

Design and Testbed Deployment of Frequency-Domain Equalization Full Duplex Radios

Manav Kohli, Mahmood Baraani Dastjerdi, Jin Zhou, Ivan Seskar,
Harish Krishnaswamy, Gil Zussman, Tingjun Chen

Abstract—Full-duplex (FD) wireless can significantly enhance spectrum efficiency but requires effective self-interference (SI) cancellers. RF SI cancellation (SIC) via frequency-domain equalization (FDE), where bandpass filters channelize the SI, is suited for integrated circuits (ICs). In this paper, we explore the limits and higher layer challenges associated with using such cancellers. We evaluate the performance of a custom FDE-based canceller using two testbeds; one with mobile FD radios and the other with upgraded, static FD radios in the PAWR COSMOS testbed. The latter is a lasting artifact for the research community, alongside a dataset containing baseband waveforms captured on the COSMOS FD radios, facilitating FD-related experimentation at the higher networking layers. We evaluate the performance of the FDE-based FD radios in both testbeds, with experiments showing 95 dB overall achieved SIC (52 dB from RF SIC) across 20 MHz bandwidth. We conduct network-level experiments for (i) uplink-downlink networks with inter-user interference, and (ii) heterogeneous networks with half-duplex and FD users, showing FD gains of $1.14\times$ – $1.25\times$ and $1.25\times$ – $1.73\times$, respectively, confirming analytical results. We also evaluate the performance of an FD jammer-receiver, demonstrating a strong dependence on relative transmit power levels and modulation schemes.

Index Terms—Full-duplex wireless, frequency-domain equalization, self-interference cancellation, software-defined radios, wireless experimentation testbeds

I. INTRODUCTION

Full-duplex (FD) wireless – simultaneous transmission and reception on the same frequency channel – can significantly improve spectrum efficiency at the physical (PHY) layer and provide many other benefits at the higher layers [2]–[4]. The main challenge associated with FD is the extremely strong self-interference (SI) signal that needs to be suppressed, requiring 80–110 dB of SI cancellation (SIC).

Prior work leveraging off-the-shelf components and software-defined radios (SDRs) has established the feasibility

This work was supported in part by NSF grants CNS-1827923, EEC-2133516, DGE-2036197, CNS-2148128 and by funds from federal agency and industry partners as specified in the NSF Resilient & Intelligent NextG Systems (RINGS) program. A partial and preliminary version of this paper appeared in ACM MobiCom’19, Oct. 2019 [1].

M. Kohli and M. Baraani Dastjerdi were with the Department of Electrical Engineering, Columbia University. They are now with Apple Inc., Sunnyvale, CA 94085, USA (email: {mpk2138, b.mahmood}@columbia.edu).

H. Krishnaswamy, and G. Zussman are with the Department of Electrical Engineering, Columbia University, New York, NY 10027, USA (email: {harish, gil}@ee.columbia.edu).

J. Zhou was with the Department of Electrical Engineering, Columbia University. He is now with MediaTek, San Jose, CA 95134, USA (e-mail: jz2495@columbia.edu).

I. Seskar is with WINLAB, Rutgers University, New Brunswick, NJ 08902, USA (e-mail: seskar@winlab.rutgers.edu)

T. Chen is with the Department of Electrical and Computer Engineering, Duke University, Durham, NC 27705, USA (e-mail: tingjun.chen@duke.edu).

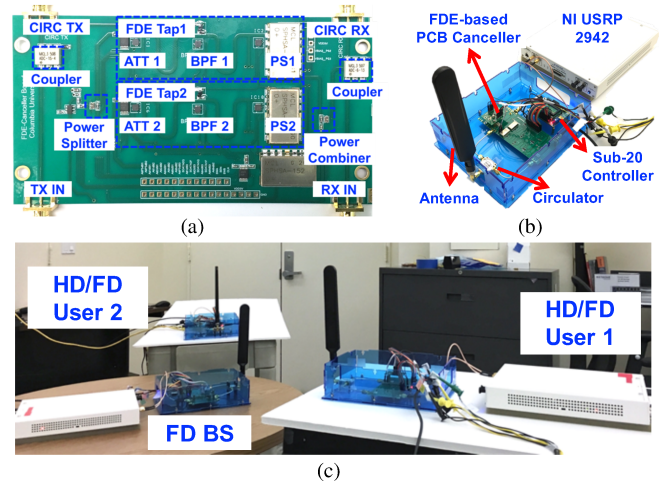


Fig. 1: (a) The frequency-domain equalization- (FDE-) based wideband RF canceller implemented using discrete components on a printed circuit board (PCB), (b) the implemented FDE-based mobile full-duplex (FD) radio, and (c) the mobile testbed consisting of an FD base station (BS) and 2 users mounted on carts that can operate in either half-duplex (HD) or FD mode [1].




Fig. 2: Integration of the FDE-based FD radios in COSMOS Sandbox 2. (a) labelled diagram of the “canceller box”, showing the various system components. (b) Two out of the four canceller boxes mounted in the testbed.

of FD wireless through SI suppression at the antenna interface, and SIC in analog/RF and digital domains [5]–[9]. However, RF cancellers achieving wideband SIC at sub-6 GHz operating frequencies (e.g., [7], [8]) commonly rely on transmission-line delays, which cannot be realized in small-form-factor nodes and/or integrated circuits (ICs) due to the required length for generating nanosecond-scale time delays and the lossy nature of the silicon substrate.¹

A compact IC-based design is necessary for supporting FD in hand-held devices (e.g., handsets and tablets) [3], [9]–[14]. In this paper, we focus on the FDE method initially presented in [11], which in contrast to the delay line-based approaches that essentially perform time-domain equalization, utilizes tunable, reconfigurable, high quality factor 2nd-order

¹For instance, obtaining a nanosecond delay in silicon typically requires a 15 cm-long delay line.

bandpass filters (BPFs) with amplitude and phase controls to emulate the frequency-selective antenna interface. In general, tunable BPFs with high quality factors can be as hard to implement on an IC as nanosecond-scale delay lines; recent advances in N -path filters [15] has simplified such designs.

While major advances have been made at the RFIC level, existing work has several limitations: (i) the limits of achievable RF SIC for FDE-based FD systems have not been fully understood, (ii) RF canceller configuration schemes need to be developed in order to achieve optimized and adaptive RF SIC in real-world environments, and (iii) the system-level performance of such IC-based FD radios has not been evaluated in different network settings. Therefore, in this paper we focus on the evaluation of FDE-based RF cancellers in two experimental testbeds: a lab-based *mobile FD testbed* (Fig. 1) and the *COSMOS FD testbed* (Fig. 2).

As interfacing an RFIC canceller to an SDR presents numerous technical challenges, we design and implement two iterations of a printed circuit board (PCB) emulation of the FDE-based RFIC canceller presented in [11] using discrete components. The first iteration of the PCB canceller appears in the *mobile FD radios* shown in Fig. 1. The second, more robust iteration of the PCB canceller has been integrated in the PAWR COSMOS testbed [16], shown in Fig. 2. The *COSMOS FD radios* are intended to be a lasting artifact for the research community, as they are openly and remotely accessible for use in FD experimentation [17].

The COSMOS FD testbed improves on the first such effort by the Columbia FlexICoN project [18] using the ORBIT testbed [19]. Specifically, the COSMOS FD testbed contains four FDE-based FD radios operating at $4\times$ the bandwidth as ORBIT's single narrowband FD radio [20]. The COSMOS FD testbed has been improved from its prior iteration [19] through the uniform use of higher performance SDRs for each FD radio and integration with the COSMOS servers, which provide significantly greater compute capability. Additionally, we provide a publicly available dataset containing baseband I/Q samples for OFDM-modulated packets.

We create the FDE-based FD radios for each testbed by integrating the PCB canceller with an SDR, as depicted in Figs. 1(b) and 2(b). Both the mobile and COSMOS FD radios achieve 85–95 dB overall SIC across 20 MHz real-time bandwidth, enabling an FD link budget of 0–10 dBm average TX power level and -85 dBm RX noise floor. In particular, up to 52 dB RF SIC is achieved, from which 20 dB is obtained from the antenna interface isolation. We present a realistic model of the PCB canceller, and use it to develop a configuration scheme based on an optimization problem for the mobile FD radios, allowing efficient adaption of the canceller to environmental changes. The PCB canceller model is experimentally validated and is shown to have high accuracy.

Using both testbeds, we extensively evaluate the network-level FD gain and confirm analytical results in two types of networks: (i) *UL-DL networks* consisting of one FD base station (BS) and two half-duplex (HD) users with inter-user interference (IUI), and (ii) *heterogeneous HD-FD networks* consisting of one FD BS and co-existing HD and FD users. For the UL-DL network on the mobile testbed, we show an

empirical throughput gain between $1.14\times$ – $1.25\times$ compared to $1.22\times$ – $1.30\times$ predicted by analysis. For heterogeneous HD-FD networks, we demonstrate the impact of different user SNR values and the number of FD users on both the FD gain and throughput fairness. For example, in a 4-node network consisting of an FD BS and 3 users with various user SNR values, the median network throughput can be improved by $1.25\times$ and $1.52\times$ when one and two users become FD-capable, respectively. The COSMOS FD testbed provides very similar results for these two types of networks under appropriately defined experimental conditions.

To the best of our knowledge, this is the first experimental study of FDE-based FD radios in such networks using testbeds comprised of both HD and FD radios. The results demonstrate the practicality and validate the performance of FDE-based FD radios, which are suitable for small-form-factor devices. The presented results, as well as openly-accessible COSMOS FD radios, can also serve as building blocks for developing higher layer (e.g. MAC) protocols.

To summarize, the main contributions of the paper are:

1. We present the design, implementation, modeling, and validation of the FDE-based PCB canceller, as well as an optimized canceller configuration scheme;
2. We experimentally evaluate the FD throughput gain using the FDE-based FD radios in various network settings with different user duplexing and SNR values.
3. We provide the open-access COSMOS FD testbed as a lasting research artifact supporting experimentation with FD wireless based on the principle of FDE, suitable for small form-factor implementations.

The rest of the paper is organized as follows. Section II reviews related work. In Section III, we present the design, implementation, and model of the FDE-based PCB canceller, as well as the optimized canceller configuration scheme. We evaluate the performance of FD radios equipped with the FDE PCB cancellers in a mobile testbed in Section IV and in the static COSMOS testbed in Section V. We also provide information about the publicly available RF SIC waveform dataset in Section V. Lastly, we conclude and discuss future directions in Section VI.

II. RELATED WORK

Extensive research related to FD wireless is summarized in [2], [9], [21], [22], including implementations of FD radios and systems, analysis of rate gains, and resource allocation at the higher layers. Below, we briefly review the related work.

