FD^2: A Directional Full Duplex Communication System for Indoor Wireless Networks

Ehsan Aryafar and Alireza Keshavarz-Haddad

***Abstract*—We present the design and implementation of FD^2, a directional full-duplex (FD) communication system for indoor wireless networks. An FD^2 AP uses directional transmit and receive antennas to reduce self-interference, and to combat AP-AP and client-client interferences that arise due to FD operation in multi-cell networks. FD^2 addresses the joint problem of scheduling and beam selection by proposing efficient practical algorithms. FD^2 is implemented on the WARP platform, and its performance is compared against CSMA/CA and other FD and directional communication systems.**

Our experimental results reveal that: *(i)* Simple application of FD to multi-cell networks can result in significant loss of capacity due to high FD induced interference, while FD^2 can effectively overcome the problem and provide an average gain of nine-fold; *(ii)* FD^2's performance depends on the hardware capture properties and the corresponding rate table, and increases when packets can be captured at lower SINR margins, or when dynamic range of the rate table is high; and *(iii)* FD^2's uplink and downlink performances are susceptible to channel dynamics, and are impacted differently due to mobility. However, we show that training FD^2's rates according to traffic direction, mobility, and feedback rate, increases its robustness to channel dynamics.

I. INTRODUCTION

The increased adoption and use of smart consumer electronic devices (*e.g.*, smartphones, tablets, video surveillance equipment, etc.) is causing an exponential growth in wireless traffic. To meet this demand, various technologies are being investigated to build the next generation broadband networks. Among them, one key technology that has recently received significant attention is full duplex (FD) wireless.

FD is the process of sending and receiving data on a wireless device at the same time and frequency. During FD operation, the transmitting antenna causes a significant amount of self-interference (SI) on the receiving antenna. Given the colocation of the interfering source compared to the farther desired transmission source, this would completely drawn the weaker intended signal in interference, rendering its decoding impossible. Thus, the key challenge in enabling FD so far has been to reduce/eliminate SI. Recent advances in RF and digital cancellation techniques, now allow us to build FD capable radios and double the link capacity [1], [2], [3], [4]. However, when FD is employed in a multi-cell network (*e.g.*, an infrastructure WLAN), additional interference is introduced which limits its potential achievable capacity gain.

In order to better understand this, consider the 2 client, 2 AP network as depicted in Fig. 1, and assume that a Time Division Duplexing (TDD) MAC is employed by the devices. If FD functionality is not employed, during downlink (DL) transmission APs project interference to one another's clients, while during uplink (UL) clients project interference to one another's APs (*i.e.* cross-talk interference).

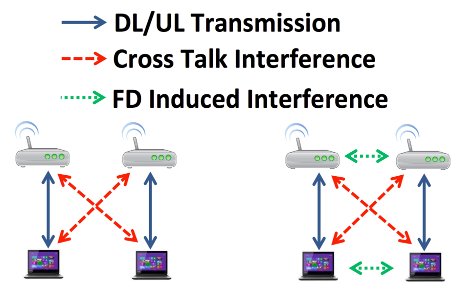

Fig. 1. **Interference patterns in TDD (left) and FD (right) systems.**

With FD, UL and DL happen at the same time. Therefore, in addition to cross-talk interference *strong co-channel AP-AP* and *client-client* interference is present, which can significantly reduce FD capacity gains, specifically in dense deployments or when APs have Line-of-Sight (LoS) links to one another (*e.g.*, when APs are deployed on ceilings).

Directional antennas can alleviate these problems by focusing transmission/reception towards the intended receiver/transmitter, and have been successfully used in a variety of applications: *(i)* to increase outdoor link capacity and stability [5], [6]; *(ii)* to increase spatial reuse in indoor environments [7], [8], [9]; and *(iii)* to build FD radios [10], [11], [12]. The high degree of spatial isolation provided by these antennas and their increased directionality, can be leveraged by full duplex APs to reduce AP-AP interference and increase AP-client signal strength (thereby, combating client-client interference), respectively.

However, translating the potential of directional antennas into practically realizable gains in indoor FD networks is a highly challenging task. In particular, *(i)* given the interference information between APs and clients, an optimal solution needs to find optimal transmit and receive beams for each AP and serve AP-client traffic requests in a small number of timeslots; *(ii)* if such a solution can be realized and implemented in practice to overcome FD multi-cell problems, and *what* are the factors affecting its performance; and *(iii)* in practice obtaining interference information constitutes overhead and it may not be up-to-date to reduce overhead. Further, mobility of UL/DL clients introduces varying interference patterns that can reduce FD gains to worse than half duplex *even with only a single AP*. Thus, the solution needs to incorporate mechanisms to increase its robustness to channel dynamics and client mobility.

978-1-4799-8381-0/15/$31.00 ©2015 IEEE

Towards addressing these challenges, we present FD^2, a wireless network design that increases both UL and DL capacity in indoor environments. FD^2 is designed for an enterprise wireless network, in which FD capable APs have access to directional transmit and receive antennas, and are coordinated through a central controller. FD^2 decouples the joint scheduling and beam selection problem into to two subproblems, and develops practical low complexity algorithms that maintain fairness among traffic flows.

FD^2 is implemented on the WARP platform and its performance is extensively evaluated indoors. Our experimental results reveal that *(i)* FD protocols that use omni antennas can lose significant portion of the capacity, while FD^2 can effectively overcome FD induced interference and provide an average gain of nine-fold; *(ii)* FD^2's gains are more when the underlying hardware can capture packets at lower SINR margins, or when the dynamic range of rate table is high, yielding close to 12 times increase in capacity over omni FD.

Finally, with controlled experiments we show that UL and DL clients are affected differently due to mobility. We also show that a mobile client's performance depends on 4 factors: directionality characteristics of APs' antennas, speed of the mobile client, feedback rate, and UL/DL direction of traffic. Hence, FD^2 categorizes the clients based on these factors, and employs appropriate rate tables when determining the transmission rate, thereby increasing its robustness to both client mobility and channel dynamics.

The rest of this paper is organized as follows. We discuss the related work in Section II. Section III describes the challenges in realizing a practical directional FD communication system. Section IV describes the components of FD^2. Section V describes its implemntation followed by detailed evaluation in Section VI. Finally, we conclude the paper in Section VII.

