

Integrated Full Duplex Radios

Jin Zhou, Negar Reiskarimian, Jelena Diakonikolas, Tolga Dinc, Tingjun Chen, Gil Zussman, and Harish Krishnaswamy

Full duplex wireless has drawn significant interest in the recent past due to the potential for doubling network capacity in the physical layer and offering numerous other benefits at higher layers. However, the implementation of integrated full duplex radios is fraught with several fundamental challenges.

ABSTRACT

Full duplex wireless has drawn significant interest in the recent past due to the potential for doubling network capacity in the physical layer and offering numerous other benefits at higher layers. However, the implementation of integrated full duplex radios is fraught with several fundamental challenges. Achieving the levels of self-interference cancellation required over the wide bandwidths mandated by emerging wireless standards is challenging in an integrated circuit implementation. The dynamic range limitations of integrated electronics restrict the transmitter power levels and receiver noise floor levels that can be supported in integrated full duplex radios. Advances in compact antenna interfaces for full duplex are also required. Finally, networks employing full duplex nodes will require a complete rethinking of the medium access control layer as well as cross-layer interaction and co-design. This article describes recent research results that address these challenges. Several generations of full duplex transceiver ICs are described that feature novel RF self-interference cancellation circuits, antenna cancellation techniques, and a non-magnetic CMOS circulator. Resource allocation algorithms and rate gain/improvement characterizations are also discussed for full duplex configurations involving IC-based nodes.

INTRODUCTION

One of the long-held precepts of wireless communication has been that it is impossible for a wireless device to transmit and receive at the same time and at the same frequency because of the resulting self-interference (SI). Recent efforts have challenged this wisdom, opening the door to *full duplex* (FD) wireless, which has the potential to immediately double network capacity at the physical (PHY) layer and offers many other benefits at the higher layers.

The concept of FD communication is certainly not entirely new. The earliest pre-electronic telephone handsets used hybrid transformers to isolate the earpiece from the microphone, enabling FD communication on a two-wire loop to the central office. In the realm of wireless, the Plessey Groundsat system of the 1970s was a military-specification FD wireless system capable of operating over radio channels within the 30–76 MHz VHF band.

However, achieving FD operation in commer-

cial wireless applications, such as cellular communications and WiFi, is fraught with several challenges. The fundamental challenge associated with FD wireless is the tremendous transmitter (TX) SI, or echo, that appears at the receiver (RX). This SI can be anywhere between 90–120 dB (a *billion* to a *trillion* times) more powerful than the desired signal depending on the application. This powerful SI is further susceptible to the uncertainties of the wireless channel (e.g., frequency selectivity and time variance) and the imperfections of the transceiver electronics (nonlinear distortion and phase noise, to name a couple). These challenges are further exacerbated when integrated implementations targeting cost-sensitive and form-factor-constrained mobile devices are considered. Finally, to fully utilize the benefits of FD communication, wireless systems will require a fundamental rethinking of not only the PHY layer but also the medium access control (MAC) layer, and a careful co-design of the two.

Initial research performed a few years ago established the feasibility of SI cancellation and FD operation in commercial wireless applications using laboratory bench-top equipment and off-the-shelf components [1]. More than 100 dB SI cancellation has also been demonstrated in [2] for military applications. However, the self-interference cancellation (SIC) techniques utilized in these works are fundamentally not compatible with small-form-factor/integrated circuit (IC) implementations. More recently, research on integrated FD radios and associated SI cancellation techniques has emerged [3–7].

This article reviews recent research at Columbia University on integrated FD radios spanning RF, analog and digital SIC, FD antenna interfaces, and non-magnetic complementary metal oxide semiconductor (CMOS) circulators that enable single-antenna FD [3–5, 8, 9]. This article also touches on the FlexICoN project at Columbia, which is taking a holistic cross-layered approach to develop FD radios and networks from PHY to MAC. It covers resource allocation algorithms and rate gate/improvement characterizations for FD configurations involving IC-based FD nodes.

CHALLENGES ASSOCIATED WITH INTEGRATED FULL DUPLEX RADIOS

Figure 1 depicts the block diagram of an FD radio that incorporates antenna, RF, and digital SI suppression. The indicated transmitter and minimum

Jin Zhou is with the University of Illinois at Urbana-Champaign; Jelena Diakonikolas is with Boston University; Harish Krishnaswamy, Gil Zussman, Negar Reiskarimian, Tolga Dinc, and Tingjun Chen are with Columbia University. This research was performed while Jin Zhou and Jelena Diakonikolas were at Columbia University.

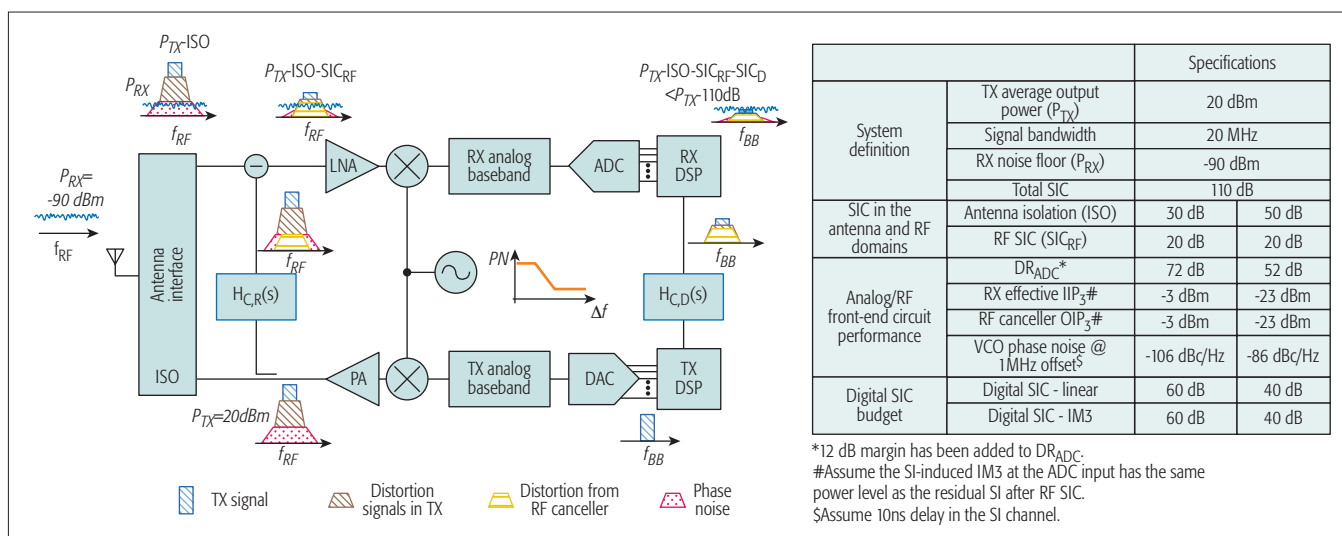


Figure 1. Block diagram of an FD radio featuring antenna, RF and digital SI suppression, along with a depiction of the various transceiver non-idealities that must also be managed for effective FD operation.

received signal power levels are typical of WiFi applications. Also depicted are various transceiver non-idealities that further complicate the SI suppression problem.

ACHIEVING > 100 dB SI SUPPRESSION

The power levels indicated in Fig. 1 necessitate > 110 dB SI suppression for WiFi-like applications. Such an extreme amount of cancellation must necessarily be achieved across multiple domains (here, antenna, RF and digital), as > 100 dB precision from a single stage or circuit is prohibitively complex and power inefficient. The suppression must be judiciously distributed across the domains, as suppression in one domain relaxes the dynamic range requirements of the domains downstream. Furthermore, all cancellation circuits must be adaptively configured together — optimization of the performance of a single cancellation stage alone can result in residual SI that is sub-optimal for the cancellers downstream.

