DuoRelay: Parallel Interference Nulling using Full-duplex Relaying

Ananya Mahanti*, Wei-Han Chen† and Prasun Sinha‡

*Department of Electrical and Computer Engineering, The Ohio State University, USA

†‡Department of Computer Science and Engineering, The Ohio State University, USA

Email: *mahanti.1@osu.edu, †chen.7323@osu.du, ‡sinha.43@osu.edu,

Abstract—**In recent years, the volume of uplink traffic in our networks has grown and become increasingly diverse. Typically, concurrent uplink transmissions can be made possible by allowing the APs to collaborate over the wired backbone for control signalling as well as data decoding. In our paper, we remove the dependency over the wired backbone for data decoding without requiring any pre-coding constraints at the clients' end. Our paper contrasts with other existing works which can decode concurrent uplink data-streams without the backbone as they either need multiple antennas at the transmitters/receivers, pre-coding at the transmitters or antennas to be physically moved at the receivers. In our paper, we address the problem space of single antenna transmitters and receivers and enable concurrent uplink transmissions by leveraging the recent developments in full-duplex wireless communication. We make each of our APs act as full-duplex relay nodes. Each AP cancels its own client's signal at the other AP, thus enabling themselves to clearly hear their own clients without requiring any pre-coding at the clients' end. We implemented and evaluated our scheme using USRP N210 software defined radios and achieved upto 1.48x throughput gains over TDMA. The throughput improvement over TDMA is upto 1.7x in our system level simulations with our greedy node selection algorithm.**

I. Introduction

Present times have witnessed an exponential increase in uplink traffic in our networks. This is due to the growing popularity of a broad variety of new applications [1] such as video conferencing, cloud computing, VoIP (Voice Over Internet Protocol), online gaming and the traffic produced from different IoT devices (e.g., sensor readings from home appliances). The nature of these applications is getting increasingly diverse and a large section of them cannot benefit from the traditional strategies of MIMO due to them being space constrained at the transmitters and/or receivers. The use of mmWave technology could enable the use of multiple smaller antennas for space constrained wireless nodes. However, present day mmWave technology needs to place additional constraints on the clients/transmitters due to the highly directional nature of the mmWave links. We need to thus devise proper solutions to address the problem space of uplink traffic for single antenna transmitters and receivers.

We consider a system model of 2 APs and 2 clients, all in a single collision domain. Client1 wants to transmit to AP1 and Client2 wants to transmit to AP2. To enable concurrent transmissions, we need to counter the interference from Client1 at AP2 and the interference from Client2 at AP1. Specifically, AP1 receives $y_1 = x_1 h_{11} + x_2 h_{21}$ while AP2 receives $y_2 = x_2 h_{22} + x_1 h_{12}$ where x_1 and x_2 are transmissions from Client1 and Client2, respectively. Previous works have addressed this problem of concurrent transmissions among single antenna nodes [1, 2, 3]. The underlying concept is that the APs collaborate over the wired backbone to solve the two equations to find the two unknowns x_1 and x_2. A problem with decoding data over the wired backbone is the high cost involved in exchanging raw samples. While [1, 2, 3] alleviate the above problem by exchanging decoded packets, all of them depend on cascaded decoding of data packets over the wired backbone. This results in jitter or variable delays in decoding and the decoding of a client's packet becomes majorly dependent on the decoding accuracy of the clients' packets prior to it. Further, for large topologies, such schemes could end up consuming significant resources of the wired backbone. Another set of works that has enabled concurrent transmissions require pre-coding at the transmitters [4] and/or multiple antennas at the transmitters or receivers [5]. In [6], interference alignment is achieved with single antenna transmitters without any cooperation from senders but it requires antennas at the receiver to be physically moved while [7] requires the channels to change every slot. These assumptions are not practical as for stationary environments, the channels may not vary from one packet to another.

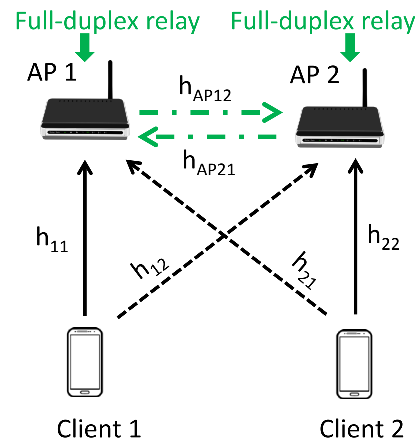

Figure 1: System model with 2 clients transmitting to 2 APs

In our paper we do away with the wired backbone for the purpose of data decoding and cancel the interference in air. We achieve concurrent uplink transmissions with single antenna transmitters without requiring any pre-coding at the transmitters' end. In our design, we leverage recent advances in full-duplex wireless communication to enable each AP to serve as a full-duplex relay node. AP1 forwards the signal it receives in such a way that Client1's transmission gets canceled at AP2 while AP2 forwards the signal it receives such that Client2's transmission gets canceled at AP1. The relayed path via one AP at the other AP essentially appears as another multipath which aligns destructively with the signal from the direct path (interfering client-to-AP). Specifically, the component of Client1's transmission received at AP1 is $h_{11}x_1$; AP1 scales this by a factor of $-h_{12}/h_{11}h_{AP12}$ and retransmits it. The signal when received at AP2 is now -$h_{12}x_1$ which thus cancels the component of Client1's transmission at AP 2 (interference) which is $h_{12}x_1$. Similar process is followed at AP2. However, we do not eliminate the wired backbone completely and use it for our control signalling. The overhead from control signalling is typically much smaller than that from data packets and the control signals also do not have a cascading feature which justifies it's delegation to the wired backbone.

