FlexRadio: Fully Flexible Radios and Networks

Bo Chen, Vivek Yenamandra, and Kannan Srinivasan, *The Ohio State University*

https://www.usenix.org/conference/nsdi15/technical-sessions/presentation/chen

This paper is included in the Proceedings of the
12th USENIX Symposium on Networked Systems
Design and Implementation (NSDI '15).
May 4–6, 2015 • Oakland, CA, USA
ISBN 978-1-931971-218

FlexRadio: Fully Flexible Radios and Networks

Bo Chen†*, Vivek Yenamandra*† *and Kannan Srinivasan*
Department of Computer Science and Engineering
The Ohio State University, Columbus, OH 43210
{*chebo, yenamand, kannan*}*@cse.ohio-state.edu*
†*Co-primary Authors*

Abstract

When a wireless node has multiple RF chains, there are several techniques that are possible; MIMO, full-duplex and interference alignment. This paper aims to unify these techniques into a single wireless node. It proposes to make a wireless node fully flexible such that it can choose any number of its RF chains for transmission and the remaining for simultaneous reception. Thus, MIMO and full duplex are subset configurations in our design. Surprisingly, this flexibility performs better than MIMO or full duplex or interference alignment or multi-user MIMO.

This paper presents the design and implementation of FlexRadio, the first system enabling flexible RF resource allocation. We implement FlexRadio on the NI PXIe 1082 platform using XCVR2450 radio front-ends. FlexRadio node networks achieves a median gain of 47% and 38% over same networks with full duplex and MIMO nodes respectively.

1 Introduction

When a wireless node has multiple radio frequency (RF) chains, the state-of-the-art technology has been to use either all of them for transmission or reception, as in multiple-input multiple-output (MIMO). Recently, many research groups have shown that a node can transmit and receive simultaneously and thus, be full-duplex. Under full-duplex operation, a node activates equal number of RF chains for transmission as it does for simultaneous reception. Thus, when a node has N RF chains, under full duplex, N/2 RF chains are active transmitting RF chains while the remaining N/2 RF chains are receiving RF chains. Under MIMO, all N RF chains are either active transmitting RF chains or active receiving RF chains. There is much work in the wireless community studying which of these techniques are better and *when* [2, 3, 7]. Fundamentally, the capacity achieved by MIMO and full-duplex between a pair of nodes, is the same. The main difference is that MIMO supports N simultaneous transmissions in one direction, while full duplex supports N/2 in both the directions. Still, the total number of transmissions is only N in both the cases[1].

From the above discussion, it is clear that there is no significant difference in the capacity between MIMO and full duplex. However, this paper shows that when we unify MIMO and full duplex, and make the design fully flexible then, surprisingly, the capacity can be improved by 2x when compared to MIMO or full duplex. By flexible, we mean that, out of N active RF chains, our system allows M of them to be transmit RF chains and (N-M) of them to be receive RF chains, where $0 \leq M \leq N$. We call our system, FlexRadio.

Although choosing between MIMO and full duplex configurations gives no improvement in throughput between a pair of nodes, it does improve the overall *network* throughput. This improvement comes from the difference in the interference footprint between MIMO and full duplex, in a network [18, 19]. During a MIMO transmission, a secondary transmission around the receiver and a secondary reception around the transmitter is prohibited. However, a secondary reception around the receiver and a secondary transmission around the transmitter is allowed as long as they do not affect the ongoing transmission. Similarly, during a full duplex transmission, transmission around both the nodes is prohibited, while another reception is possible. FlexRadio's flexibility allows a network to exploit this difference to increase the number of parallel transmissions in a network. Section 3 shows that when every node can choose between full duplex or MIMO operation, the total network throughput increases by 50% compared to the case when all the nodes are either MIMO or full duplex.

---

[1] Note that in both the cases, a node has 2N (N transmit and N receive) RF chains but, only N of them are active. For rest of the paper, by an N RF-chain node, we imply a node with N active RF chains (either transmit or receive or both) unless explicitly stated otherwise

The gain in FlexRadio is not simply from choosing between MIMO and full-duplex configurations. But, it is from choosing from all available configurations within FlexRadio. Section 3 shows one example where a configuration that is not MIMO or full duplex improves the throughput by 2x, even between a pair of nodes.

Thus, the unified architecture with its adaptability makes it more powerful than the traditional (inflexible) configurations. This is a fundamental improvement in throughput for a multi-RF chain wireless node. Section 3 motivates the need for a flexible architecture and gives some guidelines on choosing the optimal configuration. The optimal configuration depends on the topology, flow demands, wireless channel and the number of RF chains available at other neighbouring nodes.

This paper makes the following contributions;

1. It proposes flexibility as a new radio capability. In Section 3 we motivate this need based on different network properties. Further, we show that FlexRadio can outperform MIMO, full-duplex and interference alignment techniques.
2. It presents the *first* fully flexible FlexRadio prototype. This prototype has multiple novel mechanisms to reduce implementation complexity. First, an antenna placement design that reduces the number of RF cancellation elements needed (Section 4). Second, a novel non-linearity mitigation strategy to reduce complexity of digital cancellation. A naive non-linear elimination technique would require $\mathscr{O}(M^2)$ modules, where M is the number of transmitting RF chains. We eliminate the non-linear components at the transmitter by using a preconditioning module at each transmitter itself. We reduce the number of non-linearity mitigating modules to $\mathscr{O}(M)$ (Section 4).

The flexibility proposed in this paper is a new feature for a wireless node. This has not been studied in information theory or network theory or wireless systems. This new capability has deep implications to wireless networking: A wireless routing protocol can take into account the number of RF chains available at every node and choose the number of RF chains for transmission (and reception) at different nodes so as to maximize end-to-end throughput.

2 A Primer on MIMO and Full Duplex

This section gives a brief overview of capacity, the maximum achievable throughput. The overview helps motivate flexibility as shown in the following section. Capacity is a function of the quality of wireless link. This quality is measured as the ratio between the received signal strength and the local noise at a receiver (SNR).

Since the generic capacity equations are not easy to interpret, often, approximations are used in literature [17]. For the generic case, when node 1 (transmitter) has n_{tx} RF chains and node 2 (receiver) has n_{rx} RF chains, at high SNR, with a well-conditioned channel matrix, the capacity for fading channel is approximated by:

$$C_{High\_SNR} \approx min(n_{tx}, n_{rx}) * log_2(1 + SNR) \qquad (1)$$

Here, the capacity is equivalent to having $min(n_{tx}, n_{rx})$ parallel streams. Thus, at high SNR, the capacity scales linearly with $min(n_{tx}, n_{rx})$ [17].

At low SNR, with a well-conditioned channel matrix, the capacity for the fast fading channel is approximated by:

$$C_{Low\_SNR} \approx n_{rx} * log_2(1 + SNR) \approx n_{rx} * SNR \qquad (2)$$

Here, the capacity is only a function of the number of receive RF chains. It linearly increases with the number of receive RF chains [17]. These approximations are valid for MIMO and full-duplex[2].

Takeaways: When the SNR is high, equalizing the number of transmitting RF chains at the sender and the number of receiving RF chains at the receiver node gives the maximum throughput. When the SNR is low, on the other hand, maximizing the number of receive RF chains at the receiver maximizes the throughput. Note that the low SNR approximation is for very low SNRs (≈ -15dB) at which WiFi node do not operate. However, the intuition applies to SNRs that are reasonably low for WiFi, as shown in Section 5.4.3.

3 The Need for Flexibility

In this section we highlight the benefit of FlexRadio nodes in a wireless network.