II. Related Work

There is a large body of work on both directional communication and FD radios. We discuss the most relevant ones to our work.

Directional Communication. Directional antennas have been used in prior works to improve wireless network performance. Directional antennas are traditionally used in outdoor environments to enhance link level throughput [5], [6]. Other works have employed directional antennas on top of vehicles to improve throughput and link reliability between APs and moving vehicles [13].

Recent studies have shown that directional antennas are very useful indoors [7], [14], [15]. These works have provided measurement data on directional antenna channels and link qualities, but do not investigate system design aspects with directional antennas. In [8], [9], authors used directional antennas to increase indoor spatial reuse. In [8], directional antennas are used only by the APs, whereas in [9] directional antennas are used by both APs and client devices. In contrast to these works that are designed for HD APs and address only DL capacity optimization, FD^2 is designed and built for FD APs and addresses both UL and DL capacity optimization.

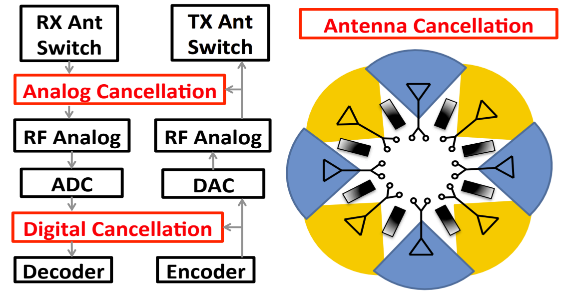

Fig. 2. **Different components of a FD architecture that can contribute to SI suppression: antenna, analog, and digital cancellation. Here antenna cancellation is achieved by using dual polarized directional Tx and Rx antennas with RF absorption employed between different antennas.**

Full Duplex Radios. Recent works have proposed several techniques to enable FD. For a complete discussion on pros and cons of each design, please refer to [16]. Fig. 3(a) shows the various components of a FD system that contribute to self-interference (SI) suppression, namely: antenna cancellation, analog cancellation, and digital cancellation.

Antenna cancellation arranges Tx and Rx antennas in a manner such that SI is reduced at the Rx antennas. In [1], [10], [17] omni antennas are leveraged and placed carefully at specific locations providing up to 45 dB of SI suppression. In [11], [12] directional antennas and other passive SI suppression techniques are leveraged to achieve up to 80 dB of SI suppression. Analog and digital cancellation require knowledge of the transmitted signal and the SI channel to create an inverse copy of the SI signal in the RF and digital domains, respectively. Recent works [2], [3], [4] have proposed several analog and digital cancellation techniques that provide a tradeoff between different objectives such as bandwidth of cancellation, scalability to higher number of antennas, and the level of SI suppression, among others [16].

FD^2 APs leverage directional antennas and other passive SI suppression techniques such as antenna polarization and RF absorption (similar to FD designs in [11], [12]) to significantly reduce SI. However, unlike all these prior work that address radio design, we address the higher level problem of FD communication in a multi-cell network, and more specifically an indoor enterprise wireless network.

III. Design Challenges

In this Section, we describe the system model and challenges in realizing a practical directional full duplex communication system.

A. System Model

We consider an infrastructure wireless network with N FD capable APs and M clients. Each AP has access to K directional Tx antennas (transmit beams) and K directional receive antennas (receive beams). Clients in an FD^2 system could be either half duplex (HD) or FD. For ease of presentation, we assume that each client is HD and has only a single omni antenna. The APs are connected to each other through a

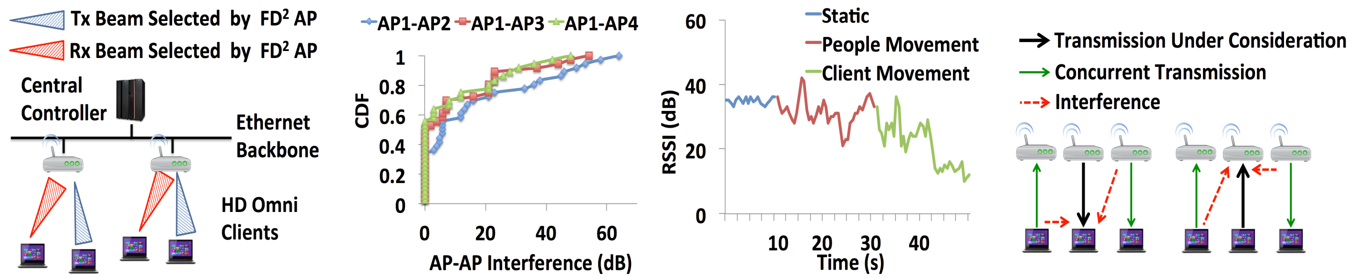

Fig. 3. FD^2 **architecture with FD APs and HD clients (a), CDF of AP-AP interference due to directional transmission and directional reception between AP1 and other APs (b), Impact of channel dynamics on a single link RSSI (c), concurrently scheduled links that cause interference to an on-going downlink (left) or uplink (right) transmission (d).**

separate backhaul network such as the wired Internet, and are coordinated by a central controller (Fig. 3(a)). Such architectures can be commonly found in enterprise wireless networks. The central controller uses interference information between APs and clients in order to schedule concurrent transmissions, and find optimal transmit and receive beams for each AP. Scheduling is done in the context of a TDMA MAC protocol. During each TDMA slot, a set of uplink and downlink links are activated simultaneously.

Our design relies on the *capture effect* property of many wireless radios [18]. Here, a receiver can successfully decode a packet in the presence of an on-going transmission, as long as the desired transmission has a sufficiently higher signal strength (*e.g.*, 10dB for 18Mbps [18]). Hence, with appropriate schedules and beams, intended receivers can capture their packets in the presence of concurrent transmissions.

B. Antenn Configuration

Benefits of Directional APs. Directional antennas can eliminate FD induced interference by focusing transmission/reception towards the intended receiver/transmitter. For example, a $10cm \times 10cm$ patch antenna can provide spatial isolation of more than 15dB within a sector of $90°$ [19], *i.e.*, the signal outside the sector is at least 15 dB weaker than the signal within the sector. This high degree of spatial isolation and the increased directionality, can be leveraged by full duplex APs to reduce AP-AP interference and increase AP-client signal strength, respectively.