The design, although simple poses non-trivial challenges to solve. First, the relayed path needs to be within the cyclic prefix in order to prevent inter-symbol interference (ISI) at the destination AP. Any processing of the signal received at the relay AP thus needs to be within a certain delay budget. Second, we are really dealing with multi-carrier transmissions and thus the scaling factor for each constituent sub-carrier is different. The processing at the relay AP thus needs to account for this fact. [8] explores the possibility of constructive relaying so that the signal from a full-duplex relay does not increase noise and/or add up destructively at the destination with the direct path. In [9], the authors propose destructive full-duplex relaying so as to cancel an uplink client's interference at the downlink client. Our paper contrasts with [8] and [9] as we propose the parallel forwarding of two canceling signals from two APs in order to simultaneously null the interference from the clients at the two APs. When two full-duplex relays function simultaneously, there is a possibility of signal looping between them in an infinite loop [10], [11] and is thus also a concern that we address. Specifically our contributions are:

- DuoRelay achieves parallel interference nulling using full-duplex signal forwarding. Two single antenna APs forward the signal they receive to cancel each others interference.
- We model the relayed path as a composite filter, thereby integrating the delay spread of the client-AP channels into the available delay budget. This modeling brings greater clarity to the feasibility of full-duplex signal forwarding.
- We analyze the inter-looping of signal between the relays in our context of parallel nulling and show the feasibility DuoRelay.
- We evaluate our solution using USRP software radios and get throughput gains of upto 1.48x over TDMA. The throughput improvement is upto 1.7x using our node selection algorithm in our system level simulations.

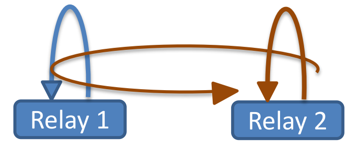

Figure 2: Training of Relay1 followed by that of Relay2

II. Background

A. Cyclic Prefix

Due to multipath propagation in a wireless medium, a receiver receives multiple copies of the transmitted signal. This causes consecutively transmitted symbols to interfere with each other at the destination resulting in inter-symbol interference (ISI). In an OFDM based system, a guard interval called cyclic prefix is placed between two symbols to combat ISI. The length of this guard interval thus needs to be at least as long as the channel delay spread.

In order for the IFFT/FFT algorithm (the key to making OFDM realizable) to work without inter-carrier interference (ICI), cyclic prefix is made to be a copy of the last portion the OFDM symbol inserted at the beginning of the same symbol. Mathematically, this converts the linear convolution with the channel response into a circular convolution to enable simple frequency domain processing. As the cyclic prefix preserves the orthogonality of the sub-carriers, the spectral components of all the multipaths arriving within CP add up linearly. In our paper we measure the total frequency domain spectral component of a multipath channel between a source and destination for different sub-carriers and artificially add another multipath such that we achieve destructive addition amounting to nulling at the destination.

B. Loop-back Interference

When two full-duplex relays function in the same collision domain, the transmitted signal from one relay is received at the other and gets relayed further to the original relay. This signal then keeps looping from one relay to the other forming an infinite loop. It is important to cut short this looping as it can potentially saturate the radios. [10] solves this problem by having the relays to train with the cross-interference channel. The training, as shown in Fig. 2 is done in a sequential manner where Relay1 trains first with respect to its self-interference channel and Relay2 trains next with respect to the composite channel comprising its own self-interference channel and the cross-interference loop through the other relay. The relays design their cancellation filters with respect to the channels that they train on. This implies that the signal from Relay2 can go to Relay1 but is canceled the next time it is received at Relay2 while any signal from Relay1 goes to Relay2 first, comes back to Relay1 and is finally canceled when it reaches Relay2 again.

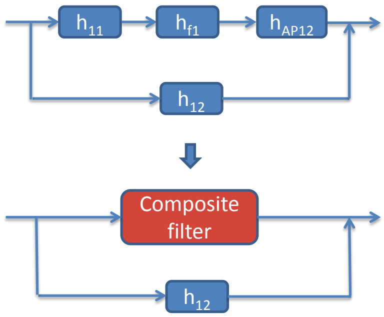

Figure 3: Equivalent filter-modeling of the relayed path

III. Relay Design at the Access Points

A. Blind Interference nulling over single slot

In this paper, we propose parallel interference nulling at two APs using full-duplex relaying. In addition to their regular job of decoding the signal they receive, each AP acts as a full-duplex relay node (which implies that they can forward the signal that they receive in the same frequency band by being able to cancel the self-interference that results from it). In our design, each AP swiftly passes the signal that it receives through an alignment filter before transmitting it again. For instance, between Client1 and AP2, h_{12} is the path for the interfering signal which we refer to as the direct path and the path from Client1 to AP2 via AP1 is the path for the canceling signal which we refer to as the relayed path. For the relayed path to cancel the direct path without causing ISI at the destination AP, we need the relayed signal to be within the cyclic prefix. The delay spread of the effective channel at a destination AP for a relayed signal is longer than the delay spread of a regular single hop channel between a transmitter-receiver pair. As shown in Fig. 3, we model the relayed path as a **composite filter**. This composite filter comprises of a cascade of three FIR filters: (1) channel from client to AP; (2) alignment filter at the AP and (3) channel from one AP to the other AP. The number of taps of this composite filter times the sampling interval gives the delay spread of the relayed path. This delay spread should be within CP. Considering the relayed path: Client1-AP1-AP2, the number of taps of the composite filter is then $T(h_{11}) + T(h_{f1}) + T(h_{AP12}) - 2$ where $T(h_{11}), T(h_{f1}), T(h_{AP12})$ are the number of taps of the channel between Client1 to AP1, the alignment filter at AP1 and the channel between AP1 to AP2, respectively. The relayed path of Client2-AP2-AP1 can be modeled in a similar way. Channels in typical indoor environments have relatively smaller delay spreads as compared to macro-cellular channels, with the measured rms delay spread ranging between 16-52 ns [12]. With a length of 400 ns, the cyclic prefix of WiFi is thus sufficiently over-provisioned (a fact also recently leveraged in [13]) to accommodate a two-hop cancellation architecture as proposed in DuoRelay. Our concern then is to keep the number of taps of the alignment filter small enough such that the total delay spread of the relayed path stays within CP.