3.1 Topology Needs Flexibility

Consider the topology shown in Figure 1(a). It has four nodes with two RF chains each. Nodes N1 and N4 cannot see each other and all other nodes can see each other. This is a common network topology. For example, consider N1 and N4 as APs in an enterprise wireless network that cannot listen to each other. Consider N2 and N3 as clients that can listen to both these APs and to each other. In this topology, if the nodes support fixed MIMO, MU-MIMO or full-duplex functionality, only two packet transmissions can be enabled simultaneously. For example, under MU-MIMO, N1 can simultaneously send one

---

[2] When multiple RF chains are involved, by full-duplex, we refer to the case that half of the chains are operating as transmitters and the others are receivers.

packet to N2 and another to N3. During this slot, N4 cannot transmit to N2 or N3 to avoid causing interference at these nodes. Similarly, if all the nodes are full-duplex nodes, N1 can send a packet to N2, while N2 sends to N4. At this time N3 cannot transmit as it causes interference at N4. Thus, the maximum number of packets transmitted simultaneously is only two. Thus, enabling a third transmission stream in addition to the two transmission streams causes destructive interference at one of the participating nodes. However, in the above topology, if each node supports flexible functionality, it presents them with the required spatial dimensions (antennas) to explore interference alignment solutions [1, 10] to allow a third simultaneous transmission. It must be noted that interference alignment does not require additional capability for MU-MIMO capable wireless nodes.

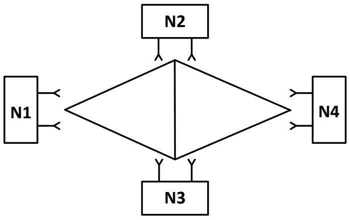

(a) Topology

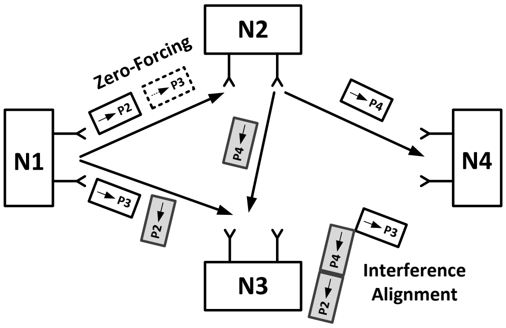

(b) FlexRadio

Figure 1: Topology Needs Flexibility: An example of FlexRadio outperforming MU-MIMO, MIMO and full-duplex, without any flow restrictions. FlexRadio can enable 3 packets to be simultaneously transmitted, while MU-MIMO, MIMO or full-duplex can only enable 2.

In more explicit terms, N1 can send one packet (P2) to N2 and one more (P3) to N3 (as shown in Figure 1(b). Simultaneously, N2 can send a packet (P4) to N4. Since N1 has two antennas, it can null (zero-force) P3 at N2, while aligning P2 with P4 at N3. Since P3 is nulled at N2 and P2 is not, N2 can decode P2. Since N3 is using both the antennas for receiving, it can decode two packets. But, it receives 3 packets. However, since P2 and P4 are aligned, N3 can decode P3 without any interference. At the same time, N4 receives P4 from N2 without any interference. Thus, there are 3 successful packet transmissions. This was possible because of flexibility enabled by FlexRadio and interference alignment techniques. Even when MIMO, MU-MIMO and full-duplex work with interference alignment, they cannot transmit more than 2 packets, while FlexRadio achieves *1.5X* throughput gain.

To understand how FlexRadio was invoked, note that N1 was using its two RF chains to transmit, N2 was using one to transmit and the other to receive, N3 was using both to receive, and N4 was using one to receive. This example shows that FlexRadio can *fundamentally improve capacity of interference limited wireless networks with multi-RF chain nodes.*

3.2 Flow Demand Needs Flexibility

Performance gains of FlexRadio can be seen in other networks as well. Consider a simple network of 3 nodes; say node 1 has 4 RF chains, node 2 has 6 RF chains and node 3 has 2 RF chains.This is a heterogeneous network with different nodes having different number of RF chains. Assume that each hop has the same, but high SNR. The MIMO scenario is shown in Figure 2(a). In this case, MIMO can support $\frac{1}{2} * 4 + \frac{1}{2} * 2$ parallel streams. Here, the first term corresponds to the performance of the link between node 1 and 2, and the second term corresponds to the link between node 2 and 3. Since only one of the two can be active at any time, their overall performance are scaled by half. From network point of view, three streams are enabled simultaneously.

Note that if full-duplex is used, every node would have to split its RF chains equally to transmit and receive. This is shown in Figure 2(b). For full-duplex also, the number of streams that can be enabled simultaneously is $\frac{1}{2} * 4 + \frac{1}{2} * 2$. The capacity is same as that of MIMO even though the flows are in both directions.

In FlexRadio, however, node 1 can transmit on all 4 of its RF chains and node 2 can receive on 4 RF chains. Simultaneously, node 2 can forward packets using the remaining 2 RF chains to node 3, while node 3 uses all of its RF chains for receiving. This is shown in Figure 2(c). Here, node 2 is able to transmit (forward) while receiving because FlexRadio supports full-duplex operation. Now, the number of stream supported in this setup is $4 + 2$. As before, the first term is for the link between node 1 and 2, and the second is between node 2 and 3. There is no scaling for these quantities because these flows happen simultaneously. Therefore, the combined system can support 6 streams. This is *twice* as much as a traditional MIMO or full-duplex system.

However, when the PHY is MU-MIMO capable (such as APs for 802.11n), the same capacity as FlexRadio can

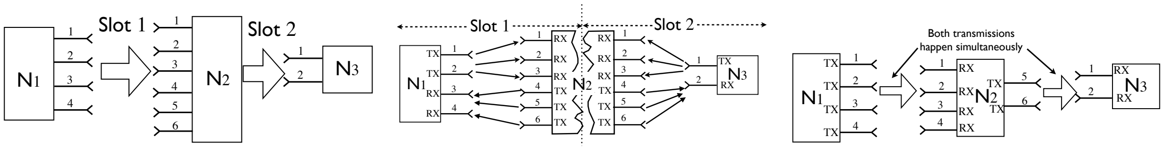

(a) MIMO (b) Full-Duplex (c) FlexRadio

Figure 2: A Heterogeneous Network with different nodes having different number of RF chains.

be achieved where Node 2 uses 4 RF chains to transmit to Node 1 and the remaining to transmit to Node 3 simultaneously. However, when there is a desired flow demand, say Node 1 to Node 2 to Node 3, FlexRadio can improve the throughput of a MU-MIMO system. For this flow demand, the MU-MIMO operation does not provide over MIMO operation.

3.3 Channel Needs Flexibility

Consider nodes 1 and 2 each with M RF chains. Assume, both of them want to transmit to each other. Also, assume a very low SNR channel.

When MIMO alone is used, Node 1 uses all M RF chains to transmit, while Node 2 uses all M RF chains to receive. In the low SNR region (for poor channel conditions), the capacity is simply proportional to the number of receivers used, as shown in Equation 2. Thus, the capacity is $C_{MIMO} \approx M * SNR$.

When fullduplex is used, node 1 uses $\frac{M}{2}$ RF chains to transmit and $\frac{M}{2}$ RF chains to receive, same as Node 2. In this case, we compute the capacity for both transmission directions. The total capacity in the low SNR regime is $C_{FD} \approx \frac{M}{2} * SNR + \frac{M}{2} * SNR$. This capacity is same for both MIMO and full-duplex.

When the flexibility is provided, nodes 1 and 2 can choose the number of RF chains they wish to transmit and receive over. Note that, at low SNR, the nodes should maximize the number of receive RF chains. Therefore, when nodes 1 and 2 use only one RF chain to transmit and the remaining (M-1) RF chains to receive, the sum capacity, in the low SNR region, is $C_{FlexRadio} \approx (M-1) * SNR + (M-1) * SNR$. This is almost double the sum capacity compared to MIMO and full-duplex.