RF canceller and FD radio designs. RF SIC typically involves two stages: (i) isolation at the antenna interface, and (ii) SIC in the RF domain using cancellation circuitry. While a separate TX/RX antenna pair can provide good isolation and can be used to achieve cancellation [5], [23]–[25], a shared antenna interface such as a circulator is more appropriate for single-antenna implementations [7], [26] and is compatible with FD MIMO systems. Existing designs of analog/RF SIC circuitry are mostly based on a time-domain interpolation approach [7], [8], [27], [28]. Several FD MIMO radio designs based on this approach are presented [25], [29]–[32]. FD relays have also been successfully demonstrated

from the circulator TX port and is then split to the two FDE taps through a Wilkinson power divider. Both FDE taps are identical, and consist of an attenuator, a second-order bandpass filter (BPF), and a phase shifter. Lastly, the two FDE taps are then combined after the phase shifters, and this combined signal is added at the circulator RX port.

Unlike the RFIC implementation, which uses an N -path design, the PCB BPF is implemented as an RLC tank, as shown in Fig. 5. The PCB BPF is matched to the attenuator and phase shifter through two identical impedance transformation networks and transmission lines with $Z_0 = 50\Omega$ characteristic impedance around the target 900–915 MHz frequency band of the PCB canceller. This is within the Region 2 902–928 MHz ISM band, but we note that this canceller design can be extended to other bands (e.g. 2.4 GHz) with appropriate adjustment to the frequency-sensitive components.

The center frequency of the PCB BPF in the i^{th} FDE tap is adjusted through the capacitor $C_{F,i}$ in the RLC tank as shown in Figure 5. In order to achieve a high and adjustable BPF quality factor, another variable capacitance $C_{Q,i}$ within the left and right impedance transformation networks can be adjusted. Fig. 4(b) demonstrates the effect of varying these two capacitance values on the center frequency and quality factor of the PCB BPF. The adjustable attenuator and phase shifter provide additional shaping of the PCB BPF response.

Compared to the N -path filter used in the RFIC canceller [11] that has constant DC power draw, the PCB BPF has zero DC power consumption. It also supports higher TX power levels and has a lower noise figure, providing additional benefit when interfacing the PCB canceller with an SDR.

C. PCB Cancellable Model

The RFIC FDE canceller implementation based on N -path filters has the frequency response $H^1(f)$ given in Eq. 1 below [11], equivalent to the response of a combined bank of M 2nd-order BPFs:

$$H^1(f) = \sum_{i=1}^M H_i^1(f) = \sum_{i=1}^M \frac{A_i^1 \cdot e^{-j\phi_i^1}}{1 - jQ_i \cdot (f_{c,i}/f - f/f_{c,i})}. \quad (1)$$

Ideally, the PCB BPF would also have a 2nd-order BPF frequency response from the RLC tank. However, in practical implementation, its response deviates from the FDE-based RFIC canceller (1), and a more precise model is required to develop an optimization problem for the configurable PCB values. Based on the circuit diagram in Fig. 5, we derive a realistic model for the frequency response of the PCB BPF, $H_i^B(f)$, given by Eq. 2.

$$H_i^B(f) = R_s^{-1} \left[j \sin(2\beta l) Z_0 Y_{F,i}(f) Y_{Q,i}(f) + \cos^2(\beta l) Y_{F,i}(f) + 2 \cos(2\beta l) Y_{Q,i}(f) + j \sin(2\beta l) / Z_0 + 0.5 j \sin(2\beta l) Z_0 (Y_{Q,i}(f))^2 - \sin^2(\beta l) Z_0^2 Y_{F,i}(f) (Y_{Q,i}(f))^2 \right]^{-1}, \quad (2)$$

where $Y_{F,i}(f)$ and $Y_{Q,i}(f)$ are the admittance of the RLC resonance tank and impedance transformation networks, i.e.,

$$\begin{aligned} Y_{F,i}(f) &= 1/R_F + j2\pi C_{F,i}f + 1/(j2\pi L_F f), \\ Y_{Q,i}(f) &= 1/R_Q + j2\pi C_{Q,i}f + 1/(j2\pi L_Q f). \end{aligned} \quad (3)$$

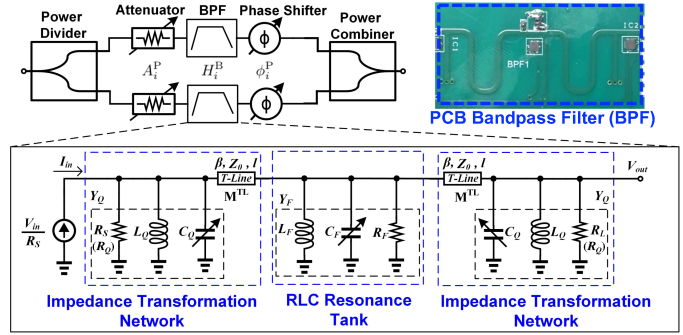


Fig. 5: Block diagram of the implemented 2 FDE taps in the PCB canceller (see Fig. 4(a)), each of which consists of an RLC bandpass filter (BPF), an attenuator for amplitude control, and a phase shifter for phase control [1].

R_S and R_L are set to the same value of $R_Q = 50\Omega$ to match the source and load impedances of the RLC resonance tank (see Fig. 5). The implemented PCB BF transmission line (see Fig. 5) has phase $\beta l = 1.37$ rad and characteristic impedance $Z_0 = 50\Omega$. Fixed inductance values of $L_F = 1.65$ nH and $L_Q = 2.85$ nH are provided by a discrete component and shorted stub line, respectively. The variable capacitance C_F has a constant component of 8.2 pF.

Recall that each FDE tap is also associated with amplitude and phase controls, A_i^P and ϕ_i^P , which act on the PCB BPF response $H_i^B(f)$. Lastly, we quantified implementation losses and group delays for the entire PCB canceller, denoted by A_0^P and τ_0^P . Through measurement of the PCB as well as S-parameters for the components used, we empirically set $A_0 = -4.1$ dB and $\tau_0 = 4.2$ ns. Altogether, the model for the PCB canceller with two PCB BPF taps is given as

$$H^P(f) = A_0^P e^{-j2\pi f \tau_0^P} \left[\sum_{i=1}^2 A_i^P e^{-j\phi_i^P} H_i^B(f) \right], \quad (4)$$

where $H_i^B(f)$ is the PCB BPF model given by (2). As a result, the i^{th} FDE tap in the PCB canceller (4) has configuration parameters $\{A_i^P, \phi_i^P, C_{F,i}, C_{Q,i}\}$, giving four degrees of freedom to optimize over. Table I summarizes the notation for these parameters as well as other parts of the model.

D. PCB Cancellable Optimization

The PCB canceller model can be used within an optimization problem to find the ideal values of the four component values for each FDE tap. Denote by $H_{SI}(f)$ the SI channel from the radio antenna interface (i.e. with the PCB canceller turned off). As shown in Fig. 4(a), the negative frequency response of the FDE canceller $H^{\text{FDE}}(f)$ is summed at the input to the RX LNA. Considering the use of the PCB canceller, the residual SI at the radio RX, $H_{\text{res}}(f)$, is given as

$$H_{\text{res}}(f) = H_{SI}(f) - H^P(f), \quad (5)$$

where $H^P(f)$ is the PCB canceller model given by (4). Naturally, the goal of the PCB canceller would be to minimize $H_{\text{res}}(f)$ across the operational bandwidth B .

Many practical implementations make use of an OFDM-based PHY layer, as do the experiments presented in Sections IV and V. In these cases, the operational bandwidth B is divided into K orthogonal frequency channels, each indexed

TABLE I: Nomenclature

$ z , \angle z$	Amplitude and phase of a complex value $z = x + jy$ ($x, y \in \mathbb{R}$), where $ z = \sqrt{x^2 + y^2}$ and $\angle z = \tan^{-1}(\frac{y}{x})$
B	Total wireless bandwidth/desired RF SIC bandwidth
K, k	Total number of frequency channels and channel index
f_k	Center frequency of the k^{th} frequency channel
M	Number of FDE taps in an FDE-based RF canceller
$H_{\text{SI}}(f_k)$	Frequency response of the antenna interface
$H^{\text{P}}(f_k)$	Frequency response of the FDE-based PCB canceller
$H_i^{\text{P}}(f_k)$	Frequency response of the i^{th} FDE tap in the PCB canceller
$A_i^{\text{P}}, \phi_i^{\text{P}}$	Amplitude and phase controls of the i^{th} FDE tap in the PCB canceller
$C_{\text{F},i}, C_{\text{Q},i}$	Digitally tunable capacitors that control the center frequency and quality factor of the i^{th} FDE tap in the PCB canceller

by $k \in \{1, \dots, K\}$. The center frequency of the k^{th} frequency channel is denoted by f_k ; henceforth this discrete variable will replace the continuous frequency parameter f in the model expressions. With this in mind, the ideal component values can be solved for using the minimization problem (P1):

$$\begin{aligned}
 \text{(P1) } \min : & \sum_{k=1}^K |H_{\text{res}}^{\text{P}}(f_k)| = \sum_{k=1}^K |H_{\text{SI}}(f_k) - H^{\text{P}}(f_k)|^2 \\
 \text{s.t.: } & A_i^{\text{P}} \in [A_{\text{min}}^{\text{P}}, A_{\text{max}}^{\text{P}}], \phi_i^{\text{P}} \in [-\pi, \pi], \\
 & C_{\text{F},i} \in [C_{\text{F,min}}, C_{\text{F,max}}], C_{\text{Q},i} \in [C_{\text{Q,min}}, C_{\text{Q,max}}], \forall i.
 \end{aligned}$$

Note that (P1) is challenging to solve due to its non-convexity and non-linearity, caused by several properties of (4) such as (i) the higher-order terms introduced by f_k , and (ii) the trigonometric term introduced by the phase control, ϕ_i^{P} . In addition, the antenna interface response, $H_{\text{SI}}(f_k)$, is also frequency-selective and time-varying.

However, in practical scenarios, it is unnecessary to find the true global minimum of (P1). As described in III-A, the amount of overall SIC that must be achieved depends on multiple factors, and there are several stages of SIC that may be performed. As long as the RF SIC is sufficient for the digital SIC stage to suppress the residual SI to the noise floor of the RX, overall system performance has been optimized. In practice, we find an RF SIC of 45 dB to be sufficient. This quality also alleviates the effect of quantization error on the overall SIC, as the tunable components only support a discrete set of values. Additionally, we find that the components used in the PCB canceller have high enough resolution to avoid excessive degradation in performance compared to the ideal, non-quantized case [1].