DuoRelay keeps the magnitude of the relayed signal same as that of the signal via the direct path while aligning the relayed signal 180 degree out of phase with the direct path. This amounts to canceling the direct path or the interfering signal. It is to be noted that the decoding at the AP can function at the same time as the full duplex signal forwarding. This is so because the relaying technique adopted is a variant of amplify and forward [8] and does not need the signal to be decoded before forwarding. The design of the alignment filter at the relay is explained in the following section.

We term this interference nulling as "blind" because the cancellation technique at the APs does not require any knowledge of the transmitted symbols at the clients' end. The alignment filter operates on the raw samples from the ADC and feeds its output directly to the DAC for further transmission. Hence, decoding for the purpose of alignment is not required. We also note here that a decode and forward version for OFDM signals would be infeasible for our problem definition because of the latency involved in the FFT computation for decoding. For decoding, the FFT resolution required is equal to the sub-carrier spacing. Therefore the signal acquisition time needed for the FFT is equal to one OFDM symbol duration which clearly exceeds our delay budget.

B. Destructive alignment of multi-carrier transmissions

As described earlier, the AP applies an alignment filter to make the relayed signal 180 degree out of phase with the interfering signal at the destination. As the sub-carriers in OFDM can be considered to be independent flat fading channels, we describe the filter design first in the context of a single sub-carrier. Let $H_{ij,k}$ be the frequency domain response for sub-carrier k between Clienti and APj, where $\mid H_{ij,k} \mid$ is the attenuation on sub-carrier k and $\angle H_{ij,k}$ is the amount of rotation that the sub-carrier goes through. Similarly, $H_{DFi,k}$ and $H_{APij,k}$ are the frequency domain responses for sub-carrier k for the desired filter at APi and the channel between APi and APj, respectively. To model the filter at AP1, we need:

$$H_{11,k} H_{DF1,k} H_{AP12,k} = -H_{12,k}$$
$$\implies H_{DF1,k} = -\frac{H_{12,k}}{H_{AP12,k} H_{11,k}} \qquad (1)$$

Following similar logic for all sub-carriers, we can construct H_{DF1}.

We need to rotate and scale each sub-carrier such that it aligns destructively at the destination, that is, we need to synthesize the alignment filter at AP1 to have the response H_{DF1}. Similarly, the desired response for the alignment filter at AP2 would be H_{DF2}. At each AP, we adopt a frequency domain filter design for this alignment filter. As explained before, the filter operates in the baseband on the digital samples. To achieve our desired response, we can employ a N-tap digital filter but this would exceed our delay budget as each additional digital filter tap adds a delay equivalent to that of the sampling interval [8],[9]. We hence design an approximate digital filter with taps ($< N$) small enough to satisfy our delay budget. Further, this alignment filter must not amplify the received signal beyond the self-interference cancellation that is possible at the access point.

Let K denote the set of sub-carriers used for transmission with k denoting the index of the sub-carrier used (where the DC sub-carrier is indexed 0). $h_{fi}(n)$ is the alignment filter at APi with N_{ti} number of taps, $P_{\text{Ci},k}$, the transmitting power of the Clienti on the kth sub-carrier and I is the residual interference at the AP due to imperfect self-interference cancellation and loop-back interference from the other AP. N_o is the noise at the relay and P_{max} is the minimum of the maximum transmit power allowed at the AP and the self-interference cancellation possible at the AP. The power of the received signal at AP1 on sub-carrier k can be given by:

$$Pr_{AP1,k} = P_{\text{C1},k}|H_{11,k}|^2 + I + N_o \tag{2}$$

The signal, after it is received at AP1 goes through the alignment filter before being transmitted again by AP1. The power of the transmitted signal needs to be less than P_{max}. The filter at AP1 is thus determined by the following constrained minimization problem:

$$\min_{h_{f1}(n)} \sum_{k\in K} | \sum_{n=0}^{N_{t1}-1} h_{f1}(n)e^{\frac{-j2\pi nk}{N}} - H_{DF1}(k)|^2 \tag{3}$$

$$\text{s.t.} \sum_{k\in K} | \sum_{n=0}^{N_{t1}-1} h_{f1}(n)e^{\frac{-j2\pi nk}{N}} |^2 Pr_{AP1,k} \leq P_{max}$$

The filter at AP2 can be derived in a similar way. The intuition behind the minimization is that we really have $| K |$ equations (one equation for each sub-carrier) while we have fewer variables; the number of variables being equal to N_{ti}. This gives us an over-constrained system and we adopt a linear regression to minimize the squared error. The functioning of the APs has been described independently in our analysis so far. We explain the parallel functioning of the APs in the following section.

C. Simultaneous working of the two APs

In our scheme, the two relays work in tandem to null each other's interference. In this section, we analyze the scheme from the perspective of the APs' simultaneous functioning. Considering Client1's information on sub-carrier k to be $x_{1,k}$ and Client2's information on sub-carrier k to be $x_{2,k}$; AP1 should receive $m_1x_{1,k}$ and AP2 should receive $m_2x_{2,k}$ where m_1 and m_2 are some known scaling factors. Let k_1 and k_2 be the modification factors applied to the sub-carrier at AP1 and AP2 respectively. In steady state, at AP2:

$$H_{22,k}x_{2,k} + H_{12,k}x_{1,k} + k_1(m_1x_{1,k})H_{AP12,k} = m_2x_{2,k}$$

As at AP2, x_1 terms should be zero:

$$H_{12,k} + k_1m_1H_{AP12,k} = 0$$

$$k_1 = -\frac{H_{12,k}}{m_1H_{AP12,k}} \tag{4}$$

Equating the x_2 terms:

