FastForward: Fast and Constructive Full Duplex Relays

Dinesh Bharadia
Stanford University
dineshb@stanford.edu

Sachin Katti
Stanford University
skatti@stanford.edu

ABSTRACT

This paper presents, **F**ast**F**orward (FF), a novel full-duplex relay that constructively forwards signals such that wireless network throughput and coverage is significantly enhanced. FF is a Layer 1 in-band full-duplex device, it receives and transmits signals directly and simultaneously on the same frequency. It cleanly integrates into existing networks (both WiFi and LTE) as a separate device and does not require changes to the clients. FF's key invention is a constructive filtering algorithm that transforms the signal at the relay such that when it reaches the destination, it constructively combines with the direct signals from the source and provides a significant throughput gain. We prototype FF using off-the-shelf software radios running a stock WiFi PHY and show experimentally that it provides a $3\times$ median throughput increase and nearly a $4\times$ gain at the edge of the coverage area.

Categories and Subject Descriptors

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

Keywords: Full Duplex Relay; Cancellation; Full Duplex ; low-latency cancelation

1. INTRODUCTION

We have all often experienced perplexingly poor wireless performance. For example, it's not uncommon to find that one's connection is flaky and offers very low throughput even when one is the only user of the WiFi AP in a home. Similarly, for LTE networks, even at nights when the network is lightly loaded, performance can be poor indoors or in urban concrete jungles, with raw link speeds varying between a few hundred Kbps to a couple of Mbps. This is despite continuous evolution of wireless standards over the last few years to provide very high link bitrates. For example, the 802.11ac WiFi standard promises bitrates of up to 1.3Gbps, while LTE downlink speeds are expected to be up to 300Mbps [5, 1]. These gains are coming from two factors: use of higher modulation (up to 256QAM for both LTE and WiFi) and higher MIMO spatial multiplexing (up to 4 parallel streams for both LTE and WiFi). Both these features should work well when there is little to no contention/interference and a single or a few users are connected to the WiFi AP or the LTE basestation. Yet often users don't realize these benefits in practice, experiencing raw speeds that are one to two orders of magnitude less than the advertised speeds.

There are two fundamental reasons for the poor performance described above: propagation loss and MIMO rank degradation. Propagation loss is a natural and expected cause of the drop in link rates. Fig. 1 shows a typical 2000 sq. ft. home with a WiFi AP at one corner of the house in the living room. We model propagation and other effects using commercial grade wireless ray propagation modeling software [21, 4] that is used for planning wireless deployments. As we can see, except for the immediate area around the AP, most of the coverage area in the middle of the home experiences SNRs between 10-15dB (as seen in Fig. 1), and at the edge the performance is worse, with SNRs between 0-6dB. This cuts down the highest modulation that can be used to QAM/16-QAM from 256-QAM, a $4\times$ reduction in bitrate. An analogous argument can be made for LTE networks where the coverage area is larger, and signals often have to propagate through large buildings in urban areas which further cause signal loss due to shadowing effects.

The second fundamental reason is MIMO rank degradation as seen in Fig. 2. To send multiple data streams via MIMO spatial multiplexing, the channel between the AP and the client needs to have several independent strong propagation paths available (in other words, the MIMO channel matrix needs to be full rank and have strong eigenvalues [9]). But in most indoor and urban scenarios, often we find that only a single strong path exists between the AP and the client, and the rest are weak or non-existent. This happens because of the geometry of homes, offices and hotels which typically have a single or few corridors with rooms off the corridors. The corridor acts like an RF pinhole [9, 17] since it is typically the only strong path available between the AP and the client, and focuses all the signals to go through a single path which makes all of the paths correlated at the destination. The consequence is that the MIMO channel rank is reduced, and the AP cannot send multiple independent streams, reducing the bitrate significantly. LTE signals behave analogously, in that the only path indoors for the signal is through windows or doors (walls tend to block signals almost completely), and the doors/windows acts as RF pinholes. Combined with the propagation loss described above, this results in nearly a 6-10x reduction in bitrate in the middle and edge of the coverage from the AP.

Our goal in this paper is to design a general, practical and easily deployable system that provides high-throughput uniform wireless coverage. By general, we mean the fundamental technique should be applicable to any OFDM based standard. By practical and easily deployable, we mean that the system should require minimal to no changes to the existing infrastructure of APs , clients and/or standards.

We design and implement a novel system called **F**ast**F**orward (FF) that achieves the above goals. FF's core operation is simple to describe. It is a single device that operates independently listening to the signal from the source, digitizing it to IQ samples, processing it by passing the IQ stream through a filter (in both digital and RF domains), and up-converting and amplifying the processed IQ stream to RF signals that are then transmitted to the destination on the same frequency. The filtering and amplification are done in such a way that the SNR of the signal at the destination is signifi-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGCOMM'14, August 17–22, 2014, Chicago, Illinois, USA.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-2836-4/14/08…$15.00.
http://dx.doi.org/10.1145/2619239.2626327 .

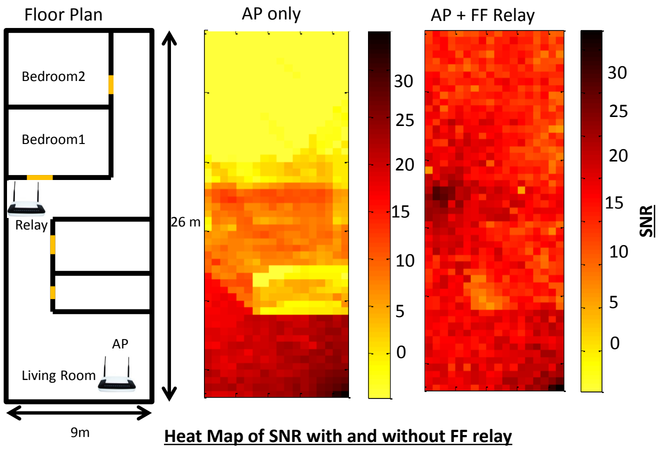

Figure 1: Heatmap of SNR with AP alone and with AP and FF relay. A majority of the home has poor SNR due to propagation loss in the AP only scenario.

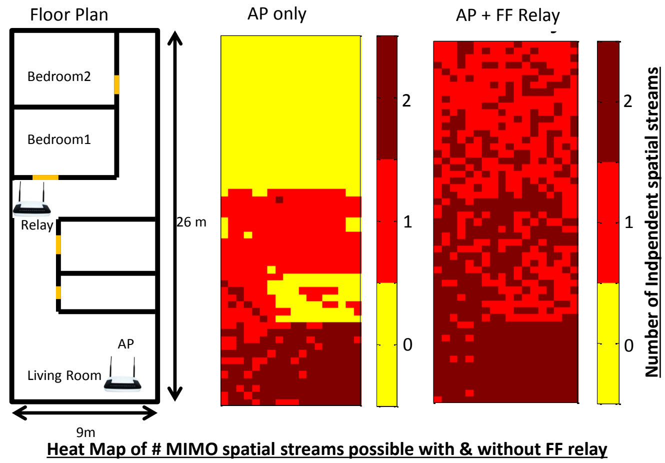

Figure 2: Heatmap of number of MIMO spatial streams possible with AP alone and with AP and FF relay. A majority of the home has poor MIMO channel rank due to pinhole effects and poor link propagation through walls.

cantly increased and the number of independent MIMO paths at the destination is also increased, enabling a significantly higher bitrate. Thus it acts as a controlled strong multi-path creator of the signals that is completely transparent to the AP and the client, they do not even realize that an FF device exists. A glimpse of its performance is shown in the heatmaps (Figs. 1,2).

The key insight behind FF is a novel technique that we invent called **construct-and-forward full-duplex relaying**. The basic idea is best described in terms of a single SISO transmission from the AP to the client. With a simple full-duplex amplify-and-forward relay that has been discussed in the literature [20], the client would receive two signals: one directly from the AP and the other amplified version from the relay. A naive implementation of the relay will result in both these signals acting as destructive interference to each other, and the relay potentially amplifying noise. FF's innovation is to control the properties of the relayed multi-path signal to in fact turn such potential interference into a constructive SNR gain. The design relies on the fact that if an OFDM receiver receives multiple reflected copies of a signal, then as long as the reflections are within the OFDM cyclic prefix (CP) interval (around 400ns for WiFi and 4.69μs for LTE), they do not cause inter-symbol interference (ISI) to each other. If we can ensure that the processing delay through the FF relay is minimized so that the relayed signal does not fall outside the OFDM CP interval at the receiver, we can achieve no inter symbol interference. FF's low latency cancellation technique achieve this purpose. Thus, FF's relay acts as an amplified multi-path component at the receiver.

While limiting the processing delay ensures that inter-symbol interference is avoided, it still does not provide a constructive SNR gain. The second aspect of construct-and-forward relaying is to intelligently process the received signal at the relay before transmission such that the relayed signal adds up constructively with the other signals that the destination is directly receiving from the source to significantly enhance the effective SNR. The basic idea is that the relay first collects the channel state information about three links: source-relay, relay-destination and source-destination. Now, when it receives the transmission from the source, it passes the signal through a filter such that cumulative effect for the received signal at the destination (which has now gone through the channel from the AP to the FF relay, the filter at the relay and the channel from the relay to the client) is such that it adds coherently (in almost complete alignment) at the destination with the direct signal received by the destination from the source. Fig. 5 shows the effect visually, the relay rotates the incoming signal such that it aligns up with the vector representing the channel between the source and the destination. The constructive addition significantly increases the SNR at the destination (client), enabling a higher bitrate to be used by the source (AP). A similar effect happens when the FF relay is combating the pinhole effect, it computes a filter that increases the number of spatial streams and the SNR at the destination (client), enabling the source (AP) to use a higher level of spatial multiplexing and therefore higher bitrates. Note that the relay can be used to improve the link from the client to the AP as well.