In all these examples, we assumed a central node is made aware of the RF resources of all nodes in the network and their respective traffic demands. We discuss the MAC implications briefly in Sec. 7. In summary, flexibility enables FlexRadio nodes to achieve significant performance gains based on topology, flow and channel constraints.

4 Design Overview

Based on FlexRadio's configuration, the self-interference constituents change. A FlexRadio *self-interference* cancellation circuitry should hence support all these configurations. The challenge in designing FlexRadio's self-inteference cancellation circuitry is the following. It should include cancellation circuitry that accounts for every TX RF chain, a potential source of self-interference, at every RX RF chain. This leads to $M * (M-1)$ cancellation circutry elements for an M RF chain system. This can make implementing FlexRadio node highly expensive. This section presents a design that significantly reduces the number of cancellation elements. For example, for a four RF-chain FlexRadio, our design only requires 2 elements, while the naive approach needs 12.

A self-interference channel between two antennas consists of two components at RF frequencies: line-of-sight and non-line-of-sight component. The line-of-sight component of the interference is simply a function of the distance between the two antennas.[3]This component can be estimated and accounted for using free-space path loss equations. The non-line-of-sight component is a function of the environment. The transmitted signal can reflect off objects in the environment and contribute to the self-interference at the receiver. We account for the self-interference in two stages. In the first stage, majority of the line-of-sight self-interference component is accounted for by RF cancellation (Sec. 4.1). The residual self-interference including the entire non-line-of-sight component is accounted for by digital cancellation (Sec. 4.2).

Finally, a recent work showed that self-interference has non-linear components due to the power amplifier [4] that needs to be accounted for. Extending their non-linear mitigation strategy to an M RF chain system naively requires $\mathcal{O}(M^2)$ non-linear mitigation modules. This section presents a technique that reduces this number to $\mathcal{O}(M)$.

---

[3] Assuming omni-directional antennas and no obstruction between the two antennas.

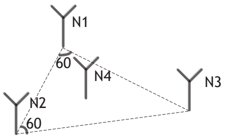

Figure 3: Antenna placement for a four RF-chain FlexRadio system; Three antennas are placed on the vertices of an equilateral triangle with N_4's antenna placed on the centroid

4.1 RF Cancellation

RF cancellation circuitry accounts for the line-of-sight component of self-interference. This component of self-interference signal typically experiences delay and attenuation that is *only* a function of the distance between the TX and RX antenna. Every such link between a transmit and receive RF chain in a FlexRadio node needs a self-interference cancellation block that matches the delay and attenuation experienced by the self-interference over air. We refer to this block as the *delay and attenuation block*. To design an efficient self-cancellation circuitry, we propose an antenna placement scheme that leverages its geometrical symmetry to alleviate the complexity of the RF cancellation circuitry. Symmetric antenna placement makes it possible to combine multiple self-interference signals that have the same delay and attenuation. By doing so, the combined self-interference needs only one *delay and attenuation* block. It must be noted that while the line-of-sight component has the same delay and attenuation as long as the distance between the transmit and receive antenna is the same. the multipath (non-line-of-sight) component can be different. However, our experiments (in Sec. 5) show that these multipath components are not as large as the line-of-sight component and therefore, can be cancelled in the digital domain (explained in the next subsection).

4.1.1 Antenna Placement Scheme (APS)

Figure 3 illustrates the antenna placement scheme for a four RF-chain FlexRadio node. Three antennas, N_1, N_2 and N_3, are on the vertices of an equilateral triangle with the fourth antenna, N_4, at the centroid. In addition to placing the antennas as illustrated, we define an order in assigning which RF-chain to transmit (receive) for a given configuration of FlexRadio. The order of transmission for a four RF-chain FlexRadio node in descending order is: N_1, N_2, N_3 and N_4. For example, N_1 is assigned as the only transmitter when FlexRadio is configured in (1/3) mode[4]. The *advantage* of biasing the order of transmission (reception), together with the symmetry of the proposed antenna placement scheme is the following: *The attenuation and delay of the transmitted signal at a given receiver is independent of the transmitter chain.* In other words, the *delay and attenuation block* in the cancellation path of a given receiver is decoupled from the configuration of the FlexRadio node. For example, the attenuation and delay of the *self interference* signal at N_4 is the same whether originating from N_1, N_2 or N_3. This is true because of biasing the transmission order as this eliminates the possibility of a *self-interference* signal at N_2 or N_3 to originate from N_4.

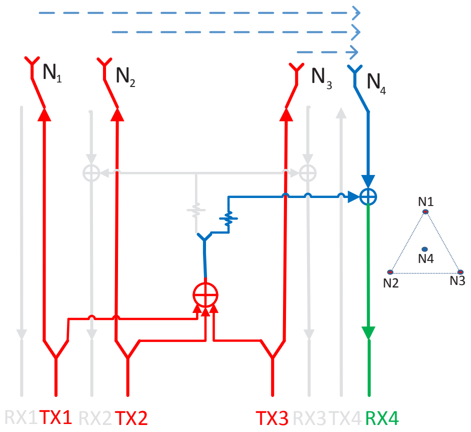

Figure 4: Simplified block diagram of the RF cancellation circuit for a four RF-chain FlexRadio. N_1,N_2,N_3 and N_4 are the 4 antennas with associated TX/RX chains. The figure highlights the active paths in the *self interference* cancellation circuitry for a 3/1 configuration. The cancellation signals from TX1, TX2 and TX3 are combined, inverted (π phase shifter not shown in figure for simplicity) and fed through the *delay and attenuation block* associated to receiver RX4. The *delay and attenuation block* matches the identical attenuation and delay of the *self interference* signals. The dashed lines directed from the TX antennas to the RX antennas illustrate the link in air traversed by the *self interference* signals. The top view of the antenna placement scheme is shown next to the block diagram.

4.1.2 Cancellation Design

Figure 4 illustrates a simplified block diagram of the *self-interference* RF cancellation circuitry for a four RF-chain FlexRadio node. It illustrates the RF signal paths connecting the antennas with the respective RF chain. Specifically, it highlights the *active RF paths* when the

---

[4] We define a configuration, n_t/n_r, of FlexRadio as a mode of operation in which it commits n_t of its RF-chains to transmit and the remaining n_r RF-chains to receive simultaneously.

node is configured in 3/1 mode. The *inactive* RF paths are greyed. The notation for the antennas in Figure 4 is consistent with that in Figure 3. The TX/RX RF chains are labelled as TXi/RXi respectively, where i is the index of the associated antenna.

As illustrated in Figure 4, in the 3/1 mode, the switches on antennas N_1, N_2 and N_3 are toggled towards the transmit RF chains TX1, TX2 and TX3 respectively, while switch on antenna N_4 is toggled towards the receive RF chain, RX4. This is in accordance with the transmission order given in Sec. 4.1.1. We explain the cancellation circuitry design by first looking at the active RF paths from the transmit RF chains and then the active RF paths to the receive RF chains. Specifically we will consider 3/1 scenario illustrated in Figure 4, to underline *how* our symmetric antenna placement design enables us to reduce complexity of the design.

The TX Chains. As indicated in Figure 4, the power from each of the Tx chains, $TX1$, $TX2$ and $TX3$, is split into two paths - *transmit path* and *cancellation path.*

The *transmit path* from each TX chain feeds the power to its corresponding antenna. As indicated in Figure 4, the path from $TX4$ to the switch is *not* split. In other words, there is not cancellation path from TX4. This is because of the biasing order in Sec 4.1.1. When N_4 is the transmitter, FlexRadio is configured as 4/0 and thus the FlexRadio node has no active receive RF chains and thus no self-interference.