We perform experiments in an indoor environment to quantify the gains achieved by employing directional APs. The experiment setup consists of 2 APs each with six Tx and six Rx directional patch antennas. Each AP has also access to a 5 dBi omni directional antenna. We fix AP1 at a location close to a real WiFi AP and vary the place of the second AP at three other locations (denoted as AP2, AP3, and AP4 in Fig. 3(b)), each close to a real WiFi AP.

We first measure the baseline AP-AP interference when omni Tx-Rx antennas are employed. The resulting interference caused by AP1 at AP2, AP3, and AP4 locations are 31, 38, and 52 dB, respectively. Next, we measure the AP-AP interference when directional APs are employed. Fig. 3(b) depicts the CDF of interference due to directional transmission by AP1 and directional reception by other APs across all possible Tx and Rx beam pairings. The results in Fig. 3(b) show that for close to 60% of Tx-Rx beam pairings, the resulting AP-AP interference is less than 3 dB, while for 10% of the pairings interference increases by more than 10 dB.

Results in Fig. 3(b) show that appropriate orientation of Tx-Rx beams can effectively eliminate AP-AP interference, maintaining FD multi-cell UL capacity. Further, directional antennas usually provide stronger signal strength than omni-directional antennas. Thus, employing directional Tx antennas can help a DL client combat the interference caused by UL transmissions. This, combined with appropriate UL-DL client pairing can substantially increase FD multi-cell DL capacity.

FD UL-DL Scheduling and Beam Selection. Solving the joint scheduling and beam selection problem is challenging at multiple levels: *(i)* determining appropriate Tx-Rx beams to cater to a set of UL-DL requests simultaneously, *(ii)* determining *if* and *how* the set of requests must be divided into into sub-groups, where FD UL-DL transmission is executed separately within each group, and *(iii)* achieving a level of fairness among the traffic flows.

In order to better understand the complexity of the problem, consider the scheduling problem to optimize capacity with M clients and N APs, where each client has only a single omni-directional antenna. Finding the optimal schedule is equivalent to comparing the throughput values across all possible partitions of the clients and APs, which is an NP-Hard problem and grows exponentially with the number of APs and clients.

The FD scheduling and beam selection problem takes the complexity to a whole new level, since each AP has access to K Tx and K Rx beams, can operate in FD mode, and overall has $O((K+1)^{2N})$ computational complexity[1]. With the increasing demand for capacity, networks are deployed more and more densely, increasing N. In parallel, advances in directional antenna technology [20] now allow us to pack a high number of directional antennas (in small form factors) with very small beamwidths, increasing K. Designing an efficient algorithm to address these issues, while being amenable

[1]Problem size is $\sum_{M_d+M_u \le M} \binom{N}{M_d} K^{M_d} \times \binom{N}{M_u} K^{M_u}$, in which M_u denotes the number of UL clients and M_d denotes the number of DL clients.

to practical implementation is all the more challenging.

C. Robustness to Channel Dynamics and Feedback Rate

The challenging problem of finding optimal schedules and beams is with respect to determination of a solution under the assumption of accurate RSSI and interference information available at the central controller. In practice, this information constitutes overhead and increases with number of APs, number of Tx/Rx beams on each AP, and number of clients. Hence, accurate interference information may not be available at the central controller. The mobility of clients can also significantly change interference patterns, specifically when FD is employed. The reasons are two fold: *(i)* a mobile client can quickly get out of the Tx/Rx beam selected by the AP, and *(ii)* mobility can rapidly change the interference pattern caused by UL clients on DL clients.

We conduct an experiment in an indoor environment to verify variations in received RSSI of a single client due to changes in the environment or user mobility. We perform an experiment in which an AP transmits back to back packets to a single client on a fixed beam. The variation in received RSSI as a function of time is plotted in Fig. 3(c). Up to time $t = 10s$, user is static and there is no movement in the environment. From $t = 10s$ to $30s$, we ask the people to move around the client and APs. Finally at $t = 30s$, we move the client around the AP. It can be seen that the channel dynamics are almost negligible for a static client, indicating the ability to stand reduced feedback frequencies. However, the situation with mobility in environment or client is quite the contrary, where the average RSSI value drops significantly (specifically with mobility) and has a high variance.

While our experiment captured RSSI variation for a single link due to directional transmission, the problem is further exacerbated in FD multi-cell networks due to UL to DL interference. Hence, it is important to understand the sensitivity of the joint scheduling and beam selection solution to channel dynamics, and devise solution to increase robustness.

IV. Design of FD^2

In this Section, we present the design of FD^2. We first provide an overview of the FD^2 system, and then describe our solution to the joint scheduling and beam selection problem that is amenable to practical implementation. We defer the extension to handle channel dynamics to Section VI.

A. System Overview

FD^2 is a centralized system designed for indoor networks and optimizes both DL and UL traffic. We assume that each client is HD, and that the clients have already associated with their APs (*e.g.*, based on the closeness to APs). The main components in the design of FD^2 are:

- **Information Collection:** During this phase central controller obtains information regarding clients' uplink traffic and interference between directional APs and clients.
- **Building the Conflict Graph:** The central controller uses the interference information to build a *conflict graph*.
- **Scheduling and beam Selection:** The central controller uses the *conflict graph*, and UL-DL queue information of each client to determine the optimal UL and DL beams for each AP, and the links that should be active simultaneously.
- **Data Transmission:** Central controller informs the clients that are selected for UL transmission. Next, selected APs and clients transmit for a pre-determined duration (defined as data transmission *timeslot*).
- **Acknowledging Packets:** At the end of the cycle, APs and clients sequentially transmit block ACKs for correctly received packets.

The main algorithmic components of FD^2 are AP beam selection and scheduling. We next describe this process.

