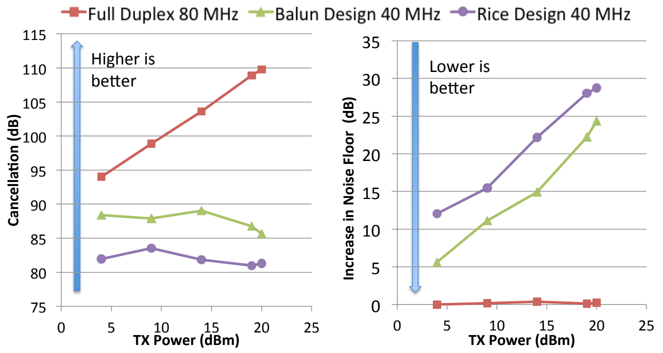

Figure 7: Cancellation and increase in noise floor vs TX power for different cancellation techniques with transmission of WiFi 802.11 signal. Our full duplex system can cancel to the noise floor standard WiFi signals of 20dBm at highest WiFi bandwidth of 80MHz, while prior techniques still leave 25dB of self interference residue, even for the narrower bandwidth of 40MHz.

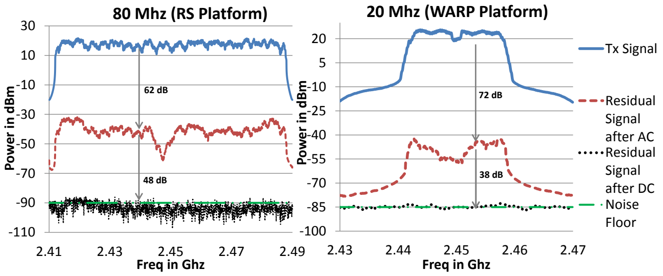

Figure 8: Spectrum Response for our cancellation with the Rohde-Schwarz (RS) radios and the WARP radios. The figure shows the amount of cancellation achieved by different stages of our design. It also shows that our design provides the same 110dB of cancellation even with WARP radios.

each TX power and location (in total 100) we conduct 20 runs and compute the average cancellation across those runs and locations. The goal is to show that we can cancel to the noise floor for a variety of transmit powers up to and including the max average TX power of 20dBm. Fig. 7 plots the average cancellation as a function of TX power. It also plots the corresponding observed increase in noise floor on the other axis.

Fig. 7 shows that our design essentially cancels the entire self interference almost to the noise floor. In the standard case of 20dBm transmit power, the noise floor is increased by at most 1dB over the receiver noise floor. The amount of cancellation increases with increasing TX power, reaching the required 110dB for the 20dBm TX power. The takeaway is that as the TX power increases, self interference increases at the same rate and we need a correspondingly larger amount of cancellation, which our design provides.

PAPR: Note that these are average cancellation numbers, in practice our WiFi transmissions exhibit transient PAPR as high as 10dB, so the peak transmit power we see is around 30dBm. We do not report the specific numbers for these due to lack of space, but our cancellation system scales up and also cancels these temporary peaks in the self interference signal to the noise floor.

The prior balun and Rice designs however fare far worse. Further, since these designs perform very poorly at 80MHz, we only report their results for the smaller 40MHz WiFi bandwidth and 20dBm TX power. As we can see, these designs can at best provide 85dB and 80dB of cancellation respectively. In other words they increase the noise floor by 25dB and 30dB respectively. The reasons for this are the ones we discussed in Sec. 2.2, the inability to adequately cancel transmitter noise in analog and the inability to model nonlinear distortions produced by radios. To check if these designs could be made to work with larger antenna separation, we repeated the experiment with an antenna separation of 40cm instead of 20cm. We found that even with an impractical rough half meter separation in antennas, the noise floor increase is at least 20dB.

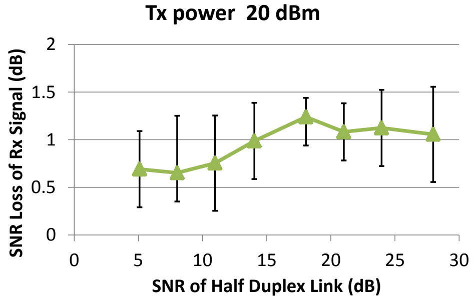

Figure 9: SNR loss vs half duplex SNR at fixed TX power = 20 dBm, constellation = 64 QAM, bandwidth = 80MHz with transmission of WiFi 802.11 signal. Our full duplex system ensures that the received signal suffers negligible SNR loss regardless of the SNR it was received at.