IV. EXPERIMENTATION IN THE MOBILE TESTBED

In this section, we discuss the integration of the PCB canceller described in Section III-B with a mobile FD testbed. Then, we present an extensive testbed evaluation of the mobile FD radios at the node and network levels; link-level results are available in [1].

A. Node-Level Implementation and Performance

The overall design of the mobile FD radios and SDR testbed is shown in Figs. 1(b) and 1(c). Each FD radio uses an 860-960 MHz circulator for the antenna interface, and the PCB canceller serves as the frontend of a USRP-2942 SDR with an SBX-120 daughterboard. Each of the three FD radios in




Fig. 6: An example experimental setup for: (a) the UL-DL networks with varying γ_{UL} and γ_{DL} , (b) heterogeneous 3-node network with one FD BS and 2 FD users, and (c) heterogeneous 4-node networks with one FD BS, 2 FD users, and one HD user [1].

TABLE II: Average FD Gain in UL-DL Networks with IUI.

UL SNR, γ_{UL}	Analytical FD Gain	Experimental FD Gain
10 dB	1.30×	1.25×
15 dB	1.23×	1.16×
20 dB	1.22×	1.14×

the mobile testbed operates at 900 MHz carrier frequency, and further radios without the PCB canceller are used as HD users.

The mobile FD radios are supported by an OFDM PHY layer running at 20 MHz bandwidth implemented in NI LabVIEW on a host PC. This PC also runs the optimized PCB canceller configuration scheme implemented in MATLAB, given in detail in [1].

Achievable SIC. With +10 dBm average TX power, the PCB canceller achieves up to 52 dB RF SIC across 20 MHz bandwidth, from which 20 dB is obtained from the circulator. Furthermore, we find that the frequency profile of the achieved RF SIC is similar to that presented in the sensitivity analysis of the FDE-based cancellers [1].

In addition to the 52 dB RF SIC, a further 40–43 dB is experimentally achieved in the digital domain [1]. This gives a total of 95 dB SIC across 20 MHz, sufficient to reduce residual SI to a level below the -85 dBm noise floor of the USRP-2942 and achieve the full dynamic range of 30–50 dB depending on the modulation and coding scheme (MCS) used. This performance is consistent for all experiments in this section.

B. Network-Level FD Gain

We now experimentally evaluate the network-level throughput gain introduced by FD-capable BS and users. Recall that user equipment (UEs) can significantly benefit from FDE-based FD implementations, suitable for RFIC implementations in hand-held devices. We compare experimental and analytical (e.g., [45]) results and demonstrate practical FD gain in different network settings. Specifically, we consider two types of networks as depicted in Fig. 6: (i) *UL-DL networks* with one FD BS and two HD users with inter-user interference (IUI), and (ii) *heterogeneous HD-FD networks* with HD and FD users. We implement a simple TDMA MAC layer in software where HD or FD users take turns to be activated for equally sized time slots. In Section V-B, we describe a similar set of experiments conducted on the COSMOS FD testbed.

1) *UL-DL Networks with IUI:* We first consider UL-DL networks consisting of one FD BS and two HD users (indexed 1 and 2). User 1 activates a UL to the BS, and the BS activates a DL to user 2, as shown in Fig. 6(a).

Analytical FD gain. We use Shannon's capacity formula $r(\gamma) = B \log_2(1 + \gamma)$ to compute the *analytical throughput* of a link under bandwidth B and HD link SNR γ . If the BS is




Fig. 7: Analytical (colored surface) and experimental (filled circles) network throughput gain for UL-DL networks consisting of one FD BS and two HD users with varying UL and DL SNR values, and inter-user interference (IUI) levels: (a) $\gamma_{UL} = 10$ dB, (b) $\gamma_{UL} = 15$ dB, and (c) $\gamma_{UL} = 20$ dB. The baseline is the network throughput when the BS is HD [1].

only HD-capable, the network throughput in a UL-DL network where UL and DL share the channel in a TDMA manner with equal fraction of time each is given by

$$r_{UL-DL}^{HD} = \frac{B}{2} \log_2(1 + \gamma_{UL}) + \frac{B}{2} \log_2(1 + \gamma_{DL}), \quad (6)$$

where γ_{UL} and γ_{DL} are the UL and DL SNRs, respectively. If the BS is FD-capable, the UL and DL can be simultaneously activated with an analytical network throughput of

$$r_{UL-DL}^{FD} = B \log_2\left(1 + \frac{\gamma_{UL}}{1 + \gamma_{Self}}\right) + B \log_2\left(1 + \frac{\gamma_{DL}}{1 + \gamma_{IUI}}\right), \quad (7)$$

where: (i) $\left(\frac{\gamma_{DL}}{1 + \gamma_{IUI}}\right)$ is the signal-to-interference-plus-noise ratio (SINR) at the DL HD user, and (ii) γ_{Self} is the residual self-interference-to-noise ratio (XINR) at the FD BS. We set $\gamma_{Self} = 1$ when computing the analytical throughput as the FDE-based FD radio can suppress the residual SI below the radio noise floor, as described in Section IV-A. The *analytical FD gain* is then defined as the ratio ($r_{UL-DL}^{FD}/r_{UL-DL}^{HD}$). Note that the FD gain depends on the coupling between γ_{UL} , γ_{DL} , and γ_{IUI} , which depend on the BS/user locations, their TX power levels, and whether the link is LOS or NLOS.

Experimental FD gain. The experimental setup is depicted in Fig. 6(a), where the TX power levels of the BS and user 1 are set to be 10 dBm and -10 dBm, respectively. We fix the location of the BS and consider different UL SNR values of $\gamma_{UL} = 10/15/20$ dB by placing user 1 at three different locations. For each value of γ_{UL} , user 2 is placed at ten different locations, resulting in varying γ_{DL} and γ_{IUI} values.

Fig. 7 shows the analytical (colored surface) and experimental (filled circles) FD gain, where the analytical gain is extracted using (6) and (7), and the experimental gain is computed using the measured UL and DL throughput. It can be seen that smaller values of γ_{UL} and lower ratios between γ_{DL} and γ_{IUI} lead to higher throughput gains in both analysis and experiments. The average analytical and experimental FD gains are summarized in Table II. The experimental FD gain is within 93% of the analytical FD gain, which shows a good match to the analysis in the presence of practical effects such as imperfect PRR. The results confirm the analysis in [45] and demonstrate the a FD gain achieved in UL-DL networks without any major changes in the network stack other than enabling FD capability at the BS and reorganizing the UL and DL schedule.

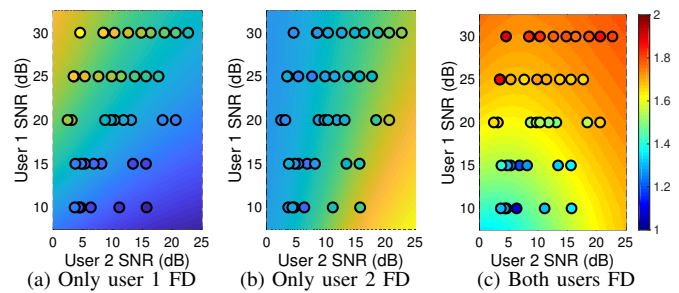


Fig. 8: Analytical (colored surface) and experimental (filled circles) network throughput gain for 3-node networks consisting of one FD BS and two users with varying link SNR values: (a) only user 1 is FD, (b) only user 2 is FD, and (c) both users are FD. The baseline is the network throughput when both users are HD [1].

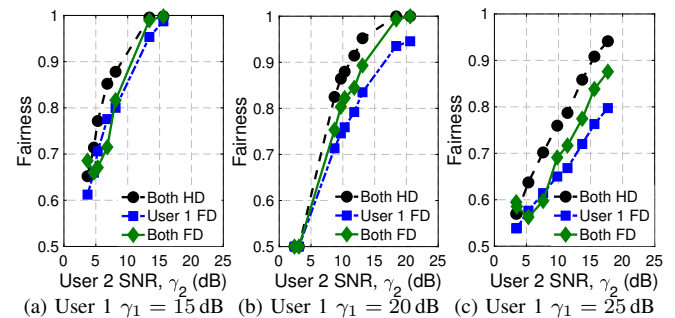


Fig. 9: Measured Jain's fairness index (JFI) in 3-node networks where both users are HD, User 1 is FD, and both users are FD, with varying (γ_1, γ_2) .

2) *Heterogeneous 3-Node Networks:* We consider heterogeneous HD-FD networks with 3 nodes: one FD BS and two users that can operate in either HD or FD mode (see an example experimental setup in Figs. 1(c) and 6(b)). All 3 nodes have the same 0 dBm TX power so that each user has symmetric UL and DL SNR values of γ_i ($i = 1, 2$). We place user 1 at 5 different locations and place user 2 at 10 different locations for each location of user 1, resulting in a total number of 50 pairs of (γ_1, γ_2) . We also consider a similar experiment running on the COSMOS testbed in Section V-B2.

Analytical FD gain. We set the users to share the channel in a TDMA manner. The analytical network throughput in a 3-node network when zero, one, and two users are FD-capable is given, respectively, by

$$r^{HD} = \frac{B}{2} \log_2(1 + \gamma_1) + \frac{B}{2} \log_2(1 + \gamma_2), \quad (8)$$

$$r_{User\ i\ FD}^{HD-FD} = B \log_2\left(1 + \frac{\gamma_i}{1 + \gamma_{Self}}\right) + \frac{B}{2} \log_2(1 + \gamma_{\bar{i}}), \quad (9)$$

$$r^{FD} = B \log_2\left(1 + \frac{\gamma_1}{1 + \gamma_{Self}}\right) + B \log_2\left(1 + \frac{\gamma_2}{1 + \gamma_{Self}}\right), \quad (10)$$

where $\gamma_{Self} = 1$ is set (similar to Section IV-B1). We consider both FD gains of ($r_{User\ i\ FD}^{HD-FD}/r^{HD}$) (i.e., user i is FD and user $\bar{i} \neq i$ is HD), and (r^{FD}/r^{HD}) (i.e., both users are FD).

Experimental FD gain. For each pair of (γ_1, γ_2) , experimental FD gain is measured in three cases: (i) only user 1 is FD, (ii) only user 2 is FD, and (iii) both users are FD. Fig. 8 shows the analytical (colored surface) and experimental (filled circles) FD gain for each case. We exclude the results with $\gamma_i < 3$ dB since the packets cannot be decoded, resulting in a




Fig. 10: CDF of the experimental network throughput in 4-node networks when zero, one, or two users are FD-capable [1].




Fig. 11: CDF of the measured Jain's fairness index (JFI) in 4-node networks when zero, one, or two users are FD-capable.

throughput of zero.