The challenge in realizing such construct-and-forward relaying while ensuring that processing delays is much smaller than the OFDM CP is the full duplex nature of the relay. The FF relay is transmitting and receiving signals at the same time on the same frequency. Further, the transmitted signal is essentially a slightly delayed and amplified version of the received signal. To receive the signal from the AP, the FF relay has to cancel the transmitted signal. The amount of cancellation puts a limit on the amount of amplification that we can apply at the relay, since if we amplify more than the cancellation, residual signal is left over and is recycled for transmission, creating an unstable positive feedback loop. Maximizing the amount of cancellation is therefore crucial to maximizing amplification. However, unlike prior work on full duplex, the cancellation has to be performed within a time budget as small as possible (e.g. within 100ns for WiFi since the CP is only 400ns long) to ensure that the relayed signal can take advantage of FF's constructive relaying capability. A second key contribution of this paper is a novel cancellation technique for relays that achieves nearly 110dB of cancellation while operating within a processing time budget of 100ns.

We design and implement FF on the WARP radio platform [8] and by designing our own self-interference cancellation RF boards. We evaluate FF in an indoor testbed and show that FF provides a 3$\times$ median increase in throughput and nearly 4$\times$ at the edge of the coverage area. The gains come from different aspects for different clients. For clients with decent SNR already, the gains come from MIMO rank expansion. For clients at the edge of the coverage area where the SNR is already quite poor, the gains come from the SNR gain constructive relaying provides. We also compare against the half duplex packet-level relay (e.g. the Apple Airport Express) and show than FF provides at least 2$\times$ better throughput and coverage.

2. RELATED WORK

A natural question is whether there are other approaches that can be used to solve the problem of coverage and capacity that FF aims

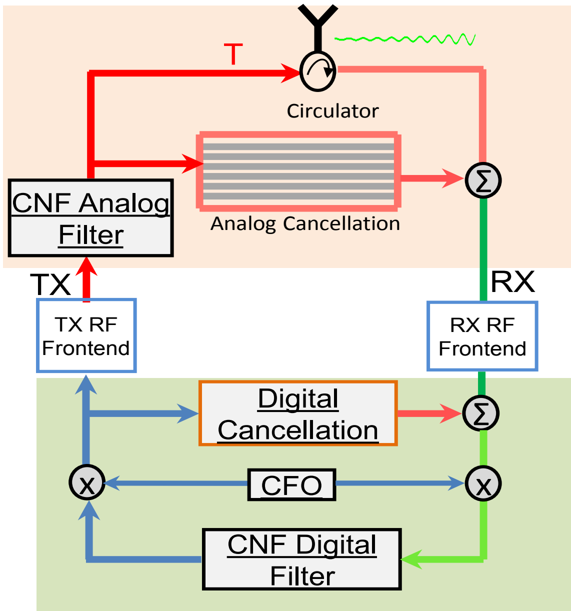

Figure 3: Overall Block Diagram of a FF relay. There are two key pieces: construct-and-forward (CNF) analog and digital filters, and self-interference cancellation.

to? There has been of course a large body of work in recent years that have proposed several PHY and MAC layer enhancements to increase network capacity and robustness, FF however is operating on signals directly and is therefore orthogonal to those approaches.

However there is one approach that could help and is immediately deployable: a half-duplex mesh router like the Apple Airport Express. These devices help extend WiFi coverage by connecting to the AP as a client, and then turning around and transmitting to the actual client in the next slot (hence the name half duplex router). Theoretical literature on relaying refers to such techniques as **decode-and-forward** relaying. However, as we show in Sec. 5 these devices do not provide capacity gains except in the edge of the coverage area. This is because they essentially require close to twice the number of time slots for transmitting the same amount of traffic. Further for clients with decent SNRs to the AP, the half-duplex mesh router is a bad option, it is better to have a single-hop medium-SNR link rather than using two hops over high-SNR links.

There are several products in the market that are called repeaters. These devices are simple **amplify-and-forward** relays. They receive a signal, and then immediately amplify it and transmit it. Such devices are available for both WiFi and LTE networks. However these devices cannot amplify too much, they are severely limited by the amount of isolation between the signals they are receiving and relaying as we show in Sec. 3.5. Second, since they are blindly amplifying signals, they amplify noise and often hurt performance as we show in Sec. 5.5. FF also belongs to the class of **amplify-and-forward** relays, however this paper makes three novel contributions:

- FF is selective and smart about relaying, it exploits the knowledge of channel state information to intelligently filter and amplify signals such that they appear as a constructive multipath component at the destination, rather than increase noise and/or add up destructively like a standard repeater would.

- FF designs a novel low-latency self-interference cancellation technique which ensures that relayed signals fall within the CP for OFDM signals and do not cause inter-symbol interference. The technique is applicable to standard repeaters too and they can benefit from being able to use a higher amplification factor due to the increased amount of cancellation.

- This paper also provides a full design, implementation and evaluation of full-duplex relays, to the best of our knowledge we are not aware of prior work that provides an experimental characterization of how well other kinds of relays work in practice.

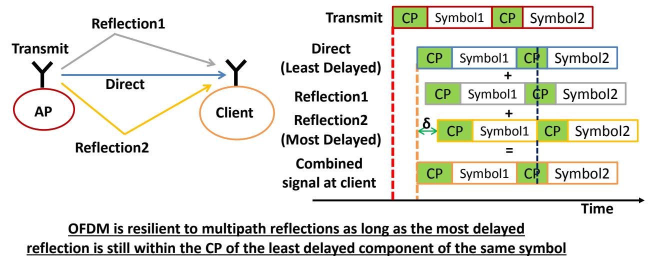

Figure 4: OFDM is resilient to multipath reflections as long as the extra delay experienced by the slowest reflection compared to the quickest arriving signal at the destination is less than the cyclic prefix (CP).

Finally, there is a large body of theoretical work on relays in the information theory literature [24, 14, 27]. Starting from early work by Shannon, there have been several proposals on relaying [13, 22, 12, 26]. Apart from the amplify-and-forward and decode-and-forward relaying techniques; a third well known class of techniques is **compress-and-forward**: this is an intermediate version between the above two relays. Here the relay may not decode the entire packet, but only the symbols and re-encodes them in a more efficient way [30, 16, 29, 19]. The destination has to combine the relayed information with the direct transmission from the source to recover the original packet. This method is typically quite complex to implement since it requires changes at the client with techniques such as soft interference cancellation and combining, as well as sophisticated processing at the relay.

3. DESIGN

FF is a layer 1 full-duplex relay, i.e. it receives signals from the source, processes them both in the analog and digital domains, and then converts them back to RF signals and transmits them on the same channel they were received on. Fig. 3 shows the high-level block diagram of an FF single-antenna relay. Note that an FF relay can have multiple antennas and can relay MIMO signals, however we use the single-antenna SISO FF relay for describing the key ideas in a concise manner. However the techniques and algorithms naturally translate to a MIMO relay implementation.

As we can see there are three main components in the design: cancellation, constructive filtering (CNF) and amplification. The insight underpinning these components is exploiting the structure of OFDM such that relaying can produce a constructive SNR gain at the receiver. We start by describing first the basics of OFDM.

3.1 OFDM Background

OFDM was introduced to combat the negative effects of multipath and inter-symbol interference. The basic idea is widely known and described in textbooks [15], but we include it here because it helps explain some of FF's algorithmic design choices later.

The basic idea of OFDM is to divide the available bandwidth B into N smaller subcarriers (e.g. 802.11ac with 80MHz bandwidth is divided into 512 subcarriers whereas LTE divides into subcarriers of width 15KHz). Each subcarrier can be conceptually treated as an independent orthogonal channel carrying independent symbols. Hence the symbol time is N/B, i.e. the symbol is N times longer, as compared to a typical communication system transmitting symbol at $1/B$, for bandwidth B. Further to each symbol, a

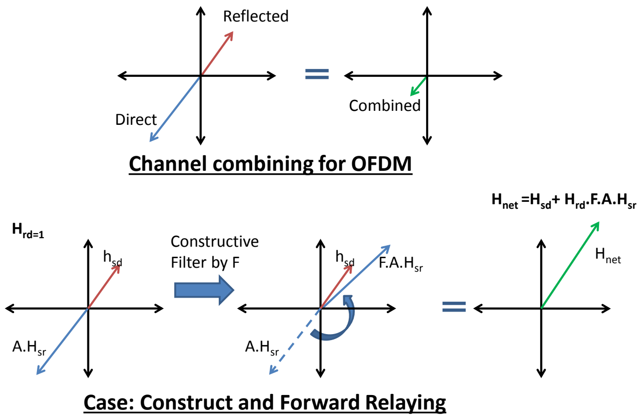

Figure 5: FF's construct-and-forward relaying rotates the relayed signal such that it aligns with the direct signal from the source to the destination and provides a constructive SNR gain. The top figure shows what happens with normal OFDM where instead of the relay there is a normal reflection of the same delay. The channel gains add up destructively and reduce SNR at the destination.