TRANSCIVER NON-IDEALITIES

The extremely powerful nature of the SI exacerbates the impact of non-idealities such as non-linearity and phase noise, particularly for IC implementations. For instance, nonlinearity along the transmitter chain will introduce distortion products. Antenna and RF cancellation that tap from the output of the transmitter will suppress these distortion products, but linear digital cancellation will not as it operates on the undistorted digital signal. Depending on the amount of antenna and RF cancellation achieved, the analog receiver front-end may introduce distortion products as well, as may the RF cancellation circuitry. Nonlinear digital cancellation may be employed to recreate and cancel these distortion products, but the associated complexity and power consumption must be considered. Local oscillator (LO) phase noise can pose problems as well. If a common LO is used for the transmitter and the receiver, the phase noise in the transmitted and the received SI will be completely correlated, enabling its cancellation in the receiver downmixer. However, delay in the SI channel will decor-

relate the phase noise, resulting in residual SI that cannot be cancelled. Figure 1 depicts transceiver performance requirements calculated for two different SIC allocations across domains, one antenna-heavy and the other digital-heavy.

SI CHANNEL FREQUENCY SELECTIVITY AND WIDEBAND RF/ANALOG SI CANCELLATION

The wireless SI channel can be extremely frequency selective. Compact antennas can be quite narrowband, and the front-end filters that are commonly used in today's radios even more so. The wireless SI channel also includes reflections off nearby objects, which will feature a delay that depends on the distance of the object from the radio. Performing wideband RF/analog cancellation requires recreating the wireless SI channel in the RF/analog domain. Conventional analog/RF cancellers feature a frequency-flat magnitude and phase response, and will therefore achieve cancellation only over a narrow bandwidth (BW). Wideband SIC at RF based on time-domain equalization (essentially an RF finite impulse response [FIR] filter) has been reported in [1] using discrete components. However, the integration of nanosecond-scale RF delay lines on an IC is a formidable (perhaps impossible) challenge, and therefore alternate wideband analog/RF SIC techniques are required.

COMPACT FD ANTENNA INTERFACES

FD radios employing a pair of antennas, one for transmit and one for receive, experience a direct trade-off between form factor and transmit-receive isolation arising from the antenna spacing and design. Therefore, techniques that can maintain or even enhance transmit-receive isolation, possibly through embedded cancellation, while maintaining a compact form factor are highly desirable. Compact FD antenna interfaces are also more readily compatible with multiple-input multiple-output (MIMO) and diversity applications, and promote channel reciprocity, which is useful at the higher layers. For highly form-factor-constrained mobile applications, single-antenna FD is required, necessitating the use of circulators. Traditionally, circulators have been implemented

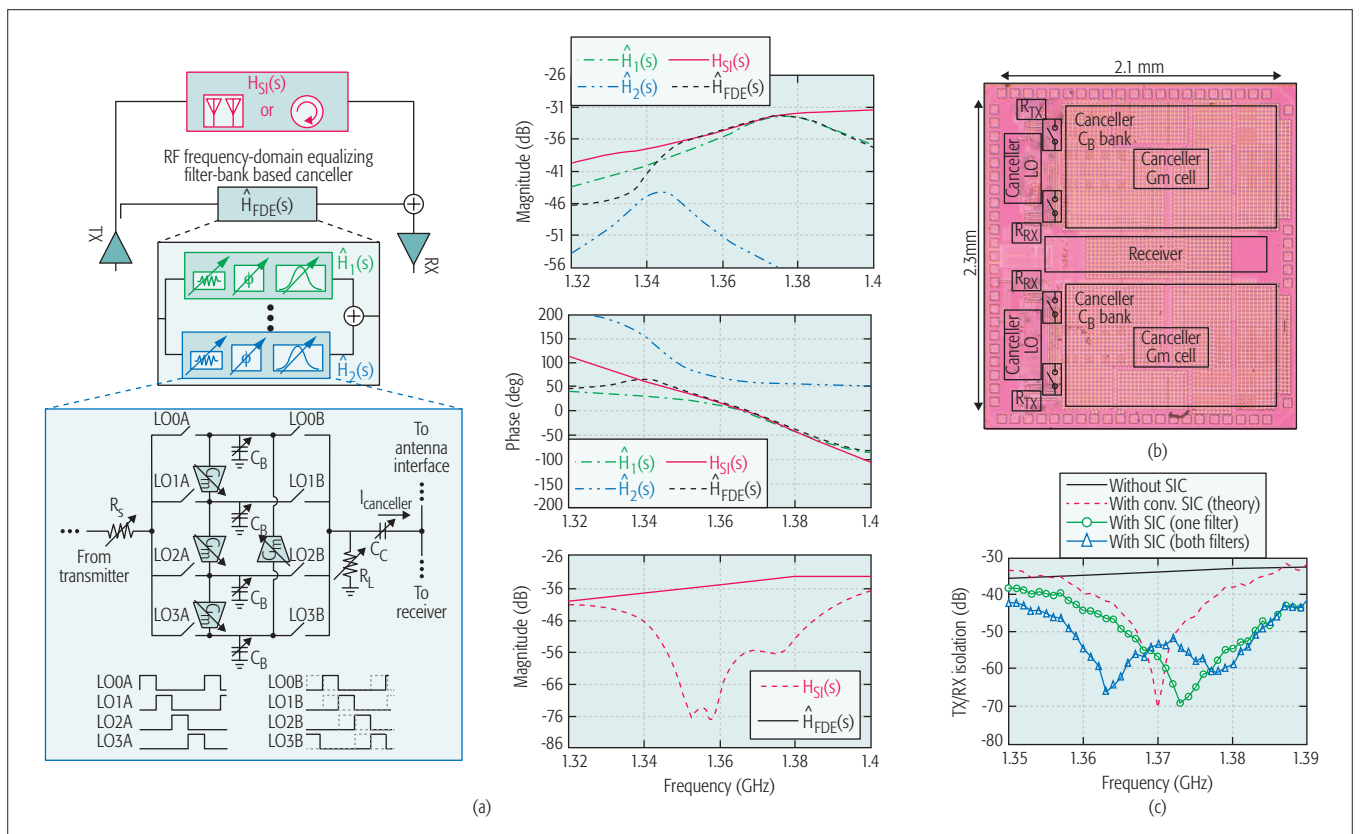


Figure 2. Integrated wideband RF SIC based on frequency-domain equalization (FDE): a) FDE concept and two-port Gm-C N-path filter with embedded variable attenuation and phase shift; b) chip photo of the implemented 0.8-to-1.4 GHz 65 nm CMOS FD receiver with FDE-based SIC in the RF domain featuring a bank of two filters; c) transmit-receive isolation of an antenna pair without SIC, with conventional SIC (theoretical), and with the proposed FDE-based SIC.

using ferrite materials, and are costly, bulky, and not compatible with IC technology. Novel techniques for high-performance non-magnetic integrated circulators are of high interest.

ADAPTIVE CANCELLATION

The SIC in all domains must be reconfigurable and automatically adapt to changing operation conditions (e.g., supply voltage and temperature) and, most importantly, a changing electromagnetic (EM) environment (i.e., wireless SI channel), given the high level of cancellation required. This requires the periodic (or perhaps even continuous) usage of pilot signals to characterize the SI channel, the implementation of reconfigurable cancellers (which is more challenging in the antenna and RF domains), and the development of canceller adaptation algorithms.

RESOURCE ALLOCATION AND RATE GAINS FOR NETWORKS WITH INTEGRATED FD RADIOS, AND RETHINKING MAC PROTOCOLS

The benefits of enabling FD are clear: the uplink (UL) and downlink (DL) rates can theoretically be doubled (in both random access networks, e.g., WiFi, and small cell networks). That, of course, is true, provided that the SI is cancelled such that it becomes negligible at the receiver. Hence, most of the research on FD at the higher layers has focused on designing protocols and assessing the capacity gains while using models of recent laboratory bench-top FD implementations (e.g., [1]) and assuming perfect SI cancellation. How-

ever, given the special characteristics of IC-based SI cancellers, there is a need to understand the capacity gains and develop resource allocation algorithms while taking into account these characteristics. These algorithms will then serve as building blocks for the redesign of MAC protocols for FD networks with integrated FD radios.