The *cancellation paths* from the TX chains feeds part of the power to the receive RF paths to enable self-interference cancellation. Self-interference cancellation at a given receiver is achieved by subtracting the self-interference signal it receives (on its antenna) with an exact copy of it. The cancellation path is responsible for *generating an exact copy of the self-interference signal* to each receive RF path. We call this the cancellation signal. The cancellation path draws part of the transmit power to generate a copy of the transmitted signal. This cancellation signal is then subjected to delay and attenuation to match that experienced by the self-interference over the air.

Exploting Symmetric Antenna Placement and Biased Transmission Order: Symmetric antenna placement coupled with transmission biasing order decouples the self-interference channel at a given receiver from the potential source of self-interference. For example, the delay and attenuation of the self-interference channel at receiver N_4 is the same irrespective of whether the source of self-interference is N_1, N_2 or N_3. This allows us to combine the cancellation signals and subject the combination of these cancellation signals to a *delay and attenuation block* that matches that experienced at that receiver[5]. Thus, as indicated in Figure 4, the cancellation signals from TX1, TX2 and TX3 are combined and are collectively subjected to match the delay and attenuation experienced at receiver N_4.

The receiver's perspective. As indicated in Figure 4 each RF path between the RF switch and the receivers RX2, RX3 and RX4 has a combiner. The combiner adds the received signal from the antenna with the inverted copy of the generated cancellation signal to implement self-interference cancellation in the RF domain. Consider $RX4$. $RX4$ is subject to *self interference* from N_1, N_2 and N_3. One input to the combiner in the RF path from N_4 to RX4 is the signal received by the antenna, N_4, itself. This signal is a combination of self-interference and the desired signal intended for the receiver RX4. The other input is the internally generated inverted copy of the combined self-interference signal as discussed previously. Thus, ideally at the combiner output, while the desired signal passes through unchanged, the self-interference signal received at the antenna is cancelled by its internally generated inverted copy.[6] As an aside, $RX1$ does not need a combiner in its path since when N_1 is the receiver, so are all the other RF-chains of the FlexRadio node.

***Delay and Attenuation Block*: Beneath the abstraction.** Each *delay and attenuation block* consists of a variable attenuator and a variable delay block that are controlled by from the baseband. By controlling the attenuator and the phase shifter, the cancellation signal can be conditioned to be an inverted replica of the signal received at the corresponding receiver.

Finally, the switch, illustrated in Figure 4 is used to connect either TX or RX path to the antenna depending on the configuration of the RF-chain. Figure 4 illustrates the active signal paths when FlexRadio is configured as 3/1. For example, when changing from mode 3/1 to mode 2/2, the RF switch associated with N_3 switches to the receive RF path. Simultaneously, TX3 is deactivated while RX3 is activated. Deactivating TX3 renders its corresponding transmit and cancellation paths in the cancellation circuitry inactive. At the same time, the RF path from N_3 to RX3 is active with its associated combiner and *delay and attenuation block.*

Is the symmetry assumption realizable? The requirement of high self-interference cancellation required ($\approx 110dB$) implies that the symmetrical placement is strictly observed. For this, we need to ensure that the omni-directional antennas are parallel to each other and are exactly placed as indicated in Fig. 3. We implement the cancellation circuitry on a PCB and couple the antennas

---

[5] Before passing the combined signal through the *delay and attenuation block*, we invert the signal to enable subtraction at the receiver using just a combiner

[6] This is called RF cancellation since the self-interference cancellation is performed completely in the RF domain.

to the PCB using SMA cables. Existing PCB manufacturing tolerances enable us to place objects on the PCB within an accuracy of 2 mils (1 mil = $\frac{1}{1000}$ inch). While the antennas are not perfectly omni-directional, we observe that inaccuracy in this modeling is accounted in digital cancellation where the self-interference channel is explicitly measured.

4.2 Digital Cancellation

Digital cancellation is used to capture the multipath components of the self-interference. The self-interference from equidistant transmit antennas to a receive antenna likely experience different multipath profiles. Our digital cancellation design is similar, in principle, to previously proposed techniques [6, 16]. This cancellation module estimates the coefficients of the multipath components using a *finite impulse response* (FIR) filter. Unlike the RF cancellation technique, an M RF chain FlexRadio system needs $M*(M-1)$ FIR-based digital cancellation modules. However, joint channel estimation techniques have been proposed to reduce the complexity of the digital cancellation implementation [3]. These techniques can be applied here as well to reduce resource utilization of digital cancellation implementation.

A recent work [4] showed that FIR-based digital cancellation alone does not suffice to achieve the 110dB total cancellation needed for a WiFi full duplex system. This work identified non-linear components of self-interference that cannot be estimated using FIR filters. It proposed modeling the non-linear component using a polynomial function at each receiving RF chain to mitigate its effect. Thus, in an M RF-chain FlexRadio node, each receiver models the non-linearities of M-1 transmitters. Since every antenna can be configured as a receiver, we would require $\mathcal{O}(M^2)$ such modules.

Can we reduce the number of non-linear mitigation modules from $\mathcal{O}(M^2)$? We present a technique to reduce this number from $\mathcal{O}(M^2)$ to $\mathcal{O}(M)$. The key insight here is that the non-linear components arise from the transmit RF chain's power amplifier [4]. Therefore, instead of estimating and correcting for this non-linearity at the receiver RF chain, we estimate it at the transmitter RF chain and correct for it even before transmission. This *pre-conditioning* needs to be done only at the transmit RF chains. This reduces the complexity from $\mathcal{O}(M^2)$ to $\mathcal{O}(M)$. While joint channel estimation techniques have been proposed to further reduce the complexity of digital cancellation implementation [3], decoupling the digital cancellation and non-linear mitigation from FlexRadio's configuration assists in supporting the flexibility desired.

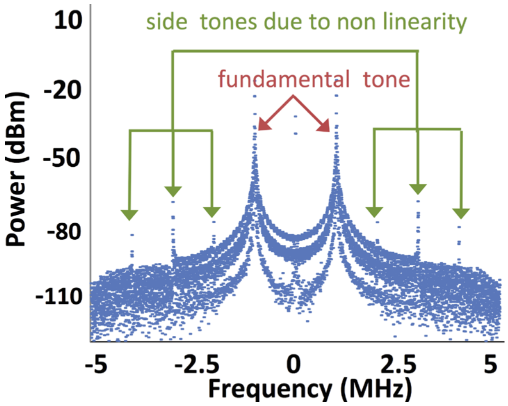

Figure 5: The effect of non-linearity on the transmitted PSD. In addition to the fundamental tones, the side tones prop up due to non-linearity of the transmitter.

4.2.1 Dealing with non-linearities

The distortion caused by transmitter non-linearity on the transmitted signal is illustrated in Fig. 5 when the transmitter sends two single tone frequencies. Similarly, for a wideband OFDM type symbol, the non-linearity results in increased power in the side-bands (adjacent band).

The observed non-linearity can be understood by looking at the received signal (without pre-conditioning):

$$Y(x) = \sum_i \alpha_i x^i \tag{3}$$

where x is the voltage of the analog signal input to the power amplifier. This simple model models the power amplifier non-linearity using a polynomial. Estimating the non-linearity is equivalent to finding the coefficients of the polynomial. Contrary to the technique proposed in [4], we tackle this phenomenon by *pre-conditioning* the input signal of the power amplifier at transmitter itself.

Thus, instead of transmitting the signal x, we transmit the following,

$$f(x) = \alpha_1 (x - \sum_{i=3,5,7,9,11} (\alpha_i/\alpha_1) x^i) \tag{4}$$

Thus, when the input signal is preconditioned, the output of the power amplifier is approximately linear. In effect, the signal preconditioning block lowers the input signal power to the power amplifier thus preventing its high gain from saturating the output, thus reducing nonlinearities.

