Computer Networks 199 (2021) 108420

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Open-access full-duplex wireless in the ORBIT and COSMOS testbeds ☆

Manav Kohli [a,*], Tingjun Chen [a], Mahmood Baraani Dastjerdi [a], Jackson Welles [a], Ivan Seskar [b], Harish Krishnaswamy [a], Gil Zussman [a]

[a] Electrical Engineering, Columbia University, United States of America
[b] WINLAB, Rutgers University, United States of America

ARTICLE INFO

Keywords:
Full duplex wireless
Open-access testbeds
Remote-access experimentation

ABSTRACT

In order to support experimentation with full-duplex (FD) wireless, we integrated two generations of FD radios in the open-access ORBIT and COSMOS testbeds. In the indoor ORBIT testbed, we integrated a customized 1st generation (Gen-1) narrowband FD radio. In the city-scale PAWR COSMOS testbed, we integrated four 2nd generation (Gen-2) wideband FD radios, as well as static and mobile Gen-1 FD radios.

Each integrated FD radio consists of an antenna, a customized RF self-interference (SI) canceller box, a USRP software-defined radio (SDR), and a remotely accessible compute node. The Gen-1/Gen-2 RF SI canceller box includes an RF canceller printed circuit board (PCB) which emulates a customized integrated circuit (IC) RF canceller implementation. The amplitude- and phase-based Gen-1 narrowband RF canceller achieves 45 dB RF SIC across 5 MHz bandwidth. The Gen-2 wideband canceller is based on the technique of frequency-domain equalization (FDE) and achieves 50 dB RF SI cancellation (SIC) across 20 MHz bandwidth. In this paper, we present the design of the two generations of FD radios, and their integration in various forms in the ORBIT and COSMOS testbeds. We then present example experiments that can be remotely run and modified by experimenters. Finally, we discuss future improvements and potential FD wireless experiments that can be supported by these open-access FD radios integrated in the COSMOS testbed.

1. Introduction

Due to its potential to double the data rate at the Physical (PHY) layer and to provide many other benefits at the higher layers of the networking stack, full-duplex (FD) wireless has drawn significant attention [2–5] as an enabler of next-generation wireless networks. One of the major challenges associated with enabling FD wireless is the extremely strong self-interference (SI) on top of the desired signal. At transmit power levels greater than 10 dBm, more than 90 dB of SI cancellation (SIC) across the antenna interface, the RF/analog, and digital domains is required to sufficiently cancel the SI to below the radio's noise floor.

Within the Columbia FlexICoN project [6], we focus on the design of and experimentation with FD radios and systems grounded in integrated circuit (IC) implementations, which are suitable for hand-held and form-factor-constrained devices [7]. In [8], we presented the 1st-generation (Gen-1) *narrowband* FD radio and an FD link, featuring 40 dB RF SIC across 5 MHz bandwidth. This Gen-1 RF SI canceller emulates its RFIC counterpart presented in [9], whose modeling and analysis were presented in [10]. In [11], we developed a 2nd-generation (Gen-2) *wideband* FD radio which leverages the technique of frequency-domain equalization (FDE) to achieve over 50 dB RF SIC across 20 MHz bandwidth (4× improved bandwidth over to the Gen-1 canceller), which emulates its RFIC counterpart presented in [12]. Both generations of RF canceller achieve these levels of RF SIC while consuming less than 300 mW of power.

In order to allow the broader community to experiment with FD wireless, we integrated the two generations of FD radios in the open-access ORBIT [13] and COSMOS [14,15] wireless testbeds. We provide several variants of programmable FD radios, including those suitable for static use and those suitable for mobile experimentation. Since interfacing an RFIC canceller with a software-defined radio (SDR) presents numerous technical challenges, we implemented the RF cancellers on printed circuit boards (PCBs) to facilitate cross-layered experiments with an SDR platform. Specifically, we integrated an improved version of the Gen-1 RF canceller with a USRP N210 SDR in the ORBIT

☆ Preliminary version Kohli et al. (2020) of this paper was presented at the 14th International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization (WiNTECH'20), Sept. 2020. [1].

* Corresponding author.
E-mail addresses: mpk2138@columbia.edu (M. Kohli), tingjun@ee.columbia.edu (T. Chen), mb4038@columbia.edu (M. Baraani Dastjerdi), jw3350@columbia.edu (J. Welles), seskar@winlab.rutgers.edu (I. Seskar), harish@ee.columbia.edu (H. Krishnaswamy), gil@ee.columbia.edu (G. Zussman).

https://doi.org/10.1016/j.comnet.2021.108420
Received 19 November 2020; Received in revised form 12 July 2021; Accepted 16 August 2021
Available online 28 August 2021

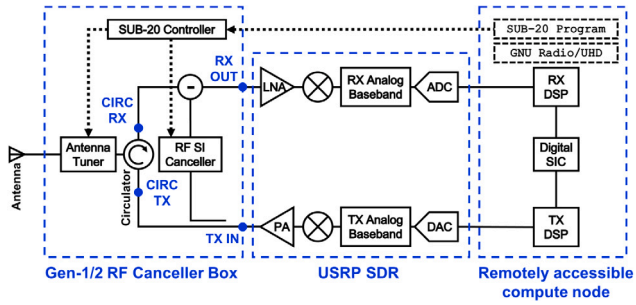

Fig. 1. Block diagram of the integrated FD radios with three main components: (i) a FlexICoN RF canceller box, (ii) a USRP software-defined radio (SDR), and (iii) a remotely accessible compute node. The host PC can range from a powerful server with GPU and FPGA resources to a palm-sized Intel NUC mini PC.

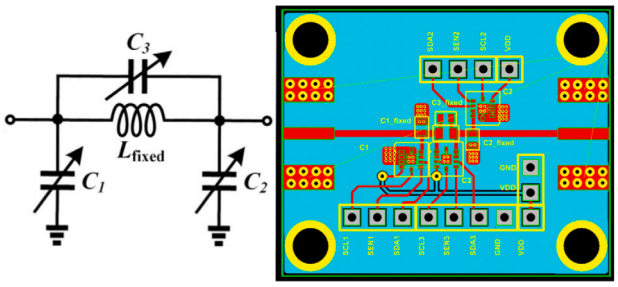

Fig. 2. Circuit diagram and PCB layout of the programmable antenna tuner as part of the Gen-1/Gen-2 RF canceller box.

testbed [13]. In the COSMOS testbed [14,15], we integrated four Gen-2 RF cancellers (an improved version of that presented in [11]) with USRP2 and USRP-2974 SDRs, one Gen-1 RF canceller with a USRP N210 SDR, and a mobile FD radio based on the Gen-1 RF canceller.

In this paper, we first present the overall design of the Gen-1 and Gen-2 FD radios. We then describe the various integrated FD radios in ORBIT and COSMOS. Finally, we present three remotely accessible example FD experiments that may be used by researchers: (i) real-time digital SIC, (ii) measurements of the link-level FD packet reception ratio (PRR), and (iii) measurements of the node-level throughput gains in a small FD ALOHA network. These experiments show that the Gen-1/Gen-2 FD radios integrated in ORBIT/COSMOS can achieve over 80 dB of overall SIC across 5/16.7–20 MHz bandwidth, respectively, up to 2× FD rate gain at the link layer, and up to 50% throughput gain at the network layer. Finally, we discuss further improvements to the testbeds, as well as potential future experimentation.

In summary, the main contributions of the paper are the following: (i) four additional FD radios were integrated into the COSMOS testbed, including two Gen-2 radios, one Gen-1 radio, and a mobile Gen-1 radio (see Section 6), (ii) the RF SIC performance of the integrated Gen-1 and Gen-2 radios was thoroughly characterized (see Section 4.2), (iii) the node-level SIC performance of the Gen-2 radios using the USRP 2974 in COSMOS was characterized (see Section 7.1), (iv) a simple FD ALOHA network experiment was demonstrated (see Section 7.3), and (v) our software stack was migrated to an open-source 802.11-like PHY (see Section 7.1), allowing the use of different modulation and coding schemes (see Section 7.2).

2. Related work

Extensive reviews of research in the area of FD wireless were presented in [2,3], including various implementations of FD systems, and analysis of the influence of FD wireless at the higher layers of the networking stack. Recent system-level design and evaluation of FD systems has included mmWave FD systems [16], the use of FD in wireless relay devices [17], and the deployment of long-range FD backscatter tags [18].

While [19–21] involved a pair of Tx and Rx antennas to achieve Tx/Rx isolation at the antenna interface, we focus on FD radio designs using a shared antenna interface such as a circulator, which are more appropriate for single-antenna systems [4,22]. In addition, existing circuit designs for RF/analog SIC often utilize a time-domain interpolation approach utilizing parallel delay lines with amplitude and phase control that are more suitable for discrete implementations on PCBs [4,23]. We present two RF cancellers which utilize techniques suitable for achieving RF SIC in compact IC implementations, while maintaining the robust integration with an SDR that a PCB allows.

The development of FD radios and systems in previous research and associated evaluation of PHY, medium access control (MAC) and higher layer algorithms and techniques has largely taken place in laboratory testbeds or in simulation [16,17,22,24–28]. The design and fabrication of FD hardware can itself be a daunting challenge, especially when networks of multiple FD devices are concerned. The integration of our FD hardware in ORBIT and COSMOS aims to alleviate this challenge. *To the best of our knowledge, this paper presents the design and implementation of the first open-access, remotely accessible FD radios, together with open-source example experiments and FD software, that can be used by the community for experimentation with FD wireless.*

3. Design of the FlexICoN FD radios

The overall design of all FD radios integrated in ORBIT and COSMOS, irrespective of generation or form factor, is shown in Fig. 1. The integrated FD radio design consists of three main components: an RF canceller box, a USRP SDR, and a remotely accessible compute node. Each Gen-1/Gen-2 canceller box has four components: a circulator, an antenna tuner, a SUB-20 controller, and an RF canceller PCB which differs between the two generations. Below, we briefly discuss these components.