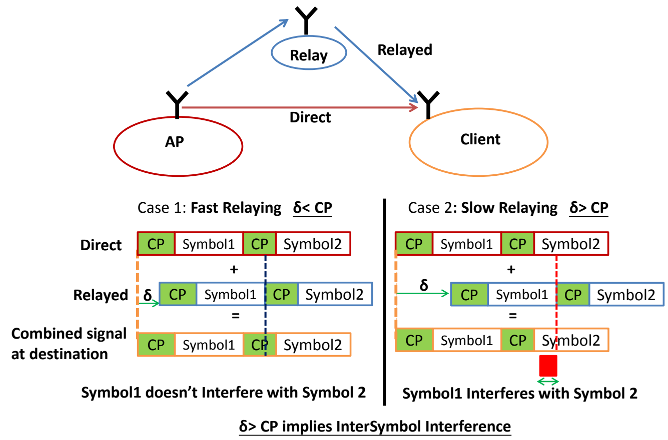

Figure 6: Low latency processing at the FF relay is critical. If the delay of processing in the FF relay is greater than the OFDM CP, then the relayed signal will cause inter-symbol interference at the destination [15].

guard period known as the cyclic prefix (CP) (typically 25% of the symbol time) is added. Hence as long as the extra delay of a multipath reflection of an OFDM symbol w.r.t. the first arriving version at the destination is less than the CP length, no inter-symbol interference is caused as seen in Fig.4. The length of the cyclic prefix is 400ns in WiFi and 4.69μs in LTE. Hence in WiFi there is tolerance for a distance spread of 400 feet whereas for LTE its almost 5000 feet, which is expected since WiFi is designed for covering homes whereas LTE is designed for covering larger outdoor areas.

Given the above fact, how does the effective channel look at the receiver? In other words how do the multipath reflections add up if they are not causing ISI with each other? To visualize this, consider Fig. 5. We are plotting the channel gains for a single OFDM subcarrier (i.e. the attenuation and phase shift applied by the direct path channel to any signal on that subcarrier). Now suppose there is another multipath reflection with a slightly longer path and higher attenuation. The channel gain for this second path shows up as a second vector that is rotated w.r.t the first channel gain. Assuming the extra delay is within the CP, the overall channel perceived by the receiver is the sum of these two channel gains. The effective SNR at the client therefore depends on the relative orientation and gains of the direct and reflected channel paths, if they are aligned with each other in the same direction SNR increases, if they are in opposite directions SNR decreases.

3.2 Construct-and-Forward Relaying

FF's construct-and-forward relaying builds on top of OFDM. Our basic insight is to make FF look like another strong multipath reflector, albeit with the ability to amplify and modify the signals. Since FF operates at the signal level, at the receiver the signal from the relay looks like yet another multipath component, albeit a strong one. As long as the extra delay of this component is within the CP, the receiver will not perceive any inter-symbol interference. The constraint then is that the overall delay of the signal going through the FF relay has to be as low as possible, and definitely well within the CP interval. Since we still have to account for normal propagation delay from the source to the relay and then from the relay to the destination, ideally we want to completely minimize the processing delay in the FF relay.

As we see in Fig. 6, by minimizing the relative delay below cyclic prefix between direct and reflected (or relayed) path we can avoid inter-symbol interference. However, depending on the relative phase of the channel gains from the relay to the destination (relayed) and from the source to the destination (direct), we might hurt overall SNR at the receiver as shown in of Fig. 5.a. So might FF be hurting the SNR by relaying ?

FF's key invention is a novel technique that leverages its relaying capability in a way to actually significantly enhance the SNR at the client. Remember that the relay has the opportunity to modify the signal before it amplifies and sends it to the destination. FF's novel idea is to apply a filter before amplifying and relaying the signal such that it adds up *constructively* at the destination to maximize the SNR gain. Mathematically, let us say the channel from the source to the destination is h_{sd}, and from the source to the relay is h_{sr} and from the relay to the destination is h_{rd}, for a particular subcarrier. Further the noise at the destination is n_d, and at the relay is n_r. The relay would amplify the signal by a factor A and then pass the signal through a constructive filter whose response is F at that subcarrier. The SNR at the destination for that subcarrier, is given by

$$SNR_d = \left| \frac{h_{sd} + h_{rd} F A h_{sr}}{N_o} \right|^2 \quad (1)$$

where $N_o = n_d + h_{rd} F A n_r$. The second term in N_o is small, since the amplification (A) is controlled as described in Sec. 3.5, which makes sure that noise is not amplified at the destination. For now we will assume the controlled amplification is represented by, $A < A_{max}$ and we will ignore N_o in optimization of Eq. 1. Visually this is demonstrated in Fig. 5.

Note that the constructive filter can introduce additional processing delay, however as before the overall delay still has to be well within the CP so that we can take advantage of OFDM. Further constructive relaying assumes that the relay knows all three channels. The channels from the source to the relay and from the relay to the destination are easy to measure by the relay itself. However the channel from the source to the destination cannot be measured by the relay and has to be explicitly fed to it. We discuss in Sec. 4.2 how this can be done in both WiFi and LTE using existing mechanisms in the standards.

The above discussion has focused on the SISO case. However the same arguments hold for the MIMO case. In effect the relay adds a separate independent strong MIMO path which increases the rank of the MIMO matrix. For constructive relaying, instead of optimizing the above equation, the relay would perform the following optimization

$$\max_{F,A} \quad \det(H_{sd} + H_{rd} F A H_{sr}) \quad (2)$$
$$\text{subject to} \quad A < A_{max}$$

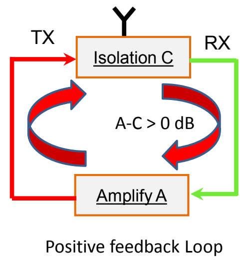

Figure 7: Amplification A is limited by the amount of isolation C. Amplifying more than the isolation implies there is still some residual left over after isolation by C dB, which is then again amplified and relayed in the next time instant and so on. This creates an unstable positive feedback loop.

where H_{sd} is an $N \times M$ channel matrix where the source and destination have M and N antennas respectively, H_{sr} is a $K \times M$ channel matrix to the relay (the relay has K antennas) and H_{rd} is an $N \times K$ matrix, A is again the scalar amplification factor (power) and F is the constructive filter which is a $K \times K$ rotation matrix in this case. Intuitively, the path through the relay acts as a strong independent MIMO path and adds rank to the overall matrix. Since a K antenna relay has only K dimensions, it can increase the MIMO rank at the destination at most by K. The filter again in this case acts as a mechanism to maximize the SNR. The optimization problem described in Eqn. 2 is non-convex and is solved using non-linear optimization technique. Note that it can be solved for $F.A$ as a single variable, and only needs to be solved whenever any of the three channels are updated, and not for every packet. The solution to this problem is referred to as $H_c(f_i)$ in the later sections, overall filter response is referred as H_c.

The takeaway from the above algorithm is that FF needs to implement two key blocks: **amplification** and **constructive filtering**. Note that both these blocks need to be as low latency as possible, ideally within a 100ns budget given that the WiFi CP is 400ns. If we can design it with that delay then the techniques will work for LTE too since it has a longer CP. We turn to the design and implementation of these blocks next.

3.3 FF: Low-Latency Amplification

As we saw in the previous section, FF enables constructive relaying by applying an amplification A and a filter F to the received signal at the relay. As expected the relay cannot receive a signal if it is also transmitting an amplified signal at the same time on the same frequency. Hence to build a relay we need to isolate the received signal from the transmitted signal, i.e. remove the transmitted signal from the received signal. Further, the amount of isolation directly dictates how much amplification the relay can apply on the received signal, which in turn dictates how much the relay expands the range and capacity of the network.

To see why, consider what happens if we amplify beyond the achievable isolation as seen in Fig. 7. In effect this means that some of the signal that is being transmitted is still left over in the received signal after isolation since amplification is greater than isolation. But remember that the transmitted signal is simply a delayed version of the received signal. So in the next instant the transmitted signal would contain a copy of the transmitted signal that was left over in the previous instant. This iteratively accumulates and creates a positive feedback loop where ultimately the relayed signal simply consists of leftover copies of the same signal from previous time instants. The positive feedback loop is unstable and ultimately leads to poor performance. On the other hand if the amplification is lower than the isolation, then all of the transmitted signal is removed from the received signal, and the relay operation proceeds smoothly.

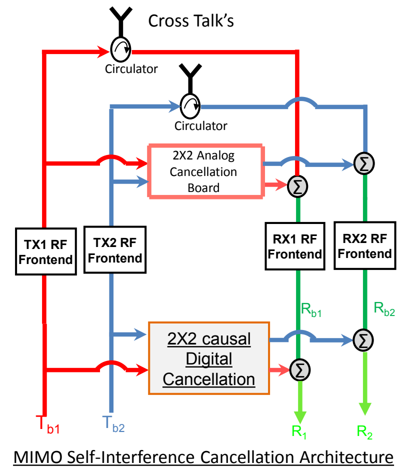

Figure 8: Self-interference cancellation architecture for a 2×2 MIMO FF relay.

Our goal therefore is to maximize the isolation from the TX to the RX. We turn to recent work on self-interference cancellation for full duplex radios [11, 10] to provide the isolation between RX and TX signals. These techniques enable a radio to almost completely cancel the transmitted signal and enable clean reception of the received signal. However there is a catch which prevents us from being able to directly apply the cancellation techniques, the self-interference cancellation has to be performed with as little latency as possible (e.g. much smaller than 400ns for WiFi signals). Self-interference cancellation in the prior work has two components, an analog and a digital cancellation stage. Analog cancellation has negligible delay (around 10ns). However the digital cancellation stage (including the ADC and DAC delays) has a delay of nearly 400ns which would put us out of range for the relay requirements for WiFi. The ADCs and the DACs contribute around 50ns of delay, hence the digital cancellation stage adds nearly 350ns of delay.