B. Collecting Information at the Central Controller

We obtain the following information during the measurement phase: *(i)* The signal strength observed by clients as a function of APs and their transmit beams; *(ii)* Client UL queue information; and *(iii)* The interference levels among clients. We collect all these information in three rounds:

In the first round, each directional AP sequentially transmits a known preamble on each of its beams. Clients record the RSSI readings for each preamble and calculate $S(i, j, k, b)$ as the mean received RSSI by client j from AP i on beam k and traffic direction b (*i.e.*, DL). Similarly, all other APs record the RSSI readings on all their Rx beams and calculate $S(i_1, i_2, k_1, k_2)$ as the mean interference level observed by AP i_2 from AP i_1, when i_1 uses Tx beam k_1 and i_2 uses Rx beam k_2. APs use backbone wired connection to transfer AP-AP interference and other information to the central controller.

In the second round, clients sequentially send back the measured RSSI information and their UL queue status (in Bytes). While a client transmits this information, each AP measures the mean received RSSI for that client on each of its received beams (*i.e.*, $S(i, j, k, b)$ for each receive beam k, and b denoting the UL direction). At the same time, the rest of the clients record the RSSI readings and calculate $S(j_1, j_2)$ as the mean interference level observed by client j_2 from client j_1.

In the third round, clients sequentially feedback the client-client interference that was measured in the previous round.

C. Building the Conflict Graph

A conflict graph is a concise way to capture interference relations in wireless networks and has been used extensively in wireless networking problems [8], [9], [21], [22]. In a conflict graph, a node represents a transmission and an edge between two vertices shows that the two transmissions interfere with each other. We now describe how to use the conflict graph concept for our FD scheduling problem.

In our conflict graph, each (i, j, k, b) vertex represents a transmission/reception between AP i and client j on beam k and DL/UL traffic direction b. We use the signal measurements described in the previous Section to generate the graph.

Fig. 3(d) (left) shows the links that can cause interference to the receiver of a scheduled DL link. There is a directional edge to (i_1, j_1, k_1, DL) from (i_2, j_2, k_2, DL) if interference level $S(i_2, j_1, k_2, DL) > 0$ (*i.e.*., there is interference from a concurrently scheduled AP with DL transmission). Further, there is a potential directional edge to (i_1, j_1, k_1, DL) from (i_2, j_2, k_2, UL) if interference level $S(j_2, j_1) > 0$ (*i.e.*, there is interference from a concurrently scheduled client with UL transmission). Similar procedure can be used to find interfering links to an UL transmission leveraging Fig. 3(d) (right).

Once the conflict graph is constructed, the expected throughput of a given set of concurrent transmissions can be calculated by measuring the aggregate interference impact on each link, and calculating the expected throughput according to the SINR-Rate mapping [8], [9], [23]. We use this feature in our joint scheduling and beam selection algorithm.

D. Joint Scheduling and beam Selection (JSBS)

FD^2's scheduling and beam selection algorithm reduces the complexity in time, while maintaining fairness among traffic flows. In order to achieve this, we make the following two approximations:

- **Scheduling.** We order the list of transmission requests at the central controller (*i.e.*, $R = \{r\}$, in which $r = (i, j, b)$ is the transmission request between AP i and client j on traffic direction b) based on a priority (*e.g.*, based on the traffic load of each queue, arrival time of the HOL packet, or both), and try to accommodate requests in the given order.
- **beam Pairing.** Unlike brute-force search algorithms that enumerate all possible combinations or a greedy algorithm that keeps track of the best beam at any given time, we maintain a parallel thread of good potential solutions and iteratively refine them to arrive at a desirable final pairing. By maintaining parallel threads of possible solutions, we increase the chance of finding a good solution at the expense of additional computational complexity.

Leveraging the above approximations, the key steps of JSBS algorithm are as follows:

Step 1: We visit each request r sequentially, and verify if we can find appropriate AP Tx/Rx beams that increases the capacity.

Step 2: If such Tx/Rx beams are found, we add r to the set of scheduled requests and maintain a parallel thread of potential beams for r. We use the conflict graph to calculate the throughput of a potential group combination, and to decide whether to maintain a beam in the thread associated with r.

Step 3: Any request that cannot be accommodated in the next timeslot will remain in the queue with increased priority, and will be served in a future timeslot. By increasing the priority of unscheduled transmissions, we maintain fairness among the traffic flows.

Note that since the main focus of this paper is on the implementation and experimental evaluation, we have omitted the details of the the JSBS algorithm and presented them in the technical report [23].

V. Experimental Setup

In this Section, we describe the implementation of FD^2 as well as other solutions studied in this paper.

A. Hardware and Software

Our implementation is based on the WARP platform [24]. We use the OFDM scheme with 10MHz channel bandwidth. All of our experiments are conducted on the 802.11 2.4 GHz channel 11 and during night, in order to avoid uncontrolled interference. Each WARP board sends/receives backlogged UDP traffic with 1000-Byte data packets. Unless otherwise specified, we perform experiments with symmetric UL-DL traffic and set the data transmission timeslot to $25ms$.

In order to implement FD functionality, we use two separate Tx and Rx WARP boards which have access to separate Tx and Rx antennas, respectively. We use a combination of SI suppression techniques (antenna cancellation, polarization, RF shielding, and directionality) in order to reduce interference and realize FD APs.

We use L-Com HG2614 [19] patch antennas with 8 dBi gain and 75° half power beamwidth for directional communication. Each FD^2 AP has access to four Tx and four Rx such directional antennas. The antennas are mounted on a circular array structure with a 20cm distance between neighboring Tx and Rx antennas. We use 5 dBi omni directional antennas when omni antennas are employed.

B. Enterprise WLAN Solutions

We implement the following mechanisms on our testbed.

Omni: This mechanism uses the 802.11 based MAC and PHY with CSMA/CA and with RTS/CTS disabled. This mechanism uses a single omni antenna and operates in half duplex (HD) mode. *Omni* provides a baseline for comparison against directional and FD systems.

DIRC: This is a centralized directional communication system for indoor wireless networks [8]. DIRC divides time into DL and UL periods and uses a centralized scheduling and beam selection algorithm for DL, and omni directional reception with CSMA/CA for UL. We use equal UL-DL timeslots for DIRC with symmetric traffic. We use the same Tx directional antenna setup for DIRC and FD^2. However, unlike FD^2 DIRC uses a single Omni Rx antenna for UL.