The results show that with small link SNR values, the experimental FD gain is lower than the analytical value due to the inability to decode the packets. On the other hand, with sufficient link SNR values, the experimental FD gain exceeds the analytical FD gain. This is because setting $\gamma_{\text{Self}} = 1$ in (9) and (10) results in a 3 dB SNR loss in the analytical FD link SNR, and thereby in a lower throughput. However, in practice, the packets can be decoded with a link PRR of 1 with sufficient link SNRs, resulting in exact twice number of packets being successfully sent over an FD link. Moreover, the FD gain is more significant when enabling FD capability for users with higher link SNR values.

Another important metric we consider is the fairness between users, which is measured by the Jain's fairness index (JFI). In the considered 3-node networks, the JFI ranges between 0.5 (worst case) and 1 (best case). Fig. 9 shows the measured JFI when both users operate in HD mode, user 1 operates in FD mode, and both users operate in FD mode, with varying user SNR values (γ_1, γ_2). The results show that introducing FD capability to both users results in an average degradation in the network JFI of only 5.6/4.4/7.4% for $\gamma_1 = 15/20/25$ dB (averaged across varying γ_2), while the average network FD gains are 1.32/1.58/1.73 \times (see Fig. 8), respectively. In addition, the JFI increases with more balanced user SNR values, which is as expected. For example, under the same value of γ_1 , increased value of γ_2 (with $\gamma_2 < \gamma_1$) leads to improved JFI, whose value approaches 1 as $\gamma_2 \rightarrow \gamma_1$.

3) *Heterogeneous 4-Node Networks*: We experimentally study 4-node networks consisting of an FD BS and three users with 10 dBm TX power (see an example experimental setup in Fig. 6(c)). The experimental setup is similar to that described in Section IV-B2, and the SNR value for the three users is varied by placing them at different locations. For each experiment, the network throughput is measured in three cases, where (i) zero, (ii) one, and (iii) two users are FD-capable.




Fig. 12: CDF of the analytical and experimental network throughput gains in 4-node networks when one or two users are FD-capable.

Fig. 10 shows the CDF of the network throughput of the three cases, where the measured link SNR varies between 5–45 dB. Overall, the median network throughput is increased by 1.25/1.52 \times when one/two FD users become FD-capable. Fig. 11 plots the CDF of the corresponding JFI, where although introducing FD-capable users results in lower values of the experimental JFI, the median JFI values degrade by only 0.06 and 0.10 with one and two FD users, respectively. Moreover, Fig. 12 shows the CDF of both the analytical and experimental network throughput gains in 4-node networks when one or two users are FD-capable. In particular, the median analytical and experimental network throughput gains have a difference of only 4% and 3% when one and two users are FD-capable. These trends and results show that in a real-world environment, the total network throughput increases as more users become FD-capable, and the improvement is more significant with higher user SNR values. Note that we only apply a TDMA scheme and a more advanced MAC layer (e.g. [54], [58]) has the potential to improve the FD gain and fairness performance. As with the 3-node experiment, a similar version is run on the COSMOS testbed in Section V-B3.

V. THE COSMOS FD RADIOS

The mobile FD testbed described in Section IV was developed for the purpose of evaluating the FDE-based FD radio performance under user mobility scenarios. While the mobile testbed served this purpose sufficiently, its ephemeral nature limits experimental reproducibility. To address this problem, we developed a second FD testbed deployed within the NSF PAWR COSMOS testbed [16] consisting of four improved FDE-based FD radios. Being a part of COSMOS, these FD radios are *open-access* and *remotely accessible*, so that they may be used by the broader research community to experiment with FD radios using custom hardware suitable for implementation on an RFIC.

In this section, we first describe the integration of various components into the FD testbed within COSMOS Sandbox 2 [64]. Next, we present experimental results for the UL-DL and heterogeneous TDMA network experiments, demonstrating similar FD performance to the mobile FD testbed over a range of similar scenarios. We then describe an experiment to evaluate a FD jammer-receiver. Lastly, we describe an available dataset with waveform traces to support the development of digital SIC algorithms.

A step-by-step procedure on how to access and use the testbed is provided by a tutorial [17] which makes use of the custom COSMOS server image `flexicon-cosmos.ndz`.




Fig. 13: Architecture of the FD testbed integrated in COSMOS Sandbox 2.

A. Design of the COSMOS FD Testbed

An initial version of the COSMOS FD testbed was presented in [19]. The testbed has since been extensively re-designed, with the overall architecture shown in Fig. 13. We detail the individual components making up the testbed below. **COSMOS FD Radio.** The COSMOS FD radios include a “canceller box”, as shown in Fig. 2(a), representing an upgrade over the one used in the mobile testbed in Fig. 1(b). The canceller box contains the PCB canceller as well as a circulator to support the single antenna interface. In addition, an antenna tuner with 20–30 MHz bandwidth is included to improve the antenna impedance matching. The box also contains a SUB-20 device for interfacing the PCB canceller with a computer.

Four canceller boxes are mounted on the corners of a 4.5×3.5 m rectangle within a square laboratory room (the same room housing the FD BS in Fig. 6), where they serve as the frontends of USRP X310 SDRs with SBX-120 RF daughterboards. Two of the four radios are shown in Fig. 2 (b); the same setup with another two canceller boxes is duplicated on another ceiling rack directly opposite.

Improved PCB canceller. The COSMOS FD radios use an improved PCB canceller design compared to the version used in Section IV. This improved design utilizes the same circuit with two BPF taps described in Section III, with some key modifications to enhance performance and reliability.

There are three main changes to the PCB: (i) the phase shifter control circuit was consolidated onto the same PCB as the BPF taps, absolving the need for a daughterboard and board-to-board connectors; (ii) the SPI control lines from the SUB-20 to the PCB canceller were simplified leading to reduced wiring within the canceller box; and (iii) the values for L_F and the constant value of the capacitance C_F in the analytical model from Section III-C were changed to 1 nH and 6.8 pF, respectively.

The COSMOS FD radios may be configured using the optimization problem defined in Section III-D, however we find that manual configuration of the FD radios is sufficient for evaluating the hardware within the testbed. This is primarily due to the static nature of the laboratory environment and the placement of the FD radios. Each of the four integrated improved PCB cancellers can achieve over 50 dB of SIC across 20 MHz within the 900 MHz band [19]. With digital SIC included, over 85 dB SIC may be achieved across 20 MHz. The PCB canceller design is publicly available, including the layout and bill of materials [65].

USRP X310. The prior iteration of the COSMOS FD testbed utilized a single USRP 2974 and two USRP N210s. The two types of SDR feature vastly different performance characteristics, predominantly caused by the different RF front ends.




Fig. 14: FD throughput gains as a function of the DL radio SINR for IUI values of 5, 10, and 15 dB.

These radios were replaced with two USRP X310s with SBX-120 daughterboards, similar to the ones used in Section IV. The USRP X310s provide uniform performance and operation for all four radio nodes, and permit integration with the COSMOS servers via their 10 Gb/s interface.

Integration with COSMOS servers. The single largest improvement to the COSMOS FD testbed has been to connect the USRP X310 SDRs and FDE canceller boxes to two Dell R740 servers via 10 Gb/s network links. Each server is equipped with two Xeon 12-core CPUs and 192 GB of memory. An experimenter can use these servers to control the SDRs and PCB canceller and develop experiments in GNU Radio or using the UHD driver directly [66]; progress may be saved by creating an image of the server that can be loaded at a later time. In addition, the USRP X310s and PCB cancellers may be accessed through a less powerful, dedicated-purpose FD Compute Node is also present within Sandbox 2 and may be used to run simple benchmarks without using the servers.

Experimentation software. The mobile FD testbed from Section IV uses a LabVIEW-based software stack. For the COSMOS FD testbed, we migrated all experimental software to GNU Radio, facilitating the imaging process used for the COSMOS servers to save and reload experimental state. The experimental software is contained within a custom out-of-tree (OOT) module for GNU Radio [65]. This custom C++ code supports real-time FD experimentation, including implementation of the linear digital SIC algorithm, and the MAC layer scheduling evaluated in Section V-B.

B. Network-Level Experimentation

In this subsection, we describe the experiments conducted on the COSMOS FD testbed.

1) *UL-DL Networks with IUI:* We re-implemented the UL-DL network from Section IV-B1 to run on the COSMOS FD nodes. As these FD radios are fixed in place, the IUI caused by user A on another user B is a function of A 's transmit power only. This reduces the explorable state space in this version of the experiment compared to the version with freely moving radios in the mobile testbed, where the IUI can be understood as a function of both radio transmit power and position.

We present the experimental FD rate gains for the UL-DL network as achieved on the testbed integration in Fig. 14. We considered three IUI levels at the DL user; 5, 10, and 15 dB. Using an 802.11-like PHY layer [67], we correspond these IUIs to specific MCS of QPSK 3/4, 16-QAM 1/2, and 16-QAM 3/4 respectively, which are used in the link between the UL user and BS. We sweep the transmit power of the

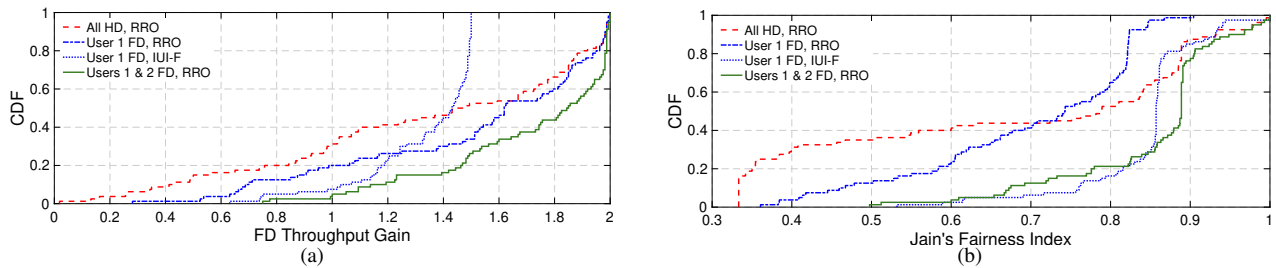


Fig. 15: Experimental results for the heterogeneous 3-node network. Results for (a) the FD throughput gain and (b) the JFI are presented as CDFs. Zero, one, or two FD users are considered, as well as the RRO and IUI-F schedules.




Fig. 16: Experimental results for the heterogeneous 4-node network. Results for (a) the FD throughput gain and (b) the JFI are presented as CDFs. Zero to three FD users are considered, along with the RRO and IUI-F schedules.