Will the non-linearity introduced in Eq. 4, violate linearity assumptions of communication systems design? It must be noted that here, we introduce preconditioning at the signal level in an effort to *balance* the nonlinearity of the power amplifier and make the resulting output signal linear. This is equivalent to preconditioning

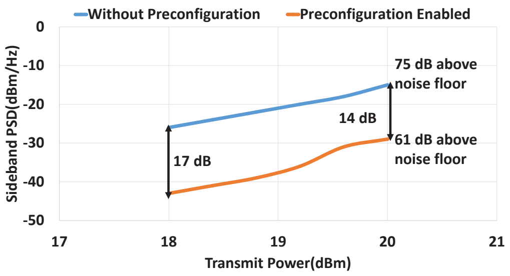

Figure 6: The PSD of the transmitter sidebands reduces after enabling the preconfiguration module.

the signal at the receiver side *after* the signal experiences non-linearity of the power amplifier. The preconditioning in effect, reduces the power of the non-linear components in the channel and makes the linear approximation of communication systems more valid.

We model the non-linearity of the transmitter in the training phase. We send a training series of analog inputs of known power to the power amplifier and derive the coefficients of the polynomial by measuring the output power. Once we model the non-linearity, we precondition the signal using equation 4. We transmit a wideband OFDM signal by sweeping the transmit power from close to its maximum power to its maximum power. When transmitting this OFDM signal, we measure the power of the sidebands, when the preconfiguration module is disabled and again when the preconfiguration module is enabled. We use an external power amplifier to boost the power up to 20 dBm.

Figure 6 plots the findings from our experiment. We vary the transmit power from 18dBm to 20 dBm. This power range captures the strongest non-linear behaviour of the power amplifier. The preconditioning module decreases the PSD of the sidebands by 17dB at transmit power of 18dBm and by 14 dB at the highest transmit power. The decrease in reduction of the PSD of the sidebands at higher power suggests that the fundamental tone is more saturated, i.e. the power amplifier exhibits a stronger non-linear characteristic. However, across the entire power range of the transmitter, enabling the preconfiguration module limits the PSD of the sidebands to at most **61 dB** above the noise floor at the receiver.

5 Implementation and Evaluation

The antenna placement design assumed that its symmetric design implied equal attenuation and delay for line-of-sight self-interference from equidistant transmit RF chains. This lead to the reduction in the number of programmable attenuators needed for RF cancellation. When this assumption does not hold, cancellation performance degrades potentially below the 110dB cancellation needed for WiFi. This section evaluates the design principles presented in the previous sections. We achieve the desired 110dB cancellation with our design.

5.1 FlexRadio Implementation

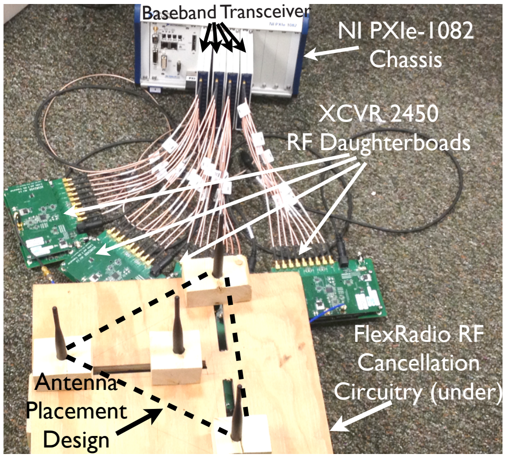

Figure 7: Four RF-chain FlexRadio system

Figure 7 shows our four RF-chain FlexRadio system implementation. It can be viewed as a cascade of three high-level modules connected to each other using SMA cables: *The Antenna Placement site, RF cancellation circuitry, RF/baseband chains*.

The antennas are held in position by sliding them through slotted wooden blocks. They are connected to the cancellation circuitry using SMA cables. The distance between the antennas on the vertices and the centroid antenna is set to 5.5”.

The RF chains are implemented using the XCVR 2450 (RF front end), the NI-5781 (data converter module with a 14 bit ADC and 16bit DAC) and the NI PXIe-7965R (a Xilinx Virtex-5 based FPGA) for baseband processing including digital cancellation implementation. The FPGAs are housed in a chassis that contains communication and clock backplanes to facilitate synchronization and communication among the FPGAs.

Figure 8 shows the designed FlexRadio RF cancellation circuitry. The TX and RX ports labelled in Figure 8 are consistent with the labelling used in Figure 4. The cancellation circuit employs the PE43704, a 0-31.75dB attenuator that can be programmed in 0.25dB steps. The attenuators are controlled with on-board switches. We match the delay between the cancellation and self-interference paths with a symmetrical copper trace design on the PCB board. We built the circuit on Rogers 4350 PCB material. The board dimensions are 9”x8”.

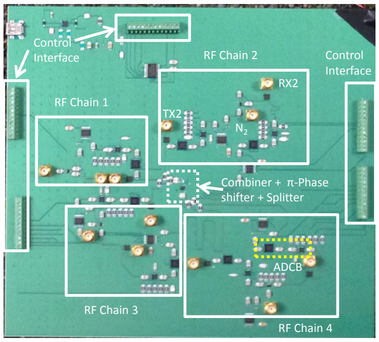

Figure 8: FlexRadio RF cancellation circuitry. Each RF chain contains three ports: Antenna, TX and RX port (indicated in Figure 4)The block labelled ADCB is the *delay and attenuation* block described in Section 4

5.2 Self-Interference Cancellation Evaluation

FlexRadio's self-interference cancellation has three distinct modules: RF, digital cancellation module and the transmitter preconditioning module. These modules, in unison, enable FlexRadio to nullify its self-interference in each of its operating configurations.

The RF cancellation cancels the line-of-sight component of the self-interference. Digital cancellation module estimates the channel and nulls the multipath component of self-interference. However, the digital cancellation module cannot predict the non-linearity of the transmitter. As indicated in Sec. 4.2.1, the preconditioning module limits the power in the sidebands to 61 dB over the noise floor. Thus, FlexRadio needs to provide RF cancellation of at least 61 dB to eliminate the non-linear components introduced by the transmitter.

Is the symmetric design effective? We evaluate the self-interference cancellation of FlexRadio over all of its operating modes. We place our four RF-chain FlexRadio prototype inside our lab - a typical indoor environment with metallic cubicles and furniture. We transmit 20 MHz OFDM signal at the transmitters in each of these modes. Figure 10 illustrates the PSD at the centroid at different stages of self-interference cancellation for different configurations of FlexRadio. The RF cancellation at the centroid is constant across different modes of operation and is 68 dB. As illustrated in Figure 10, this is sufficient to reduce the power in the side bands (and thus significant portion of the non-linear component) to the noise floor. The RF cancellation is a function of only the antenna placement (since we do not place any objects between the antennas) and we observe it to be at least 68 dB at all RF chains in our prototype.

Evaluating Digital Cancellation Effectiveness: Digital cancellation effectiveness relies on the accuracy of the self-interference channel estimation. Intuitively, measuring the channel response over a longer duration helps in estimating the channel better. Fig. 9 illustrates the digital cancellation performance as a function of the channel estimation time. As seen in Fig. 9, for a channel estimation time of 7.4μ seconds, 42 dB of digital cancellation is achieved. Our digital cancellation module cancels the residue signal from RF cancellation down to the noise floor for all operating modes of FlexRadio.