Janus: This is a centralized FD MAC protocol for indoor wireless networks with omni antennas [22]. Janus uses interference information between clients to simultaneously schedule UL and DL packets. We realize FD functionality in Janus by employing two omni Tx antennas with antennas cancellation [17] and a single omni Rx antenna. We employ RF absorbers between Tx and Rx antennas, and place Tx and Rx antennas on orthogonal polarization planes to minimize the residual SI.

FD^2: We have implemented the components of FD^2 based on our discussion in the previous Sections. We realize FD functionality by carefully placing Rx directional antennas at the Null of Tx antennas. Further, we employ RF shielding and polarization to minimize the residual SI.

C. System Implementation

We now describe the most important components of our implementation.

FD APs. We realize FD APs by using two separate WARP boards. One board is solely used for DL transmission, while the other board is used for UL reception. Each board is next connected to the corresponding Tx/Rx antennas. We place the antennas in a circular array structure and employ shielding and polarization to minimize SI.

AP Antenna Selection. Each WARP board can accommodate up to 8 antennas using 4 radio cards. Each radio card has two antenna ports and can select from two antennas. In our implementation we use all 4 radio cards and connect each to a single antenna. We select the antenna by selecting the appropriate radio card and turning off the rest of the radio cards. Thus, in our implementation there is only a single radio active at any given time.

Synchronization. FD^2 operates in cycles and contains separate information collection and data transmission phases. In our implementation, a central node (controller) is responsible to synchronize and inform other nodes (WARP boards) to start transmission or reception. The controller has several dedicated output debug header pins which are connected through wire to an input debug header pin of each other node. The controller generates pulse signals on its output debug header pins. The rest of the nodes continuously monitor their input header pins and create a high priority interrupt whenever a rising edge is detected on their input header pins. A rising edge instructs a node to operate in transmit or receive mode.

Information Collection. During this phase, the controller instructs: *(i)* each AP to transmit 4 50-Byte control packets with 6Mbps rate on each of its beams; and *(ii)* the rest of the APs and clients to measure the RSSI. The number and duration of training packets is selected such that APs can correctly measure RSSI across all of their Rx beams.

Modulation and Coding Scheme (MCS) Selection. FD^2, DIRC, and Janus rely on capture property to decode packets in presence of interference caused by concurrently scheduled transmissions. Thus we need to measure WARP's packet delivery ratio when competing against another WARP board at certain relative SINR ratio. Fig. 4(a) depicts the measured capture induced packet delivery ratio (PDR) for different modulation and coding schemes implemented in WARP. We select the rate for a given SINR value, as the highest MCS scheme such that the given SINR achieves 100% PDR.

VI. Performance Evaluation

In this Section, we compare the performance of FD^2 to CSMA/CA (Omni) and other centralized WLAN solutions (Janus and DIRC).

Scenario. Fig. 4(b) depicts our experimental setup in which we deployed nine nodes (3 APs and 6 clients) in an office environment. There are two clients associated with each AP (based on closeness to each AP). One client is configured to transmit UL traffic, whereas the other client is configured to receive DL traffic. Each client has only a single omni-directional antenna. There are 7 sub-topologies in our setup. Topology indices 1-3 denote the scenarios in which only a single AP and its associated clients are active. Topology indices 4-6 denote the scenario in which two APs (and their clients) out of the three are active. Topology index 7, denotes the scenario in which all APs and their clients are active. Each experiment is run for 30 seconds. We repeat each experiment for five times and present the average value.

A. Benchmarking

In this subsection, we first evaluate how effectively we can realize a Full Duplex (FD) AP. Next, we evaluate the effectiveness of our proposed algorithms in Section IV to find close to optimal Tx-Rx beams and user schedules.

SI Suppression Evaluation. Fig. 4(c) depicts the average effect of various SI suppression techniques that we have leveraged to realize a FD AP. The bars on the left show the resulting RSSI at the Rx antenna when we use an omni Tx antenna. The bars on the right, show the corresponding results with a directional Tx antenna. For each scheme, we first measure the baseline SI when an Omni-Rx antenna is used (OR). Fig. 4(c) shows that with an appropriate placement of a dir-Tx antenna, baseline omni SI reduces by 12 dB.

In order to realize FD with omni antennas, we evaluate several passive cancellation methods such as antenna cancellation (AC [1], [10], [17]), AC plus RF shielding (ACS), and ACS while placing antennas on different polarization planes (ACSP). In order to realize FD with directional antennas, we evaluate passive suppression techniques such as leveraging a directional Rx antenna (DR), shielding in addition to DR (DRS), and finally DRS plus polarization (DRSP).

Our results show that ACSP and DRSP provide the highest levels of SI suppression. Thus, we use these schemes to construct FD APs for Janus and FD^2, respectively. Comparing the resulting RSSI after SI suppression with WARP noise floor, shows that there is only 2-3dB of SI remaining above the noise floor. This remaining SI is due to the presence of multipath components in our environment, which can be suppressed by employing digital cancellation techniques [1], [2], [3], [4]. However, given the remaining small margin for SI suppression with WARP, we do not consider it in our implementation. Our results would thus be a lower bound when FD is employed.

Algorithm Sub-Optimality. Fig. 4(d) compares the performance of FD^2's joint scheduling and beam selection algorithm to the optimum found through exhaustive search. We consider both when a single omni Rx antenna is used (OR), and when multiple directional Rx antennas are used (DR). We observe that leveraging directional Rx antennas increases the gains compared to using a single omni Rx antenna. Further, the gains substantially increase with increasing number of active APs.

Careful investigation of SINR traces reveals the following reasons: *(i)* directional Rx antennas help eliminate AP-AP interference which cannot be fully overcome by only employing directional Tx antennas; and *(ii)* directional reception increases uplink SINR due to higher receive directionality.