We invent a novel self-interference cancellation technique that performs the cancellation with a near-zero delay (excluding than the latency of implementation, which is a few ns or lesser). In prior work on cancellation, the delay is primarily due to the fact that digital cancellation is non-causal [11]. In other words, the digital cancellation filters like to peek ahead into the future of the signal and use that information to cancel the signal at the present. In this relay, we could do this by buffering the received signal, so when we are canceling the self-interference signal at any instant, we know the future of the transmitted signal is going to be. However buffering means delay, for example buffering even 5 digital IQ samples at a 100Msps sampling rate means a delay of 50ns. Hence in FF, we invent a digital cancellation technique that is causal, i.e. it only uses information about what has been already transmitted to cancel the self-interference and does not do any buffering of the received signal before transmission. So received samples are passed in a streaming fashion to the transmit side without any delay. However causal cancellation results in digital cancellation filters which are slightly longer, they need to use more taps to recreate the self-interference for cancellation. However these taps do not add delay, they are for signal samples that have already been transmitted.

Fig. 8 shows the cancellation architecture for a 2×2 MIMO FF relay. Analog cancellation is implemented as discussed in prior work [11, 10] using a tunable FIR analog filter. Digital cancellation is slightly different, it uses a FIR filter like before but there is no buffering and delay, it is a causal filter as shown in Fig. 9.a. The samples that are used for cancellation are only the samples that are

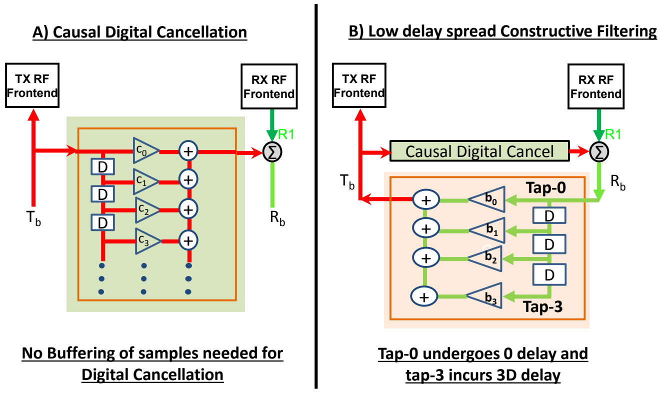

Figure 9: a) Digital cancellation in FF is causal, i.e. the cancellation is performed only using the current and past transmitted samples. No buffering of received samples is performed, which minimizes processing delay through the relay. b) The larger the number of tap delay, higher is the probability that it would cause inter-symbol interference at the destination.

currently being or have already been transmitted, indicating causality.

The coefficients for both the analog and digital cancellation filter are dynamically tuned to maximize cancellation. However, dynamically tuning cancellation in a full duplex relay is more complex than standard full duplex. The reason is that the signal that is being transmitted is a slightly delayed version of the signal being received. To see how this impacts the tuning algorithm, we can look at what happens during analog cancellation. The cancellation problem is given by:

$$\begin{aligned} y(t) &= x_R(t) + h(t) * x_T(t) - \widehat{h}(t) * x_T(t) \\ &= x_T(t+\tau) + h(t) * x_T(t) - \widehat{h}(t) * x_T(t) \end{aligned}$$

where $x_R(t)$ is the signal relay is receiving from the source, $x_T(t)$ is the signal the relay is transmitting to the destination, $h(t)$ is the time domain transformation applied by the channel before the transmitted signal from the relay causes self-interference to the received signal, $\widehat{h}(t)$ represents the filter that is being used by the analog cancellation block to approximate H and implement cancellation, and $y(t)$ is the combined signal that is being received by the relay. Clearly cancellation is maximized when $h(t) = \widehat{h}(t)$. In the second part of the above equation we substitute $x_R(t)$ with $x_T(t+\tau)$ because the relayed signal is a future version of the received signal, where the delay is represented by τ.

Prior work on analog cancellation solve the above estimation problem in the frequency domain. So the above problem can be rewritten in the frequency domain as:

$$\begin{aligned} Y(f) &= X_R(f) + H(f)X_T(f) - \widehat{H}(f)X_T(f) \\ &= \alpha(f)X_T(f) + H(f)X_T(f) - \widehat{H}(f)X_T(f) \\ &= \{\alpha(f) + H(f)\}X_T(f) - \widehat{H}(f)X_T(f) \end{aligned}$$

Where $\alpha(f) = exp(j2\pi f\tau)$ The above equation shows why correlation is a problem, in effect its quite likely that the tuning algorithm adapts $\widehat{H}(f)$ to approximate $\alpha(f) + H(f)$ which will end up canceling the received signal from the source too! We may end up with no received signal at the relay in this case.

To solve this challenge, we invent a novel cancellation tuning mechanism: we artificially inject Gaussian noise at a very low power, which is similar to the transmitter noise of the transmission, only this is known to us (30dB below the transmitted signal or 80dB above the noise floor in the worst case). Gaussian noise only undergoes the channel $H(f)$, as it is not part of received signal. Hence to figure out the response $H(f)$, i.e. the tuning parameters, we compute the correlation of the received signal with the Gaussian noise that was transmitted, and estimate the self-interference channel parameters. However once cancellation is tuned, we know that analog cancellation provides around 70dB of cancellation, and digital cancellation takes both the transmitted signal and Gaussian noise as input to eliminate all the remaining self-interference. So as soon as the cancellation is turned on, all of the Gaussian noise is canceled and is not left over in the canceled signal. Finally this injected noise doesn't affect the client data rate (since the maximum SNR required is 28dB for the highest data rate) and very likely by the time the relayed signal reaches it, the injected noise is quite likely attenuated to below the receiver noise floor of the client's receiver.

Experimental Results: We prototype the above cancellation design using WARP software radios with setup similar to the one used in [10], which is used in the evaluation Sec. 5. We experimentally evaluated the amount of cancellation when the FF relay node is placed at different locations in our indoor testbed, while its receiving the signal from another and re-transmitting the same signal after the constructive filtering. We observe that our design consistently achieves between 108-110dB of cancellation. Note that the maximum cancellation expected is 110dB, since the maximum transmit power is 20dBm and the noise floor is -90dBm.

3.4 FF: Low-delay Constructive Filter

As noted before, the relay can apply a filter such that the relayed signal adds up constructively at the receiver, as seen in Sec. 3.2. A typical implementation of this filter consists of a series of delay lines, each with its own gain, as shown in Fig. 9.b. Note that the signal at Tap-N of the filter ($N = 3$ in Fig. 9.b.) has an ND extra delay with respect to the signal at Tap-0, where D is the delay introduced by each tap. It is important to note that we have a constraint on the number of taps we can employ in our filter because the filter delay ND (which dominantly dictates the maximum delay at the relay) needs to be such that the relayed signal does not fall outside the cyclic prefix at the destination. This section describes how the ideal filter H_c can be implemented with as less filter delay as possible.

Recall that the basic intuition behind this filter is to rotate (i.e., change the phase of) the relayed signal such that it aligns with the direct signal at the destination, as we saw in Fig. 5. For example, to rotate a relayed signal at 2.45GHz by 90 degrees, the constructive filter needs to introduce a 100ps delay (400 ps is the time period of one wave at 2.45GHz which corresponds to 360 degrees, hence 100ps corresponds to 90 degrees). It is extremely hard to implement such fine-grained delays on the order of a hundred picoseconds in the digital domain. For example, if we have a sampling bandwidth of 100MHz, successive digital IQ samples are spaced 10ns apart, in other words two orders of magnitude greater than the delay resolution desired. Figuring out the minute variation in the signal that is delayed by 100ps (which is an intermediate point between two consecutive digital samples) is possible, but extremely complex [28, 18] and defeats our filter delay requirement. The reason is that figuring out the value that an analog signal will take at an intermediate point between digital samples requires us to use sinc interpolation that spans many more future and past digital samples. Using a large number of past digital samples implies that our filter needs to have a large number of taps, which in turn increases the filter delay and thus increases the chances that the relayed signal falls outside the OFDM CP at the destination.

To tackle this problem, FF therefore designs a programmable analog filter that can provide the fine-grained delay adjustment constructive filtering needs as seen in Fig.10, without introducing significantly delay multi-path. We design a tunable analog FIR filter structure with four fixed delays and tunable gains on each delay.

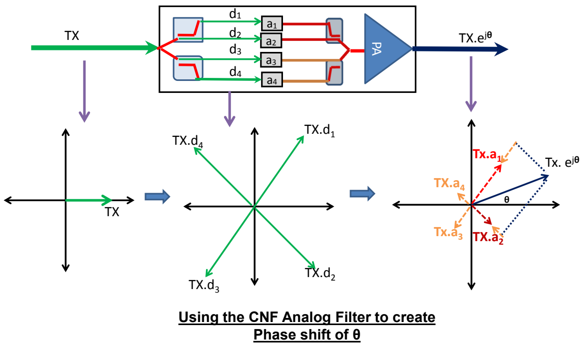

Figure 10: FF's constructive analog filter. The filter enables us to rotate the input signal TX by a fixed angle θ by appropriately adjusting the gains on the four taps of the analog filter. The four taps are placed 90 degrees apart, which at 2.45GHz implies that the tap delays are in increasing multiples of 100 picoseconds.