BS radio to generate SINR values at the DL user sufficient for decoding the BPSK 1/2, BPSK 3/4, and QPSK 1/2 MCS. This testbed setup is representative of a densely populated, low power network prone to IUI. We used the dedicated FD compute node to run these experiments; we note that it is possible to run them in an identical manner using a server.

Fig. 14 shows that throughput gains relative to the ideal HD case are achievable for all considered IUI values. However, the gain drops significantly as the IUI increases; at 5 dB IUI a gain of up to $1.6\times$ is achievable, whereas for 10 dB IUI the gain is reduced to a maximum of $1.3\times$. This effect matches the analytical surfaces shown in Fig. 7 which show a decrease in theoretical FD gain as a function of the IUI at the DL radio. We note that the specific FD rate gains shown in this experiment are comparable to those achieved in Section IV-B1, demonstrating the FD testbed’s ability to produce reliable and repeatable results using the FDE-based PCB canceller.

2) *Heterogeneous 3-Node Networks:* The experiment described in Section IV-B2 is also replicated on the COSMOS FD testbed. We consider two TDMA schedules; the first is similar to the one used in Section IV-B2, henceforth referred to as the “round-robin opportunistic” (RRO) schedule, and the second additional schedule is IUI-free (IUI-F).

The RRO schedule works by opportunistically scheduling two links whenever possible in a time slot in a round-robin manner. For example, if the current time slot is scheduled for a DL to user 1, then there are two scheduling choices for the second link based on user 1’s duplexing mode: (i) if user 1 is FD, user 1 can perform a UL to the BS in the same slot, and (ii) if user 1 is in HD mode, a different user can perform a UL to the BS. The RRO schedule will always prefer choice (i) over choice (ii) as choice (ii) will create an IUI situation analogous to that evaluated in Section V-B1.

The IUI-F schedule is essentially an enhancement on top of the RRO schedule that will not permit any time slots where choice (ii) is taken. In the IUI-F schedule, two links are

scheduled in the same time slot only if the scheduled UL or DL is between the BS and an FD user. Intuitively, the IUI-F schedule should maintain a higher throughput fairness between the users, at the expense of FD throughput gain, as some time slots will have only an HD link.

Algorithm 1 summarizes the operation of the scheduler used in this network. This algorithm presumes the presence of a control channel through which a user may request a UL slot. As mentioned in Section IV, implementing a MAC capable of such control on an SDR is a significant challenge outside of the scope of this paper. We therefore evaluate each schedule under the specific case where the DL packet queues at the BS and UL packet queue at each user are saturated, meaning that a deterministic schedule may be precomputed and implemented by each user and the BS independently.

In order to generate a range of SINRs at each radio, we sweep the transmit power of the BS across a 10 dB range, such that in HD mode, all packets may be decoded from the BS by the users. The transmit power of the users is fixed, and also chosen such that for any MCS, all packets may be received by the BS in HD mode. We consider a range of MCS between BPSK 1/2 and 64-QAM 3/4 available in the 802.11-like PHY layer [67]; all radios will use the same MCS in a given experimental run.

The results of this experiment are summarized in Fig. 15. Fig. 15(a) shows the CDF of FD throughput gain for the different schedules, and Fig. 15(b) shows the JFI. We make two observations; (i) the IUI-F schedule with zero FD users is equivalent to the HD TDMA schedule, and (ii) the IUI-F schedule with two FD users is equivalent to the RRO schedule. These two cases are therefore omitted from the figures. Fig. 15(a) shows how zero FD users can impair throughput for the RRO schedule in 30% of trials; this is likely due to the additional SNR requirement in FD mode [1], [19]. This effect is still observed when one or two users are in FD mode, but with reduced severity. The three RRO schedules

Algorithm 1 RRO and IUI-F Scheduler

```

procedure SCHEDULEDL( $U, N, C$ )
    while  $C < N$  do
        if packet in DL queue for  $U$  then return  $U$ 
        else
             $U \leftarrow$  next user
             $C \leftarrow C + 1$ 
        return  $\emptyset$ 
procedure SCHEDULEUL( $U, N, C$ )
    while  $C < N$  do
        if  $U$  has requested UL then return  $U$ 
        else
             $U \leftarrow$  next user
             $C \leftarrow C + 1$ 
        return  $\emptyset$ 
 $N \leftarrow$  # of users
 $U \leftarrow$  ▷ Initialize to first user
while True do
     $DL \leftarrow$  SCHEDULEDL( $U, N$ )
    if  $U$  is in FD mode then
         $UL \leftarrow$  SCHEDULEUL( $U, N, 0$ )
    else
         $UL \leftarrow$  SCHEDULEUL(next user,  $N, 1$ )
    if  $DL \neq UL$  and IUI-F then SCHEDULESLOT( $DL$ )
    else SCHEDULESLOT( $DL, UL$ )
     $U \leftarrow$  next user
    
```

provide median throughput gains of $1.43\times$, $1.6\times$, and $1.85\times$ when zero, one, or two users are in FD mode, respectively. By comparison, the IUI-F schedule for one FD user has a median throughput gain of $1.43\times$.

Fig. 15(b) gives the median JFI for the RRO schedule as 0.79, 0.74, and 0.89 for zero, one, or two FD users, respectively. The median JFI is lower when there is one FD user present on account of the ability to schedule the FD user whenever the HD user is scheduled, leading to the FD user receiving a higher throughput compared to the HD user. This is also the reason the one FD user experiment never reaches a fairness value of 1. The IUI-F schedule for one FD user has a median fairness of 0.86, showing how the fairness can be improved at the cost of FD throughput gain.

The results from this experiment suggest that in a scenario with high IUI between users, an RRO schedule can provide a $1.6\times$ median FD throughput gain, but with the potential of low fairness between users. If an IUI-F schedule is used, fairness can be improved, though the median and maximum FD throughput gains will be lower. We also note that 20% of experiments with one FD user on the RRO schedule experience a FD throughput *loss*, compared to only 8% of IUI-free schedule experiments. The IUI-free schedule therefore leads to a distribution of FD throughput gains with lower mean, but also lower variance.

3) *Heterogeneous 4-Node Networks*: A similar experiment, also considering the RRO and IUI-F schedules, was conducted with an additional FD-capable user. The corresponding results are presented in Fig. 16. In Fig. 16(a), the RRO schedule

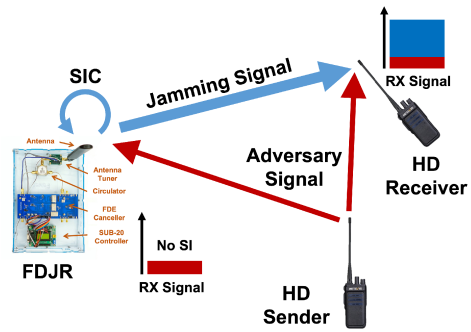


Fig. 17: FD jammer-receiver (FDJR) operation principle

provides FD throughput gains of $1.11\times/1.28\times/1.52\times/1.93\times$ when used with zero/one/two/three FD users.

The IUI-F schedules provide an FD throughput gain of $1.32\times$ and $1.66\times$ when one or two users are in FD mode, respectively. Unlike the 3-node experiment in Section V-B2, the median FD throughput gain for the IUI-F schedule is actually higher than for the RRO schedule, although the maximum achievable FD throughput gain is naturally lower for the IUI-F schedule. The results show that when two users are FD capable, only 20% of RRO experiments had a larger FD throughput gain compared to the best IUI-F experiment.

Fig. 16(b) shows that the fairness value for the one and two FD user cases improves when the IUI-F schedule is used. For one FD user, the median fairness improves from 0.43 to 0.73, and for two FD users it improves from 0.56 to 0.74. Therefore, the 4-node experiment is in a situation where the IUI-F schedule provides an advantage for both median throughput and median fairness. This result suggests that when developing a MAC protocol for FD, the IUI can play a significant hand in the performance, and scheduling both HD and FD users in the same time slot of a TDMA system may lead to impaired performance should the IUI between the HD and FD users be high.

We note that in both the 3- and 4-node networks, the transmit power of each user is kept the same regardless of whether it is in HD or FD mode, and we are sweeping the BS transmit power in the same manner for any number of FD users. Therefore, the median throughput gain increases simply by switching more users from HD to FD mode. In other words, the median energy per bit remains the same while the median throughput gain is improving, assuming that the additional power consumption for the FDE-based RF canceller (~ 40 mW for the RFIC [11] and ~ 20 mW for the PCB canceller) is negligible compared to the power consumption of the rest of the FD radio.

The heterogeneous 4-node experiment is provided on the `flexicon-cosmos.ndz` COSMOS server image, and directions on how to run this experiment and reproduce the results in this section are provided in the testbed tutorial [17].

C. Full-Duplex Jammer-Receiver

In addition to providing throughput improvements at the link and network layers, FD wireless can enable new applications that are not possible with HD radios [68]. In this section, we present an experiment that we run on the COSMOS FD testbed

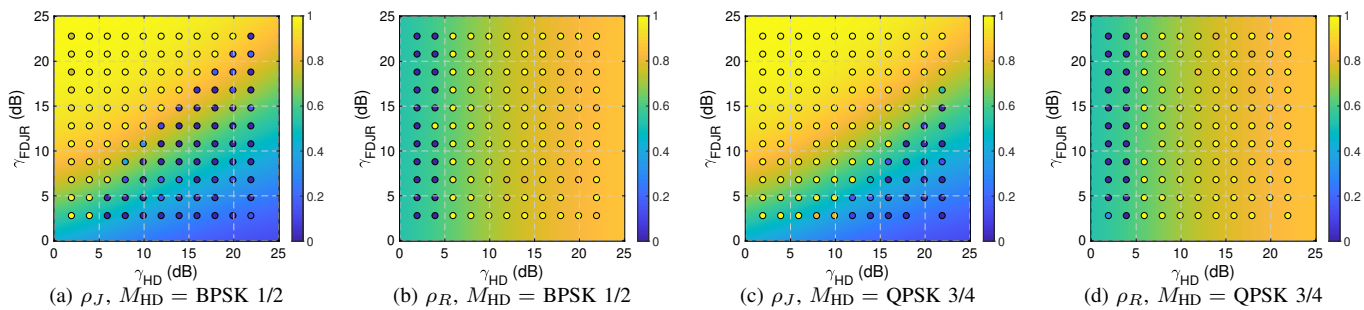


Fig. 18: Experimental results for ρ_J and ρ_R for the FDJR experiment represented as filled circles overlaid on the modeled colored surface. Results for $M_{HD} = \text{BPSK } 1/2$ and $\text{QPSK } 3/4$ shown.

to evaluate the performance of one such novel device: an FD jammer-receiver (FDJR). Recall that the primary challenge associated with FD is the need for SIC; with this in mind, the SIC capabilities of the FDJR allow it not only to jam a targeted adversarial signal, but also receive it. This principle of operation is outlined in Fig. 17.