Comparing FD^2's results in Fig. 4(d) to the optimum shows that FD^2 has a performance that is close to the optimum. On

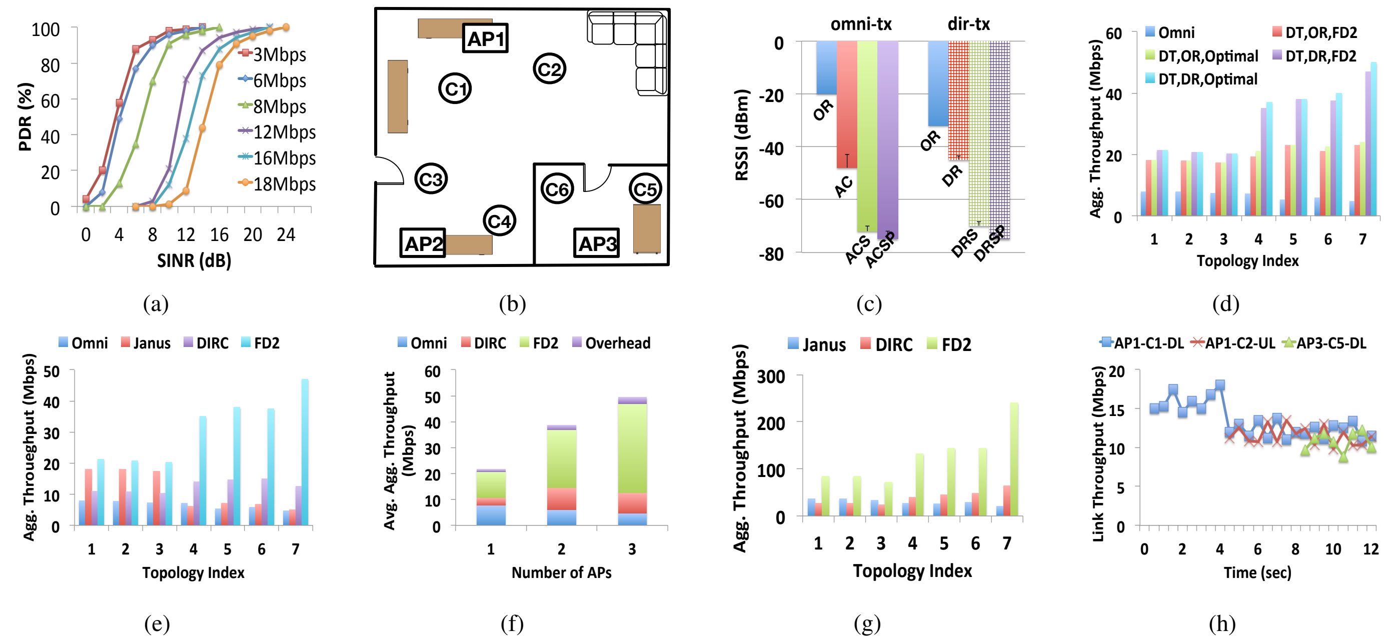

Fig. 4. **Capture probability of WARP (a), Layout (b), SI suppression evaluation: OR (Omni Rx), AC (Tx Antenna Cancellation), ACS (AC + RF Shielding), ACSP (ACS + Polarization), DR (Directional Rx), DRS (DR + RF Shielding), DRSP (DRS + Polarization) (c), FD^2 scheduling and beam selection compared to optimal (d), FD^2 vs. Omni, Janus, and DIRC over WARP (e), Aggregate throughput as a function of number of APs (f), FD^2's performance with 802.11 rates (g), FD^2's performance under traffic dynamics (h).**

average, the optimal solution increases the throughput by less than 6% compared to FD^2.

B. FD^2's Performance Evaluation

We next evaluate FD^2's gains against other solutions.

Performance Gains. Fig. 4(e) and 4(f) show the average aggregate throughput of different schemes as a function of topology index and number of active APs, respectively. Results show that higher number of active APs and clients reduces aggregate Omni throughput due to increased number of collision and presence of hidden terminals. Janus more than doubles the throughput when only a single AP is employed (Fig. 4(e)). Since Janus is a synchronized scheme, it eliminates 802.11's CSMA/CA contention. Combined with FD operation, Janus more than doubles the throughput with single AP even in the presence of UL to DL interference, caused by the UL client's transmission on the DL client. However, increasing the number of APs, reduces Janus's performance to that of Omni. This is because of high level of AP-AP and client-client interference, which in addition to cross talk interference, severely decreases Janus's performance.

DIRC achieves a lower throughput than Janus with only a single active AP, but more than double's the throughput with increasing number of APs. DIRC uses directional transmit antennas to increase spatial reuse and number of concurrent DL transmissions. However, DIRC uses 802.11's omni transmission for UL, limiting its gains with increasing number of active APs due to uplink collisions and hidden terminals.

FD^2 provides the highest gains compare to all schemes across all of the sub-topologies, with less than 3% loss in throughput due to remaining multi-cell interference and feedback overhead. On average, in these experiments FD^2 increases throughput by a factor of 9 compared to Omni.

Finding: FD^2 uses scheduled transmission for both UL and DL, eliminating contention and hidden terminals. FD^2 operates in FD mode and uses directional Tx and RX antennas not only to increase directionality gain, but most importantly to eliminate AP-AP and client-client interference through appropriate scheduling and beam selection.

Impact of Hardware Capture Properties. FD^2's client scheduling and beam selection and its overall throughput depend on the capture properties of its hardware. We now explore FD^2's performance when we select the rates according to 802.11a's measured capture table [18]. In order to measure the throughput, we measure clients' and APs' received RSSI and interference information, and derive the corresponding schedules and beams. Next, we map the resultant SINR to 802.11a's capture rate [18], and calculate the resulting throughput. We consider a time shared TDMA scheme when calculating DIRC's expected uplink throughput.

Fig. 4(g) shows that FD^2 has significantly increased the aggregate throughput compared to both Janus and DIRC. 802.11a supports rates from 6 to 54 Mbps with OFDM modulation. It also supports basic rates of 1 and 2 Mbps with DSSS modulation. Thus, FD^2 has the potential to achieve single link throughput as high as 54Mbps. This in turn results in additional increase in throughput as compared to WARP.