INTEGRATED

RF SELF-INTERFERENCE CANCELLATION

To address the challenge related to integrated RF cancellation across a wide bandwidth (BW), we proposed a frequency-domain approach in contrast to the conventional time-domain delay-based RF FIR approach [3]. To enhance the cancellation BW, second-order reconfigurable bandpass filters (BPFs) with amplitude and phase control are introduced in the RF canceller (Fig. 2a). An RF canceller with a reconfigurable second-order RF BPF features four degrees of freedom (amplitude, phase, quality factor, and center frequency of the BPF). This enables the replication of not just the amplitude and phase of the antenna interface isolation $[H_{SI}(s)]$ at a frequency point, but also the slope of the amplitude and the slope of the phase (or group delay). The use of a bank of filters with independent BPF parameters enables such replication at multiple points in different sub-bands, further enhancing SIC BW. Essentially, this approach is frequency-domain equalization (FDE) in the RF domain. In Fig. 2a, which represents a theoretical computation on the measured isolation of a

pair of 1.4 GHz antennas that are described in greater detail below, two BPFs with transfer functions $\hat{H}_1(s)$ and $\hat{H}_2(s)$ emulate the antenna interface isolation in two sub-bands, resulting in an order-of-magnitude improvement in the SIC BW over a conventional frequency-flat RF canceller.

For FDE, reconfigurable RF filters with very sharp frequency response (or high quality factor) are required. The achievable quality factor of conventional LC-based integrated RF filters has been limited by the quality factor of the inductors and capacitors that are available on silicon. However, recent research advances have revived a switched-capacitor circuit-design technique known as the N -path filter that enables the implementation of reconfigurable, high-quality filters at RF in nanoscale CMOS IC technology [10]. Figure 2a depicts a two-port N -path filter, where R_S and R_L are the resistive loads at the transmit and receive sides, respectively. C_C weakly couples the cancellation signal to the receiver input for SIC. The quality factor of an N -path filter may be reconfigured via the baseband capacitor C_B , given fixed R_S and R_L . Through clockwise/counter-clockwise (only counter-clockwise connection is shown in Fig. 2a for simplicity) connected reconfigurable transconductors (G_m), an upward/downward frequency offset with respect to the switching frequency can be introduced without having to change the clock frequency [10]. Variable attenuation can be introduced by reconfiguring R_S and R_L relative to each other. *Interestingly, phase shifts can be embedded in a two-port N -path filter by phase shifting the clocks driving the switches on the output side relative to the input-side clocks as shown in Fig. 2a [3].* All in all, the ability to integrate reconfigurable high-quality RF filters on chip using switches and capacitors uniquely enables synthesis of nanosecond-scale delays through FDE over time-domain equalization.

A 0.8–1.4 GHz FD receiver IC prototype with the FDE-based RF SI canceller was designed and fabricated in a conventional 65 nm CMOS technology (Fig. 2b) [3]. For the SIC measurement results shown in Fig. 2c, we used a 1.4 GHz narrowband antenna-pair interface with peak isolation magnitude of 32 dB, peak isolation group delay of 9 ns, and 3 dB of isolation magnitude variation over 1.36 to 1.38 GHz. The SI canceller achieves a 20 dB cancellation BW of 15/25 MHz (one/two filters) in Fig. 2c. When a conventional frequency-flat amplitude- and phase-based canceller is used, the SIC BW is about 3 MHz ($> 8\times$ lower). The 20 MHz bandwidth over which the cancellation is achieved allows our FD receiver IC to support many advanced wireless standards including small-cell LTE and WiFi.

COMPACT FULL DUPLEX ANTENNAS EMPLOYING RECONFIGURABLE POLARIZATION-BASED SI CANCELLATION

The suppression of the SI within the antenna interface itself has significant advantages, as it relaxes the dynamic range requirements on the RF, analog, and digital blocks in the receiver chain as well as the RF/analog and digital SIC circuits. *While contemplating SI suppression at the antenna, it is important to keep in mind that FD antenna interfaces must also exhibit a compact form factor, preserve*

their radiation patterns, and maintain SI suppression in the presence of a changing EM environment. Prior antenna-domain SI suppression approaches only partially satisfy these requirements [11]. In particular, they are typically static approaches that are unable to respond to a changing environment.

The antenna-electronics interface offers a unique opportunity to blend EM, RF, analog, and digital concepts to create “smart” antennas that achieve novel functionality. *In particular, the antenna domain offers another degree of freedom, wave polarization, apart from the conventional amplitude, phase, and frequency, which are the workhorses of the electronic domain.* Orthogonal polarizations can be exploited to enhance transmit-receive isolation, and, *more interestingly, embed reconfigurable SIC.*

Using this insight, we recently developed a wideband reconfigurable polarization-based antenna cancellation technique for FD. The technique is depicted in Fig. 3a and employs a pair of compact co-located antennas for the TX and the RX that use orthogonal polarizations to enhance the initial TX-to-RX isolation. Additionally, an auxiliary port that is co-polarized with the TX antenna is introduced on the RX antenna and terminated with a reconfigurable reflective termination (essentially a programmable filter). Since this port is co-polarized with the TX antenna, it “steals” a small portion of the transmitted signal, thus creating an indirect coupling path between the TX and RX ports. The signal in the indirect path is conditioned through the reflective termination and then couples to the RX port. By configuring the reflective termination appropriately, SIC can be achieved at the RX port. Through the implementation of a higher-order reflection termination, our technique can mimic the direct path’s magnitude and phase as well as their slopes at multiple frequency points to achieve wideband cancellation in a manner similar to the FDE technique described earlier. The electronically programmable nature of this reflective termination also allows SIC to be reconfigured in-field in the face of a changing EM environment.

A 4.6 GHz antenna prototype employing this technique was built (Fig. 3b), and achieves more than 50 dB isolation over 300 MHz bandwidth [8]. This represents a $20\times$ improvement in isolation bandwidth over conventional techniques. A strong reflector (a metallic plate) was also placed near the antenna during measurement, and the ability to reconfigure and recover the cancellation despite the reflector’s presence was demonstrated (Fig. 3b).

Such antenna-electronics co-design techniques readily translate to higher frequencies where antennas are naturally smaller and interface more tightly with the IC. Recently, we reported a transceiver IC that combines FD with millimeter-waves [4]. Merging millimeter-waves with FD can potentially offer the dual benefits of wide bandwidths and improved spectral efficiency, a step toward delivering the tremendous increase in capacity demanded by emerging wireless standards. As shown in Fig. 3c, our 60 GHz FD transceiver IC employs the reconfigurable wideband polarization-based antenna cancellation discussed earlier. The reconfigurable reflective termination is implemented on the chip. To improve the SI suppression further, an RF canceller from the TX out-

The suppression of the SI within the antenna interface itself has significant advantages, as it relaxes the dynamic range requirements on the RF, analog, and digital blocks in the receiver chain as well as the RF/analog and digital SIC circuits.

put to the low-noise amplifier (LNA) output is also included in the transceiver. The complete 60 GHz FD transceiver architecture and its chip microphotograph are shown in Fig. 3d. It is implemented in a 45 nm silicon-on-insulator (SOI) CMOS process and achieves the highest integration level among FD transceivers irrespective of the operation frequency. In conjunction with digital

SIC implemented in MATLAB after capturing the baseband (BB) signals using an Agilent 54855A oscilloscope (essentially an 8-bit 20 GSps ADC), a total SI suppression of nearly 80 dB was achieved over 1 GHz BW, enabling the world's first millimeter-wave FD link over a distance of almost 1 m. Figure 3e shows the demonstration setup using a 100 MHz offset continuous wave (CW) signal and