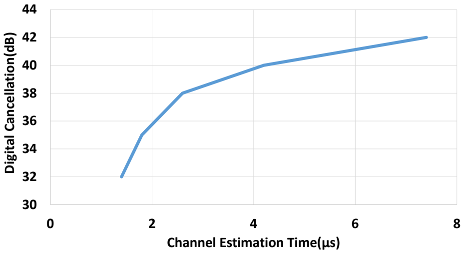

Figure 9: Digital Cancellation as a function of time taken to estimate the self-interference channel.

Figure 10 explicitly illustrates the spectrum at the centroid antenna after RF cancellation. When FlexRadio is operating in mode, 1/3, the effect of multipath is more pronounced after RF cancellation indicated by the trough in the residual spectrum after RF cancellation. However, the depth of this trough decreases as the number of transmitters increases i.e the effect of multipath is lesser. In the mode 3/1, the spectrum after the RF cancellation is almost flat. This is because when the number of transmitters increases, the multipath component decreases as the number of line-of-sight components increase.

The RF cancellation at the centroid includes 26 dB attenuation of the self-interference signal over air. Due to space constraints, the power spectral density at each of the other vertices is not included. The RF cancellation at the vertices is 70 dB, due to the the increase in attenuation of the self-interference over the air (FlexRadio's priority ensures that a receiver at the vertex experiencing self-interference only from transmitters positioned at other vertices of the equilateral triangle).

5.3 Configuration Switching Time

When switching from one FlexRadio configuration to another, the switching time can include the time needed for carrying out some, if not all, of the following events: Switching of the RF switches to change receive chains to transmit chains or vice versa; Channel estimation between all the transmit and receive links in the baseband - this event loads the coefficients of the FIR filters used to model the self-interference channels required for digital cancellation; Switching the baseband state to make

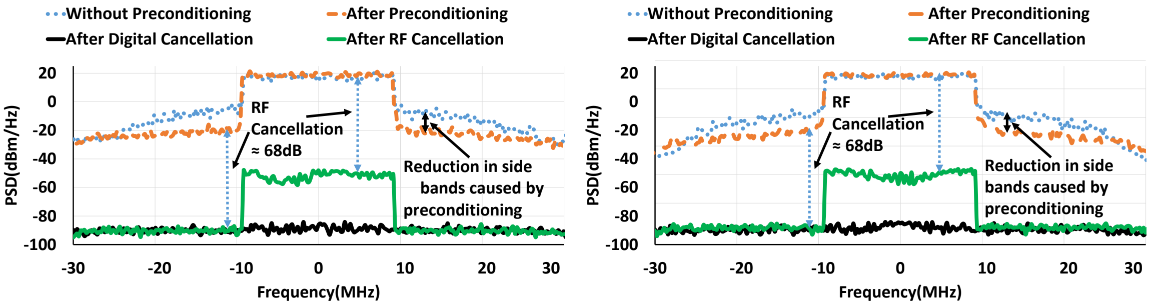

(a) 1/3 Configuration (b) 2/2 Configuration

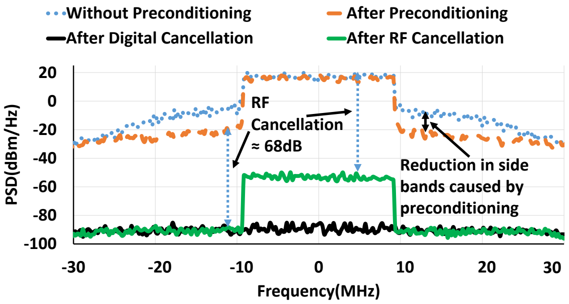

(c) 3/1 Configuration

Figure 10: PSD at the centroid for different operating modes of FlexRadio. Preconditioning reduces non-linear components, RF cancellation achieves 68dB cancellation, FIR-based digital cancellation brings the remaining self-interference to the noise floor achieving a fully working FlexRadio.

the additional transmit (receive) FIFO available (For instance, when switching from mode 2/2 to 3/1 an additional transmit FIFO is required). Explicitly, to switch between transmission modes 4/0 and 0/4, FlexRadio only needs to switch the RF switches at each RF chain from the transmit RF chain to the receive RF chain. However, when FlexRadio switches from mode 0/4 to mode 3/1, all the events listed above have to be accomplished to transition between the two modes.

The switching and settling times of the programmable attenuator used in FlexRadio are 1.1μs and 2μs respectively. The symmetric antenna placement of FlexRadio decouples the *delay and attenuation block* at each receiver chain from the configuration of FlexRadio. Thus, switching between different configurations of FlexRadio does not require reprogramming the attenuator. None the less, the preconfiguration module and the attenuators used in RF cancellation are tuned periodically to account for changes in circuit behavior due to change in temperatures, humidity etc. However, these tuning requirements are independent from switching FlexRadio configurations and are infrequent.

In our implementation, the maximum switching time occurs when FlexRadio switches from transmission mode 0/4 to 3/1, as the digital cancellation module needs to estimate three channels - between three transmitters to the receiver - in a sequential manner. As indicated in Figure 9, channel estimation time of 7.4μs yields 42dB of digital cancellation in our implementation. Thus the total time to estimate all the channels when FlexRadio switches to 3/1 transmission mode is $\approx$ 22.5μs. The switching time for off-the-shelf RF switches is of the order of tens of nanoseconds. Further, the time to make the required FIFOs available (either a transmit FIFO or a receive data) is of the order of hundreds of nanoseconds. Thus, the maximum time to switch between different transmission modes of FlexRadio is within **25 μs**.

Is the switching time overhead significant? FlexRadioconfiguration changes are motivated by changing topology or flow constraints. Many factors can affect flow constraints. Typical channel coherence time is an ultra-agressive rate estimate of changing topology constraints. However. coherence times even for mobile channels can be hundreds of milliseconds. Thus, under most circumstances, switching between different FlexRadio configurations presents negligible overhead.

5.4 FlexRadio in a network: Experiment setup and evaluation

Having evaluated the effectiveness of FlexRadio's self-interference cancellation strategies and its configurability, in this section, we evaluate the performance of FlexRadio nodes in a network. We perform a set of experiments using different network topologies, flow constraints and channel conditions. We compare the performance of FlexRadio nodes in these networks with the performance of wireless nodes having a fixed functionality (MIMO, full duplex and Multi-User MIMO (MU-MIMO)) in these networks. For fixed full-duplex radios mentioned in this section, half of their RF chains are used for transmission while the rest are assigned for signal reception. So we refer to these as half-half full-duplex.

All modes of radio operation, i.e. FlexRadio operation or fixed function, use standard modulation and coding schemes of WiFi's 802.11g transmissions; 1/2 BPSK, QPSK, QAM16 and QAM64, 2/3 BPSK, QPSK, QAM16 and 3/4 QAM64. All the experiments are conducted in the 2.4GHz ISM band over a bandwidth of 20MHz. Theoretically, FlexRadio nodes should be able to operate on different frequencies as it is based on the symmetry components placement. However, due to the manufacturing limitation of the frequency selective RF components on our PCB board (programmable attenuator, balun and switches), we operate in the 2.4GHz band for which these components have been designed.

5.4.1 FlexRadio in Interference-limited Networks

We evaluate the performance benefits of FlexRadio nodes in interference limited networks as discussed in Section. 3. For this experiment, we place four wireless radio nodes according to the topology as shown in the Figure 1(a). Each radio is implemented on the NI software radio defined platform described in the previous section. For this topology, we compare the performance of FlexRadio nodes with MIMO and half-half full duplex nodes. In both MIMO and full-duplex networks, enabling any two transmission streams simultaneously causes interference at the remaining passive nodes thus preventing another transmission stream. FlexRadio nodes can be configured to make the necessary spatial dimensions (antennas) available to align interference to enable a third stream.