Finding: FD^2's gains are highly dependent on the SNR-rate table and capture properties of the underlying hardware. The gains can significantly increase when the dynamic range of a rate table is high, or when the hardware can implement *capture with lower SINR margins.*

Handling Traffic Dynamics. Fig. 4(h) shows how FD^2 responds to changes in traffic pattern by changing its beams and schedules. In this experiment, initially only client 1 receives DL traffic from the AP achieving more than 16Mbps throughput. At time t=4s, client C2 starts transmitting UL traffic. FD^2 adapts its beam and schedule to accommodate the request by the UL user, achieving balanced throughput for both UL and DL. At time t=8s, AP3 starts receiving DL traffic for client 5. FD^2 adapts its beams to easily fit in the new user with minimal impact on the other users.

C. Impact of Channel Dynamics.

In this Section, we evaluate impact of channel dynamics on FD^2's performance. We first explore the relation between scanning interval and channel dynamics on the performance of FD^2. Next, we propose solutions to increase FD^2's robustness to client mobility and channel dynamics.

Scenario. We consider a topology consisting of a single AP and two clients: C1 with DL traffic and C2 with UL traffic. We measure FD^2's performance in the following scenarios: *(i)* a scenario when two people move around the two clients, but clients remain static; and *(ii)* a scenario in which the clients move towards each other with controllable speeds. We developed a controllable platform to move the two clients towards each other. This allows us to emulate client movements in a controllable and repeatable manner.

Scanning Interval and Channel Dynamics. Fig. 5(a) shows FD^2's performance with changing channel dynamics as a function of scanning interval. Specifically, we vary the time scale of interference/RSSI information that is available at the central controller from 20ms up to 2.5s. We observe that FD^2's performance drops as the time scale of information feedback is increased, or as the users start to move for a fixed feedback time scale. The drop in performance is significant with a client speed of 1$\frac{m}{s}$ (typical pedestrian speed). With 2.5s scanning interval, FD^2's aggregate throughput (both UL and DL) drops to even lower than half duplex Omni.

In order to better understand the reason for the drop in throughput with mobility, we show a snapshot of links' throughputs in Figs. 5(b) (from t= 0s to 6s) and 5(c) (from t = 2.5s to 3.5s), and a snapshot of PDR in Fig. 5(d) (from t = 2.5s to 3.5s). In all, the scanning interval is set to 500ms.

Prior to t = 2s, the two clients are static and achieve similar throughput values with an average of 18Mbps. At t = 2s, clients move towards each other. Fig. 5(b) shows that UL throughput drops to 10 Mbps and remains the same for the rest of the experiment. DL throughput drops to 5Mbps and decreases as the users move close to each other due to increased interference.

Figs. 5(c) and (d) show that UL throughput and PDR remain high for the selected transmission rate and during the first 300ms of the interval, but rapidly drop to zero after that. This is because the selected rate is only valid for a limited time which depends on the radiation pattern of the transmitting directional antenna and user mobility. On the other hand, DL throughput only remains high for the first 100ms interval and drops to zero after that. DL throughput is specifically susceptible to mobility, not only because of rate dependence on Tx antenna pattern and user speed, but also because of increasing interference caused by the UL client(s).

Finding: Channel dynamics reduce the effective SINR of FD^2. The level of reduction depends on the radiation pattern of directional antennas, scanning interval, user speed, and UL/DL direction of traffic. Outdated RSSI information due to high mobility or long scanning intervals can reduce FD^2's performance to even lower than Omni.

TABLE I
REQUIRED SINR TO SELECT A 12MBPS RATE IN AN UNTRAINED SYSTEM VS. A TRAINED SYSTEM WITH t_f = 500ms AND s = 1 $\frac{m}{s}$

	Untrained	Trained-UL	Trained-DL
SINR	(20, 22]	(25.4, 30.1]	(27, 34.3]

Training, and Impact on Robustness. As seen in Figs. 5(a)-(d), mobility of the users can significantly reduce FD^2's performance in terms of both throughput and PDR. This could result in significant scanning and feedback overhead especially with a high number of clients, APs, or the number of directional antennas. Thus, in any practical system it is desirable to reduce the overhead.

Since we have no control over users' mobility and would like to keep the scanning interval to a fixed value to minimize the overhead, the resulting outdated information quickly reduces the effective SINR, which in turn decreases both throughput and PDR. Since the inaccuracy in RSSI information is related to the antenna pattern (g), scanning interval (t_f), user mobility speed (s), and UL/DL direction of client traffic (b), we propose to train FD^2's SINR-rate profile to these factors. FD^2 then categorizes clients based on their (g, t_f, s, b) values and applies the appropriate rate when selecting the corresponding schedules and beams.

To train FD^2's rate table according to (g, t_f, s, b), we perform experiments in an indoor environment. For a fixed initial distance between the clients and for a given t_f, we move the clients towards each other with speed s. Next, we employ FD^2 and transmit UL and DL traffic to both clients. We repeat the experiment with all rates supported by the WARP platform and for different initial distances between the clients and between clients and APs to achieve different UL-DL SINR values. Table 1 shows the required SINR to select 12Mbps rate for different schemes. Table 1 shows that a trained system requires a higher SINR for a given rate in order to compensate for infrequent feedback and channel dynamics. The required SINR for DL transmission is higher to compensate for potential increased UL to DL interference due to mobility.

We next quantify the gains of training FD^2 based on (g, t_f, s, b) values. We use the same mobility setup of Fig. 5(a). We plot the corresponding throughput and PDR results in Figs. 5(c) and (d), respectively. We observe that the achieved throughput values remain similar for the entire of the two

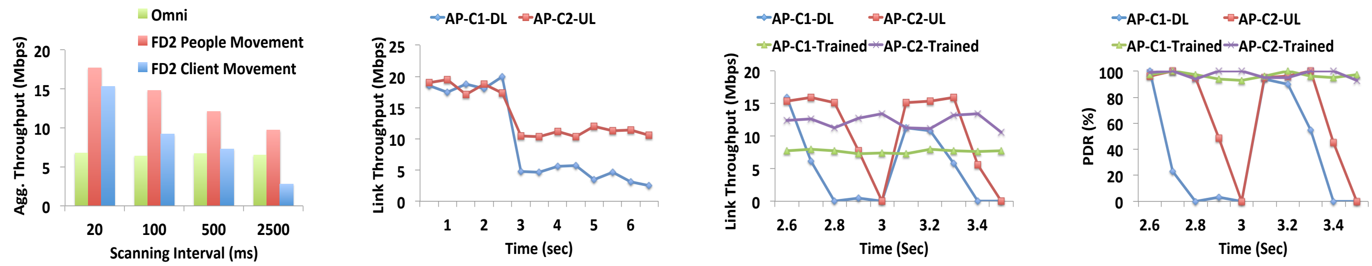

(a) (b) (c) (d)