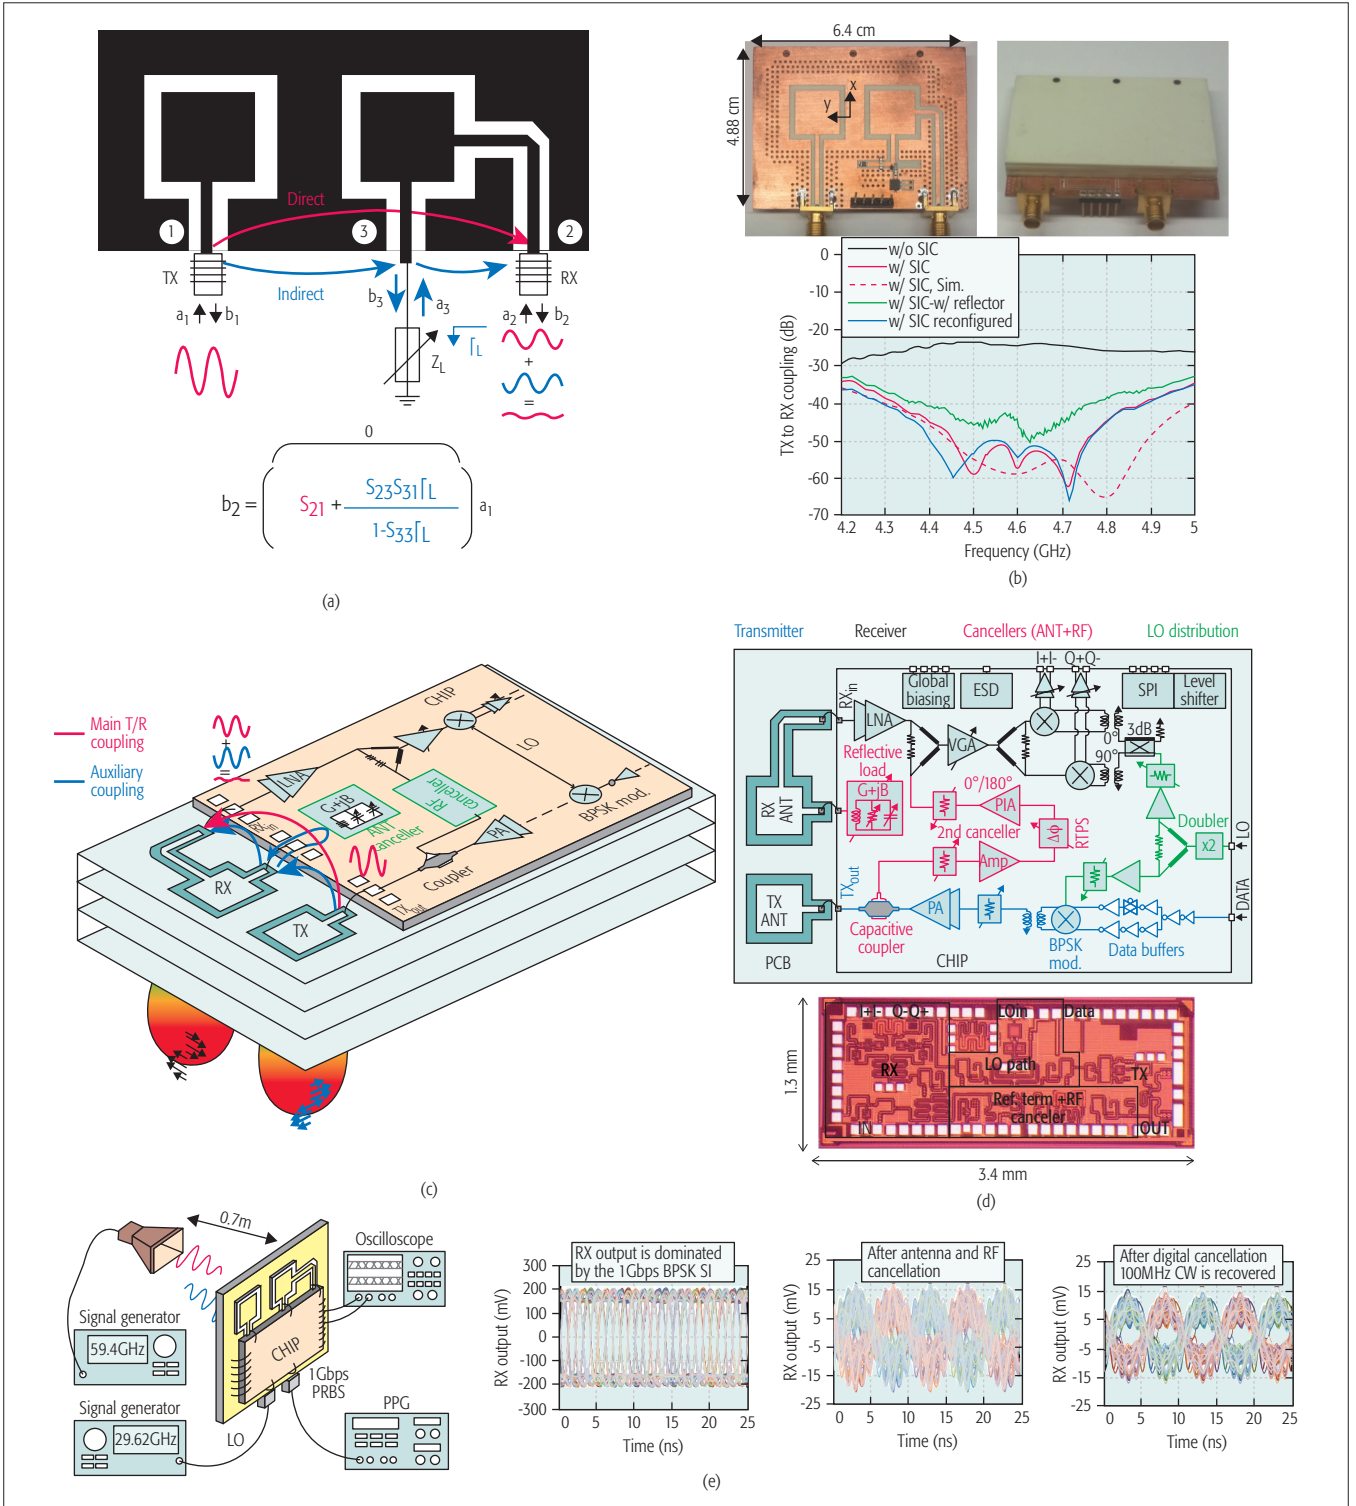


Figure 3. Polarization-based reconfigurable wideband antenna cancellation: a) concept; b) 4.6 GHz TX/RX antenna pair prototype and associated measurement results; c) 3D implementation view of a 60 GHz FD transceiver employing the proposed antenna cancellation; d) 60 GHz fully integrated FD transceiver architecture and IC microphotograph; e) 60 GHz FD link setup and demonstration.

1 Gb/s binary phase shift keying (BPSK) as the received and the transmitted signal (SI), respectively. In the absence of antenna and RF SIC, the RX output is dominated by SI. Engaging the antenna and RF SIC enables the discerning of the desired signal. Digital cancellation in MATLAB (a 100-tap adaptive LMS filter with a settling time of 16 μ s) further suppresses the SI, resulting in an even cleaner received signal with a signal-to-interference-noise-and-distortion ratio (SINDR) of 7.2 dB.

INTEGRATED NON-MAGNETIC CIRCULATOR FOR SINGLE-ANTENNA FULL DUPLEX

Highly-form-factor-constrained mobile applications, particularly at RF frequencies where the wavelength is considerably higher, demand single-antenna solutions. Single-antenna FD also ensures channel reciprocity and compatibility with antenna diversity and MIMO concepts. However, *conventional single-antenna FD interfaces, that is, non-reciprocal circulators, rely on ferrite materials and biasing magnets, and are consequently bulky, expensive and incompatible with silicon integration.* Reciprocal circuits, such as electrical-balance duplexers, have been considered, but are limited by the fundamental minimum 3 dB loss in both TX-antenna (ANT) and ANT-RX paths.

As mentioned earlier, non-reciprocity and circulation have conventionally been achieved using the magneto-optic Faraday effect in ferrite materials. However, it has recently been shown that violating time invariance within a linear, passive material with symmetric permittivity and permeability tensors can introduce non-reciprocal wave propagation, enabling the construction of non-magnetic circulators [12]. However, these initial efforts have resulted in designs that are either lossy, highly nonlinear, or comparable in size to the wavelength, and are fundamentally not amenable to silicon integration. Recently, we introduced a new non-magnetic CMOS-compatible circulator concept based on the phase-non-reciprocal behavior of linear, periodically time-varying (LPTV) two-port N -path filters that utilize staggered clock signals at the input and output [9].