5.1.1 Does our design work with commodity radios?

We repeat the above experiment, but instead of the Rohde-Schwarz test equipment, we use off-the-shelf WARP radios in the setup. The goal is to show that our design can work with cheap commodity radios and does not depend on the precision of test equipment. Since the widest bandwidth that the WARP can support is 20MHz, we only report results for that bandwidth. Fig. 8 shows the spectrum plot of canceled signals after different stages of cancellation. For comparison, we also plot the spectrum plot of cancellation using the Rohde-Schwarz equipment.

As we can see, our cancellation completely eliminates self-interference even with commodity WARP radios. The WARP has a worse noise floor of -85dBm compared to the -90dBm of the RS equipment. Hence if we used 20dBm transmit power, then a slightly smaller 105dB of self-interference cancellation is required to eliminate it to the noise floor. However for consistency, for the WARP experiments we increase the transmit power to 25dBm to show that our design can still achieve 110dB of cancellation and eliminate self-interference to the noise floor.

5.1.2 SNR loss of the Received Signal in Full Duplex Mode

The previous section provided evidence for the amount of cancellation and increase in noise floor. However the experiments had only one radio transmitting. A natural question is how well does the system work when we are in true full duplex mode, i.e. the radio is transmitting and simultaneously receiving a signal. In this section, we evaluate the SNR loss for the received signal when operating in full duplex mode.

The experiment is conducted as follows. We setup two nodes capable of full duplex operation in our building. The two nodes first send 20 WiFi packets (with the following PHY parameters: 80MHz bandwidth, 20dBm TX power, 64QAM constellation) to each other one after the other, i.e. they take turns and operate in half duplex mode. They then send 20 WiFi packets to each other simultaneously, i.e. they operate in full duplex mode. For each run we measure the average SNR of the received packets across the 20 packets in half duplex mode, and then with full duplex mode. We then compute the SNR loss which is defined as the absolute difference between the average half duplex SNR and full duplex SNR measured above. We repeat the experiment at several different locations of the two nodes in our testbed. We plot the SNR loss as a function of the half duplex SNR in Fig, 9.

As Fig. 9 shows the SNR loss is uncorrelated with the half duplex SNR value and is almost identical to the increase in noise floor value we saw in the previous experiment. The takeaway is that self inter-

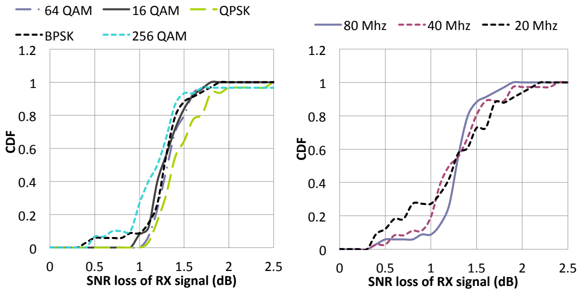

Figure 10: Shows CDF of SNR loss with changing bandwidths and constellations. Left: we see the SNR loss for different constellations with TX power = 20 dBm and bandwidth = 80MHz. Right: we see the SNR loss for different bandwidths (20 MHz, 40 MHz and 80 MHz) for TX power = 20 dBm and constellation = 64 QAM. Observe we can support all WiFi modulation schemes and bandwidths with low SNR loss.

ference cancellation is not impacted by the received signal's strength, whether it is weak or strong. Further, the SNR loss is typically around or less than 1dB which implies that even in full duplex mode the received signal should retain almost the same throughput as in clean half duplex mode.

5.2 Digging Deeper

5.2.1 *Impact of Constellation and Bandwidth*

We conduct two experiments. First we use the same setup as the SNR loss experiments and fix the bandwidth to 80MHz, but vary the constellation for the transmitted signal for the full duplex node from QPSK to the densest constellation in WiFi 256-QAM. Once again we calculate the SNR loss of the received signal across different measurements and locations from the half duplex node. In the second experiment we fix the constellation to 64-QAM but vary the bandwidth from 20 to 40 to 80MHz and once again calculate the SNR loss of the received signal. We repeat this experiment for different locations of the two nodes. Fig. 10 plot the CDFs of the SNR losses for different choices of constellations and bandwidth.