Coaxial Circulator. An RF-CI RFCR3204 coaxial circulator [29] is used, with operating frequency between 860 and 960 MHz.

Programmable Antenna Tuner. To achieve better matching between the antenna and the circulator, we also designed and implemented a programmable antenna tuner at around 900 MHz frequency. Fig. 2 shows the circuit diagram and PCB implementation of the antenna tuner. In particular, a π-network with lossless inductor (L) and digitally tunable capacitors (C_i) is used for impedance transformation. We use a fixed chip inductor with inductance $L_{\text{fixed}} = 5.1$ nH and Peregrine Semiconductor PE64909 5-bit digitally tunable capacitors for C_i ($i = 1, 2, 3$). By programming capacitor C_i with code values `CAPi`, different antenna interface impedance matching can be achieved. The configuration ranges of the tunable capacitors are:

$$\texttt{CAPi} \in \{0, 1, \ldots, 15\}, \ \forall i = 1, 2, 3.$$

SUB-20 Controller. As Fig. 1 shows, a DIMAX SUB-20 multi-interface USB adaptor [30] connected to the remotely accessible compute node is used to program the tunable components on the RF canceller and antenna tuner through the serial peripheral interface (SPI). The SUB-20 SPI is configured to operate at the maximal SPI clock frequency of 8 MHz. The SUB-20 control for both the Gen-1 and Gen-2 canceller PCBs is handled through a customized GNU Radio out-of-tree (OOT) module (see Section 7.1 for details).

RF Canceller PCB. In order to meet the USRP Rx front-end linearity and the analog-to-digital converter (ADC) dynamic range requirements, sufficient SIC in the RF domain is needed before digital SIC is performed. Therefore, the RF canceller PCB needs to provide up to 30 dB RF SIC in addition to the 20–25 dB provided by the circulator and antenna tuner. The RF canceller taps a reference signal from the output of the power amplifier (PA) at the Tx side, and SIC is performed at the input of the low-noise amplifier (LNA) at the Rx side. The difference between the Gen-1 and Gen-2 canceller boxes is the circuit design of the RF canceller PCB, as discussed in Sections 4.1 and 4.2 .

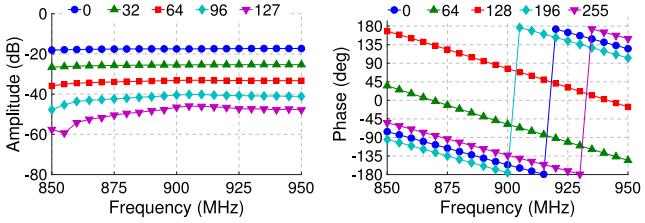

Fig. 3. Measured amplitude and phase of the Gen-1 RF canceller PCB with varying (a) attenuation values, and (b) phase shift values.

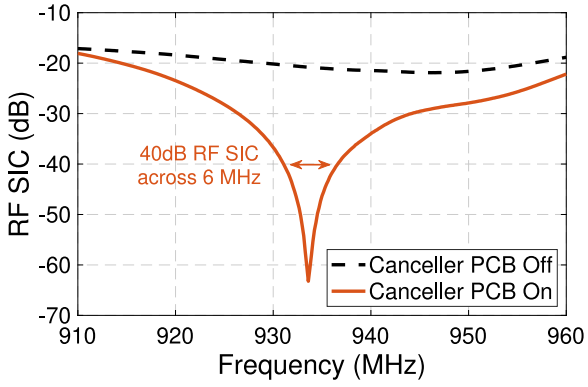

Fig. 4. Representative RF SIC achieved by the Gen-1 RF canceller box with and without turning on the RF canceller. The RF canceller box with the circulator and the RF canceller PCB provides 40 dB RF SIC across 5 MHz bandwidth.

4. Gen-1 and Gen-2 canceller PCBs

While the Gen-1 and Gen-2 FD radios share the same overall design described in Section 3, the core part of each radio, the canceller PCB, is different. In this section, we will describe the design of each generation of RF canceller PCB.

4.1. Gen-1 RF canceller PCB

The Gen-1 RF canceller PCB is a *narrowband* frequency-flat amplitude- and phase-based canceller, which is an improved version of that presented in [8]. The RF canceller is implemented using discrete components on a PCB and is optimized around 900 MHz operating frequency.[1] The reference signal is tapped through a 6 dB Mini-Circuits ADC-6-13+ directional coupler whose amplitude and phase are subsequently adjusted before SIC is performed at the Rx.

For amplitude adjustment, a SKY12343-364LF 7-bit digital attenuator is used, and for phase adjustment, a Mini-Circuits passive SPHSA-152+ phase-shifter is used. The phase shifter is controlled by a TI-DAC081S101 8-bit digital-to-analog converter (DAC). The attenuator and DAC have a 3 V supply voltage, and the phase shifter has a reference voltage of 12 V. The attenuator and DAC are programmed through the SUB-20 controller over SPI. The available parameter configuration ranges for the attenuator ATT (ATTenuation) and phase shifter PS (Phase Shift) are:

$$\texttt{ATT} \in \{0, 1, \ldots, 127\}, \ \texttt{PS} \in \{0, 1, \ldots, 255\}.$$

The attenuator and phase shift values can be input by the experimenter using the SUB-20 API via a graphical user interface (GUI) described in Section 7.1. Fig. 3 shows the measured amplitude and phase responses of the RF canceller with varying ATT values (under fixed PS = 0) and with varying PS values (under fixed ATT = 0). The Gen-1 RF canceller has an amplitude tuning range between −48 dB and −17 dB, and a phase tuning range spanning the entire 360° range. Fig. 4 shows an example of the measured Tx/Rx isolation (measured between Tx IN and Rx OUT ports shown in Fig. 1), where 40 dB RF SIC is achieved across 5 MHz bandwidth.

[1] In this implementation, we select operating frequencies around the 902–928 MHz ISM band. However, the implementation can be easily extended to other frequencies (e.g., the 2.45 GHz and 5.8 GHz ISM bands).

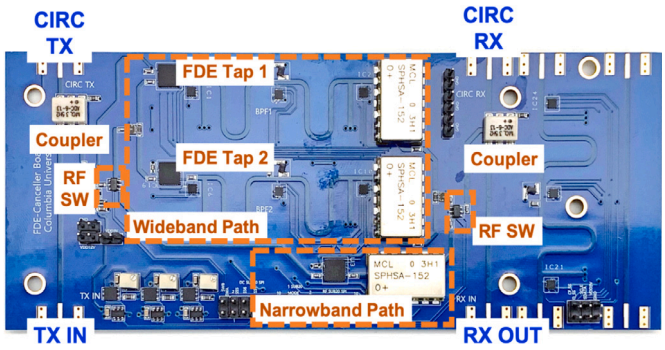

Fig. 5. The Gen-2 RF canceller PCB.

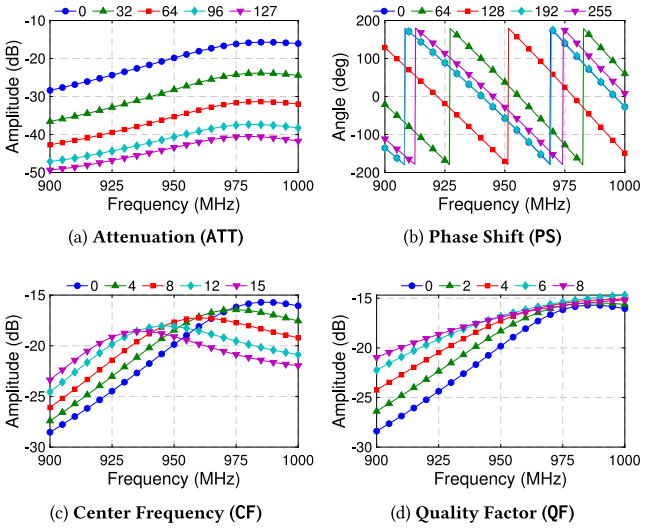

Fig. 6. Measured response of one FDE tap of the Gen-2 RF canceller with varying ATT, PS, CF, and QF values. The second FDE tap is turned off, and shares the same behavior.

4.2. Gen-2 RF canceller PCB

The Gen-2 RF canceller box includes an FDE-based RF canceller PCB which is an improved version of that described in [11].[2] This canceller PCB can achieve enhanced cancellation performance over a significantly wider bandwidth compared to the Gen-1 canceller PCB, and thus allows for experimentation with wider band signals. In this section, we present the design and implementation of the Gen-2 FD radio, including the Gen-2 RF canceller PCB.

The Gen-2 RF canceller illustrated in Fig. 5 is implemented using discrete components on a PCB and is optimized around a 900 MHz operating frequency.[3] The reference signal coupled from the Tx is first fed through an Analog Devices (ADI) HMC374 LNA before passing through

[2] The detailed design and evaluation of this canceller can be found in [11].

[3] As before, the design can be modified to other sub-6 GHz frequencies.