N -path filters, described earlier in the context of FDE, are a class of LPTV networks where the signal is periodically commutated through a bank of linear, time-invariant (LTI) networks. We found that when the non-overlapping clocks driving the input and output switch sets of a two-port N -path filter are phase shifted with respect to each other, a nonreciprocal phase shift is produced for signals traveling in the forward and reverse directions as they see a different ordering of the commutating switches (Fig. 4a). The magnitude response remains reciprocal and low-loss, similar to traditional N -path filters. To create non-reciprocal wave propagation, an N -path filter with $\pm 90^\circ$ phase shift is placed inside a $3\lambda/4$ transmission line loop (Fig. 4b). This results in satisfaction of the boundary condition in one direction (-270° phase shift from the loop added with -90° from the N -path filter) and suppression of wave propagation in the other direction ($-270^\circ + 90^\circ = -180^\circ$), effectively producing unidirectional circulation. Additionally, a three-port circulator can be realized by placing ports anywhere along the loop as long as they maintain a $\lambda/4$ circumferential distance

between them. Interestingly, maximum linearity with respect to the TX port is achieved if the RX port is placed adjacent to the N -path filter ($l = 0$), since the inherent TX-RX isolation suppresses the voltage swing on either side of the N -path filter, enhancing its linearity.

A prototype circulator based on these concepts operating over 610-850 MHz was implemented in a 65 nm CMOS process. Measurements reveal 1.7 dB loss in TX-ANT and ANT-RX transmission, and broadband isolation better than 15 dB between TX and RX (the narrowband isolation can be as high as 50 dB). The in-band ANT-RX IIP3 is +8.7 dBm while the in-band TX-ANT IIP3 is +27.5 dBm (OIP3 = +25.8 dBm), two orders of magnitude higher due to the suppression of swing across the N -path filter. The measured clock feedthrough to the ANT port is -57 dBm, and IQ image rejection for TX-ANT transmission is 49 dB. Techniques such as device stacking in SOI CMOS can be explored to further enhance the TX-ANT linearity to meet the stringent requirements of commercial wireless standards. Clock feedthrough and IQ mismatch can be calibrated by sensing and injecting appropriate BB signals through the N -path filter capacitor nodes as shown previously in the literature.

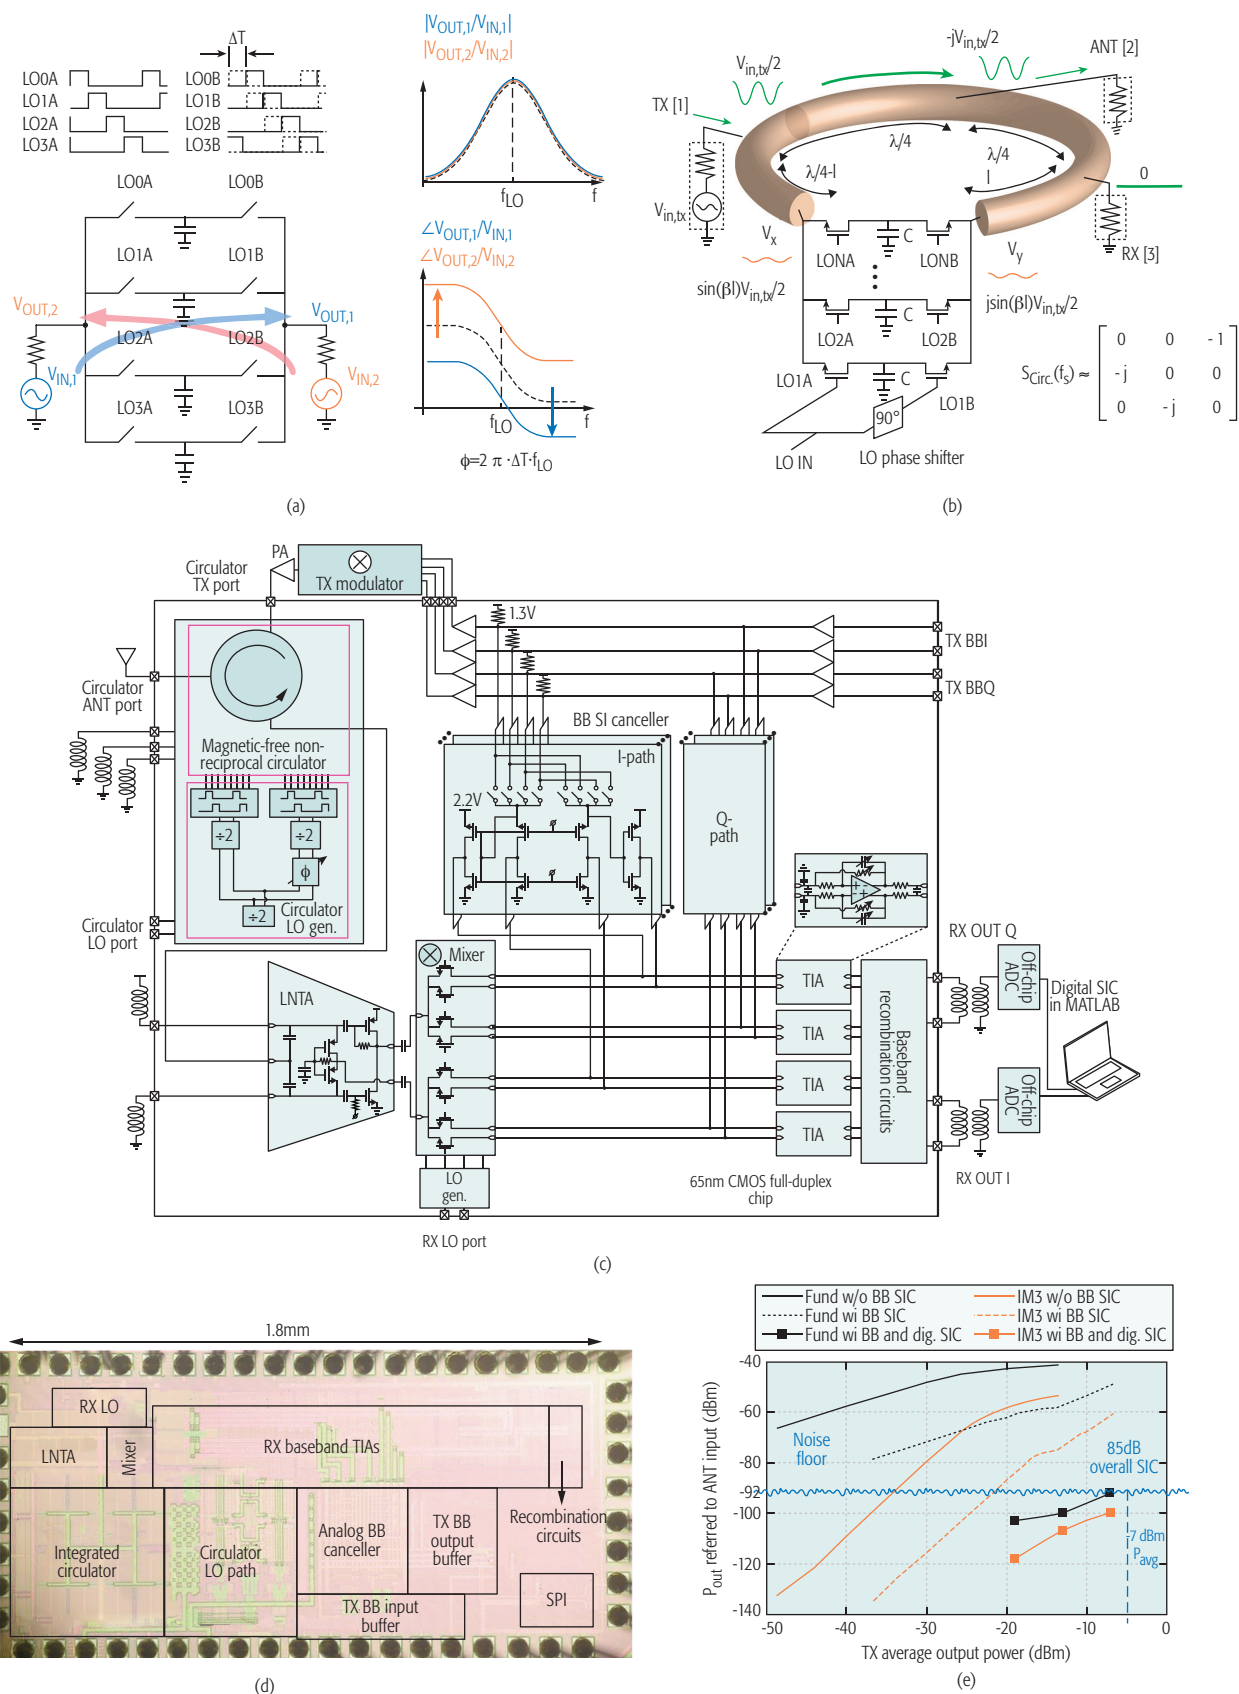
A 610-850 MHz FD receiver IC prototype incorporating the non-magnetic N -path-filter-based passive circulator and additional analog BB SI cancellation (shown in Figs. 4c and 4d) was also designed and fabricated in the 65 nm CMOS process [5]. SI suppression of 42 dB is achieved across the circulator and analog BB SIC over a BW of 12 MHz. Digital SIC has also been implemented in MATLAB after capturing the BB signals using an oscilloscope (effectively an 8-bit 40 MS/s ADC) (Fig. 4e). The digital SIC cancels not only the main SI but also the IM3 distortion generated on the SI by the circulator, receiver, and canceller. A total of 164 canceller coefficients are trained by 800 sample points. After digital SIC, the main SI tones are at the -92 dBm noise floor, while the SI IM3 tones are 8 dB below for -7 dBm TX average power. This corresponds to an overall SI suppression of 85 dB for the FD receiver.