The delays are spaced 100 picoseconds apart (quarter wavelength of center frequency). To delay a signal by some intermediate value (between 0 to 400ps), the signal is split and passed through all the taps and the gains applied on each tap are adjusted such that the eventual signal has the right phase. Fig. 10 shows the basic idea with four delay lines separated by 100 ps and tunable gains on each line. The incoming signal is at 2.45GHz, hence the two copies of the signal after going through the filter have a relative phase shift of 90 degrees. Now, by adjusting the gain on each delay line, we can rotate the vector to any intermediate phase between 0 and 90 degrees. FF's constructive analog filter applies the same idea using 4 delay lines and spans the entire 360 degrees.

However the above discussion applies to only a single subcarrier, but the signals we are relaying are wider band and have multiple subcarriers. The challenge is that typically each subcarrier needs a different phase shift because channels are frequency selective. The analog filter applies the same delays to all subcarriers, so almost all of them will be rotated by the different phase shift and which wont lead to constructive filtering on all the subcarriers.

To tackle this challenge, we use a pre-filter that is implemented in the digital domain as seen in block diagram Fig. 3 (called as CNF Digital Filter). The intuition is that this pre-filter pre-rotates the phase in each subcarrier by different amounts such that after the analog rotation occurs, all the subcarrier phases are almost lined up for constructive relaying. Note that this rotation in digital is coarse on the order of a few nanoseconds and hence is much less complex to implement, the analog CNF filter is still responsible for the fine-grained rotation necessary for constructive filtering.

However the pre-filter is limited in the number of taps it can use because each tap adds delay (e.g. for a 80Msps sampling rate, each extra tap adds 12.5ns of delay). To build a reasonable low-delay spread filter, we therefore allow only a delay budget of 50ns which would imply a 4-tap filter at 80Msps. To compute the optimal values of the coefficients for this limited pre-filter, we solve the following optimization problem:

$$\min_{h_D(n), H_A(f_i)} \left| H_A(f_i).\left\{ \sum_{n=0}^{4} h_D(n) e^{j2\pi f_i n} \right\} - H_c(f_i) \right|^2$$

where, $H_A(f_i)$ is the response of the analog constructive filter, $h_D(n)$ represents the pre-filter as described above and $H_c(f_i)$ is the desired overall constructive filter response as computed in Sec. 3.2. The above problem is essentially trying to divide up the work of rotation for constructive filtering between the digital and analog CNF filter stages in an optimal manner so as to best approximate the desired constructive filtering response. We omit the details of how to solve this optimization problem for brevity, we use a standard convex optimization technique called sequential convex programming (SCP) to solve it [7].

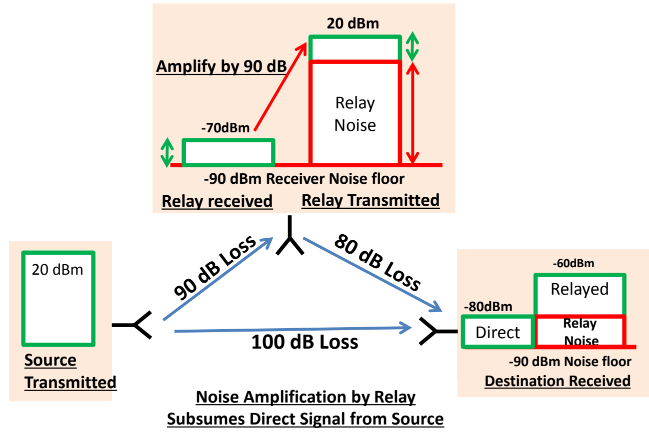

Figure 11: Naive amplification at relay can amplify and relay noise to the destination, which can subsume the direct signal from the source to the destination and negate the benefits of construct-and-forward relaying.

At this point the constructive filtering is complete. We incur a 50ns delay in the digital domain, and a negligible delay (3ns) in the analog constructive filter. Fig. 3 shows the overall block diagram.

3.5 Does the relay amplify noise?

A natural concern is whether the relay amplifies noise. For example, let's say the relay is receiving a signal at 20dB SNR. If the actual noise floor is -90dBm, the signal received is at -70dBm. Lets say it applies the 90dB amplification and transmits a 20dBm signal, in that 20dBm signal, noise is at 0dBm. If the path from the relay to the destination attenuates the signal by 80dB, then even at the receiver the noise from the relay is at -80dBm. This can overwhelm any signal directly received by the destination from the source if the SNR on that direct link is less than 10dB. So in effect the direct signal from the source is drowned out by the noise that is amplified by the relay. Fig. 11 shows how this visually.

Our key insight is that this can be prevented by smartly leveraging the relay's knowledge of the channels. The idea is to compute the amplification factor that ensures that the noise from the relay, by the time it is attenuated by the relay-destination channel, is well below the destination's noise floor. To accomplish this, let's say the attenuation applied by the channel from the relay to the destination is a dB, the maximum amplification factor is given by $(a - 3)$ dB (the 3dB is extra margin for safety). In other words amplification is dictated by how much the signal is attenuated from the relay to the destination, the higher the attenuation, the higher the amplification that can be applied. Remember however that amplification is limited at the top by the amount of cancellation achievable.

In the above example where the relay-destination channel attenuation is 80dB, if we use a maximum amplification of 77dB, the relay would transmit a 7dBm signal, and noise would be at -13dBm. This signal after being attenuated by the channel would be received at the destination at -73dBm and noise would be -93dBm. Since the destination's own noise floor is at -90dBm, higher than the noise received in the relayed signal, it doesn't hurt performance. Now the direct signal from the source is not washed out, and assuming constructive and forward filtering has been applied, it should add up to provide a SNR gain.

4. IMPLEMENTATION

A full design of FF has to grapple with several engineering challenges, we describe a few prominent ones below. Note that we defer the discussion of how the relay knows the identity of the source and destination of the packet it is relaying to Sec. 6, it needs this information to use the right CNF filter. For now, we assume that the FF relay knows the identities of the source and destination to simplify description.

4.1 Carrier Frequency Offset and other issues

As with any radio, inevitably there is a carrier frequency offset between the radios at the source and the destination. Relaying should not introduce another carrier offset into the relayed signal, this would break the assumption of it being another multi-path from the source and can confuse the receiver's CFO correction algorithms. Ideally to avoid confusion, the receiver should get the relayed signal also with the same CFO as the signal it is receiving from the source. So the relay should in fact try to relay the signal such that the original CFO offset from the source is preserved.

This would be easy to achieve if the relay itself did not need to process the signal. However for the relay's own processing and constructive filtering, the CFO w.r.t the source has to be removed. Hence the relay applies the following trick: It computes its CFO wrt to the source. When it receives a signal from the source, it first corrects for that CFO [23]. After that it performs its processing, including digital cancellation and constructive filtering. Before transmission however, it applies the reverse of the CFO correction it applied earlier. In effect it restores the CFO that existed in the signal from the source.

4.2 How does the relay know the channels for construct and forward relaying?

For construct and forward relaying, the relay needs to know the channel from the source to the destination which it cannot directly measure, as well as channels from the source to itself and from itself to the destination. The channel between itself and the source can be easily measured using received signals, and the channel from the destination to the relay can be measured by snooping on ACK packets and estimating the channel. However the direct channel between the source and destination cannot be measured by the relay, it needs to be explicitly informed of it.

Direct Channel: In cellular systems such as LTE, clients measure the channel from the basestation to themselves and feed it back explicitly to help with scheduling [6], our relay can snoop on this feedback and learn the channel. However WiFi has historically been passive, there is no explicit channel feedback from the receiver to the source. To obtain this information for WiFi at the relay, we use recent enhancements in the WiFi standards. Specifically 802.11n/ac implements an explicit channel sounding phase [25, 2, 3] where the AP sends a pre-defined packet which each client uses to measure their channels from the AP. The clients respond with the compressed channel state measurement later when polled by the AP. This is known as the Very High Throughput (VHT) beacon packet [2] in the 802.11ac standard. When FF relays are deployed, we make the corresponding AP send out the HT sounding packet every 50ms.

FF relays then take advantage of this mechanism to obtain the channel estimates from the source to each destination in the network. We make the FF relay spoof the AP and send a polling packet to all clients in the network periodically (every 50ms). The relay then listens to the replies from the clients which contain the channel estimates of the channel from the AP to themselves. Further the relay uses these packets to also measure the channel between the relay and each client in the network. The relay also keeps track of the channel between itself and the AP whenever it receives a packet from the AP.

Note that once the relay computes the constructive filter to use in the downlink direction for a particular AP-client pair, it can use the same filter in the uplink direction for the same client-AP pair. The reason is that by reciprocity the environment between the AP and the client is the same in the reverse direction. Further, the cumulative effect of the channel from the AP to the relay, the constructive filter and from the relay to the client is the same even if the order of channels and filter is permuted and multiplied in a different order by commutativity. Hence the same constructive filter can be used in both directions[1].

4.3 Hardware Prototype

We have built a prototype of the FF relay using the WARP software radio boards [8]. We build on prior full duplex radio implementations [10], but modify them appropriately to implement the relaying functionality. For all our experiments, we have built a MIMO full duplex 2×2 FF relay building on the self-interference cancellation design from [11, 10]. The prototype has 2 antennas and uses the MIMO analog cancellation design described in recent work [10]. The analog cancellation circuit has 8 taps that are spaced around 100-200ps apart as well as taps for canceling the cross-talk between MIMO antennas. Each tap has tunable digital step attenuators [11] which can be adjusted in increments of 0.25dB from 0 to 31.75dB. The couplers get a copy of the signal from the transmit side, and couple it back in to cancel it on the receive side as seen in Fig.8. The cancellation circuit is tuned from baseband after observing the residual using the algorithm described in Sec. 3.3.