When evaluating FlexRadio nodes in this topology, all nodes compute their channels to neighboring nodes (For instance, node N1, in Figure 1(a), computes the channel between itself and nodes N2 and N3 and so on.). This is required to implement interference alignment. In our implementation, the nodes share the computed channel information over Ethernet. Further, we use the communication backplane of our NI platform to synchronize the distributed nodes in time. There are other techniques in literature to achieve the same requirement [15, 20]. We transmit 200 packets over each enabled transmission stream. We measure the throughput over all active links at their highest possible data rates. We repeat this experiment for 50 different locations of nodes N2 and N3.

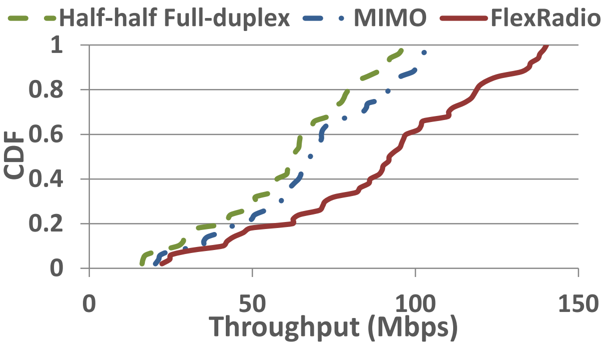

Figure 11: Throughput comparison between MIMO, fixed full-duplex and FlexRadio for the topology shown in Figure 1(a)

Figure 11 plots the CDF of the throughput measured at these locations. FlexRadio outperforms full-duplex and MIMO performance by **47%** and **38%** respectively. This is slightly below the 50% gain anticipated in Section 3. The slight drop in gain can be attributed to the additional channel measurement required between nodes N2 and N3. None the less, the gain is significant over existing MIMO and full-duplex technologies without requiring significant hardware overhead (over full-duplex nodes) or configuration switching overhead.

5.4.2 Adjusting Configuration Based on Flow Demand

We evaluate the benefits of flexibility in networks with flow constraints. We perform this experiment in a three-node network. The radio in the middle has four RF chains. The other two radios with two RF chains cannot hear each other (similar to the topology in Figure 2(a)). We repeat the experiment at 50 different locations to capture different channel conditions. The flow constraint is defined similar to that in Figure 2(a). We measure the throughput for each experiment in a method similar to that described in the previous subsection. We compare the throughput of the network between fixed full-duplex, MIMO, MU-MIMO and FlexRadio nodes. Figure 12(a) plots the CDF of the throughput. We can see that, as expected in section 3, when the middle node operates under 2Tx/2Rx FlexRadio configuration and the other two nodes operate as MIMO receiver (0/2) and MIMO transmitter (2/0) separately the optimal network throughput is achieved. This configuration achieves *twice* the throughput of the other configurations. Note that, for this flow constraint, MU-MIMO does not outperform MIMO.

We repeat the experiment for each of these 50 locations. However, this time we have the middle four RF

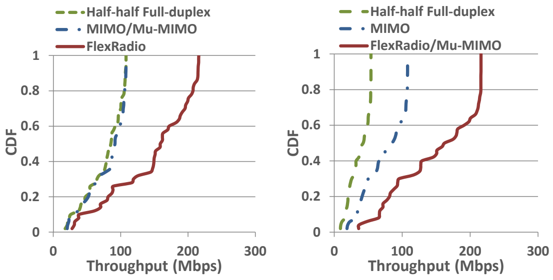

(a) (b)

Figure 12: In the flow demand (a), the four RF chain radio wants to receive some packets from one 2-RF chain radio and transmit to another one. In (b), all the two 2-RF chain radios want to transmit to the middle one.

chain node receive from the other two nodes all the time. We plot the CDF of the throughput distribution for this flow constraint in Figure 12(b). In this scenario, MU-MIMO presents the throughput optimal solution, which is the configuration that FlexRadio adopts. Through this experiment, we verify that FlexRadio enables each node in a network to *adapt* to a configuration that achieves optimal network performance.

5.4.3 Varying Channel Conditions

Finally, we seek to evaluate FlexRadio nodes in different channel conditions. Theoretically, it has been deduced that when the SNR of the channel is low, maximizing the number of RF chains/antennas at the receiver maximizes the throughput [17]. However, in the theoretic perspective, this phenomena is observed at really low SNR (around -20dB), where WiFi transmission does not occur.

None the less, we perform an experiment where two radios with four RF chains wish to transmit to each other. Under a reasonable WiFi channel, the SNR varies around 5dB. At this SNR, all the radios choose the lowest data rate (5.5Mbps) corresponding to 1/2 BPSK. For this experiment, we run MIMO in two configurations: One Stream MIMO and Two Stream MIMO. Under One Stream MIMO, all of the transmitting RF chains send the same data. This is usually the optimal strategy under very low SNR conditions. The Two Stream MIMO is the typical MIMO configuration where two RF chains send different data streams, the normal MIMO operation.

We measure the throughput for this scenario for half-half full-duplex, the two MIMO configurations and FlexRadio (1/3) configuration.

The experiment is repeated 50 times and the CDF of the throughput is plotted in Figure 13. Surprisingly, FlexRadio outperforms other configurations ≈ 85% of the times even when the channel SNR varies around 5dB. On some instances one stream MIMO performs better.

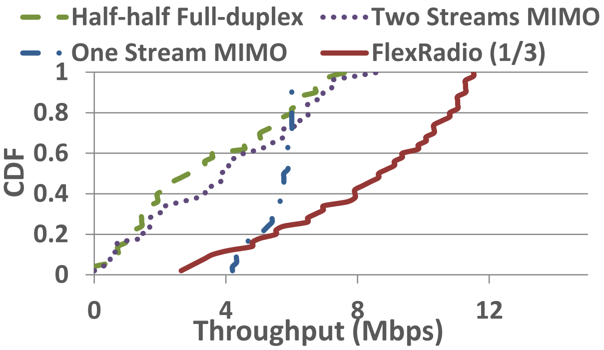

Figure 13: Throughput comparison between MIMO, half-half full-duplex and FlexRadio (1/3) configuration. One stream MIMO refers to the all TX chain in the MIMO transmitter transmit the same data while in the second setting, they are divided into two groups so that two streams are transmitted along the transmission.

Note that, under one stream MIMO, one node transmits the same data on all four RF chains and the other node receives on all of its four RF chains. While for 1/3 FlexRadio, only 3 RF chains are used for receiving by both the nodes. At very low SNR, the received SNR scales linearly with the number of receiver RF chains. This gives one stream MIMO a slight edge since it has one receive RF chain more than 1/3 FlexRadio. On average, FlexRadio provides a median gain of **1.51x over MIMO** and **2.85x over full-duplex.**

6 Related Work

Single RF-chain cancellation techniques. Prior RF cancellation techniques in existing full-duplex implementations [2, 3, 5–7, 9, 11–14] can be broadly classified into: *Passive (self interference suppression) and Active (Antenna cancellation, Analog cancellation)*. Passive suppression techniques provide electromagnetic isolation between the Tx and Rx antennas to minimize self interference, for instance, by using directional antennas, [9]. Active cancellation methods create a null at the receive antenna by sending an inverted copy of the transmitted signal, either over air (Antenna cancellation [6]) or through transmission line (Analog cancellation [12]). Antenna cancellation techniques typically require additional antennas (either for Tx, or Rx or both). FlexRadio's symmetrical RF cancellation design draws from these designs to reduce implementation complexity.

Multi-RF chain full-duplex systems. Recently, many researchers have demonstrated FD capability on multi-RF chains systems [2,3,8]. MIDU [2] employs two-level antenna cancellation. The authors propose a symmetric arrangement of Tx and Rx antennas such that the transmitted signals from a pair of TX antennas are offset by π at a given Rx as well as the received signals at a pair of Rx antennas from a given Tx antenna are offset by π. MIDU needs 2× the number of antennas needed for a MIMO-FD node with the same number of RF chains.