RESOURCE ALLOCATION AND ACHIEVABLE RATE IMPROVEMENTS IN FD

Cancelling SI to a negligible level is extremely challenging. Therefore, we have been considering the following questions: *How much can be gained given a realistic canceller residual SI profile and signal strength? And when does it make sense to use FD over legacy time-division duplex (TDD)?* These questions need to be addressed in the context of a hybrid network with both FD and legacy half-duplex (HD) nodes, as illustrated in Fig. 5a. As a first step and in order to obtain fundamental understanding, in [13, 14] we focused on a single FD bidirectional link with orthogonal channels (e.g., orthogonal frequency division multiplexing – OFDM), and obtained analytic and algorithmic resource (power, time, and frequency) allocation results that quantify the achievable rate improvements as a function of SI-to-noise ratios (XINRs) and signal-to-noise ratios (SNRs).

To model achievable rates on the UL and DL, we used Shannon's capacity formula. We

Conventional single-antenna FD interfaces, that is, non-reciprocal circulators, rely on ferrite materials and biasing magnets, and are consequently bulky, expensive, and incompatible with silicon integration. Reciprocal circuits, such as electrical-balance duplexers, have been considered, but are limited by the fundamental minimum 3 dB loss in both TX-antenna (ANT) and ANT-RX paths.



assumed that the residual SI is a constant fraction of the transmitted signal, where the “constant fraction” can be different for different channels. Under such a model, assuming that there are K orthogonal channels, the rate on the DL can be written as

$$r_b = \sum_{k=1}^K \log \left(1 + \frac{\text{SNR}_{\text{bm},k}}{1 + \text{XINR}_{\text{mm},k}} \right),$$

where $\text{SNR}_{\text{bm},k}$ is the SNR at the mobile station (MS) on channel k and $\text{XINR}_{\text{mm},k}$ is the XINR at the MS on channel k . Similarly, the UL rate is

$$r_m = \sum_{k=1}^K \log \left(1 + \frac{\text{SNR}_{\text{mb},k}}{1 + \text{XINR}_{\text{bb},k}} \right).$$

Note that the noise levels at the MS and the BS and over orthogonal channels are not assumed to be equal.

We now briefly outline a few main results. As illustrated in Fig. 5b, for a given FD rate pair (r_b, r_m) , we defined the *FD rate improvement* p as the (positive) number for which $(r_b/p, r_m/p)$ is at the boundary of the corresponding TDD capacity region.

First, we focused on maximizing the sum of the UL and DL rates (referred to as the *sum rate*) when *only a single channel is considered*. We showed that *if any FD rate pair has higher sum than the maximum TDD rate, the maximum sum rate is obtained when both UL and DL TX power levels are set to their maximum values*. Let (s_b, s_m) denote the rate pair corresponding to the maximum TX power levels. Figures 5c and 5d show the rate improvements at (s_b, s_m) when the XINR at the BS is 0 dB, as a function of the XINR at the MS and SNRs at the UL and DL. As Figs. 5c and 5d suggest, *to obtain non-negligible rate improvements from FD, SNRs need to be sufficiently high compared to the XINRs*.

For a *general number of channels*, the problem of allocating power levels to orthogonal UL and DL channels such that the sum rate is maximized is non-convex. However, we showed that under mild restrictions, the problem is, in fact, amenable to efficient optimization methods. The restrictions impose a lower bound on the amount of SIC that needs to be obtained for given levels of SNR. We show that if these restrictions are not satisfied, FD cannot provide appreciable rate improvements. Based on realistic models of residual SI at the BS [1] and at the MS [3, 15], we demonstrated in [13] that *simple power allocation methods, such as the high SINR approximation power allocation, are near-optimal whenever the gains from FD are non-negligible*.

In [14] we focused on determining the *capacity region* of an FD link. This is equivalent to the problem of maximizing one of the (UL and DL) rates when the other is fixed. An FD capacity region is not convex in general. However, the region can be “convexified” through time sharing between different FD rate pairs. We refer to a convexified capacity region as the *time-division full-duplex (TDFD) capacity region*. A TDFD capacity region generally provides higher rates than its corresponding FD region. Moreover, a *convex (TDFD) capacity region is desirable, since most resource allocation and scheduling algorithms rely on convexity, as providing performance guarantees for a non-convex region is hard*.

Although the problem of determining either the FD or the TDFD capacity region is non-con-

vex in general, we developed an algorithm (AltMax) that determines the TDFD region under mild restrictions and is guaranteed to converge to a stationary point, which in practice is a global optimum. The restrictions lead to suboptimality mainly in the region where, on average, XINR over channels is high compared to the average SNR over channels. Building on insights from our numerical experiments, we also developed a simple heuristic with similar performance but lower running time. The rate improvements obtained by AltMax and the heuristic are illustrated in Figs. 5e and 5f for the residual SI of the FD receiver implemented in [15], and Figs. 5g and 5h for the residual SI of the FD receiver implemented in [3]. In the figures, we assume that 110 dB is required for the TX signal to be cancelled to the noise level, as in, for example, [1]. We can observe from Figs. 5e–5h that as the average SNR increases, the rate improvements increase.¹

Comparing the rate improvements for the canceller from [16] with the rate improvements for the canceller from [3] (i.e., comparing Fig. 5e with Fig. 5g and Fig. 5f with Fig. 5h, it is not difficult to observe, as expected, that more broadband cancellation ([3] vs. [15]) provides higher rate improvements. Moreover, we can observe that in the regions of higher SNRs, where the rate improvements are higher (Figs. 5f and 5h), AltMax and the heuristic provide almost indistinguishable solutions. Therefore, *in the regions of high rate improvements, the TDFD capacity region can be determined near optimally with a simple heuristic*.

Finally, we note that FD has other advantages in addition to potential $2\times$ rate improvement. For example, in WiFi systems, FD can reduce collisions and, consequently, reduce the packet loss.

CONCLUSION

In summary, the Columbia FlexICoN project has been focusing on IC-based FD transceivers spanning RF to millimeter-wave, reconfigurable antenna cancellation, non-magnetic CMOS circulators, and MAC layer algorithms based on realistic hardware models. While exciting progress has been made in the last few years by the research community as a whole, several problems remain to be solved before FD wireless can become a widely deployed reality. Continued improvements are necessary in IC-based FD transceivers toward increased total cancellation over wide BWs and support for higher TX power levels through improved RX and circulator linearity. Incorporation of FD in large-scale phased-array transceivers is another open research problem. The extension of IC-based SIC concepts to MIMO transceivers is an important and formidable challenge. In a MIMO transceiver, SI will exist between every TX-RX pair, and a brute force implementation will cause canceller complexity to scale quadratically with the number of MIMO elements. At the higher layers, while extensive recent research has been devoted to this area, the implications of different possible PHY layer implementations are still not fully understood. Moreover, there are still several important open problems related to MAC layer design for both cellular and random access networks. Specifically, it is necessary to design algorithms that support asymmetric UL and DL traffic requirements and provide fairness

While exciting progress has been made in the last few years by the research community as a whole, several problems remain to be solved before FD wireless can become a widely-deployed reality.

Continued improvements are necessary in IC-based FD transceivers towards increased total cancellation over wide BWs and support for higher TX power levels through improved RX and circulator linearity.