The baseband implementation is relatively simple. We implement a 4 tap digital construct and forward filter, as well as implement digital cancellation using a 120 tap causal filter. Further CFO correction and re-distortion blocks are also located in baseband. The overall extra delay introduced by baseband process is less than 100ns in our prototype, of which nearly 50ns is from the digital CNF pre-filter and the rest are ADC and DAC delays, the digital cancellation itself doesn't introduce any delay because its causal. Finally, the signature technique identifying the source destination discussed in Sec. 6 is implemented which lets the relay know the source and the destination of the packet and allows the right constructive filter to use for relaying.

All of our experiments are run with a standard 20MHz OFDM PHY that is based on the WiFi PHY. The PHY uses 56 subcarriers and a 400ns cyclic prefix interval (this is the faster version of WiFi which uses a smaller CP). The numbers we report in our evaluation are all PHY layer throughputs and do not include MAC layer effects. Since the relay is operating at Layer 1, we are orthogonal to any MAC layer effects, so we expect the relative gains should carry over.

5. EVALUATION

We evaluate the performance of FF using experiments in an indoor setting. We place the AP and a FF relay in *various different indoor* settings, those are, open wide office space, L-shaped corridor and a wide room, two large wide room and including the one shown in the Fig. 1. The AP is a 2×2 MIMO AP, and the relay and the client are also equipped with two antennas can 2×2 MIMO. We were limited to 2-antenna devices primarily because of the availability of analog cancellation boards at the relay, we require four of them for implementing MIMO full duplex. We also require four RF analog construct-and-forward boards. However the qualitative conclusions from the experiments apply to any MIMO setup since the constructive filtering technique for improving SNR and MIMO

[1]Note that, the amplification applied is different in both direction, as the noise introduced at the relay receiver, is asymmetric in uplink and downlink directions.

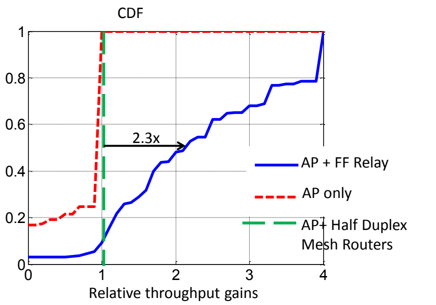

Figure 12: FF's overall throughput gains. FF provides a 3× increase in median throughput, and nearly a 4× gain in dead spot scenarios. Further, it significantly outperforms half duplex mesh routers, almost by a factor of 2.3×.

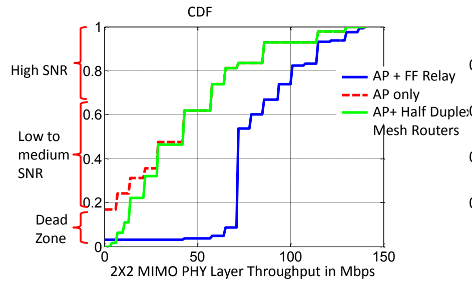

Figure 13: PHY Layer absolute throughputs achieved by different schemes. FF provides a significant throughput for nodes that were previously almost getting no connectivity or very low throughput.

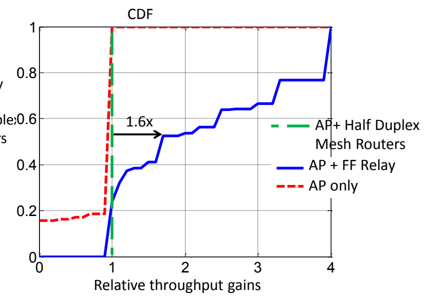

Figure 14: FF's throughput gains due to SNR amplification from construct-and-forward relaying for a SISO AP, FF relay and client. FF provides a median gain if 1.6× even without the benefit of MIMO rank expansion.

rank is independent of the number of antennas. We assume relay knows the source and destination for every transmission.

We compare the following three approaches:

- **AP only**: In this approach we only assume that an AP is deployed without any relays.
- **AP + Half-Duplex Mesh Routers**: This is akin to the approach where we have an AP and a half duplex router such as the Apple Airport Express. To make sure the gains are reported correctly, we assume that the AP and the mesh router are perfectly synchronized and transmit in alternative time slots to eliminate any MAC layer contention effects. Hence the numbers reported are PHY layer throughputs assuming perfect MAC coordination. The half-duplex mesh router also has two antennas. Also, AP is smart enough to figure out when it should use the half-duplex router and when not to use it.
- **AP + FF Relay**: This is the design proposed in this paper. We place it at the same location as the half duplex mesh node. Here too we pick the optimal bitrate to use at the AP assuming the construct-and-forward relaying is in place.

The metric we use is PHY layer throughput which is defined as the optimal bitrate that can be used at any location given the SNR and the MIMO rank. Hence we eliminate any impact of bitrate adaptation algorithms, MAC layer artifacts etc and the experiments purely quantify the impact of relaying. Further to make relative comparisons across the compared approaches, we use a relative throughput gain metric where the baseline scenario is the AP and the half duplex mesh router case. We do not use the AP only scenario because we have dead spots in this scenario where the throughput is zero and we cannot compute relative gain. So all relative gain numbers are wrt to the throughput achieved by using the AP and half duplex mesh router.

Our experiments show that:

- FF provides a median throughput gain of 3× in our experiments wrt to the AP only case. For the bottom 20th percentile of the locations, the throughput gain is as high as 4×.
- FF's gains from MIMO rank increase and SNR gains affect different nodes. For clients that had a decent SNR but low MIMO rank, the majority of the gains are from the addition of a separate independent MIMO path. For clients that are located in dead spots or with very low SNRs, the big gains are from the SNR gain.
- Construct-and-forward relaying has significant benefits, especially for clients with low SNRs. An amplify-and-forward relay without FF's constructive filtering capability performs worse, the median gain wrt to the AP only scenario drops to 1.5×.
- Low latency cancellation and constructive filtering are critical, without them relaying can actually hurt overall performance due to inter-symbol interference, in some cases the performance is worse than no relaying.

5.1 Overall Performance Gains

We begin with the basic question: how much does the FF relay help in improving throughput and coverage. We conduct the experiment as follows. We place clients at different locations in the testbed relative to the AP and relay placement as shown before. We measure the channels from the AP to the client and feed it to the relay. We also measure the channel between the AP to the relay and from the relay to the destination. These measurements are all made available to the relay and the measurements are repeated every 50ms. The relay uses these measurements to compute the construct-and-forward filter. We then conduct an experiment where the AP transmits directly to the client without any assistance from the relay. We then repeat this experiment assuming the relay is a half duplex mesh router, and then with the FF relay. We compute the relative throughput gain and plot the two CDFs in Fig. 12.

The FF relay provides a 3× increase in median throughput over the AP alone, and a 2.3× increase over half duplex relays. The reasons are as expected, the SNR gain we get from construct-and-forward relaying, as well as the increase in MIMO rank due to the additional independent path from the relay. Consequently the AP is able to use very high bitrates. Further at the edge of the coverage area where performance is typically poor, a FF relay improves performance by a factor of 4×. Compared to the half duplex router, the gains are primarily because a full duplex relay does not need an additional time slot to relay. The half duplex relay definitely helps in the edge of the coverage area, where the direct channel from the AP to the client is so poor, that it is better to take the extra hop with the half duplex mesh node.

A natural question is how much of the gains are coming from the SNR gain due to construct and forward, and how much are due to MIMO rank enhancement. We evaluate this question next.

5.2 Performance gain with SISO

We conduct this experiment by using a SISO WiFi AP, a SISO client and a SISO relay (both for the HD and the FF cases). The rest of the experiment is conducted the same way as above. We plot the CDFs of the relative throughput gains in Fig. 14. The gains in this experiment should be purely from the SNR gain from construct-

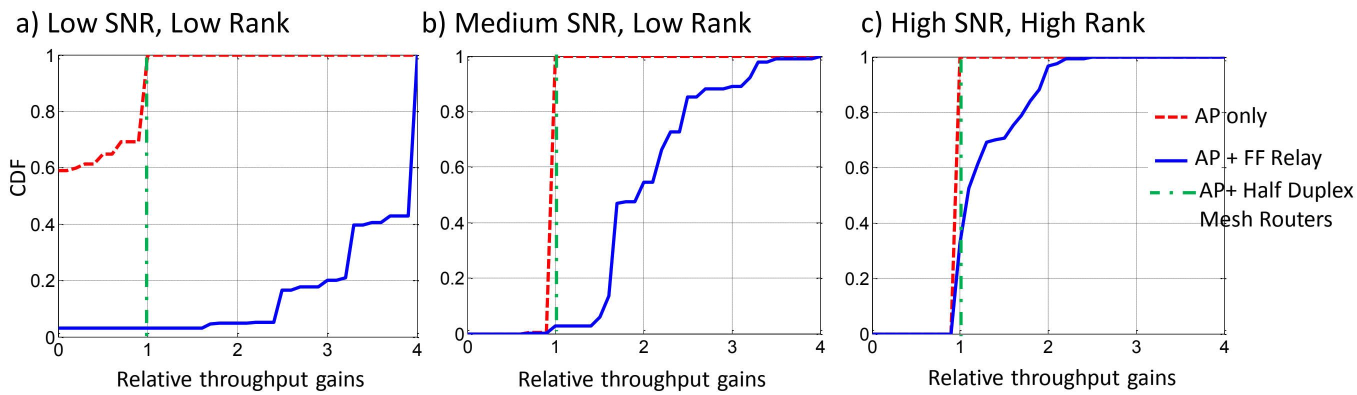

Figure 15: FF's performance gains in different scenarios. In low SNR and low MIMO rank scenario (figure a) the gains are significant because FF provides both a SNR gain as well as MIMO rank expansion, leading to a $4\times$ increase in throughput. FF's performance gains in medium SNR and low MIMO rank scenarios (figure b) leads to a $1.7\times$ increase in throughput. FF's gains in the scenarios where the clients already had high SNR and good MIMO rank (figure c) are minor as expected.