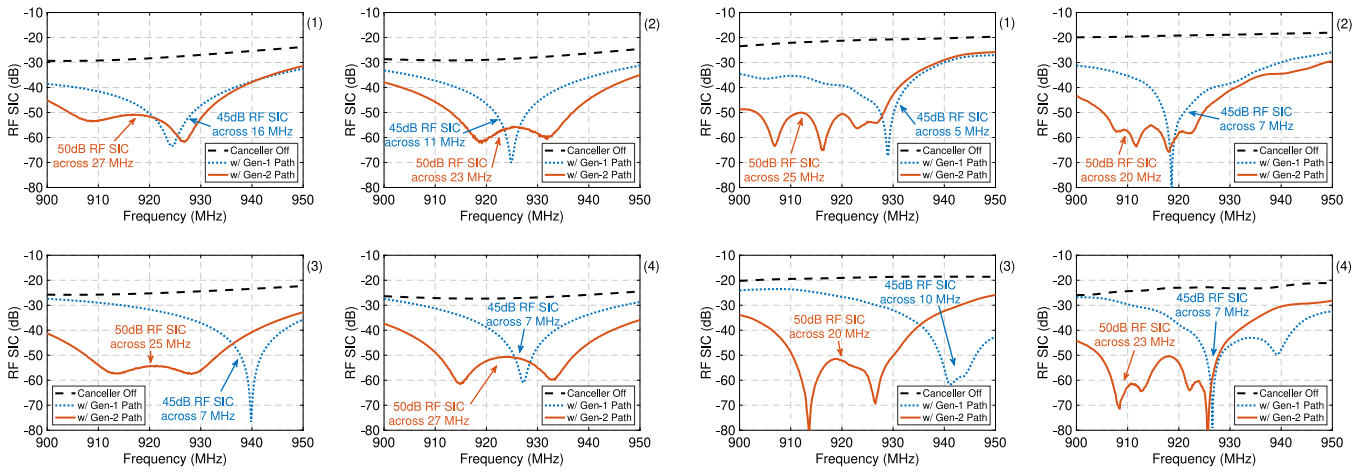

(a) **Antenna tuner terminated** (b) **Antenna tuner connected to antenna**

Fig. 7. Measured RF SIC around 925 MHz for four wideband Gen-2 RF canceller boxes when the antenna tuner is (a) terminated and (b) connected to an antenna. Shown is the RF SIC when (i) the canceller box is turned off, (ii) the Gen-1 path is used, and (iii) the Gen-2 path is used. All four canceller PCBs provide similar performance with similar configuration values for the tunable components around 925 MHz operating frequency.

Table 1
Summary of tunable components for the Gen-2 RF Canceller. $i \in \{0, 1\}$ represents the two FDE taps on the Gen-2 path.

Scope	Control	Name	Tuning range
PCB	RF Switch	`SW`	$\{0, 1\}$
Gen-2 Path	Attenuation	`ATTi`	$\{0, \ldots, 127\}$
	Phase	`PSi`	$\{0, \ldots, 255\}$
	Center Frequency	`CFi`	$\{0, \ldots, 15\}$
	Quality Factor	`QFi`	$\{0, \ldots, 31\}$
Gen-1 Path	Amplitude	`ATT2`	$\{0, \ldots, 127\}$
	Phase	`PS2`	$\{0, \ldots, 255\}$

an ADI HMC221B RF switch. This RF switch switches the reference signal between two paths: a Gen-1 narrowband path (identical to the Gen-1 RF canceller described in Section 4.1) and the wideband FDE path. In particular, the FDE path contains two parallel FDE taps, split and combined using Anaren PD0810J5050S2HF power dividers. Each FDE tap consists of a tunable bandpass filter (BPF) with amplitude and phase control.

Table 1 summarizes the tunable components on the Gen-2 canceller PCB. Each of the two paths contains an amplitude and phase control, identical to those on the Gen-1 canceller. Furthermore, each BPF is controlled by three tunable capacitors, one to control the center frequency of the BPF, and two identically controlled capacitors to control the quality factor of the BPF. For controlling the center frequency, we use the Peregrine Semiconductor PE64102 4-bit digitally tunable capacitor with a tuning range of 1.88–14.0 pF. For controlling the quality factor, we use Peregrine Semiconductor PE64909 5-bit digitally tunable capacitors with a tuning range of 0.6–2.35 pF. Together, these components provide a total of 2^{48} possible configurations for the Gen-2 path. The Gen-1 path's amplitude and phase control can be considered fully independent of the FDE path, as the RF switch provides up to 30 dB of isolation between paths. Similarly to the Gen-1 canceller PCB, the experimenter can control the values in Table 1 via the GUI described in Section 7.1.

Fig. 6 shows the effect of varying each parameter on the frequency response of one individual FDE tap. In each case, only one parameter is varied, while the others are kept at their lowest values. Fig. 7 shows the achievable RF SIC of the wideband FDE path, where >50 dB RF SIC is achieved across 20 MHz bandwidth. This is 5 dB higher RF SIC achieved over 2–4× the bandwidth when compared to the Gen-1 path, and means that higher bandwidth signals, such as orthogonal frequency-division multiplexing (OFDM) signals, may be used in experimentation.

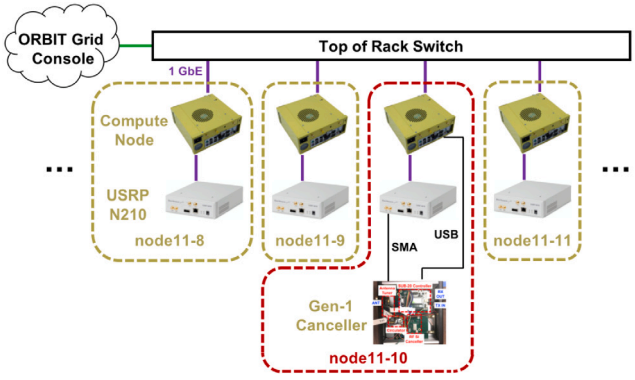

Fig. 8. Architecture and topology of the ORBIT main grid [13], where the Gen-1 canceller box is integrated with `node11-10`.

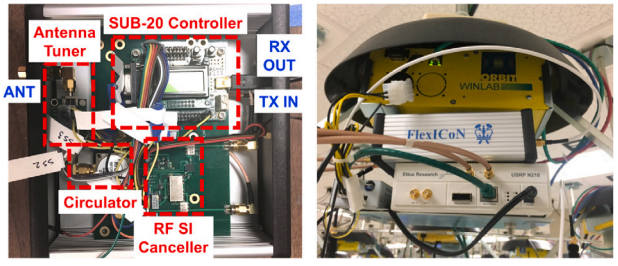

Fig. 9. Gen-1 narrowband FD radio integrated with a USRP N210 SDR in the ORBIT testbed.

5. Integration in the ORBIT testbed

Recall the goal of making the FD radios openly accessible to remote experimenters. An ORBIT node equipped with the Gen-1 RF canceller box is depicted in Fig. 9. We use `node11-10` in the ORBIT main `grid` with a USRP N210 SDR and remotely accessible compute node. The organization of the ORBIT nodes along with `node11-10` is shown in Fig. 8, and a diagram of the integrated FD radio is shown in Fig. 9. In particular, the RF canceller box Tx IN/Rx OUT ports are connected to the USRP Tx/Rx ports respectively, and the RF canceller box ANT port is connected to an Apex II multi-band antenna.

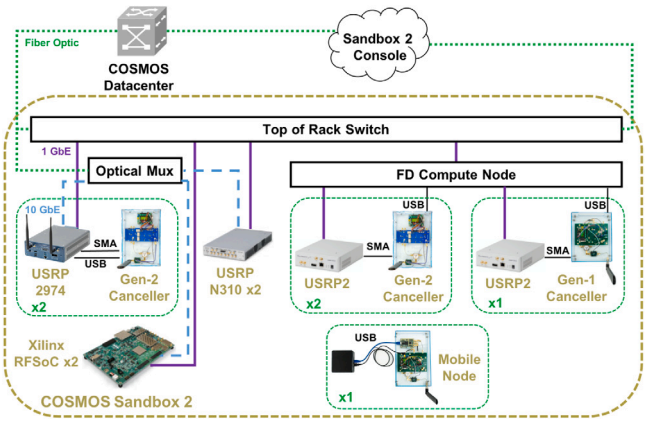

Fig. 10. COSMOS Sandbox 2 architecture, which includes (i) the remotely accessible console and integrated FD radios, connected to different SDRs and remotely accessible compute nodes, and (ii) other remotely accessible SDRs used for experimentation with various technologies.

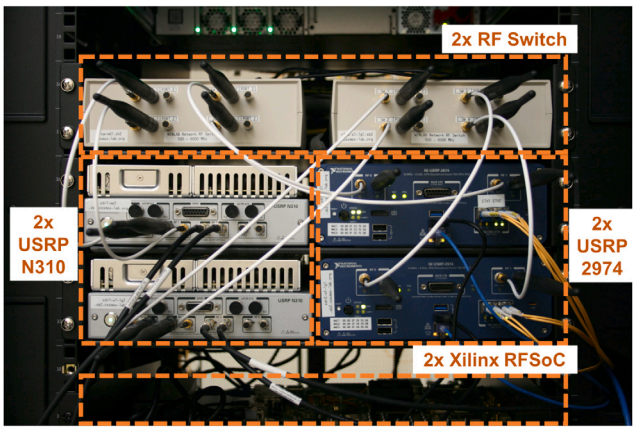

Fig. 11. Various SDRs in COSMOS Sandbox 2, including the USRP 2974s and USRP N310s, and Xilinx RFSoCs. The RF switches are omitted in Fig. 10 but are used to connect different hardware, including the Gen-2 canceller boxes, to the USRP-2974s and USRP N310s.

The Gen-1 RF canceller is mounted inside a metal box and placed below the yellow node, allowing for long-term reliability and a reproducible, well-controlled experimental environment. While using the Gen-1 FD radio, experimenters may also utilize other nodes in the testbed to create a wireless link where one radio operates in FD mode. The USRP has a receiver noise floor of −86 dBm at 10 MHz bandwidth.[4]

We developed a node image running Ubuntu (`flexicon-orbit-v4.ndz`), which contains the two example GNU Radio FD experiments described in Sections 7.1 and 7.2 as well as GNU Radio and UHD preinstalled [31,32]. While there is only one FD radio integrated in ORBIT, the large number of controllable half-duplex (HD) radios can be used to replicate many different experimentation scenarios on the node-level. For example, the node-level SIC performance can be investigated as described in Section 7.1. Moreover, the performance in the presence of many interferers can be studied.