$$m_2 = H_{22,k} \tag{5}$$

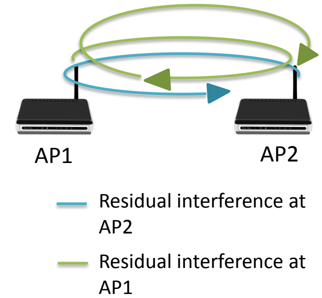

Figure 4: Looping of residual interference between the APs

Similarly at AP1:

$$H_{11,k}x_{1,k} + H_{21,k}x_{2,k} + k_2(m_2x_{2,k})H_{AP21,k} = m_1x_{1,k}$$

As at AP1, x_2 terms should be zero:

$$H_{21,k} + k_2m_2H_{AP21,k} = 0$$

$$k_2 = -\frac{H_{21,k}}{m_2H_{AP21,k}} \tag{6}$$

Equating the x_1 terms:

$$m_1 = H_{11,k} \tag{7}$$

We see that the modification factors at AP1 and AP2 come out to be the same as the destructive filter response we calculated before, which implies the two APs can indeed work in tandem to decode their respective client's transmissions. However, the cancellation of the interfering signal by the relayed path is not perfect and there is some residual interference at each AP. So each AP doesn't just forward the scaled and rotated version of its own client's signal but also the residual interference from the other client. Specifically, the residual interference from Client1 at AP2 gets relayed to AP1 and gets further relayed back to AP2. This residual interference hence keeps looping between the APs. Similar looping would happen for the residual interference at AP1 from Client2. Fortunately, [10] addresses this issue of inter-looping between the relays for a pair of source and destination for multi-hop transmission. Similar training can be employed in our problem setup to cut off the looping of the uncancelled interference between the APs.

Consider, AP1 trains first followed by AP2 using the sequential training described in section II-B. Then as illustrated in Fig. 4, residual interference from Client2 at AP1 goes to AP2, is relayed back to AP1, and is finally canceled when it reaches again at AP2 while the residual interference from Client1 at AP2 goes to AP1 and is canceled when it reaches again at AP2. This training for the APs needs to be done only once before the beginning of transmission and needs to be repeated only if the self-interference channel or the cross-interference channel between the APs changes. Instead of an infinite looping of the residual interference, AP1 thus suffers from this loop-back interference twice while AP2 suffers from it only once.

IV. DuoRelay in Networks

In our paper so far, we have discussed the working of DuoRelay with 2 clients and 2 APs. In this section we analyze how DuoRelay might work in a network of multiple clients and multiple APs. Our objective is to obtain high throughput from the network while not starving any client. We will let 2 pairs of clients and APs transmit together at a time using the technique detailed out in previous sections. Naively, we can exhaustively cycle through a collection comprising all such pairs of clients and APs but this would require $O(N^2)$ computations. To efficiently select two client-AP pairs, DuoRelay uses a centralized server to compute scheduling decisions based on a greedy algorithm, as shown in Algorithm 1. We describe the algorithm in section IV-A and how it is used in section IV-B. In our analysis, we consider our nodes to be in a single collision domain but our algorithm can be made to operate in a multi-collision domain using the scheduling algorithm stated in [1], where nodes which can hear each other can be efficiently grouped together and the groups are scheduled to operate such that neighboring groups don't transmit simultaneously.

A. Greedy selection algorithm

The goal of this selection algorithm is to achieve throughput gain and fairness. Let each client be associated to the AP to which it has the best average channel strength. A centralized server keeps track of the waiting time of the clients in the client waiting pool; it selects upto M (say M=10) clients each time out of the pool as candidates for uplink transmission based on the clients' waiting time (line 5). The longer they wait, the higher their priority. If two or more of the short listed clients are associated with the same AP, then ties are broken randomly, i.e., only unique client-AP short listed pairs are retained (line 6). Finally, two clients (C_i and C_j) out of the unique short listed pairs are selected for transmission based on their cross link interference. Specifically, as an individual client is already paired with the AP to which it has the best average channel strength, we choose the weakest diagonal channel between the shortlisted client-AP pairs to form our 2x2 network (line 10) and schedule the corresponding clients and APs for transmission. After the scheduling decision is made, the waiting times for all the other clients except for (C_i and C_j) are increased (line 11-12). In the case where all the short listed clients may be associated with a single AP, a single client may be selected randomly out of the short listed clients to transmit to the designated AP for that time slot (line 6-8).

B. Execution of the selection algorithm

To make scheduling decisions, a lead AP is first elected to act as a centralized scheduler for the other APs in the same network. In densely deployed Enterprise Wireless LAN (EWLAN), the control signaling between the APs can be done over the wired backbone in the infrastructure.

Before the clients transmit data, the lead AP has to first acquire the information of which set of clients have uplink traffic and make scheduling decisions accordingly. For this purpose, the APs may jointly broadcast a polling message to all clients and the clients can respond with a small packet to indicate their request for transmitting uplink traffic. Then, this information can be sent to the lead AP so that it can schedule client-AP pairs based on the selection algorithm. The polling process might introduce significant overhead for large number of clients and decrease the throughput gain of DuoRelay. Practically, fast polling algorithms [14] can be used to reduce this overhead. Further, the overhead from channel measurements is discussed in section VI.

Algorithm 1: Greedy Client-AP pairs selection scheme.

Input: Clients (C_i, where i - client index) waiting for uplink transmission, their associated access points (A_i) and their current waiting times (W_i).

Output: A set of Client-AP pairs that will transmit in the current time slot.