1) *Experiment Design*: We use three out of the four COSMOS FD radios in this experiment. One FD radio will be configured as the FDJR, with the PCB canceller configured so as to provide sufficient RF and digital SIC to cancel the jamming signal to the noise floor. The other two radios will be configured in HD mode, replicating a link between two “standard” wireless devices. Of these two HD radios, one will transmit 802.11-like OFDM data packets while the other receives and decodes them.

The FDJR operates by first synchronizing its transmitter to the HD sender’s packet transmission schedule. For this experiment, we assume that the HD sender will transmit packets in an orderly manner. Once the FDJR is synchronized, it begins transmitting a jamming signal, in this case a simultaneous BPSK 1/2 OFDM data packet. To evaluate the performance of the FDJR, we define three experimental variables: (i) the FDJR transmit power setting P_{FDJR} , (ii) the HD sender transmit power setting P_{HD} , and (iii) the HD sender MCS M_{HD} .

The use of an OFDM data packet as a jamming signal achieves two goals. First, it allows for evaluation of the FDJR operating at moderate bandwidths, as opposed to recent work on jammer-receivers making use of a swept tone with small instantaneous bandwidth [69]. This is important as OFDM transmission schemes may be highly resilient to narrowband signals causing interference on only a small subset of subcarriers. Second, it would allow the FDJR to also operate as a data communications device, in theory being capable of fully overriding the signal from the HD sender should the SINR be sufficient at the HD receiver.

2) *Analytical Model of FDJR Performance*: We define two metrics to describe the performance of the FJDR: (i) the *jamming ratio* ρ_J , defined as the ratio of packets sent by the HD sender that were not decodable by the HD receiver, and (ii) the *reception ratio* ρ_R , defined as the ratio of the same packets that were successfully decoded by the FDJR. Given these metrics, if both ρ_J and ρ_R are close to 1, it means the FDJR is achieving close to maximum performance.

For the purpose of modeling the FDJR performance, ρ_J and ρ_R can be equivalently considered as received bitrates for the jamming target and FDJR respectively, R_T and R_{FDJR} , with

ideal FDJR performance achieved when $R_T \approx 0$ and R_{FDJR} maximized. The bitrates are themselves defined in terms of the SNR, SINR, and XINR values, as in Section IV-B. Altogether, ρ_J may be modeled as:

$$\rho_J \approx 1 - \frac{R_T}{R_{HD}} = 1 - \frac{1}{R_{HD}} \left(B \log_2 \left(1 + \frac{\gamma_{HD}}{1 + \gamma_{FDJR}} \right) \right), \quad (11)$$

where B is the transmit bandwidth, γ_{HD} is the SNR for the signal received at the HD receiver from the HD sender, and γ_{FDJR} is the SNR for the signal received at the HDR receiver from the FDJR. R_{HD} is the maximum rate that could be achieved by the HD sender: $R_{HD} = B \log_2(1 + \gamma_{HD})$. Similarly, ρ_R may be modeled as:

$$\rho_R \approx \frac{R_{FDJR}}{R_{HD}} = \frac{1}{R_{HD}} \left(B \log_2 \left(1 + \frac{\gamma_{HD}}{1 + \gamma_{Self}} \right) \right), \quad (12)$$

where γ_{Self} is the residual FDJR SI after RF and digital SIC.

3) *Results*: The experimental results are summarized in Fig. 18. Here, the filled circles represent ρ_J and ρ_R calculated simply as the ratio of packets jammed and received by the FDJR, respectively, while the colored surfaces are the analytical models in Equations 11 and 12. Instead of P_{FDJR} and P_{HD} , we use the experimentally measured γ_{FDJR} and γ_{HD} . We set $\gamma_{Self} = 1$ dB for the ρ_R model, which is the typical residual SI when the transmitter operates within its linear region.

These results show that ρ_J and ρ_R depend on all three of the experimental variables. For example, consider the region with $\gamma_{FDJR} > 20$ dB and $\gamma_{HD} < 5$ dB in Figs. 18(a) and 18(b). In this region, the jamming ratio ρ_J is high but the reception ratio ρ_R is low; this is because the HD sender’s transmit power is insufficient for the FDJR to be able to decode the packets. Should the HD sender increase its transmit power, thus increasing γ_{HD} , ρ_J will remain high at the same time as ρ_R becomes high. The value of M_{HD} considerably impacts ρ_J as seen in Figs. 18(a) and 18(c); essentially, the higher MCS used by the HD sender, the less robust it is against the FDJR’s jamming signal. ρ_R , on the other hand, is unaffected.

The experimental models generally follow the trend predicted by the Shannon capacity models in Equations 11 and 12, with the key difference being the sharp dropoff in ρ_J and ρ_R that is not reflected in the model. This is primarily an artifact of the packetized data transmission scheme used in this experiment, which does not allow for partial recovery of packets that were incorrectly received. Should we be able to investigate the jamming and reception ratios at the bit level rather than the packet level, we would expect a closer match

to the model results. Alternatively, the use of an empirical model based on the packet reception ratios reported in prior work [1], [19] would show a better fit to the experimental data.

D. Dataset

We provide a dataset which consists of transmitted data packet waveforms and the corresponding residual SI after RF SIC on the COSMOS dataset repository [70]. The data packets are QPSK 3/4-modulated 802.11a packets, similar to those used in the experiments presented in Sections IV and V². The received power after RF SIC is between -40 to -50 dBm and the USRP noise floor is -80 dBm at 20 MHz bandwidth.

The data files are in binary format representing the complex-valued baseband signals sent and received by the USRP SDR. In other words, this data is tapped after the ADC and before the DAC on the RX and TX paths within the NI USRP block in Fig. 3. Every experimental run is represented by two data files, one each for the TX and RX, and every run contains at least 100 data packets. There are a total of twelve experimental runs across three bandwidths (5 MHz, 10 MHz, and 20 MHz) using each of the four COSMOS FD radios. A document describing the experimental parameters in further detail is provided alongside the baseband data files.

The objective of this dataset is to provide ready-to-use baseband data traces to support the development of FD-related DSP algorithms, for example those for digital SIC. We provide an example MATLAB script that implements the same linear digital SIC algorithm as used in the custom GNU Radio block for the testbed experiments in Section V-B. Furthermore, the GNU radio flowgraph used for data collection is provided on the `flexicon-cosmos.ndz` COSMOS server image. Further data files may be recorded from this flowgraph.

VI. CONCLUSION

In this paper, we presented the system-level design, implementation, and evaluation of FDE-based FD radios. Experimentation conducted on two testbeds shows that such radios, given an appropriate configuration optimization algorithm or a sufficiently stable physical environment, will achieve rate gains and overall performance that matches closely to analytical results. This shows potential in the use of FDE-based RF cancellers to enable FD wireless in small-form-factor devices. Future research directions for FDE-based FD radios could include (i) better design and implementation of FDE-based canceller to support higher TX power handling and RF SIC bandwidth and (ii) extension of the FDE technique to multi-antenna systems.

The COSMOS FD testbed is openly accessible to researchers, with a tutorial describing its use. We anticipate that the COSMOS FD testbed will be used to explore several directions of future research, including the development and experimental evaluation of resource allocation and scheduling algorithms tailored for FDE-based FD radios. Additionally, the server-focused architecture of the COSMOS FD testbed effectively provides an edge compute environment, enabling

future investigation into how FD will make use of this key component of next-generation cellular networks [40]. Lastly, we will integrate newer, wider-band time-domain equalization-based RFIC cancellers in the COSMOS FD testbed [71], as well as perform FD experimentation with the outdoors COSMOS testbed infrastructure.

REFERENCES

- [1] T. Chen, M. B. Dastjerdi, J. Zhou, H. Krishnaswamy, and G. Zussman, "Wideband full-duplex wireless via frequency-domain equalization: Design and experimentation," in *Proc. ACM MobiCom'19*, 2019.
- [2] A. Sabharwal, P. Schniter, D. Guo, D. W. Bliss, S. Rangarajan, and R. Wichman, "In-band full-duplex wireless: Challenges and opportunities," *IEEE J. Sel. Areas Commun.*, vol. 32, no. 9, pp. 1637–1652, 2014.
- [3] H. Krishnaswamy and G. Zussman, "1 chip 2x the bandwidth," *IEEE Spectrum*, vol. 53, no. 7, pp. 38–54, 2016.
- [4] Z. Zhang, K. Long, A. V. Vasilakos, and L. Hanzo, "Full-Duplex Wireless Communications: Challenges, Solutions, and Future Research Directions," *Proc. IEEE*, vol. 104, no. 7, pp. 1369–1409, 2016.
- [5] J. I. Choi, M. Jain, K. Srinivasan, P. Levis, and S. Katti, "Achieving single channel, full duplex wireless communication," in *Proc. ACM MobiCom'10*, 2010.
- [6] M. Duarte and A. Sabharwal, "Full-duplex wireless communications using off-the-shelf radios: Feasibility and first results," in *Proc. Asilomar Conference on Signals, Systems, and Computers*, 2010.
- [7] D. Bharadia, E. McMillin, and S. Katti, "Full duplex radios," in *Proc. ACM SIGCOMM'13*, 2013.
- [8] D. Korpi, J. Tamminen, M. Turunen, T. Huusari, Y.-S. Choi, L. Anttila, S. Talwar, and M. Valkama, "Full-duplex mobile device: Pushing the limits," *IEEE Commun. Mag.*, vol. 54, no. 9, pp. 80–87, 2016.
- [9] T. Chen, S. Garikapati, A. Nagulu, A. Gaonkar, M. Kohli, I. Kadota, H. Krishnaswamy, and G. Zussman, "A Survey and Quantitative Evaluation of Integrated Circuit-Based Antenna Interfaces and Self-Interference Cancellers for Full-Duplex," *IEEE Open J. Commun. Soc.*, vol. 2, pp. 1753–1776, 2021.
- [10] D. Yang, H. Yuksel, and A. Molnar, "A wideband highly integrated and widely tunable transceiver for in-band full-duplex communication," *IEEE J. Solid-State Circuits*, vol. 50, no. 5, pp. 1189–1202, 2015.
- [11] J. Zhou, T.-H. Chuang, T. Dinc, and H. Krishnaswamy, "Integrated wideband self-interference cancellation in the RF domain for FDD and full-duplex wireless," *IEEE J. Solid-State Circuits*, vol. 50, no. 12, pp. 3015–3031, 2015.
- [12] J. Zhou, N. Reiskarimian, J. Diakonikolas, T. Dinc, T. Chen, G. Zussman, and H. Krishnaswamy, "Integrated full duplex radios," *IEEE Commun. Mag.*, vol. 55, no. 4, pp. 142–151, 2017.
- [13] X. Yi, J. Wang, C. Wang, K. E. Kolodziej, and R. Han, "A 3.4–4.6GHz In-Band Full-Duplex Front-End in CMOS Using a Bi-Directional Frequency Converter," in *Proc. IEEE RFIC'20*, 2020.
- [14] C. Wang, W. Li, F. Chen, W. Zuo, Y. Pu, and H. Xu, "A 0.5–4GHz Full-Duplex Receiver with Multi-Domain Self-Interference Cancellation Using Capacitor Stacking Based Second-Order Delay Cells in RF Canceller," in *Proc. IEEE RFIC'22*, 2022.
- [15] N. Reiskarimian, J. Zhou, T.-H. Chuang, and H. Krishnaswamy, "Analysis and design of two-port N-path bandpass filters with embedded phase shifting," *IEEE Trans. Circuits Syst. II, Exp. Briefs*, vol. 63, no. 8, pp. 728–732, 2016.
- [16] D. Raychaudhuri and I. Seskar et. al, "Challenge: COSMOS: A City-Scale Programmable Testbed for Experimentation with Advanced Wireless," in *Proc. ACM MobiCom*, 2020.
- [17] "Tutorial: Full-duplex wireless in the ORBIT and COSMOS testbeds," <https://wiki.cosmos-lab.org/wiki/Tutorials/Wireless/FullDuplex>, 2024.
- [18] "FlexiCoN: open-access full-duplex wireless," <https://flexicon.ee.columbia.edu>, 2024.
- [19] M. Kohli, T. Chen, M. Baraani Dastjerdi, J. Welles, I. Seskar, H. Krishnaswamy, and G. Zussman, "Open-access full-duplex wireless in the ORBIT and COSMOS testbeds," *Computer Networks*, vol. 199, 2021.
- [20] J. Zhou, A. Chakrabarti, P. Kinget, and H. Krishnaswamy, "Low-noise active cancellation of transmitter leakage and transmitter noise in broadband wireless receivers for FDD/co-existence," *IEEE J. Solid-State Circuits*, vol. 49, no. 12, pp. 1–17, 2014.
- [21] K. E. Kolodziej, B. T. Perry, and J. S. Herd, "In-Band Full-Duplex Technology: Techniques and Systems Survey," *IEEE Trans. Microw. Theory Techn.*, vol. 67, no. 7, pp. 3025–3041, 2019.