As the figures show, our design performs consistently well for all constellation choices and bandwidths. Our cancellation technique makes no assumptions about what constellation and other parameters the PHY is using: for us all of them are a self interference signal and hence the design is unaffected by constellation choice. Our design also works equally well for all the bandwidths used by 802.11ac in the 2.4GHz band. The reason is that our analog cancellation, as we will show in the next section, has sufficient flexibility to provide an almost flat wideband cancellation, while prior designs are extremely narrow-band and cancellation tapers off quickly with wider and wider bandwidths.

5.2.2 *Deconstructing Analog Cancellation*

In this section we dig deeper into the analog cancellation component of our design. The key parameter in our analog cancellation circuit board is the number of fixed delay lines as discussed in Sec. 3.1. We conduct an experiment to examine the impact of the number of such lines. However since these are circuit boards, we do not have the flexibility to vary the number of lines in increments of one. The granularity of our board design allows us to only test two configurations, one with 8 lines and one with 16 lines. We conduct the same self interference cancellation experiment as described in Sec. 5. We measure the signal after analog cancellation (without digital cancellation) and plot the frequency response of the canceled signal for the two cases in Fig. 11. The plot should be read as the power of the self-interference signal after analog cancellation as a function of the frequency.

As Fig. 11 shows, with 8 lines we can achieve 45dB of cancellation over 80MHz, while we can achieve 63dB of cancellation with 80MHz. The reason for the difference is the higher capability of 16 lines in canceling signal reflections in addition to the main self interference component that is leaking through the circulator. When the full duplex node is transmitting, the response from the circulator and antenna in the RX chain has two primary leakage components from the TX signal: one due to the direct leakage from the TX port of the circulator to the ("isolated") RX port of the circulator, and one due to reflections from impedance mismatch between the circulator and the antenna. Because these two components travel different paths in the circulator from TX port to RX port, they undergo different delays as deduced from time domain measurements. These delays are fixed and are a function of the particular circulator and antenna we choose to use. In our implementation we find the delay of the direct leakage component is 400 picoseconds, while the reflected component is centered around 1.4 nanoseconds. With 16 lines we have the capability to center the first 8 lines to have delays around 400 picoseconds, and the other 8 lines around 1.4 nanoseconds. We can then use the interpolation trick discussed in Sec. 3.1 to cancel both the direct and reflected self interference components precisely. As expected with 8 lines, our flexibility is reduced in terms of placing our delay lines around the actual delays experienced by the self interference and consequently cancellation is reduced.

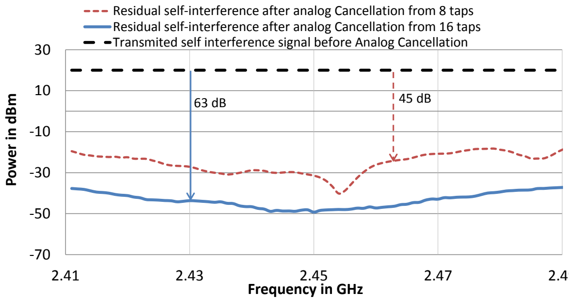

Figure 11: Frequency domain representation of self interference before analog cancellation and self interference after analog cancellation using 8 taps and 16 taps. Note that with 16 taps we can provide at least 63 dB of analog cancellation over the entire 80 MHz of bandwidth.

5.2.3 *Deconstructing Digital Cancellation*

After 62dB of analog cancellation, digital cancellation needs to clean up 48dB and 16dB of linear and non-linear self-interference components respectively. In this section, we deconstruct the amount of linear and non-linear cancellation achieved by our design. To conduct this experiment, we tune our analog cancellation circuit to provide 62dB of cancellation. We then progressively add more components to our digital cancellation design. We first implement only our "linear " digital cancellation which cancels only the linear main self interference components and multipath reflections from the environment. We then add the capability to model non-linear components which we christen "non-linear cancellation" . We calculate the cancellation achieved by these two variants of digital cancellation techniques. For comparison with prior work, we also implement only the digital cancellation technique described in the balun based design [11]. We plot the increase in noise floor for all the techniques as a function of Transmit power in Fig. 12.