// The lead AP keeps track of the waiting time W_i for each client C_i

1 $ClientPool \leftarrow \{C_i : \forall i : 1 \leq i \leq N\}$

2 $S \leftarrow \{\}$

3 **if** $ClientPool$ == $\emptyset$ **then**

4   **return** S

// Choose upto M clients with the longest waiting times

5 $Candidate \leftarrow \{C_i$: W_i is among top M maximum in the $ClientPool$ $\}$

6 $Candidate \leftarrow \{C_i : A_i \neq A_j, \forall j \neq i, C_j \in Candidate\}$

7 **if** $sizeof(Candidate)$ == 1 **then**

8   S $\leftarrow (C_i, A_i)$

9 **else**

10   S $\leftarrow$ S $\cup \{(C_i, A_i), (C_j, A_j)$: (C_i, A_j) or (C_j, A_i) is the smallest cross link interference, and $C_i, C_j \in Candidate$ $\}$

11 **foreach** $C_i \in ClientPool$ *and* $C_i \notin S$ **do**

12     $W_i \leftarrow W_i + 1$

13 **return** S

V. Performance Evaluation

A. Testbed Setup

The experiments are performed on USRP N210 software radios. Each USRP is mounted with a SBX daughterboard and operates in the 1200 MHz band. We use a 64 point FFT OFDM system and 52 subcarriers are used for data transmission. The cyclic prefix length is set to 16 and the sampling rate taken is 1MHz. All USRP nodes use a common external clock source generated by OctoClock-G for clock and time synchronization. Practically, existing algorithms, such as SourceSync [15], can be leveraged to synchronize APs in a network with upto nanosecond level of accuracy.

Baseband processing is done offline. We implemented our signal processing blocks in MATLAB, including our self-interference cancellation filter and destructive alignment filer. For self-interference cancellation, our digital filter provides up to 25dB cancellation.

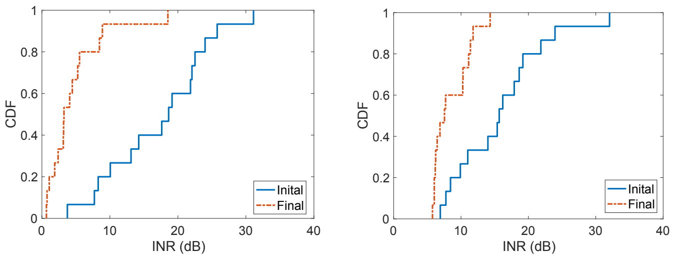

(a) Nulling in LOS (b) Nulling in NLOS

Figure 5: Interference nulling for the 3 node setup.

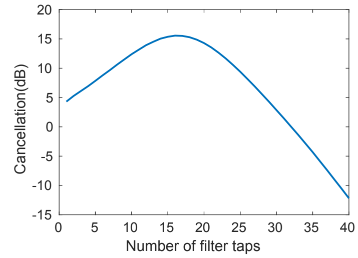

Figure 6: Cancellation versus number of filter taps.

B. Interference Nulling

We evaluate interference nulling with a three node setup, shown as Client1, AP1, and AP2 in Fig. 1. Specifically, we let AP2 be the receiver and Client1 be the interferer. AP1 acts as the relay, which receives the signal from Client1 and forwards the nulling signal to AP2.

As software processing is latency intensive, we emulate the functioning of real time full duplex signal forwarding in our experiments. To capture the effect of self-interference at AP1 while it forwards the nulling signal, we conduct our experiment in 3 steps. First, Client1 transmits data and AP1 receives. Second, Client1 and AP1 transmit simultaneously. Client1 transmits the same symbols (as transmitted in the first slot) while AP1 transmits the symbols it had received in the first slot. AP1 thus receives the samples from Client1 with self-interference captured. The samples are then processed for self-interference cancellation and destructive filtering. Third, Client1 transmits the same symbols (as transmitted in the first slot) while AP1 transmits the nulling signal, and AP2 receives. We process the samples received at AP2 at the end of third slot. Practically, implementations as detailed out in [8] can be used to implement a real-time full-duplex relay which can perform signal alignment.

The experiments are done in an indoor office environment, which cover both line-of-sight (LOS) and non-line-of-sight (NLOS) scenarios. We repeat our experiments at 15 randomly selected locations in our testbed. To evaluate DuoRelay's interference nulling performance, we measure the metric *interference to noise ratio (INR)*. We plot the INR before with the INR achieved after implementing our scheme. The distributions of INR are as shown in Fig. 5. On average, we achieve a cancellation of 12.55dB for LOS and 7.39dB for NLOS. Note that we are emulating a real-time hardware implementation by software, which inherently introduces some level of inaccuracies for the channel state information (CSI). For LOS scenarios, wireless channels tend to be more stable, which is the reason why we achieve better cancellation. While the median final INR for LOS is slightly above noise floor, the same for NLOS is higher at 7.54dB due to CSI variations.

C. Performance of the Alignment Filter

We evaluate the interference cancellation that can be achieved by our alignment filter by varying its number of taps. As we would like to investigate the effect of the number of taps on the interference cancellation, it would be unfair to evaluate it against varying channel conditions. To keep the channel parameter constant, we collect 30 channel traces from the testbed and evaluate the cancellation achievable versus the number of taps for each distinct channel trace. The result is as shown in Fig. 6. We observe that the cancellation increases initially as we increase the number of taps. This is expected as we have more dimensions for our approximation. The cancellation later starts decreasing and becomes negative as inter-symbol interference (ISI) starts dominating our cancellation performance. This is also expected, as when the length of the filter exceeds our delay budget, it ends up contributing to interference at the destination instead of canceling it.

D. Decoding Performance

To evaluate the decoding performance of DuoRelay at the two APs, we measure *signal to interference plus noise ratio (SINR)* at each AP and the total achievable throughput. We use the same 30 channel traces collected from our testbed as before, and feed them into our simulator. In the simulation, two clients transmit simultaneously to two APs, as shown in Fig. 1. Each AP acts as a full duplex relay and is trained hierarchically to cancel residual interference looping. As a result, AP1 receives two looping streams while AP2 receives only one. Each AP forwards nulling signals while decoding out message from its intended client.