²Data for other MCS may be collected from the COSMOS FD testbed

- [22] D. Kim, H. Lee, and D. Hong, "A Survey of In-Band Full-Duplex Transmission: From the Perspective of PHY and MAC Layers," *IEEE Comm. Surveys & Tutorials*, vol. 17, no. 4, pp. 2017–2046, 2015.
- [23] B. Radunovic, D. Gunawardena, P. Key, A. Proutiere, N. Singh, V. Balan, and G. Dejean, "Rethinking indoor wireless mesh design: Low power, low frequency, full-duplex," in *Proc. IEEE WIMESH'10*, 2010.
- [24] 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 *Proc. ACM MobiCom'11*, 2011.
- [25] E. Aryafar, M. A. Khojastepour, K. Sundaresan, S. Rangarajan, and M. Chiang, "MIDU: Enabling MIMO full duplex," in *Proc. ACM MobiCom'12*, 2012.
- [26] M. K. Chung, M. S. Sim, J. Kim, D. K. Kim, and C.-B. Chae, "Prototyping real-time full duplex radios," *IEEE Commun. Mag.*, vol. 53, no. 9, pp. 56–63, 2015.
- [27] L. Zhang, M. Ma, and B. Jiao, "Design and Implementation of Adaptive Multi-Tap Analog Interference Canceller," *IEEE Trans. Wireless Commun.*, vol. 18, no. 3, pp. 1698–1706, 2019.
- [28] J. Tamminen, M. Turunen, D. Korpi, T. Huusari, Y.-S. Choi, S. Talwar, and M. Valkama, "Digitally-controlled RF self-interference canceller for full-duplex radios," in *Proc. EUSIPCO'16*, 2016.
- [29] M. Duarte, A. Sabharwal, V. Aggarwal, R. Jana, K. Ramakrishnan, C. W. Rice, and N. Shankaranarayanan, "Design and characterization of a full-duplex multi-antenna system for WiFi networks," *IEEE Trans Veh. Technol.*, vol. 63, no. 3, pp. 1160–1177, 2013.
- [30] D. Bharadia and S. Katti, "Full duplex MIMO radios," in *Proc. USENIX NSDI'14*, 2014.
- [31] B. Chen, V. Yenamandra, and K. Srinivasan, "FlexRadio: Fully flexible radios and networks," in *Proc. USENIX NSDI'15*, 2015.
- [32] M. Chung, M. S. Sim, D. Kim, and C.-B. Chae, "Compact full duplex MIMO radios in D2D underlaid cellular networks: From system design to prototype results," *IEEE Access*, vol. 5, pp. 16 601–16 617, 2017.
- [33] D. Bharadia and S. Katti, "Fastforward: Fast and constructive full duplex relays," *ACM SIGCOMM CCR*, vol. 44, no. 4, pp. 199–210, 2015.
- [34] L. Chen, F. Wu, J. Xu, K. Srinivasan, and N. Shroff, "BiPass: Enabling end-to-end full duplex," in *Proc. ACM MobiCom'17*, 2017.
- [35] E. Everett, C. Shepard, L. Zhong, and A. Sabharwal, "Softnull: Many-antenna full-duplex wireless via digital beamforming," *IEEE Trans. Wireless Commun.*, vol. 15, no. 12, pp. 8077–8092, 2016.
- [36] E. Aryafar and A. Keshavarz-Haddad, "PAFD: Phased array full-duplex," in *Proc. IEEE INFOCOM'18*, 2018.
- [37] V. Singh, S. Mondal, A. Gadre, M. Srivastava, J. Paramesh, and S. Kumar, "Millimeter-Wave Full Duplex Radios," in *Proc. ACM MobiCom'20*, 2020.
- [38] X. Chen, M. I. W. Khan, X. Yi, X. Li, W. Chen, J. Zhu, Y. Yang, K. E. Kolodziej, N. M. Monroe, and R. Han, "A 140GHz Transceiver with Integrated Antenna, Inherent-Low-Loss Duplexing and Adaptive Self-Interference Cancellation for FMCW Monostatic Radar," in *Proc. IEEE ISSCC'22*, 2022.
- [39] B. Yu and C. Qian et al., "Realizing High Power Full Duplex in Millimeter Wave System: Design, Prototype and Results," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 9, pp. 2893–2906, 2023.
- [40] B. Smida, A. Sabharwal, G. Fodor, G. C. Alexandropoulos, H. A. Suraweera, and C.-B. Chae, "Full-Duplex Wireless for 6G: Progress Brings New Opportunities and Challenges," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 9, pp. 2729–2750, 2023.
- [41] H. Harada, K. Mizutani, T. Matsumura, T. Kato, and K. Shioiri, "Development of Full-Duplex Cellular System for Beyond 5G and 6G Systems," in *Proc. IEEE PIMRC'22*, 2022.
- [42] E. Ahmed, A. M. Eltawil, and A. Sabharwal, "Rate gain region and design tradeoffs for full-duplex wireless communications," *IEEE Trans. Wireless Commun.*, vol. 12, no. 7, pp. 3556–3565, 2013.
- [43] J. Bai and A. Sabharwal, "Distributed full-duplex via wireless side-channels: Bounds and protocols," *IEEE Trans. Wireless Commun.*, vol. 12, no. 8, pp. 4162–4173, 2013.
- [44] S. Goyal, P. Liu, S. S. Panwar, R. A. Difazio, R. Yang, and E. Bala, "Full duplex cellular systems: Will doubling interference prevent doubling capacity?" *IEEE Commun. Mag.*, vol. 53, no. 5, pp. 121–127, 2015.
- [45] J. Marasevic, J. Zhou, H. Krishnaswamy, Y. Zhong, and G. Zussman, "Resource allocation and rate gains in practical full-duplex systems," *IEEE/ACM Trans. Netw.*, vol. 25, no. 1, pp. 292–305, 2017.
- [46] J. Diakonikolas and G. Zussman, "On the rate regions of single-channel and multi-channel full-duplex links," *IEEE/ACM Trans. Netw.*, vol. 26, no. 1, pp. 47–60, 2018.
- [47] I. Braga, F. Lima, T. Maciel, R. Antonioli, and F. Cavalcanti, "Rate maximization in full-duplex systems: Gains and complexity of employing adaptive power allocation in downlink and/or uplink," *Transactions on Emerging Telecommunications Technologies*, vol. 33(4):e4420, 2022.
- [48] G. Zheng, "Joint beamforming optimization and power control for full-duplex MIMO two-way relay channel," *IEEE Trans. Signal Process.*, vol. 63, no. 3, 2015.
- [49] Z. Qian, F. Wu, Z. Zheng, K. Srinivasan, and N. B. Shroff, "Concurrent channel probing and data transmission in full-duplex MIMO systems," in *Proc. ACM MobiHoc'17*, 2017.
- [50] V. Singh, A. Gadre, and S. Kumar, "Full Duplex Radios: Are We There Yet?" in *Proc. ACM HotNets'20*, 2020.
- [51] M. A. Islam, G. C. Alexandropoulos, and B. Smida, "Joint Analog and Digital Transceiver Design for Wideband Full Duplex MIMO Systems," *IEEE Trans. Wireless Commun.*, vol. 21, no. 11, pp. 9729–9743, 2022.
- [52] S.-Y. Chen, T.-F. Huang, K. C.-J. Lin, Y.-W. P. Hong, and A. Sabharwal, "Probabilistic medium access control for full-duplex networks with half-duplex clients," *IEEE Trans. Wireless Commun.*, vol. 16, no. 4, pp. 2627–2640, 2017.
- [53] R. Doost-Mohammady, M. Y. Naderi, and K. R. Chowdhury, "Performance analysis of CSMA/CA based medium access in full duplex wireless communications," *IEEE Trans. Mobile Comput.*, vol. 15, no. 6, pp. 1457–1470, 2015.
- [54] T. Chen, J. Diakonikolas, J. Ghaderi, and G. Zussman, "Hybrid scheduling in heterogeneous half- and full-duplex wireless networks," in *Proc. IEEE INFOCOM'18*, 2018.
- [55] X. Xie and X. Zhang, "Does full-duplex double the capacity of wireless networks?" in *Proc. IEEE INFOCOM'14*, 2014.
- [56] X. Qin, H. Zeng, X. Yuan, B. Jalaian, Y. T. Hou, W. Lou, and S. F. Midkiff, "Impact of full duplex scheduling on end-to-end throughput in multi-hop wireless networks," *IEEE Trans. Mobile Comput.*, vol. 16, no. 1, pp. 158–171, 2016.
- [57] S. Wang, V. Venkateswaran, and X. Zhang, "Fundamental analysis of full-duplex gains in wireless networks," *IEEE/ACM Trans. Netw.*, vol. 25, no. 3, pp. 1401–1416, 2017.
- [58] J. Y. Kim, O. Mashayekhi, H. Qu, M. Kazandjieva, and P. Levis, "Janus: A novel MAC protocol for full duplex radio," *CSTR*, vol. 2, no. 7, 2013.
- [59] Y. Alkhrijah, J. Camp, and D. Rajan, "Multi-Band Full Duplex MAC Protocol (MB-FDMAC)," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 9, pp. 2864–2878, 2023.
- [60] K. Hsu, K. Lin, and H. Wei, "Inter-client interference cancellation for full-duplex networks," in *Proc. IEEE INFOCOM'17*, 2017.
- [61] T. D. Vo-Huu, E.-O. Blass, and G. Noubir, "Counter-jamming using mixed mechanical and software interference cancellation," in *Proc. ACM WISEC'13*, 2013.
- [62] Y. Liu, Y. Shen, D. Guo, and M. Z. Win, "Network localization and synchronization using full-duplex radios," *IEEE Trans. Signal Process.*, vol. 66, no. 3, pp. 714–728, 2017.
- [63] A. Nagulu, S. Garikapati, M. Essawy, I. Kadota, T. Chen, A. Natarajan, G. Zussman, and H. Krishnaswamy, "Full-Duplex Receiver with Wideband Multi-Domain FIR Cancellation Based on Stacked-Capacitor, N-Path Switched-Capacitor Delay Lines Achieving >54dB SIC Across 80MHz BW and >15dBm TX Power-Handling," in *Proc. IEEE ISSCC'21*, 2021.
- [64] "Sandbox 2 (sb2.cosmos-lab.org)," https://wiki.cosmos-lab.org/wiki/Architecture/Domains/cosmos_sb2.
- [65] "The Columbia FlexICoN project COSMOS Github," https://github.com/Wimnet/flexicon_cosmos/, 2023.
- [66] "UHD (USRP Hardware Driver)," <https://www.ettus.com/sdr-software/uhd-usrp-hardware-driver/>.
- [67] B. Bloessl, M. Segata, C. Sommer, and F. Dressler, "Performance Assessment of IEEE 802.11p with an Open Source SDR-based Prototype," *IEEE Trans. Mobile Comput.*, vol. 17, no. 5, pp. 1162–1175, May 2018.
- [68] A. Nagulu, N. Reiskarimian, T. Chen, S. Garikapati, I. Kadota, T. Dinc, S. L. Garimella, M. Kohli, A. S. Levin, G. Zussman, and H. Krishnaswamy, "Doubling Down on Wireless Capacity: A Review of Integrated Circuits, Systems, and Networks for Full Duplex," *Proc. IEEE*, vol. 112, no. 5, pp. 405–432, 2024.
- [69] J. Marin, M. Bernhardt, and T. Riihonen, "Full-Duplex Constant-Envelope Jammer and Self-Interference Suppression by Highpass Filter: Experimental Validation for Wi-Fi Security," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 9, pp. 2937–2950, 2023.
- [70] "COSMOS Full-Duplex Dataset," https://wiki.cosmos-lab.org/wiki/Datasets/FD_Baseband_IQ, 2023.
- [71] A. S. Levin, I. Kadota, S. Garikapati, B. Zhang, A. Jolly, M. Kohli, M. Seok, H. Krishnaswamy, and G. Zussman, "Demo: Experimentation with Wideband Real-Time Adaptive Full-Duplex Radios," in *Proc. ACM SIGCOMM'23*, 2023.