As we can see, our full digital cancellation technique cancels everything to the receiver noise floor. Further, notice that just our linear digital cancellation stage leaves 16 dB of self interference residue above the receiver noise floor. Being able to model the non-linear harmonics allows us to reduce self interference by a further 16 dB and cleans out the non-linear distortions almost to the receiver noise floor. In comparison, the prior work's digital cancellation technique falls far short, leaving nearly 18dB of self interference residue over the noise floor since it cannot model non-linear distortions. Note

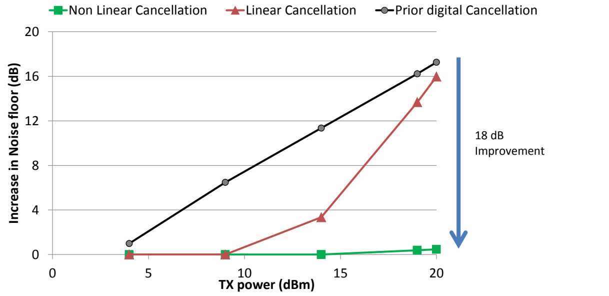

Figure 12: Performance of digital cancellation showing impact of different components of the algorithm vs TX power with fixed constellation = 64 QAM, bandwidth = 80MHz. Our algorithm cancels the main component, reflections and harmonics, thus ensuring that self interference is completely eliminated, and the increase in noise floor less the 1dB. Prior techniques can not cancel harmonics, and therefore increase the noise by 18dB.

that we have given prior work the benefit of an analog cancellation of 62dB from our circuit, as we saw before in Sec. 5.1 if we used their implementation of analog cancellation the numbers are worse.

5.2.4 Dynamic Adaptation

As environmental conditions change, the level of cancellation drops since the values of the attenuators used will be off w.r.t to the new conditions. In this section, we evaluate how long it takes to re-tune analog cancellation, as well as how often it needs to be re-tuned in our indoor environment. Note that digital cancellation is tuned on a per-packet basic, hence it is not a concern. Analog cancellation has to be tuned via a special tuning period during which no data is transmitted, hence quantifying that overhead is important.

We conduct this experiment in our busy indoor environment with other WiFi radios and students moving around. Note that an indoor environment is the worst case scenario for full duplex, because of the presence of a large number of reflectors near the transmitter. Outdoor LTE scenarios are less likely to have such strong near-field reflectors, hence we believe our design extends relatively easily to outdoor LTE scenarios. We place the full duplex node and conduct analog cancellation tuning as described in Sec. 3.3. Specifically, we use the WiFi preamble to determine the initial settings of the attenuators to be used to match the frequency response of the circulator and antenna. Next we run a gradient descent algorithm to further improve the cancellation from that initial point. Each iteration of the gradient descent consumes 92μs since we have 16 different directions to compute the gradient one (corresponding to the 16 different attenuators). We compute the time it takes for the analog cancellation to converge. We repeat this experiment several times for different node placements and environmental conditions and plot the average convergence time. We also conduct an experiment where we do not use the initial frequency based tuning and only use gradient descent from a random starting point for the attenuator values. We show the cancellation achieved as function of tuning time on right side of Fig. 13.

As we can see in right side of Fig. 13, our analog tuning converges in around 920μs, compared to the 40 or more milliseconds it takes for a pure gradient descent based approach. The reason is that the frequency based initial point estimation provides a point very close to optimal, and from that point a few gradient descent iterations allow us to find the optimal point. Our cancellation algorithm therefore tunes an order of magnitude faster than a simple gradient descent based approach.

But an important question is how often do we have to tune? Analog cancellation has to re-tuned when there is a change in the near-field reflections, since it cancels only the strong components (components 50 dB above noise floor, farther out reflections are weaker than this 50dB threshold). Hence the question is how often do the near-field reflections change? As expected, this depends on the environment, for the indoor office deployments we used in our experiments we found that we needed to retune once every 100ms on average (outdoor scenarios would be easier since changes in near field occur less frequently, and we leave mobile hand-held scenarios to future work). We show this experimentally in Fig. 13, the left plot shows the amount of cancellation observed as a function of time after we have found the optimal operating point from a large collection of different experimental runs in our testbed. We define the "near field coherence time" of analog cancellation as the time upto which the receiver remains unsaturated from when it was tuned, which we also use as the trigger to rerun the tuning algorithm. As we can see the near field coherence time for the cancellation is roughly 100 milliseconds. In other words, we have to retune the analog cancellation once every 100 milliseconds, which leads to an overhead of less than 1%.