6. Integration in the COSMOS testbed

In COSMOS Sandbox 2, located at Columbia University, we integrated both generations of RF canceller boxes. Along with different SDRs and remotely accessible compute nodes, these have been used to fabricate several variants of FD radio. The integration of this hardware is shown within the context of COSMOS Sandbox 2 in Fig. 10, and a picture of the hardware available in Sandbox 2 is shown in Fig. 11. Overall, the testbed includes the following: (i) one Gen-1 FD radio with a USRP2 SDR and a remotely accessible desktop compute node, (ii) four Gen-2 FD radios, two with USRP2 SDRs and the same remotely accessible desktop compute node, and two with USRP-2974s whose embedded PCs act as the remotely accessible compute nodes, and (iii) a mobile Gen-1 FD radio using a USRP B205mini-i SDR and Intel NUC portable PC. In this section, we will describe these three different forms of integrated FD radio. The presence of numerous FD radios allows for link and network-level FD experimentation. This capability is unavailable in the ORBIT testbed.

6.1. Gen-1 FD radio

As mentioned above, the Gen-1 RF canceller box was integrated in COSMOS. The Gen-1 RF canceller box is connected to a USRP2 SDR and a remotely accessible desktop compute node. This remotely accessible compute node is equipped with an Intel i7 4.0 GHz quad-core processor and 16 GB RAM, and runs Ubuntu, GNU Radio, and UHD. The USRP is connected to the remotely accessible compute node over Ethernet, and the SUB-20 controller is connected over USB. The USRP2 SDR supports up to 25 MHz bandwidth over the 1 Gbps Ethernet connection to the remotely accessible compute node. Together, this configuration provides real-time FD operation up to 5 MHz bandwidth, the maximum operating bandwidth of the Gen-1 RF canceller box. The Gen-1 RF canceller box and its integration in COSMOS is shown in Fig. 12(a). While the Gen-2 FD radios all contain a Gen-1 path on the RF canceller PCB, the Gen-1 FD radio provides experimenters with simpler hardware with only two tunable parameters (amplitude and phase) as described in Section 4.1. As a result, the Gen-1 FD radio is easier to control and can provide sufficient RF SIC in narrowband use cases.

6.2. Gen-2 FD radios

As illustrated in Fig. 10, two of the Gen-2 FD radios are connected to USRP2 SDRs, which are synchronized over a MIMO cable. allowing them to be used in a synchronized FD link. These two Gen-2 FD radios are shown in Fig. 12(b) with a 5 ft distance between antennas. The same remotely accessible compute node that controls the Gen-1 FD radio is also used to control these two Gen-2 FD radios. The higher bandwidth supported by the Gen-2 FD radio allows for running real-time FD experiments at 10 MHz bandwidth with a GUI. With the GUI disabled and when using the command line terminal, the supported bandwidth can be increased to the full 20 MHz supported by the Gen-2 RF canceller box.

Two more identical Gen-2 FD radios were also integrated in COSMOS, which are connected to two USRP-2974 SDRs, synchronized using an OctoClock-G CDA-2990. The USRP-2974 SDR contains a built-in PC, and thus serves both as the SDR and the remotely accessible compute node. This PC has an Intel i7 2.0 GHz quad-core processor, 16 GB RAM and runs Ubuntu, GNU Radio and UHD. The two FD radios connected to the USRP-2974s can also be used as an FD link, similar to the two Gen-2 FD radios connected to USRP2 SDRs.

The inclusion of four wideband FD radios in Sandbox 2 is an important step in the development and experimentation of network-level algorithms for FD and heterogeneous FD/HD networks. While two radios are sufficient for simple link-level experimentation, involving a network of more nodes facilitates investigation of the impacts of FD on the higher layers of the network stack (e.g. [33,34]) that is backed up by the use of real FD hardware. The four integrated Gen-2 FD radios

[4] This USRP receiver noise floor is limited by the existence of environmental interference at 900 MHz frequency. The USRP has a true noise floor of around −95 dBm at the same receiver gain setting, when not connected to an antenna.

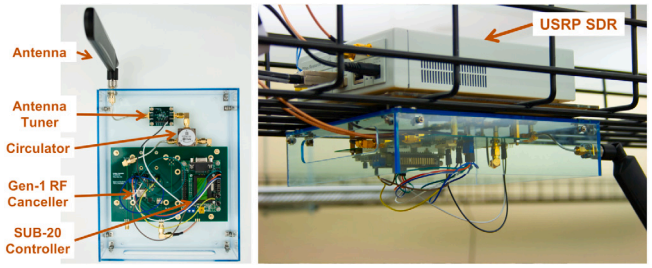

(a) **Gen-1 narrowband FD radio.**

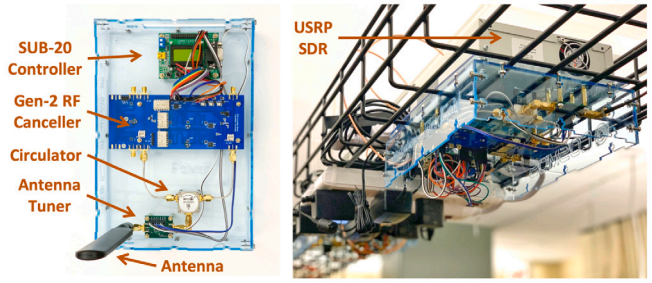

(b) **Gen-2 wideband FD radio.**

Fig. 12. The (a) Gen-1 and (b) Gen-2 FD radios integrated in the COSMOS testbed.

all support similar RF SIC performance around 925 MHz, shown in Fig. 7, meaning experiments involving multiple FD radios can be well controlled.

6.3. A prototype mobile full-duplex radio

The available Gen-1 and Gen-2 FD radios enable a wide variety of FD experiments, and are statically placed in order for remote experimenters to be able to conduct repeatable experiments. As experimentation with mobile user equipment (UE) operating in FD mode is of interest (e.g. [35,36]), we have fabricated a prototype mobile FD radio consisting of a Gen-1 RF canceller box similar to that in Fig. 12(a), a USRP B205mini-i SDR, and an Intel NUC 8i7BEH. The complete prototype mobile FD radio is shown in Fig. 13, and the RF SIC performance of the integrated Gen-1 RF canceller is shown in Fig. 4. The USRP B205mini-i is embedded within the RF box, leading to a portable FD UE that can be powered with portable batteries.

For remote-controlled mobility, we will utilize an X–Y table whose setpoint can be configured remotely through the COSMOS Sandbox 2 console (see [37] for such an X–Y table installed in COSMOS Sandbox 1). The X–Y table has as $1.3 \times 1.3\,m^2$ operating area and also can rotate the mobile node between −45° and +45°. Being in the smaller-scale Sandbox 2, this operating area will result in sufficient variation of the close-range environment to cause a significant effect on the RF SIC. The nodes can also be moved by the experimenters themselves or students recruited to help in experimentation, a similar process that is used for other portable nodes used in the COSMOS testbed [38].

The Intel NUC PC has a mobile-spec Intel i7 2.6 GHz quad-core processor and 16 GB RAM. This configuration supports real-time FD operation up to 5 MHz bandwidth, making it suitable for the Gen-1 RF canceller, as it supports 40 dB RF SIC across this bandwidth. Furthermore, the Gen-1 canceller PCB's two-component configuration described in Section 4.1 helps simplify mobile experiments.

7. Remote experimentation

Recall that a core feature of the integrated FD radios is that they can be remotely accessed by experimenters. The process for remotely accessing the FD radios is the same for both ORBIT and COSMOS: the experimenter logs into the testbed console from a local machine using Secure Shell (SSH). X11 forwarding is used to enable the GNU Radio GUI on local machines. The open-source code with detailed instructions can be found at [39,40].

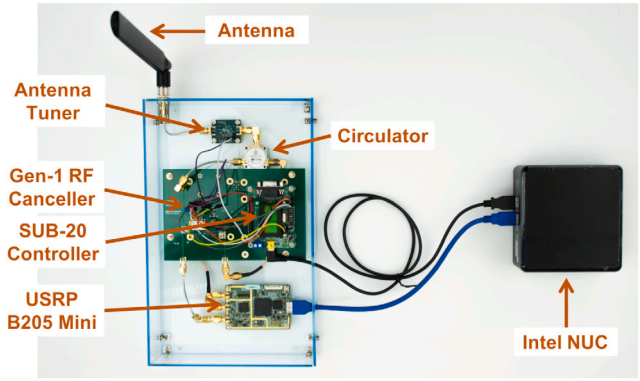

Fig. 13. Prototype mobile FD radio, consisting of a Gen-1 RF canceller box with a USRP B205mini-i and an Intel NUC PC.

The developed example experiments run in *real-time*, where the experimenter can observe results visualized without offline processing, and employ an orthogonal OFDM PHY layer with varying bandwidth and modulation and coding schemes (MCS). These example experiments, described below, can be used to benchmark several performance metrics of the Gen-1 and Gen-2 radios integrated in ORBIT and COSMOS, including node-level SIC and link-level PRR.

Real-time performance is achieved by the use of customized C++ out-of-tree (OOT) blocks. The implemented open-source OOT blocks are available at [39] and described in Section 7.1 below, and are also used in the experiments described in Sections 7.2 and 7.3 . Experimenters may use GNU Radio to develop their own FD experiments, either by leveraging the available OOT blocks or creating their own. For further improved performance, the experimenters can also develop experiments by directly invoking the UHD driver in C++ code.

7.1. Experiment 1: Node-level digital SIC

In this experiment, the experimenter can observe visualized real-time performance of a single FD radio while transmitting OFDM packets. The following data is visualized: (i) the time domain Rx signal after RF SIC and after digital SIC, (ii) the power spectrum of the Rx signal after RF SIC, (iii) the power spectrum of the Rx signal after both RF and digital SIC, and (iv) the digital SIC filter taps. The results were achieved on the Gen-1 FD radio in ORBIT, and two Gen-2 FD radios in COSMOS, one connected to a USRP2 and the other to a USRP 2974. This experiment makes use of many customized OOT blocks, which are described below.