and-forward relaying since there is no MIMO. As we can see, the median gain is $1.6\times$ and the gain at the tail is $4\times$. The experiment demonstrates the fact that in this case the clients at the edge of the coverage area benefit the most. This is expected, since without the AP these clients probably have an SNR in the range of 2-8dB. The relay significantly improves the SNR to about 15-20dB. This translates to allowing the AP to use a 64-256QAM modulation compared to BPSK or QAM before, leading to a $3-4\times$ increase in throughput. On the other hand clients that had medium to high SNR with the AP already don't benefit as much, the gains for them are marginal. The reason is that going from 64QAM to 256QAM doesn't help much, it only increases the bitrate by 33%. The intuitive reason is that the Shannon capacity curve is concave with SNR, there are diminishing returns in terms of capacity as SNR increases. For example, going from 64QAM to 256QAM requires a 6dB increase in SNR, but it only increases the bitrate by 33%.

5.3 Performance gains due to MIMO rank expansion

Next we turn to evaluating the impact of FF's ability to expand MIMO rank. We conduct the same experiment as in Sec. 5.1. However we divide the results into three classes according to how the MIMO channel matrix looked between the AP and the client without any relaying. The first category is when the SNR and the MIMO channel rank are both low, this corresponds to clients at the edge of the coverage area where both propagation losses and MIMO rank degradation are severe. The second category is when the SNR is medium to good, but the MIMO channel rank is low. This corresponds to clients which are suffering from the pinhole effect, they only have one strong path to the AP which reduces MIMO rank but the SNR is still decent. Finally, the last category is high SNR and full MIMO rank, this of course corresponds to clients which are close to the AP and enjoy strong, multiple independent links to the AP. Fig. 15 plots the CDFs of throughput gains in those categories.

Fig. 15.c shows that the benefits from FF for the last scenario (high rank, strong SNR) are small, only around 15%. This is as expected, since FF cant increase rank any more and benefits from SNR gains are small. For the second category Fig. 15.b, where there is good SNR but low rank due to pinholes, the benefits are substantial from using the FF relay. In effect these relays end up providing an additional strong MIMO path and increase the rank to full rank for MIMO, thus providing close to a $1.7\times$ increase in throughput.

The last category shows the (Fig. 15.a) maximum gains, because the relays end up providing a rank of at least two between the AP and the client, as well as enhancing SNR. Given the low baseline these clients are starting from, the gains are therefore significant, showing a $4\times$ increase in throughput.

5.4 Impact of Processing Latency

As we discussed earlier, processing latency at the FF relay has a significant impact. In this experiment we quantify the impact. We artificially introduce some buffering to vary the processing delay at the FF relay from 100ns to 400ns. We then repeat the same throughput experiments as before and plot the median throughput gain as a function of processing latency at the relay in Fig. 16. As we can see, the median throughput gains drop significantly and is in fact worse than having no relay when the processing latency exceeds 300ns. The reason is as expected, above a certain latency we hinder OFDM's ability to absorb highly delayed multipath reflections into the current symbol and avoid inter-symbol interference.

5.5 Impact of No Construct-and-Forward Relaying

In this experiment, we turn off construct-and-forward filtering at the relay and let it simply amplify the received signal to the maximum extent, i.e. as much as the amount of cancellation we obtain. The rest of the throughput experiment is the same as before. We plot the CDF in Fig. 17. As we can see, there are still significant gains at the tail. These correspond to client which were at the edge of the coverage area of the AP, and benefit significantly from the amplified relaying. However the median gain is small to non-existent. This is because for the clients that have medium to good SNRs, blind amplification ends up amplifying noise and washing out the direct signal from the AP to the client. Hence the gains are limited and in some cases are worse than before because of the enhanced noise.

5.6 Impact of Reduced Cancellation

We conduct an experiment where we vary the amount of cancellation at the relay. Remember that cancellation sets an upper limit on the amount of amplification that the relay could use. We plot the median throughput gain as a function of the amount of cancellation obtained in Fig. 18. As expected, with reduced cancellation, overall median throughput gains drop significantly. The reason is that at the edge of the coverage area, being able to use high amplification factors is crucial. A reduced amount of cancellation means the

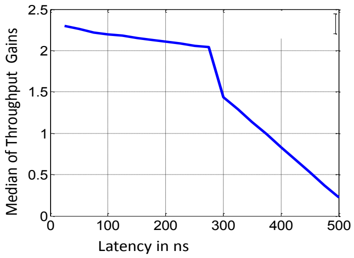

Figure 16: Relaying performance suffers as processing latency increases at the relay. Higher latency means that the relayed OFDM symbol falls outside the CP of the quickest arriving OFDM symbol at the destination, leading to inter-symbol interference and poor performance.

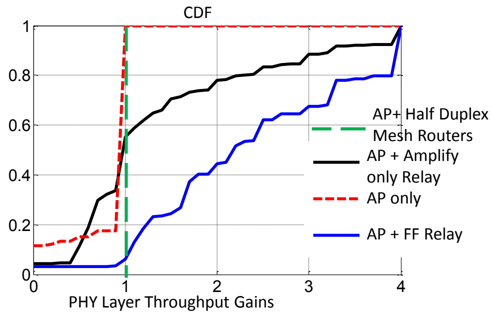

Figure 17: FF's construct-and-forward relaying is crucial for obtaining good performance. If we disable it and implement simple amplify-and-forward relaying, sometimes the performance is worse than no relaying because noise gets amplified.

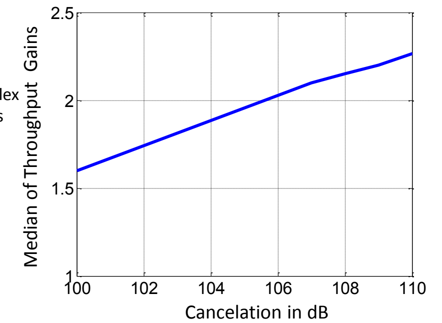

Figure 18: Reduced cancellation means reduced amplification, which leads to significantly reduced throughput gains for FF relays.

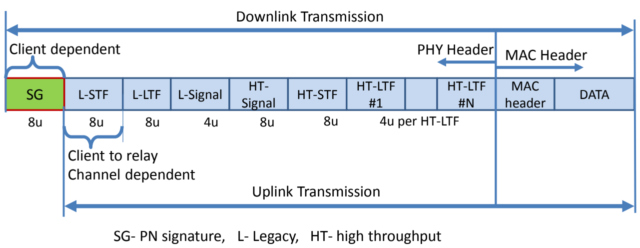

Figure 19: WiFi Header with the amendment on the downlink and for uplink we use standard WiFi header, but we use the STF to find the source of the uplink.

relay's amplification factor is reduced and consequently clients in dead spots see reduced throughput.

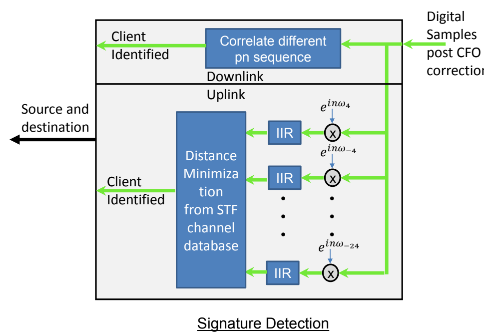

Figure 20: Signature Detection technique showing both uplink and Downlink. For Downlink we use correlation based client identifier. For uplink we extract the 10 subcarriers of STF (using the complex exponent and low latency IIR filters) to run distance minimizing on the database of client estimation, which is simply finding minimum distance vector with a phase compensation.

6. HOW CAN WE DEPLOY FF?

A final implementation question is how selective is the FF relay. Should it relay any packet it detects? Further, which construct and forward filter should it apply? If it did just an amplify and forward, the FF device might relay packets from a different network and AP (neighbor's WiFi for example) and cause destructive multi-path. Even within the network, if the channel between source and destination is strong, relaying may hurt performance by adding noise. Hence we make a conscious design decision that FF should only constructively relay the packets from its own network.

Further to achieve construct and forward filtering, the relay needs to learn the identities of the source and destination pair to apply the correct constructive filter. In the last section, relay knew the identity of source and destination to know whether to relay or not, and which filter to apply. In scheduled systems such as LTE, this information is known in advance to the AP and can be communicated to the relay explicitly, hence this isn't an issue in LTE. However systems such as WiFi are random access and at any point of time any of the clients or the AP could be transmitting.

One approach could be for the relay to just decode the MAC header (as seen in Fig. 19)it is receiving, identify the source and destination and use that to then apply the right constructive filter. However this won't work in practice, because the MAC header is after the PHY header and channel estimation at the destination is performed using the PHY header. Hence the destination would use an incorrect channel estimate in decoding, if relay waits for MAC header to start construct and forward relaying. Therefore in WiFi, we need to find a mechanism for the relay to start applying the right constructive filter before the PHY header itself.