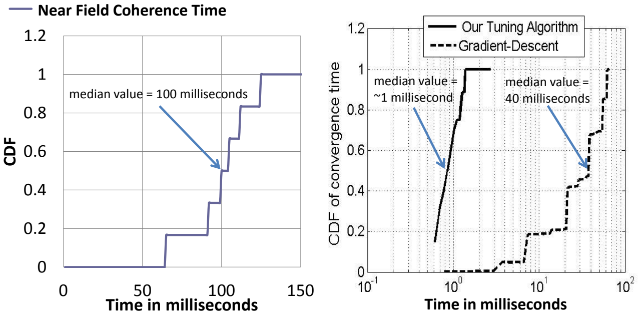

Figure 13: Left figure shows CDF of near field coherence time. This implies that we have to retune analog cancellation on an average of every 100 milliseconds. Right figure shows how long it takes for our tuning algorithm to converge to the required cancellation, after the initiation of tuning. We observe exponential improvement compared to the gradient descent algorithm which takes an order of magnitude longer.

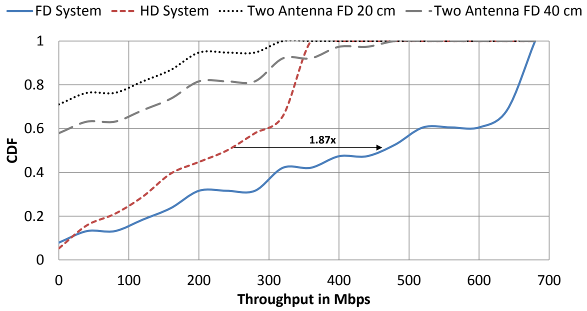

Figure 14: CDF of throughput for full duplex link using TX power = 20 dBm, bandwidth = 80MHz. We see a median gain of 87% using full duplex as compared half duplex. Further, prior full duplex with two antenna's separated by 40cm show gains, only in 8% of cases.

5.3 Does Full Duplex Double Throughput?

This section demonstrates experimentally that our design delivers close to the theoretically expected doubling of throughput for a full duplex WiFi link. Note that this is a PHY layer experiment, a full MAC design for full duplex WiFi is beyond the scope of this paper.

We conduct these experiments as follows. We place the two full duplex nodes at different locations and send trains of 1000 packets in full duplex mode, and then similar trains for each direction of the half duplex mode. Each train uses a particular bitrate (from WiFi) and we cycle through all the bitrates for each location. We pick the bitrate with the best overall throughput for full duplex, two antenna full duplex and half duplex respectively. We repeat this experiment for different locations. We found the SNRs of the links varied uniformly between $0 - 45$dB across locations as we would find in a typical indoor deployment. We plot the CDF of the throughput for half duplex and full duplex link in Fig. 14. 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 $1.87\times$ over the standard half duplex mode. As we known from the experimental analysis in Sec. 5.1.2 that there is a small SNR loss due to a small amount of self interference residue. This SNR loss is the reason that instead of the theoretical $2\times$, we see a slightly reduced gain of $1.87\times$.

How do prior designs perform? We found that in 60% of the scenarios, the throughput with prior full duplex techniques was zero. This is because these designs leave at least 25dB of self-interference residue that acts as noise and if the link SNR is below 30dB no signal is decoded (WiFi requires a minimum of $4-5$dB to decode even the lowest rate packet). As the half-duplex link SNR increases, 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 prior techniques was when the link SNR was greater than 40dB.

6. DISCUSSION & CONCLUSION

We believe this paper marks an important step in proving that full duplex is not only possible, but feasible and practical. Further, it can be deployed with no overhead in terms of antennas used and yet achieve the theoretical doubling of throughput. Below we discuss the current design's limitations, potential avenues of future work and then conclude.

Size of circuit: The current analog cancellation circuit is large, our prototype board measures 10×10 cm. Such a design is fine for APs and base-stations which is our initial focus, however this design is not implementable on phones and other portable devices where size is at a premium. To realize full duplex on such devices, we need to design an RFIC that is sufficiently small (at best $20 - 30$sq.mm for current phones). The key consumers of space on our circuit are delay lines, which we currently realize via traces on the board. For an RFIC we expect to use different techniques to realize the same delays, such as LC ladders and acoustic technologies such as SAW and BAW [12]. These techniques operate by slowing the speed of light, and thus true time delays are obtained in very short form factors that can be integrated on chip. However the above discussion is speculative and is part of our future work.