SUB-20 Control. A core part of this experiment is the ability of the experimenter to manually configure the canceller PCB to achieve different RF SIC profiles. The experimenter has access to a GUI with controls that allow the full range of input described in Table 1. The numerical values set by these GUI controls are input to the SUB-20 control OOT block, which configures the canceller box using the SUB-20 application programming interface (API) [30] over SPI. For the Gen-1 and Gen-2 RF canceller boxes, a total of 7 and 23 bytes, respectively, are sent over SPI to program the canceller PCB and antenna tuner. With an 8 MHz SPI clock, one configuration of the Gen-1 and Gen-2 canceller PCBs takes 7 μs and 23 μs, respectively. Therefore, experimenters can observe different RF SIC profiles in real time.

In the stable indoor environment provided by COSMOS Sandbox 2 and the ORBIT grid, a given configuration can reliably provide

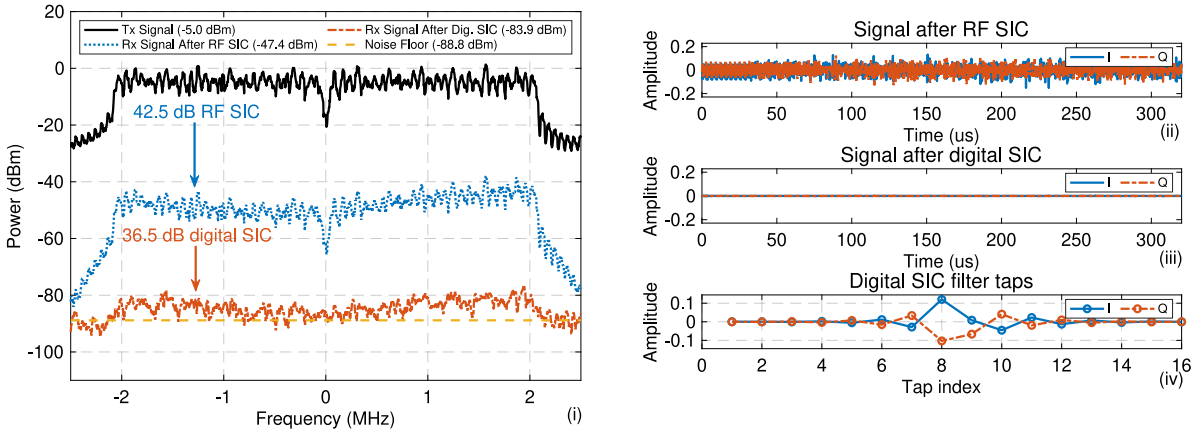

(a) **Gen-1 FD radio with USRP N210 at 916 MHz carrier frequency with 5 MHz bandwidth, −89 dBm noise floor, and −5 dBm transmit power.**

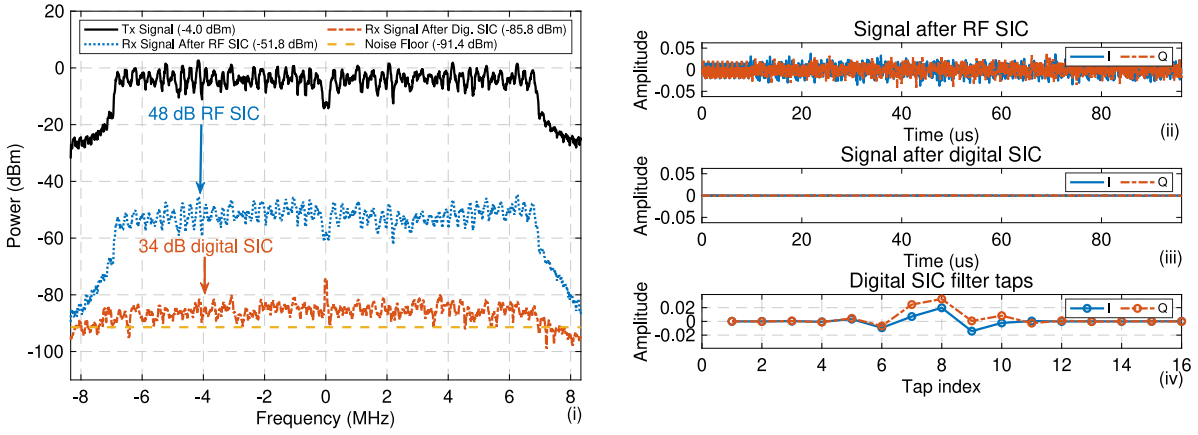

(b) **Gen-2 FD radio with USRP2 at 900 MHz carrier frequency with 16.7 MHz bandwidth, −91 dBm noise floor, and −4 dBm transmit power.**

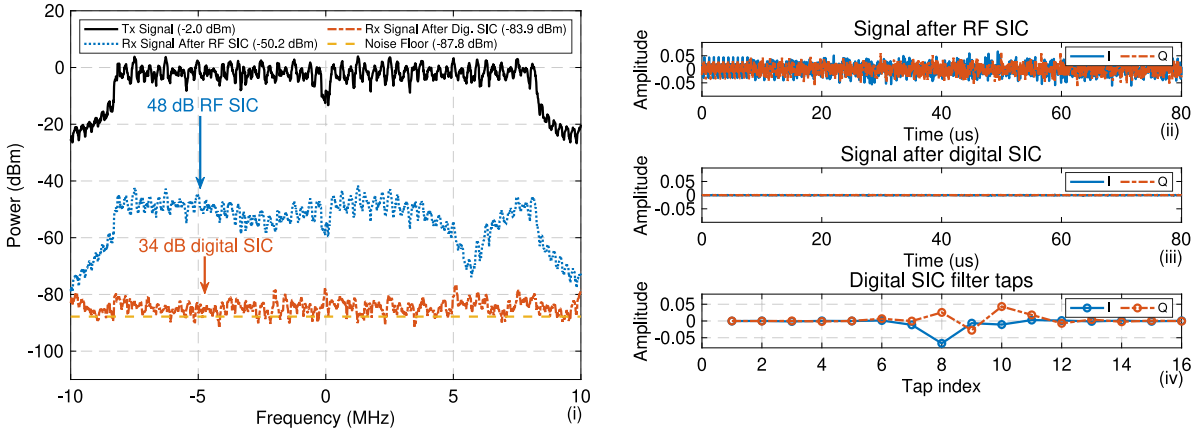

(c) **Gen-2 FD radio with USRP 2974 at 962 MHz carrier frequency with 20 MHz bandwidth, −88 dBm noise floor, and −2 dBm transmit power.**

Fig. 14. Node-level self-interference cancellation (SIC) performance using WiFi-like QPSK 3/4 MCS OFDM packets, each of length 1,600 samples, for (a) the ORBIT Gen-1 FD radio, (b) the COSMOS Gen-2 FD radios integrated with USRP2 SDRs, and (c) the COSMOS Gen-2 FD radios integrated USRP 2974 SDRs. Results shown are (i) power spectrum of the received signal after SIC in the RF and digital domains; (ii), (iii) time domain signals after RF and after digital SIC; (iv) digital SIC filter taps.

the same SIC performance across different experiment runs. A C++-based implementation of the adaptive canceller configuration method described in [11] is also considered as a subject of our future work.

WiFi-like OFDM Packets. The available example experiments make use of WiFi-like data packets for transmission. The data packets are generated a priori using an open-source GNU Radio 802.11 framework [41] and are saved in binary data files. The framework utilizes the four different modulation schemes used by 802.11 data packets: BPSK, QPSK, 16-QAM, and 64-QAM. The first three modulation schemes use 1/2 and 3/4 coding rates, and 64-QAM uses 2/3 and 3/4, giving a total number of 8 available MCS. The transmitted data packets can then be successfully decoded in real time using a receiver from the aforementioned 802.11 framework. Together, the use of this framework in this method allows for real-time transmission of realistic 802.11 data packets.

Packet Encapsulation. The implemented digital SIC algorithm requires additional pilot symbols prepended to every packet. We implemented a customized OOT block to prepend these symbols, and to add zero padding between packets.

Digital SIC. Digital SIC is performed on the packet-level using a least-squares algorithm to estimate the SI channel in the time domain with additional pilot OFDM symbols. The OOT block performs both the

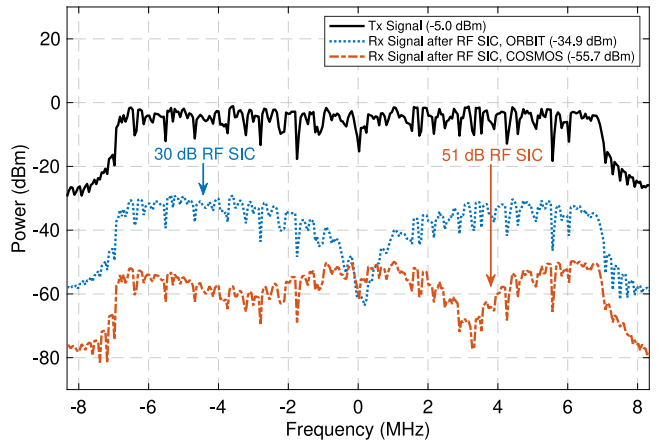

Fig. 15. RF SIC achieved over 16.7 MHz using WiFi-like OFDM packets for the Gen-1 FD radio integrated in ORBIT and Gen-2 FD radios integrated with USRP2 SDRs in COSMOS. The Gen-2 wideband FD radio can achieve a greater RF SIC over a wider bandwidth.

least-squares SI channel estimation and digital SIC with the following computation: $\mathbf{r}_{\text{Rx}} = \mathbf{y}_{\text{Rx}} - \mathbf{A}\hat{\mathbf{h}}$. Here, $\mathbf{r}_{\text{Rx}}$ is the residual SI after RF and digital SIC, $\mathbf{y}_{\text{Rx}}$ is the signal after RF SIC, $\mathbf{A}$ is the Toeplitz matrix constructed from the transmitted signal, and $\hat{\mathbf{h}}$ is the estimated SI channel (i.e., the digital SIC filter taps).