Single-Antenna full-duplex systems. All the above implementations use at least one antenna for each active RF chain. In the case of antenna cancellation, or MIDU, multiple antennas are used *per* active RF chain. However, recent work [4] implements a full-duplex node (with one active TX RF chain and one active RX RF chain) using only a single antenna. This technique uses a circulator to provide isolation between the Tx and Rx paths. They achieve further cancellation using analog cancellation techniques implemented with passive delay lines and variable attenuators on the cancellation signal. The work in [3] extends this full-duplex design to MIMO radios. In [3], the authors implement a *six RF-chain full-duplex node (3 transmit RF-chains and 3 receive RF-chains)* using only 3 antennas. Since an equivalent 3 antenna MIMO node can activate *at most* 3 transmit or 3 receive RF chains, a 3 antenna MIMO node is essentially a 3 RF-chain MIMO node. While the full-duplex design in [3] almost doubles the capacity between two nodes over that of MIMO nodes with the same number of antennas, this comes at the cost of having *more active RF-chains*. On the other hand, FlexRadio exploits flexibility to realize a fundamental performance increase while *not using any additional active RF chains.*

7 Discussion and Conclusion

MAC layer Implications. The examples in Sec. 3 assumed the presence of a central node with knowledge of RF resource capabilities of the all the nodes in the network. Nodes can piggyback information of their RF resource capabilities (in terms of number of antennas, RF chains etc.) with packets exchanging channel state information. For example, in enterprise wireless networks APs can collect information from their respective clients and forward this information to a designated server over the backbone. The server can then determine optimum configuration for all the nodes in the network. Designing algorithms to exploit FlexRadio capability to maximize network performance is an open problem.

Extending beyond four RF chains. The four RF-chain FlexRadio prototype is applicable to many existing MIMO systems (the standard LTE system, for instance). However FlexRadio's design principles can extend to nodes with more than four RF chains. The extended design can leverage the geometrical symmetry of the symmetric antenna placement design to minimize cost, area and power consumption of the FlexRadio node.

The complexity reduction of the RF cancellation is based on the following observation: If multiple transmitter antennas are equidistant to a given receiver antenna, the cancellation signal of these TX chains can be *combined* before passing through a *single delay and attenuation block* to cancel out their *self interference* at the given receiver. If multiple sets of transmitters are equidistant, at different distances, to a given receiver, then the receiver needs an independent *delay and attenuation block* for each such set of transmitters. In general terms, for an N RF-chain FlexRadio system with a biased transmission order defined as: $N_1, N_2, \cdots, N_k$ are the K transmitters in K/N-K mode, the number of *delay and attenuation blocks* is bounded by:

$\sum_{i=1}^{n-1}$ (distinct distances from set $\{N_{i+1}, N_{i+2}, \cdots N_n\}$ to N_i).

This governs the minimum complexity RF cancellation circuitry for our design. We present an antenna placement scheme for FlexRadio system with larger number of RF-chains by simply extending the antenna placement scheme of the four RF-chain FlexRadio node along all sides. For a FlexRadio system with more than four RF chains, additional antennas can be added to the four antenna arrangement using the following priority:

- On the centroids of the triangle formed by the excenter and the vertices of the four antenna arrangement on the side closest to the excenter.
- On the excenters along the three sides of the four antenna arrangement.
- The above steps extend the four antenna arrangement by making copies of its geometry along all its sides. The process can be repeated on the newly created copy until all the antennas corresponding to its respective RF-chains of the node are placed.

Conclusion. FlexRadio is a fundamentally new capability for a wireless node. By choosing the number of RF chains to transmit and receive, network-wide throughput gains are possible. These opportunities can be potentially recognized either centrally or in a distributed fashion. Further, the symmetric antenna placement design of FlexRadio ensures that realization of FlexRadio does not present a significant hardware overhead compared to a full-duplex design with same number of RF chains. Thus, we believe the performance gains of FlexRadio nodes are promising.

Acknowledgements

We would like to thank our shepherd, Lakshmi Subramanian, and the anonymous reviewers for their comments and suggestions to improve our paper. This work is supported by NSF CNS-1254032.

References

[1] Adib, F., Kumar, S., Aryan, O., Gollakota, S., and Katabi, D. Interference alignment by motion. In *MOBICOM* (2013).

[2] Aryafar, E., Khojastepour, M. A., Sundaresan, K., Rangarajan, S., and Chiang, M. Midu: enabling mimo full duplex. In *MOBICOM* (2012).

[3] Bharadia, D., and Katti, S. Full duplex mimo radios. In *NSDI* (2014).

[4] Bharadia, D., McMilin, E., and Katti, S. Full duplex radios. In *SIGCOMM* (2013).

[5] Bliss, D., Parker, P., and Margetts, A. Simultaneous transmission and reception for improved wireless network performance. In *SSP* (2007).

[6] Choi, J. I., Jain, M., Srinivasan, K., Levis, P., and Katti, S. Achieving single channel, full duplex wireless communication. In *MOBICOM* (2010).

[7] Duarte, M., and Sabharwal, A. Full-duplex wireless communications using off-the-shelf radios: Feasibility and first results. In *ASILOMAR* (2010).

[8] Duarte, M., Sabharwal, A., Aggarwal, V., Jana, R., Ramakrishnan, K. K., Rice, C. W., and Shankaranarayanan, N. K. Design and characterization of a full-duplex multi-antenna system for wifi networks. *CoRR abs/1210.1639* (2012).

[9] Everett, E., Duarte, M., Dick, C., and Sabharwal, A. Empowering full-duplex wireless communication by exploiting directional diversity. In *ASILOMAR* (2011).

[10] Gollakota, S., Perli, S. D., and Katabi, D. Interference alignment and cancellation. *SIGCOMM* (2009).

[11] Hong, S. S., Mehlman, J., and Katti, S. R. Picasso: flexible rf and spectrum slicing. In *SIGCOMM* (2012).

[12] Jain, M., Choi, J. I., Kim, T., Bharadia, D., Seth, S., Srinivasan, K., Levis, P., Katti, S., and Sinha, P. Practical, real-time, full duplex wireless. In *MOBICOM* (2011).

[13] Khojastepour, M. A., Sundaresan, K., Rangarajan, S., Zhang, X., and Barghi, S. The case for antenna cancellation for scalable full-duplex wireless communications. In *HoTNets-X* (2011).

[14] Radunovic, B., Gunawardena, D., Key, P., Proutiere, R., Singh, N., Balan, V., and Dejean, G. Rethinking indoor wireless mesh design: Low power, low frequency, full-duplex. In *WiMesh* (2010).

[15] Rahul, H., Hassanieh, H., and Katabi, D. Sourcesync: a distributed wireless architecture for exploiting sender diversity. *SIGCOMM* (2011).

[16] Sen, S., Choudhury, R. R., and Nelakuditi, S. Csma/cn: carrier sense multiple access with collision notification. *IEEE/ACM Transactions on Networking (TON) 20*, 2 (2012), 544–556.

[17] Tse, D., and Viswanath, P. *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.

[18] Xie, X., and Zhang, X. Does full-duplex double the capacity of wireless networks? In *INFOCOM* (2014).

[19] Yang, Y., Chen, B., Srinivasan, K., and Shroff, N. Characterizing the achievable throughput in wireless networks with two active rf chains. In *INFOCOM* (2014).

[20] Yenamandra, V., and Srinivasan, K. Vidyut: exploiting power line infrastructure for enterprise wireless networks. In *SIGCOMM* (2014).