LTE: Our current prototype targets WiFi frequencies in the 2.4GHz band. However our prototype can also be used for the 2.3GHz and 2.5GHz LTE bands found in Asia and Europe. However the general design of our system is frequency independent, the dependence in our prototype comes from the fact that several analog components in our cancellation board work only in specific frequency ranges (our tunable attenuators operate only between 2-2.6GHz). However, the same design can be used for different frequencies with corresponding components that work in those frequency ranges. Further, unlike WiFi, LTE uses smaller channels, the widest channel is 20MHz and this makes the cancellation problem somewhat simpler. Hence we believe our current design can be adapted to work with LTE, and this remains future work.

MIMO: The current design targets SISO scenarios. For MIMO we could use the same design, but a key challenge is that the cross-talk between different antennas also has to be canceled in analog. Hence, an analog cancellation circuit has to be designed that models not just the distortions through a circulator and a single antenna, but also the distortions that happen when signals travel across antennas. Designing an efficient space-compact circuit for this problem is part of our current research focus.

Finally, we would like to comment that full duplex radio design is a problem that spans three different research areas: RF circuit & system design, digital signal processing and networking. The problem cannot be solved in any one domain alone, the solution in our opinion requires understanding trade-offs across all these domains and architecting it appropriately. Historically however, these communities have been separate, RF system designers expect baseband IQ samples as the interface and view their job as sending and receiving signals in RF from these baseband IQ samples. DSP designers view their job as converting between bits and IQ samples efficiently in the presence of noise. Finally, networking researchers transact in bits and packets and design medium access while abstracting out the underlying details. Realizing and taking advantage of full duplex requires research that spans across these domains, and this work represents a step in that direction.

Acknowledgments: We would like to thank all the Stanford Networked Systems Group members for their insightful comments. This work was supported by a Thomas and Sarah Kailath Stanford Graduate Fellowship & NSF grant #1064504.

7. REFERENCES

[1] D. W. Bliss, P. A. Parker, and A. R. Margetts. Simultaneous transmission and reception for improved wireless network performance. In *Proceings of the 2007 IEEE Workshop on Statistical Signal Processing*, 2007.

[2] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004.

[3] 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.

[4] L. Ding. Digital predistortion of power amplifiers for wireless applications. 2004.

[5] M. Duarte, C. Dick, and A. Sabharwal. Experiment-driven characterization of full-duplex wireless systems. *CoRR*, abs/1107.1276, 2011.

[6] M. Duarte and A. Sabharwal. Full-duplex wireless communications using off-the-shelf radios: Feasibility and first results. In *Forty-Fourth Asilomar Conference on Signals, Systems, and Components*, 2010.

[7] E. Everett, M. Duarte, C. Dick, and A. Sabharwal. Empowering full-duplex wireless communication by exploiting directional diversity. In *Signals, Systems and Computers (ASILOMAR), 2011 Conference Record of the Forty Fifth Asilomar Conference on*, pages 2002 –2006, nov. 2011.

[8] A. Goldsmith. *Wireless Communications*. Cambridge University Press, New York, NY, USA, 2005.

[9] 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.

[10] 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.

[11] 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.

[12] T. Lee. *The Design of CMOS Radio-Frequency Integrated Circuits*. Cambridge University Press, 2004.

[13] A. V. Oppenheim, R. W. Schafer, and J. R. Buck. *Discrete-time signal processing (2nd ed.)*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1999.

[14] 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.

[15] Rohde & Schwarz. *Rohde & Schwarz FSW Signal and Spectrum Analyzer User Manual*, 2012.

[16] Rohde & Schwarz. *Rohde & Schwarz SMBV 100A Signal Generator User Manual*, 2012.

[17] *PE 47303 Data-sheet*. http://www.psemi.com/pdf/datasheets/pe43703ds.pdf.

[18] *Power Amplifier Data-sheet*. http://www.minicircuits.com/pdfs/ZHL-30W-262+.pdf.

[19] *US Patent 5444864*. http://www.google.com/patents/US5444864.

[20] *US Patent 6539204*. http://www.google.com/patents/US6539204.

[21] J. Bardwell. *Tech Report*. http://www.connect802.com/download/techpubs/2005/commercial_radios_E0523-15.pdf.