SIC Performance. Fig. 14 shows the node-level SIC performance of the Gen-1 narrowband FD radio in ORBIT and the Gen-2 wideband FD radio using the USRP2s and USRP 2974 in COSMOS. By configuring the RF cancellers appropriately, the Gen-1 FD radio can achieve 40–45 dB RF SIC across 3–5 MHz, and the Gen-2 FD radios can achieve 45–50 dB RF SIC across 16.7–20 MHz. In both cases, the RF SIC is followed by 34–37 dB of digital SIC, resulting in an overall SIC of over 80 dB. Given the low transmit power used in these experiments, 80 dB of overall SIC is sufficient for successful FD operation.

For higher transmit power levels, an overall SIC of closer to 90 dB is required for cancellation to the noise floor. The primary limitation of the Gen-1 and Gen-2 FD radio SIC performance is the nonlinear performance of the RF chain, most predominantly due the power amplifier of the USRP and the RF canceller itself. These nonlinearities manifest themselves as imperfect SIC at a level that is, in our case, 4–5 dB above the noise floor. When a higher transmit power is used, this excess above the noise floor increases further.

A comparison of the RF SIC performance for the integrated Gen-1 and Gen-2 radios is illustrated over a wider bandwidth in Fig. 15, which clearly shows the narrowband and wideband performances. Across the 16.6 MHz bandwidth shown in the figure, the Gen-1 FD radio can achieve 30 dB RF SIC, whereas the Gen-2 FD radio can achieve 50 dB RF SIC.

7.2. Experiment 2: Link-level FD packet reception ratio (PRR)

In this experiment, the experimenter can use the two Gen-2 FD radios connected to USRP2 SDRs to perform link-level experimentation. The link PRR as a function of the signal-to-noise ratio (SNR) is measured while the two Gen-2 FD radios operate in HD or FD mode. To achieve a successful FD link, the experimenter can configure each FD radio via the methods described in Section 7.1 to achieve the desired SIC across the RF and digital domains. The Tx power of the FD radios is swept during the experiment. For each Tx power value, the FD radios transmit 1,000 WiFi-like 100-byte data packets. The experiment is repeated for all 8 MCS described in Section 7.1.

We consider two performance metrics. The first is the *HD or FD link SNR*, which is measured as the ratio between the average Rx signal power level (across all 1,000 packets) and the Rx noise floor when the link operates in HD or FD mode. The second metric is the *HD or FD link PRR*, which is calculated for each FD radio as the percentage of the 1,000 transmitted packets that are successfully received and decoded. Fig. 16 plots the measured HD and FD link PRR as a function of the HD link SNR. The results show that with sufficient link SNR values, both Gen-2 FD radios in a link can achieve a PRR of 1 while operating in FD mode. This corresponds to a link-level FD rate gain of exactly 2×. With insufficient link SNR values, there is a roughly 2 dB degradation in achieved SNR when operating in FD mode, caused by imperfect SIC, which causes a residual SI that is 1–2 dB above the noise floor. This leads to a reduction in PRR by 30%–50%, leading to an FD rate gain of 1.0–1.4×.

In particular, with an HD link SNR of 20 dB, the PRR for 64-QAM 3/4 in HD mode is 0.45, and the PRR for 64-QAM 2/3 in FD mode is 0.35. The data rate supported by the 802.11a 64-QAM 3/4 packet transmissions used in our experiments is 54 Mbps, and 48 Mbps for 64-QAM 2/3. Therefore, using the empirical PRR values from the results in Fig. 16, we can compute an effective data rate for 64-QAM 3/4 in HD mode of 24.3 Mbps. For 64-QAM 2/3 in FD mode, the computation is similar, except for an additional factor of 2 as the radios operate in FD mode, giving an effective data rate of 33.6 Mbps. This represents a data rate gain in FD mode of roughly 1.4×. If we perform a similar comparison at 20 dB SNR but with 64-QAM 3/4 also used for FD mode, we find a data rate reduction of almost 3×. This suggests that in insufficient SNR regimes, FD link gain may be achieved using a lower data rate MCS.

7.3. Experiment 3: FD gains in an ALOHA network

The experimenter can make use of multiple nodes (in HD or FD mode) for network-level experimentation. In Section 7.2, a link-level FD rate gain of up to 2× is demonstrated with sufficient link SNR values, and it is possible to investigate the impact of this result in a network of multiple FD radios. To demonstrate this capability, a simple slotted ALOHA-based protocol is used for medium access. Up to four Gen-2 FD radios can be used in such a network. Fig. 10 shows these four radios within Sandbox 2, while Fig. 17 shows the physical topology of the four radios, and the SNR between each radio pair is shown in Table 2 for −10 dBm transmit power and 5 MHz bandwidth at 915 MHz carrier frequency. The USRP 2974 and two USRP2s that are connected to the four Gen-2 FD radios are synchronized using an external clock source, allowing for the time-slotted operation.

In this experiment, we utilize three Gen-2 FD radios out of the available four shown in Fig. 17: radios #1 and #2 connected to the USRP-2974 and radio #3 connected to a USRP2. In each time slot with a duration of 22 ms, a radio may only send one packet with an over-the-air transmission time of 2 ms. For each radio, the average node throughput S_i, $i \in \{1, 2, 3\}$, is defined as the ratio between the number of packets successfully *decoded* by radio i, n_i, divided by the total number of time slots T: $S_i = \frac{n_i}{T}$.[5] Similarly to Section 7.2, the experiment is run in two modes: (i) HD mode: all radios operate in HD and utilize a slotted ALOHA protocol with an equal transmission probability of $\frac{1}{3}$, and (ii) FD mode: all radios operate in FD and utilize a slotted ALOHA protocol with an equal transmission probability of $\frac{1}{3}$. The latter mode can result in scenarios where the network can now support two simultaneous transmissions that can each be correctly decoded at the receiver. In particular, a radio i can successfully decode a packet in a given time slot if i transmits *and* one other radio transmits.

Each experiment run consists of 20,000 time slots, during which each radio sends WiFi-like 100-byte data packets with BPSK 1/2 MCS.

[5] Notice that in order to avoid the need for acknowledgment packets, the throughput is computed based on the number of *received* packets. While this is not the "standard" definition, it is easy to analytically evaluate such a small network.

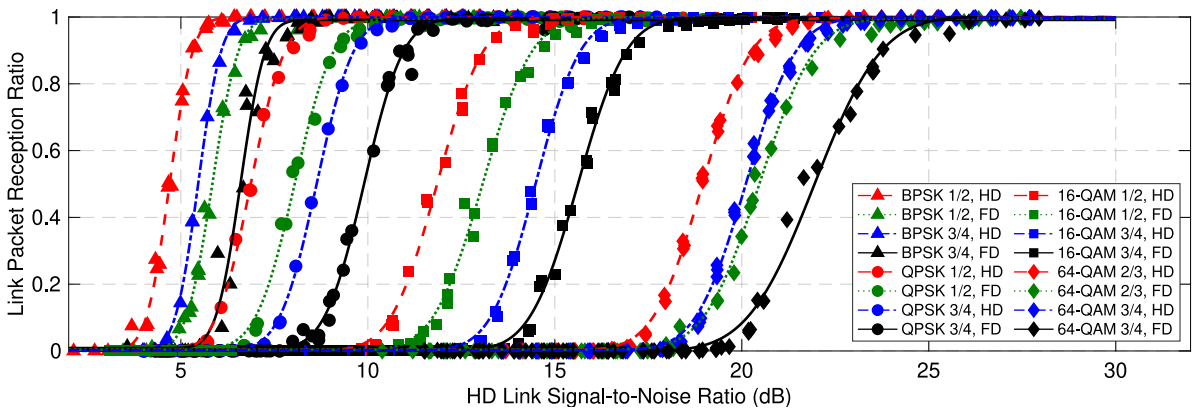

Fig. 16. HD and FD link packet reception ratio (PRR) as a function of varying HD link SNR values for 802.11-like packets with different modulation and coding schemes (MCSs).

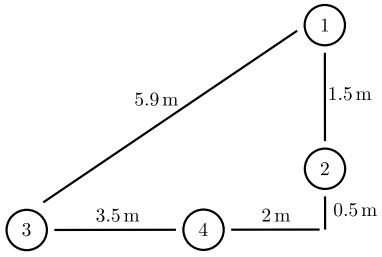

Fig. 17. Physical network topology of the four Gen-2 FD radios in the ALOHA-based network experiment.

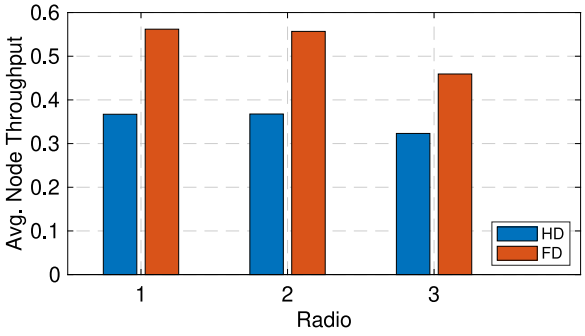

Fig. 18. Node-level throughput results for each radio in HD and FD mode for the network of three Gen-2 FD radios.

Table 2
Summary of the SNR values between each radio pair used in the ALOHA-based network experiment.