Manav Kohli received the Sc.B. degree in electrical engineering from Brown University in 2018. He defended his Ph.D. thesis on the topic of full-duplex and millimeter-wave wireless in September 2024, covering research done at the Wireless and Mobile Networking (WiMNet) laboratory at Columbia University. Manav's work was supported by the NSF GRFP and NPSC GFSD fellowships, and he was the recipient of the 2024 Millman Teaching Award from the Columbia University Department of Electrical Engineering.



Harish Krishnaswamy received the B.Tech. degree in electrical engineering from Indian Institute of Technology Madras, Chennai, India, in 2001, and the M.S. and Ph.D. degrees in electrical engineering from the University of Southern California (USC), Los Angeles, CA, USA, in 2003 and 2009.


In 2009, he joined the Electrical Engineering Department, Columbia University, New York, NY, USA, where he is currently a Professor and the Director of the Columbia High-Speed and Millimeter-Wave IC Laboratory (CoSMIC). In 2017, he co-founded MixComm Inc., Chatham, NJ, USA, a venture-backed start-up, to commercialize CoSMIC Laboratory's advanced wireless research. MixComm was acquired in February 2022 by Sivers Semiconductors for \$155M, where he is currently the Managing Director of the Wireless Business Unit. His research interests include integrated devices, circuits, and systems for a variety of RF, millimeter-wave (mmWave), and sub-mmWave applications.

Dr. Krishnaswamy has been a member of the Technical Program Committee of several conferences, including the IEEE International Solid-State Circuits Conference and the IEEE Radio Frequency Integrated Circuits Symposium. He is a member of the DARPA Microelectronics Exploratory Council. He was a recipient of the IEEE International Solid-State Circuits Conference Lewis Winner Award for Outstanding Paper in 2007; the Best Thesis in Experimental Research Award from the USC Viterbi School of Engineering in 2009; the DARPA Young Faculty Award in 2011; the 2014 IBM Faculty Award; the Best Demo Award at the 2017 IEEE ISSCC; the Best Student Paper Awards at the 2015, 2018, and 2020 IEEE Radio Frequency Integrated Circuits Symposium and the 2020 IEEE International Microwave Symposium; the 2021 IEEE MTT-S Microwave Magazine Best Paper Award; and the 2019 IEEE MTT-S Outstanding Young Engineer Award. He has served as a Distinguished Lecturer for the IEEE Solid-State Circuits Society.



Mahmood Baraani Dastjerdi received the M.S. degree from the Sharif University of Technology in 2014 and the Ph.D. degree in electrical engineering from Columbia University in 2020. His career is focused on the theory, design, and experimental validation of analog/RF/millimeter-wave integrated circuits and systems. He was a recipient of the 2020 Columbia University Electrical Engineering Department Jury Award, the 2019 IEEE SSCS Predoctoral Achievement Award, the 2018 ADI ISSCC Outstanding Student Designer Award, and the Creative

Tech Award in Engineering at the NYC Media Lab'18.



Jin Zhou received the B.E. degree in electronics science and technology from Wuhan University, Wuhan, China, in 2008, the M.S. degree in microelectronics from Fudan University, Shanghai, China, in 2011, and the Ph.D. degree in electrical engineering from Columbia University, New York, NY, USA, in 2017. From 2011 to 2012, he was an RFIC Design Engineer with MediaTek, Singapore. From 2017 to 2022, he was an Assistant Professor with the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Cham-

paign, IL, USA. Since 2023, he has rejoined MediaTek, San Jose, CA, USA, developing RFIC products. Dr. Zhou has been serving on the Technical Program Committee of the IEEE RFIC Symposium since 2022. He also serves as a reviewer for several IEEE journals, including Journal of Solid-State Circuits.



Ivan Seskar is the Chief Technologist and Director/IT at WINLAB, Rutgers University responsible for experimental systems and prototyping projects. He is currently the program director for the COSMOS project responsible for the New York City NSF PAWR deployment. He has also been the co-PI and project manager for all three phases of the NSF-supported ORBIT mid-scale testbed project at WINLAB, successfully leading technology development and operations since the testbed was released as a community resource in 2005 and for which the

team received the 2008 NSF Alexander Schwarzkopf Prize for Technological Innovation. Ivan is a co-chair of the IEEE Future Networks Testbed Working Group, member of the IEEE Standardization Programs Development Board and the co-founder and CTO of Upside Wireless Inc.



Gil Zussman received the Ph.D. degree in electrical engineering from Technion, Haifa, Israel, in 2004. He was a Postdoctoral Associate at Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, from 2004 to 2007. He has been with Columbia University, New York, NY, USA, since 2007, where he is currently the Kenneth Brayer Professor of Electrical Engineering, Professor (affiliated faculty) of Computer Science, and the Electrical Engineering Department Chair. His research interests are in the area of networking, in

particular in the areas of wireless, mobile, and resilient networks. Dr. Zussman received the Fulbright Fellowship, two Marie Curie Fellowships, the DTRA Young Investigator Award, and the NSF CAREER Award. He was a co-recipient of seven paper awards, including the ACM SIGMETRICS'06 Best Paper Award, the 2011 IEEE Communications Society Award for Advances in Communication, and the ACM CoNEXT'16 Best Paper Award. He was an Associate Editor of IEEE/ACM Transactions on Networking, IEEE Transactions on Control of Network Systems, and IEEE Transactions on Wireless Communications, and the TPC Chair of IEEE INFOCOM 2023, ACM MobiHoc 2015, and IFIP Performance 2011. He is the Columbia PI of the NSF PAWR COSMOS Testbed.



Tingjun Chen received the B.Eng. degree in electronic engineering from Tsinghua University, Beijing, China, in 2014, and the Ph.D. degree in electrical engineering from Columbia University, New York, NY, USA, in 2020.

From 2020 to 2021, he was a Postdoctoral Associate at Yale University, New Haven, CT, USA. Since Fall 2021, he has been with Duke University, Durham, NC, USA where he is currently an Assistant Professor with the Department of Electrical and Computer Engineering and the Department of

Computer Science (secondary appointment). His research interests are in the area of networking and communications with a specific focus on next-generation wireless, optical, mobile networks, and IoT systems.

Dr. Chen received the IBM Academic Award, the Google Research Scholar Award, the Columbia Engineering Morton B. Friedman Memorial Prize for Excellence, the Columbia University Eli Jury Award, and the Facebook Fellowship. He was a co-recipient of several paper awards, including the ACM CoNEXT'16 Best Paper Award, ECOC'23 Best Paper Award, and Top-Scored Papers from OFC 2023 and OFC 2024. His Ph.D. thesis received the ACM SIGMOBILE Dissertation Award Runner-Up.