Fig. 5. **Impact of scanning interval on FD^2's performance (a), Snapshot of link throughput variation due to client mobility (b), impact of training on FD^2's link throughput (c), impact of training on FD^2's Packet Delivery Ratio (d).**

$500ms$ intervals (prior to the time when a new scanning is performed). Compared to an untrained system, trained UL and DL throughputs increase by 15% and 51%, respectively. Moreover, PDR remains stable and close to 100% during the experiment.

Finding: Training a rate table according to (g, t_f, s, b) in addition to capture properties allows FD^2 to combat channel dynamics due to mobility or long scanning intervals.

VII. Conclusions

In this paper, we presented the design and implementation of FD^2, a directional FD communication system for indoor multi-cell networks. We proposed efficient algorithms to solve the joint scheduling and beam selection problem. We implemented FD^2 on the WARP platform, and showed significant gains compared to both CSMA/CA and competing centralized solutions. We also investigated the impact of client mobility on UL and DL performance, and proposed solutions to increase robustness to channel dynamics.

References

[1] J. Choi, M. Jain, K. Srinivasan, P. Levis, and S. Katti, "Achieving single channel, full duplex wireless communication," in *MobiCom*, 2010.

[2] M. Jain, T.M. Kim, D. Bharadia, S. Seth, K. Srinivasan, P. Levis, S. Katti, and P. Sinha, "Practical, real-time, full duplex wireless," in *MobiCom*, 2011.

[3] M. Duarte, C. Dick, and A. Sabharwal, "Experiment-driven Characterization of Full-duplex wireless systems," in *IEEE Transactions on Wireless Communications*, 2012.

[4] D. Bharadia, E. McMilin, and S. Katti, "Full duplex radios," in *SIGCOMM*, 2013.

[5] R. K. Patra, S. Nedevschi, S. Surana, A. Sheth, L. Subramanian, and E. A. Brewer, "WiLDNet: design and implementation of high performance WiFi based long distance networks," in *NSDI*, 2007.

[6] B. Raman and K. Chebrolu, "Design and evaluation of a new MAC protocol for long-distance 802.11 mesh networks," in *MobiCom*, 2005.

[7] A. P. Subramanian, H. Lundgren, and T. Salonidis, "Experimental characterization of sectorized antennas in dense 802.11 wireless mesh networks," in *MobiHoc*, 2009.

[8] X. Liu, A. Sheth, M. Kaminsky, K. Papagiannaki, S. Seshan, and P. Steenkiste, "DIRC: increasing indoor wireless capacity using directional antennas," in *SIGCOMM*, 2009.

[9] X. Liu, A. Sheth, M. Kaminsky, K. Papagiannaki, S. Seshan, and P. Steenkiste, "Pushing the envelope of indoor wireless spatial reuse using directional access points and clients," in *MobiCom*, 2010.

[10] A. K. Khandani, "Methods for spatial multiplexing of wireless two-way channels," United States Patents, Patent No.: US 7,817,641 B1, Oct 2010.

[11] E. Everett, A. Sahai, and A. Sabharwal, "Passive self-interference suppression for full-duplex infrastructure nodes," in *IEEE Transactions on Wireless Communications*, 2013.

[12] M. A. Khojastepour, E. Aryafar, K. Sundaresan, R. Mahindra, and S. Rangarajan, "Exploring the potential for full duplex in legacy LTE systems," in *Secon*, 2014.

[13] V. Navda, A. P. Subramanian, K. Dhanasekaran, A. Timm-Giel, and S. Das, "MobiSteer: using steerable beam directional antenna for vehicular network access," in *MobiSys*, 2007.

[14] M. Blanco, R. Kokku, K. Ramachandran, S. Rangarajan, and K. Sundaresan, "On the effectiveness of switched beam antennas in indoor environments," in *PAM*, 2008.

[15] T. H. Kim, T. Salonidis, and H. Lundgren, "MIMO wireless networks with directional antennas in indoor environments," in *Infocom*, 2012.

[16] A. Sabharwal, P. Schniter, D. Guo, D. Bliss, S. Rangarajan, and R. Wichman, "In-band full duplex wireless: challenges and opportunities," in *JSAC Special Issue on Full-Duplex Wireless Communications and Networking*, 2014.

[17] E. Aryafar, M.A. Khojastepour, K. Sundaresan, S. Rangarajan, and M. Chiang, "MIDU: Enabling MIMO full duplex," in *MobiCom*, 2012.

[18] J. Lee, W. Kim, S. J. Lee, D. Joe, J. Ryu, T. Kwon, and Y. Choi, "An experimental study on the capture effect in 802.11a networks," in *WinTECH*, 2007.

[19] "L-com: 2.4 ghz 60 degree dual-polarized antenna," Available at: http://www.l-com.com.

[20] S. Zhang, G. Huff, J. Feng, and J. Bernhard, "A pattern reconfigurable microstrip parasitic array," in *IEEE Transactions on Antenna and Propagations*, 2004.

[21] V. Shrivastava, N. Ahmed, S. Rayanchu, S. Banerjee, S. Keshav, K. Papagiannaki, and A. Mishra, "CENTAUR: realizing the full potential of centralized WLANs through a hybrid data path," in *MobiCom*, 2009.

[22] J. Y. Kim, O. Mashayekhi, H. Qu, M. Kaz, and P. Levis, "Janus: a novel MAC protocol for full duplex radio," in *Available at: http://hci.stanford.edu/cstr/reports/2013-02.pdf*, 2013.

[23] "Technical Report," Available at: https://dl.dropboxusercontent.com/u/22524598/FD2-TechReport.pdf.

[24] "Rice University WARP project," Available at: http://warp.rice.edu.