		SNR when receiving from...			
Radio #	**SDR**	**1**	**2**	**3**	**4**
1	USRP 2974		29 dB	24 dB	30 dB
2	USRP 2974	40 dB		31 dB	33 dB
3	USRP2	19 dB	20 dB		32 dB
4	USRP2	36 dB	28 dB	40 dB	

The average node throughput for each radio is then computed after the experiment completes. The SNR values between each radio pair described in Table 2 imply that with BPSK 1/2 MCS, it is possible in some cases for a radio to decode a packet even in the presence of multiple transmissions from two other radios. This occurs in the following cases: (i) radio #1 and radio #3 transmit; radio #2 will decode radio #1's packet, (ii) radio #2 and radio #3 transmit; radio #1 will decode radio #2's packet. Therefore, it is straightforward to see that the throughput in HD mode for radio i, $S_{\mathrm{HD,i}}(Rx)$ is $\frac{10}{27}$ for $i = \{1,2\}$ and $\frac{8}{27}$ for $i = 3$. In FD mode, the throughput values $S_{\mathrm{FD,i}}(Rx)$ are $\frac{15}{27}$ for $i = \{1,2\}$ and $\frac{12}{27}$ for $i = 3$.

The results of operating the ALOHA network in HD and FD modes with the network topology in Fig. 17 are shown in Fig. 18. In HD mode, the average node throughput of radios #1 and #2 is close to the analytically computed throughput value, while for radio #3 it is slightly higher, at 0.32. This is likely due to radio #3 having a nonzero probability of successfully decoding one of radio #1 or radio #2's packets if they transmit in the same slot. In FD mode, the average node throughput of radios #1 and #2 approaches $\frac{15}{27}$, and for radio #3 it approaches $\frac{12}{27}$, each showing a throughput increase of up to 50%. Therefore, it is possible to the demonstrate that by simply enabling FD operation, the throughput of a small slotted ALOHA network can be increased by up to 50%.

FD wireless can benefit the MAC layer of a network in ways other than throughput improvement. For example, the ability of an FD radio to receive a signal while it transmits would enable collision detection during an ongoing transmission, a control mechanism traditionally limited to wired networks. Experimenters may use the synchronized Gen-2 radios to investigate such impacts on various network topologies, including the broadcast network used in the ALOHA experiment, as well as centralized networks where one node acts as a base station or access point. The integrated FD radios can also support experimentation with heterogeneous HD/FD networks, where relative priority of HD and FD nodes can impact network throughput and fairness [42]. In such experiments, especially those using a carrier sense-based MAC protocol [33], an important factor is the host-to-SDR latency, the optimization of which is part of our future work on additional customized C++ FD software.

8. Conclusion

In this paper, we presented our cross-layered (hardware and software) design and implementation of the first open-access, remotely accessible FD radios integrated in the ORBIT and COSMOS wireless testbeds. We make openly accessible a variety of FD radios based upon wideband and narrowband RF SI canceller designs, including fixed radios connected to powerful remotely accessible compute nodes, and a mobile radio connected to a portable PC. The presented example experiments along with the tutorial and open-source software [39,40] can be used on any integrated FD radio and expanded to different network scenarios. Therefore, we anticipate that the integrated FD radios and example experiments can facilitate further hands-on research in FD wireless.

In future work, we plan to make several improvements to the integrated FD radios. In particular, we will fully replace the USRP2s with USRP 2974s, which have superior performance that includes supporting a higher signal bandwidth and better RF chain linearity. This will

facilitate a higher level of overall achievable SIC and the use of the 802.11a PHY layer waveforms [41], as well as support RF SI canceller designs of higher bandwidth. We will also make the radios accessible to the servers in COSMOS Sandbox 2, which provide heterogeneous computing resources, including CPUs, GPUs, and FPGAs, as well as configurable connectivity to COSMOS' optical network. Lastly, we will make 28 GHz millimeter-wave (mmWave) phased array antenna modules (PAAMs) [43] connected to the USRP N310 SDRs available to researchers. These 28 GHz PAAMs will support research into FD in mmWave bands [16] and using phased arrays [44].

We anticipate that these improvements will allow for more complex experiments, such as those at the MAC layer [42], which require precise timing and control. Integration with the Sandbox 2 servers will also allow for remote experimentation across the radio, compute, and optical domains [45].

Acknowledgments

This work was supported in part by NSF grants ECCS-1547406, CNS-1827923, and OAC-2029295, NSF-BSF grant CNS-1910757, the DARPA RF-FPGA, SPAR, and WARP programs, an NSF GRFP Fellowship, a GFSD Fellowship, and a Facebook Fellowship. We thank Jelena Diakonikolas, Guy Farkash, Jakub Kolodziejski, Prasanthi Maddala, Jonathan Ostrometzky, Michael Sherman, and Jin Zhou for their contributions to various aspects of the project.

References

[1] Manav Kohli, Tingjun Chen, Mahmood Baraani Dastjerdi, Jackson Welles, Ivan Seskar, Harish Krishnaswamy, Gil Zussman, Open-access full-duplex wireless in the ORBIT and COSMOS testbeds, in: Proc. ACM 14th International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization (WiNTECH'20), 2020.

[2] Ashutosh Sabharwal, Philip Schniter, Dongning Guo, Daniel W. Bliss, Sampath Rangarajan, Risto Wichman, In-band full-duplex wireless: Challenges and opportunities, IEEE J. Sel. Areas Commun. 32 (9) (2014) 1637–1652.

[3] Kenneth E. Kolodziej, Bradley T. Perry, Jeffrey S. Herd, In-band full-duplex technology: Techniques and systems survey, IEEE Trans. Microw. Theory Tech. 67 (7) (2019) 3025–3041.

[4] Dinesh Bharadia, Emily McMilin, Sachin Katti, Full duplex radios, in: Proc. ACM SIGCOMM'13, 2013.

[5] Liang Zhang, Nirwan Ansari, A framework for 5G networks with in-band full-duplex enabled drone-mounted base-stations, IEEE Wirel. Commun. 26 (5) (2019) 121–127.

[6] The Columbia FlexICoN project, https://flexicon.ee.columbia.edu/.

[7] Jin Zhou, Negar Reiskarimian, Jelena Diakonikolas, Tolga Dinc, Tingjun Chen, Gil Zussman, Harish Krishnaswamy, Integrated full duplex radios, IEEE Commun. Mag. 55 (4) (2017) 142–151.

[8] Tingjun Chen, Jin Zhou, Nicole Grimwood, Rel Fogel, Jelena Marašević, Harish Krishnaswamy, Gil Zussman, Demo: Full-duplex wireless based on a small-form-factor analog self-interference canceller, in: Proc. ACM MobiHoc'16, 2016.

[9] Jin Zhou, Anandaroop Chakrabarti, Peter Kinget, Harish Krishnaswamy, Low-noise active cancellation of transmitter leakage and transmitter noise in broadband wireless receivers for FDD/Co-Existence, IEEE J. Solid-State Circuits 49 (12) (2014) 1–17.

[10] Jelena Marasevic, Jin Zhou, Harish Krishnaswamy, Yuan Zhong, Gil Zussman, Resource allocation and rate gains in practical full-duplex systems, IEEE/ACM Trans. Netw. 25 (1) (2017) 292–305.

[11] Tingjun Chen, Mahmood Baraani Dastjerdi, Jin Zhou, Harish Krishnaswamy, Gil Zussman, Wideband full-duplex wireless via frequency-domain equalization: Design and experimentation, in: Proc. ACM MobiCom'19, 2019.

[12] Jin Zhou, Tsung-Hao Chuang, Tolga Dinc, Harish Krishnaswamy, Integrated wideband self-interference cancellation in the RF domain for FDD and full-duplex wireless, IEEE J. Solid-State Circuits 50 (12) (2015) 3015–3031.

[13] Open-access research testbed for next-generation wireless networks (ORBIT), https://www.orbit-lab.org/.

[14] Dipankar Raychaudhuri, Ivan Seskar, Gil Zussman, Thanasis Korakis, Dan Kilper, Tingjun Chen, Jakub Kolodziejski, Michael Sherman, Zoran Kostic, Xiaoxiong Gu, Harish Krishnaswamy, Sumit Maheshwari, Panagiotis Skrimponis, Craig Gutterman, Challenge: COSMOS: A city-scale programmable testbed for experimentation with advanced wireless, in: Proc. ACM MobiCom'20, 2020.

[15] Cloud Enhanced Open Software Defined Mobile Wireless Testbed for City-Scale Deployment (COSMOS), 2020, https://cosmos-lab.org/.

[16] Vaibhav Singh, Susnata Mondal, Akshay Gadre, Milind Srivastava, Jeyanandh Paramesh, Swarun Kumar, Millimeter-wave full duplex radios, in: Proc. ACM MobiCom'20, 2020.

[17] Muhammad Sohaib Amjad, Falko Dressler, Software-based real-time full-duplex relaying: An experimental study, IEEE Trans. Green Commun. Netw. 4 (3) (2020) 647–656.

[18] Mohamad Katanbaf, Anthony Weinand, Vamsi Talla, Simplifying backscatter deployment: Full-duplex LoRa backscatter, in: Proc. NSDI'21, 2021.

[19] Jung Il Choi, Mayank Jain, Kannan Srinivasan, Phil Levis, Sachin Katti, Achieving single channel, full duplex wireless communication, in: Proc. ACM MobiCom'10, 2010.

[20] Bozidar Radunovic, Dinan Gunawardena, Peter Key, Alexandre Proutiere, Nikhil Singh, Vlad Balan, Gerald Dejean, Rethinking indoor wireless mesh design: Low power, low frequency, full-duplex, in: Proc. IEEE WIMESH'10, 2010.

[21] Mayank Jain, Jung Il Choi, Taemin Kim, Dinesh Bharadia, Siddharth Seth, Kannan Srinivasan, Philip Levis, Sachin Katti, Prasun Sinha, Practical, real-time, full duplex wireless, in: Proc. ACM MobiCom'11, 2011.