The distribution of SINR at each AP is as shown in Fig. 7a and Fig. 7b. After interference nulling, AP1 and AP2 have an average SINR of 9.51dB and 12.36dB, respectively. As expected, AP1 has more residual interference looping streams and thus it has lower final SINR. Through our scheme of interference nulling, AP1's SINR increases by 5.2dB and AP2's SINR increases by 9.46dB.

To evaluate the throughput improvement, we compare the total achievable throughput of DuoRelay against TDMA. As we would like to investigate the throughput gains that we have from our physical layer design, we eliminate the effect of bitrate adaption algorithms and plot the absolute physical layer throughput from the Shannon capacity equation. The throughput plots are as shown in Fig. 7c. With our design, we obtain an average throughput improvement of 1.22x and a maximum throughput improvement of upto 1.48x over TDMA. Intuitively, the throughput performance degrades in scenarios where the strength of the cross-link interference is greater than that of the direct links. We leverage this observation for designing our greedy node selection algorithm which we evaluate in the next section.

E. DuoRelay in Networks

We evaluate the performance of DuoRelay in a topology of multiple clients and multiple APs. We simulate the topology by deploying varying number of clients and APs with randomly generated channel variations. We evaluate our client selection scheme by first fixing the number of clients at 60 and varying the number of APs from 10 to 80 and next by fixing the number of APs at 60 while varying the number of clients from 20 to 80. We implement Omni-TDMA in our simulation by scheduling a single pair of client-AP at a time for transmission in a round-robin fashion where each client transmits to the AP to which it has the maximum average channel strength. The plots of throughput and fairness are as shown in Fig. 8. Specifically, we create topologies of three different average SNR levels- low (˜10dB), medium(˜20dB) and high(˜30dB) and plot the throughput of DuoRelay along with that of Omni-TDMA in Fig. 8a and Fig. 8b. Our throughput improvement is upto 1.7x in low SNR and 1.4x in medium SNR topologies. We find the throughput gains to decrease in a topology of high SNR values. This is expected as all the cross-links in the topology are much stronger than the amount of cancellation achievable and hence the decrease in SINR offsets the amount of gains achieved by enabling two clients to transmit together. We plot the fairness of our client selection scheme against that of Omni-TDMA using Jain's fairness index[16]. Fig. 8c and Fig. 8d are the plots for fairness measure in a medium SNR topology. The fairness of the scheme stays consistently above 0.98 when the number of clients is fixed at 60 and the number of APs is varied. The fairness is lower when the number of clients is at 20 as the algorithm implementation considered shortlists as many as 10 clients based on their waiting times, after which it greedily chooses client-AP pairs for transmission from among them. The fairness steadily improves to well over 0.98 when the number of clients increases. The fairness measure shows similar trends in low and high snr environments and the corresponding graphs are omitted for brevity. The overall throughput performance is better than the results obtained from our testbed channel traces as the design benefits from the node selection algorithm.

VI. Discussion

Processing latency at the relay: In our filter design, we have assumed no buffering or additional processing latency at the relay. If there is additional latency at the relay in the transfer of samples from the RX RF-Frontend to the baseband processing unit (FPGA or the host) and from the baseband processing unit to the TX RF-Frontend, then this disturbs the alignment even if the latency is within the cyclic prefix [17]. This is so because we now need to account for the additional delay into our channel estimation between the Client to the AP and from the AP to the other AP. This would amount to multiplying the desired response H_D with e^{jwd} where d is the processing latency at the relay. For the simplest case of flat fading channels, consider a H_D whose IFFT would have a time centre towards its beginning. Multiplying with the phase term e^{jwd} in frequency domian amounts to an advancement in time domain and thus shifts this time centre towards the end, making it harder to approximate the resulting response using a filter of few taps. We can however *still undo the effect of this latency* but we would require more taps for our alignment filter. As noted before, its possible to align perfectly with a N tap filter, however this would require the CP to be extended to accommodate the length of the filter.

Scaling to N clients: We have proposed and evaluated our design for enabling concurrent transmissions between 2 clients and 2 APs but we believe it can be extended to serve multiple client-AP topologies. Previous works [5, 18] have exploited MIMO to scale concurrent transmissions. Using multiple antennas increases the number of dimensions we can exploit and hence brings benefits over SISO, and DuoRelay can benefit from such capabilities as well. For instance, if each AP has two antennas, then each AP has one extra dimension that can be used to accommodate a third client [7, 19]. Therefore, by adding one more antenna at each AP, three concurrent uplink transmissions can become possible. We think DuoRelay can be extended to serve more clients in a single collision domain as we put more antennas at each AP.

Channnel measurements: To compute the alignment filter in DuoRelay, the APs need to know the channel between itself and the other AP, the channel between its own client and the other AP and the channel between its own client and itself. The overhead of computing the channels between clients and APs has been addressed by previous works [5, 20]. We can adopt similar version of the method noted in [1] and gather channel estimates from multiple transmitters simultaneously by the use of PN sequences [21].

Dynamic environments: DuoRelay has been evaluated in largely static environments. In dynamic environments, the channels may change rapidly which might require the filter design to adapt to the variations. Techniques from adaptive filtering can be incorporated to better tackle such variations.

Wideband Transmissions: Traditionally wireless full duplex technology was limited to narrowband transmissions. But recent advances [8, 22, 23] have demonstrated the practicality of full duplex for wideband transmissions. We have conducted our experiments over a bandwidth of 1MHz but our scheme can be extended to wide band transmissions with the help of analog components as demonstrated in previous works.

VII. Related Works

The related works with respect to our paper can be categorized mainly into hitherto unrelated domains of full-duplex relays and uplink multi-user MIMO.