To do so, we make each AP explicitly prepend a pseudo-random sequence of length $4\mu s$, repeated twice, to each packet they transmit. A separate pseudo-random sequence is used for each associated client, and these sequences are learned by the relay on the fly, as AP transmits packets to these clients. The relay continuously looks for these sequences via simple correlation as seen in Fig. 20, and whenever it finds a match, picks up the right constructive filter and applies it to the rest of the packet. The pseudo-random sequence at the start of the packet does not affect the client since its decoding kicks in only after it recognizes the standard WiFi preamble. Fig. 19 shows the structure of the downlink packet.

The above technique clearly requires the APs to change, and we believe thats reasonable to expect since its relatively easier to upgrade them. However, we cannot use this technique at the clients, since it will be far harder to expect them all to be upgraded with this new feature. Therefore the above technique only works in the downlink. So, what could we do about the uplink?

We make a key observation here, unlike the downlink, on the uplink the identity of the destination is fixed, its the AP. All we need to do is identify the source, i.e. the transmitting client. To do so, we design a fingerprinting technique as seen in Fig. 20. The idea is that every WiFi packet has a short preamble at the start of the packet that is known in advance and when it is transmitted it undergoes a transformation governed by the channel between the client and the relay. Remember that the relay already knows the channel between every client in the network to itself, so it can try and match the received preamble to a set of pre-transformed preambles

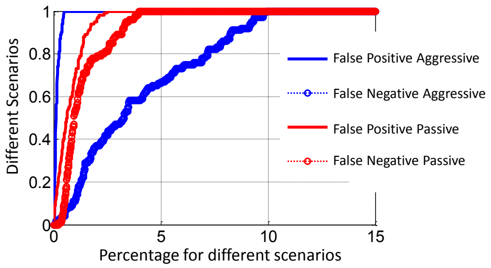

Figure 21: Performance of two channel fingerprinting technique, the aggressive one is more suitable.

corresponding to all the clients. This is once again similar to the pseudo-random sequence correlation idea used in the downlink, but in this case we are simply using the transformed standard preambles itself as the sequences to correlate against. Note that this technique does not require any changes to the clients.

This technique will have false positives , since the WiFi preamble even after transformations corresponding to different channels does not have the same differentiating properties as a set of carefully designed pseudo-random sequences. In Sec. 6.1 we evaluate the false negative and positive rates of this technique. A false negative is relatively harmless, since it just means that no constructive filtering will be applied and the situation will be no worse than a standard WiFi network. A false positive (defined as mistaking one client for another) could in some cases worsen the SNR by applying the wrong filter. Hence we tune our identification thresholds to have nearly a zero false positive rate at the expense of a higher false negative rate.

6.1 Sender Identity from Channel Fingerprints

We evaluate how well our correlation based technique to identify the identity of the sender in the uplink direction as described in Sec. 6 works in practice. We place 4 different client in 100 different locations in our testbed, and for each location calculate the accuracy of sender identification over a time period of five minutes and atleast 1000 packets per client. The extended time period allows us to also account for any channel fluctuations over time. Fig. 21 plots the CDF of false positive and negative rates. A false negative means that no sender is identified, whereas a false positive means that some other sender from the actual sender is identified. We see that the technique does have a 5% false negative rate, but essentially a zero false positive rate. The reason for the false negatives is the aggressive threshold applied for identification, sometimes legitimate senders are missed because of these stringent requirement. The conservative trade-off does ensure a zero false positive rate however and prevents the relay from doing any harm.

7. CONCLUSION

This paper demonstrated how we can design powerful yet simple relaying techniques that can greatly improve throughput and coverage, yet are minimally invasive and do not require sophisticated changes to clients. FF operates within the framework of the current network architecture and design, and we believe can be easily deployed.

Acknowledgments: We would like to thank Brad Karp, Manu Bansal, Rakesh Misra, Aaron Schulman and the Stanford Networked Systems Group members and anonymous reviewers for their insightful comments. This work was supported by a Thomas and Sarah Kailath Stanford Graduate Fellowship.

8. REFERENCES

[1] *LTE Advanced Speeds*. http://en.wikipedia.org/wiki/4G#LTE_Advanced.

[2] *802.11ac: The Fifth Generation of Wi-Fi*. http://www.cisco.com/en/US/prod/collateral/wireless/ps5678/ps11983/white_paper_c11-713103.pdf.

[3] *Meraki White Paper: 802.11n Technology*. https://meraki.cisco.com/lib/pdf/meraki_whitepaper_802_11n.pdf.

[4] *Modeling Indoor Propagation*. http://www.remcom.com/examples/modeling-indoor-propagation.html.

[5] *Next Generation Gigabit WiFi - 802.11ac*. http://www.netgear.com/landing/80211ac/images/wp_netgear_802_11ac_wifi.pdf.

[6] *Physical layer procedures(FDD)*. http://www.qtc.jp/3GPP/Specs/25214-890.pdf.

[7] *Sequential Convex Programming*. http://www.stanford.edu/class/ee364b/lectures/seq_slides.pdf.

[8] *WARP Project*. http://warpproject.org.

[9] P. Almers, F. Tufvesson, and A. Molisch. Keyhole effect in mimo wireless channels: Measurements and theory. *Wireless Communications, IEEE Transactions on*, 5(12):3596–3604, December 2006.

[10] D. Bharadia and S. Katti. Full duplex mimo radios. In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*, pages 359–372, Seattle, WA, Apr. 2014. USENIX Association.

[11] D. Bharadia, E. McMilin, and S. Katti. Full duplex radios. *SIGCOMM '13: To appear in the Proceedings of the ACM SIGCOMM 2013 conference*, 2013.

[12] A. Carleial. Multiple-access channels with different generalized feedback signals. *Information Theory, IEEE Transactions on*, 28(6):841–850, Nov 1982.

[13] T. Cover and A. Gamal. Capacity theorems for the relay channel. *Information Theory, IEEE Transactions on*, 25(5):572–584, Sep 1979.

[14] M. Duarte, A. Sengupta, S. Brahma, C. Fragouli, and S. Diggavi. Quantize-map-forward (qmf) relaying: An experimental study. In *Proceedings of the Fourteenth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, MobiHoc '13, pages 227–236, New York, NY, USA, 2013. ACM.

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

[16] D. Gunduz, A. Goldsmith, and H. Poor. Mimo two-way relay channel: Diversity-multiplexing tradeoff analysis. In *Signals, Systems and Computers, 2008 42nd Asilomar Conference on*, pages 1474–1478, Oct 2008.

[17] V. V. Khang and N. D. Thong. Rank-deficiency in indoor mimo. In *TENCON 2007 - 2007 IEEE Region 10 Conference*, pages 1–4, Oct 2007.

[18] T. Laakso, V. Valimaki, M. Karjalainen, and U. Laine. Splitting the unit delay [fir/all pass filters design]. *Signal Processing Magazine, IEEE*, 13(1):30–60, Jan 1996.

[19] S.-H. Lee and S.-Y. Chung. When is compress-and-forward optimal? In *Information Theory and Applications Workshop (ITA), 2010*, pages 1–3, Jan 2010.

[20] A. Lo and P. Guan. Performance of in-band full-duplex amplify-and-forward and decode-and-forward relays with spatial diversity for next-generation wireless broadband. In *Information Networking (ICOIN), 2011 International Conference on*, pages 290–294, Jan 2011.

[21] P. Mededovic, M. Veletic, and Z. Blagojevic. Wireless insite software verification via analysis and comparison of simulation and measurement results. In *MIPRO, 2012 Proceedings of the 35th International Convention*, pages 776–781, May 2012.

[22] E. C. V. D. Meulen. Three-terminal communication channels. *Advances in Applied Probability*, 3(1):pp. 120–154, 1971.

[23] P. Murphy and A. Sabharwal. Design, implementation and characterization of a cooperative communications system. *CoRR*, abs/1102.0485, 2011.

[24] A. Ozgur and S. Diggavi. Approximately achieving gaussian relay network capacity with lattice codes. In *Information Theory Proceedings (ISIT), 2010 IEEE International Symposium on*, pages 669–673, June 2010.

[25] R. Porat, E. Ojard, N. Jindal, M. Fischer, and V. Erceg. Improved mu-mimo performance for future 802.11 systems using differential feedback. In *Information Theory and Applications Workshop (ITA), 2013*, pages 1–5, Feb 2013.

[26] B. Rankov and A. Wittneben. Achievable rate regions for the two-way relay channel. In *Information Theory, 2006 IEEE International Symposium on*, pages 1668–1672, July 2006.

[27] S. Simoens, O. Muñoz Medina, J. Vidal, and A. Del Coso. Compress-and-forward cooperative mimo relaying with full channel state information. *Trans. Sig. Proc.*, 58(2):781–791, Feb. 2010.

[28] V. VÃd'limÃd'ki and T. I. Laakso. Principles of fractional delay filters. In *PROCEEDINGS OF THE IEEE INTERNATIONAL CONFERENCE ON ACOUSTICS, SPEECH AND SIGNAL PROCESSING*, pages 5–9, 2000.

[29] M. Yuksel and E. Erkip. Diversity-multiplexing tradeoff in cooperative wireless systems. In *Information Sciences and Systems, 2006 40th Annual Conference on*, pages 1062–1067, March 2006.

[30] M. Yuksel and E. Erkip. Multiple-antenna cooperative wireless systems: A diversity x2013;multiplexing tradeoff perspective. *Information Theory, IEEE Transactions on*, 53(10):3371–3393, Oct 2007.