¹ In fact, it can be shown that for any finite values of XINR over channels, if it was possible to increase the SNR to infinity, the rate improvements would reach the theoretical upper bound of 2.

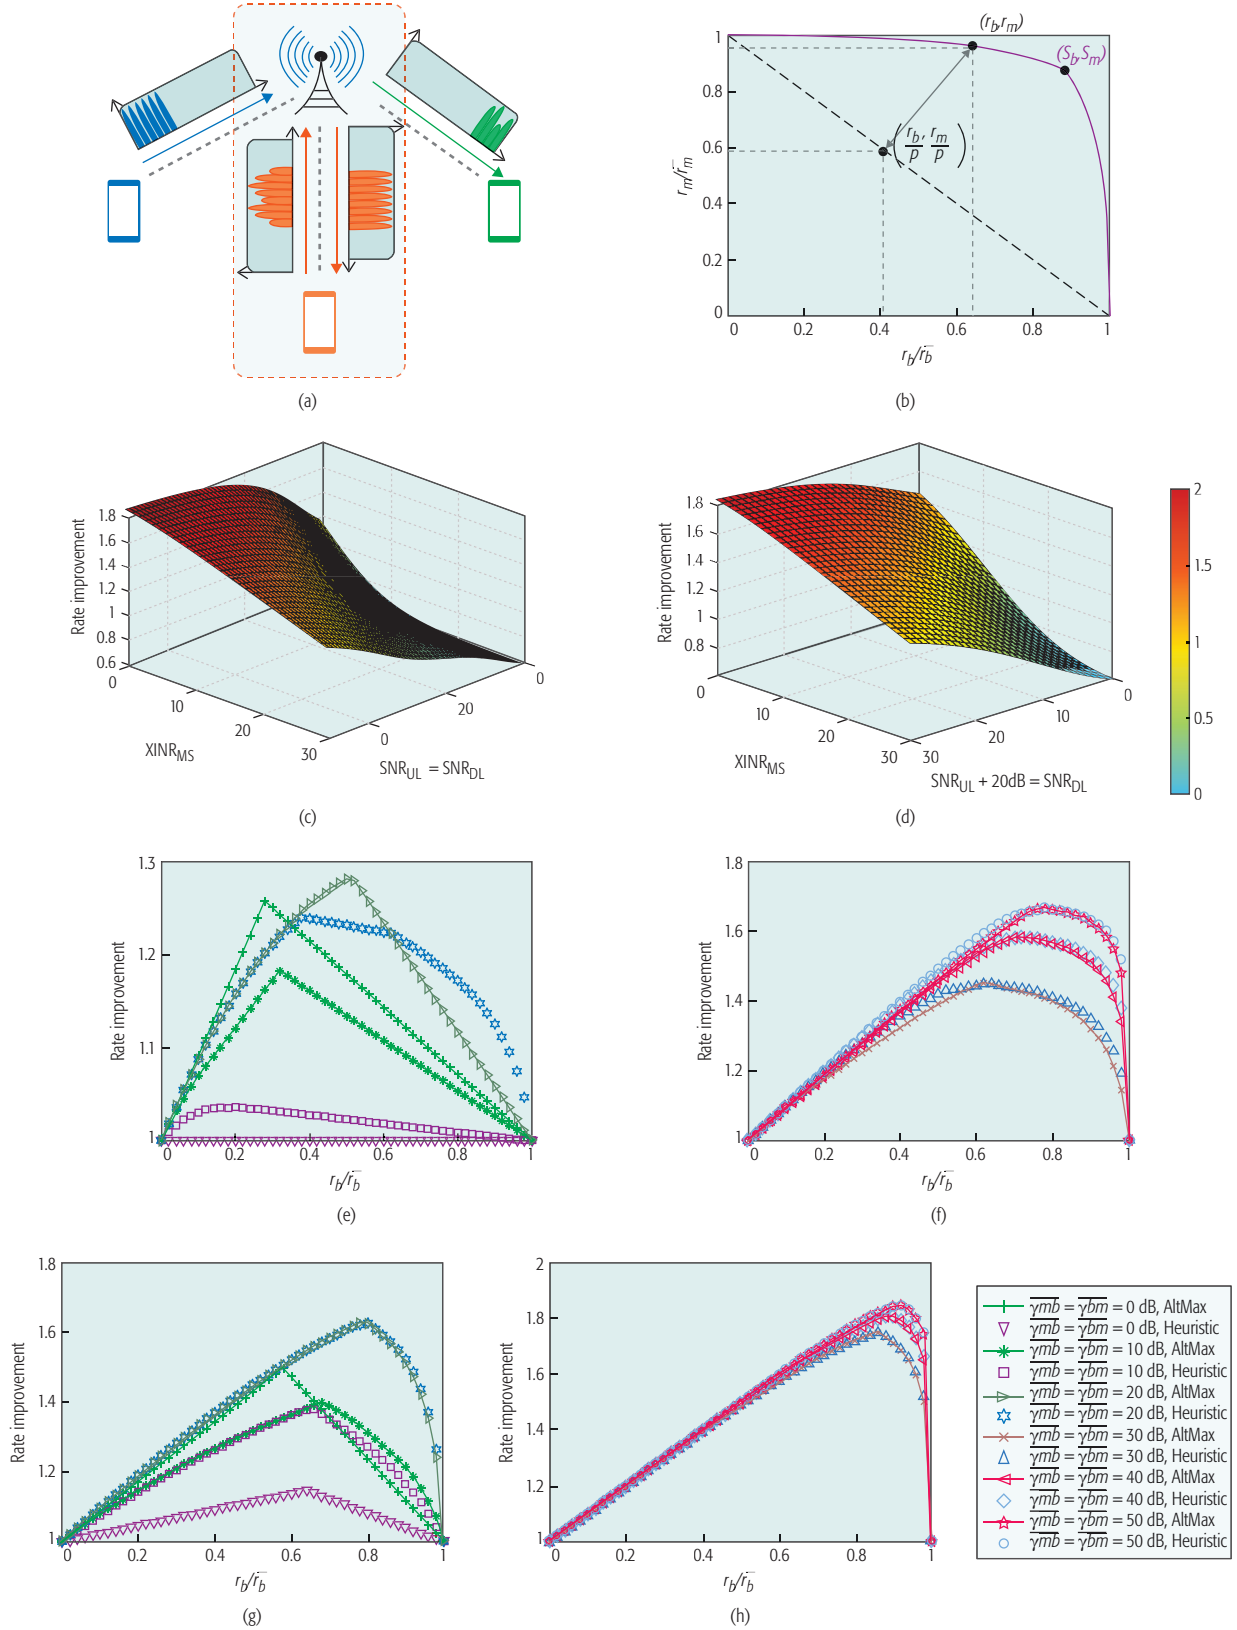


Figure 5. Resource allocation and rate gains in FD systems: a) illustration of a network with FD and HD users, with the FD link outlined by a box; b) definition of rate improvement; c) rate improvement at (s_b, s_m) for equal SNRs on UL and DL; d) rate improvement at (s_b, s_m) for 20 dB higher SNR at the DL than at the UL; e) rate improvements for fixed r_b and maximized r_m , for the FD receiver from [15] and average SNR (denoted as $\gamma_{mb} = \gamma_{bm}$) up to 20 dB; f) rate improvements for fixed r_b and maximized r_m , for the FD receiver from [15] and average SNR from 30 dB to 50 dB; g) rate improvements for fixed r_b and maximized r_m , for the FD receiver from [3] and average SNR up to 20 dB; h) rate improvements for fixed r_b and maximized r_m , for the FD receiver from [3] and average SNR from 30 dB to 50 dB.

in networks composed of both FD and legacy HD nodes. In addition, it is important to consider interference management in OFDM networks jointly at the PHY and MAC layers.

Integrated FD is suitable for many applications ranging from backhaul, relays, and WiFi to small-cell LTE (where cancellation requirements are relaxed compared to macrocells). Although antenna cancellation techniques provide wider BW and relax dynamic range requirements, they tend to have large form factors, making them more practical for point-to-point/less form-factor-constrained applications like backhaul and relays. On the other hand, RF cancellation and on-chip circulator-based FD works generally result in smaller form factors, making them more feasible for mobile devices for WiFi and small cells with future improvements in cancellation BW and power handling capability.