Full-duplex relays- In [8], the authors explore the importance of constructive alignment of relayed signals as opposed to blind relaying when the signal from the direct path is present at the destination. The relay engages in filtering and amplification of the signal it receives such that the relayed

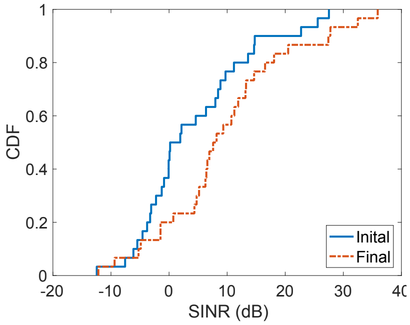

(a) SINR at AP1.

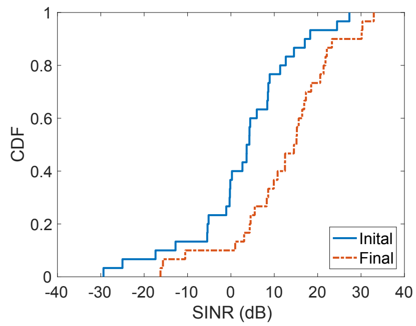

(b) SINR at AP2.

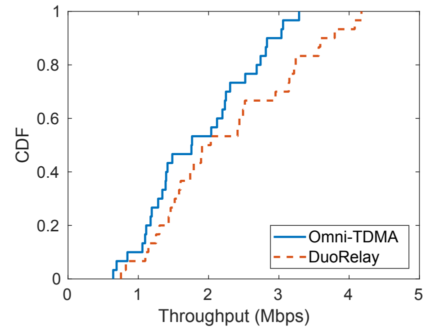

(c) Throughput gain of DuoRelay.

Figure 7: SINR performance and throughput of DuoRelay with interference looping.

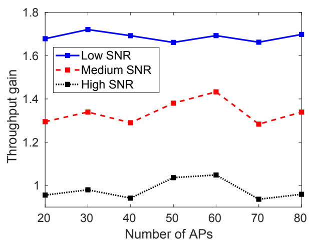

(a) Throughput gain versus number of APs.

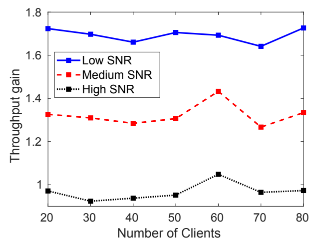

(b) Throughput gain versus number of clients.

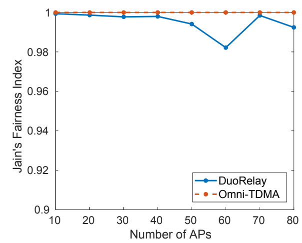

(c) Jain's fairness index versus number of APs.

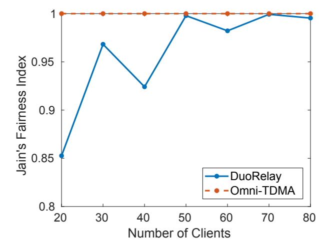

(d) Jain's fairness index versus number of clients.

Figure 8: DuoRelay's performance in a network with varying number of APs and clients.

path does not increase noise and/or add up destructively with the direct path at the destination. The authors report blind relaying to be worse than no relaying. [9] analyses the idea of destructive relaying to cancel the interference from an uplink client at the downlink client. Our work differs from [8] and [9] as we investigate the parallel working two full-duplex APs for achieving parallel interference nulling. [10, 11] deal with designs involving simultaneous operation of multiple full duplex relays. [10] enables in-band wireless cut-through transmissions where the signal from the source goes through a multi-hop path involving a chain of full duplex relays to reach the destination. [11] enables concurrent bi-directional wireless transmissions on the same frequency channel between two nodes where the signal to and from one node to the other traverses a cluster of full-duplex relays. Our work contrasts with [10, 11] as the works do not explore the consequences of blind relaying when the signals from different relays reach a particular destination node.

Uplink Multi-user MIMO- Prior works [1, 2, 3] have dealt with uplink transmissions from clients with single antenna by decoding packets in a cascading fashion over the wired backbone between the APs. In our paper, we design our decoding scheme to be independent of the wired backbone. Further, [2] receives multiple collided packets from transmitters at APs and then stops some of the transmitters from transmitting in the next slot until finally the number of active transmitters are small enough for the APs to collaboratively decode them over the wired backbone. It however assumes a setup where all of the clients and APs are not in a single collision domain. [2] thus does not work for our problem setup where the clients and APs are in a single collision domain. Motivated by the significantly higher bandwidth required to exchange raw samples over the wired backbone, [1] designs a scheme which requires the APs to only exchange decoded packets over the wired backbone to enable multiple clients to transmit to multiple APs, all in a single collision domain. However, [1] does not give any improvement over TDMA for a 2x2 network as both the schemes would require 2 time slots for operation. [3] decodes multiple concurrently transmitted uplink packets using Successive Interference Cancellation, but at different APs which gives it a SNR improvement of 1.5 over TDMA and SIC. Our design sits in clear contrast with all Successive Interference Cancellation (SIC) based schemes in general including [3] as SIC requires difference in power levels received at any AP from the clients present in a single collision domain. Without a power control scheme among the transmitters, this may often not be the case as the channels between client to APs can be quite similar. Our design poses no such constraint on different client's transmit power levels.

Finally, Interference Alignment schemes have been used before to achieve concurrent transmissions. However, unlike our work, [6] needs physical displacement of antennas for achieving alignment, while [5] needs multiple antennas at transmitter or receivers to enable simultaneous transmissions. [7] can enable alignment assuming channels always vary over time. However, this assumption isn't reasonable as when clients are at rest, the channels may be static between two slots. Motivated by this, [24] designs an interference alignment scheme for stationary channels for single antenna transmitters but uses pre-coding at the transmitter, wired backbone for cancellation and multiple time slots for operation. [4] deals with concurrent transmission on the downlink channel but needs the transmitters to jointly precode their transmissions in order to cancel their interference. Our design is for uplink traffic and removes all such precoding constraints on the client/transmitter's end.