[22] MinKeun Chung, Min Soo Sim, Jaeweon Kim, Dong Ku Kim, Chan-Byoung Chae, Prototyping real-time full duplex radios, IEEE Commun. Mag. 53 (9) (2015) 56–63.

[23] Dani Korpi, Joose Tamminen, Matias Turunen, Timo Huusari, Yang-Seok Choi, Lauri Anttila, Shilpa Talwar, Mikko Valkama, Full-duplex mobile device: Pushing the limits, IEEE Commun. Mag. 54 (9) (2016) 80–87.

[24] Kate Ching-Ju Lin, Kai-Cheng Hsu, Hung-Yu Wei, Inter-client interference cancellation for full-duplex networks with half-duplex clients, IEEE/ACM Trans. Netw. 27 (5) (2019) 2150–2163.

[25] Ehsan Aryafar, Alireza Keshavarz-Haddad, FD2: A directional full duplex communication system for indoor wireless networks, in: Proc. IEEE INFOCOM'15, 2015.

[26] Kai-Cheng Hsu, Kate Ching-Ju Lin, Hung-Yu Wei, Inter-client interference cancellation for full-duplex networks, in: Proc. IEEE INFOCOM'17, 2017.

[27] Shih-Ying Chen, Ting-Feng Huang, Kate Ching-Ju Lin, Y-W Peter Hong, Ashutosh Sabharwal, Probabilistic medium access control for full-duplex networks with half-duplex clients, IEEE Trans. Wirel. Commun. 16 (4) (2017) 2627–2640.

[28] Aimin Tang, Xudong Wang, A-duplex: Medium access control for efficient coexistence between full-duplex and half-duplex communications, IEEE Trans. Wirel. Commun. 14 (10) (2015) 5871–5885.

[29] RF-CI SMA-F Circulator Model RFCR3204, mechanical drawing and data sheet, 2020. http://www.rf-ci.com/wp-content/themes/rfci/pdf/sma-connector-thin-circulator-communication/CR3204-OSrA.pdf. (Accessed on 06/20/2020).

[30] DIMAX SUB-20 user manual, 2019. https://www.xdimax.com/sub20/doc/sub20-man.pdf. (Accessed on 05/15/2019).

[31] GNU Radio, http://gnuradio.org/.

[32] USRP Hardware Driver (UHD) software, 2020. https://github.com/EttusResearch/uhd. (Accessed on 07/01/2020).

[33] Shengbo Liu, Liqun Fu, Wei Xie, Hidden-node problem in full-duplex enabled CSMA networks, IEEE Trans. Mob. Comput. 19 (2) (2020) 347–361.

[34] Xi Zhang, Wenchi Cheng, Hailin Zhang, Full-duplex transmission in PHY and MAC layers for 5G mobile wireless networks, IEEE Wirel. Commun. 22 (5) (2015) 112–121.

[35] Claudia Campolo, Antonella Molinaro, Antoine. O. Berthet, Alexey Vinel, Full-duplex radios for vehicular communications, IEEE Commun. Mag. 55 (6) (2017) 182–189.

[36] Zhifeng Yuan, Yihua Ma, Yuzhou Hu, Weimin Li, High-efficiency full-duplex V2V communication, in: Proc. 2020 2nd 6G Wireless Summit (6G SUMMIT), 2020.

[37] Service: XY Table, 2020, https://wiki.cosmos-lab.org/wiki/Resources/Services/XYTable.

[38] Nodes: Portable Nodes, 2020, https://wiki.cosmos-lab.org/wiki/Hardware/Nodes#Portable.

[39] The Columbia FlexICoN project: Instructions and Code, 2020, https://github.com/Wimnet/flexicon_orbit.

[40] Tutorial: Full-duplex wireless in the ORBIT and COSMOS testbeds, 2020, https://wiki.cosmos-lab.org/wiki/Tutorials/Wireless/FullDuplex.

[41] Bastian Bloessl, Michele Segata, Christoph Sommer, Falko Dressler, An IEEE 802.11a/g/p OFDM receiver for GNU radio, in: Proc. ACM 2nd Workshop on Software Radio Implementation Forum (SRIF'13), 2013.

[42] Tingjun Chen, Jelena Diakonikolas, Javad Ghaderi, Gil Zussman, Hybrid scheduling in heterogeneous half-and full-duplex wireless networks, IEEE/ACM Trans. Netw. 28 (2) (2020) 764–777.

[43] Bodhisatwa Sadhu, Yahya Tousi, Joakim Hallin, Stefan Sahl, Scott K. Reynolds, Örjan Renström, Kristoffer Sjögren, Olov Haapalahti, Nadav Mazor, et al., A 28-GHz 32-element TRX phased-array IC with concurrent dual-polarized operation and orthogonal phase and gain control for 5G communications, IEEE J. Solid-State Circuits 52 (12) (2017) 3373–3391.

[44] Tingjun Chen, Mahmood Baraani Dastjerdi, Harish Krishnaswamy, Gil Zussman, Wideband full-duplex phased array with joint transmit and receive beamforming: Optimization and rate gains, in: Proc. ACM MobiHoc'19, 2019.

[45] Artur Minakhmetov, Craig Gutterman, Tingjun Chen, Cedric Ware, Luigi Iannone, Dan Kilper, Gil Zussman, Experiments on cloud-RAN wireless handover using optical switching in a dense urban testbed, in: Proc. OSA OFC'20, 2020.

Manav Kohli is a Ph.D. student with the Department of Electrical Engineering at Columbia University. He received the Sc.B. degree in electrical engineering from Brown University in 2018. His research focuses on technologies for future wireless networks, including full-duplex wireless, mmWave, and testbeds for open-access wireless experimentation, such as the NSF Platforms for Advanced Wireless Research COSMOS testbed in West Harlem, New York City. Manav's work has been supported by the NSF Graduate Research Fellowship Program and the National Physical Science Consortium's Graduate Fellowships for STEM Diversity.

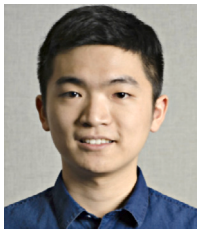

Tingjun Chen received the Ph.D. degree in electrical engineering from Columbia University in 2020, and the B.Eng. degree in electronic engineering from Tsinghua University in 2014. He is an Assistant Professor in the Department of Electrical and Computer Engineering at Duke University and was a postdoctoral associate at Yale University from 2020–2021. His research interests are in the area of networking and communications with a specific focus on next-generation wireless networks and Internet-of-Things systems. He received the Google Research Scholars Award, the IBM Faculty Award, the Facebook Fellowship, the Wei Family Private Foundation Fellowship, the Columbia Engineering Morton B. Friedman Memorial Prize for Excellence, the Columbia University Eli Jury Award and Armstrong Memorial Award, the ACM SIGMOBILE Doctoral Dissertation Award Runner-up, and the ACM CoNEXT'16 Best Paper Award.

Mahmood Baraani Dastjerdi received the M.S. degree from Sharif University of Technology, Tehran, Iran, in 2014, and the Ph.D. degree in electrical engineering from Columbia University, New York, NY in 2020. His career is focused on the theory, design and experimental validation of analog/RF/millimeter-wave integrated circuits and systems. Dr. Baraani Dastjerdi joined MixComm Inc. as a member of technical staff in 2019. He is recipient of 2020 Columbia University Electrical Engineering Department Jury Award, 2019 IEEE SSCS Predoctoral Achievement Award, 2018 ADI ISSCC outstanding student designer award and creative tech award in engineering at the 2018 NYC Media Lab.

Jackson Welles received the B.S. degree in Electrical Engineering from Columbia University, New York, NY, in 2019, where he worked on the development of PCB-based RF self-interference cancellers. After graduating, he joined Latch in NYC as a firmware engineer, specializing in the development of building automation systems. Jackson received the EE Department Research Award for outstanding contribution to research from Columbia University in 2019.

Ivan Seskar is the Chief Technologist 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, the PI for the NSF GENI Wireless project, which resulted in campus deployments of LTE/WiMAX base stations at several US universities, and the PI for the NSF CloudLab deployment at Rutgers. 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, a Senior Member of the IEEE, a member of ACM and the co-founder and CTO of Upside Wireless Inc.

Harish Krishnaswamy received the B.Tech. degree in electrical engineering from IIT 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, respectively. In 2009, he joined the Electrical Engineering Department, Columbia University, New York, NY, USA, where he is currently an Associate 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. His research interests include integrated devices, circuits, and systems for a variety of RF, millimeter-wave (mmWave), and sub-mmWave applications. Dr. Krishnaswamy 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 Defense Advanced Research Projects Agency 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 been a member of the technical program committee of several conferences, including the IEEE International Solid-State Circuits Conference since 2015 and the IEEE Radio Frequency Integrated Circuits Symposium since 2013. He has also served as a Distinguished Lecturer for the IEEE Solid-State Circuits Society and is a member for the DARPA Microelectronics Exploratory Council.

Gil Zussman received the Ph.D. degree in electrical engineering from the Technion in 2004 and was a postdoctoral associate at MIT in 2004–2007. He has been with Columbia University since 2007, where he is a Professor of Electrical Engineering and Computer Science (affiliated faculty). His research interests are in the area of networking, and in particular in the areas of wireless, mobile, and resilient networks. He has been an associate editor of IEEE/ACM Transactions on Networking, IEEE Transactions on Control of Network Systems, and IEEE Transactions on Wireless Communications. He is a co-recipient of 7 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 received the Fulbright Fellowship, the DTRA Young Investigator Award, and the NSF CAREER Award. He was the PI of a team that won the 1st place in the 2009 Vodafone Americas Foundation Wireless Innovation Project competition and is the Columbia PI of the NSF PAWR COSMOS testbed.