ACKNOWLEDGMENTS

This work was supported in part by the DARPA RF-FPGA program, the DARPA ACT program, NSF grant ECCS-1547406, the People Programme (Marie Curie Actions) of the European Union's Seventh Framework Programme (FP7/2007-2013) under REA grant agreement no. [PIIF-GA-2013-629740].11, and a Qualcomm Innovation Fellowship.

REFERENCES

- [1] D. Bharadia, E. McMillin, and S. Katti, "Full Duplex Radios," *Proc. ACM SIGCOMM '13*, Oct. 2013, pp. 375–86.
- [2] S. Enserink *et al.*, "Joint Analog and Digital Interference Cancellation," *Proc. 2014 IEEE MTT-S Int'l. Microwave Symp.*, June 2014, pp. 1–3.
- [3] J. Zhou *et al.*, "Integrated Wideband Self-Interference Cancellation in the RF Domain for FDD and Full-Duplex Wireless," *IEEE J. Solid-State Circuits*, vol. 50, no. 12, Dec. 2015, pp. 3015–31.
- [4] T. Dinc, A. Chakrabarti, and H. Krishnaswamy, "A 60GHz CMOS Full-Duplex Transceiver and Link with Polarization-Based Antenna and RF Cancellation," *IEEE J. Solid-State Circuits*, vol. 51, no. 5, 2016, pp. 1125–40.
- [5] J. Zhou, N. Reiskarimian, and H. Krishnaswamy, "Receiver with Integrated Magnetic-Free N-Path-Filter-Based Non-Reciprocal Circulator and Baseband Self-Interference Cancellation for Full-Duplex Wireless," *Proc. IEEE ISSCC '16*, Feb. 2016, pp. 178–80.
- [6] 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, May 2015, pp. 1189–1202.
- [7] D. J. van den Broek, E. A. M. Klumperink, and B. Nauta, "An in-Band Full-Duplex Radio Receiver with a Passive Vector Modulator Downmixer for Self-Interference Cancellation," *IEEE J. Solid-State Circuits*, vol. 50, no. 12, Dec. 2015, pp. 3003–14.
- [8] T. Dinc and H. Krishnaswamy, "A T/R Antenna Pair with Polarization-Based Wideband Reconfigurable Self-Interference Cancellation for Simultaneous Transmit and Receive," *Proc. IEEE IMS'15*, 2015, pp. 1–4.
- [9] N. Reiskarimian and H. Krishnaswamy, "Magnetic-Free Non-Reciprocity Based on Staggered Commutation," *Nature Commun.*, vol. 7, no. 4, Apr. 2016.
- [10] M. Darvishi *et al.*, "Widely Tunable 4th Order Switched Gm-C Band-Pass Filter Based on N-Path Filters," *IEEE J. Solid-State Circuits*, vol. 47, no. 12, Dec 2012, pp. 3105–19.
- [11] E. Yetisir, C.-C. Chen, and J. Volakis, "Low-Profile UWB 2-Port Antenna With High Isolation," *IEEE Antennas Wireless Propag. Lett.*, vol. 13, 2014, pp. 55–58.
- [12] N. A. Estep, D. L. Sounas, and A. Alu, "Magnetless Microwave Circulators Based on Spatiotemporally Modulated Rings of Coupled Resonators," *IEEE Trans. Microwave Theory Tech.*, vol. 64, no. 2, Feb. 2016, pp. 502–18.
- [13] J. Marasevic *et al.*, "Resource Allocation and Rate Gains in Practical Full-Duplex Systems," *IEEE/ACM Trans. Net.*, vol. 25, no. 1, Feb. 2017, pp. 292–305.
- [14] J. Marasevic and G. Zussman, "On the Capacity Regions of Single-Channel and Multi-Channel Full-Duplex Links," *Proc. ACM MobiHoc '16*, 2016.

- [15] J. Zhou *et al.*, "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, Dec. 2014, pp. 3046–62.

BIOGRAPHIES

JIN ZHOU [S'11, M'17] received his B.S. degree in electronics science and technology from Wuhan University, China, in 2008, his M.S. degree in microelectronics from Fudan University, Shanghai, China, in 2011, and his Ph.D. degree in electrical engineering from Columbia University, New York, in 2017. From 2011 to 2012, he worked as an RF integrated circuits design engineer with MediaTek Singapore. In 2017, he joined the Department of Electrical and Computer Engineering at the University of Illinois at Urbana-Champaign as an assistant professor.

NEGAR REISKARIMIAN [S'15] received her B.Sc. and M.Sc. degrees in telecommunication systems and microelectronic circuits, both from Sharif University of Technology, Tehran, Iran, in 2011 and 2013. She is currently pursuing her Ph.D. degree in electrical engineering at Columbia University. She is the recipient of a 2016 Qualcomm Innovation Fellowship and the 2017 IEEE SSCS Pre-doctoral Achievement Award, the Analog Devices Outstanding Student Designer Award, and an IEEE MTT-S Graduate Fellowship.

JELENA DIAKONIKOLAS received her B.Sc. degree from the University of Belgrade in 2011, and her M.S. and Ph.D. degrees from Columbia University in 2012 and 2016. She is currently a post-doctoral associate at Boston University. Her research focuses on principled design of networked systems. She is a recipient of the M.S. Award of Excellence and the Jacob Millman Prize for Excellence in Teaching Assistance from Columbia University. She is a co-winner of the Qualcomm Innovation Fellowship 2015 award, and was listed as one of the "10 Women in Networking/Communications That You Should Watch" in 2016.

TOLGA DINC [S'09] received his B.S. and M.S. degrees in electrical engineering from Sabanci University, Istanbul, Turkey, in 2010 and 2012, respectively. He is currently working toward a Ph.D. degree in electrical engineering at Columbia University. He is a recipient of several honors and awards including the IEEE RFIC Symposium Best Student Paper Award (1st Place) in 2015.

TINGJUN CHEN [S'16] received his B.Eng. degree in electronic engineering from Tsinghua University, Beijing, China, in 2014, and his M.S. degree in electrical engineering from Columbia University in 2015. He is currently a Ph.D. student in electrical engineering at Columbia University. His research interests are in algorithms, optimization, and system design in the Internet of Things, energy harvesting networks, full-duplex networks, and 5G networks. He received the Wei Family Private Foundation Fellowship, the Columbia University Electrical Engineering Armstrong Memorial Award, and the ACM CoNEXT 2016 Best Paper Award.

GIL ZUSSMAN [S'02, M'05, SM'07] received his Ph.D. degree in electrical engineering from the Technion in 2004 and was a postdoctoral associate at MIT in 2004–2007. Since 2007 he has been with Columbia University where he is currently an associate professor of electrical engineering. His research interests are in the areas of wireless, mobile, and resilient networks. He received the Fulbright Fellowship, DTRA Young Investigator Award, NSF CAREER Award, and two Marie Curie Fellowships. He was the PI of a team that won 1st place in the 2009 Vodafone Americas Foundation Wireless Innovation Project competition. He is a co-recipient of seven best paper awards, including the ACM SIGMETRICS/IFIP Performance '06 Best Paper Award, the 2011 IEEE Communications Society Award for Advances in Communication, and the ACM CoNEXT'16 Best Paper Award.

HARISH KRISHNASWAMY [S'03, M'09] received his B.Tech. degree in electrical engineering from the Indian Institute of Technology, Madras, India, in 2001, and his M.S. and Ph.D. degrees in electrical engineering from the University of Southern California (USC), Los Angeles, in 2003 and 2009, respectively. In 2009, he joined the Electrical Engineering Department, Columbia University, where he is currently an associate professor. His research interests broadly span integrated devices, circuits, and systems for a variety of RF, mmWave and sub-mmWave applications. He serves as a member of the TPC of several conferences. He is currently serving as an IEEE SSCS Distinguished Lecturer for 2017 and 2018. He was the 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, a 2014 IBM Faculty Award, and the 2015 IEEE RFIC Symposium Best Student Paper Award, 1st Place.

It is necessary to design algorithms that support asymmetric UL and DL traffic requirements and provide fairness in networks composed of both FD and legacy HD nodes. In addition, it is important to consider interference management in OFDM networks jointly at the PHY and MAC layers.