VIII. Conclusions and Future Work

In this paper we have proposed and evaluated DuoRelay which integrates the benefits of full-duplex wireless technology into the problem space of MU-MIMO. The scheme enables parallel interference nulling using the full-duplex APs to enable simultaneous decoding of the uplink clients' transmissions. The design removes the dependency over the wired backbone for data decoding and cancels the interferring streams in air without requiring any precoding at the clients' end. We consider this is as an important step in extending the benefits of full-duplex wireless technology towards catering to the needs of an increasingly diversified problem space of uplink traffic. Salient topics for future work include studying the gains from the system in topologies involving mobile clients and analysing the extension of the design for multiple antenna nodes.

IX. Acknowledgement

This project was partially funded by NSF grant CNS-1409336.

References

[1] W. Zhou, T. Bansal, P. Sinha, and K. Srinivasan, "BBN: throughput scaling in dense enterprise WLANs with Bind Beamforming and Nulling," in *Proceedings of the 20th annual international conference on Mobile computing and networking*. ACM, 2014, pp. 165–176.

[2] T. Bansal, B. Chen, P. Sinha, and K. Srinivasan, "Symphony: Cooperative packet recovery over the wired backbone in enterprise WLANs," in *Proceedings of the 19th annual international conference on Mobile computing & networking*. ACM, 2013, pp. 351–362.

[3] W. Zhou, T. Das, L. Chen, K. Srinivasan, and P. Sinha, "BASIC: backbone-assisted successive interference cancellation," in *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*. ACM, 2016, pp. 149–161.

[4] H. S. Rahul, S. Kumar, and D. Katabi, "Jmb: scaling wireless capacity with user demands," in *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*. ACM, 2012, pp. 235–246.

[5] S. Gollakota, S. D. Perli, and D. Katabi, "Interference alignment and cancellation," in *ACM SIGCOMM Computer Communication Review*, vol. 39, no. 4. ACM, 2009, pp. 159–170.

[6] F. Adib, S. Kumar, O. Aryan, S. Gollakota, and D. Katabi, "Interference alignment by motion," in *Proceedings of the 19th annual international conference on Mobile computing & networking*. ACM, 2013, pp. 279–290.

[7] V. R. Cadambe and S. A. Jafar, "Interference alignment and spatial degrees of freedom for the k user interference channel," in *Communications, 2008. ICC'08. IEEE International Conference on*. IEEE, 2008, pp. 971–975.

[8] D. Bharadia and S. Katti, "Fastforward: Fast and constructive full duplex relays," *Proc. ACM SIGCOMM*, vol. 44, no. 4, pp. 199–210, 2015.

[9] K.-C. Hsu, K. C.-J. Lin, and H.-Y. Wei, "Inter-client interference cancellation for full-duplex networks," in *INFOCOM 2017-IEEE Conference on Computer Communications, IEEE*. IEEE, 2017, pp. 1–9.

[10] B. Chen, Y. Qiao, O. Zhang, and K. Srinivasan, "Airexpress: Enabling seamless in-band wireless multi-hop transmission," in *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*. ACM, 2015, pp. 566–577.

[11] L. Chen, F. Wu, J. Xu, K. Srinivasan, and N. Shroff, "Bipass: Enabling end-to-end full duplex," in *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking*. ACM, 2017, pp. 114–126.

[12] I. Propagation, 2010. [Online]. Available: http://www.wirelesscommunication.nl/reference/about.htm

[13] S. Rathinakumar, B. Radunovic, and M. K. Marina, "Cprecycle: Recycling cyclic prefix for versatile interference mitigation in ofdm based wireless systems," in *Proceedings of the 12th International on Conference on emerging Networking EXperiments and Technologies*. ACM, 2016, pp. 67–81.

[14] W. Zhou, D. Li, K. Srinivasan, and P. Sinha, "Domino: relative scheduling in enterprise wireless lans," in *Proceedings of the ninth ACM conference on Emerging networking experiments and technologies*. ACM, 2013, pp. 381–392.

[15] H. Rahul, H. Hassanieh, and D. Katabi, "Sourcesync: A distributed wireless architecture for exploiting sender diversity," 2010.

[16] Wikipedia, 2018. [Online]. Available: https://en.wikipedia.org/wiki/Fairness_measure

[17] V. Yenamandra, "Exploiting Synchronization to Enhance Wireless Networks and Sensing," PhD dissertation, The Ohio State University, 2017.

[18] K. C.-J. Lin, S. Gollakota, and D. Katabi, "Random access heterogeneous mimo networks," in *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 4. ACM, 2011, pp. 146–157.

[19] D. Tse and P. Viswanath, *Fundamentals of wireless communication*. Cambridge university press, 2005.

[20] X. Xie, X. Zhang, and K. Sundaresan, "Adaptive feedback compression for mimo networks," in *Proceedings of the 19th annual international conference on Mobile computing & networking*. ACM, 2013, pp. 477–488.

[21] T. Li, M. K. Han, A. Bhartia, L. Qiu, E. Rozner, Y. Zhang, and B. Zarikoff, "Crma: Collision-resistant multiple access," in *Proceedings of the 17th annual international conference on Mobile computing and networking*. ACM, 2011, pp. 61–72.

[22] M. Jain, J. I. Choi, T. Kim, D. Bharadia, S. Seth, K. Srinivasan, P. Levis, S. Katti, and P. Sinha, "Practical, real-time, full duplex wireless," in *Proceedings of the 17th annual international conference on Mobile computing and networking*. ACM, 2011, pp. 301–312.

[23] D. Bharadia, E. McMilin, and S. Katti, "Full duplex radios," in *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 4. ACM, 2013, pp. 375–386.

[24] B. Chen, V. Yenamandra, and K. Srinivasan, "Interference alignment using shadow channel," in *Computer Communications (INFOCOM), 2015 IEEE Conference on*. IEEE, 2015, pp. 2128–2136.